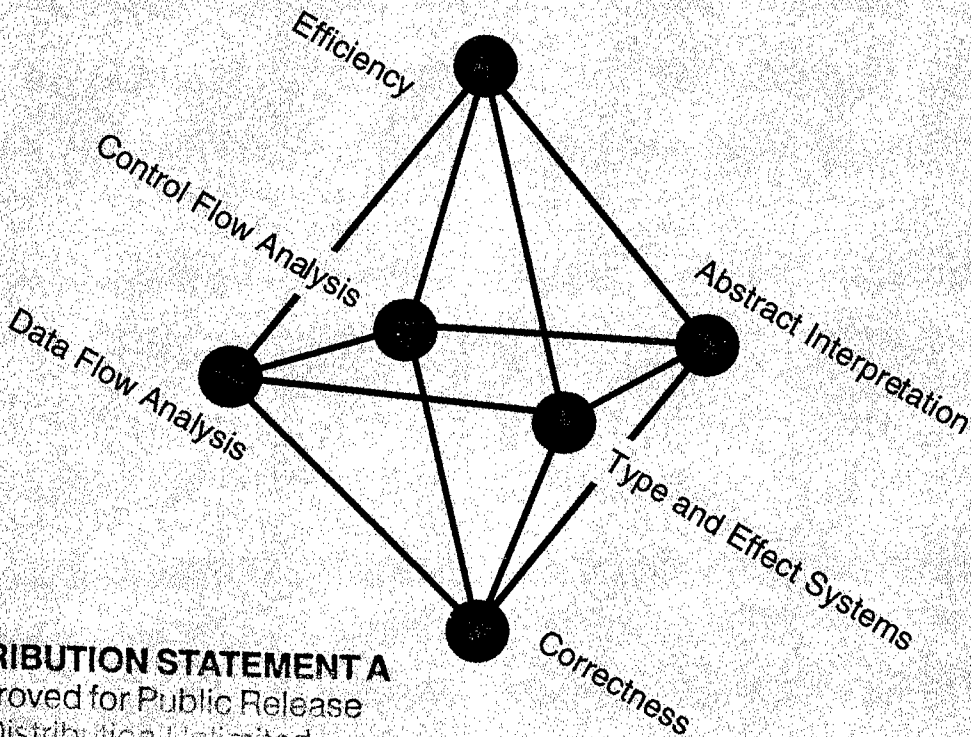


*Flemming Nielson
Hanne Riis Nielson
Chris Hankin*

Principles of Program Analysis

19990728 019



DISTRIBUTION STATEMENT A
Approved for Public Release
Distribution Unlimited

© Flemming Nielson
© Hanne Riis Nielson
© Chris Hankin

*Do not distribute
without permission.
Dagstuhl, Nov.1998.*

<http://www.daimi.au.dk/~hrn/PPA/ppa.html>

AGF 99-10-1860

REPORT DOCUMENTATION PAGE

Form Approved OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE 1999	3. REPORT TYPE AND DATES COVERED Conference Proceedings	
4. TITLE AND SUBTITLE International Workshop on Principles of Program Analysis			5. FUNDING NUMBERS F61775-99-WF003	
6. AUTHOR(S) Conference Committee				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Aarhus Univ (Bldg 540), Computer Science Department Ny Munkegrade Aarhus C DK-8000 Denmark			8. PERFORMING ORGANIZATION REPORT NUMBER N/A	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) EOARD PSC 802 BOX 14 FPO 09499-0200			10. SPONSORING/MONITORING AGENCY REPORT NUMBER CSP 99-5003	
11. SUPPLEMENTARY NOTES				
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.			12b. DISTRIBUTION CODE A	
13. ABSTRACT (Maximum 200 words) The Final Proceedings for International Workshop on Principles of Program Analysis, 9 November 1998 - 13 November 1998 This is an international workshop to foster research in program analysis under the discipline of computer science. Topics include: data flow analysis, control flow analysis, and abstract interpretation of computer programs.				
14. SUBJECT TERMS EOARD, Program Analysis, Computers			15. NUMBER OF PAGES 399	
			16. PRICE CODE N/A	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL	

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. Z39-18
298-102

We wish to thank the following for their contribution to the success of this conference:

- United States Air Force European Office of Aerospace Research and Development.
- Center for Information Technology.
- European Union: Training and Mobility of Researchers.

Contents

Preface	ix
1 Introduction	1
1.1 The Nature of Program Analysis	1
1.2 Setting the Scene	3
1.3 Data Flow Analysis	5
1.3.1 The Equational Approach	5
1.3.2 The Constraint Based Approach	8
1.4 Control Flow Analysis	10
1.5 Abstract Interpretation	13
1.6 Type and Effect Systems	17
1.6.1 Annotated Type Systems	18
1.6.2 Effect Systems	22
1.7 Algorithms	25
1.8 Transformations	26
Concluding Remarks	29
Mini Projects	29
Exercises	31
2 Data Flow Analysis	33
2.1 Intraprocedural Analysis	33
2.1.1 Available Expressions Analysis	37
2.1.2 Reaching Definitions Analysis	41
2.1.3 Very Busy Expressions Analysis	44

2.1.4	Live Variables Analysis	47
2.1.5	Derived Data Flow Information	50
2.2	Theoretical Properties	52
2.2.1	Structural Operational Semantics	52
2.2.2	Correctness of Live Variables Analysis	57
2.3	Monotone Frameworks	63
2.3.1	Basic Definitions	65
2.3.2	The Examples Revisited	68
2.3.3	A Non-distributive Example	70
2.4	Equation Solving	72
2.4.1	The MFP Solution	72
2.4.2	The MOP Solution	76
2.5	Interprocedural Analysis	80
2.5.1	Structural Operational Semantics	83
2.5.2	Intraprocedural versus Interprocedural Analysis	85
2.5.3	Making Context Explicit	88
2.5.4	Call Strings as Context	93
2.5.5	Assumption Sets as Context	96
2.5.6	Flow-Sensitivity versus Flow-Insensitivity	98
2.6	Shape Analysis	101
2.6.1	Structural Operational Semantics	103
2.6.2	Shape Graphs	106
2.6.3	The Analysis	112
	Concluding Remarks	127
	Mini Projects	130
	Exercises	134
3	Control Flow Analysis	139
3.1	Abstract 0-CFA Analysis	139
3.1.1	The Analysis	141
3.1.2	Well-definedness of the Analysis	148
3.2	Theoretical Properties	150
3.2.1	Structural Operational Semantics	150

3.2.2	Semantic Correctness	155
3.2.3	Existence of Solutions	159
3.2.4	Coinduction versus Induction	162
3.3	Syntax Directed 0-CFA Analysis	165
3.3.1	Syntax Directed Specification	166
3.3.2	Preservation of Solutions	168
3.4	Constraint Based 0-CFA Analysis	170
3.4.1	Preservation of Solutions	172
3.4.2	Solving the Constraints	173
3.5	Adding Data Flow Analysis	179
3.5.1	Abstract Values as Powersets	179
3.5.2	Abstract Values as Complete Lattices	182
3.6	Adding Context Information	186
3.6.1	Uniform k -CFA Analysis	188
3.6.2	Interprocedural Analysis Revisited	192
	Concluding Remarks	193
	Mini Projects	197
	Exercises	199
4	Abstract Interpretation	205
4.1	A Mundane Approach to Correctness	205
4.1.1	Correctness Relations	208
4.1.2	Representation Functions	210
4.1.3	A Modest Generalisation	213
4.2	Approximation of Fixed Points	215
4.2.1	Widening Operators	218
4.2.2	Narrowing Operators	223
4.3	Galois Connections	227
4.3.1	Properties of Galois Connections	231
4.3.2	Galois Insertions	235
4.4	Systematic Design of Galois Connections	239
4.4.1	Component-wise Combinations	242
4.4.2	Other Combinations	246

4.5	Induced Operations	251
4.5.1	Inducing along the Abstraction Function	251
4.5.2	Application to Data Flow Analysis	254
4.5.3	Inducing along the Concretisation Function	260
	Concluding Remarks	263
	Mini Projects	266
	Exercises	268
5	Type and Effect Systems	277
5.1	Control Flow Analysis	277
5.1.1	The Underlying Type System	278
5.1.2	The Analysis	281
5.2	Theoretical Properties	285
5.2.1	Natural Semantics	286
5.2.2	Semantic Correctness	288
5.2.3	Existence of Solutions	291
5.3	Inference Algorithms	294
5.3.1	An Algorithm for the Underlying Type System	294
5.3.2	An Algorithm for Control Flow Analysis	300
5.3.3	Syntactic Soundness and Completeness	306
5.3.4	Existence of Solutions	311
5.4	Effects	313
5.4.1	Side Effect Analysis	313
5.4.2	Exception Analysis	319
5.4.3	Region Inference	324
5.5	Behaviours	333
5.5.1	Communication Analysis	333
	Concluding Remarks	343
	Mini Projects	347
	Exercises	349
6	Algorithms	355
A	Partially Ordered Sets	357

A.1 Basic Definitions	357
A.2 Construction of Complete Lattices	361
A.3 Chains	362
A.4 Fixed Points	366
Concluding Remarks	368
B Induction and Coinduction	369
B.1 Proof by Induction	369
B.2 Introducing Coinduction	371
B.3 Proof by Coinduction	375
Concluding Remarks	377
C Graphs and Regular Expressions	379
Index of Notation	381
Index	385
Bibliography	389

Preface

(This material remains to be integrated.)

Chapter 1

Introduction

In this book we shall introduce four of the main approaches to program analysis: Data Flow Analysis, Control Flow Analysis, Abstract Interpretation, and Type and Effect Systems. Each of Chapters 2 to 5 deals with one of these approaches to some length and generally treats the more advanced material in later sections. Throughout the book we aim at stressing the many similarities between what may at a first glance appear to be very unrelated approaches. To help getting this idea across, and to serve as a gentle introduction, this chapter treats all of the approaches at the level of examples. The technical details are worked out but it may be difficult to apply the techniques to related examples until some of the material of later chapters have been studied.

1.1 The Nature of Program Analysis

Program analysis offers static compile-time techniques for predicting safe and computable approximations to the set of values or behaviours arising dynamically at run-time when executing a program on a computer. A main application is to allow compilers to generate code avoiding *redundant* computations, e.g. by reusing available results or by moving loop invariant computations out of loops, or avoiding *superfluous* computations, e.g. of results known to be not needed or of results known already at compile-time. Among the more recent applications is the validation of software (possibly purchased from sub-contractors) to reduce the likelihood of malicious or unintended behaviour. Common for these applications is the need to combine information from different parts of the program.

A main aim of this book is to give an overview of a number of approaches to program analysis, all of which have a quite extensive literature, and to show




Figure 1.1: The nature of approximation: erring on the safe side.

that there is a large amount of *commonality* among the approaches. This should help in cultivating the ability to choose the right approach for the right task and in exploiting insights developed in one approach to enhance the power of other approaches.

One common theme behind all approaches to program analysis is that in order to remain computable one can only provide *approximate answers*. As an example consider a simple language of statements and the program

```
read(x); (if x>0 then y:=1 else (y:=2;S)); z:=y
```

where S is some statement that does not contain an assignment to y . Intuitively, the values of y that can reach $z:=y$ will be 1 or 2.

Now suppose an analysis claims that the only value for y that can reach $z:=y$ is in fact 1. While this seems intuitively wrong, it is in fact correct in the case where S is known never to terminate for $x \leq 0$ and $y = 2$. But since it is *undecidable* whether or not S terminates, we normally do not expect our analysis to attempt to detect this situation. So in general, we expect the program analysis to produce a possibly *larger set* of possibilities than what will ever happen during execution of the program. This means that we shall also accept a program analysis claiming that the values of y reaching $z:=y$ are among 1, 2 or 27, although we will clearly prefer the analysis that gives the more precise answer that the values are among 1 or 2. This notion of safe approximation is illustrated in Figure 1.1. Clearly the challenge is not to

produce the safe “ $\{d_1, \dots, d_N\}$ ” too often as the analysis will then be utterly useless. Note, that although the analysis does not give precise information it may still give useful information: knowing that the value of y is one of 1, 2 and 27 just before the assignment $z:=y$ still tells us that z will be positive, and that z will fit within 1 byte of storage etc. To avoid confusion it may help to be precise in the use of terminology: it is better to say “the values of y possible at $z:=y$ are among 1 and 2” than the slightly shorter and more frequently used “the values of y possible at $z:=y$ are 1 and 2”.

Another common theme, to be stressed throughout this book, is that all program analyses should be *semantics based*: this means that the information obtained from the analysis can be proved to be safe (or correct) with respect to a semantics of the programming language. It is a sad fact that new program analyses often contain subtle bugs, and a formal justification of the program analysis will help finding these bugs sooner rather than later. However, we should stress that we do *not* suggest that program analyses be *semantics directed*: this would mean that the structure of the program analysis should reflect the structure of the semantics and this will be the case only for a few approaches which are not covered in this book.

1.2 Setting the Scene

Syntax of the WHILE language. We shall consider a simple imperative language called WHILE. A program in WHILE is just a statement which may be, and normally will be, a sequence of statements. In the interest of simplicity, we will associate data flow information with single assignment statements, the tests that appear in conditionals and loops, and skip statements. We will require a method to identify these. The most convenient way of doing this is to work with a labelled program – as indicated in the syntax below. We will often refer to the labelled items (assignments, tests and skip statements) as *elementary blocks*. In this chapter we will assume that distinct elementary blocks are initially assigned distinct labels; we could drop this requirement, in which case some of the examples would need to be slightly reformulated and the resultant analyses would be less accurate.

We use the following syntactic categories:

$$\begin{array}{ll} a \in \mathbf{AExp} & \text{arithmetic expressions} \\ b \in \mathbf{BExp} & \text{boolean expressions} \\ S \in \mathbf{Stmt} & \text{statements} \end{array}$$

We assume some countable set of variables is given; numerals and labels will not be further defined and neither will the operators:

$$\begin{array}{ll} x, y \in \mathbf{Var} & \text{variables} \\ n \in \mathbf{Num} & \text{numerals} \\ \ell \in \mathbf{Lab} & \text{labels} \end{array}$$

ℓ	$RD_{entry}(\ell)$	$RD_{exit}(\ell)$
1	$(x, ?), (y, ?), (z, ?)$	$(x, ?), (y, 1), (z, ?)$
2	$(x, ?), (y, 1), (z, ?)$	$(x, ?), (y, 1), (z, 2)$
3	$(x, ?), (y, 1), (y, 5), (z, 2), (z, 4)$	$(x, ?), (y, 1), (y, 5), (z, 2), (z, 4)$
4	$(x, ?), (y, 1), (y, 5), (z, 2), (z, 4)$	$(x, ?), (y, 1), (y, 5), (z, 4)$
5	$(x, ?), (y, 1), (y, 5), (z, 4)$	$(x, ?), (y, 5), (z, 4)$
6	$(x, ?), (y, 1), (y, 5), (z, 2), (z, 4)$	$(x, ?), (y, 6), (z, 2), (z, 4)$

Table 1.1: Reaching Definitions information for the factorial program.

$op_a \in Op_a$ arithmetic operators
 $op_b \in Op_b$ boolean operators
 $op_r \in Op_r$ relational operators

The syntax of the language is given by the following *abstract syntax*:

$$\begin{aligned}
 a &::= x \mid n \mid a_1 op_a a_2 \\
 b &::= \text{true} \mid \text{false} \mid \text{not } b \mid b_1 op_b b_2 \mid a_1 op_r a_2 \\
 S &::= [x := a]^\ell \mid [\text{skip}]^\ell \mid S_1; S_2 \mid \\
 &\quad \text{if } [b]^\ell \text{ then } S_1 \text{ else } S_2 \mid \text{while } [b]^\ell \text{ do } S
 \end{aligned}$$

One way to think of the abstract syntax is as specifying the parse trees of the language; it will then be the purpose of the *concrete syntax* to provide sufficient information to enable unique parse trees to be constructed. In this book shall *not* be concerned with concrete syntax: whenever we talk about some syntactic entity we will always be talking about the abstract syntax so there will be no ambiguity with respect to the form of the entity. We shall use a textual representation of the abstract syntax and to disambiguate it we shall use parentheses. For statements one often writes `begin ... end` or `{...}` for this but we shall feel free to use `(...)`. Similarly, we use brackets `(...)` to resolve ambiguities in other syntactic categories. To cut down on the number of brackets needed we shall use the familiar relative precedences of arithmetic, boolean and relational operators.

Example 1.1 An example of a program written in this language is the following which computes the factorial of the number stored in `x` and leaves the result in `z`:

$$[y:=x]^1; [z:=1]^2; \text{while } [y>1]^3 \text{ do } ([z:=z*y]^4; [y:=y-1]^5); [y:=0]^6 \quad \blacksquare$$

Reaching Definitions Analysis. The use of distinct labels allows us to identify the primitive constructs of a program without explicitly constructing a flow graph (or flow chart). It also allows us to introduce a program analysis to be used throughout the chapter: *Reaching Definitions Analysis*, or as it more properly should be called, reaching assignments analysis:

An assignment (called a definition in classical literature) of the form $[x := a]^\ell$ may reach a certain program point (typically the entry or exit of an elementary block) if there is an execution of the program where x was last assigned a value at ℓ when the program point is reached.

Consider the factorial program of Example 1.1. Here $[y:=x]^1$ reaches the entry to $[z:=1]^2$; to allow a more succinct presentation we shall say that $(y,1)$ reaches the entry to 2. Also we shall say that $(x,?)$ reaches the entry to 2; here “?” is a special label not appearing in the program and it is used to record the possibility of an uninitialised variable reaching a certain program point.

Full information about reaching definitions for the factorial program is then given by the pair $RD = (RD_{entry}, RD_{exit})$ of functions in Table 1.1. Careful inspection of this table reveals that the entry and exit information agree for elementary blocks of the form $[b]^\ell$ whereas for elementary blocks of the form $[x := a]^\ell$ they may differ on pairs (x, ℓ') . We shall come back to this when formulating the analysis in subsequent sections.

Returning to the discussion of safe approximation note that if we modify Table 1.1 to include the pair $(z,2)$ in $RD_{entry}(5)$ and $RD_{exit}(5)$ we still have safe information about reaching definitions but the information is more approximate. However, if we remove $(z,2)$ from $RD_{entry}(6)$ and $RD_{exit}(6)$ then the information will no longer be safe – there exists a run of the factorial program where the set $\{(x,?), (y,6), (z,4)\}$ does not correctly describe the reaching definitions at the exit of label 6.

1.3 Data Flow Analysis

In *Data Flow Analysis* it is customary to think of a program as a graph: the nodes are the elementary blocks and the edges describe how control might pass from one elementary block to another. Figure 1.2 shows the flow graph for the factorial program of Example 1.1. We shall first illustrate an *equational approach* to Data Flow Analysis and then a *constraint based approach*.

1.3.1 The Equational Approach

The equation system. An analysis like Reaching Definitions can be specified by extracting a number of equations from a program. There are two classes of equations. One class of equations relate exit information of a node to entry information for the same node. For the factorial program

$$[y:=x]^1; [z:=1]^2; \text{ while } [y>1]^3 \text{ do } ([z:=z*y]^4; [y:=y-1]^5); [y:=0]^6$$

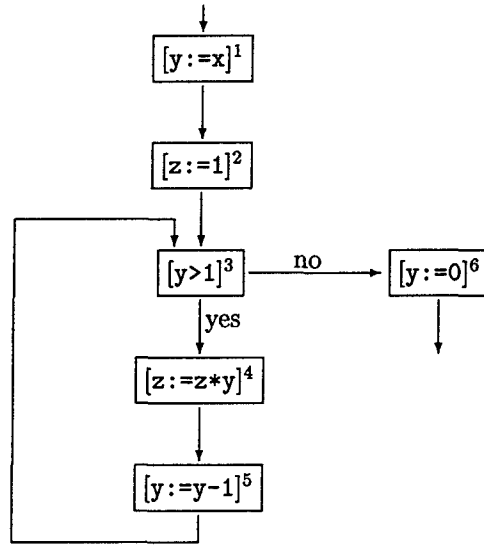


Figure 1.2: Flow graph for the factorial program.

we obtain the following six equations:

$$\begin{aligned}
 RD_{exit}(1) &= (RD_{entry}(1) \setminus \{(y, \ell) \mid \ell \in \mathbf{Lab}\}) \cup \{(y, 1)\} \\
 RD_{exit}(2) &= (RD_{entry}(2) \setminus \{(z, \ell) \mid \ell \in \mathbf{Lab}\}) \cup \{(z, 2)\} \\
 RD_{exit}(3) &= RD_{entry}(3) \\
 RD_{exit}(4) &= (RD_{entry}(4) \setminus \{(z, \ell) \mid \ell \in \mathbf{Lab}\}) \cup \{(z, 4)\} \\
 RD_{exit}(5) &= (RD_{entry}(5) \setminus \{(y, \ell) \mid \ell \in \mathbf{Lab}\}) \cup \{(y, 5)\} \\
 RD_{exit}(6) &= (RD_{entry}(6) \setminus \{(y, \ell) \mid \ell \in \mathbf{Lab}\}) \cup \{(y, 6)\}
 \end{aligned}$$

These are instances of the following schema: for an assignment $[x := a]^{\ell'}$ we exclude all pairs (x, ℓ) from $RD_{entry}(\ell')$ and add (x, ℓ') in order to obtain $RD_{exit}(\ell')$ – this reflects that x is redefined at ℓ' . For all other elementary blocks $[\dots]^{\ell'}$ we let $RD_{exit}(\ell')$ equal $RD_{entry}(\ell')$ – reflecting that no variables are changed.

The other class of equations relate entry information of a node to exit information of nodes from which there is an edge to the node of interest; that is, entry information is obtained from all the exit information where control could have come from. For the example program we obtain the following equations:

$$RD_{entry}(2) = RD_{exit}(1)$$

$$\begin{aligned}
RD_{entry}(3) &= RD_{exit}(2) \cup RD_{exit}(5) \\
RD_{entry}(4) &= RD_{exit}(3) \\
RD_{entry}(5) &= RD_{exit}(4) \\
RD_{entry}(6) &= RD_{exit}(3)
\end{aligned}$$

In general, we write $RD_{entry}(\ell) = RD_{exit}(\ell_1) \cup \dots \cup RD_{exit}(\ell_n)$ if ℓ_1, \dots, ℓ_n are all the labels from which control might pass to ℓ . We shall consider more precise ways of explaining this in Chapter 2. Finally, let us consider the equation

$$RD_{entry}(1) = \{(x, ?) \mid x \text{ is a variable in the program}\}$$

that makes it clear that the label “?” is to be used for uninitialised variables; so in our case

$$RD_{entry}(1) = \{(x, ?), (y, ?), (z, ?)\}$$

The least solution. The above system of equations defines the twelve sets

$$RD_{entry}(1), \dots, RD_{exit}(6)$$

in terms of each other. Writing \vec{RD} for this twelve-tuple of sets we can regard the equation system as defining a function F and demanding that:

$$\vec{RD} = F(\vec{RD})$$

To be more specific we can write

$$F(\vec{RD}) = (F_{entry}(1)(\vec{RD}), F_{exit}(1)(\vec{RD}), \dots, F_{entry}(6)(\vec{RD}), F_{exit}(6)(\vec{RD}))$$

where e.g.:

$$F_{entry}(3)(\dots, RD_{exit}(2), \dots, RD_{exit}(5), \dots) = RD_{exit}(2) \cup RD_{exit}(5)$$

It should be clear that F operates over twelve-tuples of sets of pairs of variables and labels; this can be written as

$$F : (\mathcal{P}(\mathbf{Var}_* \times \mathbf{Lab}_*))^{12} \rightarrow (\mathcal{P}(\mathbf{Var}_* \times \mathbf{Lab}_*))^{12}$$

where it might be natural to take $\mathbf{Var}_* = \mathbf{Var}$ and $\mathbf{Lab}_* = \mathbf{Lab}$. However, it will simplify the presentation in this chapter to let \mathbf{Var}_* be a *finite* subset of \mathbf{Var} that contains the variables occurring in the program S_* of interest and similarly for \mathbf{Lab}_* . So for the example program we might have $\mathbf{Var}_* = \{x, y, z\}$ and $\mathbf{Lab}_* = \{1, \dots, 6\}$.

It is immediate that $(\mathcal{P}(\mathbf{Var}_* \times \mathbf{Lab}_*))^{12}$ can be partially ordered by setting

$$\vec{RD} \subseteq \vec{RD}' \quad \text{iff} \quad \forall i : RD_i \subseteq RD'_i$$

where $\vec{RD} = (RD_1, \dots, RD_{12})$ and similarly $\vec{RD}' = (RD'_1, \dots, RD'_{12})$. This turns $(\mathcal{P}(\mathbf{Var}_* \times \mathbf{Lab}_*))^{12}$ into a complete lattice (see Appendix A) with least element

$$\vec{\emptyset} = (\emptyset, \dots, \emptyset)$$

and binary least upper bounds given by:

$$\vec{RD} \sqcup \vec{RD}' = (RD_1 \cup RD'_1, \dots, RD_{12} \cup RD'_{12})$$

It is easy to show that F is in fact a monotone function (see Appendix A) meaning that:

$$\vec{RD} \sqsubseteq \vec{RD}' \text{ implies } F(\vec{RD}) \sqsubseteq F(\vec{RD}')$$

This involves calculations like

$$\begin{aligned} & RD_{exit}(2) \subseteq RD'_{exit}(2) \text{ and } RD_{exit}(5) \subseteq RD'_{exit}(5) \\ \text{imply} & \\ & RD_{exit}(2) \cup RD_{exit}(5) \subseteq RD'_{exit}(2) \cup RD'_{exit}(5) \end{aligned}$$

and the details are left to the reader.

Consider the sequence $(F^n(\vec{\emptyset}))_n$ and note that $\vec{\emptyset} \sqsubseteq F(\vec{\emptyset})$. Since F is monotone, a straightforward mathematical induction (see Appendix B) gives that $F^n(\vec{\emptyset}) \sqsubseteq F^{n+1}(\vec{\emptyset})$ for all n . All the elements of the sequence will be in $(\mathcal{P}(\mathbf{Var}_* \times \mathbf{Lab}_*))^{12}$ and since this is a finite set it cannot be the case that all elements of the sequence are distinct so there must be some n such that:

$$F^{n+1}(\vec{\emptyset}) = F^n(\vec{\emptyset})$$

But since $F^{n+1}(\vec{\emptyset}) = F(F^n(\vec{\emptyset}))$ this just says that $F^n(\vec{\emptyset})$ is a *fixed point* of F and hence that $F^n(\vec{\emptyset})$ is a solution to the above equation system.

In fact we have obtained the *least solution* to the equation system. To see this suppose that \vec{RD} is some other solution, i.e. $\vec{RD} = F(\vec{RD})$. Then a straightforward mathematical induction shows that $F^n(\vec{\emptyset}) \sqsubseteq \vec{RD}$. Hence the solution $F^n(\vec{\emptyset})$ contains the fewest pairs of reaching definitions that is consistent with the program, and intuitively, this is also the solution we want: while we can add additional pairs of reaching definitions without making the analysis semantically unsound, this will make the analysis less usable as discussed in Section 1.1. In Exercise 1.7 we shall see that the least solution is in fact the one displayed in Table 1.1.

1.3.2 The Constraint Based Approach

The constraint system. An alternative to the equational approach above is to use a *constraint based approach*. The idea is here to extract a number of inclusions (or in-equations or constraints) out of a program. We

shall present the constraint system for Reaching Definitions in such a way that the relationship to the equational approach becomes apparent; however, it is not a general phenomenon that the constraints are naturally divided into two classes as was the case for the equations.

For the factorial program

$$[y:=x]^1; [z:=1]^2; \text{ while } [y>1]^3 \text{ do } ([z:=z*y]^4; [y:=y-1]^5); [y:=0]^6$$

we obtain the following constraints for expressing the effect of elementary blocks:

$$\begin{aligned} RD_{exit}(1) &\supseteq RD_{entry}(1) \setminus \{(y, \ell) \mid \ell \in \mathbf{Lab}\} \\ RD_{exit}(1) &\supseteq \{(y, 1)\} \\ RD_{exit}(2) &\supseteq RD_{entry}(2) \setminus \{(z, \ell) \mid \ell \in \mathbf{Lab}\} \\ RD_{exit}(2) &\supseteq \{(z, 2)\} \\ RD_{exit}(3) &\supseteq RD_{entry}(3) \\ RD_{exit}(4) &\supseteq RD_{entry}(4) \setminus \{(z, \ell) \mid \ell \in \mathbf{Lab}\} \\ RD_{exit}(4) &\supseteq \{(z, 4)\} \\ RD_{exit}(5) &\supseteq RD_{entry}(5) \setminus \{(y, \ell) \mid \ell \in \mathbf{Lab}\} \\ RD_{exit}(5) &\supseteq \{(y, 5)\} \\ RD_{exit}(6) &\supseteq RD_{entry}(6) \setminus \{(y, \ell) \mid \ell \in \mathbf{Lab}\} \\ RD_{exit}(6) &\supseteq \{(y, 6)\} \end{aligned}$$

By considering this system a certain methodology emerges: for an assignment $[x := a]^{\ell'}$ we have one constraint that excludes all pairs (x, ℓ) from $RD_{entry}(\ell')$ in reaching $RD_{exit}(\ell')$ and we have one constraint for incorporating (x, ℓ') ; for all other elementary blocks $[\cdot]^{\ell'}$ we just have one constraint that allows everything in $RD_{entry}(\ell')$ to reach $RD_{exit}(\ell')$.

Next consider the constraints for more directly expressing how control may flow through the program. For the example program we obtain the constraints:

$$\begin{aligned} RD_{entry}(2) &\supseteq RD_{exit}(1) \\ RD_{entry}(3) &\supseteq RD_{exit}(2) \\ RD_{entry}(3) &\supseteq RD_{exit}(5) \\ RD_{entry}(5) &\supseteq RD_{exit}(4) \\ RD_{entry}(6) &\supseteq RD_{exit}(3) \end{aligned}$$

In general, we have a constraint $RD_{entry}(\ell) \supseteq RD_{exit}(\ell')$ if it is possible for control to pass from ℓ' to ℓ . Finally, the constraint

$$RD_{entry}(1) \supseteq \{(x, ?), (y, ?), (z, ?)\}$$

records that we cannot be sure about the definition point of uninitialised variables.

The least solution revisited. It is not hard to see that a solution to the equation system presented previously will also be a solution to the above constraint system. To make this connection more transparent we can rearrange the constraints by *collecting* all constraints with the same left hand side. This means that for example

$$\begin{aligned} \text{RD}_{\text{exit}}(1) &\supseteq \text{RD}_{\text{entry}}(1) \setminus \{(y, \ell) \mid \ell \in \mathbf{Lab}\} \\ \text{RD}_{\text{exit}}(1) &\supseteq \{(y, 1)\} \end{aligned}$$

will be replaced by

$$\text{RD}_{\text{exit}}(1) \supseteq (\text{RD}_{\text{entry}}(1) \setminus \{(y, \ell) \mid \ell \in \mathbf{Lab}\}) \cup \{(y, 1)\}$$

and clearly this has no consequence for whether or not $\overrightarrow{\text{RD}}$ is a solution. In other words we obtain a version of the previous equation system except that all equalities have been replaced by inclusions. Formally, whereas the equational approach demands that $\overrightarrow{\text{RD}} = F(\overrightarrow{\text{RD}})$, the constraint based approach demands that $\overrightarrow{\text{RD}} \supseteq F(\overrightarrow{\text{RD}})$ for the *same* function F . It is therefore immediate that a solution to the equation system is also a solution to the constraint system whereas the converse is not necessarily the case.

Luckily we can show that both the equation system and the constraint system have the same *least solution*. Recall that the least solution to $\overrightarrow{\text{RD}} = F(\overrightarrow{\text{RD}})$ is constructed as $F^n(\vec{\emptyset})$ for a value of n such that $F^n(\vec{\emptyset}) = F^{n+1}(\vec{\emptyset})$. If $\overrightarrow{\text{RD}}$ is a solution to the constraint system, that is $\overrightarrow{\text{RD}} \supseteq F(\overrightarrow{\text{RD}})$, then $\vec{\emptyset} \subseteq \overrightarrow{\text{RD}}$ is immediate and the monotonicity of F and mathematical induction then gives $F^n(\vec{\emptyset}) \subseteq \overrightarrow{\text{RD}}$. Since $F^n(\vec{\emptyset})$ is a solution to the constraint system this shows that it is also the least solution to the constraint system.

In summary, we have thus seen a very strong connection between the equational approach and the constraint based approach. This connection is not always as apparent as it is here: one of the characteristics of the constraint based approach is that often constraints with the same left hand side are generated at many different places in the program and therefore it may require serious work to collect them.

1.4 Control Flow Analysis

The purpose of *Control Flow Analysis* is to determine information about what “elementary blocks” may lead to what other “elementary blocks”. This information is immediately available for the WHILE language unlike what is the case for more advanced imperative, functional and object-oriented languages.

Consider the following functional program:

```

let  f = fn x => x 1;
     g = fn y => y+2;
     h = fn z => z+3
in   (f g) + (f h)

```

It defines a higher-order function f with formal parameter x and body $x\ 1$; then it defines two functions g and h that are given as actual parameters to f in the body of the `let`-construct. Semantically, x will be bound to each of these two functions in turn so both g and h will be applied to 1 and the result of the computation will be the value 7 .

An application of f will transfer control to the body of f , i.e. to $x\ 1$, and this application of x will transfer control to the body of x . The problem is that we cannot immediately point to the body of x : we need to know what parameters f will be called with. This is exactly the information that the Control Flow Analysis gives us:

For each function application, which functions may be applied.

As is typical of functional languages, the labelling scheme used would seem to have a very different character than the one employed for imperative languages because the “elementary blocks” may be nested. We shall therefore label *all* subexpressions as in the following simple program that will be used to illustrate the analysis.

Example 1.2 Consider the program:

$$[[\text{fn } x \Rightarrow [x]^1]^2 [\text{fn } y \Rightarrow [y]^3]^4]^5$$

It calls the identity function $\text{fn } x \Rightarrow x$ on the argument $\text{fn } y \Rightarrow y$ and clearly evaluates to $\text{fn } y \Rightarrow y$ itself (omitting all $[\dots]^\ell$). ■

We shall now be interested in associating information with the labels themselves, rather than with the entries and exits of the labels – thereby we exploit that there are no side-effects in our simple functional language. The Control Flow Analysis will be specified by a pair $(\widehat{C}, \widehat{\rho})$ of functions where $\widehat{C}(\ell)$ is supposed to contain the values that the subexpression (or “elementary block”) labelled ℓ may evaluate to and $\widehat{\rho}(x)$ contain the values that x can be bound to.

The constraint system. One way to specify the Control Flow Analysis then is by means of a collection of constraints and we shall illustrate this for the program of Example 1.2. There are three classes of constraints. One class of constraints relate the values of function abstractions to their labels:

$$\begin{aligned} \{\text{fn } x \Rightarrow [x]^1\} &\subseteq \widehat{C}(2) \\ \{\text{fn } y \Rightarrow [y]^3\} &\subseteq \widehat{C}(4) \end{aligned}$$

These constraints state that a function abstraction evaluates to a closure containing the abstraction itself. So the general pattern is that an occurrence of $[\text{fn } x \Rightarrow e]^\ell$ in the program gives rise to a constraint $\{\text{fn } x \Rightarrow e\} \subseteq \widehat{C}(\ell)$.

The second class of constraints relate the values of variables to their labels:

$$\begin{aligned}\widehat{\rho}(x) &\subseteq \widehat{C}(1) \\ \widehat{\rho}(y) &\subseteq \widehat{C}(3)\end{aligned}$$

The constraints state that a variable always evaluates to its value. So for each occurrence of $[x]^\ell$ in the program we will have a constraint $\widehat{\rho}(x) \subseteq \widehat{C}(\ell)$.

The third class of constraints concerns function application: for each application point $[e_1 e_2]^\ell$, and for each possible function $[\text{fn } x \Rightarrow e]^\ell$ that could be called at this point, we will have: (i) a constraint expressing that the formal parameter of the function is bound to the actual parameter at the application point, and (ii) a constraint expressing that the result obtained by evaluating the body of the function is a possible result of the application.

Our example program has just one application $[[\cdot\cdot]^\ell [\cdot\cdot]^\ell]^\ell$, but there are two candidates for the function, i.e. $\widehat{C}(2)$ is a subset of the set $\{\text{fn } x \Rightarrow [x]^\ell, \text{fn } y \Rightarrow [y]^\ell\}$. If the function $\text{fn } x \Rightarrow [x]^\ell$ is applied then the two constraints are $\widehat{C}(4) \subseteq \widehat{\rho}(x)$ and $\widehat{C}(1) \subseteq \widehat{C}(5)$. We express this as *conditional constraints*:

$$\begin{aligned}\{\text{fn } x \Rightarrow [x]^\ell\} \subseteq \widehat{C}(2) &\Rightarrow \widehat{C}(4) \subseteq \widehat{\rho}(x) \\ \{\text{fn } x \Rightarrow [x]^\ell\} \subseteq \widehat{C}(2) &\Rightarrow \widehat{C}(1) \subseteq \widehat{C}(5)\end{aligned}$$

Alternatively, the function being applied could be $\text{fn } y \Rightarrow [y]^\ell$ and the corresponding conditional constraints are:

$$\begin{aligned}\{\text{fn } y \Rightarrow [y]^\ell\} \subseteq \widehat{C}(2) &\Rightarrow \widehat{C}(4) \subseteq \widehat{\rho}(y) \\ \{\text{fn } y \Rightarrow [y]^\ell\} \subseteq \widehat{C}(2) &\Rightarrow \widehat{C}(3) \subseteq \widehat{C}(5)\end{aligned}$$

The least solution. As in Section 1.3 we shall be interested in the least solution to this set of constraints: the smaller the sets of values given by \widehat{C} and $\widehat{\rho}$, the more precise the analysis is in predicting which functions are applied. In Exercise 1.2 we show that the following choice of \widehat{C} and $\widehat{\rho}$ gives a solution to the above constraints:

$$\begin{aligned}\widehat{C}(1) &= \{\text{fn } y \Rightarrow [y]^\ell\} \\ \widehat{C}(2) &= \{\text{fn } x \Rightarrow [x]^\ell\} \\ \widehat{C}(3) &= \emptyset \\ \widehat{C}(4) &= \{\text{fn } y \Rightarrow [y]^\ell\} \\ \widehat{C}(5) &= \{\text{fn } y \Rightarrow [y]^\ell\} \\ \widehat{\rho}(x) &= \{\text{fn } y \Rightarrow [y]^\ell\} \\ \widehat{\rho}(y) &= \emptyset\end{aligned}$$

Among other things this tells us that the function abstraction $\text{fn } y \Rightarrow y$ is never applied (since $\hat{\rho}(y) = \emptyset$) and that the program may only evaluate to the function abstraction $\text{fn } y \Rightarrow y$ (since $\hat{C}(5) = \{\text{fn } y \Rightarrow [y]^3\}$).

Note the similarities between the constraint based approaches to Data Flow Analysis and Control Flow Analysis: in both cases the syntactic structure of the program gives rise to a set of constraints whose least solution is desired. The main difference is that the constraints for the Control Flow Analysis have a more complex structure than those for the Data Flow Analysis.

1.5 Abstract Interpretation

The theory of *Abstract Interpretation* is a general methodology for calculating analyses rather than just specifying them and then rely on a posteriori validation. To some extent the application of Abstract Interpretation is independent of the specification style used for presenting the program analysis and so applies not only to the Data Flow Analysis formulation to be used here.

Collecting semantics. As a preliminary step we shall formulate a so-called *collecting semantics* that records the set of *traces* tr that can reach a given program point:

$$tr \in \text{Trac}_{\bar{c}} = (\text{Var} \times \text{Lab})^*$$

Intuitively, a trace will record where the variables have obtained their values in the course of the computation. So for the factorial program

$[y:=x]^1; [z:=1]^2; \text{ while } [y>1]^3 \text{ do } ([z:=z*y]^4; [y:=y-1]^5); [y:=0]^6$

we will e.g. have the trace

$((x, ?), (y, ?), (z, ?), (y, 1), (z, 2), (z, 4), (y, 5), (z, 4), (y, 5), (y, 6))$

corresponding to a run of the program where the body of the while-loop is executed twice.

The traces contain sufficient information that we can extract a set of *semantically reaching definitions*:

$$\text{SRD}(tr)(x) = \ell \quad \text{iff} \quad \text{the rightmost pair } (x, \ell') \text{ in } tr \text{ has } \ell = \ell'$$

In order for the Reaching Definitions Analysis to be correct (or safe) we shall require that it captures the semantic reaching definitions, that is, if tr is a possible trace just before entering the elementary block labelled ℓ then we shall demand that

$$\forall x \in \text{Var} : (x, \text{SRD}(tr)(x)) \in \text{RD}_{\text{entry}}(\ell)$$

in order to trust the information in $RD_{entry}(\ell)$ about the set of definitions that may reach the entry to ℓ . In later chapters, we will conduct proofs of results like this.

The collecting semantics will specify a *superset* of the possible traces at the various program points. We shall specify the collecting semantics CS in the style of the Reaching Definitions Analysis in Section 1.3; more precisely, we shall specify a twelve-tuple of elements from $(\mathcal{P}(\mathbf{Trace}))^{12}$ by means of a set of equations. First we have

$$\begin{aligned} CS_{exit}(1) &= \{tr : (y, 1) \mid tr \in CS_{entry}(1)\} \\ CS_{exit}(2) &= \{tr : (z, 2) \mid tr \in CS_{entry}(2)\} \\ CS_{exit}(3) &= CS_{entry}(3) \\ CS_{exit}(4) &= \{tr : (z, 4) \mid tr \in CS_{entry}(4)\} \\ CS_{exit}(5) &= \{tr : (y, 5) \mid tr \in CS_{entry}(5)\} \\ CS_{exit}(6) &= \{tr : (y, 6) \mid tr \in CS_{entry}(6)\} \end{aligned}$$

showing how the assignment statements give rise to extensions of the traces. Here we write $tr : (x, \ell)$ for appending an element (x, ℓ) to a list tr , that is $((x_1, \ell_1), \dots, (x_n, \ell_n)) : (x, \ell)$ equals $((x_1, \ell_1), \dots, (x_n, \ell_n), (x, \ell))$. Furthermore, we have

$$\begin{aligned} CS_{entry}(2) &= CS_{exit}(1) \\ CS_{entry}(3) &= CS_{exit}(2) \cup CS_{exit}(5) \\ CS_{entry}(4) &= CS_{exit}(3) \\ CS_{entry}(5) &= CS_{exit}(4) \\ CS_{entry}(6) &= CS_{exit}(3) \end{aligned}$$

corresponding to the flow of control in the program; more detailed information about the value of the variables would allow us to define the sets $CS_{entry}(4)$ and $CS_{entry}(6)$ more precisely but the above definitions are sufficient for illustrating the approach. Finally, we take

$$CS_{entry}(1) = \{(x, ?), (y, ?), (z, ?)\}$$

corresponding to the fact that all variables are uninitialised in the beginning. In the manner of the previous sections we can rewrite the above system of equations in the form

$$\vec{CS} = G(\vec{CS})$$

where \vec{CS} is a twelve-tuple of elements from $(\mathcal{P}(\mathbf{Trace}))^{12}$ and where G is a monotone function of functionality:

$$G : (\mathcal{P}(\mathbf{Trace}))^{12} \rightarrow (\mathcal{P}(\mathbf{Trace}))^{12}$$




Figure 1.3: The adjunction (α, γ) .

As is explained in Appendix A there is a body of general theory that ensures that the equation in fact has a least solution; we shall write it as $\text{lfp}(G)$. However, since $(\mathcal{P}(\text{Trace}))^{12}$ is not finite we cannot simply use the methods of the previous sections in order to construct $\text{lfp}(G)$.

Galois connections. As we have seen the collecting semantics operates on sets of traces whereas the Reaching Definitions Analysis operates on sets of pairs of variables and labels. To relate these “worlds” we define an abstraction function α and a concretisation function γ as illustrated in:

$$\mathcal{P}(\text{Trace}) \begin{array}{c} \xleftarrow{\gamma} \\ \xrightarrow{\alpha} \end{array} \mathcal{P}(\text{Var} \times \text{Lab})$$

The idea is that the *abstraction function* α extracts the reachability information present in a set of traces; it is natural to define

$$\alpha(X) = \{(x, \text{SRD}(tr)(x)) \mid x \in \text{Var} \wedge tr \in X\}$$

where we exploit the notion of semantically reaching definitions.

The *concretisation function* γ then produces all traces tr that are consistent with the given reachability information:

$$\gamma(Y) = \{tr \mid \forall x \in \text{Var} : (x, \text{SRD}(tr)(x)) \in Y\}$$

Often it is demanded that α and γ satisfy the condition

$$\alpha(X) \subseteq Y \Leftrightarrow X \subseteq \gamma(Y)$$

and we shall say that (α, γ) is an *adjunction*, or a *Galois connection*, whenever this condition is satisfied; this is illustrated on Figure 1.3. We shall leave it to the reader to verify that (α, γ) as defined above does in fact fulfil this condition.

Induced analysis. We shall now show how the collecting semantics can be used to *calculate* (as opposed to “guess”) an analysis like the one in Section 1.3; we shall say that the analysis is an *induced analysis*. For this we define

$$\begin{aligned}\bar{\alpha}(X_1, \dots, X_{12}) &= (\alpha(X_1), \dots, \alpha(X_{12})) \\ \bar{\gamma}(Y_1, \dots, Y_{12}) &= (\gamma(Y_1), \dots, \gamma(Y_{12}))\end{aligned}$$

where α and γ are as above and we consider the function $\bar{\alpha} \circ G \circ \bar{\gamma}$ of functionality:

$$(\bar{\alpha} \circ G \circ \bar{\gamma}) : (\mathcal{P}(\mathbf{Var} \times \mathbf{Lab}))^{12} \rightarrow (\mathcal{P}(\mathbf{Var} \times \mathbf{Lab}))^{12}$$

This function defines a Reaching Definitions analysis in an indirect way. Since G is specified by a set of equations (over $\mathcal{P}(\mathbf{Trace})$) we can use $\bar{\alpha} \circ G \circ \bar{\gamma}$ to calculate a new set of equations (over $\mathcal{P}(\mathbf{Var} \times \mathbf{Lab})$). We shall illustrate this for one of the equations:

$$CS_{exit}(4) = \{tr : (z, 4) \mid tr \in CS_{entry}(4)\}$$

The corresponding clause in the definition of G is:

$$G_{exit}(4)(\dots, CS_{entry}(4), \dots) = \{tr : (z, 4) \mid tr \in CS_{entry}(4)\}$$

We can now calculate the corresponding clause in the definition of $\bar{\alpha} \circ G \circ \bar{\gamma}$:

$$\begin{aligned}\alpha(G_{exit}(4)(\bar{\gamma}(\dots, RD_{entry}(4), \dots))) \\ &= \alpha(\{tr : (z, 4) \mid tr \in \gamma(RD_{entry}(4))\}) \\ &= \{(x, SRD(tr : (z, 4))(x)) \\ &\quad \mid x \in \mathbf{Var}, \forall y \in \mathbf{Var} : (y, SRD(tr)(y)) \in RD_{entry}(4)\} \\ &= \{(x, SRD(tr)(x)) \\ &\quad \mid x \in \mathbf{Var}, x \neq z, (x, SRD(tr)(x)) \in RD_{entry}(4)\} \cup \{(z, 4)\} \\ &= (RD_{entry}(4) \setminus \{(z, \ell) \mid \ell \in \mathbf{Lab}\}) \cup \{(z, 4)\}\end{aligned}$$

This clause is equivalent to the equation for $RD_{exit}(4)$ in Section 1.3; similar calculations can be performed for the other clauses.

The least solution. As explained in Appendix A the equation system

$$\vec{RD} = (\bar{\alpha} \circ G \circ \bar{\gamma})(\vec{RD})$$

has a least solution; we shall write it as $lfp(\bar{\alpha} \circ G \circ \bar{\gamma})$. It is interesting to note that if one replaces the infinite sets \mathbf{Var} and \mathbf{Lab} with finite sets \mathbf{Var}_* and \mathbf{Lab}_* as before, then the least fixed point of $\bar{\alpha} \circ G \circ \bar{\gamma}$ can be obtained as $(\bar{\alpha} \circ G \circ \bar{\gamma})^n(\vec{\emptyset})$ just as was the case for F previously.

In Exercise 1.4 we shall show that $\bar{\alpha} \circ G \circ \bar{\gamma} \sqsubseteq F$ and it follows that

$$\bar{\alpha}(lfp(G)) \sqsubseteq lfp(\bar{\alpha} \circ G \circ \bar{\gamma}) \sqsubseteq lfp(F)$$

where the second inclusion should not be surprising because mathematical induction suffices for showing that $(\bar{\alpha} \circ G \circ \bar{\gamma})^n(\bar{\theta}) \subseteq F^n(\bar{\theta})$. The above inclusions just say that the least solution to the equation system defined by $\bar{\alpha} \circ G \circ \bar{\gamma}$ is correct with respect to the collecting semantics, and similarly that the least solution to the equation system of Section 1.3 is also correct with respect to the collecting semantics. Thus it follows that we will only need to show that the collecting semantics is correct – the correctness of the induced analysis will follow for free.

For some analyses one is able to prove the stronger result $\bar{\alpha} \circ G \circ \bar{\gamma} = F$. Then the analysis is *optimal* (given the choice of approximate properties it operates on) and clearly $\text{lfp}(\bar{\alpha} \circ G \circ \bar{\gamma}) = \text{lfp}(F)$. In Exercise 1.4 we shall study whether or not this is the case here.

1.6 Type and Effect Systems

A simple type system. The ideal setting for explaining *Type and Effect Systems* is to consider a typed functional or imperative language. However, even our simple toy language can be considered to be typed: a statement S maps a state to a state (in case it terminates) and may therefore be considered to have type $\Sigma \rightarrow \Sigma$ where Σ denotes the type of states; we write this as the judgement:

$$S : \Sigma \rightarrow \Sigma$$

One way to formalise this is by the following utterly trivial system of axioms and inference rules:

$$[x := a]^\ell : \Sigma \rightarrow \Sigma$$

$$[\text{skip}]^\ell : \Sigma \rightarrow \Sigma$$

$$\frac{S_1 : \Sigma \rightarrow \Sigma \quad S_2 : \Sigma \rightarrow \Sigma}{S_1; S_2 : \Sigma \rightarrow \Sigma}$$

$$\frac{S_1 : \Sigma \rightarrow \Sigma \quad S_2 : \Sigma \rightarrow \Sigma}{\text{if } [b]^\ell \text{ then } S_1 \text{ else } S_2 : \Sigma \rightarrow \Sigma}$$

$$\frac{S : \Sigma \rightarrow \Sigma}{\text{while } [b]^\ell \text{ do } S : \Sigma \rightarrow \Sigma}$$

Often a Type and Effect System can be viewed as the amalgamation of two ingredients: an *Effect System* and an *Annotated Type System*. In an Effect System we typically have judgements of the form $S : \Sigma \xrightarrow{\varphi} \Sigma$ where the *effect* φ tells something about what happens when S is executed: this may for example be which errors might occur, which exceptions might be raised, or which files might be modified. In an Annotated Type System we typically

[<i>ass</i>]	$[x := a]^{\ell'} : \text{RD} \rightarrow ((\text{RD} \setminus \{(x, \ell) \mid \ell \in \text{Lab}\}) \cup \{(x, \ell')\})$
[<i>skip</i>]	$[\text{skip}]^{\ell'} : \text{RD} \rightarrow \text{RD}$
[<i>seq</i>]	$\frac{S_1 : \text{RD}_1 \rightarrow \text{RD}_2 \quad S_2 : \text{RD}_2 \rightarrow \text{RD}_3}{S_1; S_2 : \text{RD}_1 \rightarrow \text{RD}_3}$
[<i>if</i>]	$\frac{S_1 : \text{RD}_1 \rightarrow \text{RD}_2 \quad S_2 : \text{RD}_1 \rightarrow \text{RD}_2}{\text{if } [b]^{\ell} \text{ then } S_1 \text{ else } S_2 : \text{RD}_1 \rightarrow \text{RD}_2}$
[<i>wh</i>]	$\frac{S : \text{RD} \rightarrow \text{RD}}{\text{while } [b]^{\ell} \text{ do } S : \text{RD} \rightarrow \text{RD}}$
[<i>sub</i>]	$\frac{S : \text{RD}_2 \rightarrow \text{RD}_3}{S : \text{RD}_1 \rightarrow \text{RD}_4} \text{ if } \text{RD}_1 \subseteq \text{RD}_2 \text{ and } \text{RD}_3 \subseteq \text{RD}_4$

Table 1.2: Reaching Definitions: annotated base types.

have judgements of the form $S : \Sigma_1 \rightarrow \Sigma_2$ where the Σ_i describe certain *properties of states*: this may for example be that a variable is positive or that a certain invariant is maintained. We shall first illustrate the latter approach for the WHILE language and then illustrate the Effect Systems using the functional language.

1.6.1 Annotated Type Systems

Annotated base types. To obtain our first specification of Reaching Definitions we shall focus on a formulation where the *base types* are annotated. Here we will have judgements of the form

$$S : \text{RD}_1 \rightarrow \text{RD}_2$$

where $\text{RD}_1, \text{RD}_2 \in \mathcal{P}(\text{Var} \times \text{Lab})$ are sets of reaching definitions. Based on the trivial axioms and rules displayed above we then obtain the more interesting ones in Table 1.2.

To explain these rules let us first explain the meaning of $S : \text{RD}_1 \rightarrow \text{RD}_2$ in terms of the developments performed in Section 1.3. For this we first observe that any statement S will have one elementary block at its entry, denoted $\text{init}(S)$, and one or more elementary blocks at its exit, denoted $\text{final}(S)$; for a statement like $\text{if } [x < y]^1 \text{ then } [x := y]^2 \text{ else } [y := x]^3$ we thus get $\text{init}(\dots) = 1$ and $\text{final}(\dots) = \{2, 3\}$.

Our first (and not quite successful) attempt at explaining the meaning of

$S : RD_1 \rightarrow RL_2$ then is to say that:

$$RD_1 = RD_{entry}(init(S)) \\ \cup \{RD_{exit}(\ell) \mid \ell \in final(S)\} = RD_2$$

This suffices for explaining the axioms for assignment and skip: here the formulae after the arrows correspond exactly to the similar equations in the equational formulation of the analysis in Section 1.3. Also the rule for sequencing now seems rather natural. However, the rule for conditional is more dubious: considering the statement `if [x<y]1 then [x:=y]2 else [y:=x]3` once more, it seems impossible to achieve that the then-branch gives rise to the same set of reaching definitions as the else-branch does.

Our second (and successful) attempt at explaining the intended meaning of $S : RD_1 \rightarrow RD_2$ then is to say that:

$$RD_1 \subseteq RD_{entry}(init(S)) \\ \forall \ell \in final(S) : RD_{exit}(\ell) \subseteq RD_2$$

This formulation is somewhat closer to the development in the constraint based formulation of the analysis in Section 1.3 and it explains why the last rule, called a *subsumption rule*, is unproblematic. Actually, the subsumption rule will solve our problem with the conditional because even when the then-branch gives a different set of reaching definitions than the else-branch we can enlarge both results to a common set of reaching definitions. Finally, consider the rule for the iterative construct. Here we simply express that RD is a consistent guess concerning what may reach the entry and exits of S – this expresses a fixed point property.

Example 1.3 To analyse the factorial program

$$[y:=x]^1; [z:=1]^2; \text{ while } [y>1]^3 \text{ do } ([z:=z*y]^4; [y:=y-1]^5); [y:=0]^6$$

of Example 1.1 we will proceed as follows. We shall write RD_f for the set $\{(x, ?), (y, 1), (y, 5), (z, 2), (z, 4)\}$ and consider the body of the while-loop. The axiom `[ass]` gives

$$[z:=z*y]^4: RD_f \rightarrow \{(x, ?), (y, 1), (y, 5), (z, 4)\} \\ [y:=y-1]^5 : \{(x, ?), (y, 1), (y, 5), (z, 4)\} \rightarrow \{(x, ?), (y, 5), (z, 4)\}$$

so the rule `[seq]` gives:

$$([z:=z*y]^4; [y:=y-1]^5): RD_f \rightarrow \{(x, ?), (y, 5), (z, 4)\}$$

Now $\{(x, ?), (y, 5), (z, 4)\} \subseteq RD_f$ so the subsumption rule gives:

$$([z:=z*y]^4; [y:=y-1]^5): RD_f \rightarrow RD_f$$

$[ass]$	$[x := a]^{\ell} : \Sigma \xrightarrow{\{(x, \ell)\}} \Sigma$
$[skip]$	$[skip]^{\ell} : \Sigma \xrightarrow{\emptyset} \Sigma$
$[seq]$	$\frac{S_1 : \Sigma \xrightarrow{X_1} \Sigma \quad S_2 : \Sigma \xrightarrow{X_2} \Sigma}{S_1; S_2 : \Sigma \xrightarrow{(RD_1 \setminus X_2) \cup X_2} \Sigma}$
$[if]$	$\frac{S_1 : \Sigma \xrightarrow{X_1} \Sigma \quad S_2 : \Sigma \xrightarrow{X_2} \Sigma}{\text{if } [b]^{\ell} \text{ then } S_1 \text{ else } S_2 : \Sigma \xrightarrow{X_1 \cap X_2} \Sigma}$
$[wh]$	$\frac{S : \Sigma \xrightarrow{X} \Sigma}{\text{while } [b]^{\ell} \text{ do } S : \Sigma \xrightarrow{\emptyset} \Sigma}$
$[sub]$	$\frac{S : \Sigma \xrightarrow{X} \Sigma}{S : \Sigma \xrightarrow{X'} \Sigma} \quad \text{if } X' \subseteq X \text{ and } RD \subseteq RD'$

Table 1.3: Reaching Definitions: annotated type constructors.

We can now apply the rule $[wh]$ and get:

$$\text{while } [y > 1]^3 \text{ do } ([z := z * y]^4; [y := y - 1]^5) : RD_f \rightarrow RD_f$$

Using the axiom $[ass]$ we get:

$$\begin{aligned} [y := x]^1 &: \{(x, ?), (y, ?), (z, ?)\} \rightarrow \{(x, ?), (y, 1), (z, ?)\} \\ [z := 1]^2 &: \{(x, ?), (y, 1), (z, ?)\} \rightarrow \{(x, ?), (y, 1), (z, 2)\} \\ [y := 0]^6 &: RD_f \rightarrow \{(x, ?), (y, 6), (z, 2), (z, 4)\} \end{aligned}$$

Since $\{(x, ?), (y, 1), (z, 2)\} \subseteq RD_f$ we can apply the rules $[seq]$ and $[sub]$ to get

$$\begin{aligned} ([y := x]^1; [z := 1]^2; \text{while } [y > 1]^3 \text{ do } ([z := z * y]^4; [y := y - 1]^5); [y := 0]^6) : \\ \{(x, ?), (y, ?), (z, ?)\} \rightarrow \{(x, ?), (y, 6), (z, 2), (z, 4)\} \end{aligned}$$

corresponding to the result in Table 1.1. ■

The system in Table 1.2 suffices for manually analysing a given program. To obtain an implementation it will be natural to extract a set of constraints similar to those considered in Section 1.3, and then solve them in the same way as before. This will be the idea behind the approach taken in Chapter 5.

Annotated type constructors. Another approach to Reaching Definitions has a little bit of the flavour of Effect Systems in that it is the type constructors (arrow in our case) that are annotated. Here we will have judgements of the form

$$S : \Sigma \xrightarrow[\text{RD}]{X} \Sigma$$

where X denotes the set of variables that *definitely* will be assigned in S and RD denotes the set of reaching definitions that S *might* produce. The axioms and rules are shown in Table 1.3 and are explained below.

The axiom for assignment simply expresses that the variable x definitely will be assigned and that the reaching definition (x, ℓ) is produced. In the rule for sequencing the notation $\text{RD} \setminus X$ means $\{(x, \ell) \in \text{RD} \mid x \notin X\}$. The rule expresses that we take the union of the reaching definitions after having removed entries from S_1 that are definitely redefined in S_2 . Also we take the union of the two sets of assigned variables. In the rule for conditional we take the union of information about reaching definitions whereas we take the intersection (rather than the union) of the assigned variables because we are not completely sure what path was taken through the conditional. A similar comment holds for the rule for the `while`-loop; here we can think of \emptyset as the intersection between \emptyset (when the body is not executed) and X .

We have included a *subsumption rule* because this is normally the case for such systems as we shall see in Chapter 5. However, in the system above there is little need for it, and if one excludes it then implementation becomes very straightforward: simply perform a syntax directed traversal of the program where the sets X and RD are computed for each subprogram.

Example 1.4 Let us once again consider the analysis of the factorial program

$$[y:=x]^1; [z:=1]^2; \text{while } [y>1]^3 \text{ do } ([z:=z*y]^4; [y:=y-1]^5); [y:=0]^6$$

For the body of the `while`-loop we get

$$\begin{aligned} [z:=z*y]^4: \Sigma &\xrightarrow[\{(z,4)\}]{\{z\}} \Sigma \\ [y:=y-1]^5: \Sigma &\xrightarrow[\{(y,5)\}]{\{y\}} \Sigma \end{aligned}$$

so the rule `[seq]` gives:

$$([z:=z*y]^4; [y:=y-1]^5): \Sigma \xrightarrow[\{(y,5),(z,4)\}]{\{y,z\}} \Sigma$$

We can now apply the rule `[wh]` and get:

$$\text{while } [y>1]^3 \text{ do } ([z:=z*y]^4; [y:=y-1]^5): \Sigma \xrightarrow[\{(y,5),(z,4)\}]{\emptyset} \Sigma$$

In a similar way we get

$$([y:=x]^1; [z:=1]^2): \Sigma \frac{\{y,z\}}{\{(y,1),(z,2)\}} \rightarrow \Sigma$$

$$[y:=0]^6: \Sigma \frac{\{y\}}{\{(y,6)\}} \rightarrow \Sigma$$

so using the rule *[seq]* we obtain

$$([y:=x]^1; [z:=1]^2; \text{while } [y>1]^3 \text{ do } ([z:=z*y]^4; [y:=y-1]^5); [y:=0]^6): \\ \Sigma \frac{\{y,z\}}{\{(y,6),(z,2),(z,4)\}} \rightarrow \Sigma$$

showing that the program definitely will assign to y and z and that the final value of y will be assigned at 6 and the final value of z at 2 or 4. ■

Compared with the previous specifications of Reaching Definitions analysis the flavour of Table 1.3 is rather different: the analysis of a statement expresses how information present at the entry will be *modified* by the statement – we may therefore view the specification as a higher-order formulation of Reaching Definitions analysis.

1.6.2 Effect Systems

A simple type system. To give the flavour of Effect Systems let us once more turn to the functional language. As above, the idea is to annotate a traditional type system with analysis information, so let us start by presenting a simple type system for a language with variables x , function abstraction $\text{fn}_\pi x \Rightarrow e$ (where π is the name of the abstraction), and function application $e_1 e_2$. The judgements have the form

$$\Gamma \vdash e : \tau$$

where Γ is a *type environment* that gives types to all free variables of e and τ is the *type* of e . For simplicity we shall assume that types are either base types as `int` and `bool` or they are function types written $\tau_1 \rightarrow \tau_2$. The type system is given by the following axioms and rules:

$$\Gamma \vdash x : \tau_x \quad \text{if } \Gamma(x) = \tau_x$$

$$\frac{\Gamma[x \mapsto \tau_x] \vdash e : \tau}{\Gamma \vdash \text{fn}_\pi x \Rightarrow e : \tau_x \rightarrow \tau}$$

$$\frac{\Gamma \vdash e_1 : \tau_2 \rightarrow \tau, \Gamma \vdash e_2 : \tau_2}{\Gamma \vdash e_1 e_2 : \tau}$$

So the axiom for variables just expresses that the type of x is obtained from the assumptions of the type environment. The rule for function abstraction requires that we “guess” a type τ_x for the formal parameter x and we

determine the type of the body of the abstraction under that additional assumption. The rule for function application requires that we determine the type of the operator as well as the argument and it implicitly expresses that the operator must have a function type by requiring the type of e_1 to have the form $\tau_2 \rightarrow \tau$. Furthermore the two occurrences of τ_2 in the rule implicitly express that the type of the actual parameter must equal the type expected by the formal parameter of the function.

Example 1.5 Consider the following version of the program of Example 1.2

$$(\text{fn}_X x \Rightarrow x) (\text{fn}_Y y \Rightarrow y)$$

where we now have given $\text{fn } x \Rightarrow x$ the name X and $\text{fn } y \Rightarrow y$ the name Y . To see that this program has type $\text{int} \rightarrow \text{int}$ we first observe that $[y \mapsto \text{int}] \vdash y : \text{int}$ so:

$$[] \vdash \text{fn}_Y y \Rightarrow y : \text{int} \rightarrow \text{int}$$

Similarly, we have $[x \mapsto \text{int} \rightarrow \text{int}] \vdash x : \text{int} \rightarrow \text{int}$ so:

$$[] \vdash \text{fn}_X x \Rightarrow x : (\text{int} \rightarrow \text{int}) \rightarrow (\text{int} \rightarrow \text{int})$$

The rule for application then gives:

$$[] \vdash (\text{fn}_X x \Rightarrow x) (\text{fn}_Y y \Rightarrow y) : \text{int} \rightarrow \text{int} \quad \blacksquare$$

Effects. The analysis we shall consider is a *Call-tracking Analysis*:

For each subexpression, which function abstractions may be applied during its evaluation.

The set of function names constitutes the *effect* of the subexpression. To determine this information we shall annotate the function types with their *latent effect* so we shall e.g. write $\text{int} \xrightarrow{\{X\}} \text{int}$ for the type of a function mapping integers to integers and with effect $\{X\}$ meaning that when executing the function it may apply the function named X . More generally, the annotated types $\hat{\tau}$ will either be base types or they will have the form

$$\hat{\tau}_1 \xrightarrow{\varphi} \hat{\tau}_2$$

where φ is the effect, i.e. the names of the function abstractions that we might apply when applying a function of this type.

We specify the analysis using judgements of the form

$$\hat{\Gamma} \vdash e : \hat{\tau} \& \varphi$$

$[var]$	$\widehat{\Gamma} \vdash x : \widehat{\tau}_x \ \& \ \emptyset \text{ if } \widehat{\Gamma}(x) = \widehat{\tau}_x$
$[fn]$	$\frac{\widehat{\Gamma}[x \mapsto \widehat{\tau}_x] \vdash e : \widehat{\tau} \ \& \ \varphi}{\widehat{\Gamma} \vdash \text{fn}_\pi x \Rightarrow e : \widehat{\tau}_x \xrightarrow{\varphi \cup \{\pi\}} \widehat{\tau} \ \& \ \emptyset}$
$[app]$	$\frac{\widehat{\Gamma} \vdash e_1 : \widehat{\tau}_2 \xrightarrow{\varphi} \widehat{\tau} \ \& \ \varphi_1 \quad \widehat{\Gamma} \vdash e_2 : \widehat{\tau}_2 \ \& \ \varphi_2}{\widehat{\Gamma} \vdash e_1 e_2 : \widehat{\tau} \ \& \ \varphi_1 \cup \varphi_2 \cup \varphi}$

Table 1.4: Call-tracking Analysis: Effect System.

where $\widehat{\Gamma}$ is the type environment that now gives the annotated type of all free variables, $\widehat{\tau}$ is the annotated type of e , and φ is the effect of evaluating e . The analysis is specified by the axioms and rules in Table 1.4 which will be explained below.

In the axiom $[var]$ for variables we produce an empty effect because we assume that the parameter mechanism is call-by-value so no evaluation takes place when mentioning a variable. Similarly, in the rule $[fn]$ for function abstractions we produce an empty effect: no evaluation takes place because we only construct a closure. The body of the abstraction is analysed in order to determine its annotated type and effect. This information is needed to annotate the function arrow: all the names of functions in the effect of the body and the name of the abstraction itself may be involved when this particular abstraction is applied.

Next, consider the rule $[app]$ for function application $e_1 e_2$. Here we obtain annotated types and effects from the operator e_1 as well as the operand e_2 . The effect of the application will contain the effect φ_1 of the operator (because we have to evaluate it before the application can take place), the effect φ_2 of the operand (because we employ a call-by-value semantics so this expression has to be evaluated too) and finally we need the effect φ of the function being applied. But this is exactly the information given by the annotation of the arrow in the type $\widehat{\tau}_2 \xrightarrow{\varphi} \widehat{\tau}$ of the operand. Hence we produce the union of these three sets as the overall effect of the application.

Example 1.6 Returning to the program of Example 1.5 we have:

$$\begin{aligned}
[] \vdash \text{fn}_Y y \Rightarrow y : \text{int} \xrightarrow{\{Y\}} \text{int} \ \& \ \emptyset \\
[] \vdash \text{fn}_X x \Rightarrow x : (\text{int} \xrightarrow{\{Y\}} \text{int}) \xrightarrow{\{X\}} (\text{int} \xrightarrow{\{Y\}} \text{int}) \ \& \ \emptyset \\
[] \vdash (\text{fn}_X x \Rightarrow x) (\text{fn}_Y y \Rightarrow y) : \text{int} \xrightarrow{\{Y\}} \text{int} \ \& \ \{X\}
\end{aligned}$$

This shows that our example program may (in fact it will) apply the function $\text{fn } x \Rightarrow x$ but that it will not apply the function $\text{fn } y \Rightarrow y$. ■

INPUT:	Example equations for Reaching Definitions
OUTPUT:	The least solution: \vec{RD}
METHOD:	Step 1: Initialisation $RD_1 := \emptyset; \dots; RD_{12} := \emptyset$ Step 2: Iteration while $RD_j \neq F_j(RD_1, \dots, RD_{12})$ for some j do $RD_j := F_j(RD_1, \dots, RD_{12})$

Table 1.5: Chaotic Iteration for Reaching Definitions.

For a more general language we will also need to introduce some form of subsumption rule in the manner of Tables 1.2 and 1.3; there are different approaches to this and we shall return to this later. Effect Systems are often implemented as extension of type inference algorithms and depending on the form of the effects it may be possible to calculate them on the fly; alternatively, sets of constraints can be generated and solved subsequently. We refer to Chapter 5 for more details.

1.7 Algorithms

Let us now reconsider the problem of computing the least solution to the program analysis problems considered in Data Flow Analysis and Control Flow Analysis.

Recall from Section 1.3 that we consider twelve-tuples $\vec{RD} \in (\mathcal{P}(\mathbf{Var}_* \times \mathbf{Lab}_*))^{12}$ of pairs of variables and labels where each label indicates an elementary block in which the corresponding variable was last assigned. The equation or constraint system gives rise to demanding the least solution to an equation $\vec{RD} = F(\vec{RD})$ or inclusion $\vec{RD} \sqsupseteq F(\vec{RD})$ where F is a monotone function over $(\mathcal{P}(\mathbf{Var}_* \times \mathbf{Lab}_*))^{12}$. Due to the finiteness of $(\mathcal{P}(\mathbf{Var}_* \times \mathbf{Lab}_*))^{12}$ the desired solution is in both cases obtainable as $F^n(\vec{\emptyset})$ for any n such that $F^{n+1}(\vec{\emptyset}) = F^n(\vec{\emptyset})$ and we know that such n does in fact exist.

Chaotic Iteration. Naively implementing the above procedure soon turns out to require an overwhelming amount of work. In later chapters we shall see much more efficient algorithms and in this section we shall illustrate the principle of *Chaotic Iteration* that lies at the heart of many of them. For this let us write

$$\vec{RD} = (RD_1, \dots, RD_{12})$$

$$F(\vec{RD}) = (F_1(\vec{RD}), \dots, F_{12}(\vec{RD}))$$

and consider the non-deterministic algorithm in Table 1.5. It is immediate that there exists j such that $RD_j \neq F_j(RD_1, \dots, RD_{12})$ if and only if $\vec{RD} \neq F(\vec{RD})$. Hence if the algorithm terminates it will produce a fixed point of F ; that is, a solution to the desired equations or constraints.

Properties of the algorithm. To further analyse the algorithm we shall exploit that

$$\vec{\emptyset} \subseteq \vec{RD} \subseteq F(\vec{RD}) \subseteq F^n(\vec{\emptyset})$$

holds at all points in the algorithm (where n is determined by $F^{n+1}(\vec{\emptyset}) = F^n(\vec{\emptyset})$): it clearly holds initially and as will be shown in Exercise 1.6 it is maintained during iteration. This means that if the algorithm terminates we will have obtained not only a fixed point of F but in fact the least fixed point (i.e. $F^n(\vec{\emptyset})$).

To see that the algorithm terminates note that if j satisfies

$$RD_j \neq F_j(RD_1, \dots, RD_{12})$$

then in fact $RD_j \subset F_j(RD_1, \dots, RD_{12})$ and hence the size of \vec{RD} increases by at least one as we perform each iteration. This ensures termination since we assumed that $(\mathcal{P}(\text{Var}_* \times \text{Lab}_*))^{12}$ is finite.

The above algorithm is suitable for *manually* solving data flow equations and constraint systems. To obtain an algorithm that is suitable for implementation we need to give more details about the choice of j so as to avoid an extensive search for the value; we shall return to this in Chapters 2 and 6.

1.8 Transformations

A major application of program analysis is to transform the program (at the source level or at some intermediate level inside a compiler) so as to obtain better performance. To illustrate the ideas we shall show how Reaching Definitions can be used to perform a transformation known as *Constant Folding*. There are two ingredients in this. One is to replace the use of a variable in some expression by a constant if it is known that the value of the variable will always be that constant. The other is to simplify an expression by partially evaluating it: subexpressions that contain no variables can be evaluated.

Source to source transformation. Consider a program S_* and let RD be a solution (preferable the least) to the Reaching Definitions Analysis for S_* . For a sub-statement S of S_* we shall now describe how to transform it into a “better” statement S' . We do so by means of judgements of the form

$$RD \vdash S \triangleright S'$$

[<i>ass</i> ₁]	RD ⊢ [x := a] ^ℓ ▷ [x := a[y ↦ n]] ^ℓ
	if $\int y \in FV(a) \wedge (y, ?) \notin RD_{entry}(\ell) \wedge$ $\{ \forall (z, \ell') \in RD_{entry}(\ell) : (z = y \Rightarrow [\dots]^{\ell'} \text{ is } [y := n]^{\ell'}) \}$
[<i>ass</i> ₂]	RD ⊢ [x := a] ^ℓ ▷ [x := n] ^ℓ
	if $FV(a) = \emptyset \wedge a \notin \mathbf{Num} \wedge a$ evaluates to n
[<i>seq</i> ₁]	$\frac{RD \vdash S_1 \triangleright S'_1}{RD \vdash S_1; S_2 \triangleright S'_1; S_2}$
[<i>seq</i> ₂]	$\frac{RD \vdash S_2 \triangleright S'_2}{RD \vdash S_1; S_2 \triangleright S_1; S'_2}$
[<i>if</i> ₁]	$\frac{RD \vdash S_1 \triangleright S'_1}{RD \vdash \text{if } [b]^{\ell} \text{ then } S_1 \text{ else } S_2 \triangleright \text{if } [b]^{\ell} \text{ then } S'_1 \text{ else } S_2}$
[<i>if</i> ₂]	$\frac{RD \vdash S_2 \triangleright S'_2}{RD \vdash \text{if } [b]^{\ell} \text{ then } S_1 \text{ else } S_2 \triangleright \text{if } [b]^{\ell} \text{ then } S_1 \text{ else } S'_2}$
[<i>wh</i>]	$\frac{RD \vdash S \triangleright S'}{RD \vdash \text{while } [b]^{\ell} \text{ do } S \triangleright \text{while } [b]^{\ell} \text{ do } S'}$

Table 1.6: Constant Folding transformation.

expressing *one step* of the transformation process. We may define the transformation using the axioms and rules in Table 1.6; they are explained below.

The first axiom [*ass*₁] expresses the first ingredient in Constant Folding as explained above – the use of a variable can be replaced with a constant if it is known that the variable always will be that constant; here we write $a[y \mapsto n]$ for the expression that is as a except that all occurrences of y have been replaced by n ; also we write $FV(a)$ for the set of variables occurring in a .

The second axiom [*ass*₂] expresses the second ingredient of the transformation – expressions can be partially evaluated; it uses the fact that if an expression contains no variables then it will always evaluate to the same value.

The five rules in Table 1.6 simply say that if we can transform a sub-statement then we can transform the statement itself. Note that the rules (e.g. [*seq*₁] and [*seq*₂]) do not prescribe a specific transformation order and hence many different transformation sequences may exist. Also note that the relation $RD \vdash \cdot \triangleright \cdot$ is neither reflexive nor transitive because there are no rules that forces it to be so. Hence we shall often want to perform an entire sequence of transformations.

Example 1.7 To illustrate the use of the transformation consider the program:

$$[x:=10]^1; [y:=x+10]^2; [z:=y+10]^3$$

The least solution to the Reaching Definitions Analysis for this program is:

$$\begin{aligned} RD_{entry}(1) &= \{(x, ?), (y, ?), (z, ?)\} \\ RD_{exit}(1) &= \{(x, 1), (y, ?), (z, ?)\} \\ RD_{entry}(2) &= \{(x, 1), (y, ?), (z, ?)\} \\ RD_{exit}(2) &= \{(x, 1), (y, 2), (z, ?)\} \\ RD_{entry}(3) &= \{(x, 1), (y, 2), (z, ?)\} \\ RD_{exit}(3) &= \{(x, 1), (y, 2), (z, 3)\} \end{aligned}$$

Let us now see how to transform the program. From the axiom $[ass_1]$ we have

$$RD \vdash [y:=x+10]^2 \triangleright [y:=10+10]^2$$

and therefore the rules for sequencing gives:

$$RD \vdash [x:=10]^1; [y:=x+10]^2; [z:=y+10]^3 \triangleright [x:=10]^1; [y:=10+10]^2; [z:=y+10]^3$$

We can now continue and obtain the following transformation sequence:

$$\begin{aligned} RD \quad \vdash \quad & [x:=10]^1; [y:=x+10]^2; [z:=y+10]^3 \\ \triangleright \quad & [x:=10]^1; [y:=10+10]^2; [z:=y+10]^3 \\ \triangleright \quad & [x:=10]^1; [y:=20]^2; [z:=y+10]^3 \\ \triangleright \quad & [x:=10]^1; [y:=20]^2; [z:=20+10]^3 \\ \triangleright \quad & [x:=10]^1; [y:=20]^2; [z:=30]^3 \end{aligned}$$

after which no more steps are possible. Another transformation sequence is

$$\begin{aligned} RD \quad \vdash \quad & [x:=10]^1; [y:=x+10]^2; [z:=y+10]^3 \\ \triangleright \quad & [x:=10]^1; [y:=10+10]^2; [z:=y+10]^3 \\ \triangleright \quad & [x:=10]^1; [y:=10+10]^2; [z:=10+10+10]^3 \\ \triangleright \quad & [x:=10]^1; [y:=10+10]^2; [z:=30]^3 \\ \triangleright \quad & [x:=10]^1; [y:=20]^2; [z:=30]^3 \end{aligned}$$

which, in this particular case, yields the same resulting program. ■

Successive transformations. The above example shows that we shall want to perform many successive transformations:

$$RD \vdash S_1 \triangleright S_2 \triangleright \cdots \triangleright S_{n+1}$$

This could be costly because once S_1 has been transformed into S_2 we might have to *recompute* Reaching Definitions Analysis for S_2 before the transformation can be used to transform it into S_3 etc. It turns out that it is

sometimes possible to use the analysis for S_1 to obtain a reasonable analysis for S_2 without performing the analysis from scratch. In the case of Reaching Definitions and Constant Folding this is very easy: if RD is a solution to Reaching Definitions for S_i and $\text{RD} \vdash S_i \triangleright S_{i+1}$ then RD is also a solution to Reaching Definitions for S_{i+1} – intuitively, the reason is that the transformation only changed things that were not observed by the Reaching Definitions Analysis.

Concluding Remarks

In this chapter we have briefly illustrated a few approaches (but by no means all) to program analysis. Clearly there are many differences between the four approaches. However, the main aim of the chapter has been to suggest that there are also more *similarities* than one would perhaps have expected at first sight: in particular, the interplay between the use of equations versus constraints. It is also interesting to note that some of the techniques touched upon in this chapter have close connections to other approaches to reasoning about programs; especially, some versions of Annotated Type Systems are closely related to Hoare's logic for partial correctness assertions.

As mentioned earlier, the approaches to program analysis covered in this book are *semantics based* rather than *semantics directed*. The semantics directed approaches include the denotational based approaches [18, 66, 87, 89] and logic based approaches [12, 13, 61, 62].

Mini Projects

Mini Project 1.1 Correctness of Reaching Definitions

In this mini project we shall increase our faith in the Type and Effect System of Table 1.3 by proving that it is correct. This requires knowledge of regular expressions and homomorphisms to the extent covered in Appendix C.

First we shall show how to associate a regular expression with each statement. We define a function \mathcal{S} such that $\mathcal{S}(S)$ is a regular expression for each statement $S \in \text{Stmt}$. It is defined by structural induction (see Appendix B) as follows:

$$\begin{aligned} \mathcal{S}([x:=a]^t) &= !_x^t \\ \mathcal{S}([\text{skip}]^t) &= \Lambda \\ \mathcal{S}(S_1; S_2) &= \mathcal{S}(S_1) \cdot \mathcal{S}(S_2) \\ \mathcal{S}(\text{if } [b]^t \text{ then } S_1 \text{ else } S_2) &= \mathcal{S}(S_1) + \mathcal{S}(S_2) \end{aligned}$$

$$S(\text{while } b \text{ do } S) = (S(S))^*$$

The alphabet is $\{!_x^\ell \mid x \in \mathbf{Var}_*, \ell \in \mathbf{Lab}_*\}$ where \mathbf{Var}_* and \mathbf{Lab}_* are finite and non-empty sets that contain all the variables and labels, respectively, of the statement S_* of interest. As an example, for S_* being

$$\text{if } [x > 0]^1 \text{ then } [x := x + 1]^2 \text{ else } ([x := x + 2]^3; [x := x + 3]^4)$$

we have $S(S_*) = !_x^2 + (!_x^3 \cdot !_x^4)$.

Correctness of X. To show the correctness of the X component in $S : \Sigma \xrightarrow{X} \Sigma$ we shall for each $y \in \mathbf{Var}_*$ define a homomorphism

$$h_y : \{!_x^\ell \mid x \in \mathbf{Var}_*, \ell \in \mathbf{Lab}_*\} \rightarrow \{!\}^*$$

as follows:

$$h_y(!_x^\ell) = \begin{cases} ! & \text{if } y = x \\ \Lambda & \text{if } y \neq x \end{cases}$$

As an example $h_x(S(S_*)) = ! + (! \cdot !)$ and $h_y(S(S_*)) = \Lambda$ using that $\Lambda \cdot \Lambda = \Lambda$ and $\Lambda + \Lambda = \Lambda$. Next we write

$$h_y(S(S)) \subseteq ! \cdot !^*$$

to mean that the language $\mathcal{L}[h_y(S(S))]$ defined by the regular expression $h_y(S(S))$ is a subset of the language $\mathcal{L}[! \cdot !^*]$ defined by $! \cdot !^*$; this is equivalent to

$$\neg \exists w \in \mathcal{L}[h_y(S(S))] : h_y(w) = \Lambda$$

and intuitively says that y is always assigned in S . Prove that

$$\text{if } S : \Sigma \xrightarrow{X} \Sigma \text{ and } y \in X \text{ then } h_y(S(S)) \subseteq ! \cdot !^*$$

by induction on the shape of the inference tree establishing $S : \Sigma \xrightarrow{X} \Sigma$ (see Appendix B for an introduction to the proof principle).

Correctness of RD. To show the correctness of the RD component in $S : \Sigma \xrightarrow{RD} \Sigma$ we shall for each $y \in \mathbf{Var}_*$ and $\ell' \in \mathbf{Lab}_*$ define a homomorphism

$$h_y^{\ell'} : \{!_x^\ell \mid x \in \mathbf{Var}_*, \ell \in \mathbf{Lab}_*\} \rightarrow \{!, ?\}^*$$

as follows:

$$h_y^{\ell'}(!_x^\ell) = \begin{cases} ! & \text{if } y = x \wedge \ell = \ell' \\ ? & \text{if } y = x \wedge \ell \neq \ell' \\ \Lambda & \text{if } y \neq x \end{cases}$$

As an example $h_x^2(S(S_*)) = ! + (? \cdot ?)$ and $h_y^5(S(S_*)) = \Lambda$. Next

$$h_y^{\ell'}(S(S)) \subseteq ((! + ?)^* \cdot ?) + \Lambda$$

is equivalent to

$$\neg \exists w \in \mathcal{L}[S(S)] : h_y^{\ell'}(w) \text{ ends in !}$$

and intuitively means that the last assignment to y could not have been performed at the statement labelled ℓ' . Prove that

$$\text{if } S : \Sigma \xrightarrow{\overline{RD}} \Sigma \text{ and } (y, \ell') \notin \text{RD then } h_y^{\ell'}(S(S)) \subseteq ((!+?)^* \cdot ?) + \Lambda$$

by induction on the shape of the inference tree establishing $S : \Sigma \xrightarrow{\overline{RD}} \Sigma$. ■

Exercises

Exercise 1.1 A variant of Reaching Definitions replaces $\text{RD} \in \mathcal{P}(\text{Var} \times \text{Lab})$ by $\text{RL} \in \mathcal{P}(\text{Lab})$; the idea is that given the program, a label should suffice for finding the variables that may be assigned in some elementary block bearing that label. Use this as the basis for modifying the equation system given in Section 1.3 for $\overline{\text{RD}}$ to an equation system for $\overline{\text{RL}}$. (Hint: It may be appropriate to think of $\text{RD} = \{(x_1, ?), \dots, (x_n, ?)\}$ as meaning $\text{RD} = \{(x_1, ?_{x_1}), \dots, (x_n, ?_{x_n})\}$ and then use $\text{RL} = \{?_{x_1}, \dots, ?_{x_n}\}$.) ■

Exercise 1.2 Show that the solution displayed for the Control Flow Analysis in Section 1.4 is a solution. Also show that it is in fact the least solution. (Hint: Consider the demands on $\widehat{C}(2)$, $\widehat{C}(4)$, $\widehat{\rho}(x)$, $\widehat{C}(1)$ and $\widehat{C}(5)$.) ■

Exercise 1.3 Show that if (α, γ) is a Galois connection then α uniquely determines γ (meaning that if also (α, γ') is a Galois connection then $\gamma = \gamma'$) and similarly that γ uniquely determines α . ■

Exercise 1.4 For F as in Section 1.3 and G as in Section 1.5 show that $\overline{\alpha} \circ G \circ \overline{\gamma} \sqsubseteq F$; this involves showing that

$$\alpha(G_j(\gamma(\text{RD}_1), \dots, \gamma(\text{RD}_{12}))) \subseteq F_j(\text{RD}_1, \dots, \text{RD}_{12})$$

for all j and $(\text{RD}_1, \dots, \text{RD}_{12})$. Determine whether or not $F = \overline{\alpha} \circ G \circ \overline{\gamma}$. ■

Exercise 1.5 Consider the Annotated Type System for Reaching Definitions defined in Table 1.2 in Section 1.6 and suppose that we want to stick to the first (and unsuccessful) explanation of what $S : \text{RD}_1 \rightarrow \text{RD}_2$ means in terms of Data Flow Analysis. Can you change Table 1.2 (by modifying or removing axioms and rules) such that this becomes possible? ■

Exercise 1.6 Consider the Chaotic Iteration algorithm of Section 1.7 and suppose that

$$\vec{\emptyset} \sqsubseteq \vec{RD} \sqsubseteq F(\vec{RD}) \sqsubseteq F^n(\vec{\emptyset}) = F^{n+1}(\vec{\emptyset})$$

holds immediately before the assignment to RD_j ; show that is also holds afterwards. (Hint: Write \vec{RD}' for $(RD_1, \dots, F_j(\vec{RD}), \dots, RD_{12})$ and use the monotonicity of F and $\vec{RD} \sqsubseteq F(\vec{RD})$ to establish that $\vec{RD} \sqsubseteq \vec{RD}' \sqsubseteq F(\vec{RD}) \sqsubseteq F(\vec{RD}')$.) ■

Exercise 1.7 Use the Chaotic Iteration scheme of Section 1.7 to show that the information displayed in Table 1.1 is in fact the least fixed point of the function F defined in Section 1.3. ■

Exercise 1.8 Consider the following program

$$[z:=1]^1; \text{while } [x>0]^2 \text{ do } ([z:=z*y]^3; [x:=x-1]^4)$$

computing the x -th power of the number stored in y . Formulate a system of data flow equations in the manner of Section 1.3. Next use the Chaotic Iteration strategy of Section 1.7 to compute the least solution and present it in a table (like Table 1.1). ■

Exercise 1.9 Perform Constant Folding upon the program

$$[x:=10]^1; [y:=x+10]^2; [z:=y+x]^3$$

so as to obtain

$$[x:=10]^1; [y:=20]^2; [z:=30]^3$$

How many ways of obtaining the result are there? ■

Exercise 1.10 The specification of Constant Folding in Section 1.8 only considers arithmetic expressions. Extend it to deal also with boolean expressions. Consider adding axioms like

$$RD \vdash ([\text{skip}]^\ell; S) \triangleright S$$

$$RD \vdash (\text{if } [\text{true}]^\ell \text{ then } S_1 \text{ else } S_2) \triangleright S_1$$

and discuss what complications arise. ■

Exercise 1.11 Consider adding the axiom

$$RD \vdash [x := a]^\ell \triangleright [x := a[y \mapsto a']]^\ell$$

$$\text{if } \begin{cases} y \in FV(a) \wedge (y, ?) \notin RD_{\text{entry}}(\ell) \wedge \\ \forall (z, \ell') \in RD_{\text{entry}}(\ell) : (y = z \Rightarrow [\cdot \cdot]^\ell \text{ is } [y := a']^{\ell'}) \end{cases}$$

to the specification of Constant Folding given in Section 1.8 and discuss whether or not this is a good idea. ■

Chapter 2

Data Flow Analysis

In this chapter we introduce techniques for Data Flow Analysis. Data Flow Analysis is the traditional form of program analysis which is described in many textbooks on compiler writing. We will present analyses for the simple imperative language WHILE that was introduced in Chapter 1. This includes a number of classical Data Flow Analyses: Available Expressions, Reaching Definitions, Very Busy Expressions and Live Variables. We introduce an operational semantics for WHILE and demonstrate the correctness of the Live Variables Analysis. We then present the notion of Monotone Frameworks and show how the examples may be recast as such frameworks. We continue by presenting a worklist algorithm for solving flow equations and we study its termination and correctness properties. The chapter concludes with a presentation of some advanced topics, including Interprocedural Data Flow Analysis and Shape Analysis.

2.1 Intraprocedural Analysis

In this section we present a number of example Data Flow Analyses for the WHILE language. The analyses are each defined by pairs of functions that map labels to the appropriate sets; one function in each pair specifies information that is true on *entry* to the block, the second specifies information that is true at the *exit*.

Initial and final labels. When presenting examples of Data Flow Analyses we will use a number of operations on programs and labels. The first of these is

$$\text{init} : \text{Stmt} \rightarrow \text{Lab}$$

which returns the *initial label* of a statement:

$$\begin{aligned}
\text{init}([x := a]^\ell) &= \ell \\
\text{init}([\text{skip}]^\ell) &= \ell \\
\text{init}(S_1; S_2) &= \text{init}(S_1) \\
\text{init}(\text{if } [b]^\ell \text{ then } S_1 \text{ else } S_2) &= \ell \\
\text{init}(\text{while } [b]^\ell \text{ do } S) &= \ell
\end{aligned}$$

We will also need a function

$$\text{final} : \text{Stmt} \rightarrow \mathcal{P}(\text{Lab})$$

which returns the set of *final labels* in a statement; whereas a sequence of statements has a single entry, it may have multiple exits (as for example in the conditional):

$$\begin{aligned}
\text{final}([x := a]^\ell) &= \{\ell\} \\
\text{final}([\text{skip}]^\ell) &= \{\ell\} \\
\text{final}(S_1; S_2) &= \text{final}(S_2) \\
\text{final}(\text{if } [b]^\ell \text{ then } S_1 \text{ else } S_2) &= \text{final}(S_1) \cup \text{final}(S_2) \\
\text{final}(\text{while } [b]^\ell \text{ do } S) &= \{\ell\}
\end{aligned}$$

Note that the *while*-loop terminates immediately after the test has evaluated to false.

Blocks. To access the statements or test associated with a label in a program we use the function

$$\text{blocks} : \text{Stmt} \rightarrow \mathcal{P}(\text{Blocks})$$

where **Blocks** is the set of statements, or *elementary blocks*, of the form $[x := a]^\ell$ or $[\text{skip}]^\ell$ as well as the set of tests of the form $[b]^\ell$. It is defined as follows:

$$\begin{aligned}
\text{blocks}([x := a]^\ell) &= \{[x := a]^\ell\} \\
\text{blocks}([\text{skip}]^\ell) &= \{[\text{skip}]^\ell\} \\
\text{blocks}(S_1; S_2) &= \text{blocks}(S_1) \cup \text{blocks}(S_2) \\
\text{blocks}(\text{if } [b]^\ell \text{ then } S_1 \text{ else } S_2) &= \{[b]^\ell\} \cup \text{blocks}(S_1) \cup \text{blocks}(S_2) \\
\text{blocks}(\text{while } [b]^\ell \text{ do } S) &= \{[b]^\ell\} \cup \text{blocks}(S)
\end{aligned}$$

Then the set of *labels* occurring in a program is given by

$$\text{labels} : \text{Stmt} \rightarrow \mathcal{P}(\text{Lab})$$

where

$$\text{labels}(S) = \{\ell \mid [B]^\ell \in \text{blocks}(S)\}$$

Clearly $\text{init}(S) \in \text{labels}(S)$ and $\text{final}(S) \subseteq \text{labels}(S)$.

Flows and reverse flows. We will need to operate on edges, or *flows*, between labels in a statement. We define a function

$$\text{flow} : \text{Stmt} \rightarrow \mathcal{P}(\text{Lab} \times \text{Lab})$$

which maps statements to sets of flows:

$$\begin{aligned} \text{flow}([x := a]^\ell) &= \emptyset \\ \text{flow}([\text{skip}]^\ell) &= \emptyset \\ \text{flow}(S_1; S_2) &= \text{flow}(S_1) \cup \text{flow}(S_2) \\ &\quad \cup \{(\ell, \text{init}(S_2)) \mid \ell \in \text{final}(S_1)\} \\ \text{flow}(\text{if } [b]^\ell \text{ then } S_1 \text{ else } S_2) &= \text{flow}(S_1) \cup \text{flow}(S_2) \\ &\quad \cup \{(\ell, \text{init}(S_1)), (\ell, \text{init}(S_2))\} \\ \text{flow}(\text{while } [b]^\ell \text{ do } S) &= \text{flow}(S) \cup \{(\ell, \text{init}(S))\} \\ &\quad \cup \{(\ell', \ell) \mid \ell' \in \text{final}(S)\} \end{aligned}$$

Thus $\text{labels}(S)$ and $\text{flow}(S)$ will be a representation of the *flow graph* of S .

Example 2.1 Consider the following program, *power*, computing the x -th power of the number stored in y :

```
[z:=1]1; while [x>0]2 do ([z:=z*y]3; [x:=x-1]4)
```

We have $\text{init}(\text{power}) = 1$, $\text{final}(\text{power}) = \{2\}$ and $\text{labels}(\text{power}) = \{1, 2, 3, 4\}$. The function *flow* produces the following set

$$\{(1, 2), (2, 3), (3, 4), (4, 2)\}$$

which corresponds to the flow graph in Figure 2.1. ■

The function *flow* is used in the formulation of *forward analyses*. Clearly $\text{init}(S)$ is the (unique) entry node for the flow graph with nodes $\text{labels}(S)$ and edges $\text{flow}(S)$. Also

$$\text{labels}(S) = \{\text{init}(S)\} \cup \{\ell \mid (\ell, \ell') \in \text{flow}(S)\} \cup \{\ell' \mid (\ell, \ell') \in \text{flow}(S)\}$$

and for composite statements (meaning those not simply of the form $[B]^\ell$) the equation remains true when removing the $\{\text{init}(S)\}$ component.

In order to formulate *backward analyses* we require a function that computes *reverse flows*:

$$\text{flow}^R : \text{Stmt} \rightarrow \mathcal{P}(\text{Lab} \times \text{Lab})$$

It is defined by:

$$\text{flow}^R(S) = \{(\ell, \ell') \mid (\ell', \ell) \in \text{flow}(S)\}$$




Figure 2.1: Flow graph for the power program.

Example 2.2 For the power program of Example 2.1, $flow^R$ produces

$$\{(2, 1), (2, 4), (3, 2), (4, 3)\}$$

which corresponds to a modified version of the flow graph in Figure 2.1 where the direction of the arcs has been reversed. ■

In case $final(S)$ contains just one element that will be the unique entry node for the flow graph with nodes $labels(S)$ and edges $flow^R(S)$. Also

$$labels(S) = final(S) \cup \{\ell \mid (\ell, \ell') \in flow^R(S)\} \cup \{\ell' \mid (\ell, \ell') \in flow^R(S)\}$$

and for composite statements the equation remains true when removing the $final(S)$ component.

The program of interest. We will use the notation S_* to represent the program that we are analysing (the “top-level” statement), Lab_* to represent the labels ($labels(S)$) appearing in S_* , Var_* to represent the variables ($FV(S_*)$) appearing in S_* , $Blocks_*$ to represent the elementary blocks ($blocks(S_*)$) occurring in S_* , and $AExp_*$ to represent the set of *non-trivial* arithmetic subexpressions in S_* ; an expression is trivial if it is a single variable or constant. We will also write $AExp(a)$ and $AExp(b)$ to refer to the set of non-trivial arithmetic subexpressions of a given arithmetic, respectively boolean, expression.

To simplify the presentation of the analyses, and to follow the traditions of the literature, we shall frequently assume that the program S_* has *isolated*

entries; this means that:

$$\forall \ell \in \mathbf{Lab} : (\ell, \text{init}(S_*)) \notin \text{flow}(S_*)$$

This is the case whenever S_* does not start with a `while`-loop. Similarly, we shall frequently assume that the program S_* has *isolated exits*; this means that:

$$\forall \ell_1 \in \text{final}(S_*) \forall \ell_2 \in \mathbf{Lab} : (\ell_1, \ell_2) \notin \text{flow}(S_*)$$

A statement, S , is *label consistent* if and only if:

$$[B_1]^\ell, [B_2]^\ell \in \text{blocks}(S) \text{ implies } B_1 = B_2$$

Clearly, if all blocks in S are uniquely labelled (meaning that each label occurs only once), then S is label consistent. When S is label consistent the clause “where $[B]^\ell \in \text{blocks}(S)$ ” is unambiguous in defining a partial function from labels to elementary blocks; we shall then say that ℓ *labels* the block B . We shall exploit this when defining the example analyses below.

Example 2.3 The power program of Example 2.1 has isolated entries but not isolated exits. It is clearly label consistent as well as uniquely labelled. ■

2.1.1 Available Expressions Analysis

The *Available Expressions Analysis* will determine:

For each program point, which expressions *must* have already been computed, and not later modified, on all paths to the program point.

This information can be used to avoid the re-computation of an expression. For clarity, we will concentrate on arithmetic expressions.

Example 2.4 Consider the following program:

$$[x:=a+b]^1; [y:=a*b]^2; \text{while } [y>a+b]^3 \text{ do } ([a:=a+1]^4; [x:=a+b]^5)$$

It should be clear that the expression $a+b$ is available every time execution reaches the test (label 3) in the loop; as a consequence, the expression need not be recomputed. ■

The analysis is defined in Table 2.1 and explained below. An expression is *killed* in a block if any of the variables used in the expression are modified in the block; we use the function

$$\text{kill}_{\text{AE}} : \mathbf{Blocks}_* \rightarrow \mathcal{P}(\mathbf{AExp}_*)$$

kill and gen functions	
$kill_{AE}([x := a]^\ell)$	$= \{a' \in AExp_* \mid x \in FV(a')\}$
$kill_{AE}([skip]^\ell)$	$= \emptyset$
$kill_{AE}([b]^\ell)$	$= \emptyset$
$gen_{AE}([x := a]^\ell)$	$= \{a' \in AExp(a) \mid x \notin FV(a')\}$
$gen_{AE}([skip]^\ell)$	$= \emptyset$
$gen_{AE}([b]^\ell)$	$= AExp(b)$
data flow equations: $AE^\#$	
$AE_{entry}(\ell)$	$= \begin{cases} \emptyset & \text{if } \ell = init(S_*) \\ \bigcap \{AE_{exit}(\ell') \mid (\ell', \ell) \in flow(S_*)\} & \text{otherwise} \end{cases}$
$AE_{exit}(\ell)$	$= (AE_{entry}(\ell) \setminus kill_{AE}(B^\ell)) \cup gen_{AE}(B^\ell)$ where $B^\ell \in blocks(S_*)$

Table 2.1: Available Expressions Analysis.

to produce the set of non-trivial expressions killed in the block. Test and skip blocks do not kill any expressions and assignments kill any expression that uses the variable that appears in the left hand side of the assignment. Note that in the clause for $[x:=a]^\ell$ we have used the notation $a' \in AExp_*$ to denote the fact that a' is a non-trivial arithmetic expression appearing in the program.

A *generated* expression is an expression that is evaluated in the block and where none of the variables used in the expression are later modified in the block. The set of non-trivial generated expressions is produced by the function:

$$gen_{AE} : Blocks_* \rightarrow \mathcal{P}(AExp_*)$$

The analysis itself is now defined by the functions AE_{entry} and AE_{exit} that each map labels to sets of expressions:

$$AE_{entry}, AE_{exit} : Lab_* \rightarrow \mathcal{P}(AExp_*)$$

For a label consistent program S_* (with isolated entries) the functions can be defined as in Table 2.1.

The analysis is a *forward analysis* and, as we shall see, we are interested in the *largest* sets satisfying the equation for AE_{entry} – an expression will be considered available if no path kills it. No expression is available at the start of the program. Subsequently, the expressions that are available at the entry to a block are any expressions that are available at *all* of the exits




Figure 2.2: A schematic flow graph.

from blocks that flow to the block; if there are no such blocks the formula evaluates to \mathbf{AExp}_* . Given a set of expressions that are available at the entry, the expressions available at the exit of the block are computed by removing killed expressions and adding any new expression generated by the block.

To see why we require the largest solution, consider Figure 2.2 which shows the flow graph for a program in a schematic way. Such a flow graph might correspond to the following program:

$$[z:=x+y]^\ell; \text{while } [\text{true}]^{\ell'} \text{ do } [\text{skip}]^{\ell''}$$

The set of expressions generated by the first assignment is $\{x+y\}$; the other blocks do not generate expressions and no block kills any expressions. The equations for $\mathbf{AE}_{\text{entry}}$ and $\mathbf{AE}_{\text{exit}}$ are as follows:

$$\begin{aligned} \mathbf{AE}_{\text{entry}}(\ell) &= \emptyset \\ \mathbf{AE}_{\text{entry}}(\ell') &= \mathbf{AE}_{\text{exit}}(\ell) \cap \mathbf{AE}_{\text{exit}}(\ell'') \\ \mathbf{AE}_{\text{entry}}(\ell'') &= \mathbf{AE}_{\text{exit}}(\ell') \\ \mathbf{AE}_{\text{exit}}(\ell) &= \mathbf{AE}_{\text{entry}}(\ell) \cup \{x+y\} \\ \mathbf{AE}_{\text{exit}}(\ell') &= \mathbf{AE}_{\text{entry}}(\ell') \\ \mathbf{AE}_{\text{exit}}(\ell'') &= \mathbf{AE}_{\text{entry}}(\ell'') \end{aligned}$$

After some simplification, we find that:

$$\mathbf{AE}_{\text{entry}}(\ell') = \{x+y\} \cap \mathbf{AE}_{\text{entry}}(\ell')$$

There are two solutions to this equation: $\{x+y\}$ and \emptyset . Consideration of the example and the definition of available expressions shows that the most

informative solution is $\{x+y\}$ – the expression is available every time we enter ℓ' . Thus we require the *largest* solution to the equations.

Example 2.5 For the program

$$[x:=a+b]^1; [y:=a*b]^2; \text{while } [y>a+b]^3 \text{ do } ([a:=a+1]^4; [x:=a+b]^5)$$

of Example 2.4, $kill_{AE}$ and gen_{AE} are defined as follows:

ℓ	$kill_{AE}(\ell)$	$gen_{AE}(\ell)$
1	\emptyset	$\{a+b\}$
2	\emptyset	$\{a*b\}$
3	\emptyset	$\{a+b\}$
4	$\{a+b, a*b, a+1\}$	\emptyset
5	\emptyset	$\{a+b\}$

We get the following equations:

$$\begin{aligned}
 AE_{entry}(1) &= \emptyset \\
 AE_{entry}(2) &= AE_{exit}(1) \\
 AE_{entry}(3) &= AE_{exit}(2) \cap AE_{exit}(5) \\
 AE_{entry}(4) &= AE_{exit}(3) \\
 AE_{entry}(5) &= AE_{exit}(4) \\
 AE_{exit}(1) &= AE_{entry}(1) \cup \{a+b\} \\
 AE_{exit}(2) &= AE_{entry}(2) \cup \{a*b\} \\
 AE_{exit}(3) &= AE_{entry}(3) \cup \{a+b\} \\
 AE_{exit}(4) &= AE_{entry}(4) \setminus \{a+b, a*b, a+1\} \\
 AE_{exit}(5) &= AE_{entry}(5) \cup \{a+b\}
 \end{aligned}$$

Using an analogue of the Chaotic Iteration discussed in Chapter 1 (starting with $AExp_*$ rather than \emptyset) we can compute the following solution:

ℓ	$AE_{entry}(\ell)$	$AE_{exit}(\ell)$
1	\emptyset	$\{a+b\}$
2	$\{a+b\}$	$\{a+b, a*b\}$
3	$\{a+b\}$	$\{a+b\}$
4	$\{a+b\}$	\emptyset
5	\emptyset	$\{a+b\}$

Note that, even though a is redefined in the loop, the expression $a+b$ is re-evaluated in the loop and so it is always available on entry to the loop. On the other hand, $a*b$ is available on the first entry to the loop but is killed before the next iteration. ■

kill and gen functions	
$kill_{RD}([x := a]^\ell)$	$= \{(x, ?)\} \cup \{(x, \ell') \mid B^{\ell'} \text{ is an assignment to } x \text{ in } S_\star\}$
$kill_{RD}([\text{skip}]^\ell)$	$= \emptyset$
$kill_{RD}([b]^\ell)$	$= \emptyset$
$gen_{RD}([x := a]^\ell)$	$= \{(x, \ell)\}$
$gen_{RD}([\text{skip}]^\ell)$	$= \emptyset$
$gen_{RD}([b]^\ell)$	$= \emptyset$
data flow equations: $RD^\#$	
$RD_{entry}(\ell)$	$= \begin{cases} \{(x, ?) \mid x \in FV(S_\star)\} & \text{if } \ell = \text{init}(S_\star) \\ \cup \{RD_{exit}(\ell') \mid (\ell', \ell) \in \text{flow}(S_\star)\} & \text{otherwise} \end{cases}$
$RD_{exit}(\ell)$	$= (RD_{entry}(\ell) \setminus kill_{RD}(B^\ell)) \cup gen_{RD}(B^\ell)$ where $B^\ell \in \text{blocks}(S_\star)$

Table 2.2: Reaching Definitions Analysis.

2.1.2 Reaching Definitions Analysis

As mentioned in Chapter 1, the *Reaching Definitions Analysis* should more properly be called reaching assignments but we will use the traditional name. This analysis is analogous to the previous one except that we are interested in:

For each program point, which assignments *may* have been made and not overwritten, when program execution reaches this point along some path.

A main application of Reaching Definitions Analysis is in the construction of direct links between blocks that produce values and blocks that use them; we shall return to this in Subsection 2.1.5.

Example 2.6 Consider the following program:

$$[x:=5]^1; [y:=1]^2; \text{while } [x>1]^3 \text{ do } ([y:=x*y]^4; [x:=x-1]^5)$$

All of the assignments reach the entry of 4 (the assignments labelled 1 and 2 reach there on the first iteration); only the assignments labelled 1, 4 and 5 reach the entry of 5. ■

The analysis is specified in Table 2.2. The function

$$kill_{RD} : \text{Blocks}_\star \rightarrow \mathcal{P}(\text{Var}_\star \times \text{Lab}_\star)$$

produces the set of pairs of variables and labels of assignments that are destroyed by the block. An assignment is destroyed if the block assigns a new value to the variable, i.e. the left hand side of the assignment. To deal with uninitialised variables we shall, as in Chapter 1, use the special label "?".

The function

$$\text{gen}_{\text{RD}} : \text{Blocks}_* \rightarrow \mathcal{P}(\text{Var}_* \times \text{Lab}_*)$$

produces the set of pairs of variables and labels of assignments generated by the block; only assignments generate definitions.

The analysis itself is now defined by the pair of functions RD_{entry} and RD_{exit} mapping labels to sets of pairs of variables and labels (of assignment blocks):

$$\text{RD}_{\text{entry}}, \text{RD}_{\text{exit}} : \text{Lab}_* \rightarrow \mathcal{P}(\text{Var}_* \times \text{Lab}_*)$$

For a label consistent program S_* (with isolated entries) the functions are defined as in Table 2.2.

Similar to the previous example, this is a *forward analysis* but, as we shall see, we are interested in the *smallest* sets satisfying the equation for RD_{entry} . An assignment reaches the entry of a block if it reaches the exit of *any* of the blocks which precede it; if there are none the formula evaluates to \emptyset . The computation of the set of assignments reaching the exit of a block is analogous to the Available Expressions Analysis.

We motivate the requirement for the smallest solution by consideration of the program $[z:=x+y]^{\ell'}; \text{while } [\text{true}]^{\ell'} \text{ do } [\text{skip}]^{\ell''}$ corresponding to Figure 2.2 again. The equations for RD_{entry} and RD_{exit} are as follows:

$$\begin{aligned} \text{RD}_{\text{entry}}(\ell) &= \{(x, ?), (y, ?), (z, ?)\} \\ \text{RD}_{\text{entry}}(\ell') &= \text{RD}_{\text{exit}}(\ell) \cup \text{RD}_{\text{exit}}(\ell'') \\ \text{RD}_{\text{entry}}(\ell'') &= \text{RD}_{\text{exit}}(\ell') \\ \text{RD}_{\text{exit}}(\ell) &= (\text{RD}_{\text{entry}}(\ell) \setminus \{(z, ?)\}) \cup \{(z, \ell)\} \\ \text{RD}_{\text{exit}}(\ell') &= \text{RD}_{\text{entry}}(\ell') \\ \text{RD}_{\text{exit}}(\ell'') &= \text{RD}_{\text{entry}}(\ell'') \end{aligned}$$

Once again, we concentrate on the entry of the block labelled ℓ' , $\text{RD}_{\text{entry}}(\ell')$; after some simplification we get

$$\text{RD}_{\text{entry}}(\ell') = \{(x, ?), (y, ?), (z, \ell)\} \cup \text{RD}_{\text{entry}}(\ell')$$

but this equation has many solutions: we can take $\text{RD}_{\text{entry}}(\ell')$ to be any superset of $\{(x, ?), (y, ?), (z, \ell)\}$. However, since ℓ' does not generate any new definitions, the most precise solution is $\{(x, ?), (y, ?), (z, \ell)\}$ – we require the *smallest* solution to the equations.

Sometimes, when the Reaching Definitions Analysis is presented in the literature, one has $RD_{entry}(init(S_*)) = \emptyset$ rather than $RD_{entry}(init(S_*)) = \{(x, ?) \mid x \in FV(S_*)\}$. This is correct only for programs that always assign variables before their first use; incorrect optimisations may result if this is not the case. The advantage of our formulation, as will emerge from Mini Project 2.2, is that it is always semantically sound.

Example 2.7 The following table summarises the assignments killed and generated by each of the blocks in the program

$$[x:=5]^1; [y:=1]^2; \text{while } [x>1]^3 \text{ do } ([y:=x*y]^4; [x:=x-1]^5)$$

of Example 2.6:

ℓ	$kill_{RD}(\ell)$	$gen_{RD}(\ell)$
1	$\{(x, ?), (x, 1), (x, 5)\}$	$\{(x, 1)\}$
2	$\{(y, ?), (y, 2), (y, 4)\}$	$\{(y, 2)\}$
3	\emptyset	\emptyset
4	$\{(y, ?), (y, 2), (y, 4)\}$	$\{(y, 4)\}$
5	$\{(x, ?), (x, 1), (x, 5)\}$	$\{(x, 5)\}$

The analysis gives rise to the following equations:

$$\begin{aligned}
 RD_{entry}(1) &= \{(x, ?), (y, ?)\} \\
 RD_{entry}(2) &= RD_{exit}(1) \\
 RD_{entry}(3) &= RD_{exit}(2) \cup RD_{exit}(5) \\
 RD_{entry}(4) &= RD_{exit}(3) \\
 RD_{entry}(5) &= RD_{exit}(4) \\
 RD_{exit}(1) &= (RD_{entry}(1) \setminus \{(x, ?), (x, 1), (x, 5)\}) \cup \{(x, 1)\} \\
 RD_{exit}(2) &= (RD_{entry}(2) \setminus \{(y, ?), (y, 2), (y, 4)\}) \cup \{(y, 2)\} \\
 RD_{exit}(3) &= RD_{entry}(3) \\
 RD_{exit}(4) &= (RD_{entry}(4) \setminus \{(y, ?), (y, 2), (y, 4)\}) \cup \{(y, 4)\} \\
 RD_{exit}(5) &= (RD_{entry}(5) \setminus \{(x, ?), (x, 1), (x, 5)\}) \cup \{(x, 5)\}
 \end{aligned}$$

Using Chaotic Iteration we may compute the solution:

ℓ	$RD_{entry}(\ell)$	$RD_{exit}(\ell)$
1	$\{(x, ?), (y, ?)\}$	$\{(y, ?), (x, 1)\}$
2	$\{(y, ?), (x, 1)\}$	$\{(x, 1), (y, 2)\}$
3	$\{(x, 1), (y, 2), (y, 4), (x, 5)\}$	$\{(x, 1), (y, 2), (y, 4), (x, 5)\}$
4	$\{(x, 1), (y, 2), (y, 4), (x, 5)\}$	$\{(x, 1), (y, 4), (x, 5)\}$
5	$\{(x, 1), (y, 4), (x, 5)\}$	$\{(y, 4), (x, 5)\}$

kill and gen functions	
$kill_{VB}([x := a]^\ell)$	$= \{a' \in \mathbf{AExp}_* \mid x \in FV(a')\}$
$kill_{VB}([\text{skip}]^\ell)$	$= \emptyset$
$kill_{VB}([b]^\ell)$	$= \emptyset$
$gen_{VB}([x := a]^\ell)$	$= \mathbf{AExp}(a)$
$gen_{VB}([\text{skip}]^\ell)$	$= \emptyset$
$gen_{VB}([b]^\ell)$	$= \mathbf{AExp}(b)$
data flow equations: $VB^=$	
$VB_{exit}(\ell)$	$= \begin{cases} \emptyset & \text{if } \ell \in \text{final}(S_*) \\ \bigcap \{VB_{entry}(\ell') \mid (\ell', \ell) \in \text{flow}^R(S_*)\} & \text{otherwise} \end{cases}$
$VB_{entry}(\ell)$	$= (VB_{exit}(\ell) \setminus kill_{VB}(B^\ell)) \cup gen_{VB}(B^\ell)$ where $B^\ell \in \text{blocks}(S_*)$

Table 2.3: Very Busy Expressions Analysis.

2.1.3 Very Busy Expressions Analysis

An expression is *very busy* at the exit from a label if, no matter what path is taken from the label, the expression must always be used before any of the variables occurring in it are redefined. The aim of the *Very Busy Expressions Analysis* is to determine:

For each program point, which expressions *must* be very busy at the exit from the point.

A possible optimisation based on this information is to evaluate the expression at the block and store its value for later use; this optimisation is sometimes called *hoisting* the expression.

Example 2.8 Consider the program:

if $[a > b]^1$ then $([x := b - a]^2; [y := a - b]^3)$ else $([y := b - a]^4; [x := a - b]^5)$

The expressions $a - b$ and $b - a$ are both *very busy* at the start of the conditional; they can be hoisted to the start of the conditional resulting in a space saving in the size of the code generated for this program. ■

The analysis is specified in Table 2.3. We have already defined the notion of an expression being killed when we presented the Available Expressions




Figure 2.3: A schematic flow graph (in reverse).

Analysis; we use an equivalent function here:

$$kill_{VB} : \mathbf{Blocks}_* \rightarrow \mathcal{P}(\mathbf{AExp}_*)$$

By analogy with the previous analyses, we also need to define how a block generates additional very busy expressions. For this we use:

$$gen_{VB} : \mathbf{Blocks}_* \rightarrow \mathcal{P}(\mathbf{AExp}_*)$$

All of the expressions that appear in a block are very busy at the entry to the block (unlike what was the case for Available Expressions).

The analysis itself is defined by the pair of functions VB_{entry} and VB_{exit} mapping labels to sets of expressions:

$$VB_{entry}, VB_{exit} : \mathbf{Lab}_* \rightarrow \mathcal{P}(\mathbf{AExp}_*)$$

For a label consistent program S_* (with isolated exits) they are defined as in Table 2.3.

The analysis is a *backward analysis* and, as we shall see, we are interested in the *largest* sets satisfying the equation for VB_{exit} . The functions propagate information *against* the flow of the program: an expression is very busy at the exit from a block if it is very busy at the entry to *every* block that follows; if there are none the formula evaluates to \mathbf{AExp}_* . However, no expressions are very busy at the exit from any final block.

To motivate the fact that we require the largest set, we consider the situation where we have a flow graph as shown in Figure 2.3; this flow graph might correspond to the program:

$$(\text{while } [x > 1]^\ell \text{ do } [\text{skip}]^{\ell'}; [x := x + 1]^{\ell''})$$

The equations for this program are

$$VB_{entry}(\ell) = VB_{exit}(\ell)$$

$$\begin{aligned}
VB_{entry}(\ell') &= VB_{exit}(\ell') \\
VB_{entry}(\ell'') &= \{x+1\} \\
VB_{exit}(\ell) &= VB_{entry}(\ell') \cap VB_{entry}(\ell'') \\
VB_{exit}(\ell') &= VB_{entry}(\ell) \\
VB_{exit}(\ell'') &= \emptyset
\end{aligned}$$

and, for the exit conditions of ℓ , we calculate:

$$VB_{exit}(\ell) = VB_{exit}(\ell) \cap \{x+1\}$$

Any subset of $\{x+1\}$ is a solution but $\{x+1\}$ is the most informative. Hence we want the *largest* solution to the equations.

Example 2.9 To analyse the program

if $[a>b]^1$ then $([x:=b-a]^2; [y:=a-b]^3)$ else $([y:=b-a]^4; [x:=a-b]^5)$

of Example 2.8, we calculate the following killed and generated sets:

ℓ	$kill_{VB}(\ell)$	$gen_{VB}(\ell)$
1	\emptyset	\emptyset
2	\emptyset	$\{b-a\}$
3	\emptyset	$\{a-b\}$
4	\emptyset	$\{b-a\}$
5	\emptyset	$\{a-b\}$

We get the following equations:

$$\begin{aligned}
VB_{entry}(1) &= VB_{exit}(1) \\
VB_{entry}(2) &= VB_{exit}(2) \cup \{b-a\} \\
VB_{entry}(3) &= \{a-b\} \\
VB_{entry}(4) &= VB_{exit}(4) \cup \{b-a\} \\
VB_{entry}(5) &= \{a-b\} \\
VB_{exit}(1) &= VB_{entry}(2) \cap VB_{entry}(4) \\
VB_{exit}(2) &= VB_{entry}(3) \\
VB_{exit}(3) &= \emptyset \\
VB_{exit}(4) &= VB_{entry}(5) \\
VB_{exit}(5) &= \emptyset
\end{aligned}$$

We can then use an analogue of Chaotic Iteration (starting with \mathbf{AExp}_*

rather than \emptyset) to compute:

ℓ	$VB_{entry}(\ell)$	$VB_{exit}(\ell)$
1	{a-b, b-a}	{a-b, b-a}
2	{a-b, b-a}	{a-b}
3	{a-b}	\emptyset
4	{a-b, b-a}	{a-b}
5	{a-b}	\emptyset

2.1.4 Live Variables Analysis

A variable is *live* at the exit from a label if there exists a path from the label to a use of the variable that does not re-define the variable. The *Live Variables Analysis* will determine:

For each program point, which variables *may* be live at the exit from the point.

This analysis might be used as the basis for *Dead Code Elimination*. If the variable is not live at the exit from a label then, if the elementary block is an assignment to the variable, the elementary block can be eliminated.

Example 2.10 Consider the following expression:

$[x:=2]^1; [y:=4]^2; [x:=1]^3; (\text{if } [y>x]^4 \text{ then } [z:=y]^5 \text{ else } [z:=y*y]^6); [x:=z]^7$

The variable x is not live at the exit from label 1; the first assignment of the program is redundant. Both x and y are live at the exit from label 3. ■

The analysis is defined in Table 2.4. The variable that appears on the left hand side of an assignment is killed by the assignment; tests and skip statements do not kill variables. This is expressed by the function:

$$kill_{LV} : \text{Blocks}_* \rightarrow \mathcal{P}(\text{Var}_*)$$

The function

$$gen_{LV} : \text{Blocks}_* \rightarrow \mathcal{P}(\text{Var}_*)$$

produces the set of variables that appear in the block.

The analysis itself is defined by the pair of functions LV_{entry} and LV_{exit} mapping labels to sets of variables:

$$LV_{exit}, LV_{entry} : \text{Lab}_* \rightarrow \mathcal{P}(\text{Var}_*)$$

<i>kill and gen functions</i>	
$kill_{LV}([x := a]^\ell)$	$= \{x\}$
$kill_{LV}([\text{skip}]^\ell)$	$= \emptyset$
$kill_{LV}([b]^\ell)$	$= \emptyset$
$gen_{LV}([x := a]^\ell)$	$= FV(a)$
$gen_{LV}([\text{skip}]^\ell)$	$= \emptyset$
$gen_{LV}([b]^\ell)$	$= FV(b)$
data flow equations: $LV^=$	
$LV_{exit}(\ell)$	$= \begin{cases} \emptyset & \text{if } \ell \in final(S_*) \\ \bigcup \{LV_{entry}(\ell') \mid (\ell', \ell) \in flow^R(S_*)\} & \text{otherwise} \end{cases}$
$LV_{entry}(\ell)$	$= (LV_{exit}(\ell) \setminus kill_{LV}(B^\ell)) \cup gen_{LV}(B^\ell)$ where $B^\ell \in blocks(S_*)$

Table 2.4: Live Variables Analysis.

For a label consistent program S_* (with isolated exits) they can be defined as in Table 2.4. The equation for $LV_{exit}(\ell)$ includes a variable in the set of live variables (at the exit from ℓ) if it is live at the entry to any of the blocks that follow ℓ ; if there are none then the formula evaluates to \emptyset .

The analysis is a *backward analysis* and, as we shall see, we are interested in the *smallest* sets satisfying the equation for LV_{exit} . To see why we require the smallest set, consider once again the program

$$(\text{while } [x > 1]^\ell \text{ do } [\text{skip}]^{\ell'}); [x := x + 1]^{\ell''}$$

corresponding to the flow graph in Figure 2.3. The equations for the program are:

$$\begin{aligned} LV_{entry}(\ell) &= LV_{exit}(\ell) \cup \{x\} \\ LV_{entry}(\ell') &= LV_{exit}(\ell') \\ LV_{entry}(\ell'') &= \{x\} \\ LV_{exit}(\ell) &= LV_{entry}(\ell') \cup LV_{entry}(\ell'') \\ LV_{exit}(\ell') &= LV_{entry}(\ell) \\ LV_{exit}(\ell'') &= \emptyset \end{aligned}$$

Suppose that we are interested in $LV_{exit}(\ell)$; after some calculation we get:

$$LV_{exit}(\ell) = LV_{exit}(\ell) \cup \{x\}$$

Any superset of $\{x\}$ is a solution. Optimisations based on this analysis are based on “dead” variables – the smaller the set of live variables, the more optimisations are possible. Hence we shall be interested in the *smallest* solution $\{x\}$ to the equations. Correctness of the analysis will be established in Section 2.2.

Example 2.11 Returning to the program

$[x:=2]^1; [y:=4]^2; [x:=1]^3; (\text{if } [y>x]^4 \text{ then } [z:=y]^5 \text{ else } [z:=y*y]^6); [x:=z]^7$

of Example 2.10, we can compute $kill_{LV}$ and gen_{LV} as:

ℓ	$kill_{LV}(\ell)$	$gen_{LV}(\ell)$
1	$\{x\}$	\emptyset
2	$\{y\}$	\emptyset
3	$\{x\}$	\emptyset
4	\emptyset	$\{x, y\}$
5	$\{z\}$	$\{y\}$
6	$\{z\}$	$\{y\}$
7	$\{x\}$	$\{z\}$

We get the following equations:

$$\begin{aligned}
 LV_{entry}(1) &= LV_{exit}(1) \setminus \{x\} \\
 LV_{entry}(2) &= LV_{exit}(2) \setminus \{y\} \\
 LV_{entry}(3) &= LV_{exit}(3) \setminus \{x\} \\
 LV_{entry}(4) &= LV_{exit}(4) \cup \{x, y\} \\
 LV_{entry}(5) &= (LV_{exit}(5) \setminus \{z\}) \cup \{y\} \\
 LV_{entry}(6) &= (LV_{exit}(6) \setminus \{z\}) \cup \{y\} \\
 LV_{entry}(7) &= \{z\} \\
 \\
 LV_{exit}(1) &= LV_{entry}(2) \\
 LV_{exit}(2) &= LV_{entry}(3) \\
 LV_{exit}(3) &= LV_{entry}(4) \\
 LV_{exit}(4) &= LV_{entry}(5) \cup LV_{entry}(6) \\
 LV_{exit}(5) &= LV_{entry}(7) \\
 LV_{exit}(6) &= LV_{entry}(7) \\
 LV_{exit}(7) &= \emptyset
 \end{aligned}$$

We can then use Chaotic Iteration to compute the solution:

ℓ	$LV_{entry}(\ell)$	$LV_{exit}(\ell)$
1	\emptyset	\emptyset
2	\emptyset	$\{y\}$
3	$\{y\}$	$\{x, y\}$
4	$\{x, y\}$	$\{y\}$
5	$\{y\}$	$\{z\}$
6	$\{y\}$	$\{z\}$
7	$\{z\}$	\emptyset

Note that we have assumed that all variables are dead at the end of the program. Some authors assume that the variables of interest are output at the end of the program; in that case $LV_{exit}(7)$ should be $\{x, y, z\}$ which means that $LV_{entry}(7)$, $LV_{exit}(5)$ and $LV_{exit}(6)$ should all be $\{y, z\}$. ■

2.1.5 Derived Data Flow Information

It is often convenient to directly link labels of statements that produce values to the labels of statements that use them. Links that, for each use of a variable, associate all assignments that reach that use are called *Use-Definition chains* or *ud-chains*. Links that, for each assignment, associate all uses are called *Definition-Use chains* or *du-chains*.

In order to make these definitions more precise, we will use the notion of a *definition clear path* with respect to some variable. The idea is that ℓ_1, \dots, ℓ_n is a definition clear path for x if none of the blocks labelled ℓ_1, \dots, ℓ_n assigns a value to x and if ℓ_n uses x . Formally, for a label consistent program S_* we define the predicate *clear* :

$$\begin{aligned}
 clear(x, \ell, \ell') &= \exists \ell_1, \dots, \ell_n : \\
 &(\ell_1 = \ell) \wedge (\ell_n = \ell') \wedge (n > 0) \wedge \\
 &(\forall i \in \{1, \dots, n-1\} : (\ell_i, \ell_{i+1}) \in flow(S_*)) \wedge \\
 &(\forall i \in \{1, \dots, n-1\} : \neg def(x, \ell_i)) \wedge use(x, \ell_n)
 \end{aligned}$$

Here the predicate *use* checks whether the variable is used in a block

$$use(x, \ell) = (\exists B : [B]^\ell \in blocks(S_*) \wedge x \in gen_{LV}([B]^\ell))$$

and the predicate *def* checks whether the variable is assigned in a block:

$$def(x, \ell) = (\exists B : [B]^\ell \in blocks(S_*) \wedge x \in kill_{LV}([B]^\ell))$$

Armed with these definitions, we can define the functions

$$ud, du : \mathbf{Var}_* \times \mathbf{Lab}_* \rightarrow \mathcal{P}(\mathbf{Lab}_*)$$

as follows:

$$\begin{aligned}
 ud(x, \ell') &= \{\ell \mid \text{def}(x, \ell) \wedge \exists \ell'' : (\ell, \ell'') \in \text{flow}(S_*) \wedge \text{clear}(x, \ell'', \ell')\} \\
 &\cup \{? \mid \text{clear}(x, \text{init}(S_*), \ell')\} \\
 du(x, \ell) &= \begin{cases} \{\ell' \mid \text{def}(x, \ell) \wedge \exists \ell'' : (\ell, \ell'') \in \text{flow}(S_*) \wedge \text{clear}(x, \ell'', \ell')\} \\ \quad \text{if } \ell \neq ? \\ \{\ell' \mid \text{clear}(x, \text{init}(S_*), \ell')\} \\ \quad \text{if } \ell = ? \end{cases}
 \end{aligned}$$

So $ud(x, \ell')$ will return the labels where an occurrence of x at ℓ' might have obtained its value; this may be at a label ℓ in S_* or x may be uninitialised as indicated by the occurrence of “?”. And $du(x, \ell)$ will return the labels where the value assigned to x at ℓ might be used; again we distinguish between the case where x gets its value within the program and the case where it is uninitialised. It turns out that:

$$du(x, \ell) = \{\ell' \mid \ell \in ud(x, \ell')\}$$

Before showing how ud - and du -chains can be used, we illustrate the functions by a simple example.

Example 2.12 Consider the program:

$[x:=0]^1; [x:=3]^2; (\text{if } [z=x]^3 \text{ then } [z:=0]^4 \text{ else } [z:=x]^5); [y:=x]^6; [x:=y+z]^7$

Then we get:

$ud(x, \ell)$	x	y	z	$du(x, \ell)$	x	y	z
1	\emptyset	\emptyset	\emptyset	1	\emptyset	\emptyset	\emptyset
2	\emptyset	\emptyset	\emptyset	2	{3, 5, 6}	\emptyset	\emptyset
3	{2}	\emptyset	{?}	3	\emptyset	\emptyset	\emptyset
4	\emptyset	\emptyset	\emptyset	4	\emptyset	\emptyset	{7}
5	{2}	\emptyset	\emptyset	5	\emptyset	\emptyset	{7}
6	{2}	\emptyset	\emptyset	6	\emptyset	{7}	\emptyset
7	\emptyset	{6}	{4, 5}	7	\emptyset	\emptyset	\emptyset
				?	\emptyset	\emptyset	{3}

The table for ud shows that the occurrence of x in block 3 will get its value in block 2 and the table for du shows that the value assigned to x in block 2 may be used in block 3, 5 and 6. ■

One application of ud - and du -chains is for *Dead Code Elimination*; for the program of Example 2.12 we may e.g. remove the block labelled 1 because there will be no use of the value assigned to x before it is reassigned in the next block. Another application is in *Code Motion*; in the example program the block labelled 6 can be moved to just in front of the conditional because

it only uses variables assigned in earlier blocks and the conditional does not use the variable assigned in block 6.

The definitions of *ud*- and *du*-chains do not give any hints as to how to compute the chains – the definitions are not constructive. It is possible to give constructive definitions which re-use some of the functions that we have defined in the earlier examples. In order to define *ud*-chains we can use RD_{entry} , which records the assignments reaching a block and define

$$UD : \mathbf{Var}_* \times \mathbf{Lab}_* \rightarrow \mathcal{P}(\mathbf{Lab}_*)$$

by:

$$UD(x, \ell) = \begin{cases} \{\ell' \mid (x, \ell') \in RD_{entry}(\ell)\} & \text{if } x \in gen_{LV}(B^\ell) \\ \emptyset & \text{otherwise} \end{cases}$$

Similarly, we can define a function $DU : \mathbf{Var}_* \times \mathbf{Lab}_* \rightarrow \mathcal{P}(\mathbf{Lab}_*)$ for *du*-chains based on the functions we have seen previously. We shall leave this to Mini Project 2.1 where we also consider the formal relationship between the UD and DU functions and the functions *ud* and *du*.

2.2 Theoretical Properties

In this section we will show that the Live Variables Analysis of Subsection 2.1.4 is indeed correct; the correctness of the Reaching Definitions Analysis is the topic of Mini Project 2.2. We shall begin by presenting a formal semantics for WHILE.

The material of this section may be skimmed through on a first reading; however, it is frequently when conducting the correctness proof that the final and subtle errors in the analysis are found and corrected! In other words, proving the semantic correctness of the analysis should not be considered a dispensable development that is merely of interest for theoreticians.

2.2.1 Structural Operational Semantics

We choose to use a (so-called small step) *Structural Operational Semantics* because it allows us to reason about intermediate stages in a program execution and it also allows us to deal with non-terminating programs.

Configurations and transitions. First define a *state* as a mapping from variables to integers:

$$\sigma \in \mathbf{State} = \mathbf{Var} \rightarrow \mathbf{Z}$$

A *configuration* of the semantics is either a pair consisting of a statement and a state or it is a state; a terminal configuration is a configuration that simply is a state. The *transitions* of the semantics are of the form

$\mathcal{A} : \mathbf{AExp} \rightarrow (\mathbf{State} \rightarrow \mathbf{Z})$
$\mathcal{A}[x]\sigma = \sigma(x)$
$\mathcal{A}[n]\sigma = \mathcal{N}[n]$
$\mathcal{A}[a_1 \text{ op}_a a_2]\sigma = \mathcal{A}[a_1]\sigma \text{ op}_a \mathcal{A}[a_2]\sigma$
$\mathcal{B} : \mathbf{BExp} \rightarrow (\mathbf{State} \rightarrow \mathbf{T})$
$\mathcal{B}[\text{not } b]\sigma = \neg \mathcal{B}[b]\sigma$
$\mathcal{B}[b_1 \text{ op}_b b_2]\sigma = \mathcal{B}[b_1]\sigma \text{ op}_b \mathcal{B}[b_2]\sigma$
$\mathcal{B}[a_1 \text{ op}_r a_2]\sigma = \mathcal{A}[a_1]\sigma \text{ op}_r \mathcal{A}[a_2]\sigma$

Table 2.5: Semantics of expressions in WHILE.

$$\langle S, \sigma \rangle \rightarrow \sigma' \quad \text{and} \quad \langle S, \sigma \rangle \rightarrow \langle S', \sigma' \rangle$$

and express how the configuration is changed by *one step* of computation. So in the configuration $\langle S, \sigma \rangle$ one of two things may happen:

- the execution terminates after one step and we record that by giving the resulting state σ' , or
- the execution does not terminate after one step and we record that by a new configuration $\langle S', \sigma' \rangle$ where S' is the rest of the program and σ' is the updated state.

To deal with arithmetic and boolean expressions we require the semantic functions

$$\begin{aligned} \mathcal{A} : \mathbf{AExp} &\rightarrow (\mathbf{State} \rightarrow \mathbf{Z}) \\ \mathcal{B} : \mathbf{BExp} &\rightarrow (\mathbf{State} \rightarrow \mathbf{T}) \end{aligned}$$

whose definition are given in Table 2.5. Here we assume that op_a , op_b and op_r are the semantic counterparts of the corresponding syntax. We have also assumed the existence of $\mathcal{N} : \mathbf{Num} \rightarrow \mathbf{Z}$ which defines the semantics of numerals. For simplicity we have assumed that no errors can occur; this means that e.g. division by 0 will have to produce an integer. One can modify the definition so as to allow errors but this will complicate the correctness proof to be performed below. Note that the value of an expression is only affected by the variables appearing in it, that is:

$$\text{if } \forall x \in FV(a) : \sigma_1(x) = \sigma_2(x) \text{ then } \mathcal{A}[a]\sigma_1 = \mathcal{A}[a]\sigma_2$$

$$\text{if } \forall x \in FV(b) : \sigma_1(x) = \sigma_2(x) \text{ then } \mathcal{B}[b]\sigma_1 = \mathcal{B}[b]\sigma_2$$

These results can easily be proved by structural induction on expressions (or by mathematical induction on their size); see Appendix B for a brief introduction to these proof principles.

$[ass]$	$\langle [x := a]^{\ell}, \sigma \rangle \rightarrow \sigma[x \mapsto \mathcal{A}[a]\sigma]$	
$[skip]$	$\langle [skip]^{\ell}, \sigma \rangle \rightarrow \sigma$	
$[seq_1]$	$\frac{\langle S_1, \sigma \rangle \rightarrow \langle S'_1, \sigma' \rangle}{\langle S_1; S_2, \sigma \rangle \rightarrow \langle S'_1; S_2, \sigma' \rangle}$	
$[seq_2]$	$\frac{\langle S_1, \sigma \rangle \rightarrow \sigma'}{\langle S_1; S_2, \sigma \rangle \rightarrow \langle S_2, \sigma' \rangle}$	
$[if_1]$	$\langle [if [b]^{\ell} \text{ then } S_1 \text{ else } S_2], \sigma \rangle \rightarrow \langle S_1, \sigma \rangle$	$\text{if } \mathcal{B}[b]\sigma = \text{true}$
$[if_2]$	$\langle [if [b]^{\ell} \text{ then } S_1 \text{ else } S_2], \sigma \rangle \rightarrow \langle S_2, \sigma \rangle$	$\text{if } \mathcal{B}[b]\sigma = \text{false}$
$[wh_1]$	$\langle [while [b]^{\ell} \text{ do } S], \sigma \rangle \rightarrow \langle (S; \text{while } [b]^{\ell} \text{ do } S), \sigma \rangle$	$\text{if } \mathcal{B}[b]\sigma = \text{true}$
$[wh_2]$	$\langle [while [b]^{\ell} \text{ do } S], \sigma \rangle \rightarrow \sigma$	$\text{if } \mathcal{B}[b]\sigma = \text{false}$

Table 2.6: The Structural Operational Semantics of WHILE.

The detailed definition of the semantics of the statements may be found in Table 2.6; it is explained below.

The clause $[ass]$ specifies that the assignment $x := a$ is executed in one step; here we write $\sigma[x \mapsto \mathcal{A}[a]\sigma]$ for the state that is as σ except that x is mapped to $\mathcal{A}[a]\sigma$, i.e. the value that a will evaluate to in the state σ . Formally:

$$(\sigma[x \mapsto \mathcal{A}[a]\sigma])y = \begin{cases} \mathcal{A}[a]\sigma & \text{if } x = y \\ \sigma(y) & \text{otherwise} \end{cases}$$

The semantics of sequencing is given by the two rules $[seq_1]$ and $[seq_2]$ and the idea is as follows. The first step of executing $S_1; S_2$ is the first step of executing S_1 . It may be that only one step is needed for S_1 to terminate and then the rule $[seq_2]$ applies and says that the new configuration is $\langle S_2, \sigma' \rangle$ reflecting that we are ready to start executing S_2 in the next step. Alternatively, S_1 may not terminate in just one step but gives rise to some other configuration $\langle S'_1, \sigma' \rangle$; then the rule $[seq_1]$ applies and it expresses that the rest of S_1 and all of S_2 still have to be executed: the next configuration is $\langle S'_1; S_2, \sigma' \rangle$.

The semantics of the conditional is given by the two axioms $[if_1]$ and $[if_2]$ expressing that the first step of computation will select the appropriate branch based on the current value of the boolean expression.

Finally, the semantics of the while-construct is given by the two axioms $[wh_1]$ and $[wh_2]$; the first axiom expresses that if the boolean expressions evaluates

to true then the first step is to unroll the loop and the second axiom expresses that the execution terminates if the boolean expression evaluates to false.

Derivation sequences. A *derivation sequence* for a statement S_1 and a state σ_1 can take one of two forms:

- It is a *finite* sequence of configurations $\langle S_1, \sigma_1 \rangle, \dots, \langle S_n, \sigma_n \rangle, \sigma_{n+1}$ satisfying $\langle S_i, \sigma_i \rangle \rightarrow \langle S_{i+1}, \sigma_{i+1} \rangle$ for $i = 1, \dots, n-1$ and $\langle S_n, \sigma_n \rangle \rightarrow \sigma_{n+1}$; this corresponds to a terminating computation.
- It is an *infinite* sequence of configurations $\langle S_1, \sigma_1 \rangle, \dots, \langle S_i, \sigma_i \rangle, \dots$ satisfying $\langle S_i, \sigma_i \rangle \rightarrow \langle S_{i+1}, \sigma_{i+1} \rangle$ for all $i \geq 1$; this corresponds to a looping computation.

Example 2.13 We illustrate the semantics by showing an execution of the factorial program of Example 1.1. In the following we assume that the state $\sigma_{n_x n_y n_z}$ maps x to n_x , y to n_y and z to n_z . We then get the following finite derivation sequence:

$$\begin{aligned}
& \langle [y:=x]^1; [z:=1]^2; \text{while } [y>1]^3 \text{ do } ([z:=z*y]^4; [y:=y-1]^5); [y:=0]^6, \sigma_{300} \rangle \\
& \rightarrow \langle [z:=1]^2; \text{while } [y>1]^3 \text{ do } ([z:=z*y]^4; [y:=y-1]^5); [y:=0]^6, \sigma_{330} \rangle \\
& \rightarrow \langle \text{while } [y>1]^3 \text{ do } ([z:=z*y]^4; [y:=y-1]^5); [y:=0]^6, \sigma_{331} \rangle \\
& \rightarrow \langle [z:=z*y]^4; [y:=y-1]^5; \\
& \quad \text{while } [y>1]^3 \text{ do } ([z:=z*y]^4; [y:=y-1]^5); [y:=0]^6, \sigma_{331} \rangle \\
& \rightarrow \langle [y:=y-1]^5; \text{while } [y>1]^3 \text{ do } ([z:=z*y]^4; [y:=y-1]^5); [y:=0]^6, \sigma_{333} \rangle \\
& \rightarrow \langle \text{while } [y>1]^3 \text{ do } ([z:=z*y]^4; [y:=y-1]^5); [y:=0]^6, \sigma_{323} \rangle \\
& \rightarrow \langle [z:=z*y]^4; [y:=y-1]^5; \\
& \quad \text{while } [y>1]^3 \text{ do } ([z:=z*y]^4; [y:=y-1]^5); [y:=0]^6, \sigma_{323} \rangle \\
& \rightarrow \langle [y:=y-1]^5; \text{while } [y>1]^3 \text{ do } ([z:=z*y]^4; [y:=y-1]^5); [y:=0]^6, \sigma_{326} \rangle \\
& \rightarrow \langle \text{while } [y>1]^3 \text{ do } ([z:=z*y]^4; [y:=y-1]^5); [y:=0]^6, \sigma_{316} \rangle \\
& \rightarrow \langle [y:=0]^6, \sigma_{316} \rangle \\
& \rightarrow \sigma_{306}
\end{aligned}$$

Note that labels have no impact on the semantics: they are merely carried along and never inspected. ■

Properties of the semantics. We shall first establish a number of properties of the operations on programs and labels that we have used in the formulation of the analyses. In the course of the computation the set of

flows, the set of final labels and the set of elementary blocks of the statements of the configurations will be modified; Lemma 2.14 shows that the sets will decrease:

Lemma 2.14

- (i) If $\langle S, \sigma \rangle \rightarrow \sigma'$ then $final(S) = \{init(S)\}$.
- (ii) If $\langle S, \sigma \rangle \rightarrow \langle S', \sigma' \rangle$ then $final(S) \supseteq final(S')$.
- (iii) If $\langle S, \sigma \rangle \rightarrow \langle S', \sigma' \rangle$ then $flow(S) \supseteq flow(S')$.
- (iv) If $\langle S, \sigma \rangle \rightarrow \langle S', \sigma' \rangle$ then $blocks(S) \supseteq blocks(S')$ and if S is label consistent then so is S' . ■

Proof The proof of (i) is by induction on the shape of the inference tree used to establish $\langle S, \sigma \rangle \rightarrow \sigma'$; we refer to Appendix B for a brief introduction to the proof principle. Consulting Table 2.6 we see that there are three non-vacuous cases:

The case $[ass]$. Then $\langle [x := a]^\ell, \sigma \rangle \rightarrow \sigma[x \mapsto \mathcal{A}[a]\sigma]$ and we get:

$$final([x := a]^\ell) = \{\ell\} = \{init([x := a]^\ell)\}$$

The case $[skip]$. Then $\langle [skip]^\ell, \sigma \rangle \rightarrow \sigma$ and we get:

$$final([skip]^\ell) = \{\ell\} = \{init([skip]^\ell)\}$$

The case $[wh_2]$. Then $\langle \text{while } [b]^\ell \text{ do } S, \sigma \rangle \rightarrow \sigma$ because $\mathcal{B}[b]\sigma = false$ and we get:

$$final(\text{while } [b]^\ell \text{ do } S) = \{\ell\} = \{init(\text{while } [b]^\ell \text{ do } S)\}$$

This completes the proof of (i).

The proof of (ii) is by induction on the shape of the inference tree used to establish $\langle S, \sigma \rangle \rightarrow \langle S', \sigma' \rangle$. There are five non-vacuous cases:

The case $[seq_1]$. Then $\langle S_1; S_2, \sigma \rangle \rightarrow \langle S'_1; S_2, \sigma' \rangle$ because $\langle S_1, \sigma \rangle \rightarrow \langle S'_1, \sigma' \rangle$ and we get:

$$final(S_1; S_2) = final(S_2) = final(S'_1; S_2)$$

The case $[seq_2]$. Then $\langle S_1; S_2, \sigma \rangle \rightarrow \langle S_2, \sigma' \rangle$ because $\langle S_1, \sigma \rangle \rightarrow \sigma'$ and we get:

$$final(S_1; S_2) = final(S_2)$$

The case $[if_1]$. Then $\langle \text{if } [b]^\ell \text{ then } S_1 \text{ else } S_2, \sigma \rangle \rightarrow \langle S_1, \sigma \rangle$ because $\mathcal{B}[b]\sigma = true$ and we get:

$$final(\text{if } [b]^\ell \text{ then } S_1 \text{ else } S_2) = final(S_1) \cup final(S_2) \supseteq final(S_1)$$

The case $[if_2]$ is similar to the previous case.

The case $[wh_1]$. Then $\langle \text{while } [b]^\ell \text{ do } S, \sigma \rangle \rightarrow \langle (S; \text{while } [b]^\ell \text{ do } S), \sigma \rangle$ because $\mathcal{B}[b]\sigma = true$ and we get:

$$final(S; \text{while } [b]^\ell \text{ do } S) = final(\text{while } [b]^\ell \text{ do } S)$$

This completes the proof of (ii).

The proof of (iii) is by induction on the shape of the inference tree used to establish $\langle S, \sigma \rangle \rightarrow \langle S', \sigma' \rangle$. There are five non-vacuous cases:

The case $[seq_1]$. Then $\langle S_1; S_2, \sigma \rangle \rightarrow \langle S'_1; S_2, \sigma' \rangle$ because $\langle S_1, \sigma \rangle \rightarrow \langle S'_1, \sigma' \rangle$ and we get

$$\begin{aligned} flow(S_1; S_2) &= flow(S_1) \cup flow(S_2) \cup \{(\ell, init(S_2)) \mid \ell \in final(S_1)\} \\ &\supseteq flow(S'_1) \cup flow(S_2) \cup \{(\ell, init(S_2)) \mid \ell \in final(S_1)\} \\ &\supseteq flow(S'_1) \cup flow(S_2) \cup \{(\ell, init(S_2)) \mid \ell \in final(S'_1)\} \\ &= flow(S'_1; S_2) \end{aligned}$$

where we have used the induction hypothesis and (ii).

The case $[seq_2]$. Then $\langle S_1; S_2, \sigma \rangle \rightarrow \langle S_2, \sigma' \rangle$ because $\langle S_1, \sigma \rangle \rightarrow \sigma'$ and we get:

$$\begin{aligned} flow(S_1; S_2) &= flow(S_1) \cup flow(S_2) \cup \{(\ell, init(S_2)) \mid \ell \in final(S_1)\} \\ &\supseteq flow(S_2) \end{aligned}$$

The case $[if_1]$. Then $\langle \text{if } [b]^t \text{ then } S_1 \text{ else } S_2, \sigma \rangle \rightarrow \langle S_1, \sigma \rangle$ because $B[b]\sigma = true$ and we get:

$$\begin{aligned} flow(\text{if } [b]^t \text{ then } S_1 \text{ else } S_2) &= flow(S_1) \cup flow(S_2) \\ &\quad \cup \{(\ell, init(S_1)), (\ell, init(S_2))\} \\ &\supseteq flow(S_1) \end{aligned}$$

The case $[if_2]$ is similar to the previous case.

The case $[wh_2]$. Then $\langle \text{while } [b]^t \text{ do } S, \sigma \rangle \rightarrow \langle S; \text{while } [b]^t \text{ do } S, \sigma \rangle$ because $B[b]\sigma = true$ and we get:

$$\begin{aligned} flow(S; \text{while } [b]^t \text{ do } S) &= flow(S) \cup flow(\text{while } [b]^t \text{ do } S) \\ &\quad \cup \{(\ell', \ell) \mid \ell' \in final(S)\} \\ &= flow(S) \cup flow(S) \cup \{(\ell, init(S))\} \\ &\quad \cup \{(\ell', \ell) \mid \ell' \in final(S)\} \cup \{(\ell', \ell) \mid \ell' \in final(S)\} \\ &= flow(S) \cup \{(\ell, init(S))\} \cup \{(\ell', \ell) \mid \ell' \in final(S)\} \\ &= flow(\text{while } [b]^t \text{ do } S) \end{aligned}$$

This completes the proof of (iii).

The proof of (iv) is similar to that of (iii) and we omit the details. \blacksquare

2.2.2 Correctness of Live Variables Analysis

Preservation of solutions. Subsection 2.1.4 shows how to define an *equation system* for a label consistent program S_* (with isolated exits); we will

refer to this system as $LV^=(S_*)$. The construction of $LV^=(S_*)$ can be modified to give a *constraint system* $LV^⊆(S_*)$ of the form studied in Subsection 1.3.2:

$$LV_{exit}(\ell) \supseteq \begin{cases} \emptyset & \text{if } \ell \in \text{final}(S_*) \\ \bigcup \{LV_{entry}(\ell') \mid (\ell', \ell) \in \text{flow}^R(S_*)\} & \text{otherwise} \end{cases}$$

$$LV_{entry}(\ell) \supseteq (LV_{exit}(\ell) \setminus \text{kill}_{LV}(B^\ell)) \cup \text{gen}_{LV}(B^\ell) \\ \text{where } B^\ell \in \text{blocks}(S_*)$$

We make this definition because in the correctness proof we will want to use the *same* solution for all statements derived from S_* ; this will be possible for $LV^⊆(S_*)$ but not for $LV^=(S_*)$.

Now consider a collection *live* of functions:

$$\text{live}_{entry}, \text{live}_{exit} : \mathbf{Lab}_* \rightarrow \mathcal{P}(\mathbf{Var}_*)$$

We say that *live* solves $LV^=(S)$, and write

$$\text{live} \models LV^=(S)$$

if the functions satisfy the equations; similarly we write

$$\text{live} \models LV^⊆(S)$$

if *live* solves $LV^⊆(S_*)$. The following result shows that any solution of the equation system is also a solution of the constraint system and that the least solutions of the two systems coincide.

Lemma 2.15 Consider a label consistent program S_* . If $\text{live} \models LV^=(S_*)$ then $\text{live} \models LV^⊆(S_*)$. The least solution of $LV^=(S_*)$ coincides with the least solution of $LV^⊆(S_*)$. ■

Proof If $\text{live} \models LV^=(S_*)$ then clearly $\text{live} \models LV^⊆(S_*)$ because “ \supseteq ” includes the case of “ $=$ ”.

Next let us prove that $LV^⊆(S_*)$ and $LV^=(S_*)$ have the same least solution. We gave a constructive proof of a related result in Chapter 1 (under some assumptions about finiteness) so let us here give a more abstract proof using more advanced fixed point theory (as covered in Appendix A). In the manner of Chapter 1 we construct a function F_{LV}^S such that:

$$\begin{aligned} \text{live} \models LV^⊆(S) & \quad \text{iff} \quad \text{live} \supseteq F_{LV}^S(\text{live}) \\ \text{live} \models LV^=(S) & \quad \text{iff} \quad \text{live} = F_{LV}^S(\text{live}) \end{aligned}$$

Using Tarski's Fixed Point Theorem (Proposition A.10) we now have that F_{LV}^S has a least fixed point $\text{lfp}(F_{LV}^S)$ such that

$$\text{lfp}(F_{LV}^S) = \bigcap \{ \text{live} \mid \text{live} \supseteq F_{LV}^S(\text{live}) \} = \bigcap \{ \text{live} \mid \text{live} = F_{LV}^S(\text{live}) \}$$




Figure 2.4: Preservation of analysis result.

and since $\text{Ifp}(F_{\text{LV}}^S) = F_{\text{LV}}^S(\text{Ifp}(F_{\text{LV}}^S))$ as well as $\text{Ifp}(F_{\text{LV}}^S) \supseteq F_{\text{LV}}^S(\text{Ifp}(F_{\text{LV}}^S))$ this proves the result. ■

The next result shows that if we have a solution to the constraint system corresponding to some statement S_1 then it will also be a solution to the constraint system obtained from a sub-statement S_2 ; this result will be essential in the proof of the correctness result.

Lemma 2.16 If $\text{live} \models \text{LV}^\subseteq(S_1)$ (with S_1 being label consistent) and $\text{flow}(S_1) \supseteq \text{flow}(S_2)$ and $\text{blocks}(S_1) \supseteq \text{blocks}(S_2)$ then $\text{live} \models \text{LV}^\subseteq(S_2)$ (with S_2 being label consistent). ■

Proof If S_1 is label consistent and $\text{blocks}(S_1) \supseteq \text{blocks}(S_2)$ then also S_2 is label consistent. If $\text{live} \models \text{LV}^\subseteq(S_1)$ then live also satisfies each constraint in $\text{LV}^\subseteq(S_2)$ and hence $\text{live} \models \text{LV}^\subseteq(S_2)$. ■

We now have the following corollary expressing that the solution to the constraints of LV^\subseteq is preserved during computation; this is illustrated in Figure 2.4 for finite computations.

Corollary 2.17 If $\text{live} \models \text{LV}^\subseteq(S)$ (for S being label consistent) and if $\langle S, \sigma \rangle \rightarrow \langle S', \sigma' \rangle$ then also $\text{live} \models \text{LV}^\subseteq(S')$. ■

Proof Follows from Lemma 2.14 and 2.16. ■

We also have an easy result relating entry and exit components of a solution.

Lemma 2.18 If $\text{live} \models \text{LV}^\subseteq(S)$ (with S being label consistent) then for all $(\ell, \ell') \in \text{flow}(S)$ we have $\text{live}_{\text{exit}}(\ell) \supseteq \text{live}_{\text{entry}}(\ell')$. ■

Proof The result follows immediately from the construction of $\text{LV}^\subseteq(S)$. ■

Correctness relation. Intuitively, the correctness result for the Live Variables Analysis should express that the sets of live variables computed by




Figure 2.5: The correctness result for live.

the analysis are correct throughout the computation. Only the values of the live variables are of interest for the computation: if a variable is not live then its value in the state is irrelevant – its value cannot affect the interesting parts of the result of the computation. Assume now that V is a set of live variables and define the *correctness relation*:

$$\sigma_1 \sim_V \sigma_2 \quad \text{iff} \quad \forall x \in V : \sigma_1(x) = \sigma_2(x)$$

The relation expresses that for all practical purposes the two states σ_1 and σ_2 are equal: only the values of the live variables matters and here the two states are equal.

Example 2.19 Consider the statement $[x:=y+z]^\ell$ and let $V_1 = \{y, z\}$ and $V_2 = \{x\}$. Then $\sigma_1 \sim_{V_1} \sigma_2$ means $\sigma_1(y) = \sigma_2(y) \wedge \sigma_1(z) = \sigma_2(z)$ and $\sigma_1 \sim_{V_2} \sigma_2$ means $\sigma_1(x) = \sigma_2(x)$.

Next suppose that $([x:=y+z]^\ell, \sigma_1) \rightarrow \sigma'_1$ and $([x:=y+z]^\ell, \sigma_2) \rightarrow \sigma'_2$. Clearly $\sigma_1 \sim_{V_1} \sigma_2$ ensures that $\sigma'_1 \sim_{V_2} \sigma'_2$. So if V_2 is the set of variables live after $[x:=y+z]^\ell$ then V_1 is the set of variables live before $[x:=y+z]^\ell$. ■

The correctness result will express that the relation “ \sim ” is an invariant under the computation. This is illustrated in Figure 2.5 for finite computations and it is formally expressed by Corollary 2.22 below; to improve the legibility of formulae we write:

$$N(\ell) = \text{live}_{\text{entry}}(\ell)$$

$$X(\ell) = \text{live}_{\text{exit}}(\ell)$$

Since the live variables at the exit from a label are defined to be (a superset of) the union of the live variables at the entries of all of its successor labels, we have the following result.

Lemma 2.20 Assume $live \models LV^{\subseteq}(S)$ with S being label consistent. Then $\sigma_1 \sim_{X(\ell)} \sigma_2$ implies $\sigma_1 \sim_{N(\ell')} \sigma_2$ for all $(\ell, \ell') \in flow(S)$. ■

Proof Follows directly from Lemma 2.18 and the definition of \sim_V . ■

Correctness result. We are now ready for the main result. It states how semantically correct liveness information is preserved under each step of the execution: (i) in the case where we do not immediately terminate and (ii) in the case where we do immediately terminate.

Theorem 2.21

If $live \models LV^{\subseteq}(S)$ (with S being label consistent) then:

- (i) if $\langle S, \sigma_1 \rangle \rightarrow \langle S', \sigma'_1 \rangle$ and $\sigma_1 \sim_{N(init(S))} \sigma_2$ then there exists σ'_2 such that $\langle S, \sigma_2 \rangle \rightarrow \langle S', \sigma'_2 \rangle$ and $\sigma'_1 \sim_{N(init(S'))} \sigma'_2$, and
- (ii) if $\langle S, \sigma_1 \rangle \rightarrow \sigma'_1$ and $\sigma_1 \sim_{N(init(S))} \sigma_2$ then there exists σ'_2 such that $\langle S, \sigma_2 \rangle \rightarrow \sigma'_2$ and $\sigma'_1 \sim_{X(init(S))} \sigma'_2$

Proof The proof is by induction on the shape of the inference tree used to establish $\langle S, \sigma_1 \rangle \rightarrow \langle S', \sigma'_1 \rangle$ and $\langle S, \sigma_1 \rangle \rightarrow \sigma'_1$, respectively.

The case [ass]. Then $\langle [x := a]^{\ell}, \sigma_1 \rangle \rightarrow \sigma_1[x \mapsto \mathcal{A}[a]\sigma_1]$ and from the specification of the constraint system we have

$$N(\ell) = live_{entry}(\ell) = (live_{exit}(\ell) \setminus \{x\}) \cup FV(a) = (X(\ell) \setminus \{x\}) \cup FV(a)$$

and thus

$$\sigma_1 \sim_{N(\ell)} \sigma_2 \text{ implies } \mathcal{A}[a]\sigma_1 = \mathcal{A}[a]\sigma_2$$

because the value of a is only affected by the variables occurring in it. Therefore, taking

$$\sigma'_2 = \sigma_2[x \mapsto \mathcal{A}[a]\sigma_2]$$

we have that $\sigma'_1(x) = \sigma'_2(x)$ and thus $\sigma'_1 \sim_{X(\ell)} \sigma'_2$ as required.

The case [skip]. Then $\langle [skip]^{\ell}, \sigma_1 \rangle \rightarrow \sigma_1$ and from the specification of the constraint system

$$N(\ell) = live_{entry}(\ell) = (live_{exit}(\ell) \setminus \emptyset) \cup \emptyset = live_{exit}(\ell) = X(\ell)$$

and we take σ'_2 to be σ_2 .

The case [seq]. Then $\langle S_1; S_2, \sigma_1 \rangle \rightarrow \langle S'_1; S_2, \sigma'_1 \rangle$ because $\langle S_1, \sigma_1 \rangle \rightarrow \langle S'_1, \sigma'_1 \rangle$. By construction we have $flow(S_1; S_2) \supseteq flow(S_1)$ and also $blocks(S_1; S_2) \supseteq blocks(S_1)$. Thus by Lemma 2.16, $live$ is a solution to $LV^{\subseteq}(S_1)$ and thus by the induction hypothesis there exists σ'_2 such that

$$\langle S_1, \sigma_2 \rangle \rightarrow \langle S'_1, \sigma'_2 \rangle \text{ and } \sigma'_1 \sim_{N(init(S'_1))} \sigma'_2$$

and the result follows.

The case $[seq_2]$. Then $\langle S_1; S_2, \sigma_1 \rangle \rightarrow \langle S_2, \sigma'_1 \rangle$ because $\langle S_1, \sigma_1 \rangle \rightarrow \sigma'_1$. Once again by Lemma 2.16, $live$ is a solution to $LV^{\subseteq}(S_1)$ and thus by the induction hypothesis there exists σ'_2 such that:

$$\langle S_1, \sigma_2 \rangle \rightarrow \sigma'_2 \text{ and } \sigma'_1 \sim_{X(\text{init}(S_1))} \sigma'_2$$

Now

$$\{(\ell, \text{init}(S_2)) \mid \ell \in \text{final}(S_1)\} \subseteq \text{flow}(S_1; S_2)$$

and by Lemma 2.14, $\text{final}(S_1) = \{\text{init}(S_1)\}$. Thus by Lemma 2.20

$$\sigma'_1 \sim_{N(\text{init}(S_2))} \sigma'_2$$

and the result follows.

The case $[if_1]$. Then $\langle \text{if } [b]^{\ell} \text{ then } S_1 \text{ else } S_2, \sigma_1 \rangle \rightarrow \langle S_1, \sigma_1 \rangle$ because $\mathcal{B}[b]\sigma_1 = \text{true}$. Since $\sigma_1 \sim_{N(\ell)} \sigma_2$ and $N(\ell) = \text{live}_{\text{entry}}(\ell) \supseteq FV(b)$, we also have that $\mathcal{B}[b]\sigma_2 = \text{true}$ (the value of b is only affected by the variables occurring in it) and thus:

$$\langle \text{if } [b]^{\ell} \text{ then } S_1 \text{ else } S_2, \sigma_2 \rangle \rightarrow \langle S_1, \sigma_2 \rangle$$

From the specification of the constraint system, $N(\ell) = \text{live}_{\text{entry}}(\ell) \supseteq \text{live}_{\text{exit}}(\ell) = X(\ell)$ and hence $\sigma_1 \sim_{X(\ell)} \sigma_2$. Since $(\ell, \text{init}(S_1)) \in \text{flow}(S)$, Lemma 2.22 gives $\sigma_1 \sim_{N(\text{init}(S_1))} \sigma_2$ as required.

The case $[if_2]$ is similar to the previous case.

The case $[wh_1]$. Then $\langle \text{while } [b]^{\ell} \text{ do } S, \sigma_1 \rangle \rightarrow \langle S; \text{while } [b]^{\ell} \text{ do } S, \sigma_1 \rangle$ because $\mathcal{B}[b]\sigma_1 = \text{true}$. Since $\sigma_1 \sim_{N(\ell)} \sigma_2$ and $N(\ell) \supseteq FV(b)$, we also have that $\mathcal{B}[b]\sigma_2 = \text{true}$ and thus

$$\langle \text{while } [b]^{\ell} \text{ do } S, \sigma_2 \rangle \rightarrow \langle S; \text{while } [b]^{\ell} \text{ do } S, \sigma_2 \rangle$$

and again, since $N(\ell) = \text{live}_{\text{entry}}(\ell) \supseteq \text{live}_{\text{exit}}(\ell) = X(\ell)$ we have $\sigma_1 \sim_{X(\ell)} \sigma_2$ and then

$$\sigma_1 \sim_{N(\text{init}(S))} \sigma_2$$

follows from Lemma 2.20 because $(\ell, \text{init}(S)) \in \text{flow}(\text{while } [b]^{\ell} \text{ do } S)$.

The case $[wh_2]$. Then $\langle \text{while } [b]^{\ell} \text{ do } S, \sigma_1 \rangle \rightarrow \sigma_1$ because $\mathcal{B}[b]\sigma_1 = \text{false}$. Since $\sigma_1 \sim_{N(\ell)} \sigma_2$ and $N(\ell) \supseteq FV(b)$, we also have that $\mathcal{B}[b]\sigma_2 = \text{false}$ and thus:

$$\langle \text{while } [b]^{\ell} \text{ do } S, \sigma_2 \rangle \rightarrow \sigma_2$$

From the specification of $LV^{\subseteq}(S)$, we have $N(\ell) = \text{live}_{\text{entry}}(\ell) \supseteq \text{live}_{\text{exit}}(\ell) = X(\ell)$ and thus $\sigma_1 \sim_{X(\ell)} \sigma_2$.

This completes the proof. \blacksquare

Finally, we have an important corollary which lifts the previous result to program executions: (i) in the case where the derivation sequence has not yet terminated and (ii) in the case it has terminated:

Corollary 2.22 If $live \models LV^{\subseteq}(S)$ (with S being label consistent) then:

- (i) if $\langle S, \sigma_1 \rangle \rightarrow^* \langle S', \sigma'_1 \rangle$ and $\sigma_1 \sim_{N(\text{init}(S))} \sigma_2$ then there exists σ'_2 such that $\langle S, \sigma_2 \rangle \rightarrow^* \langle S', \sigma'_2 \rangle$ and $\sigma'_1 \sim_{N(\text{init}(S'))} \sigma'_2$, and

- (ii) if $\langle S, \sigma_1 \rangle \rightarrow^* \sigma'_1$ and $\sigma_1 \sim_{N(\text{init}(S))} \sigma_2$ then there exists σ'_2 such that $\langle S, \sigma_2 \rangle \rightarrow^* \sigma'_2$ and $\sigma'_1 \sim_{X(\ell)} \sigma'_2$ for some $\ell \in \text{final}(S)$. ■

Proof The proof is by induction on the length of the derivation sequence and uses Theorem 2.21. ■

Remark. We have now proved the correctness of Live Variables Analysis with respect to a small step operational semantics. Obviously, the correctness of the analysis can also be proved with respect to other kinds of semantics. However, note that if one relies on, say, a big step (or natural) semantics then it is not so obvious how to express (and prove) the correctness of looping computations: in a big step semantics a looping computation is modelled by the absence of an inference tree – in contrast to the small step semantics where it is modelled by an infinite derivation sequence. ■

2.3 Monotone Frameworks

Despite the differences between the analyses presented in Section 2.1, there are sufficient similarities to make it plausible that there might be an underlying framework. The advantages that accrue from identifying such a framework include the possibility of designing generic algorithms for solving the data flow equations, as we will see in Section 2.4.

The overall pattern. Each of the four classical analyses (presented in Subsection 2.1.1 to 2.1.4) considers equations for a label consistent program S_* and they take the form

$$\begin{aligned} \text{Analysis}_\circ(\ell) &= \begin{cases} \iota & \text{if } \ell \in E \\ \sqcup \{ \text{Analysis}_\bullet(\ell') \mid (\ell', \ell) \in F \} & \text{otherwise} \end{cases} \\ \text{Analysis}_\bullet(\ell) &= f_\ell(\text{Analysis}_\circ(\ell)) \end{aligned}$$

where

- \sqcup is \cap or \cup (and \sqcap is \cup or \cap),
- F is either $\text{flow}(S_*)$ or $\text{flow}^R(S_*)$,
- E is $\{\text{init}(S_*)\}$ or $\{\text{final}(S_*)\}$,
- ι specifies the initial or final analysis information, and
- f_ℓ is the transfer function associated with $B^\ell \in \text{blocks}(S_*)$.

We now have the following characterisation:

- The *forward analyses* have F to be $flow(S_*)$ and then $Analysis_o$ concerns entry conditions and $Analysis_e$ concerns exit conditions; also the equation system presupposes that S_* has isolated entries.
- The *backward analyses* have F to be $flow^R(S_*)$ and then $Analysis_o$ concerns exit conditions and $Analysis_e$ concerns entry conditions; also the equation system presupposes that S_* has isolated exits.

The principle we have seen emerging in Section 2.1 is that:

- When \sqcup is \cap we require the *greatest* sets that solve the equations and we are able to detect properties satisfied by *all* paths of execution reaching (or leaving) the entry (or exit) of a label; these analyses are often called *must analyses*.
- When \sqcup is \cup we require the *least* sets that solve the equations and we are able to detect properties satisfied by *at least one* execution path to (or from) the entry (or exit) of a label; these analyses are often called *may analyses*.

Remark. Some authors propose a typology for Data Flow Analysis, characterising each analysis by a triple from

$$\{\cap, \cup\} \times \{\rightarrow, \leftarrow\} \times \{\uparrow, \downarrow\}$$

where \rightarrow means forwards, \leftarrow means backwards, \downarrow means smallest and \uparrow means largest. This leads to eight possible types of analysis – a cube. In fact, given our association of \cap with \uparrow and \cup with \downarrow , the cube collapses to a square. We have presented analyses of the following four types: $(\cap, \rightarrow, \uparrow)$, $(\cup, \rightarrow, \downarrow)$, $(\cap, \leftarrow, \uparrow)$ and $(\cup, \leftarrow, \downarrow)$. ■

It is occasionally awkward to have to assume that forward analyses have isolated entries and that backward analyses have isolated exits. This motivates reformulating the above equations to be of the form

$$Analysis_o(\ell) = \bigsqcup \{Analysis_e(\ell') \mid (\ell', \ell) \in F\} \sqcup \iota_E^\ell$$

$$\text{where } \iota_E^\ell = \begin{cases} \iota & \text{if } \ell \in E \\ \perp & \text{if } \ell \notin E \end{cases}$$

$$Analysis_e(\ell) = f_\ell(Analysis_o(\ell))$$

where \perp satisfies $l \sqcup \perp = l$ (hence \perp is not really there). This formulation makes sense also for analyses that do not have isolated entries and exits.

In this section, we present a more formal approach to defining data flow frameworks that exploits the similarities that we have identified above. Nothing that we present in this section is dependent on the definition of elementary

blocks, or the programming language constructs; however, the techniques do not directly apply to languages with procedures (which will be addressed in Section 2.5). The view that we take here is that a program is a *transition system*; the nodes represent blocks and each block has a *transfer function* associated with it that specifies how the block acts on the “input” state. (Note that for forward analyses, the input state is the entry state, and for backward analyses, it is the exit state.)

2.3.1 Basic Definitions

Property spaces. One important ingredient in the framework is the *property space*, L , used to represent the data flow information as well as the *combination operator*, $\sqcup : \mathcal{P}(L) \rightarrow L$, that combines information from different paths; as usual $\sqcup : L \times L \rightarrow L$ is defined by $l_1 \sqcup l_2 = \sqcup\{l_1, l_2\}$ and we write \perp for $\sqcup \emptyset$. It is customary to demand that this property space is in fact a complete lattice; as discussed in Appendix A this just means that it is a partially ordered set, (L, \sqsubseteq) , such that each subset, Y , has a least upper bound, $\sqcup Y$. Looking ahead to the task of implementing the analysis one often requires that L satisfies the *Ascending Chain Condition*; as discussed in Appendix A this means that each ascending chain, $(l_n)_n$, i.e. $l_1 \sqsubseteq l_2 \sqsubseteq l_3 \sqsubseteq \dots$, eventually stabilises, i.e. $\exists n : l_n = l_{n+1} = \dots$.

Example 2.23 For Reaching Definitions we have $L = \mathcal{P}(\text{Var}_* \times \text{Lab}_*)$ and it is partially ordered by subset-inclusion, i.e. “ \sqsubseteq ” is “ \subseteq ”. Similarly, $\sqcup Y$ is $\bigcup Y$, $l_1 \sqcup l_2$ is $l_1 \cup l_2$, and \perp is \emptyset . That L satisfies the Ascending Chain Condition, i.e. that $l_1 \subseteq l_2 \subseteq \dots$ implies $\exists n : l_n = l_{n+1} = \dots$, follows because $\text{Var}_* \times \text{Lab}_*$ is finite (unlike $\text{Var} \times \text{Lab}$). ■

Example 2.24 For Available Expressions we have $L = \mathcal{P}(\text{AExp}_*)$ and it is partially ordered by superset inclusion, i.e. “ \sqsubseteq ” is “ \supseteq ”. Similarly, $\sqcup Y$ is $\bigcap Y$, $l_1 \sqcup l_2$ is $l_1 \cap l_2$, and \perp is AExp_* . That L satisfies the Ascending Chain Condition, i.e. that $l_1 \supseteq l_2 \supseteq \dots$ implies $\exists n : l_n = l_{n+1} = \dots$, follows because AExp_* is finite (unlike AExp). ■

Remark. Historically, the demands on the property space, L , have often been expressed in a different way. A *join semi-lattice* is a non-empty set, L , with a binary *join* operation, \sqcup , which is idempotent, commutative and associative, i.e. $l \sqcup l = l$, $l_1 \sqcup l_2 = l_2 \sqcup l_1$ and $(l_1 \sqcup l_2) \sqcup l_3 = l_1 \sqcup (l_2 \sqcup l_3)$. The commutativity and associativity of the operation mean that it does not matter in which order we combine information from different paths. The join operation induces a partial ordering, \sqsubseteq , on the elements by taking $l_1 \sqsubseteq l_2$ if and only if $l_1 \sqcup l_2 = l_2$. It is not hard to show that this in fact defines a partial ordering and that $l_1 \sqcup l_2$ is the least upper bound (with respect to \sqsubseteq). A unit for the join operation is an element, \perp , such that $\perp \sqcup l = l$. It is

not hard to show that the unit is in fact the least element (with respect to \sqsubseteq). It has been customary to demand that the property space, L , is a join semi-lattice with a unit and that it satisfies the Ascending Chain Condition. As proved in Lemma A.8 of Appendix A this is equivalent to our assumption that the property space, L , is a complete lattice satisfying the Ascending Chain Condition. ■

Some formulations of Monotone Frameworks are expressed in terms of property spaces satisfying a *Descending Chain Condition* and using a combination operator \sqcap . It follows from the principle of lattice duality (see the Concluding Remarks of Chapter 4) that this does not change the notion of Monotone Framework.

Transfer functions. Another important ingredient in the framework is the set of *transfer functions*, $f_\ell : L \rightarrow L$ for $\ell \in \text{Lab}_*$. It is natural to demand that each transfer function is monotone, i.e. $l \sqsubseteq l'$ implies $f_\ell(l) \sqsubseteq f_\ell(l')$. Intuitively, this says that an increase in our knowledge about the input must give rise to an increase in our knowledge about the output (or at the least that we know the same as before). Formally, we shall see that monotonicity is of importance for the algorithms we develop. To control the set of transfer functions we demand that there is a set \mathcal{F} of monotone functions over L , fulfilling the following conditions:

- \mathcal{F} contains all the transfer functions f_ℓ in question,
- \mathcal{F} contains the identity function, and
- \mathcal{F} is closed under composition of functions.

The condition on the identity function is natural because of the *skip* statement and the condition on composition of functions is natural because of the sequencing of statements. Clearly one can take \mathcal{F} to be the space of monotone functions over L but it is occasionally advantageous to consider a smaller set because it makes it easier to find compact representations of the functions.

Some formulations of Monotone Frameworks associate transfer functions with edges (or flows) rather than nodes (or labels). A similar effect can be obtained using the approach of Exercise 2.11.

Frameworks. In summary, a *Monotone Framework* consists of:

- a complete lattice, L , that satisfies the Ascending Chain Condition, and we write \sqcup for the least upper bound operator; and
- a set \mathcal{F} of monotone functions from L to L that contains the identity function and that is closed under function composition.

Note that we do *not* demand that \mathcal{F} is a complete lattice or even a partially ordered set although this is the case for the set of all monotone functions from L to L (see Appendix A).

A somewhat stronger concept is that of a *Distributive Framework*. This is a Monotone Framework where additionally all functions f in \mathcal{F} are required to be distributive:

$$f(l_1 \sqcup l_2) = f(l_1) \sqcup f(l_2)$$

Since $f(l_1 \sqcup l_2) \supseteq f(l_1) \sqcup f(l_2)$ follows from monotonicity, the only additional demand is that $f(l_1 \sqcup l_2) \subseteq f(l_1) \sqcup f(l_2)$. When this condition is fulfilled it is sometimes possible to get more efficient algorithms.

Instances. The data flow equations make it clear that more than just a Monotone (or Distributive) Framework is needed in order to specify an analysis. To this end we define an *instance*, *Analysis*, of a Monotone (or Distributive) Framework to consist of:

- the complete lattice, L , of the framework;
- the space of functions, \mathcal{F} , of the framework;
- a finite *flow*, F , that typically is $\text{flow}(S_*)$ or $\text{flow}^R(S_*)$;
- a finite set of so-called *extremal labels*, E , that typically is $\{\text{init}(S_*)\}$ or $\{\text{final}(S_*)\}$;
- an *extremal value*, $\iota \in L$, for the extremal labels; and
- a mapping, f ., from the labels Lab_* of F and E to transfer functions in \mathcal{F} .

The instance then gives rise to a *set of equations*, Analysis^\equiv , of the form considered earlier:

$$\begin{aligned} \text{Analysis}_o(\ell) &= \bigsqcup \{ \text{Analysis}_*(\ell') \mid (\ell', \ell) \in F \} \sqcup \iota_E^\ell \\ &\quad \text{where } \iota_E^\ell = \begin{cases} \iota & \text{if } \ell \in E \\ \perp & \text{if } \ell \notin E \end{cases} \\ \text{Analysis}_*(\ell) &= f_\ell(\text{Analysis}_o(\ell)) \end{aligned}$$

It also gives rise to a *set of constraints*, $\text{Analysis}^\sqsubseteq$, defined by:

$$\begin{aligned} \text{Analysis}_o(\ell) &\sqsupseteq \bigsqcup \{ \text{Analysis}_*(\ell') \mid (\ell', \ell) \in F \} \sqcup \iota_E^\ell \\ &\quad \text{where } \iota_E^\ell = \begin{cases} \iota & \text{if } \ell \in E \\ \perp & \text{if } \ell \notin E \end{cases} \\ \text{Analysis}_*(\ell) &\sqsupseteq f_\ell(\text{Analysis}_o(\ell)) \end{aligned}$$

	Available Expressions	Reaching Definitions	Very Busy Expressions	Live Variables
L	$\mathcal{P}(\mathbf{AExp}_*)$	$\mathcal{P}(\mathbf{Var}_* \times \mathbf{Lab}_*)$	$\mathcal{P}(\mathbf{AExp}_*)$	$\mathcal{P}(\mathbf{Var}_*)$
\sqsubseteq	\supseteq	\subseteq	\supseteq	\subseteq
\sqcup	\cap	\cup	\cap	\cup
\perp	\mathbf{AExp}_*	\emptyset	\mathbf{AExp}_*	\emptyset
ι	\emptyset	$\{(x, ?) \mid x \in FV(S_*)\}$	\emptyset	\emptyset
E	$\{\mathit{init}(S_*)\}$	$\{\mathit{init}(S_*)\}$	$\mathit{final}(S_*)$	$\mathit{final}(S_*)$
F	$\mathit{flow}(S_*)$	$\mathit{flow}(S_*)$	$\mathit{flow}^R(S_*)$	$\mathit{flow}^R(S_*)$
\mathcal{F}	$\{f : L \rightarrow L \mid \exists l_k, l_g : f(l) = (l \setminus l_k) \cup l_g\}$			
f_l	$f_l(l) = (l \setminus \mathit{kill}([B]^l)) \cup \mathit{gen}([B]^l)$ where $[B]^l \in \mathit{blocks}(S_*)$			

Figure 2.6: Instances for the four classical analyses.

2.3.2 The Examples Revisited

We now return to the four classical analyses from Section 2.1 and show how the analyses of a label consistent program, S_* , can be recast as an instance of a Monotone (in fact Distributive) Framework. We refer to Figure 2.6 for all the data needed to specify the Monotone Framework as well as the instance.

It is immediate that the property space, L , is a complete lattice in all cases. Given the choice of a partial ordering, \sqsubseteq , the information about the least element, \perp , and the least upper bound operation, \sqcup , is immediate. Note that we define \sqsubseteq to be \subseteq for those analyses where we used \cup (and require the least solution) in Section 2.1, and similarly, that we define \sqsubseteq to be \supseteq for those analyses where we used \cap (and required the greatest solution) in Section 2.1. To ensure that L satisfies the Ascending Chain Condition we have restricted the attention to the finite sets of expressions, labels and variables occurring in the program, S_* , under consideration.

The definition of the flow, F , is as one should expect: it is $\mathit{flow}(S_*)$ for forward analyses and $\mathit{flow}^R(S_*)$ for backward analyses. Similarly, the extremal labels, E are $\{\mathit{init}(S_*)\}$ for forward analyses and $\mathit{final}(S_*)$ for backward analyses. The only thing to note about the extremal value, ι , is that there seems to be no general pattern concerning how to define it: it is not always \top_L (nor is it always \perp_L).

It remains to show that the conditions on the set \mathcal{F} of transfer functions are

satisfied.

Lemma 2.25 Each of the four data flow analyses in Figure 2.6 is a Monotone Framework as well as a Distributive Framework. ■

Proof To prove that the analyses are Monotone Frameworks we just have to confirm that \mathcal{F} has the necessary properties:

The functions of \mathcal{F} are monotone: Assume that $l \sqsubseteq l'$. Then $(l \setminus l_k) \sqsubseteq (l' \setminus l_k)$ and, therefore $((l \setminus l_k) \cup l_g) \sqsubseteq ((l' \setminus l_k) \cup l_g)$ and thus $f(l) \sqsubseteq f(l')$ as required. Note that this calculation is valid regardless of whether \sqsubseteq is \subseteq or \supseteq .

The identity function is in \mathcal{F} : It is obtained by taking both l_k and l_g to be \emptyset .

The function of \mathcal{F} are closed under composition: Suppose $f(l) = (l \setminus l_k) \cup l_g$ and $f'(l) = (l \setminus l'_k) \cup l'_g$. Then we calculate:

$$\begin{aligned} (f \circ f')(l) &= (((l \setminus l'_k) \cup l'_g) \setminus l_k) \cup l_g \\ &= (l \setminus (l'_k \cup l_k)) \cup ((l'_g \setminus l_k) \cup l_g) \end{aligned}$$

So $(f \circ f')(l) = (l \setminus l''_k) \cup l''_g$ where $l''_k = l'_k \cup l_k$ and $l''_g = (l'_g \setminus l_k) \cup l_g$. This completes the proof of the first part of the lemma.

To prove that the analyses are Distributive Frameworks consider $f \in \mathcal{F}$ given by $f(l) = (l \setminus l_k) \cup l_g$. Then we have:

$$\begin{aligned} f(l \sqcup l') &= ((l \sqcup l') \setminus l_k) \cup l_g \\ &= ((l \setminus l_k) \sqcup (l' \setminus l_k)) \cup l_g \\ &= ((l \setminus l_k) \cup l_g) \sqcup ((l' \setminus l_k) \cup l_g) \\ &= f(l) \sqcup f(l') \end{aligned}$$

Note that the above calculation is valid regardless of whether \sqcup is \cup or \cap . This completes the proof. ■

It is worth pointing out that in order to get this result we have made the frameworks dependent upon the actual program – this is needed to enforce that the Ascending Chain Condition is fulfilled.

Example 2.26 Let us return to the Available Expressions Analysis of the program

$$[x:=a+b]^1; [y:=a*b]^2; \text{while } [y>a+b]^3 \text{ do } ([a:=a+1]^4; [x:=a+b]^5)$$

of Examples 2.4 and 2.5 and let us specify it as an instance of the associated Monotone Framework. The complete lattice of interest is

$$(\mathcal{P}(\{a+b, a*b, a+1\}), \supseteq)$$

with least element $\{a+b, a*b, a+1\}$. The set of transfer functions has the form shown in Figure 2.6.

The instance of the framework additionally has the flow $\{(1, 2), (2, 3), (3, 4), (4, 5), (5, 3)\}$ and the set of extremal labels is $\{1\}$. The extremal value is \emptyset and the transfer functions associated with the labels are

$$\begin{aligned} f_1^{AE}(Y) &= Y \cup \{a+b\} \\ f_2^{AE}(Y) &= Y \cup \{a*b\} \\ f_3^{AE}(Y) &= Y \cup \{a+b\} \\ f_4^{AE}(Y) &= Y \setminus \{a+b, a*b, a+1\} \\ f_5^{AE}(Y) &= Y \cup \{a+b\} \end{aligned}$$

for $Y \subseteq \{a+b, a*b, a+1\}$. ■

2.3.3 A Non-distributive Example

Lest the reader should imagine that all Monotone Frameworks are Distributive Frameworks, here we present one that is not. The *Constant Propagation Analysis* will determine:

For each program point, whether or not a variable has a constant value whenever execution reaches that point.

Such information can be used as the basis for an optimisation known as *Constant Folding*: all uses of the variable may be replaced by the constant value.

The Constant Propagation framework. The complete lattice used for Constant Propagation Analysis of a program, S_* , is

$$\widehat{\text{State}}_{CP} = ((\text{Var}_* \rightarrow \mathbf{Z}^T)_{\perp}, \sqsubseteq, \cup, \cap, \perp, \lambda x. \top)$$

where Var_* is the set of variables appearing in the program and $\mathbf{Z}^T = \mathbf{Z} \cup \{\top\}$ is partially ordered as follows:

$$\begin{aligned} \forall z \in \mathbf{Z}^T : z \sqsubseteq \top \\ \forall z_1, z_2 \in \mathbf{Z} : (z_1 \sqsubseteq z_2) \Leftrightarrow (z_1 = z_2) \end{aligned}$$

The top element of \mathbf{Z}^T is used to indicate that a variable is non-constant and all other elements indicate that the value is that particular constant. The idea is that an element $\hat{\sigma}$ of $\text{Var}_* \rightarrow \mathbf{Z}^T$ is a property state: for each variable x , $\hat{\sigma}(x)$ will give information about whether or not x is a constant and in the latter case which constant.

To capture the case where no information is available we extend $\text{Var}_* \rightarrow \mathbf{Z}^T$ with a least element \perp , written $(\text{Var}_* \rightarrow \mathbf{Z}^T)_{\perp}$. The partial ordering \sqsubseteq on

$$\begin{array}{c}
\mathcal{A}_{\text{CP}} : \mathbf{AExp} \rightarrow (\widehat{\text{State}}_{\text{CP}} \rightarrow \mathbf{Z}_{\perp}^{\top}) \\
\hline
\mathcal{A}_{\text{CP}}[x]\hat{\sigma} = \begin{cases} \perp & \text{if } \hat{\sigma} = \perp \\ \hat{\sigma}(x) & \text{otherwise} \end{cases} \\
\mathcal{A}_{\text{CP}}[n]\hat{\sigma} = \begin{cases} \perp & \text{if } \hat{\sigma} = \perp \\ n & \text{otherwise} \end{cases} \\
\mathcal{A}_{\text{CP}}[a_1 \text{ op}_a a_2]\hat{\sigma} = \mathcal{A}_{\text{CP}}[a_1]\hat{\sigma} \widehat{\text{op}}_a \mathcal{A}_{\text{CP}}[a_2]\hat{\sigma} \\
\hline
\text{transfer functions: } f_{\ell}^{\text{CP}} \\
\hline
[x := a]^{\ell} : f_{\ell}^{\text{CP}}(\hat{\sigma}) = \begin{cases} \perp & \text{if } \hat{\sigma} = \perp \\ \hat{\sigma}[x \mapsto \mathcal{A}_{\text{CP}}[a]\hat{\sigma}] & \text{otherwise} \end{cases} \\
[\text{skip}]^{\ell} : f_{\ell}^{\text{CP}}(\hat{\sigma}) = \hat{\sigma} \\
[b]^{\ell} : f_{\ell}^{\text{CP}}(\hat{\sigma}) = \hat{\sigma} \\
\hline
\end{array}$$

Table 2.7: Constant Propagation Analysis.

$\widehat{\text{State}}_{\text{CP}} = (\text{Var}_{\star} \rightarrow \mathbf{Z}^{\top})_{\perp}$ is defined by

$$\begin{aligned}
\forall \hat{\sigma} \in (\text{Var}_{\star} \rightarrow \mathbf{Z}^{\top})_{\perp} : \quad & \perp \sqsubseteq \hat{\sigma} \\
\forall \hat{\sigma}_1, \hat{\sigma}_2 \in \text{Var}_{\star} \rightarrow \mathbf{Z}^{\top} : \quad & \hat{\sigma}_1 \sqsubseteq \hat{\sigma}_2 \text{ iff } \forall x : \hat{\sigma}_1(x) \sqsubseteq \hat{\sigma}_2(x)
\end{aligned}$$

and the binary least upper bound operation is then:

$$\begin{aligned}
\forall \hat{\sigma} \in (\text{Var}_{\star} \rightarrow \mathbf{Z}^{\top})_{\perp} : \quad & \hat{\sigma} \sqcup \perp = \hat{\sigma} = \perp \sqcup \hat{\sigma} \\
\forall \hat{\sigma}_1, \hat{\sigma}_2 \in \text{Var}_{\star} \rightarrow \mathbf{Z}^{\top} : \quad & \forall x : (\hat{\sigma}_1 \sqcup \hat{\sigma}_2)(x) = \hat{\sigma}_1(x) \sqcup \hat{\sigma}_2(x)
\end{aligned}$$

In contrast to the earlier examples, we define the transfer functions as follows:

$$\mathcal{F}_{\text{CP}} = \{f \mid f \text{ is a monotone function on } \widehat{\text{State}}_{\text{CP}}\}$$

It is easy to verify that $\widehat{\text{State}}_{\text{CP}}$ and \mathcal{F}_{CP} satisfy the requirements of being a Monotone Framework (see Exercise 2.8).

Constant Propagation is a forward analysis, so for the program S_{\star} we take the flow, F , to be $\text{flow}(S_{\star})$, the extremal labels, E , to be $\{\text{init}(S_{\star})\}$, the extremal value, ι_{CP} , to be $\lambda x. \perp$, and the mapping, f^{CP} , of labels to transfer functions is given in Table 2.7. The specification of the transfer functions uses the function

$$\mathcal{A}_{\text{CP}} : \mathbf{AExp} \rightarrow (\widehat{\text{State}}_{\text{CP}} \rightarrow \mathbf{Z}_{\perp}^{\top})$$

for analysing expressions. Here the operations on \mathbf{Z} are lifted to $\mathbf{Z}_{\perp}^{\top} = \mathbf{Z} \cup \{\perp, \top\}$ by taking $z_1 \widehat{\text{op}}_a z_2 = z_1 \text{op}_a z_2$ if $z_1, z_2 \in \mathbf{Z}$ (and where op_a is the corresponding arithmetic operation on \mathbf{Z}), $z_1 \widehat{\text{op}}_a z_2 = \perp$ if $z_1 = \perp$ or $z_2 = \perp$ and $z_1 \widehat{\text{op}}_a z_2 = \top$ otherwise.

Lemma 2.27 Constant Propagation is a Monotone Framework that is *not* a Distributive Framework. ■

Proof The proof that Constant Propagation is a Monotone Framework is left for Exercise 2.8. To show that it is not a Distributive Framework consider the transfer function f_t^{CP} for $[y:=x*x]^t$ and let $\hat{\sigma}_1$ and $\hat{\sigma}_2$ be such that $\hat{\sigma}_1(x) = 1$ and $\hat{\sigma}_2(x) = -1$. Then $\hat{\sigma}_1 \sqcup \hat{\sigma}_2$ maps x to \top and thus $f_t^{CP}(\hat{\sigma}_1 \sqcup \hat{\sigma}_2)$ maps y to \top and hence fails to record that y has the constant value 1. However, both $f_t^{CP}(\hat{\sigma}_1)$ and $f_t^{CP}(\hat{\sigma}_2)$ map y to 1 and so does $f_t^{CP}(\hat{\sigma}_1) \sqcup f_t^{CP}(\hat{\sigma}_2)$. ■

Correctness of the analysis will be established in Section 4.5.

2.4 Equation Solving

Having set up a framework, there remains the question of how to *use* the framework to obtain an analysis result. In this section we shall consider two approaches. One is an iterative algorithm in the spirit of Chaotic Iteration as presented in Section 1.7. The other more directly propagates analysis information along paths in the program.

2.4.1 The MFP Solution

We first present a general iterative algorithm for Monotone Frameworks that computes the least solution to the data flow equations. Historically, this is called the *MFP solution* (for *Maximal Fixed Point*) although it in fact computes the least fixed point; the reason is that the classical literature tends to focus on analyses where \sqcup is \cap rather than \cup (and because the least fixed point with respect to \supseteq equals the greatest fixed point with respect to \subseteq).

The algorithm, written in pseudo-code in Table 2.8, takes as input an instance of a Monotone Framework. It uses an array, *Analysis*, which contains the *Analysis_o* information for each elementary block; the array is indexed by labels. It also uses a *worklist* W being a list of pairs; each pair is an element of the flow relation F . The presence of a pair in the worklist indicates that the analysis has changed at the exit of (or entry to – for backward analyses) the block labelled by the first component and so must be recomputed at the entry to (or exit from) the block labelled by the second component. As a final stage the algorithm presents the result (MFP_o, MFP_\bullet) of the analysis in a form close to the formulation of the data flow equations.

Example 2.28 To illustrate how the algorithm works let us return to Example 2.26 where we consider the program

$$[x:=a+b]^1; [y:=a*b]^2; \text{while } [y>a+b]^3 \text{ do } ([a:=a+1]^4; [x:=a+b]^5)$$

INPUT: An instance of a Monotone Framework:
 $(L, \mathcal{F}, F, E, \iota, f.)$

OUTPUT: MFP_o, MFP_\bullet .

METHOD: Step 1: Initialisation (of W and Analysis)
 $W := \text{nil};$
for all (ℓ, ℓ') in F do
 $W := \text{cons}((\ell, \ell'), W);$
for all ℓ in F or E do
if $\ell \in E$ then $\text{Analysis}[\ell] := \iota$
else $\text{Analysis}[\ell] := \perp_L;$

Step 2: Iteration (updating W and Analysis)
while $W \neq \text{nil}$ do
 $\ell := \text{fst}(\text{head}(W)); \ell' = \text{snd}(\text{head}(W));$
 $W := \text{tail}(W);$
if $f_\ell(\text{Analysis}[\ell]) \not\sqsubseteq \text{Analysis}[\ell']$ then
 $\text{Analysis}[\ell'] := \text{Analysis}[\ell'] \sqcup f_\ell(\text{Analysis}[\ell]);$
for all ℓ'' with (ℓ', ℓ'') in F do
 $W := \text{cons}((\ell', \ell''), W);$

Step 3: Presenting the result (MFP_o and MFP_\bullet)
for all ℓ in F or E do
 $MFP_o(\ell) := \text{Analysis}[\ell];$
 $MFP_\bullet(\ell) := f_\ell(\text{Analysis}[\ell])$

Table 2.8: Algorithm for solving data flow equations.

Writing W for the list $((2,3),(3,4),(4,5),(5,3))$ and U for the set $\{a+b, a*b, a+1\}$, step 1 of the algorithm will initialise the data structures as in the first row in Table 2.9. Step 2 will inspect the first element of the worklist and rows 2–7 represent cases where there is a change in the array Analysis and hence a new pair is placed on top of the worklist; it is inspected in the next iteration. Rows 8–12 represent cases where no modification is made in the array and hence the worklist is getting smaller – the elements of W are now inspected. Step 3 will then produce the solution we already saw in Example 2.5. ■

Properties of the algorithm. We shall first show that the algorithm computes the expected solution to the equation system.

Lemma 2.29 The worklist algorithm in Table 2.8 always terminates and it computes the least (or MFP) solution to the instance of the framework

	W	Analysis[ℓ] for ℓ being				
		1	2	3	4	5
1	$((1,2),W)$	\emptyset	U	U	U	U
2	$((2,3),W)$	\emptyset	$\{a+b\}$	U	U	U
3	$((3,4),W)$	\emptyset	$\{a+b\}$	$\{a+b,a*b\}$	U	U
4	$((4,5),W)$	\emptyset	$\{a+b\}$	$\{a+b,a*b\}$	$\{a+b,a*b\}$	U
5	$((5,3),W)$	\emptyset	$\{a+b\}$	$\{a+b,a*b\}$	$\{a+b,a*b\}$	\emptyset
6	$((3,4),W)$	\emptyset	$\{a+b\}$	$\{a+b\}$	$\{a+b,a*b\}$	\emptyset
7	$((4,5),W)$	\emptyset	$\{a+b\}$	$\{a+b\}$	$\{a+b\}$	\emptyset
8	$((2,3),\dots)$	\emptyset	$\{a+b\}$	$\{a+b\}$	$\{a+b\}$	\emptyset
9	$((3,4),\dots)$	\emptyset	$\{a+b\}$	$\{a+b\}$	$\{a+b\}$	\emptyset
10	$((4,5),\dots)$	\emptyset	$\{a+b\}$	$\{a+b\}$	$\{a+b\}$	\emptyset
11	$((5,3))$	\emptyset	$\{a+b\}$	$\{a+b\}$	$\{a+b\}$	\emptyset
12	$()$	\emptyset	$\{a+b\}$	$\{a+b\}$	$\{a+b\}$	\emptyset

Table 2.9: Iteration steps of the worklist algorithm

given as input. ■

Proof First we prove the termination result. Step 1 and 3 are bounded loops over finite sets and thus trivially terminate. Next consider step 2. Assume that there are b labels in the program. Then the worklist initially has at most b^2 elements; the worst case is that F associates every label to every label. Each iteration either deletes an element from the worklist or adds up to b new elements. New elements are added if for the pair selected in this iteration, (ℓ, ℓ') , we have $f_{\ell}(\text{Analysis}[\ell]) \not\subseteq \text{Analysis}[\ell']$; that is, $f_{\ell}(\text{Analysis}[\ell]) \sqsupset \text{Analysis}[\ell']$ or they are incomparable. In either case, the new value of $\text{Analysis}[\ell']$ is strictly greater than the previous one. Since the set of values satisfies the Ascending Chain Condition, this can only happen a finite number of times. Thus the worklist will eventually be exhausted.

Next we prove the correctness result. Let Analysis_o and Analysis_s be the least solution to the instance given as input to the algorithm. The proof is now in three parts: (i) first we show that on each iteration the values in Analysis are approximations to the corresponding values of Analysis_o , (ii) then we show that Analysis_o is an approximation to Analysis at the termination of step 2 of the algorithm, and (iii) finally we combine these results.

Part (i). We show that

$$\forall \ell : \text{Analysis}[\ell] \subseteq \text{Analysis}_o(\ell)$$

is an invariant of the loop of step 2. After step 1 we have $\text{Analysis}[\ell] \subseteq \text{Analysis}_o(\ell)$ for all ℓ because $\text{Analysis}_o(\ell) \supseteq \iota$ whenever $\ell \in E$. After each iteration through the loop either there is no change because the iteration just deletes an element from the worklist or else $\text{Analysis}[\ell']$ is unchanged for all ℓ'' except for some ℓ' . In that case there is some ℓ such that $(\ell, \ell') \in F$ and

$$\text{newAnalysis}[\ell'] = \text{oldAnalysis}[\ell'] \sqcup f_{\ell}(\text{oldAnalysis}[\ell])$$

$$\begin{aligned} &\sqsubseteq \text{Analysis}_o(\ell') \sqcup f_\ell(\text{Analysis}_o(\ell)) \\ &= \text{Analysis}_o(\ell') \end{aligned}$$

The inequality follows since f_ℓ is monotone and the last equation follows from $(\text{Analysis}_o, \text{Analysis}_o)$ being a solution to the instance.

Part (ii). On termination of the loop, the worklist is empty. We show that

$$\forall \ell, \ell' : (\ell, \ell') \in F \Rightarrow \text{Analysis}[\ell'] \sqsupseteq f_\ell(\text{Analysis}[\ell])$$

by contradiction. So suppose that $\text{Analysis}[\ell'] \not\sqsupseteq f_\ell(\text{Analysis}[\ell])$ for some $(\ell, \ell') \in F$ and let us obtain a contradiction. Consider the last time that $\text{Analysis}[\ell]$ was updated. If this was in step 1 we considered (ℓ, ℓ') in step 2 and ensured that

$$\text{Analysis}[\ell'] \sqsupseteq f_\ell(\text{Analysis}[\ell])$$

and this invariant has been maintained ever since; hence this case cannot apply. It follows that $\text{Analysis}[\ell]$ was last updated in step 2. But at that time (ℓ, ℓ') was placed in the worklist once again. When considering (ℓ, ℓ') in step 2 we then ensured that

$$\text{Analysis}[\ell'] \sqsupseteq f_\ell(\text{Analysis}[\ell])$$

and this invariant has been maintained ever since; hence this case cannot apply either. This completes the proof by contradiction.

On termination of the loop we have:

$$\forall \ell \in E : \text{Analysis}[\ell] \sqsupseteq \iota$$

This follows because it was established in step 1 and it is maintained ever since. Thus it follows that at termination of step 2:

$$\forall \ell : \text{Analysis}[\ell] \sqsupseteq \left(\bigsqcup \{f_{\ell'}(\text{Analysis}[\ell']) \mid (\ell', \ell) \in F\} \right) \sqcup \iota_E^{\ell}$$

Part (iii). By our assumptions and Proposition A.10 we have

$$\forall \ell : \text{MFP}_o(\ell) \sqsupseteq \text{Analysis}_o(\ell)$$

since $\text{Analysis}_o(\ell)$ is the least solution to the above constraint system and MFP_o equals the final value of Analysis . Together with part (i) this proves that

$$\forall \ell : \text{MFP}_o[\ell] = \text{Analysis}_o(\ell)$$

upon termination of step 2. ■

Based on the proof of termination in Lemma 2.29 we can determine an upper bound on the number of basic operations (for example an application of f_ℓ , an application of \sqcup , or an update of Analysis) performed by the algorithm. For this we shall assume that the flow F is represented in such a way (for example an array of lists) that all (ℓ', ℓ'') emanating from ℓ' can be found in time proportional to their number. Suppose that E and F contain at most $b \geq 1$ distinct labels, that F contains at most $e \geq b$ pairs, and that L has

finite height at most $h \geq 1$. Then steps 1 and 3 perform at most $O(b + e)$ basic operations. Concerning step 2 a pair is placed on the worklist at most $O(h)$ times, and each time it takes only a constant number of basic steps to process it – not counting the time needed to add new pairs to W ; this yields at most $O(e \cdot h)$ basic operations for step 2. Since $h \geq 1$ and $e \geq b$ this gives at most $O(e \cdot h)$ basic operations for the algorithm. (Since $e \leq b^2$ a potentially coarser bound is $O(b^2 \cdot h)$.)

Example 2.30 Consider the Reaching Definitions Analysis and suppose that there are at most $v \geq 1$ variables and $b \geq 1$ labels in the program, S_* , being analysed. Since $L = \mathcal{P}(\mathbf{Var}_* \times \mathbf{Lab}_*)$, it follows that $h \leq v \cdot b$ and thus we have an $O(v \cdot b^3)$ upper bound on the number of basic operations.

Actually we can do better. If S_* is label consistent then the variable of the pairs (x, ℓ) of $\mathcal{P}(\mathbf{Var}_* \times \mathbf{Lab}_*)$ will always be uniquely determined by the label ℓ so we get an $O(b^3)$ upper bound on the number of basic operations. Furthermore, F is $\text{flow}(S_*)$ and inspection of the equations for $\text{flow}(S_*)$ shows that for each label ℓ we construct at most two pairs with ℓ in the first component. This means that $e \leq 2 \cdot b$ and we get an $O(b^2)$ upper bound on the number of basic operations. ■

2.4.2 The MOP Solution

Let us now consider the other solution method for Monotone Frameworks where we more directly propagate analysis information along paths in the program. Historically, this is called the *MOP solution* (for *Meet Over all Paths*) although we do in fact take the join (or least upper bound) over all paths leading to an elementary block; once again the reason is that the classical literature focuses on analyses where \sqcup is \cap rather than \cup .

Paths. For the moment, we adopt the informal notion of a *path* to the entry of a block as the list of blocks traversed from the start of the program up to that block (but not including it); analogously, we can define a path from an exit of the block. Data Flow Analyses determine properties of such paths. Forward analyses concern paths from the initial block to the entry of a block; backward analyses concern paths from the exit of a block to a final block. The effect of a path on the state can be computed by composing the transfer functions associated with the individual blocks in the path. In the forward case we collect information about the state of affairs before the block is executed and in the backward case we collect information about the state of affairs immediately after the block has been executed. This informal description contrasts with the approach taken in Section 2.1 and earlier in this section; there we presented equations which were defined in terms of the immediate predecessors (successors) of a block (as defined by the flow and flow^R functions). We will see later that, for a large class of analyses, these

two approaches coincide.

For the formal development let us consider an instance $(L, \mathcal{F}, F, E, \iota, f.)$ of a Monotone Framework. We shall use the notation $\vec{\ell} = [\ell_1, \dots, \ell_n]$ for a sequence of $n \geq 0$ labels. We then define two sets of paths. The paths up to *but not* including ℓ are

$$\text{path}_o(\ell) = \{[\ell_1, \dots, \ell_{n-1}] \mid n \geq 1 \wedge \forall i < n : (\ell_i, \ell_{i+1}) \in F \wedge \ell_n = \ell \wedge \ell_1 \in E\}$$

and the paths up to *and* including ℓ are:

$$\text{path}_\bullet(\ell) = \{[\ell_1, \dots, \ell_n] \mid n \geq 1 \wedge \forall i < n : (\ell_i, \ell_{i+1}) \in F \wedge \ell_n = \ell \wedge \ell_1 \in E\}$$

For a path $\vec{\ell} = [\ell_1, \dots, \ell_n]$ we define the transfer function

$$f_{\vec{\ell}} = f_{\ell_n} \circ \dots \circ f_{\ell_1} \circ id$$

so that for the empty path we have $f_{[]} = id$ where id is the identity function.

By analogy with the definition of solutions to the equation system, in particular $MFP_o(\ell)$ and $MFP_\bullet(\ell)$, we now define two components of the MOP solution. The solution up to *but not* including ℓ is

$$MOP_o(\ell) = \bigsqcup \{f_{\vec{\ell}}(\iota) \mid \vec{\ell} \in \text{path}_o(\ell)\}$$

and the solution up to *and* including ℓ is:

$$MOP_\bullet(\ell) = \bigsqcup \{f_{\vec{\ell}}(\iota) \mid \vec{\ell} \in \text{path}_\bullet(\ell)\}$$

Unfortunately, the MOP solution is sometimes uncomputable (meaning that it is undecidable) even though the MFP solution is always easily computable (because of the property space satisfying the Ascending Chain Condition); the following result establishes one such result:

Lemma 2.31 The MOP solution for Constant Propagation is undecidable. ■

Proof Let u_1, \dots, u_n and v_1, \dots, v_n be strings over the alphabet $\{1, \dots, 9\}$ (see Appendix C). The *Modified Post Correspondence Problem* is to determine whether or not there exists a sequence i_1, \dots, i_m with $i_1 = 1$ such that $u_{i_1} \dots u_{i_m} = v_{i_1} \dots v_{i_m}$.

Let $|u|$ denote the length of the string u and let $[u]$ be its value interpreted as a natural number. Consider the program (omitting most labels)

```

x:=[u1]; y:=[v1];
while [...] do
  (if [...] then x:=x * 10|u1| + [u1]; y:=y * 10|v1| + [v1] else
  :
  if [...] then x:=x * 10|un| + [un]; y:=y * 10|vn| + [vn] else skip)
[z:=abs((x-y)*(x-y))]

```


where abs gives the absolute value (which is 1 for a positive argument and 0 or -1 otherwise) and where the details of $[\dots]$ are of no concern to us (and so could be taken to be $[\text{true}]$).

Then $MOP_o(\ell)$ will map z to 1 if and only if the Modified Post Correspondence Problem has no solution. Since the Modified Post Correspondence problem is undecidable [57] so is the MOP solution for Constant Propagation (assuming that our selection of arithmetic operation does indeed allow those used to be defined). ■

MOP versus MFP solutions. We shall shortly prove that the MFP solution safely approximates the MOP solution (informally, $MFP \sqsupseteq MOP$). In the case of a $(\cap, \rightarrow, \uparrow)$ or $(\cap, \leftarrow, \uparrow)$ analysis, the MFP solution is a subset of the MOP solution (\sqsupseteq is \subseteq); in the case of a $(\cup, \rightarrow, \downarrow)$ or $(\cup, \leftarrow, \downarrow)$ analysis, the MFP solution is a superset of the MOP solution. We can also show that, in the case of Distributive Frameworks, the MOP and MFP solutions coincide.

Lemma 2.32 Consider the MFP and MOP solutions to an instance $(L, \mathcal{F}, F, B, \iota, f)$ of a Monotone Framework; then:

$$MFP_o \sqsupseteq MOP_o \text{ and } MFP_o \sqsupseteq MOP_o.$$

If the framework is distributive and if $\text{path}_o(\ell) \neq \emptyset$ for all ℓ in E and F then:

$$MFP_o = MOP_o \text{ and } MFP_o = MOP_o. \quad \blacksquare$$

Proof It is straightforward to show that:

$$\begin{aligned} \forall \ell : MOP_o(\ell) &\subseteq f_\ell(MOP_o(\ell)) \\ \forall \ell : MFP_o(\ell) &= f_\ell(MFP_o(\ell)) \end{aligned}$$

For the first part of the lemma it therefore suffices to prove that:

$$\forall \ell : MOP_o(\ell) \subseteq MFP_o(\ell)$$

Note that MFP_o is the least fixed point of the functional F defined by:

$$F(A_o)(\ell) = \left(\bigsqcup \{f_{\ell'}(A_o(\ell')) \mid (\ell', \ell) \in F\} \right) \sqcup \iota_E^\ell$$

Next let us restrict the length of the paths used to compute MOP_o ; for $n \geq 0$ define:

$$MOP_o^n(\ell) = \bigsqcup \{f_{\vec{\ell}}(\iota) \mid \vec{\ell} \in \text{path}_o(\ell), |\vec{\ell}| < n\}$$

Clearly, $MOP_o(\ell) = \bigsqcup_n MOP_o^n(\ell)$ and to prove $MFP_o \sqsupseteq MOP_o$ it therefore suffices to prove

$$\forall n : MFP_o \sqsupseteq MOP_o^n$$

and we do so by numerical induction. The basis, $MFP_o \sqsupseteq MOP_o^0$, is trivial. The inductive step proceeds as follows:

$$\begin{aligned}
MFP_o(\ell) &= F(MFP_o)(\ell) \\
&= (\bigsqcup \{f_{\ell'}(MFP_o(\ell')) \mid (\ell', \ell) \in F\}) \sqcup \iota_E^{\ell} \\
&\sqsupseteq (\bigsqcup \{f_{\ell'}(MOP_o^n(\ell')) \mid (\ell', \ell) \in F\}) \sqcup \iota_E^{\ell} \\
&= (\bigsqcup \{f_{\ell'}(\bigsqcup \{f_{\vec{\ell}}(\iota) \mid \vec{\ell} \in \text{path}_o(\ell'), |\vec{\ell}| < n\}) \mid (\ell', \ell) \in F\}) \sqcup \iota_E^{\ell} \\
&\sqsupseteq (\bigsqcup \{(\bigsqcup \{f_{\ell'}(f_{\vec{\ell}}(\iota)) \mid \vec{\ell} \in \text{path}_o(\ell'), |\vec{\ell}| < n\}) \mid (\ell', \ell) \in F\}) \sqcup \iota_E^{\ell} \\
&= \bigsqcup \{(\bigsqcup \{f_{\vec{\ell}}(\iota) \mid \vec{\ell} \in \text{path}_o(\ell), 1 \leq |\vec{\ell}| \leq n\}) \sqcup \iota_E^{\ell} \\
&= MOP_o^{n+1}(\ell)
\end{aligned}$$

where we have used the induction hypothesis to get the first inequality. This completes the proof of $MFP_o \sqsupseteq MOP_o$ and $MFP_{\bullet} \sqsupseteq MOP_{\bullet}$.

To prove the second part of the lemma we now assume that the framework is distributive. Consider ℓ in E or F . By assumption f_{ℓ} is distributive, that is $f_{\ell}(\iota_1 \sqcup \iota_2) = f_{\ell}(\iota_1) \sqcup f_{\ell}(\iota_2)$, and from Lemma A.9 of Appendix A it follows that

$$f_{\ell}(\bigsqcup Y) = \bigsqcup \{f_{\ell}(\iota) \mid \iota \in Y\}$$

whenever Y is non-empty. By assumption we also have $\text{path}_o(\ell) \neq \emptyset$ and it follows that

$$\begin{aligned}
f_{\ell}(\bigsqcup \{f_{\vec{\ell}}(\iota) \mid \vec{\ell} \in \text{path}_o(\ell)\}) &= \bigsqcup \{f_{\ell}(f_{\vec{\ell}}(\iota)) \mid \vec{\ell} \in \text{path}_o(\ell)\} \\
&= \bigsqcup \{f_{\vec{\ell}}(\iota) \mid \vec{\ell} \in \text{path}_{\bullet}(\ell)\}
\end{aligned}$$

and this shows that:

$$\forall \ell : f_{\ell}(MOP_o(\ell)) = MOP_{\bullet}(\ell)$$

Next we calculate:

$$\begin{aligned}
MOP_o(\ell) &= \bigsqcup \{f_{\vec{\ell}}(\iota) \mid \vec{\ell} \in \text{path}_o(\ell)\} \\
&= \bigsqcup \{f_{\vec{\ell}}(\iota) \mid \vec{\ell} \in \bigcup \{\text{path}_o(\ell') \mid (\ell', \ell) \in F\} \cup \{\iota \mid \ell \in E\}\} \\
&= \bigsqcup \{(\bigsqcup \{f_{\ell'}(f_{\vec{\ell}}(\iota)) \mid \vec{\ell} \in \text{path}_o(\ell'), (\ell', \ell) \in F\} \cup \{\iota \mid \ell \in E\}) \\
&= (\bigsqcup \{f_{\ell'}(\bigsqcup \{f_{\vec{\ell}}(\iota) \mid \vec{\ell} \in \text{path}_o(\ell')\}) \mid (\ell', \ell) \in F\}) \sqcup \iota_E^{\ell} \\
&= (\bigsqcup \{f_{\ell'}(MOP_o(\ell')) \mid (\ell', \ell) \in F\}) \sqcup \iota_E^{\ell}
\end{aligned}$$

Together this shows that (MOP_o, MOP_{\bullet}) is a solution to the data flow equations. Using Proposition A.10 of Appendix A and the fact that (MFP_o, MFP_{\bullet}) is the least solution we get $MOP_o \sqsupseteq MFP_o$ and $MOP_{\bullet} \sqsupseteq MFP_{\bullet}$. Together with the results of the first part of the lemma we get $MOP_o = MFP_o$ and $MOP_{\bullet} = MFP_{\bullet}$. ■

We shall leave it to Exercise 2.13 to show that the condition that $path_o(\ell) \neq \emptyset$ (for ℓ in E and F) does hold when the Monotone Framework is constructed from a program S_* in the manner of the earlier sections.

It is sometimes stated that the MOP solution is the desired solution and that one only uses the MFP solution because the MOP solution might not be computable. In order to validate this we would need to prove that the MOP solution is semantically correct as was proved for the MFP solution in Section 2.2 in the case of Live Variables Analysis – in the case of Live Variables this is of course immediate since it is a Distributive Framework. However, it is always possible to formulate the MOP solution as an MFP solution over a different property space (like $\mathcal{P}(L)$) and therefore little is lost by focusing on the fixed point approach to Monotone Frameworks.

2.5 Interprocedural Analysis

The Data Flow Analysis techniques that have been presented in the previous sections are called *intraprocedural analyses* because they deal with simple languages without functions or procedures. It is somewhat more demanding to perform *interprocedural analyses* where functions and procedures are taken into account. Complications arise when ensuring that calls and returns match one another, when dealing with parameter mechanisms (and the aliasing that may result from call-by-reference) and when allowing procedures as parameters.

In this section we shall introduce some of the key techniques of interprocedural analysis. To keep things simple we just extend the WHILE language with top-level declarations of global mutually recursive procedures having a *call-by-value* parameter and a *call-by-result* parameter. The extension of the techniques to a language where procedures may have multiple call-by-value, call-by-result and call-by-value-result parameters is straightforward and so is the extension with local variable declarations (see Exercise 2.20) and we shall allow to use these extensions in examples.

Syntax of the procedure language. A program, P_* , in the extended WHILE-language has the form

$$\text{begin } D_* S_* \text{ end}$$

where D_* is a sequence of procedure declarations:

$$D ::= \text{proc } p(\text{val } x, \text{res } y) \text{ is}^{\ell_n} S \text{ end}^{\ell_x} \mid D D$$

Procedure names (denoted p) are syntactically distinct from variables (denoted x and y). The label ℓ_n of *is* marks the entry to the procedure body and the label ℓ_x of *end* marks the exit from the procedure body. The syntax

of statements is extended with:

$$S ::= \dots \mid [\text{call } p(a, z)]_{\ell_r}^{\ell_c}$$

The call statement has two labels: ℓ_c will be used for the call of the procedure and ℓ_r will be used for the associated return; the actual parameters are a and z .

The language is statically scoped, the parameter mechanism is call-by-value for the first parameter and call-by-result for the second parameter and the procedures may be mutually recursive. We shall *assume* throughout that the program is uniquely labelled (and hence label consistent); also we shall *assume* that only procedures that have been declared in D_* are ever called and that D_* does not contain two definitions of the same procedure name.

Example 2.33 Consider the following program calculating the Fibonacci number of the positive integer stored in x and returning it in y :

```
begin  proc fib(val z, u, res v) is1
        if [z<3]2 then [v:=u+1]3
          else ([call fib(z-1,u,v)]54; [call fib(z-2,v,v)]76)
        end8;
        [call fib(x,0,y)]109
end
```

It uses the procedure `fib` that returns in v the Fibonacci number of z plus the value of u . Both x and y are global variables whereas z , u and v are formal parameters and hence local variables. ■

Flow graphs for statements. The next step is to extend the definitions of the functions *init*, *final*, *blocks*, *labels*, and *flow* to specify the flow graphs also for the procedure language. For the new statement we take:

$$\begin{aligned} \text{init}([\text{call } p(a, z)]_{\ell_r}^{\ell_c}) &= \ell_c \\ \text{final}([\text{call } p(a, z)]_{\ell_r}^{\ell_c}) &= \{\ell_r\} \\ \text{blocks}([\text{call } p(a, z)]_{\ell_r}^{\ell_c}) &= \{[\text{call } p(a, z)]_{\ell_r}^{\ell_c}\} \\ \text{labels}([\text{call } p(a, z)]_{\ell_r}^{\ell_c}) &= \{\ell_c, \ell_r\} \\ \text{flow}([\text{call } p(a, z)]_{\ell_r}^{\ell_c}) &= \{(\ell_c; \ell_n), (\ell_x; \ell_r)\} \\ &\quad \text{if proc } p(\text{val } x, \text{res } y) \text{ is }^{\ell_n} S \text{ end}^{\ell_x} \\ &\quad \text{is in } D_* \end{aligned}$$

Here $(\ell_c; \ell_n)$ and $(\ell_x; \ell_r)$ are new kinds of flows:

- $(\ell_c; \ell_n)$ is the flow corresponding to *calling* a procedure at ℓ_c and with ℓ_n being the entry point for the procedure body, and

- $(l_x; l_r)$ is the flow corresponding to exiting a procedure body at l_x and returning to the call at l_r .

The definition of $flow([call\ p(a, z)]_{l_r}^{l_c})$ exploits the fact that the syntax of procedure calls only allows us to use the (constant) name of a procedure defined in the program; had we been allowed to use a variable that denotes a procedure (e.g. because it was a formal parameter to some procedure or because it was a variable being assigned some procedure) then it would be much harder to define $flow([call\ p(a, z)]_{l_r}^{l_c})$. This is often called the *dynamic dispatch problem* and we shall deal with it in Chapter 3.

Flow graphs for programs. Next consider the program P_* of the form $begin\ D_*\ S_*\ end$. For each procedure declaration $proc\ p(val\ x, res\ y)\ is^{l_n}\ S\ end^{l_z}$ we set

$$\begin{aligned} init(p) &= l_n \\ final(p) &= \{l_x\} \\ blocks(p) &= \{is^{l_n}, end^{l_z}\} \cup blocks(S) \\ labels(p) &= \{l_n, l_x\} \cup labels(S) \\ flow(p) &= \{(l_n, init(S))\} \cup flow(S) \cup \{(\ell, l_x) \mid \ell \in final(S)\} \end{aligned}$$

and for the entire program P_* we set

$$\begin{aligned} init_* &= init(S_*) \\ final_* &= final(S_*) \\ blocks_* &= \bigcup \{blocks(p) \mid proc\ p(val\ x, res\ y)\ is^{l_n}\ S\ end^{l_z}\ is\ in\ D_*\} \\ &\quad \cup blocks(S_*) \\ labels_* &= \bigcup \{labels(p) \mid proc\ p(val\ x, res\ y)\ is^{l_n}\ S\ end^{l_z}\ is\ in\ D_*\} \\ &\quad \cup labels(S_*) \\ flow_* &= \bigcup \{flow(p) \mid proc\ p(val\ x, res\ y)\ is^{l_n}\ S\ end^{l_z}\ is\ in\ D_*\} \\ &\quad \cup flow(S_*) \end{aligned}$$

as well as $Lab_* = labels_*$.

We shall also need to define a notion of *interprocedural flow*

$$\begin{aligned} inter-flow_* &= \{(l_c, l_n, l_x, l_r) \mid P_*\ contains\ [call\ p(a, z)]_{l_r}^{l_c} \\ &\quad as\ well\ as\ proc\ p(val\ x, res\ y)\ is^{l_n}\ S\ end^{l_z}\} \end{aligned}$$

that clearly indicates the relationship between the labels of a procedure call and the corresponding procedure body.

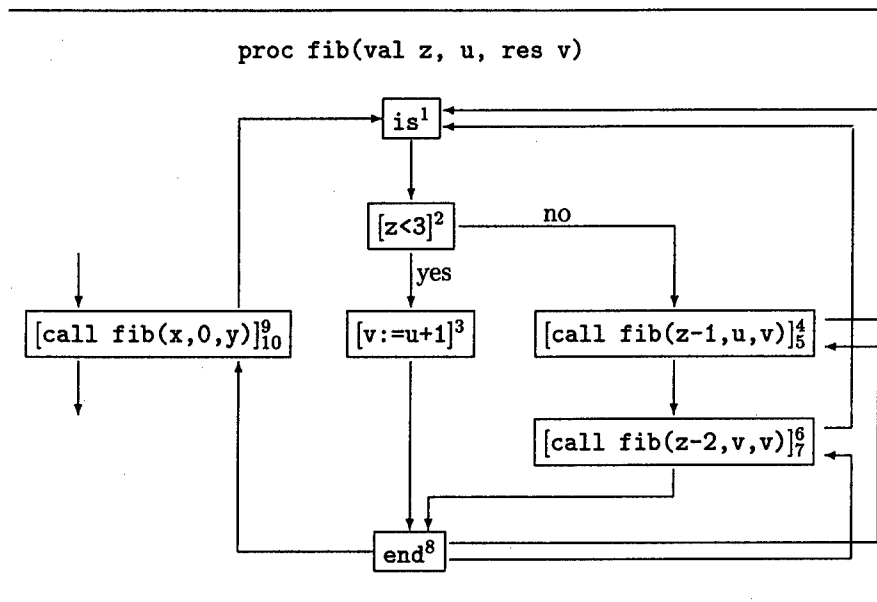


Figure 2.7: Flow graph for the Fibonacci program.

Example 2.34 For the Fibonacci program considered in Example 2.33 we have

$$\begin{aligned} flow_* = & \{(1, 2), (2, 3), (3, 8), \\ & (2, 4), (4, 1), (8, 5), (5, 6), (6, 1), (8, 7), (7, 8), \\ & (9, 1), (8, 10)\} \end{aligned}$$

$$inter-flow_* = \{(9, 1, 8, 10), (4, 1, 8, 5), (6, 1, 8, 7)\}$$

and $init_* = 9$ and $final_* = \{10\}$. The corresponding flow graph is illustrated in Figure 2.7. ■

For a *forward analysis* we use $F = flow_*$, $E = \{init_*\}$ and $IF = inter-flow_*$ whereas for a *backward analysis* we use $F = flow_*^R$, $E = final_*$ and $IF = inter-flow_*$. However, the explanations in the sequel will focus on forward analyses.

2.5.1 Structural Operational Semantics

We shall now show how the semantics of WHILE can be extended to cope with the new constructs. To ensure that the language allows *local data* in procedures we shall need to distinguish between the values assigned to different incarnations of the same variable and for this we introduce an infinite

set of *locations* (or addresses):

$$\xi \in \mathbf{Loc} \quad \text{locations}$$

An *environment*, ρ , will map the variables in the current scope to their locations, and a *store*, ς , will then specify the values of these locations:

$$\begin{aligned} \rho \in \mathbf{Env} &= \mathbf{Var}_* \rightarrow \mathbf{Loc} && \text{environments} \\ \varsigma \in \mathbf{Store} &= \mathbf{Loc} \rightarrow_{\text{fin}} \mathbf{Z} && \text{stores} \end{aligned}$$

Here \mathbf{Var}_* is the (finite) set of variables occurring in the program and $\mathbf{Loc} \rightarrow_{\text{fin}} \mathbf{Z}$ denotes the set of *partial* functions from \mathbf{Loc} to \mathbf{Z} that have a *finite* domain. Thus the previously used states $\sigma \in \mathbf{State} = \mathbf{Var}_* \rightarrow \mathbf{Z}$ have been replaced by the two mappings ρ and ς and can be reconstructed as $\sigma = \varsigma \circ \rho$: to determine the value of a variable x we first determine its location $\xi = \rho(x)$ and next the value $\varsigma(\xi)$ stored in that location. For this to work it is essential that $\varsigma \circ \rho : \mathbf{Var}_* \rightarrow \mathbf{Z}$ is a *total* function rather than a partial function; in other words, we demand that $\text{ran}(\rho) \subseteq \text{dom}(\varsigma)$ where $\text{ran}(\rho) = \{\rho(x) \mid x \in \mathbf{Var}_*\}$ and $\text{dom}(\varsigma) = \{\xi \mid \varsigma \text{ is defined on } \xi\}$.

The locations of the global variables of the program P_* are given by a *top-level environment* denoted ρ_* ; we shall *assume* that it maps all variables to unique locations. The semantics of statements is now given relative to modifications of this environment. The transitions have the general form

$$\rho \vdash_* \langle S, \varsigma \rangle \rightarrow \langle S', \varsigma' \rangle$$

in case that the computation does not terminate in one step, and the form

$$\rho \vdash_* \langle S, \varsigma \rangle \rightarrow \varsigma'$$

in case that it does terminate in one step. It is fairly straightforward to rewrite the semantics of WHILE given in Table 2.6 to have this form; as an example the clause [ass] for assignments becomes:

$$\rho \vdash_* \langle x := a, \varsigma \rangle \rightarrow \varsigma[\rho(x) \mapsto \mathcal{A}[a](\varsigma \circ \rho)] \quad \text{if } \varsigma \circ \rho \text{ is total}$$

Note that there is no need to modify the semantics of arithmetic and boolean expressions.

For procedure calls we make use of the top-level environment, ρ_* , and we take:

$$\begin{aligned} \rho \vdash_* \langle [\text{call } p(a, z)]_{\xi_r}^{\xi_c}, \varsigma \rangle \rightarrow \\ \langle \text{bind } \rho_*[x \mapsto \xi_1, y \mapsto \xi_2] \text{ in } S \text{ then } z := y, \varsigma[\xi_1 \mapsto \mathcal{A}[a](\varsigma \circ \rho), \xi_2 \mapsto v] \rangle \\ \text{where } \xi_1, \xi_2 \notin \text{dom}(\rho), v \in \mathbf{Z} \\ \text{and proc } p(\text{val } x, \text{res } y) \text{ is}^{\xi_r} S \text{ end}^{\xi_c} \text{ is in } D_* \end{aligned}$$

The idea is that we allocate new locations ξ_1 and ξ_2 for the formal parameters x and y , and we then make use of a *bind-construct* to combine the procedure

body S with the environment $\rho_*[x \mapsto \xi_1, y \mapsto \xi_2]$ in which it must be executed and we also record that the final value of y must be returned in the actual parameter z . At the same time the store is updated such that the new location for x is mapped to the value of the actual parameter a whereas we do not control the initial value of the new location for y . The `bind`-construct is only needed to ensure that we have static scope rules and its semantics is as follows:

$$\frac{\rho' \vdash_* \langle S, \varsigma \rangle \rightarrow \langle S', \varsigma' \rangle}{\rho \vdash_* \langle \text{bind } \rho' \text{ in } S \text{ then } z := y, \varsigma \rangle \rightarrow \langle \text{bind } \rho' \text{ in } S' \text{ then } z := y, \varsigma' \rangle}$$

$$\frac{\rho' \vdash_* \langle S, \varsigma \rangle \rightarrow \varsigma'}{\rho \vdash_* \langle \text{bind } \rho' \text{ in } S \text{ then } z := y, \varsigma \rangle \rightarrow \varsigma'[\rho(z) \mapsto \varsigma'(\rho'(y))]}$$

The first rule expresses that executing one step of the body of the construct amounts to executing one step of the construct itself; note that we use the local environment when executing the body. The second rule expresses that when the execution of the body finishes then so does execution of the construct itself and we update the value of the global variable z to be that of the local variable y ; furthermore, there is no need for the local environment ρ' to be retained as subsequent computations will use the previous environment ρ .

Remark. Although the semantics works with two mappings, an environment and a store, it is often the case that the analysis abstracts the state, i.e. the composition of the environment and the store. The correctness of the analysis will then have to relate the abstract state both to the environment and the store.

The correctness result will often be expressed in the style of Section 2.2: information obtained by analysing the original program will remain correct under execution of the program. The semantics presented above deviates from that of the `WHILE`-language in that it introduces the `bind`-construct which is only used in the intermediate configurations. So in order to prove the correctness result we will also need to specify how to analyse the `bind`-construct. We refer to Chapter 3 for an illustration of how to do this. ■

2.5.2 Intraprocedural versus Interprocedural Analysis

To appreciate why interprocedural analysis is harder than intraprocedural analysis let us begin by just naively using the techniques from the previous sections. For this we suppose that:

- for each procedure call $[\text{call } p(a, z)]_{\xi_r}^{\xi_c}$ we have two transfer functions f_{ξ_c} and f_{ξ_r} corresponding to calling the procedure and returning from the call, and
- for each procedure definition `proc p(val x, res y) isl_n S endl_z` we

have two transfer functions f_{e_n} and f_{e_x} corresponding to entering and exiting the procedure body.

A naive formulation. Given an instance $(L, \mathcal{F}, F, E, \iota, f.)$ of a Monotone Framework we shall now treat the two kinds of flow $((\ell_1, \ell_2)$ versus $(\ell_c; \ell_n)$ and $(\ell_x; \ell_r)$) in the same way: we interpret the semi-colon as standing for a comma. While a Monotone Framework is allowed to interpret all transfer functions freely, we shall for now, naively, assume that the two transfer functions associated with procedure definitions are the identity functions, and that the two transfer functions associated with each procedure call are also the identity functions, thus effectively ignoring the parameter-passing.

We now obtain an equation system of the form considered in the previous sections:

$$\begin{aligned} A_\bullet(\ell) &= f_\ell(A_\bullet(\ell)) \\ A_\circ(\ell) &= \bigsqcup \{A_\bullet(\ell') \mid (\ell', \ell) \in F \text{ or } (\ell'; \ell) \in F\} \sqcup \iota_E^\ell \end{aligned}$$

Here ι_E^ℓ is as in Section 2.3:

$$\iota_E^\ell = \begin{cases} \iota & \text{if } \ell \in E \\ \perp & \text{if } \ell \notin E \end{cases}$$

When inspecting this equation system it should be apparent that both procedure calls $(\ell_c; \ell_n)$ and procedure returns $(\ell_x; \ell_r)$ are treated like goto's: there is no mechanism for ensuring that information flowing along $(\ell_c; \ell_n)$ from a call to a procedure only flows back along $(\ell_x; \ell_r)$ from the procedure to the *same* call. (Indeed, nowhere does the formulation consult the interprocedural flow, *IF*.) Expressed in terms of the flow graph in Figure 2.7, there is nothing preventing us from considering a path like [9, 1, 2, 4, 1, 2, 3, 8, 10] that does not correspond to a run of the program. Intuitively, the equation system considers a much too large set of "paths" through the program and hence will be grossly imprecise (although formally on the safe side).

Valid paths. A natural way to overcome this shortcoming is to somehow restrict the attention to paths that have the proper nesting of procedure calls and exits. We shall explore this idea in the context of redefining the MOP solution of Section 2.4 to only take the proper set of paths into account, thereby defining an *MVP solution* (for *Meet over all Valid Paths*).

So consider a program P_\star of the form $\text{begin } D_\star S_\star \text{ end}$. A path is said to be a *complete path* from ℓ_1 to ℓ_2 in P_\star if it has proper nesting of procedure entries and exits and such that a procedure returns to the point where it was called. These paths are generated by the nonterminal CP_{ℓ_1, ℓ_2} according to the following productions:

$CP_{l_1, l_2} \rightarrow l_1$	whenever $l_1 = l_2$
$CP_{l_1, l_3} \rightarrow l_1, CP_{l_2, l_3}$	whenever $(l_1, l_2) \in F$; for a forward analysis this means that $(l_1, l_2) \in flow_*$
$CP_{l_c, l} \rightarrow l_c, CP_{l_n, l_x}, CP_{l_r, l}$	whenever $(l_c, l_n, l_x, l_r) \in IF$; for a forward analysis this means that P_* contains $[call\ p(a, z)]_{l_r}^{l_c}$ and $proc\ p(val\ x, res\ y)\ is^{l_n}\ S\ end^{l_x}$

The matching of calls and returns is ensured by the last kind of productions: the flows (l_c, l_n) and (l_x, l_r) are forced to obey a parenthesis structure in that l_c, l_n only will be in the generated path if there is a matching occurrence of l_x, l_r – and vice versa. Hence for a forward analysis, a terminating computation will give rise to a complete path from $init_*$ to one of the labels of $final_*$. Note that the grammar constructed above will only have a finite set of nonterminals because there only is a finite set of labels in P_* .

Example 2.35 For the Fibonacci program of Example 2.33 we obtain the following grammar (using forward flow and ignoring the parts not reachable from $CP_{9,10}$):

$CP_{9,10} \rightarrow 9, CP_{1,8}, CP_{10,10}$	$CP_{3,8} \rightarrow 3, CP_{8,8}$
$CP_{10,10} \rightarrow 10$	$CP_{8,8} \rightarrow 8$
$CP_{1,8} \rightarrow 1, CP_{2,8}$	$CP_{4,8} \rightarrow 4, CP_{1,8}, CP_{5,8}$
$CP_{2,8} \rightarrow 2, CP_{3,8}$	$CP_{5,8} \rightarrow 5, CP_{6,8}$
$CP_{2,8} \rightarrow 2, CP_{4,8}$	$CP_{6,8} \rightarrow 6, CP_{1,8}, CP_{7,8}$
	$CP_{7,8} \rightarrow 7, CP_{8,8}$

It is now easy to verify that the path $[9, 1, 2, 4, 1, 2, 3, 8, 5, 6, 1, 2, 3, 8, 7, 8, 10]$ is generated by $CP_{9,10}$ whereas the path $[9, 1, 2, 4, 1, 2, 3, 8, 10]$ is not. ■

A path is said to be a *valid path* if it starts at an extremal node of P_* and if all the procedure exits match the procedure entries but it is possible that some procedures are entered but not yet exited. This will obviously include all prefixes of the complete paths starting in E but we also have to take into account prefixes of computations that might not terminate. To specify the valid paths we therefore construct another grammar with productions:

$VP_* \rightarrow VP_{l_1, l_2}$	whenever $l_1 \in E$ and $l_2 \in Lab_*$
$VP_{l_1, l_2} \rightarrow l_1$	whenever $l_1 = l_2$
$VP_{l_1, l_3} \rightarrow l_1, VP_{l_2, l_3}$	whenever $(l_1, l_2) \in F$
$VP_{l_c, l} \rightarrow l_c, CP_{l_n, l_x}, VP_{l_r, l}$	whenever $(l_c, l_n, l_x, l_r) \in IF$
$VP_{l_c, l} \rightarrow l_c, VP_{l_n, l}$	whenever $(l_c, l_n, l_x, l_r) \in IF$

The valid paths will then be generated by the nonterminal VP_* . For a forward analysis, to come from the label ℓ_c of a procedure call to the program point ℓ there are two possibilities. One is that the call initiated at ℓ_c terminates before reaching ℓ and this corresponds to the second last kind of production where we use the nonterminal CP_{ℓ_n, ℓ_c} to generate the complete path corresponding to executing the procedure body. The other possibility is that ℓ is reached before the call terminates and this corresponds to the last kind of production where we simply use $VP_{\ell_n, \ell}$ to generate a valid path in the procedure body.

We can now modify the two sets of paths defined in Section 2.3 as follows (keeping in mind that the definitions are implicitly parameterised on F , E and IF):

$$\begin{aligned} vpath_o(\ell) &= \{[\ell_1, \dots, \ell_{n-1}] \mid n \geq 1 \wedge \ell_n = \ell \wedge [\ell_1, \dots, \ell_n] \text{ is a valid path}\} \\ vpath_\bullet(\ell) &= \{[\ell_1, \dots, \ell_n] \mid n \geq 1 \wedge \ell_n = \ell \wedge [\ell_1, \dots, \ell_n] \text{ is a valid path}\} \end{aligned}$$

Clearly the sets of paths are smaller than what would have resulted if we had merely regarded $(\ell_1; \ell_2)$ as standing for (ℓ_1, ℓ_2) and had used the notions $path_o(\ell)$ and $path_\bullet(\ell)$ of Section 2.3.

Using valid paths we now define the *MVP solution* as follows:

$$\begin{aligned} MVP_o(\ell) &= \bigsqcup \{f_{\vec{\ell}}(\ell) \mid \vec{\ell} \in vpath_o(\ell)\} \\ MVP_\bullet(\ell) &= \bigsqcup \{f_{\vec{\ell}}(\ell) \mid \vec{\ell} \in vpath_\bullet(\ell)\} \end{aligned}$$

Since the sets of paths are smaller than in the similar definitions in Section 2.3, we clearly have $MVP_o(\ell) \sqsubseteq MOP_o(\ell)$ and $MVP_\bullet(\ell) \sqsubseteq MOP_\bullet(\ell)$ for all ℓ .

2.5.3 Making Context Explicit

The MVP solution may be undecidable for lattices of finite height, just as was the case for the MOP solution, so we now have to reconsider the MFP solution and how to avoid taking too many invalid paths. An obvious approach is to encode information about the paths taken into the data flow properties themselves; to this end we introduce *context information*:

$$\delta \in \Delta \quad \text{context information}$$

The context may simply be an encoding of the path taken but we shall see in Subsection 2.5.5 that there are other possibilities. We shall now show how an instance of a Monotone Framework (as introduced in Section 2.3) can be extended to take context into account.

The intraprocedural fragment. Consider an instance $(L, \mathcal{F}, F, E, \iota, f.)$ of a Monotone Framework. We shall now construct an instance

$$(\widehat{L}, \widehat{\mathcal{F}}, F, E, \widehat{\iota}, \widehat{f}.)$$

of an *embellished monotone framework* that takes context into account. We begin with the parts of its definition that are independent of the actual choice of Δ , i.e. the parts that correspond to the intraprocedural analysis:

- $\widehat{L} = \Delta \rightarrow L$;
- the transfer functions in $\widehat{\mathcal{F}}$ are monotone; and
- each transfer function \widehat{f}_ℓ is given by $\widehat{f}_\ell(\widehat{\iota})(\delta) = f_\ell(\widehat{\iota}(\delta))$.

In other words, the new instance applies the transfer functions of the original instance in a pointwise fashion.

Ignoring procedures, the data flow equations will take the form displayed earlier:

$$A_\bullet(\ell) = \widehat{f}_\ell(A_\bullet(\ell))$$

for all labels that do not label a procedure call
(i.e. that do not occur as first or fourth components
of a tuple in IF)

$$A_\bullet(\ell) = \bigsqcup \{A_\bullet(\ell') \mid (\ell', \ell) \in F \text{ or } (\ell'; \ell) \in F\} \sqcup \widehat{\iota}_E^{\ell}$$

for all labels (including those that label procedure calls)

Example 2.36 Let $(L_{\text{sign}}, \mathcal{F}_{\text{sign}}, F, E, \iota_{\text{sign}}, f^{\text{sign}})$ be an instance of a Monotone Framework specifying a *Detection of Signs Analysis* (see Exercise 2.15) and assume that

$$L_{\text{sign}} = \mathcal{P}(\text{Var}_* \rightarrow \text{Sign})$$

where $\text{Sign} = \{-, 0, +\}$. Thus L_{sign} describes sets of abstract states σ^{sign} mapping variables to their possible signs. The transfer function f_ℓ^{sign} associated with the assignment $[x := a]^\ell$ will now be written as

$$f_\ell^{\text{sign}}(Y) = \bigcup \{\phi_\ell^{\text{sign}}(\sigma^{\text{sign}}) \mid \sigma^{\text{sign}} \in Y\}$$

where $Y \subseteq \text{Var}_* \rightarrow \text{Sign}$ and

$$\phi_\ell^{\text{sign}}(\sigma^{\text{sign}}) = \{\sigma^{\text{sign}}[x \mapsto s] \mid s \in \mathcal{A}_{\text{sign}}[a](\sigma^{\text{sign}})\}$$

Here $\mathcal{A}_{\text{sign}} : \mathbf{AExp} \rightarrow (\text{Var}_* \rightarrow \text{Sign}) \rightarrow \mathcal{P}(\text{Sign})$ specifies the analysis of arithmetic expressions. The transfer functions for tests and skip-statements are the identity functions.

Given a set Δ of contexts, the embellished monotone framework will have

$$\widehat{L}_{\text{sign}} = \Delta \rightarrow L_{\text{sign}}$$

but we shall prefer the following isomorphic definition

$$\widehat{L}_{\text{sign}} = \mathcal{P}(\Delta \times (\text{Var}_* \rightarrow \text{Sign}))$$

Thus $\widehat{L}_{\text{sign}}$ describes sets of pairs of context and abstract states. The transfer function associated with the assignment $[x := a]^t$ will now be:

$$\widehat{f}_t^{\text{sign}}(Z) = \bigcup \{ \{\delta\} \times \phi_t^{\text{sign}}(\sigma^{\text{sign}}) \mid (\delta, \sigma^{\text{sign}}) \in Z \}$$

In subsequent examples we shall further develop this analysis. ■

The interprocedural fragment. It remains to formulate the data flow equations corresponding to procedures.

For a procedure definition $\text{proc } p(\text{val } x, \text{res } y) \text{ is}^{t_n} S \text{ end}^{t_x}$ we have two transfer functions:

$$\widehat{f}_{t_n}, \widehat{f}_{t_x} : (\Delta \rightarrow L) \rightarrow (\Delta \rightarrow L)$$

In the case of our simple language we shall prefer to take both of these transfer functions to be the identity function; i.e.

$$\widehat{f}_{t_n}(\widehat{l}) = \widehat{l}$$

$$\widehat{f}_{t_x}(\widehat{l}) = \widehat{l}$$

for all $\widehat{l} \in \widehat{L}$. Hence the effect of procedure entry is handled by the transfer function for procedure call (considered below) and similarly the effect of procedure exit is handled by the transfer function for procedure return (also considered below). For more advanced languages where many semantic actions take place at procedure entry or exit it may be preferable to reconsider this decision.

For a procedure call $(l_c, l_n, l_x, l_r) \in IF$ we shall define two transfer functions. In our explanation we shall concentrate on the case of forward analyses where P_* contains $[\text{call } p(a, z)]_{l_r}^{l_c}$ as well as $\text{proc } p(\text{val } x, \text{res } y) \text{ is}^{t_n} S \text{ end}^{t_x}$. Corresponding to the actual call we have the transfer function

$$\widehat{f}_{l_c}^1 : (\Delta \rightarrow L) \rightarrow (\Delta \rightarrow L)$$

and it is used in the equation:

$$A_\bullet(l_c) = \widehat{f}_{l_c}^1(A_\circ(l_c)) \text{ for all } (l_c, l_n, l_x, l_r) \in IF$$




Figure 2.8: Analysis of procedure call: the forward case.

In other words, the transfer function modifies the data flow properties (and the context) as required for passing to the procedure entry.

Corresponding to the return we have the transfer function

$$\widehat{f_{l_c, l_r}^2} : (\Delta \rightarrow L) \times (\Delta \rightarrow L) \rightarrow (\Delta \rightarrow L)$$

and it is used in the equation:

$$A_\bullet(l_r) = \widehat{f_{l_c, l_r}^2}(A_\circ(l_c), A_\bullet(l_r)) \text{ for all } (l_c, l_n, l_x, l_r) \in IF$$

The first parameter of $\widehat{f_{l_c, l_r}^2}$ describes the data flow properties at the call point for the procedure and the second parameter describes the properties at the exit from the procedure body. Ignoring the first parameter, the transfer function modifies the data flow properties (and the context) as required for passing back from the procedure exit. The purpose of the first parameter is to recover some of the information (data flow properties as well as context information) that was available before the actual call; how this is done depends on the actual choice of the set, Δ , of context information and we shall return to this shortly. Figure 2.8 illustrates the flow of data in the analysis of the procedure call.

Variations. The functionality and use of $\widehat{f_{l_c, l_r}^2}$ (as well as Figure 2.8) is sufficiently general that it allows us to deal with most of the scenarios found in the literature. A simple example being the possibility to define

$$\widehat{f_{l_c, l_r}^2}(\hat{l}, \hat{l}') = \widehat{f_{l_r}^2}(\hat{l}')$$

thereby completely ignoring the information before the call; this is illustrated in Figure 2.9.




Figure 2.9: Analysis of procedure call: ignoring calling context.




Figure 2.10: Analysis of procedure call: merging of context.

A somewhat more interesting example is the ability to define

$$f_{l_c, l_r}^2(\hat{l}, \hat{v}) = f_{l_c, l_r}^{2A}(\hat{l}) \sqcup f_{l_c, l_r}^{2B}(\hat{v})$$

thereby allowing a simple combination of the information coming back from the call with the information pertaining before the call. This form is illustrated in Figure 2.10 and is often motivated on the grounds that f_{l_c, l_r}^{2A} copies data that is local to the calling procedure whereas f_{l_c, l_r}^{2B} copies information that is global. (It may be worth noticing that the function f_{l_c, l_r}^2 is completely additive if and only if it can be written in this form with f_{l_c, l_r}^{2A} and f_{l_c, l_r}^{2B} being completely additive.)

Context-sensitive versus context-insensitive. So far we have criticised the naive approach because it was unable to maintain the proper relationship between procedure calls and procedure returns. A related criticism of the naive approach is that it cannot distinguish between the different calls of a procedure. The information about calling states is combined for all call sites, the procedure body is analysed only once using this combined information, and the resulting information about the set of return states is used at all return points. The phrase *context-insensitive* is often used to refer to this shortcoming.

The use of non-trivial context information not only helps to avoid the first criticism but also the second: if there are two different calls but they are reached with different contexts, δ_1 and δ_2 , then all information obtained from the procedure will be clearly related to δ_1 or δ_2 and no undesired combination or “cross-over” will take place. The phrase *context-sensitive* is often used to refer to this ability.

Clearly a context-sensitive analysis is more precise than a context-insensitive analysis but at the same time it is also likely to be more costly. The choice between which technique to use amounts to a careful balance between precision and efficiency.

2.5.4 Call Strings as Context

To complete the design of the analysis of the program we must choose the set, Δ , of context information and also specify the extremal value, $\hat{\iota}$, and define the two transfer functions associated with procedure calls. In this subsection we shall consider two approaches based on call strings and our explanation will be in terms of forward analyses.

Call strings of unbounded length. As the first possibility we simply encode the path taken; however, since our main interest is with procedure calls we shall only record flows of the form $(\ell_c; \ell_n)$ corresponding to a procedure call. Formally we take

$$\Delta = \text{Lab}^*$$

where the most recent label ℓ_c of a procedure call is at the right end (just as was the case for valid paths and paths); elements of Δ are called *call strings*. We then define

$$\hat{\iota} = (\Lambda, \iota)$$

where Λ is the empty sequence corresponding to the fact that there are no pending procedure calls when the program starts execution; ι is the extremal value available from the underlying Monotone Framework.

Example 2.37 For the Fibonacci program of Example 2.33 the following call strings will be of interest:

$$\Lambda, [9], [9, 4], [9, 6], [9, 4, 4], [9, 4, 6], [9, 6, 4], [9, 6, 6], \dots$$

corresponding to the cases with 0, 1, 2, 3, ... pending procedure calls. ■

For a procedure call $(\ell_c, \ell_n, \ell_x, \ell_r) \in IF$, amounting to $[\text{call } p(a, z)]_{\ell_r}^{\ell_c}$ in the case of a forward analysis, we define the transfer function $\widehat{f}_{\ell_c}^1$ such that $\widehat{f}_{\ell_c}^1(\widehat{l})([\delta, \ell_c]) = f_{\ell_c}^1(\widehat{l}(\delta))$ where $[\delta, \ell_c]$ denotes the path obtained by appending ℓ_c to δ (so as to reflect that now we enter the body of the procedure) and the function $f_{\ell_c}^1 : L \rightarrow L$ describes how the property is modified. This is achieved by setting

$$\widehat{f}_{\ell_c}^1(\widehat{l})(\delta') = \begin{cases} f_{\ell_c}^1(\widehat{l}(\delta)) & \text{when } \delta' = [\delta, \ell_c] \\ \perp & \text{otherwise} \end{cases}$$

which takes care of the special case of empty paths.

Next we define the transfer function $\widehat{f}_{\ell_c, \ell_r}^2$ corresponding to returning from the procedure call:

$$\widehat{f}_{\ell_c, \ell_r}^2(\widehat{l}, \widehat{l}')(\delta) = f_{\ell_c, \ell_r}^2(\widehat{l}(\delta), \widehat{l}'([\delta, \ell_c]))$$

Here the information \widehat{l} from the original call is combined with information \widehat{l}' from the procedure exit using the function $f_{\ell_c, \ell_r}^2 : L \times L \rightarrow L$. However, only information corresponding to the *same* contexts for call point ℓ_c is combined: this is ensured by the two occurrences of δ in the above formula.

Example 2.38 Let us return to the Detection of Signs Analysis of Example 2.36. For a procedure call $[\text{call } p(a, z)]_{\ell_r}^{\ell_c}$ where p is declared by $\text{proc } p(\text{val } x, \text{res } y) \text{ is } \ell_n \text{ } S \text{ end } \ell_x$ we may take:

$$\widehat{f}_{\ell_c}^{\text{sign}1}(Z) = \bigcup \{ \{\delta'\} \times \phi_{\ell_c}^{\text{sign}1}(\sigma^{\text{sign}}) \mid (\delta', \sigma^{\text{sign}}) \in Z \wedge \delta' = [\delta, \ell_c] \}$$

$$\phi_{\ell_c}^{\text{sign}1}(\sigma^{\text{sign}}) = \{ \sigma^{\text{sign}}[x \mapsto s][y \mapsto s'] \mid s \in \mathcal{A}_{\text{sign}}[a](\sigma^{\text{sign}}) \wedge s' \in \{-, 0, +\} \}$$

When returning from the procedure call we take:

$$\widehat{f}_{\ell_c, \ell_r}^{\text{sign}2}(Z, Z') = \bigcup \{ \{\delta\} \times \phi_{\ell_c, \ell_r}^{\text{sign}2}(\sigma_1^{\text{sign}}, \sigma_2^{\text{sign}}) \mid (\delta, \sigma_1^{\text{sign}}) \in Z \wedge (\delta', \sigma_2^{\text{sign}}) \in Z' \wedge \delta' = [\delta, \ell_c] \}$$

$$\phi_{\ell_c, \ell_r}^{\text{sign}2}(\sigma_1^{\text{sign}}, \sigma_2^{\text{sign}}) = \{ \sigma_2^{\text{sign}}[x \mapsto \sigma_1^{\text{sign}}(x); y \mapsto \sigma_1^{\text{sign}}(y); z \mapsto \sigma_2^{\text{sign}}(y)] \}$$

Thus we extract all the information from the procedure body except for the information about the formal parameters x and y and the actual parameter

z. For the formal parameters we rely on the information available before the current call which is still correct and for the actual parameter we perform the required update of the information. Note that to facilitate this definition it is crucial that the transfer function $\widehat{f_{\ell_c, \ell_r}^2}$ takes two arguments: information from the call point as well as from the procedure exit. ■

Call strings of bounded length. Clearly the call strings can become arbitrarily long because the procedures may be recursive. It is therefore customary to restrict their length to be at most k for some number $k \geq 0$; the idea being that only the last k calls are recorded. We write this as

$$\Delta = \text{Lab}^{\leq k}$$

and we still take the extremal value to be $\widehat{\iota} = (\Lambda, \iota)$. Note that in the case $k = 0$ we have $\Delta = \{\Lambda\}$ which is equivalent to having no context information.

Example 2.39 Consider the Fibonacci program of Example 2.33 and assume that we are only interested in recording the last call, i.e. $k = 1$. Then the call strings of interest are:

$$\Lambda, [9], [4], [6]$$

Alternatively, we may choose to record the last two calls, i.e. $k = 2$, in which case the following call strings are of interest:

$$\Lambda, [9], [9, 4], [9, 6], [4, 4], [4, 6], [6, 4], [6, 6]$$

In general, we would expect an analysis using these 8 contexts to be more precise than one using the 4 different contexts displayed above. ■

We shall now present the transfer functions for the general case where call strings have length at most k . The transfer function $\widehat{f_{\ell_c}^1}$ for procedure call is redefined by

$$\widehat{f_{\ell_c}^1}(\widehat{\iota})(\delta') = \bigsqcup \{f_{\ell_c}^1(\widehat{\iota}(\delta)) \mid \delta' = [\delta, \ell_c]_k\}$$

where $[\delta, \ell_c]_k$ denotes the call string $[\delta, \ell_c]$ but possibly truncated (by omitting elements on the left) so as to have length at most k . Since the function mapping δ to $[\delta, \ell_c]_k$ is not injective (unlike the one mapping δ to $[\delta, \ell_c]$) we need to take the least upper bound over all δ that can be mapped to the relevant context δ' .

Similarly, the transfer function $\widehat{f_{\ell_c, \ell_r}^2}$ for procedure return is redefined by

$$\widehat{f_{\ell_c, \ell_r}^2}(\widehat{\iota}, \widehat{\iota}') = f_{\ell_c, \ell_r}^2(\widehat{\iota}(\delta), \widehat{\iota}'([\delta, \ell_c]_k))$$

as should be expected.

Example 2.40 Let us consider Detection of Signs Analysis in the special case where $k = 0$, i.e. where $\Delta = \{\Lambda\}$ and hence $\Delta \times (\mathbf{Var}_* \rightarrow \mathbf{Sign})$ is isomorphic to $\mathbf{Var}_* \rightarrow \mathbf{Sign}$. Using this isomorphism the formulae defining the transfer functions for procedure call can be simplified to

$$\begin{aligned}\widehat{f_{\ell_c}^{\text{sign}1}}(Y) &= \bigcup \{\phi_{\ell_c}^{\text{sign}1}(\sigma^{\text{sign}}) \mid \sigma^{\text{sign}} \in Y\} \\ \widehat{f_{\ell_c, \ell_r}^{\text{sign}2}}(Y, Y') &= \bigcup \{\phi_{\ell_c, \ell_r}^{\text{sign}2}(\sigma_1^{\text{sign}}, \sigma_2^{\text{sign}}) \mid \sigma_1^{\text{sign}} \in Y \wedge \sigma_2^{\text{sign}} \in Y'\}\end{aligned}$$

where $Y, Y' \subseteq \mathbf{Var}_* \rightarrow \mathbf{Sign}$. It is now easy to see that the analysis is context-insensitive: at procedure return it is not possible to distinguish between the different call points.

Let us next consider the case where $k = 1$. Here $\Delta = \mathbf{Lab} \cup \{\Lambda\}$ and the transfer functions for procedure call are:

$$\begin{aligned}\widehat{f_{\ell_c}^{\text{sign}1}}(Z) &= \bigcup \{\{\ell_c\} \times \phi_{\ell_c}^{\text{sign}1}(\sigma^{\text{sign}}) \mid (\delta, \sigma^{\text{sign}}) \in Z\} \\ \widehat{f_{\ell_c, \ell_r}^{\text{sign}2}}(Z, Z') &= \bigcup \{\{\delta\} \times \phi_{\ell_c, \ell_r}^{\text{sign}2}(\sigma_1^{\text{sign}}, \sigma_2^{\text{sign}}) \mid (\delta, \sigma_1^{\text{sign}}) \in Z \\ &\quad \wedge (\ell_c, \sigma_2^{\text{sign}}) \in Z'\}\end{aligned}$$

Now the transfer function $\widehat{f_{\ell_c}^{\text{sign}1}}$ will mark all data from the call point ℓ_c with that label. Thus it does not harm that the information $\widehat{f_{\ell_c}^{\text{sign}1}}(Z)$ is merged with similar information $\widehat{f_{\ell_c}^{\text{sign}1}}(Z)$ from another procedure call. At the return from the call the transfer function $\widehat{f_{\ell_c, \ell_r}^{\text{sign}2}}$ selects those pairs $(\ell_c, \sigma_2^{\text{sign}}) \in Z'$ that are relevant for the current call and combines them with those pairs $(\delta, \sigma_1^{\text{sign}}) \in Z$ that describe the situation before the call; in particular, this allows us to reset the context to be that of the call point. ■

2.5.5 Assumption Sets as Context

An alternative to describing a path directly in terms of the calls being performed is to record information about the state in which the call was made; these methods can clearly be combined but in the interest of simplicity we shall abstain from doing so.

Large assumption sets. Throughout this subsection we shall make the simplifying assumption that

$$L = \mathcal{P}(D)$$

as is the case for the Detection of Signs Analysis. Restricting the attention to only recording information about the last call (corresponding to taking

$k = 1$ above), one possibility is to take

$$\Delta = \mathcal{P}(D)$$

and we then take the extremal value to be

$$\hat{\iota} = (\{\iota\}, \iota)$$

meaning that the initial context is described by the initial abstract state. This kind of context information is often called an *assumption set* and expresses a dependency on data (as opposed to a dependency on control as in the case of call strings).

Example 2.41 Assume that we want to perform a Detection of Signs Analysis (Example 2.36) of the Fibonacci program of Example 2.33 and that the extremal value ι_{sign} is the singleton $[x \mapsto +, y \mapsto -, z \mapsto -]$. Then the contexts of primary interest will be sets consisting of some of the following abstract states

$$\begin{array}{lll} [x \mapsto +, y \mapsto 0, z \mapsto -], & [x \mapsto +, y \mapsto 0, z \mapsto 0], & [x \mapsto +, y \mapsto 0, z \mapsto +], \\ [x \mapsto +, y \mapsto +, z \mapsto -], & [x \mapsto +, y \mapsto +, z \mapsto 0], & [x \mapsto +, y \mapsto +, z \mapsto +] \end{array}$$

corresponding to the states in which the call-statements may be encountered. ■

For a procedure call $(\ell_c, \ell_n, \ell_x, \ell_r) \in IF$, i.e. $[\text{call } p(a, z)]_{\ell_c}^{\ell_r}$ in the case of forward analysis, we define the transfer function $\widehat{f}_{\ell_c}^1$ for procedure call by:

$$\widehat{f}_{\ell_c}^1(Z) = \bigcup \{ \{\delta'\} \times \phi_{\ell_c}^1(d) \mid (\delta, d) \in Z \wedge \delta' = \{d'' \mid (\delta, d'') \in Z\} \}$$

where $\phi_{\ell_c}^1 : D \rightarrow \mathcal{P}(D)$. The idea is as follows: a pair $(\delta, d) \in Z$ describes a context and an abstract state for the current call. We now have to modify the context to take the call into account, i.e. we have to determine the set of possible abstract states in which the call could happen in the current context and this is $\delta' = \{d'' \mid (\delta, d'') \in Z\}$. Given this context we proceed as in the call string formulations presented above and mark the data flow property with this context.

Next we shall consider the transfer function $\widehat{f}_{\ell_c, \ell_r}^2$ for procedure return

$$\widehat{f}_{\ell_c, \ell_r}^2(Z, Z') = \bigcup \{ \{\delta\} \times \phi_{\ell_c, \ell_r}^2(d, d') \mid (\delta, d) \in Z \wedge (\delta', d') \in Z' \wedge \delta' = \{d'' \mid (\delta, d'') \in Z\} \}$$

where $\phi_{\ell_c, \ell_r}^2 : D \times D \rightarrow \mathcal{P}(D)$. Here $(\delta, d) \in Z$ describes the situation before the call and $(\delta', d') \in Z'$ describes the situation at the procedure exit.

From the definition of $\widehat{f_{l_c}^1}$ we know that the context matching (δ, d) will be $\delta' = \{d'' \mid (\delta, d'') \in Z\}$ so we impose that condition. We can now combine information from before the call with that at the procedure exit much as in the call string approach; in particular, we can reset the context to be that of the call point.

Small assumption sets. As a simpler version of using assumption sets one may take

$$\Delta = D$$

and then use $\widehat{\iota} = (\iota, \iota)$ as the extremal value. So rather than basing the embellished monotone framework on $\mathcal{P}(D) \times D$ as above we now base it on $D \times D$. Of course, this is much less precise but, on the positive side, the size of the data flow properties has been reduced dramatically.

For a procedure call $(l_c, l_n, l_x, l_r) \in IF$, i.e. $[\text{call } p(a, z)]_{l_c}^{l_r}$ for forward analyses, the transfer function $\widehat{f_{l_c}^1}$ is now defined by

$$\widehat{f_{l_c}^1}(Z) = \bigcup \{ \{d\} \times \phi_{l_c}^1(d) \mid (\delta, d) \in Z \}$$

where, as before, $\phi_{l_c}^1 : D \rightarrow \mathcal{P}(D)$. Here the individual pieces of information concerning the abstract state of the call have their own local contexts; we have no way of grouping the abstract states corresponding to δ as we did in the approach with large assumption sets.

The corresponding definition of the transfer function $\widehat{f_{l_c, l_r}^2}$ for procedure return then is

$$\widehat{f_{l_c, l_r}^2}(Z, Z') = \bigcup \{ \{\delta\} \times \phi_{l_c, l_r}^2(d, d') \mid (\delta, d) \in Z \wedge (d, d') \in Z' \}$$

where again $\phi_{l_c, l_r}^2 : D \times D \rightarrow \mathcal{P}(D)$. Examples of how to use assumption sets will be considered in the exercises.

2.5.6 Flow-Sensitivity versus Flow-Insensitivity

All of the data flow analyses we have considered so far have been *flow-sensitive*: this just means that in general we would expect the analysis of a program $S_1; S_2$ to differ from the analysis of the program $S_2; S_1$ where the statements come in a different order.

Sometimes one considers *flow-insensitive* analyses where the order of statements is of no importance for the analysis being performed. This may sound weird at first, but suppose that the analysis being performed is like the ones considered in Section 2.1 except that for simplicity all kill components are empty sets. Given these assumptions one might expect that the programs $S_1; S_2$ and $S_2; S_1$ give rise to the same analysis. Clearly a flow-insensitive

analysis may be much less precise than its flow-sensitive analogue but also it is likely to be much cheaper; since interprocedural data flow analyses tend to be very costly, it is therefore useful to have a repertoire of techniques for reducing the cost.

Sets of assigned variables. We shall now present an example of a flow-insensitive analysis. Consider a program P_* of the form `begin D_* S_* end`. For each procedure

$$\text{proc } p(\text{val } x, \text{res } y) \text{ is}^{\ell_n} S \text{ end}^{\ell_x}$$

in D_* , the aim is to determine the set $IAV(p)$ of global variables that might be assigned directly or indirectly when p is called.

To compute these sets we need two auxiliary notions. The set $AV(S)$ of directly *assigned variables* gives for each statement S the set of variables that could be assigned in S – but ignoring the effect of procedure calls. It is defined inductively upon the structure of S :

$$\begin{aligned} AV([\text{skip}]^{\ell}) &= \emptyset \\ AV([x := a]^{\ell}) &= \{x\} \\ AV(S_1; S_2) &= AV(S_1) \cup AV(S_2) \\ AV(\text{if } [b]^{\ell} \text{ then } S_1 \text{ else } S_2) &= AV(S_1) \cup AV(S_2) \\ AV(\text{while } [b]^{\ell} \text{ do } S) &= AV(S) \\ AV([\text{call } p(a, z)]_{\ell_r}^{\ell_c}) &= \{z\} \end{aligned}$$

Similarly we shall need the set $CP(S)$ of immediately *called procedures* that gives for each statement S the set of procedure names that could be directly called in S – but ignoring the effect of procedure calls. It is defined inductively upon the structure of S :

$$\begin{aligned} CP([\text{skip}]^{\ell}) &= \emptyset \\ CP([x := a]^{\ell}) &= \emptyset \\ CP(S_1; S_2) &= CP(S_1) \cup CP(S_2) \\ CP(\text{if } [b]^{\ell} \text{ then } S_1 \text{ else } S_2) &= CP(S_1) \cup CP(S_2) \\ CP(\text{while } [b]^{\ell} \text{ do } S) &= CP(S) \\ CP([\text{call } p(a, z)]_{\ell_r}^{\ell_c}) &= \{p\} \end{aligned}$$

Both the sets $AV(S)$ and $CP(S)$ are well-defined by induction on the structure of S ; also it should be clear that they are context-insensitive in the sense that any rearrangement of the statements inside S would have given the same result. The information in $CP(\dots)$ can be presented graphically: let the graph have a node for each procedure name as well as a node called main_* for the program itself, and let the graph have an edge from p (respectively

main_*) to p' whenever the procedure body S of p has $p' \in CP(S)$ (respectively $p' \in CP(S_*)$). This graph is usually called the *procedure call graph*.

We can now formulate a system of data flow equations that specifies how to obtain the desired sets $IAV(p)$:

$$IAV(p) = (AV(S) \setminus \{x\}) \cup \bigcup \{IAV(p') \mid p' \in CP(S)\}$$

where $\text{proc } p(\text{val } x, \text{res } y) \text{ is } {}^{\prime n} S \text{ end}^{\prime z}$ is in D_*

By analogy with the considerations in Section 2.1 we want the least solution of this system of equations.

Example 2.42 Let us now consider the following version of the Fibonacci program (omitting labels):

```
begin proc fib(val z) is if z<3 then call add(1)
                        else (call fib(z-1); call fib(z-2))
end;
proc add(val u) is (y:=y+u; u:=0)
end;
y:=0; call fib(x)
end
```

We then get the following equations

$$IAV(\text{fib}) = (\emptyset \setminus \{z\}) \cup IAV(\text{fib}) \cup IAV(\text{add})$$

$$IAV(\text{add}) = \{y, u\} \setminus \{u\}$$

The associated procedure call graph is shown in Figure 2.11. The least solution to the equation system is

$$IAV(\text{fib}) = IAV(\text{add}) = \{y\}$$

showing that only the variable y will be assigned by the procedure calls. (Had we instead taken the greatest solution to the equations we would have $IAV(\text{fib}) = IAV(\text{add}) = \text{Var}_*$ for any set Var_* of variables that contains those used in the program and this would be completely unusable.) ■

Note that the formulation of the example analysis did not associate information with entries and exit of blocks but rather with the blocks (or more generally the statements) themselves. This is a rather natural space saving approach for a context-insensitive analysis. It also relates to the discussion of Type and Effect Systems in Section 1.6: the “annotated base types” in Table 1.2 versus the “annotated type constructors” in Table 1.3.




Figure 2.11: Procedure call graph for example program.

2.6 Shape Analysis

We shall now study an extension of the WHILE-language with heap allocated data structures and an interprocedural *Shape Analysis* that gives a *finite characterisation* of the *shapes* of these data structures. So while the aim of the previous sections has been to present the basic techniques of Data Flow Analysis, the aim of this section is to show how the techniques can be used to specify a rather complex analysis.

Shape analysis information is not only useful for classical compiler optimisations but also for software development tools: the Shape Analysis will allow us to statically detect errors like dereferencing a nil-pointer – this is guaranteed to give rise to a dynamic error and a warning can be issued. Perhaps more surprisingly, the analysis allows us to validate certain properties of the shape of the data structures manipulated by the program; we can for example validate that a program for in-situ list reversal does indeed transform a non-cyclic list into a non-cyclic list.

Syntax of the pointer language. We shall study an extension of WHILE that allows us to create cells in the heap; the cells are structured and may contain values as well as pointers to other cells. The data stored in a cell is accessed via selectors so we assume that a finite and non-empty set Sel of *selector names* are given:

$$sel \in Sel \quad \text{selector names}$$

As an example Sel may include the Lisp-like selectors `cdr` and `car` for selecting the first and second components of pairs. The cells of the heap can be addressed by expressions like `x.cdr`: this will first determine the cell pointed to by the variable `x` and then return the value of the `cdr` field. For the sake of simplicity we shall only allow one level of selectors although the development

generalises to several levels. Formally the *pointer expressions* are given by:

$$p ::= x \mid x.sel$$

The syntax of the WHILE-language is now extended to have:

$$\begin{aligned} a &::= p \mid n \mid a_1 \ op_a \ a_2 \mid \text{nil} \\ b &::= \text{true} \mid \text{false} \mid \text{not } b \mid b_1 \ op_b \ b_2 \mid a_1 \ op_r \ a_2 \mid op_p \ p \\ S &::= [p:=a]^\ell \mid [\text{skip}]^\ell \mid S_1; S_2 \mid \\ &\quad \text{if } [b]^\ell \text{ then } S_1 \text{ else } S_2 \mid \text{while } [b]^\ell \text{ do } S \mid \\ &\quad [\text{malloc } p]^\ell \end{aligned}$$

Arithmetic expressions are extended to use pointer expressions rather than just variables, and an arithmetic expression can also be the constant `nil`. The binary operations op_a are as before, that is, they are the standard arithmetic operations and in particular they do *not* allow pointer arithmetic. The boolean expressions are extended such that the relational operators op_r now allow testing for the *equality of pointers* and also we shall allow unary operations op_p on pointers (as for example `is-nil` and `has-sel` for each $sel \in \text{Sel}$). Note that arithmetic as well as boolean expressions can only access cells in the heap, they cannot create new cells nor update existing cells.

The assignment statement takes the general form $p:=a$ where p is a pointer expression. In the case where p is just a variable we have an extension of the ordinary assignment of the WHILE language and in the case where p contains a selector we have a destructive update of the heap. The statements of the extended language also contain a statement `malloc p` for creating a new cell pointed to by p .

Example 2.43 The following program reverses the list pointed to by x and leaves the result in y :

```
[y:=nil]1;
while [not is-nil(x)]2 do
  ([z:=y]3; [y:=x]4; [x:=x.cdr]5; [y.cdr:=z]6);
[z:=nil]7
```

Figure 2.12 illustrates the effect of the program when x points to a five element list and y and z are initially undefined. Row 0 shows the heap just before entering the `while`-loop: x points to the list and y is `nil` (denoted by \diamond); to avoid cluttering the figure we do not draw the `car`-pointers. After having executed the statements of the body of the loop the situation is as in row 1: x now points to the tail of the list, y points to the head of the list and z is `nil`. In general the n 'th row illustrates the situation just before entering the loop the $n + 1$ 'th time so in row 5 we see that x points to `nil` and the execution of the loop terminates and y points to the reversed list. The final statement `z:=nil` simply removes the pointer from z to ξ_4 and sets it to the `nil`-value. ■




Figure 2.12: Reversal of a list of five elements.

2.6.1 Structural Operational Semantics

To model the scenario described above we shall introduce an infinite set Loc of *locations* (or addresses) for the heap cells:

$$\xi \in \text{Loc} \quad \text{locations}$$

The value of a variable will now either be an integer (as before), a location (i.e. a pointer) or the special constant \diamond reflecting that it is the nil value. Thus the *states* are given by

$$\sigma \in \text{State} = \text{Var}_* \rightarrow (\mathbf{Z} + \text{Loc} + \{\diamond\})$$

where as usual Var_* is the (finite) set of variables occurring in the program of interest. As mentioned above the cells of the *heap* have multiple fields and they are accessed using the selectors. Each field can either be an integer, a pointer to another cell or it can be nil. We formalise this by taking

$$\mathcal{H} \in \text{Heap} = (\text{Loc} \times \text{Sel}) \rightarrow_{\text{fin}} (\mathbf{Z} + \text{Loc} + \{\diamond\})$$

Note that the use of *partial* functions reflects that not all selector fields need to be defined; as we shall see later, a newly created cell with location ξ will have all its fields to be uninitialised and hence the corresponding heap \mathcal{H} will have $\mathcal{H}(\xi, \text{sel})$ to be undefined for all $\text{sel} \in \text{Sel}$.

Pointer expressions. Given a state and a heap we need to determine the value of a pointer expression p as an element of $\mathbf{Z} + \text{Loc} + \{\diamond\}$. For this we introduce the function

$$\wp : \text{Var} \rightarrow (\text{State} \times \text{Heap}) \rightarrow_{\text{fin}} (\mathbf{Z} + \{\diamond\} + \text{Loc})$$

defined by

$$\wp[x](\sigma, \mathcal{H}) = \sigma(x)$$

$$\wp[x.\text{sel}](\sigma, \mathcal{H}) = \begin{cases} \mathcal{H}(\sigma(x), \text{sel}) & \text{if } \sigma(x) \in \text{Loc} \text{ and } \mathcal{H} \text{ is defined on } (\sigma(x), \text{sel}) \\ \text{undef} & \text{if } \sigma(x) \notin \text{Loc} \text{ or } \mathcal{H} \text{ is undefined on } (\sigma(x), \text{sel}) \end{cases}$$

The first clause takes care of the situation where p is a simple variable and using the state we determine its value – note that this may be an integer, a location or the special nil-value \diamond . The second clause takes care of the case where the pointer expression has the form $x.\text{sel}$. Here we first have to determine the value of x ; it only makes sense to inspect the *sel*-field in the case x evaluates to a location that has a *sel*-field and hence the clause is split into two sub-cases. In the case where x evaluates to a location we simply inspect the heap \mathcal{H} to determine the value of the *sel*-field – again we may note that this can be an integer, a location or the special value \diamond .

Example 2.44 In Figure 2.12 the *oval nodes* model the cells of the heap \mathcal{H} and they are labelled with their location (or address). The *unlabelled edges* denote the state σ : an edge from a variable x to some node labelled ξ means that $\sigma(x) = \xi$; an edge from x to the symbol \diamond means that $\sigma(x) = \diamond$. The *labelled edges* model the heap \mathcal{H} : an edge labelled *sel* from a node labelled

ξ to a node labelled ξ' means that there is a *sel* pointer between the two cells, that is $\mathcal{H}(\xi, \text{sel}) = \xi'$; an edge labelled *sel* from a node labelled ξ to the symbol \diamond means that the pointer is a nil-pointer, that is $\mathcal{H}(\xi, \text{sel}) = \diamond$.

Consider the pointer expression $x.\text{cdr}$ and assume that σ and \mathcal{H} are as in row 0 of Figure 2.12, that is $\sigma(x) = \xi_1$ and $\mathcal{H}(\xi_1, \text{cdr}) = \xi_2$. Then $\wp[x.\text{cdr}](\sigma, \mathcal{H}) = \xi_2$. ■

Arithmetic and boolean expressions. It is now straightforward to extend the semantics of arithmetic and boolean expressions to handle pointer expressions and the nil-constant. Obviously the functionality of the semantic functions \mathcal{A} and \mathcal{B} have to be changed to take the heap into account:

$$\mathcal{A} : \mathbf{AExp} \rightarrow (\mathbf{State} \times \mathbf{Heap}) \rightarrow_{\text{fin}} (\mathbf{Z} + \mathbf{Loc} + \{\diamond\})$$

$$\mathcal{B} : \mathbf{BExp} \rightarrow (\mathbf{State} \times \mathbf{Heap}) \rightarrow_{\text{fin}} \mathbf{T}$$

The clauses for arithmetic expressions are

$$\mathcal{A}[p](\sigma, \mathcal{H}) = \wp[p](\sigma, \mathcal{H})$$

$$\mathcal{A}[n](\sigma, \mathcal{H}) = \mathcal{N}[n]$$

$$\mathcal{A}[a_1 \text{ op}_a a_2](\sigma, \mathcal{H}) = \mathcal{A}[a_1](\sigma, \mathcal{H}) \text{ op}_a \mathcal{A}[a_2](\sigma, \mathcal{H})$$

$$\mathcal{A}[\text{nil}](\sigma, \mathcal{H}) = \diamond$$

where we use \wp to determine the value of pointer expressions and we explicitly write that the meaning of nil is \diamond . Also the meaning op_a of the binary operation op_a has to be suitably modified to be undefined unless both arguments are integers in which case the results are as for the WHILE-language.

The definition of the semantics of boolean expressions is similar so we only give two of the clauses:

$$\mathcal{B}[a_1 \text{ op}_r a_2](\sigma, \mathcal{H}) = \mathcal{A}[a_1](\sigma, \mathcal{H}) \text{ op}_r \mathcal{A}[a_2](\sigma, \mathcal{H})$$

$$\mathcal{B}[\text{op}_p p](\sigma, \mathcal{H}) = \text{op}_p (\wp[p](\sigma, \mathcal{H}))$$

Analogously to above, the meaning op_r of the binary relation operator op_r has to be suitably modified to give undefined in case the arguments are not integers except that the equality operation now allows testing for the equality of pointers. The meaning of the unary operation op_p is defined by op_p ; as an example:

$$\text{is-nil}(v) = \begin{cases} \text{tt} & \text{if } v = \diamond \\ \text{ff} & \text{otherwise} \end{cases}$$

Statements. Finally, the semantics of statements is extended to cope with the heap component. The configurations will now contain a state as well as a heap so we have

$$\langle [x:=a]^t, \sigma, \mathcal{H} \rangle \rightarrow \langle \sigma[x \mapsto \mathcal{A}[a](\sigma, \mathcal{H})], \mathcal{H} \rangle \\ \text{if } \mathcal{A}[a](\sigma, \mathcal{H}) \text{ is defined}$$

reflecting that for the assignment $x:=a$ the state is updated as usual and the heap is left unchanged. In the case where we assign to a pointer expression containing a selector field we shall leave the state unchanged and update the heap as follows:

$$\langle [x.sel:=a]^\ell, \sigma, \mathcal{H} \rangle \rightarrow \langle \sigma, \mathcal{H}[(\sigma(x), sel) \mapsto \mathcal{A}[a](\sigma, \mathcal{H})] \rangle$$

if $\sigma(x) \in \text{Loc}$ and $\mathcal{A}[a](\sigma, \mathcal{H})$ is defined

Here the side condition ensures that the left hand side of the assignment does indeed evaluate to a location.

The construct `malloc p` is responsible for creating a new cell. We have two clauses depending on the form of p :

$$\langle [\text{malloc } x]^\ell, \sigma, \mathcal{H} \rangle \rightarrow \langle \sigma[x \mapsto \xi], \mathcal{H} \rangle$$

where ξ does not occur in σ or \mathcal{H}

$$\langle [\text{malloc } (x.sel)]^\ell, \sigma, \mathcal{H} \rangle \rightarrow \langle \sigma, \mathcal{H}[(\sigma(x), sel) \mapsto \xi] \rangle$$

where ξ does not occur in σ or \mathcal{H} and $\sigma(x) \in \text{Loc}$

Note that in both cases we introduce a fresh location ξ but we do not specify any values for $\mathcal{H}(\xi, sel)$ – as discussed before we have settled for a semantics where the fields of ξ are undefined; obviously other choices are possible. Also note that in the last clause the side condition ensures that we already have a location corresponding to x and hence can create an edge to the new location.

Remark. The semantics only allows a limited reuse of garbage locations. For a statement like `[malloc x]1; [x:=nil]2; [malloc y]3` we will assign some location to x at the statement with label 1 and since it neither occurs in the state nor the heap after the assignment labelled 2 we are free to reuse it in the statement labelled 3 (but we do not have to). For a statement like `[malloc x]1; [x.cdr:=nil]2; [x:=nil]3; [malloc y]4` we would not be able to reuse the location allocated at 1 although it will be unreachable (and hence garbage) after the statement labelled 3. ■

2.6.2 Shape Graphs

It should be evident that there are programs for which the heap can grow arbitrarily large. Therefore the aim of the analysis will be to come up with *finite* representations of it. To do so we shall introduce a method for combining the locations of the semantics into *abstract locations* such that there is only a finite number of abstract locations. We then introduce an *abstract state* S mapping variables to abstract locations (rather than locations) and an *abstract heap* H specifying the links between the abstract locations (rather than the locations). More precisely, the analysis will operate on *shape graphs* (S, H, is) consisting of:

- an abstract state, S ,

- an abstract heap, H , and
- sharing information, is , for the abstract locations.

The last component allows us to recover some of the imprecision introduced by combining many locations into one abstract location. We shall now describe how a given state σ and heap \mathcal{h} give rise to a shape graph (S, H, is) ; in doing so we shall specify the functionality of S , H and is in detail as well as formulating a total of five invariants.

Abstract locations. The *abstract locations* have the form n_X where X is a subset of the variables of \mathbf{Var}_* :

$$\mathbf{ALoc} = \{n_X \mid X \subseteq \mathbf{Var}_*\} \quad \text{abstract locations}$$

Since \mathbf{Var}_* is finite it is clear that \mathbf{ALoc} is finite and a given shape graph will contain a subset of the abstract locations of \mathbf{ALoc} .

The idea is that if $x \in X$ then the abstract location n_X will (among others) represent the location $\sigma(x)$. The abstract location n_\emptyset is called the *abstract summary location* and it will represent all the locations that cannot be reached directly from the state without consulting the heap. Clearly n_X and n_\emptyset will represent disjoint sets of locations when $X \neq \emptyset$.

In general, we shall enforce the invariant that two distinct abstract locations n_X and n_Y always represent disjoint sets of locations. As a consequence we get that for any two abstract locations n_X and n_Y it is either the case that $X = Y$ or that $X \cap Y = \emptyset$. To see this assume that $X \neq Y$ and that $z \in X \cap Y$. From $z \in X$ we get that $\sigma(z)$ is represented by n_X and similarly $z \in Y$ gives that $\sigma(z)$ is represented by n_Y . Hence n_X and n_Y must be equal and we have a contradiction.

The invariant can be formulated as follows:

Invariant 1. If two abstract locations n_X and n_Y occur in the same shape graph then either $X = Y$ or $X \cap Y = \emptyset$.

Example 2.45 Consider the state and heap in row 2 of Figure 2.12. The variables x , y and z point to different locations (ξ_3 , ξ_2 , and ξ_1 , respectively) so in the shape graph they will be represented by different abstract locations named $n_{\{x\}}$, $n_{\{y\}}$ and $n_{\{z\}}$. The two locations ξ_4 and ξ_5 cannot be reached directly from the state so they will be represented by the abstract summary location n_\emptyset . ■

Abstract states. One of the components of a shape graph is the *abstract state*, S , that maps variables to abstract locations. To maintain the naming convention for abstract locations we shall ensure that:




Figure 2.13: Shape graphs corresponding to Figure 2.12.

Invariant 2. If x is mapped to n_x by the abstract state then $x \in X$.

From Invariant 1 it follows that there will be at most one abstract location in the shape graph containing a given variable.

We shall only be interested in the shape of the heap so we shall not distinguish between integer values, nil-pointers and uninitialised fields; hence we can view the abstract state as an element of

$$S \in \mathbf{AState} = \mathcal{P}(\mathbf{Var}_* \times \mathbf{ALoc})$$

where we have chosen to use powersets so as to simplify the notation in later parts of the development. We shall write $\mathbf{ALoc}(S) = \{n_x \mid \exists x : (x, n_x) \in S\}$ for the set of abstract locations occurring in S . (Note that \mathbf{AState} is too large in the sense that it contains elements that do not satisfy the invariants.)

Abstract heaps. Another component of the shape graph is the *abstract heap*, H , that specifies the links between the abstract locations (just as the heap specifies the links between the locations in the semantics). The links will be specified by triples (n_V, sel, n_W) and formally we take the abstract heap as an element of

$$H \in \mathbf{AHeap} = \mathcal{P}(\mathbf{ALoc} \times \mathbf{Sel} \times \mathbf{ALoc})$$

where we again do not distinguish between integers, nil-pointers and uninitialised fields. We shall write $\mathbf{ALoc}(H) = \{n_V, n_W \mid \exists sel : (n_V, sel, n_W) \in H\}$ for the set of abstract locations occurring in H .

The intention is that if $\kappa(\xi_1, sel) = \xi_2$ and ξ_1 and ξ_2 are represented by n_V and n_W , respectively, then $(n_V, sel, n_W) \in H$.

In the heap κ there will be *at most* one location ξ_2 such that $\kappa(\xi_1, sel) = \xi_2$. The abstract heap only partly shares this property because the abstract location n_\emptyset can represent several locations pointing to different locations. However, the abstract heap must satisfy:

Invariant 3. Whenever (n_V, sel, n_W) and $(n_V, sel, n_{W'})$ are in the abstract heap then either $V = \emptyset$ or $W = W'$.

Thus the target of a selector field will be uniquely determined by the source unless the source is the abstract summary location n_\emptyset .

Example 2.46 Continuing Example 2.45 we can now see that the abstract state S_2 corresponding to the state of row 2 of Figure 2.12 will be

$$S_2 = \{(x, n_{\{x\}}), (y, n_{\{y\}}), (z, n_{\{z\}})\}$$

The abstract heap H_2 corresponding to row 2 has

$$H_2 = \{(n_{\{x\}}, cdr, n_\emptyset), (n_\emptyset, cdr, n_\emptyset), (n_{\{y\}}, cdr, n_{\{z\}})\}$$

The first triple reflects that the heap maps ξ_3 and cdr to ξ_4 , ξ_3 is represented by $n_{\{x\}}$ and ξ_4 is represented by n_\emptyset . The second triple reflects that the heap maps ξ_4 and cdr to ξ_5 and both ξ_4 and ξ_5 are represented by n_\emptyset . The final triple reflects that the heap maps ξ_2 and cdr to ξ_1 , ξ_2 is represented by $n_{\{y\}}$ and ξ_1 is represented by $n_{\{z\}}$. Note that there is no triple $(n_{\{z\}}, cdr, n_\emptyset)$ because the heap maps ξ_1 and cdr to \diamond rather than a location.

The resulting abstract state and abstract heap is illustrated in Figure 2.13 together with similar information for the other states and heaps of Figure 2.12. The *square nodes* model abstract locations; the *unlabelled edges* from variables to square nodes model the abstract state and the *labelled edges* between square nodes model the abstract heap. If the abstract state does not associate an abstract location with some variable then that variable does not occur in the picture.

Note that even if the semantics uses the same locations throughout the computation it need not be the case that the locations are associated with the same abstract locations at all points in the analysis. Consider Figures 2.12 and 2.13: the abstract location n_\emptyset will in turn represent the locations $\{\xi_2, \xi_3, \xi_4, \xi_5\}$, $\{\xi_3, \xi_4, \xi_5\}$, $\{\xi_4, \xi_5\}$, $\{\xi_1, \xi_5\}$, $\{\xi_1, \xi_2\}$ and $\{\xi_1, \xi_2, \xi_3\}$. ■

Sharing information. We are now ready to introduce the third and final component of the shape graphs. Consider the top row of Figure 2.14. The abstract state and abstract heap to the right represent the state and the

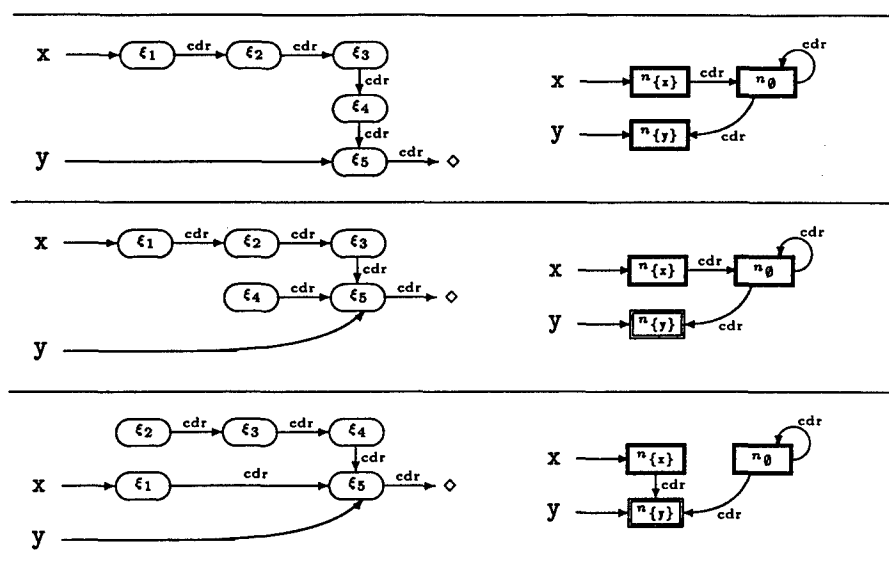


Figure 2.14: Sharing information.

heap to the left but they also represent the state and the heap shown in the second row. We shall now show how to distinguish between these two cases.

The idea is to specify a subset, is , of the abstract locations that represent locations that are shared due to pointers in the heap: an abstract location n_X will be included in is if it does represent a location that is the target of more than one pointer in the heap. In the top row of Figure 2.14, the abstract location $n_{\{y\}}$ represents the location ξ_5 and it *is not* shared (by two or more heap pointers) so $n_{\{y\}} \notin is$; the fat box indicates that the abstract location is unshared. On the other hand, in the second row ξ_5 *is* shared (both ξ_3 and ξ_4 point to it) so $n_{\{y\}} \in is$; the double box indicates that the abstract location might be shared.

Obviously, the abstract heaps themselves also contain some implicit sharing information: this is illustrated in the bottom row of Figure 2.14 where there are two distinct edges with target $n_{\{y\}}$. We shall ensure that this implicit sharing information is consistent with the explicit sharing information (as given by is) by imposing two invariants. The first ensures that information in the sharing component is also reflected in the abstract heap:

Invariant 4. If $n_X \in is$ then either

- (a) (n_0, sel, n_X) is in the abstract heap for some sel , or
- (b) there exists two distinct triples (n_V, sel_1, n_X) and (n_W, sel_2, n_X) in the abstract heap (that is either $sel_1 \neq sel_2$ or $V \neq W$).

Case 4(a) takes care of the situation where there might be several locations represented by n_\emptyset that point to n_X (as in the second row of Figure 2.14). Case 4(b) takes care of the case where two distinct pointers (with different source or different selectors) point to n_X (as in the bottom row of Figure 2.14).

The second invariant ensures that sharing information present in the abstract heap is also reflected in the sharing component:

Invariant 5. Whenever there are two distinct triples (n_V, sel_1, n_X) and (n_W, sel_2, n_X) in the abstract heap and $n_X \neq n_\emptyset$ then $n_X \in is$.

This takes care of the case where n_X represents a single location being the target of two or more heap pointers (as in the bottom row of Figure 2.14). Note that invariant 5 is the “inverse” of invariant 4(b) and that we have no “inverse” of invariant 4(a) – the presence of a pointer from n_\emptyset to n_X gives no information about sharing properties of n_X .

In the case of the abstract summary location the explicit sharing information clearly gives extra information: if $n_\emptyset \in is$ then there might be a location represented by n_\emptyset that is the target of two or more heap pointers, whereas if $n_\emptyset \notin is$ then all the locations represented by n_\emptyset will be the target of at most one heap pointer. The explicit sharing information may also give extra information for abstract locations n_X where $X \neq \emptyset$: from 4(a) alone we cannot deduce that n_X is shared – this is clearly illustrated for the node $n_{\{Y\}}$ by the top two rows of Figure 2.14.

The complete lattice of shape graphs. To summarise, a *shape graph* is a triple consisting of an abstract state S , an abstract heap H , and a set is of abstract locations that are shared:

$$\begin{aligned} S &\in \mathbf{AState} &= \mathcal{P}(\mathbf{Var}_* \times \mathbf{ALoc}) \\ H &\in \mathbf{AHeap} &= \mathcal{P}(\mathbf{ALoc} \times \mathbf{Sel} \times \mathbf{ALoc}) \\ is &\in \mathbf{IsShared} &= \mathcal{P}(\mathbf{ALoc}) \end{aligned}$$

where $\mathbf{ALoc} = \{n_Z \mid Z \subseteq \mathbf{Var}_*\}$. A shape graph (S, H, is) is a *compatible shape graph* if it fulfils the five invariants presented above:

1. $\forall n_V, n_W \in \mathbf{ALoc}(S) \cup \mathbf{ALoc}(H) \cup is : V = W \vee V \cap W = \emptyset$
2. $\forall (x, n_X) \in S : x \in X$
3. $\forall (n_V, sel, n_W), (n_V, sel, n_{W'}) \in H : V = \emptyset \vee W = W'$
4. $\forall n_X \in is : (\exists sel : (n_\emptyset, sel, n_X) \in H) \vee$
 $(\exists (n_V, sel_1, n_X), (n_W, sel_2, n_X) \in H : sel_1 \neq sel_2 \vee V \neq W)$




Figure 2.15: The extremal value ι for the list reversal program.

5. $\forall (n_V, sel_1, n_X), (n_W, sel_2, n_X) \in H :$
 $((sel_1 \neq sel_2 \vee V \neq W) \wedge X \neq \emptyset) \Rightarrow n_X \in is$

The set of compatible shape graphs is denoted

$$SG = \{(S, H, is) \mid (S, H, is) \text{ is compatible}\}$$

and the analysis, to be called *Shape*, will operate over *sets* of compatible shape graphs, i.e. elements of $\mathcal{P}(SG)$. Since $\mathcal{P}(SG)$ is a powerset, it is trivially a complete lattice with \sqcup being \cup and \sqsubseteq being \subseteq .

2.6.3 The Analysis

The analysis will be specified as an instance of a Monotone Framework with the complete lattice of properties being $\mathcal{P}(SG)$. For each label consistent program S_* with isolated entries we obtain a set of equations of the form

$$Shape_o(\ell) = \begin{cases} \iota & \text{if } \ell = \text{init}(S_*) \\ \bigcup \{Shape_o(\ell') \mid (\ell', \ell) \in \text{flow}(S_*)\} & \text{otherwise} \end{cases}$$

$$Shape_e(\ell) = f_\ell^{SA}(Shape_o(\ell))$$

where $\iota \in \mathcal{P}(SG)$ is the extremal value holding at entry to S_* and f_ℓ^{SA} are the transfer functions to be developed below. The analysis is a *forward analysis* since it is defined in terms of the set $\text{flow}(S_*)$, and it is a *may analysis* since we are using \bigcup as the combination operation.

Example 2.47 Consider again the list reversal program of Example 2.43 and assume that x initially points to an unshared list with at least two elements and that y and z are initially undefined; the singleton shape graph corresponding to this state and heap is illustrated in Figure 2.15 and will be the extremal value ι used throughout this development.

The Shape Analysis computes the sets $Shape_o(\ell)$ and $Shape_e(\ell)$ of shape graphs describing the state and heap before and after executing the elementary block labelled ℓ . The equations for $Shape_e(\ell)$ are

$$Shape_e(1) = f_1^{SA}(Shape_o(1)) = f_1^{SA}(\iota)$$

$$Shape_e(2) = f_2^{SA}(Shape_o(2)) = f_2^{SA}(Shape_e(1) \cup Shape_e(6))$$

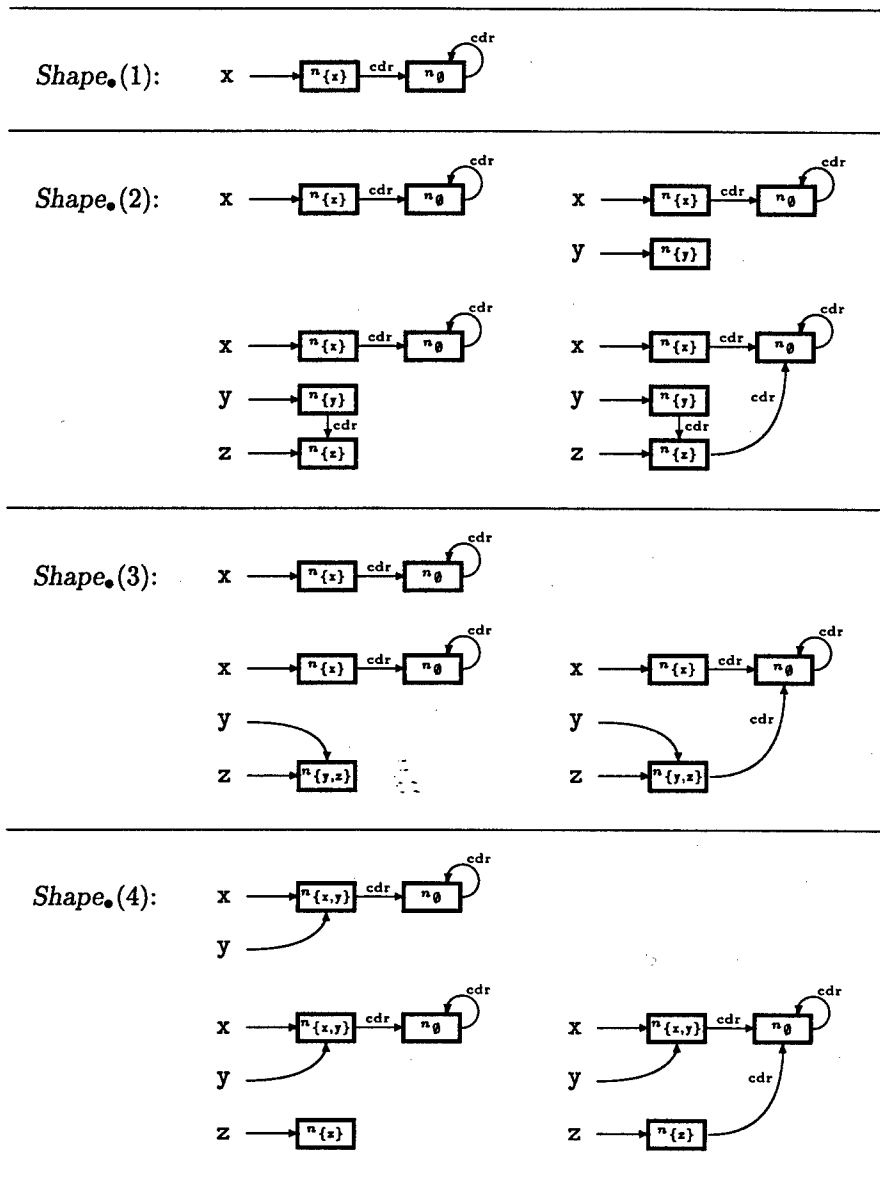


Figure 2.16: Shape graphs computed for the list reversal program (part 1).

$$\begin{aligned}
 \text{Shape}_0(3) &= f_3^{\text{SA}}(\text{Shape}_0(3)) = f_3^{\text{SA}}(\text{Shape}_0(2)) \\
 \text{Shape}_0(4) &= f_4^{\text{SA}}(\text{Shape}_0(4)) = f_4^{\text{SA}}(\text{Shape}_0(3)) \\
 \text{Shape}_0(5) &= f_5^{\text{SA}}(\text{Shape}_0(5)) = f_5^{\text{SA}}(\text{Shape}_0(4)) \\
 \text{Shape}_0(6) &= f_6^{\text{SA}}(\text{Shape}_0(6)) = f_6^{\text{SA}}(\text{Shape}_0(5))
 \end{aligned}$$

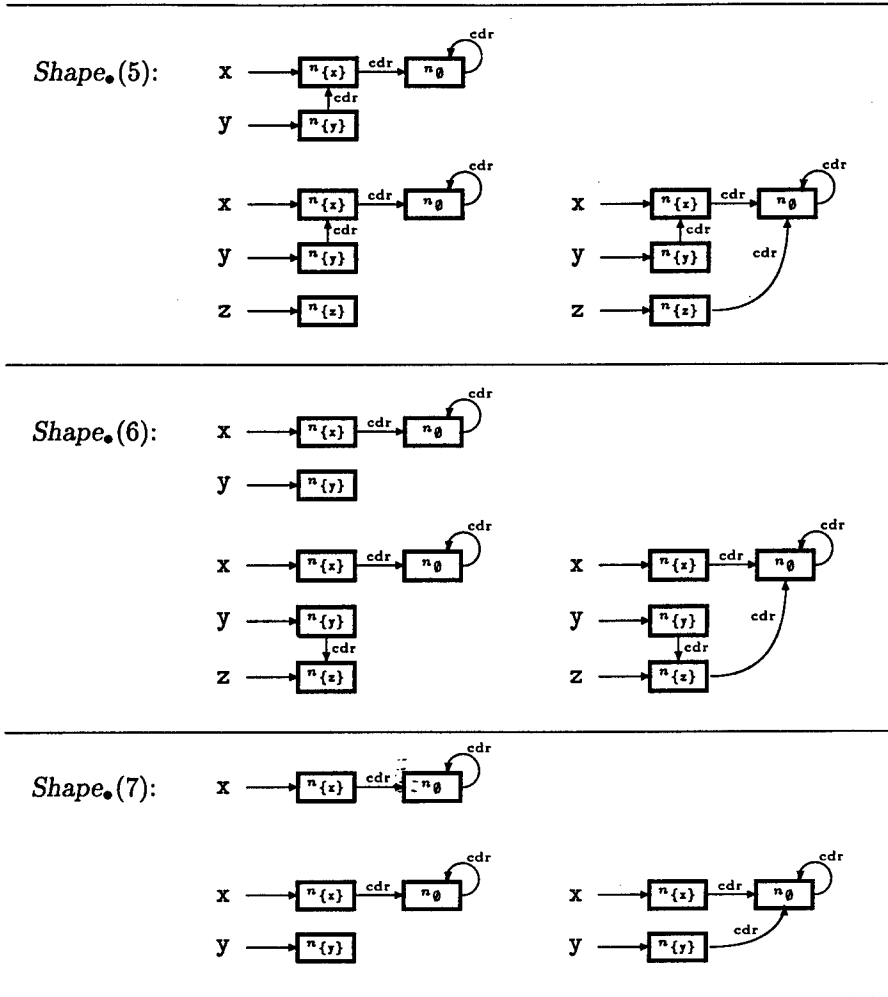


Figure 2.17: Shape graphs computed for the list reversal program (part 2).

$$Shape_{\bullet}(7) = f_7^{SA}(Shape_{\bullet}(7)) = f_7^{SA}(Shape_{\bullet}(2))$$

where the transfer functions f_l^{SA} will be specified below. The least solution to the equations is shown in Figures 2.16 and 2.17; the sets $Shape_{\bullet}(\ell)$ contain between 1 and 4 distinct shape graphs.

From $Shape_{\bullet}(2)$ we can see that at the beginning of each iteration of the loop it is the case that

- x never points to a shared list,
- y never points to a shared list,

- z may point to the second element of the list pointed to by y (if it exists), and
- x and y never point to lists that are not disjoint (because n_\emptyset is not shared).

The shape graphs in $Shape_\bullet(7)$ show that at termination y cannot point to a shared list – it will not be possible that the list is cyclic since such a list cannot be represented by any of the shape graphs in $Shape_\bullet(7)$. ■

The transfer function $f_\ell^{SA} : \mathcal{P}(\mathbf{SG}) \rightarrow \mathcal{P}(\mathbf{SG})$ associated with a label, ℓ , has the form:

$$f_\ell^{SA}(SG) = \bigcup \{ \phi_\ell^{SA}((S, H, is)) \mid (S, H, is) \in SG \}$$

where $\phi_\ell^{SA} : \mathbf{SG} \rightarrow \mathcal{P}(\mathbf{SG})$ specifies how a *single* shape graph (in $Shape_\bullet(\ell)$) may be transformed into a *set* of shape graphs (in $Shape_\bullet(\ell)$) by the elementary block labelled ℓ . We shall now inspect the various forms of elementary block and specify ϕ_ℓ^{SA} in each case. We shall first consider the boolean expressions and the skip-statement, then the different forms of assignments and finally the malloc-statement.

Transfer function for $[b]^\ell$ and $[\text{skip}]^\ell$. We are only interested in the shape of the heap and the boolean tests do not modify the heap. Hence we take

$$\phi_\ell^{SA}((S, H, is)) = \{(S, H, is)\}$$

so that the transfer function f_ℓ^{SA} will be the identity function. Similarly for the skip-statement.

Example 2.48 This case is illustrated by the test $[\text{not is-nil}(x)]^2$ of the list reversal program of Example 2.43: the transfer function f_2^{SA} is the identity function. Hence $Shape_\bullet(2) = Shape_\bullet(1) \cup Shape_\bullet(6)$ in Figures 2.16 and 2.17. ■

Transfer function for $[x := a]^\ell$ where a is of the form n , $a_1 \text{ op}_a a_2$ or nil . The effect of this assignment will be to remove the binding to x , and to rename all abstract locations so that they do not include x in their name. The renaming of abstract locations is specified by the function

$$k_x(n_Z) = n_{Z \setminus \{x\}}$$

and we then take

$$\phi_\ell^{SA}((S, H, is)) = \{kill_x((S, H, is))\}$$

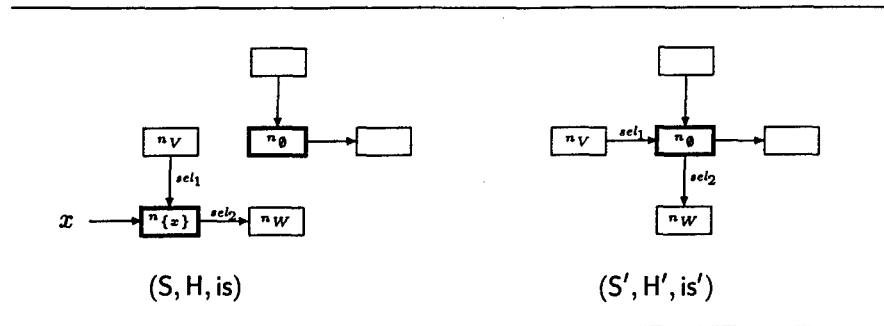


Figure 2.18: The effect of $[x:=\text{nil}]^t$.

where $\text{kill}_x((S, H, \text{is})) = (S', H', \text{is}')$ is given by

$$\begin{aligned} S' &= \{(z, k_x(n_z)) \mid (z, n_z) \in S \wedge z \neq x\} \\ H' &= \{(k_x(n_V), \text{sel}, k_x(n_W)) \mid (n_V, \text{sel}, n_W) \in H\} \\ \text{is}' &= \{k_x(n_X) \mid n_X \in \text{is}\} \end{aligned}$$

It is easy to check that if (S, H, is) is compatible then so is (S', H', is') .

Example 2.49 The statement $[y:=\text{nil}]^1$ of the list reversal program of Example 2.43 is of the form considered here. Since there is no occurrence of y in the shape graph ι of Figure 2.15, the shape graph $\text{Shape}_\bullet(1)$ in Figure 2.16 is equal to ι . ■

An interesting case is when $(x, n_{\{x\}}) \in S$ since this will cause the two abstract locations $n_{\{x\}}$ and n_{\emptyset} to be merged. The sharing information is then updated to capture that we can only be sure that n_{\emptyset} is unshared in the updated shape graph if both n_{\emptyset} and $n_{\{x\}}$ were unshared in the original shape graph. This is illustrated in Figure 2.18: the left hand picture shows the interesting parts of the shape graph (S, H, is) and the right hand picture shows the corresponding parts of (S', H', is') . We shall assume that the square boxes represent distinct abstract locations so in particular V , $\{x\}$, W and \emptyset are all distinct sets. The fat boxes represent unshared abstract locations as before, the thin boxes represent abstract locations whose sharing information is not affected by the transfer function, and unlabelled edges between abstract locations represent pointers that are unaffected by the transfer function.

Example 2.50 The statement $[z:=\text{nil}]^7$ of the list reversal program of Example 2.43 illustrates this case: for each of the shape graphs of $\text{Shape}_\bullet(2)$ in Figures 2.16 and 2.17 the abstract location $n_{\{z\}}$ is merged with n_{\emptyset} to produce one of the shape graphs of $\text{Shape}_\bullet(7)$. ■

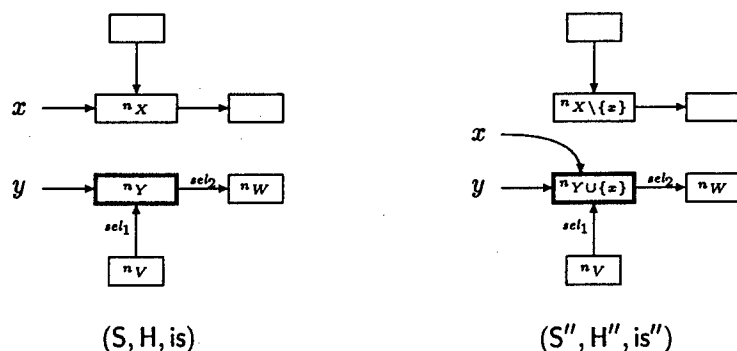


Figure 2.19: The effect of $[x:=y]^l$ when $x \neq y$.

Remark. The analysis does *not* perform garbage collection: it might be the case that there are no heap pointers to $n_{\{x\}}$ and then the corresponding location in the heap will be unreachable after the assignment. Nonetheless the analysis will still merge the two abstract locations $n_{\{x\}}$ and n_{\emptyset} and insist on a pointer from n_{\emptyset} to any abstract location that $n_{\{x\}}$ might point to. ■

Transfer function for $[x:=y]^l$. If $x = y$ then the transfer function f_l^{SA} is just the identity.

Next suppose that $x \neq y$. The first effect of the assignment is to remove the old bindings to x ; for this we use the $kill_x$ operation introduced above. Then the new binding to x is recorded; this includes renaming the abstract location that includes y in its variable set to also include x . The renaming of the abstract locations is specified by the function:

$$g_x^y(n_Z) = \begin{cases} n_{Z \cup \{x\}} & \text{if } y \in Z \\ n_Z & \text{otherwise} \end{cases}$$

We shall then take

$$\phi_l^{SA}((S, H, is)) = \{(S'', H'', is'')\}$$

where $(S', H', is') = kill_x((S, H, is))$ and

$$\begin{aligned} S'' &= \{(z, g_x^y(n_Z)) \mid (z, n_Z) \in S'\} \\ &\quad \cup \{(x, g_x^y(n_Y)) \mid (y', n_Y) \in S' \wedge y' = y\} \\ H'' &= \{(g_x^y(n_V), sel, g_x^y(n_W)) \mid (n_V, sel, n_W) \in H'\} \\ is'' &= \{g_x^y(n_Z) \mid n_Z \in is'\} \end{aligned}$$

Here the second clause in the formula for S'' adds the new binding to x . Again we note that if (S, H, is) is compatible then so is (S'', H'', is'') .

The clause is illustrated in Figure 2.19 where we assume that nodes represent distinct abstract locations; it follows from the invariants that $y \in Y$ but $y \notin V$ and $y \notin W$. Note that $n_{Y \cup \{x\}}$ inherits the sharing properties of n_Y although both x and y will point to the same cell; the reason is that the sharing information only records sharing in the heap – not sharing via the state.

Example 2.51 The statement $[y:=x]^4$ of the list reversal program of Example 2.43 is of the form considered here: each of the shape graphs of $\text{Shape}_\bullet(3)$ in Figure 2.16 is transformed into one of the shape graphs of $\text{Shape}_\bullet(4)$.

Also the statement $[z:=y]^3$ is of the form considered here: each of the shape graphs of $\text{Shape}_\bullet(2)$ is transformed into one of the shape graphs of $\text{Shape}_\bullet(3)$; indeed two of the shape graphs in $\text{Shape}_\bullet(2)$ give rise to the same shape graph in $\text{Shape}_\bullet(3)$. ■

Transfer function for $[x:=y.sel]^\ell$. First assume that $x = y$; then the assignment is semantically equivalent to the following sequence of assignments

$$[t:=y.sel]^{\ell_1}; [x:=t]^{\ell_2}; [t:=\text{nil}]^{\ell_3}$$

where t is a fresh variable and ℓ_1, ℓ_2 and ℓ_3 are fresh labels. The transfer function f_t^{SA} can therefore be obtained as

$$f_t^{\text{SA}} = f_{\ell_3}^{\text{SA}} \circ f_{\ell_2}^{\text{SA}} \circ f_{\ell_1}^{\text{SA}}$$

where the transfer functions $f_{\ell_2}^{\text{SA}}$ and $f_{\ell_3}^{\text{SA}}$ follow the pattern described above. We shall therefore concentrate on the transfer function $f_{\ell_3}^{\text{SA}}$, or equivalently, f_t^{SA} in the case where $x \neq y$.

Example 2.52 The statement $[x:=x.cdr]^5$ of the list reversal program of Example 2.43 is transformed into $[t:=x.cdr]^{51}; [x:=t]^{52}; [t:=\text{nil}]^{53}$. We shall return to the analysis of $[t:=x.cdr]^{51}$ later. ■

So assume that $x \neq y$ and let (S, H, is) be a compatible shape graph before the analysis of the statement. As in the previous case, the first step will be to remove the old binding for x and again we use the auxiliary function kill_x :

$$(S', H', \text{is}') = \text{kill}_x((S, H, \text{is}))$$

The next step will be to rename the abstract location corresponding to $y.sel$ to include x in its name and to establish the binding of x to that abstract location. We can now identify three possibilities:

1. There is no abstract location n_Y such that $(y, n_Y) \in S'$ or there is an abstract location n_Y such that $(y, n_Y) \in S'$ but no n_Z such that $(n_Y, \text{sel}, n_Z) \in H'$; in this case the shape graph will represent a state and a heap where y or $y.sel$ is an integer, nil or undefined.

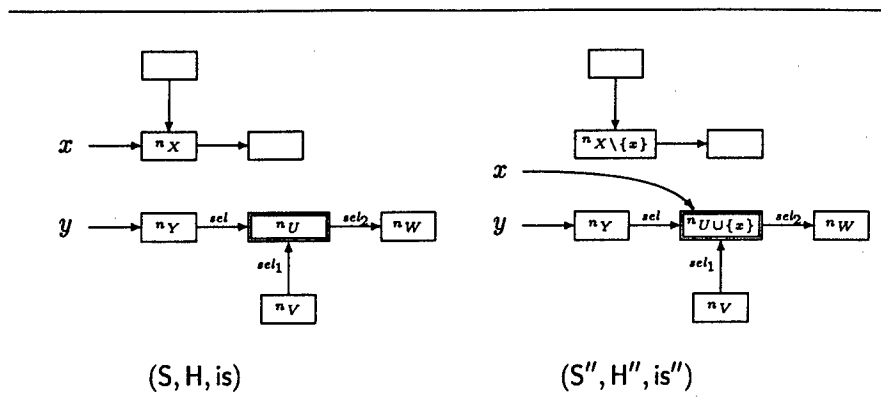


Figure 2.20: The effect of $[x:=y.sel]^l$ in Case 2 when $x \neq y$.

2. There is an abstract location n_Y such that $(y, n_Y) \in S'$ and there is an abstract location $n_U \neq n_\emptyset$ such that $(n_Y, sel, n_U) \in H'$; in this case the shape graph will represent a state and a heap where the location pointed to by $y.sel$ will also be pointed to by some other variable (in U).
3. There is an abstract location n_Y such that $(y, n_Y) \in S'$ and $(n_Y, sel, n_\emptyset) \in H'$; in this case the shape graph will represent a state and a heap where no other variable points to the location pointed to by $y.sel$.

Case 1. First consider the statement $[x:=y.sel]^l$ (where $x \neq y$) in the case where there is no abstract location n_Y such that $(y, n_Y) \in S'$. Then there is no abstract location for $y.sel$ and hence no abstract location to rename and no binding to establish. Thus we take:

$$\phi_l^{SA}((S, H, is)) = \{kill_x((S, H, is))\}$$

Note that this situation captures the case where an attempt is made to dereference a nil-pointer.

Alternatively, there is an abstract location n_Y such that $(y, n_Y) \in S'$ but there is no abstract location n such that $(n_Y, sel, n) \in H'$. From the invariants it follows that n_Y is unique but still there is no abstract location to rename and no binding to establish. So again we take:

$$\phi_l^{SA}((S, H, is)) = \{kill_x((S, H, is))\}$$

This situation captures the case where an attempt is made to dereference a non-existing selector field of pointer.

Case 2. We consider the statement $[x:=y.sel]^l$ (where $x \neq y$) in the case where there is an abstract location n_Y such that $(y, n_Y) \in S'$ and there is an

abstract location $n_U \neq n_\emptyset$ such that $(n_Y, sel, n_U) \in H'$. Both n_Y and n_U will be uniquely determined because of the invariants (and they might be equal). The abstract location n_U will be renamed to include the variable x using the function:

$$h_x^U(n_Z) = \begin{cases} n_{U \cup \{x\}} & \text{if } Z = U \\ n_Z & \text{otherwise} \end{cases}$$

We shall then take

$$\phi_t^{SA}((S, H, is)) = \{(S'', H'', is'')\}$$

where $(S', H', is') = kill_x((S, H, is))$ and

$$\begin{aligned} S'' &= \{(z, h_x^U(n_Z)) \mid (z, n_Z) \in S'\} \cup \{(x, h_x^U(n_U))\} \\ H'' &= \{(h_x^U(n_V), sel', h_x^U(n_W)) \mid (n_V, sel', n_W) \in H'\} \\ is'' &= \{h_x^U(n_Z) \mid n_Z \in is'\} \end{aligned}$$

The inclusion of $(x, h_x^U(n_U))$ in S'' reflects the assignment. The definition of is'' ensures that sharing is preserved by the operation; in particular, $n_{U \cup \{x\}}$ is shared in H'' if and only if n_U is shared in H' .

The effect of the assignment is illustrated in Figure 2.20 in the case where $n_U \in is$. As before we assume that the abstract locations shown on the figure are distinct so in particular Y, V and W are all distinct from U .

Case 3. We now consider the statement $[x:=y.sel]^t$ (where $x \neq y$) in the case where there is an abstract location n_Y such that $(y, n_Y) \in S'$ and furthermore $(n_Y, sel, n_\emptyset) \in H'$. As before the invariants ensure that n_Y is uniquely determined. The location n_\emptyset describes the location for $y.sel$ as well as a (possibly empty) set of other locations. We now have to *materialise* a new abstract location $n_{\{x\}}$ from n_\emptyset ; then $n_{\{x\}}$ will describe the location for $y.sel$ and n_\emptyset will continue to represent the remaining locations. Having introduced a new abstract location we will have to modify the abstract heap accordingly.

This is a potentially difficult operation, so let us consider the following sequence of assignments:

$$\begin{array}{ccccccc} & & [x:=nil]^{t'}; & [x:=y.sel]^t; & [x:=nil]^{t''} & & \\ & \uparrow & & \uparrow & \uparrow & & \uparrow \\ (S, H, is) & & (S', H', is') & & (S'', H'', is'') & & (S''', H''', is''') \end{array}$$

Clearly $[x:=nil]^{t'}; [x:=y.sel]^t$ is equivalent to $[x:=y.sel]^t$ both in terms of the analysis and the semantics. Indeed, $(S', H', is') = kill_x((S, H, is))$ represents the effect of removing the binding to x . We are trying to determine candidate shape graphs (S'', H'', is'') holding after the assignment $[x:=y.sel]^t$ (where

$x \neq y$) but let us first study our expectations to (S''', H''', is''') . It is immediate that $(S''', H''', is''') = kill_x((S'', H'', is''))$. Furthermore, the states and heaps possible at the point described by (S', H', is') should be the same as those possible at the point described by (S''', H''', is''') . This suggests demanding that

$$(S''', H''', is''') = (S', H', is')$$

which means that $kill_x((S'', H'', is'')) = (S', H', is')$. It is also immediate that $(x, n_{\{x\}}) \in S''$ and that $(n_Y, sel, n_{\{x\}}) \in H''$.

We shall then take

$$\phi_t^{SA}((S, H, is)) = \{(S'', H'', is'') \mid (S'', H'', is'') \text{ is compatible} \wedge \\ kill_x((S'', H'', is'')) = (S', H', is') \wedge \\ (x, n_{\{x\}}) \in S'' \wedge (n_Y, sel, n_{\{x\}}) \in H''\}$$

where $(S', H', is') = kill_x((S, H, is))$.

It is hopefully clear that we have not missed any shape graphs (S'', H'', is'') that might be the result of the assignment. What might be a worry is that we have included an excessive amount of irrelevant shape graphs. (Indeed producing all compatible shape graphs would be trivially sound but also utterly useless.) Although it is possible to do slightly better (see Exercise 2.22) we shall now argue that there is no excessive amount of imprecision in the above definition.

We first establish that

$$S'' = S' \cup \{(x, n_{\{x\}})\}$$

showing that the abstract state is fully determined. Consider $(z, n_Z) \in S''$. If $z = x$ it follows from the compatibility of (S'', H'', is'') that $n_Z = n_{\{x\}}$. If $z \neq x$ it follows from $(x, n_{\{x\}}) \in S''$ and the compatibility of (S'', H'', is'') that $x \notin Z$ and hence $(z, n_Z) = (z, k_x(n_Z))$ (where k_x is the renaming of the $kill_x$ operation). This establishes that $S'' \subseteq S' \cup \{(x, n_{\{x\}})\}$. Next consider $(u, n_U) \in S'$. We know that $u \neq x$ and $x \notin U$ from the definition of S' and from compatibility of (S', H', is') . There must exist $(u, n'_U) \in S''$ such that $k_x(n'_U) = n_U$ but since $x \neq u$ this gives $n'_U = n_U$. It follows that $S'' \supseteq S' \cup \{(x, n_{\{x\}})\}$ and we have proved the required equality.

We next establish that

$$is' \setminus \{n_\emptyset\} = is'' \setminus \{n_\emptyset, n_{\{x\}}\} \\ n_\emptyset \in is' \text{ iff } n_\emptyset \in is'' \vee n_{\{x\}} \in is''$$

showing that

- abstract locations apart from n_\emptyset retain their sharing information,
- if n_\emptyset is shared then that sharing cannot go away but must give rise to sharing of at least one of n_\emptyset or $n_{\{x\}}$, and

- if n_\emptyset is not shared then no sharing can be introduced for n_\emptyset or $n_{\{x\}}$.

Since both (S', H', is') and (S'', H'', is'') are compatible shape graphs it follows that if $n_U \in S'$ then $x \notin U$ and if $n_U \in S''$ then $x \notin U \cup \{x\} = U$. Hence $is' = \{k_x(n_U) \mid n_U \in is''\}$ establishes $is' \setminus \{n_\emptyset\} = is'' \setminus \{n_\emptyset, n_{\{x\}}\}$ because $k_x(n_U) = n_U \neq n_\emptyset$ for all $n_U \in is'' \setminus \{n_\emptyset, n_{\{x\}}\}$. Furthermore, $n_\emptyset \in is'' \vee n_{\{x\}} \in is''$ gives $n_\emptyset \in is'$, and $n_\emptyset \notin is'' \wedge n_{\{x\}} \notin is''$ gives $n_\emptyset \notin is'$. Thus we have established the required relationship.

We now turn to the abstract heap. We shall classify the labelled edges (n_V, sel', n_W) into four groups depending on whether or not the source or target may be one of the nodes n_\emptyset or $n_{\{x\}}$:

(n_V, sel', n_W) is *external* iff $\{n_V, n_W\} \cap \{n_\emptyset, n_{\{x\}}\} = \emptyset$

(n_V, sel', n_W) is *internal* iff $\{n_V, n_W\} \subseteq \{n_\emptyset, n_{\{x\}}\}$

(n_V, sel', n_W) is *going-out* iff $n_V \in \{n_\emptyset, n_{\{x\}}\} \wedge n_W \notin \{n_\emptyset, n_{\{x\}}\}$

(n_V, sel', n_W) is *going-in* iff $n_V \notin \{n_\emptyset, n_{\{x\}}\} \wedge n_W \in \{n_\emptyset, n_{\{x\}}\}$

We shall also say that two edges (n_V, sel', n_W) and (n'_V, sel'', n'_W) are *related* if and only if $k_x(n_V) = k_x(n'_V)$, $sel' = sel''$ and $k_x(n_W) = k_x(n'_W)$. Clearly an external edge is related only to itself.

Reasoning as above one can show that

- H' and H'' have the same external edges,
- each internal edge in H' is related to an internal edge in H'' and vice versa,
- each edge going-out in H' is related to an edge going-out in H'' and vice versa, and
- each edge going-in in H' is related to an edge going-in in H'' and vice versa.

One consideration is that the going-in edge $(n_Y, sel, n_\emptyset) \in H'$ should be *changed* into the going-in edge $(n_Y, sel, n_{\{x\}}) \in H''$. We clearly demanded that $(n_Y, sel, n_{\{x\}}) \in H''$ and because (S'', H'', is'') is compatible it follows that $(n_Y, sel, n_\emptyset) \notin H''$.

As a more concrete illustration consider the scenario in Figure 2.21. Here neither n_\emptyset nor n_W is shared and we assume that both n_V and n_W are distinct from n_\emptyset ; we also assume that $x \neq y$ and $sel_2 \neq sel_3$. The result of the transfer function is shown in Figure 2.22. First note that the going-in edge $(n_Y, sel, n_\emptyset) \in H$ is changed to $(n_Y, sel, n_{\{x\}}) \in H'_i$ in all shape graphs. Next note that the going-in edge labelled sel_1 can only point to n_\emptyset because $n_{\{x\}}$ is




Figure 2.21: The effect of $[x:=y.sel]^{\ell}$ in a special case (part 1).

not shared (as n_0 is not) and n_Y points to $n_{\{x\}}$. The going-out edge labelled sel_2 can start at both n_0 and $n_{\{x\}}$ but it cannot do so simultaneously because n_W is not shared. The internal edge labelled sel_3 can only point to n_0 because $n_{\{x\}}$ is not shared and n_Y points to $n_{\{x\}}$; but it can start at both n_0 and $n_{\{x\}}$ and can even do so simultaneously. This explains why there are only six shape graphs in $\phi_t^{SA}((S, H, is))$, all of which are clearly needed.

Example 2.53 The statement- $[t:=x.cdr]^{51}$ introduced in Example 2.52 is of the form considered here: the transfer function will transform each of the shape graphs of $Shape_0(4)$ in Figure 2.16 into one of the shape graphs in Figure 2.23. ■

Transfer function for $[x.sel:=a]^{\ell}$ where a is of the form $n, a_1 op_a a_2$ or nil . Again we consider a compatible shape graph (S, H, is) . First assume that there is no n_X such that $(x, n_X) \in S$; then x will not point to a cell in the heap and the statement will have no effect on the shape of the heap so the transfer function f_t^{SA} is just the identity. Next assume that there is a (unique) n_X such that $(x, n_X) \in S$ but that there is no n_U such that $(n_X, sel, n_U) \in H$; then the cell pointed to by sel does not point to another cell so the statement will not change the shape of the heap and also in this case the transfer function f_t^{SA} will be the identity.

The interesting case is when there are abstract locations n_X and n_U such that $(x, n_X) \in S$ and $(n_X, sel, n_U) \in H$; these abstract locations will be unique because of the invariants. The effect of the assignment will be to remove the triple (n_X, sel, n_U) from H :

$$\phi_t^{SA}((S, H, is)) = \{kill_{x.sel}((S, H, is))\}$$

where $kill_{x.sel}((S, H, is)) = (S', H', is')$ is given by:

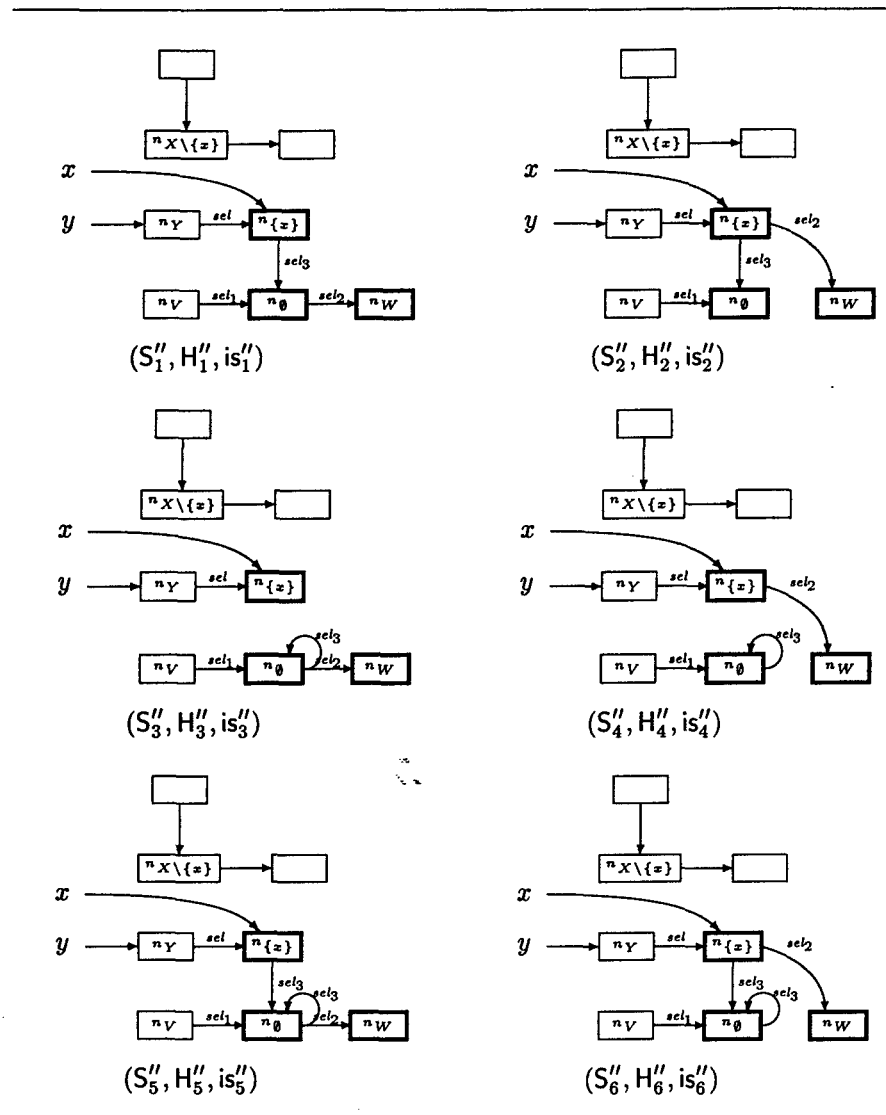


Figure 2.22: The effect of $[x:=y.sel]^l$ in a special case (part 2).

$$\begin{aligned}
 S' &= S \\
 H' &= \{(n_V, sel', n_W) \mid (n_V, sel', n_W) \in H \wedge \neg(X = V \wedge sel = sel')\} \\
 is' &= \begin{cases} is \setminus \{n_U\} & \text{if } n_U \in is \wedge \#into(n_U, H') \leq 1 \wedge \\ & \neg \exists (n_{\emptyset}, sel', n_U) \in H' \\ is & \text{otherwise} \end{cases}
 \end{aligned}$$

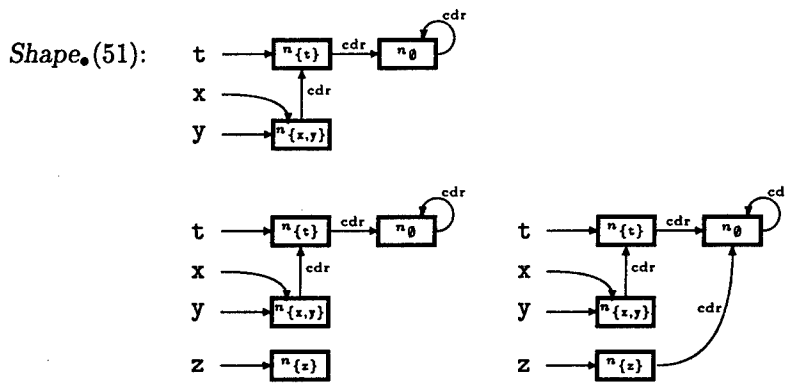


Figure 2.23: Shape graphs computed for the list reversal program.




Figure 2.24: The effect of $[x.sel:=nil]^\ell$ when $\#into(n_U, H') \leq 1$.

The sharing information is as before except that we may be able to do better for the node n_U – we have removed one of the pointers to it and in the case where there is at most one pointer left and it does not have source n_0 the corresponding location will be unshared. Here we write $\#into(n_U, H')$ for the number of pointers to n_U in H' . This clause is illustrated in Figure 2.24.

Remark. Again we shall note that the analysis does *not* incorporate garbage collection: it might be the case that there is only one pointer to the abstract location n_U and that after the assignment $x.sel:=nil$ the corresponding location will be unreachable. However, the abstract location may still be part of the shape graph. ■

Transfer function for $[x.sel:=y]^\ell$. First assume that $x = y$. The statement is then semantically equivalent to

$$[t:=y]^{\ell_1}; [x.sel:=t]^{\ell_2}; [t:=nil]^{\ell_3}$$




Figure 2.25: The effect of $[x.sel:=y]^t$ when $\#into(n_Y, H') < 1$.

where t is a fresh variable and ℓ_1 , ℓ_2 and ℓ_3 are fresh labels. The transfer function f_t^{SA} is then given by

$$f_t^{SA} = f_{\ell_3}^{SA} \circ f_{\ell_2}^{SA} \circ f_{\ell_1}^{SA}$$

The transfer functions $f_{\ell_1}^{SA}$ and $f_{\ell_3}^{SA}$ follow the pattern we have seen before so we shall concentrate on the clause for $f_{\ell_2}^{SA}$, or equivalently, f_t^{SA} in the case where $x \neq y$.

So assume that $x \neq y$ and that (S, H, is) is a compatible shape graph. It may be the case that there is no n_X such that $(x, n_X) \in S$ and in that case the transfer function will be the identity since the statement cannot affect the shape of the heap.

So assume that n_X satisfies $(x, n_X) \in S$. The case where there is no n_Y such that $(y, n_Y) \in S$ corresponds to a situation where the value of y is an integer, the nil-value or undefined and is therefore similar to the case $[x.sel:=nil]^t$:

$$\phi_t^{SA}((S, H, is)) = \{kill_{x.sel}((S, H, is))\}$$

The interesting case is when $x \neq y$, $(x, n_X) \in S$ and $(y, n_Y) \in S$. The first step will be to remove the binding for $x.sel$ and for this we can use the $kill_{x.sel}$ function. The second step will be to establish the new binding. So we take

$$\phi_t^{SA}((S, H, is)) = \{(S'', H'', is'')\}$$

where $(S', H', is') = kill_{x.sel}((S, H, is))$ and

$$S'' = S' (= S)$$

$$H'' = H' \cup \{(n_X, sel, n_Y)\}$$

$$is'' = \begin{cases} is' \cup \{n_Y\} & \text{if } \#into(n_Y, H') \geq 1 \\ is' & \text{otherwise} \end{cases}$$

Note that the node n_γ might become shared when we add a new pointer to it. The effect of the transfer function is illustrated in Figure 2.25.

Example 2.54 This transfer function is illustrated by the assignment $[y.cdr:=z]^6$ of the list reversal program of Example 2.43: each of the shape graphs of $Shape_*(5)$ in Figure 2.17 are transformed into one of the shape graphs of $Shape_*(6)$. ■

Transfer function for $[x.sel:=y.sel']^\ell$. This statement is equivalent to the sequence of statement

$$[t:=y.sel']^{\ell_1}; [x.sel:=t]^{\ell_2}; [t:=nil]^{\ell_3}$$

for t being a fresh variable and ℓ_1 , ℓ_2 and ℓ_3 being fresh labels. Thus the transfer function f_t^{SA} satisfies

$$f_t^{SA} = f_{\ell_3}^{SA} \circ f_{\ell_2}^{SA} \circ f_{\ell_1}^{SA}$$

The transfer functions $f_{\ell_1}^{SA}$, $f_{\ell_2}^{SA}$ and $f_{\ell_3}^{SA}$ all follow the patterns we have seen before so this completes the specification of the transfer function.

Transfer function for $[malloc p]^\ell$. We first consider the statement $[malloc x]^\ell$ where we have to remove the binding for x and then introduce a new (unshared) location pointed to by x . Thus we define

$$\phi_t^{SA}((S, H, is)) = \{(S' \cup \{(x, n_{\{x\}})\}, H', is')\}$$

where $(S', H', is') = kill_x(S, H, is)$.

The statement $[malloc (x.sel)]^\ell$ is equivalent to the sequence

$$[malloc t]^{\ell_1}; [x.sel:=t]^{\ell_2}; [t:=nil]^{\ell_3}$$

where t is a fresh variable and ℓ_1 , ℓ_2 and ℓ_3 are fresh labels. The transfer function f_t^{SA} is then

$$f_t^{SA} = f_{\ell_3}^{SA} \circ f_{\ell_2}^{SA} \circ f_{\ell_1}^{SA}$$

The transfer functions $f_{\ell_1}^{SA}$, $f_{\ell_2}^{SA}$ and $f_{\ell_3}^{SA}$ all follow the patterns we have seen before so this completes the specification of the transfer function.

Concluding Remarks

Data Flow Analysis for imperative languages. As mentioned in the beginning of this chapter, Data Flow Analysis has a long tradition. Most compiler textbooks contain sections on optimisation which mainly discuss Data Flow Analyses and their implementation [2, 39, 141]. The emphasis in these books is often on practical implementations of data flow analyses. A

classic textbook which provides a more theoretical treatment of the subject is by Hecht [52]; the book contains a detailed discussion of the four example Data Flow Analyses in Section 2.1, and also presents a more traditional treatment of Monotone Frameworks based on the use of semi-lattices as well as a number of algorithms (see Chapter 6 for a more thorough treatment of algorithms). Marlowe and Ryder [78] provide a survey of data flow frameworks. Steffen [129] and Schmidt [117] express data flow analyses using modal logic (rather than equations) thereby opening up for using model checking techniques for program analysis.

The examples presented in Section 2.1 are fairly standard. Alternative treatments of this material can be found in any of the books already cited. The examples may all be represented as Bit Vector Frameworks (see Exercise 2.9): the lattice elements may be represented by a vector of bits and the lattice operations efficiently implemented as boolean operations. The method used in Section 2.2 to prove the correctness of the Live Variables Analysis is adapted from [84]. The notion of faint variables, introduced in Exercise 2.4, was first introduced by Giegerich, Möncke and Wilhelm [49].

The use of semi-lattices in Data Flow Analysis was first proposed in [72]. The notion of Monotone Frameworks is due to Kam and Ullman [70]. These early papers, and much of the later literature, use the dual notions (meets and maximal fixed points) to our presentation. Kam and Ullman [70] prove that the existence of a general algorithm to compute MOP solutions would imply the decidability of the Modified Post Correspondence Problem [57]. Cousot and Cousot [27] model abstract program properties by complete semi-lattices in their paper on Abstract Interpretation (see Chapter 4).

We have associated transfer functions with elementary blocks. It would be possible to associate transfer functions with flows instead as e.g. in [113]. These two approaches have equal power: to go from the first to the second, the transfer functions may be moved from the blocks to their outgoing flows; to go from the second to the first, we can introduce artificial blocks. In fact artificial blocks can be avoided as shown in Exercise 2.11.

Most of the papers that we have cited so far concentrate on intraprocedural analysis. An early, and influential, paper on interprocedural analysis is [120] that studies two approaches to establishing context. One is based on call strings and expresses aspects of the dynamic calling context; our presentation is inspired by [139]. The other is based on data and shares some of the aims of assumption sets [74, 106, 112]; the technical formulation is different because [120] obtains the effect by calculating the transfer functions for the call statement. Most of the subsequent papers in the literature can be seen as variations and combinations over this theme; a substantial effort in this direction may be found in [34].

Pointer analysis. There is an extensive literature on the analysis of alias problems for languages with pointers. Following [46] we can distinguish between analyses of pointers to (1) *statically* allocated data (typically on the stack) and (2) *dynamically* allocated data (typically in the heap). The analysis of pointers to statically allocated data is the simplest: typically the data will have compile-time names and the analysis result can be presented as a set of points-to pairs of the form (p, x) meaning that the pointer p points to the data x or as alias pairs of the form $(*p, x)$ meaning that $*p$ and x are aliased. Analyses in this category include [36, 75, 112, 142, 127, 119].

The analysis of dynamically allocated data is more complicated since the objects of interest are inherently anonymous. The simplest analyses [28, 45] study the *connectivity* of the heap: they attempt to split the heap into disjoint parts and do not keep any information about the internal structure of the individual parts. These analyses have been found quite useful for many applications.

The more complex analyses of dynamically allocated data give more precise information about the *shape* of the heap. A number of approaches use *graphs* to represent the heap. A main distinction between these approaches is how they map a heap of potentially unbounded size to a graph of bounded size: some bound the length of paths in the heap [64, 130], others merge heap cells created at the same program point [65, 20], and yet others merge heap cells that cannot be kept apart by the set of pointer variables pointing to them [114, 115]. Another group of analyses obtain information about the shape of the heap by more directly approximating the *access paths*. Here a main distinction is the kind of properties of the access paths that are recorded: some focus on simple connectivity properties [46], others use some limited form of regular expressions [76], and yet others use monomial relations [35].

The analysis presented in Section 2.6 is based on the work of Sagiv, Reps and Wilhelm [114, 115]. In contrast to [114, 115] it uses *sets* of compatible shape graphs; [114, 115] merges sets of compatible shape graphs into a *single* summary shape graph and then use various mechanisms for extracting parts of the individual compatible shape graphs and in this way an exponential factor in the cost of the analysis can be avoided. The sharing component of the shape graphs is designed to detect list-like properties; it can be replaced by other components detecting other shape properties [116].

Static Single Assignments. Prior to analysis, many modern compilers perform a source-to-source transformation which produces a program in *static single assignment* (SSA) form. In SSA form there is a single assignment for each name; this is achieved by variable renaming. The transformation also introduces new assignments which combine the results from several assignments (using the so-called ϕ functions) at certain points. Uses of names on the right hand side of an assignment refer to the result of a unique assignment of one of the previous two forms. The advantage of transforming the pro-

gram in this way is that some analyses become simpler or more effective. An efficient algorithm for computing SSA form was introduced in [32]. The first algorithm gives conservative results for arrays; [73] presents a new approach to array SSA form that captures precise element-level data flow information for array variables.

Data Flow Analysis for other language paradigms. The analysis techniques that we have studied assume the existence of some representation of the flow of control in the program. For the class of imperative languages that we have studied, it is relatively easy to determine this control flow information. For many languages, for example functional programming languages, this is not the case. The next chapter presents techniques for determining control flow information for such languages and shows how Data Flow Analysis can be integrated with Control Flow Analysis.

The techniques we have presented can be applied directly to other language paradigms. Two examples are in object-oriented programming and a communicating processes language. In both cases the authors describe their analyses by data flow equations. Vitek, Horspool and Uhl [139] present an analysis for object-oriented languages which determines classes of objects and their lifetimes. Their analysis is an interprocedural analysis that uses a graph-based representation of the memory as data flow values. Reif and Smolka [109] apply Data Flow Analysis techniques to distributed communicating processes to detect unreachable code and to determine the values of program expressions. They apply their analysis to a language with asynchronous communication. Their reachability analysis is based on an algorithm that builds a spanning tree for each process flow graph and links matching transmits and receives between processes. They construct a Monotone Framework for determining value sets.

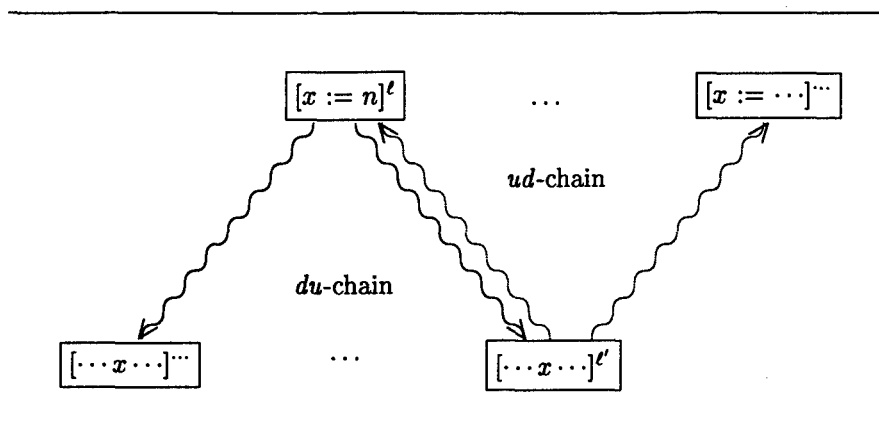
We refer to the Concluding Remarks of Chapter 6 for a discussion of systems implementing data flow analysers.

Mini Projects

Mini Project 2.1 *ud*- and *du*-chains

The aim of this mini project is to develop a more thorough understanding of the concepts of *ud*- and *du*-chains introduced in Subsection 2.1.5.

1. The function *ud* is specified in terms of definition clear paths, whilst UD re-uses functions introduced for the Reaching Definitions and Live Variables Analyses. Prove that the two functions compute the same information.

Figure 2.26: *du*- and *ud*-chains.

2. DU can be defined by analogy with UD. Starting from the definition of *du*, develop an equational definition of DU and verify its correctness.
3. A *Constant Propagation Analysis* is presented in Subsection 2.3.3; an alternative approach would be to use *du*- and *ud*-chains. Suppose there is a block $[x := n]^l$ that assigns a constant n to a variable x . By following the *du*-chain it is possible to find all blocks using the variable x . It is only safe to replace a use of x by the constant n in a block l' if all other definitions that reach l' also assign the same constant n to x . This can be determined by using the *ud*-chain. This is illustrated in Figure 2.26. Considering the program of Example 2.12, *Constant Folding* (followed by *Dead Code Elimination*) can be used to produce the following program:

```
(if [z=3]3 then [z:=0]4 else [z:=3]5); [y:=3]6; [x:=3+z]7
```

Develop a formal description of this analysis. ■

Mini Project 2.2 Correctness of Reaching Definitions

The aim of this mini project is to prove the correctness of Reaching Definitions with respect to the notion of semantic reaching definitions introduced in Section 1.5. To get a precise definition of the set of traces of interest we shall begin by introducing a so-called *instrumented semantics*: an extension of a more traditional semantics that keeps track of additional information that is mainly of interest for the program analysis.

The instrumented semantics has transitions of the forms:

$$\langle S, \sigma, tr \rangle \rightarrow \langle \sigma', tr' \rangle \text{ and } \langle S, \sigma, tr \rangle \rightarrow \langle S', \sigma', tr' \rangle$$

$[ass]$	$\langle [x := a]^{\ell}, \sigma, tr \rangle \rightarrow \langle \sigma[x \mapsto \mathcal{A}[a]\sigma], tr : (x, \ell) \rangle$
$[skip]$	$\langle [skip]^{\ell}, \sigma, tr \rangle \rightarrow \langle \sigma, tr \rangle$
$[seq_1]$	$\frac{\langle S_1, \sigma, tr \rangle \rightarrow \langle S'_1, \sigma', tr' \rangle}{\langle S_1; S_2, \sigma, tr \rangle \rightarrow \langle S'_1; S_2, \sigma', tr' \rangle}$
$[seq_2]$	$\frac{\langle S_1, \sigma, tr \rangle \rightarrow \langle \sigma', tr' \rangle}{\langle S_1; S_2, \sigma, tr \rangle \rightarrow \langle S_2, \sigma', tr' \rangle}$
$[if_1]$	$\langle \text{if } [b]^{\ell} \text{ then } S_1 \text{ else } S_2, \sigma, tr \rangle \rightarrow \langle S_1, \sigma, tr \rangle \quad \text{if } \mathcal{B}[b]\sigma = \text{true}$
$[if_2]$	$\langle \text{if } [b]^{\ell} \text{ then } S_1 \text{ else } S_2, \sigma, tr \rangle \rightarrow \langle S_2, \sigma, tr \rangle \quad \text{if } \mathcal{B}[b]\sigma = \text{false}$
$[wh_1]$	$\langle \text{while } [b]^{\ell} \text{ do } S, \sigma, tr \rangle \rightarrow \langle (S; \text{while } [b]^{\ell} \text{ do } S), \sigma, tr \rangle$ if $\mathcal{B}[b]\sigma = \text{true}$
$[wh_2]$	$\langle \text{while } [b]^{\ell} \text{ do } S, \sigma, tr \rangle \rightarrow \langle \sigma, tr \rangle$ if $\mathcal{B}[b]\sigma = \text{false}$

Table 2.10: The instrumented semantics of WHILE.

All configurations include a *trace* $tr \in \mathbf{Trace} = (\mathbf{Var} \times \mathbf{Lab})^*$ that records the elementary block in which a variable is being assigned. The detailed definition of the instrumented semantics is given in Table 2.10.

Given a program S_* and an initial state $\sigma_* \in \mathbf{State}$ it is natural to construct the trace

$$tr_* = ((x_1, ?), \dots, (x_n, ?))$$

where x_1, \dots, x_n are the variables in S_* and to consider the finite derivation sequence:

$$\langle S_*, \sigma_*, tr_* \rangle \rightarrow^* \langle \sigma', tr' \rangle$$

Intuitively, there should be a similar derivation sequence $\langle S_*, \sigma_* \rangle \rightarrow^* \sigma'$ in the Structural Operational Semantics. Similar remarks apply to infinite derivation sequences.

As in Section 2.2 we shall study the constraint system $\mathbf{RD}^{\subseteq}(S_*)$ corresponding to the equation system $\mathbf{RD}^{\supseteq}(S_*)$. Let *reach* be a collection of functions:

$$reach_{entry}, reach_{exit} : \mathbf{Lab}_* \rightarrow \mathcal{P}(\mathbf{Var}_* \times \mathbf{Lab}_*)$$

We say that *reach* solves $\mathbf{RD}^{\subseteq}(S)$, and write

$$reach \models \mathbf{RD}^{\subseteq}(S)$$

if the functions satisfy the constraints; similarly for $reach \models \mathbf{RD}^{\supseteq}(S)$.

1. Formulate and prove results corresponding to Lemmas 2.15, 2.16 and 2.18.

The correctness relation \sim will relate traces $tr \in \mathbf{Trace}$ to the information obtained by the analysis. Let $Y \subseteq \mathcal{P}(\mathbf{Var}_* \times \mathbf{Lab}_*)$ and define

$$tr \sim Y \text{ iff } \forall x \in \mathbf{Var}_* : (x, \text{SRD}(tr)(x)) \in Y$$

meaning that Y contains at least the semantically reaching definitions obtained from the trace tr by the function SRD introduced in Section 1.5.

2. Formulate and prove results corresponding to Lemma 2.20, Theorem 2.21 and Corollary 2.22. ■

Mini Project 2.3 A Prototype Implementation

In this mini project we shall implement one of the program analyses considered in Section 2.1. As implementation language we shall choose a functional language such as Standard ML or Haskell. We can then define a suitable data type for WHILE programs as follows:

```

type var      = string
type label    = int
datatype aexp = Var of var | Const of int
              | Op of string * aexp * aexp
and bexp      = True | False
              | Not of bexp | Boolop of string * bexp * bexp
              | Relop of string * aexp * aexp
datatype stat = Assign of var * aexp * label | Skip of label
              | Seq of stat * stat | If of bexp * label * stat * stat
              | While of bexp * label * stat

```

Now proceed as follows:

1. Implement the operations *init*, *final*, *flow*, *flow^R* and *blocks*.
2. Generate the data flow equations for the Live Variables Analysis of Subsection 2.1.4.
3. Solve the data flow equations; the function should be based on the algorithm of Section 2.4.

For the more ambitious: generalise your program to accept an instance of a Monotone Framework as input. ■

Exercises

Exercise 2.1 Formulate data flow equations for the Reaching Definitions Analysis of the program studied in Example 1.1 of Chapter 1 and in particular define the appropriate *gen* and *kill* functions. ■

Exercise 2.2 Consider the following program:

$$[x:=1]^1; (\text{while } [y>0]^2 \text{ do } [x:=x-1]^3); [x:=2]^4$$

Perform a Live Variables Analysis for this program using the equations of Section 2.1.4. ■

Exercise 2.3 A modification of the Available Expressions Analysis detects when an expression is available *in a particular variable*: a non-trivial expression a is available in x at a label ℓ if it has been evaluated and assigned to x on all paths leading to ℓ and if the values of x and the variables in the expression have not changed since then. Write down the data flow equations and any auxiliary functions for this analysis. ■

Exercise 2.4 Consider the following program:

$$[x:=1]^1; [x:=x-1]^2; [x:=2]^3$$

Clearly x is dead at the exits from 2 and 3. But x is live at the exit of 1 even though its only use is to calculate a new value for a variable that turns out to be dead. We shall say that a variable is a *faint variable* if it is dead or if it is only used to calculate new values for faint variables; otherwise it is *strongly live*. In the example x is faint at the exits from 1, 2 and 3. Define a Data Flow Analysis that detects strongly live variables. (Hint: For an assignment $[x := a]^\ell$ the definition of $f_\ell(l)$ should be by cases on whether x is in l or not.) ■

Exercise 2.5 A *basic block* is often taken to be a maximal group of statements such that all transfers to the block are to the first statement in the group and, once the block has been entered, all statements in the group are executed sequentially. In this exercise we shall consider basic blocks of the form

$$[x_1 := a_1; \dots; x_n := a_n; B]^\ell$$

where $n \geq 0$ and B is $x := a$, *skip* or *b*. Reformulate the analyses of Section 2.1 for this more general notion of basic block. ■

Exercise 2.6 Consider the analyses Available Expressions and Reaching Definitions. Which of the equations make sense for programs that do not have isolated entries (and how can this be improved)? Similarly, which of the equations for Very Busy Expressions and Live Variables make sense for programs that do not have isolated exits (and how can this be improved)? (Hint: See the beginning of Section 2.3.) ■

Exercise 2.7 Consider the correctness proof for the Live Variables Analysis in Section 2.2. Give a compositional definition of $LV^=(\dots)$ for a label consistent statement using

$$LV^=(\text{[skip]}^\ell) = \{LV_{\text{exit}}(\ell) = LV_{\text{entry}}(\ell)\}$$

as one of the clauses and observe that a similar development is possible for $LV^\subseteq(\dots)$. Give a formal definition of $\text{live} \models C$ where C is a set of equalities or inclusions as might have been produced by $LV^=(S)$ or $LV^\subseteq(S)$.

Prove that $\{\text{live} \mid \text{live} \models LV^\subseteq(S)\}$ is a *Moore family* in the sense of Appendix A (with \sqcap being \cap and determine whether or not a similar result holds for $\{\text{live} \mid \text{live} \models LV^=(S)\}$). ■

Exercise 2.8 Show that Constant Propagation is a Monotone Framework with the set \mathcal{F}_{CP} as defined in Section 2.3.3. ■

Exercise 2.9 A *Bit Vector Framework* is a special instance of a Monotone Framework where

- $L = (\mathcal{P}(D), \sqsubseteq)$ for some finite set D and where \sqsubseteq is either \subseteq or \supseteq , and
- $\mathcal{F} = \{f : \mathcal{P}(D) \rightarrow \mathcal{P}(D) \mid \exists Y_f^1, Y_f^2 \subseteq D : \forall Y \subseteq D : f(Y) = (Y \cap Y_f^1) \cup Y_f^2\}$

Show that the four classical analyses of Section 2.1 are Bit Vector Frameworks. Show that all Bit Vector Frameworks are indeed Distributive Frameworks. Devise a Distributive Framework that is not also a Bit Vector Framework. ■

Exercise 2.10 Consider the Constant Propagation Analysis of Section 2.3.3 and the program

(if $[\dots]^1$ then $[x:=-1]^2; [y:=1]^3$ else $[x:=1]^4; [y:=-1]^5$); $[z:=x*y]^6$

Show that $MFP_\bullet(6)$ differs from $MOP_\bullet(6)$. ■

Exercise 2.11 In our formulation of Monotone Frameworks we associate transfer functions with basic blocks. In a statement of the form

$$\text{if } [b]^\ell \text{ then } S_1 \text{ else } S_2$$

this prevents us from using the result of the test to pass different information to S_1 and S_2 ; as an example suppose that x is known to be positive or negative and that b is the test $x > 0$, then x is always positive at the entry to S_1 and always negative at the entry to S_2 . To remedy this deficiency consider writing $[b]^\ell$ as $[b]^{\ell_1, \ell_2}$ where ℓ_1 corresponds to b evaluating to true and ℓ_2 corresponds to b evaluating to false. Make the necessary changes to the development in Sections 2.1 and 2.3. (Begin by considering forward analyses.) ■

Exercise 2.12 Consider one of the analyses Available Expressions, Very Busy Expressions and Live Variables Analysis and perform a complexity analysis in the manner of Example 2.30. ■

Exercise 2.13 Let F be $\text{flow}(S_*)$ and E be $\{\text{init}(S_*)\}$ for a label consistent program S_* . Show that

$$\forall \ell \in \text{Lab}_* : \text{path}_o(\ell) \neq \emptyset$$

Prove a similar result when F is $\text{flow}^R(S_*)$ and E is $\text{final}(S_*)$. ■

Exercise 2.14 In a Detection of Signs Analysis one models all negative numbers by the symbol $-$, zero by the symbol 0 , and all positive numbers by the symbol $+$. As an example, the set $\{-2, -1, 1\}$ is modelled by the set $\{-, +\}$, that is an element of the powerset $\mathcal{P}(\{-, 0, +\})$.

Let S_* be a program and Var_* be the finite set of variables in S_* . Take L to be $\text{Var}_* \rightarrow \mathcal{P}(\{-, 0, +\})$ and define an instance $(L, \mathcal{F}, F, E, \iota, f)$ of a Monotone Framework for performing Detection of Signs Analysis.

Similarly, take L' to be $\mathcal{P}(\text{Var}_* \times \{-, 0, +\})$ and define an instance $(L', \mathcal{F}', F', E', \iota', f')$ of a Monotone Framework for Detection of Signs Analysis. Is there any difference in the precision obtained by the two approaches? ■

Exercise 2.15 In the previous exercise we defined a Detection of Signs Analysis that could *not* record the interdependencies between signs of variables (e.g. that two variables x and y always will have the same sign); this is sometimes called an *independent attribute analysis*. In this exercise we shall consider a variant of the analysis that is able to record the interdependencies between signs of variables; this is sometimes called a *relational analysis*. To do so take L to be $\mathcal{P}(\text{Var}_* \rightarrow \{-, 0, +\})$ and define an instance $(L, \mathcal{F}, F, E, \iota, f)$ of a Monotone Framework for performing Detection of Signs Analysis. Construct an example showing that the result of this relational analysis may be

more informative than that of the independent attribute analysis. The distinction between independent attribute methods and relational methods is further discussed in Chapter 4. ■

Exercise 2.16 The interprocedural analysis using bounded call strings uses contexts to record the last k call sites. Reformulate the analysis for a notion of context that records the last k *distinct* call sites. Discuss whether or not this analysis is useful for distinguishing between the call of a procedure and subsequent recursive calls. ■

Exercise 2.17 Consider the Fibonacci program of Example 2.33 and the Detection of Signs Analysis of Exercise 2.15 and Example 2.36. Construct the data flow equations corresponding to using large and small assumption sets, respectively. ■

Exercise 2.18 Choose one of the four classical analyses from Section 2.1 and formulate it as an interprocedural analysis based on call strings. (Hint: Some may be easier than others.) ■

Exercise 2.19 Extend the syntax of programs to have the form

$$\text{begin } D_*; \text{ input } x; S_*; \text{ output } y \text{ end}$$

so that it maps integers to integers rather than states to states. Consider the Detection of Signs Analysis and define the transfer functions for the input and output statements. ■

Exercise 2.20 Consider extending the procedure language such that procedures can have multiple call-by-value, call-by-result and call-by-value-result parameters as well as local variables and reconsider the Detection of Signs Analysis. How should one define the transfer functions associated with procedure call, procedure entry, procedure exit, and procedure return? ■

Exercise 2.21 In the Shape Analysis of Section 2.6 work out direct definitions of the transfer functions for elementary statements of the forms $[x:=x.sel]^t$, $[x.sel:=x]^t$, $[x.sel:=x.sel']^t$ and $[\text{malloc}(x.sel)]^t$. ■

Exercise 2.22 Consider Case 3 in the definition of the transfer function for $[x:=y.sel]^t$ (where $x \neq y$) in the Shape Analysis. Make a careful analysis of internal, going-in and going-out edges and determine whether or not some of the shape graphs (S'', H'', is'') in $\phi_t^{SA}((S, H, is))$ can be removed by placing stronger demands on the edges in H'' compared to those in H' (where $(S', H', is') = \text{kill}_x((S, H, is))$). ■

Exercise* 2.23 The Shape Analysis as presented in Section 2.6 does not take garbage collection into account. Modify the Structural Operational Semantics of the pointer language to perform garbage collection and subsequently modify the analysis to reflect this. ■

Exercise* 2.24 The use of a single abstract summary location leads to a certain amount of inaccuracy in the Shape Analysis. A more accurate analysis could associate allocation sites with the abstract locations. An abstract location would then have the form $n_{\ell, X}$ where ℓ is an allocation site (a label of a `malloc`-statement) and X is a set of variables as before. Develop the transfer functions for the new analysis. ■

Chapter 3

Control Flow Analysis

In this chapter we present the technique of Control Flow Analysis using a simple functional language, FUN. We begin by presenting an abstract specification of the analysis and then study its theoretical properties: it is correct with respect to a Structural Operational Semantics and it can be used to analyse all programs. This specification of the analysis does not immediately lend itself to an efficient algorithm for computing a solution so we proceed by developing first a syntax directed specification and then a constraint based formulation and finally we show how the constraints can be solved. We conclude by illustrating how the precision of the analysis can be improved by combining it with Data Flow Analysis and by incorporating context information thereby linking up with the development of the previous chapter.

3.1 Abstract 0-CFA Analysis

In Chapter 2 we saw how properties of data could be propagated through a program. In developing the specification we relied on the ability to identify for each program fragment all the possible successor (and predecessor) fragments – via the operator *flow* (and $flow^R$) and the interprocedural flow *inter-flow*_{*} (and $inter-flow^R$ _{*}). The usefulness of the resulting specification was due to the number of successors and predecessors being small (usually just one or two except for procedure exits). This is a typical feature of imperative programs without procedures but it usually fails for more general languages, whether imperative languages with procedures as parameters, functional languages, or object-oriented languages. In particular, the interprocedural techniques of Section 2.5 provide a solution for the simpler cases where the program text allows one to limit the number of successors, as is the case when a proce-

ture call is performed by explicitly mentioning the name of the procedure. However, these techniques are not powerful enough to handle the *dynamic dispatch problem* where variables can denote procedures. In Section 1.4 we illustrated this by the functional program

```
let f = fn x => x 1;
    g = fn y => y+2;
    h = fn z => z+3
in (f g) + (f h)
```

where the function application $x\ 1$ in the body of f will transfer control to the body of the function x , and here it is not so obvious what program fragment this actually is, since x is the formal parameter of f . The Control Flow Analysis of the present chapter will provide a solution to the dynamic dispatch problem by determining for each subexpression a hopefully small number of functions that it may evaluate to; thereby it will determine where the flow of control may be transferred to in the case where the subexpression is the operator of a function application. In short, Control Flow Analysis will determine the *interprocedural flow* information (*inter-flow** or *IF*) upon which the development of Section 2.5 is based.

Syntax of the FUN language. For the main part of this chapter we shall concentrate on a small functional language: the untyped lambda calculus extended with explicit operators for recursion, conditional and local definitions. The purpose of the Control Flow Analysis will be to compute for *each* subexpression the set of functions that it could evaluate to, and to express this it is important that we are able to label *all* program fragments. We shall be very explicit about this: a program fragment with a label is called an *expression* whereas a program fragment without a label is called a *term*. So we use the following syntactic categories:

$$\begin{aligned} e &\in \mathbf{Exp} && \text{expressions (or labelled terms)} \\ t &\in \mathbf{Term} && \text{terms (or unlabelled expressions)} \end{aligned}$$

We assume that a countable set of variables is given and that constants (including the truth values), binary operators (including the usual arithmetic, boolean and relational operators) and labels are left unspecified:

$$\begin{aligned} f, x &\in \mathbf{Var} && \text{variables} \\ c &\in \mathbf{Const} && \text{constants} \\ op &\in \mathbf{Op} && \text{binary operators} \\ \ell &\in \mathbf{Lab} && \text{labels} \end{aligned}$$

The *abstract syntax* of the language is now given by:

$$\begin{aligned} e &::= t \\ t &::= c \mid x \mid \text{fn } x \Rightarrow e_0 \mid \text{fun } f\ x \Rightarrow e_0 \mid e_1\ e_2 \\ &\quad \mid \text{if } e_0 \text{ then } e_1 \text{ else } e_2 \mid \text{let } x = e_1 \text{ in } e_2 \mid e_1\ op\ e_2 \end{aligned}$$

Here $\text{fn } x \Rightarrow e_0$ is a function definition (or function abstraction) whereas $\text{fun } f x \Rightarrow e_0$ is a recursive variant of $\text{fn } x \Rightarrow e_0$ where all free occurrences of f in e_0 refer to $\text{fun } f x \Rightarrow e_0$ itself. The construct $\text{let } x = e_1 \text{ in } e_2$ is a non-recursive local definition that semantically is equivalent to $(\text{fn } x \Rightarrow e_2)(e_1)$. As usual we shall use parentheses to disambiguate the parsing whenever needed. Also we shall *assume* throughout that in all occurrences of $\text{fun } f x \Rightarrow e_0$, f and x are *distinct* variables.

We shall need the notion of *free variables* of expressions and terms so we define the function

$$FV: (\text{Term} \cup \text{Exp}) \rightarrow \mathcal{P}(\text{Var})$$

in the following standard way. The abstractions $\text{fn } x \Rightarrow e_0$ and $\text{fun } f x \Rightarrow e_0$ contain binding occurrences of variables so $FV(\text{fn } x \Rightarrow e_0) = FV(e_0) \setminus \{x\}$ and similarly $FV(\text{fun } f x \Rightarrow e_0) = FV(e_0) \setminus \{f, x\}$. Similarly, $\text{let } x = e_1 \text{ in } e_2$ contains a binding occurrence of x so we have $FV(\text{let } x = e_1 \text{ in } e_2) = FV(e_1) \cup (FV(e_2) \setminus \{x\})$ reflecting that free occurrences of x in e_1 are bound outside the construct. The remaining clauses for FV are straightforward.

Example 3.1 The functional program $(\text{fn } x \Rightarrow x) (\text{fn } y \Rightarrow y)$ considered in Section 1.4 is now written as:

$$((\text{fn } x \Rightarrow x^1)^2 (\text{fn } y \Rightarrow y^3)^4)^5$$

Compared with the notation of Example 1.2 we have omitted the square brackets. ■

Example 3.2 Consider the following expression, loop, of FUN:

$$\begin{aligned} &(\text{let } g = (\text{fun } f x \Rightarrow (f^1 (\text{fn } y \Rightarrow y^2)^3)^4)^5 \\ &\text{in } (g^6 (\text{fn } z \Rightarrow z^7)^8)^9)^{10} \end{aligned}$$

It defines a function g that is applied to the identity function $\text{fn } z \Rightarrow z^7$. The function g is defined recursively: f is its local name and x is the formal parameter. Hence the function will ignore its actual parameter and call itself recursively with the argument $\text{fn } y \Rightarrow y^2$. This will happen again and again so the program loops. ■

3.1.1 The Analysis

Abstract domains. We shall now show how to specify 0-CFA analyses. These may be regarded as the simplest possible form of Control Flow Analysis in that *no* context information is taken into account. As will become clear in Section 3.6, this is what the number 0 is indicating.

The result of a 0-CFA analysis is a pair $(\hat{C}, \hat{\rho})$ where:

- \widehat{C} is the *abstract cache* associating abstract values with each labelled program point.
- $\widehat{\rho}$ is the *abstract environment* associating abstract values with each variable.

This is made precise by:

$$\begin{aligned} \widehat{v} \in \widehat{\mathbf{Val}} &= \mathcal{P}(\mathbf{Term}) && \text{abstract values} \\ \widehat{\rho} \in \widehat{\mathbf{Env}} &= \mathbf{Var} \rightarrow \widehat{\mathbf{Val}} && \text{abstract environments} \\ \widehat{C} \in \widehat{\mathbf{Cache}} &= \mathbf{Lab} \rightarrow \widehat{\mathbf{Val}} && \text{abstract caches} \end{aligned}$$

Here an *abstract value* \widehat{v} is an abstraction of a set of functions: it is a set of terms of the form $\text{fn } x \Rightarrow e_0$ or $\text{fun } f \ x \Rightarrow e_0$. We will not be recording any constants in the abstract values because the analysis we shall specify is a pure Control Flow Analysis with no Data Flow Analysis component; in Section 3.5 we shall show how to extend it with Data Flow Analysis components. Furthermore, we do *not* need to assume that all bound variables are distinct and that all labels are distinct, but clearly, greater precision is achieved if this is the case. As we shall see an *abstract environment* is an abstraction of a set of environments occurring in closures at run-time (see the semantics in Subsection 3.2.1). In a similar way an *abstract cache* might be considered as an abstraction of a set of execution profiles: as discussed below some texts prefer to combine the abstract environment with the abstract cache.

Example 3.3 Consider the expression $((\text{fn } x \Rightarrow x^1)^2 (\text{fn } y \Rightarrow y^3)^4)^5$ of Example 3.1. The following table contains three guesses of a 0-CFA analysis:

	$(\widehat{C}_e, \widehat{\rho}_e)$	$(\widehat{C}'_e, \widehat{\rho}'_e)$	$(\widehat{C}''_e, \widehat{\rho}''_e)$
1	$\{\text{fn } y \Rightarrow y^3\}$	$\{\text{fn } y \Rightarrow y^3\}$	$\{\text{fn } x \Rightarrow x^1, \text{fn } y \Rightarrow y^3\}$
2	$\{\text{fn } x \Rightarrow x^1\}$	$\{\text{fn } x \Rightarrow x^1\}$	$\{\text{fn } x \Rightarrow x^1, \text{fn } y \Rightarrow y^3\}$
3	\emptyset	\emptyset	$\{\text{fn } x \Rightarrow x^1, \text{fn } y \Rightarrow y^3\}$
4	$\{\text{fn } y \Rightarrow y^3\}$	$\{\text{fn } y \Rightarrow y^3\}$	$\{\text{fn } x \Rightarrow x^1, \text{fn } y \Rightarrow y^3\}$
5	$\{\text{fn } y \Rightarrow y^3\}$	$\{\text{fn } y \Rightarrow y^3\}$	$\{\text{fn } x \Rightarrow x^1, \text{fn } y \Rightarrow y^3\}$
x	$\{\text{fn } y \Rightarrow y^3\}$	\emptyset	$\{\text{fn } x \Rightarrow x^1, \text{fn } y \Rightarrow y^3\}$
y	\emptyset	\emptyset	$\{\text{fn } x \Rightarrow x^1, \text{fn } y \Rightarrow y^3\}$

Intuitively, the guess $(\widehat{C}_e, \widehat{\rho}_e)$ of the first column is acceptable whereas the guess $(\widehat{C}'_e, \widehat{\rho}'_e)$ of the second column is wrong: it would seem that $\text{fn } x \Rightarrow x^1$ is never called since $\widehat{\rho}'_e(x) = \emptyset$ indicates that x will never be bound to any closures. Also the guess $(\widehat{C}''_e, \widehat{\rho}''_e)$ of the third column would seem to be acceptable although clearly more imprecise than $(\widehat{C}_e, \widehat{\rho}_e)$. ■

Example 3.4 Let us consider the expression, loop, of Example 3.2 and introduce the following abbreviations for abstract values:

$$\begin{aligned} f &= \text{fun } f \ x \Rightarrow (f^1 (fn \ y \Rightarrow y^2)^3)^4 \\ id_y &= fn \ y \Rightarrow y^2 \\ id_z &= fn \ z \Rightarrow z^7 \end{aligned}$$

One guess of a 0-CFA analysis for this program is $(\widehat{C}_{lp}, \widehat{\rho}_{lp})$ defined by:

$$\begin{array}{lll} \widehat{C}_{lp}(1) &= \{f\} & \widehat{C}_{lp}(6) &= \{f\} & \widehat{\rho}_{lp}(f) &= \{f\} \\ \widehat{C}_{lp}(2) &= \emptyset & \widehat{C}_{lp}(7) &= \emptyset & \widehat{\rho}_{lp}(g) &= \{f\} \\ \widehat{C}_{lp}(3) &= \{id_y\} & \widehat{C}_{lp}(8) &= \{id_z\} & \widehat{\rho}_{lp}(x) &= \{id_y, id_z\} \\ \widehat{C}_{lp}(4) &= \emptyset & \widehat{C}_{lp}(9) &= \emptyset & \widehat{\rho}_{lp}(y) &= \emptyset \\ \widehat{C}_{lp}(5) &= \{f\} & \widehat{C}_{lp}(10) &= \emptyset & \widehat{\rho}_{lp}(z) &= \emptyset \end{array}$$

Intuitively, this is an acceptable guess. The choice of $\widehat{\rho}_{lp}(g) = \{f\}$ reflects that g will evaluate to a closure constructed from that abstraction. The choice of $\widehat{\rho}_{lp}(x) = \{id_y, id_z\}$ reflects that x will be bound to closures constructed from both abstractions in the course of the evaluation. The choice of $\widehat{C}_{lp}(10) = \emptyset$ reflects that the evaluation of the expression will never terminate. ■

We have already said that Control Flow Analysis computes the interprocedural flow information used in Section 2.5. It is also instructive to point out the similarity between Control Flow Analysis and *Definition-Use chains* (*du-chains*) for imperative languages (see Subsection 2.1.5): in both cases we attempt to trace how definition points reach points of use. In the case of Control Flow Analysis the *definition points* are the points where the function abstractions are created, and the *use points* are the points where functions are applied; in the case of Definition-Use chains the *definition points* are the points where variables are assigned a value, and the *use points* are the points where values of variables are accessed.

Remark. Clearly an abstract cache $\widehat{C} : \mathbf{Lab} \rightarrow \widehat{\mathbf{Val}}$ and an abstract environment $\widehat{\rho} : \mathbf{Var} \rightarrow \widehat{\mathbf{Val}}$ can be combined into an entity of type $(\mathbf{Var} \cup \mathbf{Lab}) \rightarrow \widehat{\mathbf{Val}}$. Some texts dispense with the labels altogether, simply using an abstract environment and no abstract cache, by ensuring that *all* subterms are properly “labelled” by variables. This type of expression frequently occurs in the internals of compilers in the form of “continuation passing style”, “A-normal form” or “three address code”. We have abstained from doing so in order to illustrate that the techniques not only work for compiler intermediate forms but also for general programming languages and calculi for computation; this flexibility is useful when dealing with non-standard applications (as discussed in the Concluding Remarks). ■

[con]	$(\widehat{C}, \widehat{\rho}) \models c^l$ always
[var]	$(\widehat{C}, \widehat{\rho}) \models x^l$ iff $\widehat{\rho}(x) \subseteq \widehat{C}(l)$
[fn]	$(\widehat{C}, \widehat{\rho}) \models (\text{fn } x \Rightarrow e_0)^l$ iff $\{\text{fn } x \Rightarrow e_0\} \subseteq \widehat{C}(l)$
[fun]	$(\widehat{C}, \widehat{\rho}) \models (\text{fun } f \ x \Rightarrow e_0)^l$ iff $\{\text{fun } f \ x \Rightarrow e_0\} \subseteq \widehat{C}(l)$
[app]	$(\widehat{C}, \widehat{\rho}) \models (t_1^{l_1} \ t_2^{l_2})^l$ iff $(\widehat{C}, \widehat{\rho}) \models t_1^{l_1} \wedge (\widehat{C}, \widehat{\rho}) \models t_2^{l_2} \wedge$ $(\forall (\text{fn } x \Rightarrow t_0^{l_0}) \in \widehat{C}(l_1) :$ $\quad (\widehat{C}, \widehat{\rho}) \models t_0^{l_0} \wedge$ $\quad \widehat{C}(l_2) \subseteq \widehat{\rho}(x) \wedge \widehat{C}(l_0) \subseteq \widehat{C}(l) \wedge$ $(\forall (\text{fun } f \ x \Rightarrow t_0^{l_0}) \in \widehat{C}(l_1) :$ $\quad (\widehat{C}, \widehat{\rho}) \models t_0^{l_0} \wedge$ $\quad \widehat{C}(l_2) \subseteq \widehat{\rho}(x) \wedge \widehat{C}(l_0) \subseteq \widehat{C}(l) \wedge$ $\quad \{\text{fun } f \ x \Rightarrow t_0^{l_0}\} \subseteq \widehat{\rho}(f))$
[if]	$(\widehat{C}, \widehat{\rho}) \models (\text{if } t_0^{l_0} \text{ then } t_1^{l_1} \text{ else } t_2^{l_2})^l$ iff $(\widehat{C}, \widehat{\rho}) \models t_0^{l_0} \wedge$ $(\widehat{C}, \widehat{\rho}) \models t_1^{l_1} \wedge (\widehat{C}, \widehat{\rho}) \models t_2^{l_2} \wedge$ $\widehat{C}(l_1) \subseteq \widehat{C}(l) \wedge \widehat{C}(l_2) \subseteq \widehat{C}(l)$
[let]	$(\widehat{C}, \widehat{\rho}) \models (\text{let } x = t_1^{l_1} \text{ in } t_2^{l_2})^l$ iff $(\widehat{C}, \widehat{\rho}) \models t_1^{l_1} \wedge (\widehat{C}, \widehat{\rho}) \models t_2^{l_2} \wedge$ $\widehat{C}(l_1) \subseteq \widehat{\rho}(x) \wedge \widehat{C}(l_2) \subseteq \widehat{C}(l)$
[op]	$(\widehat{C}, \widehat{\rho}) \models (t_1^{l_1} \text{ op } t_2^{l_2})^l$ iff $(\widehat{C}, \widehat{\rho}) \models t_1^{l_1} \wedge (\widehat{C}, \widehat{\rho}) \models t_2^{l_2}$

Table 3.1: Abstract Control Flow Analysis.

Acceptability relation. It remains to determine whether or not a proposed guess $(\widehat{C}, \widehat{\rho})$ of an analysis results is in fact an *acceptable 0-CFA* analysis for the program considered. We shall give an *abstract specification* of what this means; having studied its theoretical properties (in Section 3.2) we then consider how to compute the desired analysis (in Sections 3.3 and 3.4).

It is instructive to point out that the abstract specification corresponds to an *implicit* formulation of the data flow equations of Chapter 2; it will be used to determine whether or not a guess is indeed an acceptable solution to the analysis problem. The syntax directed and constraint based formulations (of Sections 3.3 and 3.4) correspond to *explicit* formulations of the data flow equations from which an iterative algorithm in the spirit of Chaotic Iteration

(Section 1.7) can be used to compute an analysis result.

For the formulation of the *abstract 0-CFA* analysis we shall write

$$(\widehat{C}, \widehat{\rho}) \models e$$

for when $(\widehat{C}, \widehat{\rho})$ is an acceptable Control Flow Analysis of the expression e . Thus the relation “ \models ” has functionality:

$$\models : (\widehat{\text{Cache}} \times \widehat{\text{Env}} \times \text{Exp}) \rightarrow \{\text{true}, \text{false}\}$$

and its defining clauses are given in Table 3.1 (writing “always” for “iff true”); they are explained below.

The clause *[con]* places no demands on $\widehat{C}(\ell)$ because we are not tracking any data values in the pure 0-CFA analysis considered here and because we assume that there are no functions among the constants; the clause can be reformulated as

$$(\widehat{C}, \widehat{\rho}) \models c^\ell \text{ iff } \emptyset \subseteq \widehat{C}(\ell)$$

thereby highlighting this point.

The clause *[var]* is responsible for linking the abstract environment into the abstract cache: so in order for $(\widehat{C}, \widehat{\rho})$ to be an acceptable analysis, everything the variable x can evaluate to has to be included in what may be observed at the program point ℓ : $\widehat{\rho}(x) \subseteq \widehat{C}(\ell)$.

The clauses *[fn]* and *[fun]* simply demand that in order for $(\widehat{C}, \widehat{\rho})$ to be an acceptable analysis, the functional term ($\text{fn } x \Rightarrow e_0$ or $\text{fun } f \ x \Rightarrow e_0$) must be included in $\widehat{C}(\ell)$; this says that the term is part of a closure that can arise at program point ℓ during evaluation. Note that these clauses do not demand that $(\widehat{C}, \widehat{\rho})$ is an acceptable analysis result for the bodies of the functions; the clause for function application will take care of that.

Before turning to the more complicated clause *[app]* let us consider the clauses *[if]* and *[let]*. They contain “recursive calls” demanding that subexpressions must be analysed in consistent ways using $(\widehat{C}, \widehat{\rho})$; additionally, the clauses explicitly link the values produced by subexpressions to the value of the overall expression, and in the case of *[let]* also the abstract cache is linked into the abstract environment. The interplay between the clauses *[var]* and *[let]* is illustrated in Figure 3.1; as in Chapter 2 an arrow indicates a flow of information. The clause *[op]* follows the same overall pattern.

Clause *[app]* also contains “recursive calls” demanding that the operator $t_1^{\ell_1}$ and the operand $t_2^{\ell_2}$ can be analysed using $(\widehat{C}, \widehat{\rho})$. For each term $\text{fn } x \Rightarrow t_0^{\ell_0}$ that may reach the operator position (ℓ_1) , i.e. where

$$(\text{fn } x \Rightarrow t_0^{\ell_0}) \in \widehat{C}(\ell_1)$$

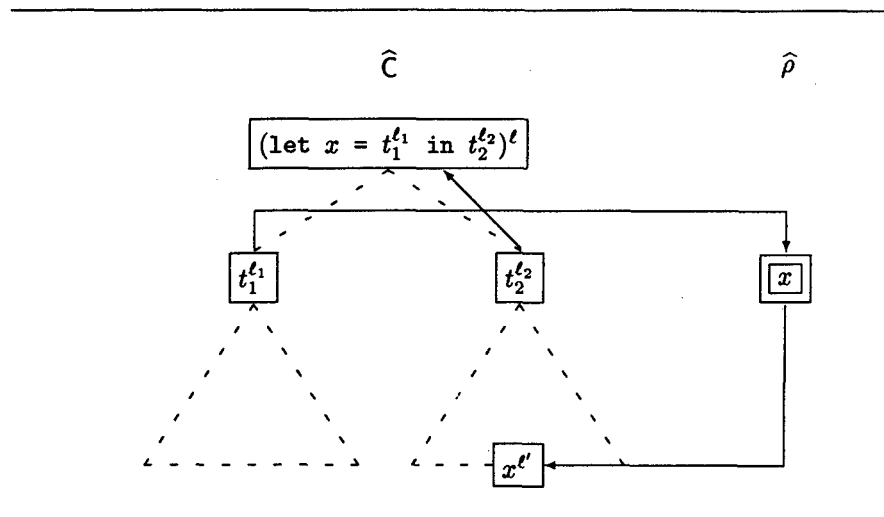


Figure 3.1: Pictorial illustration of the clauses $[let]$ and $[var]$.

it further demands that the actual parameter (labelled ℓ_2) is linked to the formal parameter (x)

$$\widehat{C}(\ell_2) \subseteq \widehat{\rho}(x)$$

and that the result of the function evaluation (labelled ℓ_0) is linked to the result of the application itself (labelled ℓ)

$$\widehat{C}(\ell_0) \subseteq \widehat{C}(\ell)$$

and finally, that the function body itself can be analysed using $(\widehat{C}, \widehat{\rho})$:

$$(\widehat{C}, \widehat{\rho}) \models t_0^{\ell_0}$$

This is illustrated in Figure 3.2. For terms $\text{fun } f \ x \Rightarrow t_0^{\ell_0}$ the demands are much the same except that the term itself additionally needs to be included in $\widehat{\rho}(f)$ in order to reflect the recursive nature of $\text{fun } f \ x \Rightarrow t_0^{\ell_0}$.

Example 3.5 Consider Example 3.3 and the guesses of a 0-CFA analysis for the expression. First we show that $(\widehat{C}_e, \widehat{\rho}_e)$ is an acceptable guess:

$$(\widehat{C}_e, \widehat{\rho}_e) \models ((\text{fn } x \Rightarrow x^1)^2 (\text{fn } y \Rightarrow y^3)^4)^5$$

Using clause $[app]$ and $\widehat{C}_e(2) = \{\text{fn } x \Rightarrow x^1\}$ we must check:

$$(\widehat{C}_e, \widehat{\rho}_e) \models (\text{fn } x \Rightarrow x^1)^2$$

$$(\widehat{C}_e, \widehat{\rho}_e) \models (\text{fn } y \Rightarrow y^3)^4$$




Figure 3.2: Pictorial illustration of the clauses $[app]$, $[fn]$ and $[var]$.

$$\begin{aligned}
 (\widehat{C}_e, \widehat{\rho}_e) &\models x^1 \\
 \widehat{C}_e(4) &\subseteq \widehat{\rho}_e(x) \\
 \widehat{C}_e(1) &\subseteq \widehat{C}_e(5)
 \end{aligned}$$

All of these are easily checked using the clauses $[fn]$ and $[var]$.

Next we show that $(\widehat{C}'_e, \widehat{\rho}'_e)$ is *not* an acceptable guess:

$$(\widehat{C}'_e, \widehat{\rho}'_e) \not\models ((fn\ x \Rightarrow x^1)^2 (fn\ y \Rightarrow y^3)^4)^5$$

We do so by proceeding as above and observing that $\widehat{C}'_e(4) \not\subseteq \widehat{\rho}'_e(x)$. ■

Note that the clauses contain a number of inclusions of the form

$$lhs \subseteq rhs$$

where rhs is of the form $\widehat{C}(\ell)$ or $\widehat{\rho}(x)$ and where lhs is of the form $\widehat{C}(\ell)$, $\widehat{\rho}(x)$, or $\{t\}$. These inclusions express how the higher-order entities may flow through the expression.

It is important to observe that the clauses $[fn]$ and $[fun]$ do *not* contain “recursive calls” demanding that subexpressions must be analysed. Instead one relies on the clause $[app]$ demanding this for all “subexpressions” that may eventually be applied. This is a phenomenon common in program analysis, where one does *not* want to analyse *unreachable* program fragments:

occasionally results obtained from these parts of the program can suppress transformations in the reachable part of the program. It also allows us to deal with *open systems* where functions may be supplied by the environment; this is particularly important for languages involving concurrency. However, note that this perspective is different from that of type inference, where even unreachable fragments must be correctly typed.

In the terminology of Section 2.5 the analysis is *flow-insensitive* because FUN contains no side effects and because we analyse the operand to a function call even when the operator cannot evaluate to any function; see Exercise 3.3 for how to improve on this. Also the analysis is *context-insensitive* because it treats all function calls in the same way; we refer to Section 3.6 for how to improve on this.

3.1.2 Well-definedness of the Analysis

Finally, we need to clarify that the clauses of Table 3.1 do indeed define a relation. The difficulty here is that the clause [app] is *not* in a form that allows us to define $(\widehat{C}, \widehat{\rho}) \models e$ by structural induction in the expression e – it requires checking the acceptability of $(\widehat{C}, \widehat{\rho})$ for an expression $t_0^{\ell_0}$ that is not a subexpression of the application $(t_1^{\ell_1} t_2^{\ell_2})^{\ell}$. This leads to defining the relation “ \models ” of Table 3.1 by *coinduction*, that is as the *greatest fixed point* of a certain functional. An alternative will be to define the analysis as the least fixed point of the functional but, as we shall see in Example 3.6 and more formally in Proposition 3.16, this may lead to problems.

The functional Q . Following the approach of Appendix B we shall view Table 3.1 as defining a function:

$$Q : ((\widehat{\text{Cache}} \times \widehat{\text{Env}} \times \text{Exp}) \rightarrow \{\text{true}, \text{false}\}) \\ \rightarrow ((\widehat{\text{Cache}} \times \widehat{\text{Env}} \times \text{Exp}) \rightarrow \{\text{true}, \text{false}\})$$

As an example we have:

$$Q(Q)(\widehat{C}, \widehat{\rho}, (\text{let } x = t_1^{\ell_1} \text{ in } t_2^{\ell_2})^{\ell}) \\ = Q(\widehat{C}, \widehat{\rho}, t_1^{\ell_1}) \wedge Q(\widehat{C}, \widehat{\rho}, t_2^{\ell_2}) \wedge \widehat{C}(\ell_1) \subseteq \widehat{\rho}(x) \wedge \widehat{C}(\ell_2) \subseteq \widehat{C}(\ell)$$

Now by inspecting Table 3.1 it is easy to verify that the function Q constructed this way is a monotone function on the complete lattice

$$((\widehat{\text{Cache}} \times \widehat{\text{Env}} \times \text{Exp}) \rightarrow \{\text{true}, \text{false}\}, \sqsubseteq)$$

where the ordering \sqsubseteq is defined by:

$$Q_1 \sqsubseteq Q_2 \text{ iff } \forall (\widehat{C}, \widehat{\rho}, e) : (Q_1(\widehat{C}, \widehat{\rho}, e) = \text{true}) \Rightarrow (Q_2(\widehat{C}, \widehat{\rho}, e) = \text{true})$$

Hence Q has fixed points and we shall define “ \models ” coinductively:

\models is the *greatest* fixed point of \mathcal{Q}

The following example motivates the use of a coinductive (i.e. greatest fixed point) definition as opposed to an inductive (i.e. least fixed point) definition.

Example 3.6 Consider the expression loop of Example 3.4

$$\begin{aligned} &(\text{let } g = (\text{fun } f \ x \Rightarrow (f^1 (fn \ y \Rightarrow y^2)^3)^4)^5 \\ &\text{in } (g^6 (fn \ z \Rightarrow z^7)^8)^9)^{10} \end{aligned}$$

and the suggested analysis result $(\widehat{C}_{lp}, \widehat{\rho}_{lp})$. To show $(\widehat{C}_{lp}, \widehat{\rho}_{lp}) \models \text{loop}$ it is, according to the clause $[let]$, sufficient to verify that

$$\begin{aligned} (\widehat{C}_{lp}, \widehat{\rho}_{lp}) &\models (\text{fun } f \ x \Rightarrow (f^1 (fn \ y \Rightarrow y^2)^3)^4)^5 \\ (\widehat{C}_{lp}, \widehat{\rho}_{lp}) &\models (g^6 (fn \ z \Rightarrow z^7)^8)^9 \end{aligned}$$

because $\widehat{C}_{lp}(5) \subseteq \widehat{\rho}_{lp}(g)$ and $\widehat{C}_{lp}(9) \subseteq \widehat{C}_{lp}(10)$. The first clause follows from $[fun]$ and for the second clause we use that $\widehat{C}_{lp}(6) = \{f\}$ so it is, according to $[app]$, sufficient to verify that

$$\begin{aligned} (\widehat{C}_{lp}, \widehat{\rho}_{lp}) &\models g^6 \\ (\widehat{C}_{lp}, \widehat{\rho}_{lp}) &\models (fn \ z \Rightarrow z^7)^8 \\ (\widehat{C}_{lp}, \widehat{\rho}_{lp}) &\models (f^1 (fn \ y \Rightarrow y^2)^3)^4 \end{aligned}$$

because $\widehat{C}_{lp}(8) \subseteq \widehat{\rho}_{lp}(x)$, $\widehat{C}_{lp}(4) \subseteq \widehat{C}_{lp}(9)$ and $f \in \widehat{\rho}_{lp}(f)$. The first two clauses now follow from $[var]$ and $[fn]$. For the third clause we proceed as above and since $\widehat{C}_{lp}(1) = \{f\}$ it is sufficient to verify

$$\begin{aligned} (\widehat{C}_{lp}, \widehat{\rho}_{lp}) &\models f^1 \\ (\widehat{C}_{lp}, \widehat{\rho}_{lp}) &\models (fn \ y \Rightarrow y^2)^3 \\ (\widehat{C}_{lp}, \widehat{\rho}_{lp}) &\models (f^1 (fn \ y \Rightarrow y^2)^3)^4 \end{aligned}$$

because $\widehat{C}_{lp}(3) \subseteq \widehat{\rho}_{lp}(x)$, $\widehat{C}_{lp}(4) \subseteq \widehat{C}_{lp}(4)$ and $f \in \widehat{\rho}_{lp}(f)$.

Again the first two clauses are straightforward using $[var]$ and $[fn]$ but in the last clause where we have encountered a *circularity*: to verify $(\widehat{C}_{lp}, \widehat{\rho}_{lp}) \models (f^1 (fn \ y \Rightarrow y^2)^3)^4$ we have to verify $(\widehat{C}_{lp}, \widehat{\rho}_{lp}) \models (f^1 (fn \ y \Rightarrow y^2)^3)^4$!

To solve this we use coinduction: basically this amounts to assuming that $(\widehat{C}_{lp}, \widehat{\rho}_{lp}) \models (f^1 (fn \ y \Rightarrow y^2)^3)^4$ holds at the "inner level" and proving that it also holds at the "outer level". This will give us the required proof. ■

3.2 Theoretical Properties

In this section we shall investigate some more theoretical properties of the Control Flow Analysis, namely:

- semantic correctness, and
- the existence of least solutions.

The semantic correctness result is important since it ensures that the information from the analysis is indeed a safe description of what will happen during the evaluation of the program. The result about the existence of least solutions ensures that all programs can be analysed and furthermore that there is a “best” or “most precise” analysis result.

As in Section 2.2, the material of this section may be skimmed through on a first reading; however, it is frequently when conducting the correctness proof that the final and subtle errors in the analysis are found and corrected!

3.2.1 Structural Operational Semantics

Configurations. We shall equip the language FUN with a *Structural Operational Semantics*. We shall choose an approach based on *explicit* environments rather than substitutions because (as discussed in the Concluding Remarks) a substitution based semantics does not preserve the identity of functions (and hence abstract values) during evaluation. So a function definition will evaluate to a *closure* containing the syntax of the function definition together with an *environment* mapping its free variables to their *values*. For this we introduce the following categories

$$\begin{aligned} v &\in \mathbf{Val} && \text{values} \\ \rho &\in \mathbf{Env} && \text{environments} \end{aligned}$$

defined by:

$$\begin{aligned} v &::= c \mid \text{close } t \text{ in } \rho \\ \rho &::= [] \mid \rho[x \mapsto v] \end{aligned}$$

A function abstraction $\text{fn } x \Rightarrow e_0$ will then evaluate to a closure, written $\text{close } (\text{fn } x \Rightarrow e_0) \text{ in } \rho$; similarly, the abstraction $\text{fun } f \ x \Rightarrow e_0$ will evaluate to $\text{close } (\text{fun } f \ x \Rightarrow e_0) \text{ in } \rho$. Our definitions do not demand that all terms t occurring in some $\text{close } t \text{ in } \rho$ in the semantics will be of the form $\text{fn } x \Rightarrow e_0$ or $\text{fun } f \ x \Rightarrow e_0$; however, it will be the case for the semantics presented below.

As in Section 2.5 we shall need intermediate configurations to handle the binding of local variables. We therefore introduce syntactic categories for

intermediate expressions and *intermediate terms*

$ie \in \mathbf{IExp}$ intermediate expressions
 $it \in \mathbf{ITerm}$ intermediate terms

that extend the syntax of expressions and terms as follows:

$$ie ::= it^l$$

$$it ::= c \mid x \mid \text{fn } x \Rightarrow e_0 \mid \text{fun } f \ x \Rightarrow e_0 \mid ie_1 \ ie_2$$

$$\quad \mid \text{if } ie_0 \text{ then } e_1 \text{ else } e_2 \mid \text{let } x = ie_1 \text{ in } e_2 \mid ie_1 \text{ op } ie_2$$

$$\quad \mid \text{bind } \rho \text{ in } ie \mid \text{close } t \text{ in } \rho$$

The rôle of the `bind`-construct is much as in Section 2.5: `bind ρ in ie` records that the intermediate expression ie has to be evaluated in an environment with the bindings ρ . (The sequence of environments of nested `bind`-constructs may be viewed as an encoding of the frames of a run-time stack.) So while `close t in ρ` is a fully evaluated value this is not the case for `bind ρ in ie` . We shall need these intermediate terms because we are based on a small step semantics; only the `close`-constructs will be needed for a big step variant of the semantics.

Alternatively, the definitions of `Val` and `Env` could have been written $\mathbf{Val} = \mathbf{Const} + (\mathbf{Term} \times \mathbf{Env})$ and $\mathbf{Env} = \mathbf{Var} \rightarrow_{\text{fin}} \mathbf{Val}$ (for a finite mapping) but it is important to stress that all entities are defined mutually recursively in the manner of context-free grammars. Formally, we defined an environment ρ as a list but nevertheless we shall feel free to regard it as a finite mapping: we write $\text{dom}(\rho)$ for $\{x \mid \rho \text{ contains } [x \mapsto \dots]\}$; we write $\rho(x) = v$ if $x \in \text{dom}(\rho)$ and the rightmost occurrence of $[x \mapsto \dots]$ in ρ is $[x \mapsto v]$, and we write $\rho \upharpoonright X$ for the environment obtained from ρ by removing all occurrences of $[x \mapsto \dots]$ with $x \notin X$. For the sake of readability we shall write $[x \mapsto v]$ for $[\] [x \mapsto v]$.

We have been very deliberate in when to use intermediate expressions and when to use expressions although it is evident that all expressions are also intermediate expressions. Since we do not evaluate the body of a function before it is applied we continue to let the body be an expression rather than an intermediate expression. Similar remarks apply to the branches of the conditional and the body of the local definitions. Note that although an environment only records the *terms* `fn $x \Rightarrow e_0$` and `fun $f \ x \Rightarrow e_0$` in the closures bound into it, we do not lose the identity of the function abstractions as e_0 will be of the form $t_0^{l_0}$ and hence l_0 may be used as the “*unique*” *identification* of the function abstraction.

Transitions. We are now ready to define the transition rules of the Structural Operational Semantics by means of judgements of the form

$$\rho \vdash ie_1 \rightarrow ie_2$$

$[var]$	$\rho \vdash x^\ell \rightarrow v^\ell$ if $x \in \text{dom}(\rho)$ and $v = \rho(x)$
$[fn]$	$\rho \vdash (\text{fn } x \Rightarrow e_0)^\ell \rightarrow (\text{close } (\text{fn } x \Rightarrow e_0) \text{ in } \rho_0)^\ell$ where $\rho_0 = \rho \upharpoonright \text{FV}(\text{fn } x \Rightarrow e_0)$
$[fun]$	$\rho \vdash (\text{fun } f \ x \Rightarrow e_0)^\ell \rightarrow (\text{close } (\text{fun } f \ x \Rightarrow e_0) \text{ in } \rho_0)^\ell$ where $\rho_0 = \rho \upharpoonright \text{FV}(\text{fun } f \ x \Rightarrow e_0)$
$[app_1]$	$\frac{\rho \vdash ie_1 \rightarrow ie'_1}{\rho \vdash (ie_1 \ ie_2)^\ell \rightarrow (ie'_1 \ ie_2)^\ell}$
$[app_2]$	$\frac{\rho \vdash ie_2 \rightarrow ie'_2}{\rho \vdash (v_1^{\ell_1} \ ie_2)^\ell \rightarrow (v_1^{\ell_1} \ ie'_2)^\ell}$
$[app_{fn}]$	$\rho \vdash ((\text{close } (\text{fn } x \Rightarrow e_1) \text{ in } \rho_1)^{\ell_1} \ v_2^{\ell_2})^\ell \rightarrow$ $(\text{bind } \rho_1[x \mapsto v_2] \text{ in } e_1)^\ell$
$[app_{fun}]$	$\rho \vdash ((\text{close } (\text{fun } f \ x \Rightarrow e_1) \text{ in } \rho_1)^{\ell_1} \ v_2^{\ell_2})^\ell \rightarrow$ $(\text{bind } \rho_2[x \mapsto v_2] \text{ in } e_1)^\ell$ where $\rho_2 = \rho_1[f \mapsto \text{close } (\text{fun } f \ x \Rightarrow e_1) \text{ in } \rho_1]$
$[bind_1]$	$\frac{\rho_1 \vdash ie_1 \rightarrow ie'_1}{\rho \vdash (\text{bind } \rho_1 \text{ in } ie_1)^\ell \rightarrow (\text{bind } \rho_1 \text{ in } ie'_1)^\ell}$
$[bind_2]$	$\rho \vdash (\text{bind } \rho_1 \text{ in } v_1^{\ell_1})^\ell \rightarrow v_1^\ell$

Table 3.2: The Structural Operational Semantics of FUN (part 1).

given by the axioms and inference rules of Tables 3.2 and 3.3; they are explained below. The idea is that *one step* of computation of the expression ie_1 in the environment ρ will transform it into ie_2 .

The value of a variable is obtained from the environment as expressed by the axiom $[var]$. The axioms $[fn]$ and $[fun]$ construct the appropriate closures; they restrict the environment ρ to the free variables of the abstraction. Note that in $[fun]$ it is only recorded that we have a recursively defined function; the unfolding of the recursion will not happen until it is called.

The clauses for application shows that the semantics is a *call-by-value* semantics: In an application we first evaluate the operator in a number of steps using the rule $[app_1]$ and then we evaluate the operand in a number of steps using the rule $[app_2]$. The next stage is to use one of the rules $[app_{fn}]$ or $[app_{fun}]$ to bind the actual parameter to the formal parameter and, in the case of $[app_{fun}]$, to unfold the recursive function so that subsequent recursive calls will be bound correctly. We shall use a *bind*-construct to contain the body of the function together with the appropriate environment. Finally, we

$[if_1]$	$\frac{\rho \vdash ie_0 \rightarrow ie'_0}{\rho \vdash (\text{if } ie_0 \text{ then } e_1 \text{ else } e_2)^\ell \rightarrow (\text{if } ie'_0 \text{ then } e_1 \text{ else } e_2)^\ell}$
$[if_2]$	$\rho \vdash (\text{if true}^{\ell_0} \text{ then } t_1^{\ell_1} \text{ else } t_2^{\ell_2})^\ell \rightarrow t_1^\ell$
$[if_3]$	$\rho \vdash (\text{if false}^{\ell_0} \text{ then } t_1^{\ell_1} \text{ else } t_2^{\ell_2})^\ell \rightarrow t_2^\ell$
$[let_1]$	$\frac{\rho \vdash ie_1 \rightarrow ie'_1}{\rho \vdash (\text{let } x = ie_1 \text{ in } e_2)^\ell \rightarrow (\text{let } x = ie'_1 \text{ in } e_2)^\ell}$
$[let_2]$	$\rho \vdash (\text{let } x = v^{\ell_1} \text{ in } e_2)^\ell \rightarrow (\text{bind } [x \mapsto v] \text{ in } e_2)^\ell$
$[op_1]$	$\frac{\rho \vdash ie_1 \rightarrow ie'_1}{\rho \vdash (ie_1 \text{ op } ie_2)^\ell \rightarrow (ie'_1 \text{ op } ie_2)^\ell}$
$[op_2]$	$\frac{\rho \vdash ie_2 \rightarrow ie'_2}{\rho \vdash (v_1^{\ell_1} \text{ op } ie_2)^\ell \rightarrow (v_1^{\ell_1} \text{ op } ie'_2)^\ell}$
$[op_3]$	$\rho \vdash (v_1^{\ell_1} \text{ op } v_2^{\ell_2})^\ell \rightarrow v^\ell \quad \text{if } v = v_1 \text{ op } v_2$

Table 3.3: The Structural Operational Semantics of FUN (part 2).

evaluate the `bind`-construct using rule $[bind_1]$ a number of times, and we get the result of the application by using rule $[bind_2]$. The interplay between these rules is illustrated by the following example.

Example 3.7 Consider the expression $((\text{fn } x \Rightarrow x^1)^2 (\text{fn } y \Rightarrow y^3)^4)^5$ of Example 3.1. It has the following derivation sequence (explained below):

$$\begin{aligned}
[] &\vdash ((\text{fn } x \Rightarrow x^1)^2 (\text{fn } y \Rightarrow y^3)^4)^5 \\
&\rightarrow ((\text{close } (\text{fn } x \Rightarrow x^1) \text{ in } [])^2 (\text{fn } y \Rightarrow y^3)^4)^5 \\
&\rightarrow ((\text{close } (\text{fn } x \Rightarrow x^1) \text{ in } [])^2 (\text{close } (\text{fn } y \Rightarrow y^3) \text{ in } [])^4)^5 \\
&\rightarrow (\text{bind } [x \mapsto (\text{close } (\text{fn } y \Rightarrow y^3) \text{ in } [])] \text{ in } x^1)^5 \\
&\rightarrow (\text{bind } [x \mapsto (\text{close } (\text{fn } y \Rightarrow y^3) \text{ in } [])] \text{ in} \\
&\quad (\text{close } (\text{fn } y \Rightarrow y^3) \text{ in } [])^1)^5 \\
&\rightarrow (\text{close } (\text{fn } y \Rightarrow y^3) \text{ in } [])^5
\end{aligned}$$

First $[app_1]$ and $[fn]$ are used to evaluate the operator, then $[app_2]$ and $[fn]$ are used to evaluate the operand and $[app_{fn}]$ introduces the `bind`-construct containing the local environment $[x \mapsto (\text{close } (\text{fn } y \Rightarrow y^3) \text{ in } [])]$ needed to evaluate its body. So x^1 is evaluated using $[bind_1]$ and $[var]$, and finally $[bind_2]$ is used to get rid of the local environment. ■

The semantics of the conditional is the usual one: first the condition is eval-

uated in a number of steps using rule $[if_1]$ and then the appropriate branch is selected by rules $[if_2]$ and $[if_3]$. For the local definitions we first compute the value of the bound variable in a number of steps using rule $[let_1]$ and then we introduce a bind-construct using rule $[let_2]$ reflecting that the body of the let-construct has to be evaluated in an extended environment. The rules $[bind_1]$ and $[bind_2]$ are now used to compute the result. For binary expressions we first evaluate the arguments using $[op_1]$ and $[op_2]$ and then the operation itself, denoted op , is performed using $[op_3]$.

As in Chapter 2 the labels have no impact on the semantics but are merely carried along. It is important to note that the outermost label never changes while inner labels may disappear; see for example the rules $[if_2]$ and $[bind_2]$. This is an important property of the semantics that is exploited by the 0-CFA analysis.

Example 3.8 Let us consider the expression, loop

$$\begin{aligned} & (\text{let } g = (\text{fun } f \ x \Rightarrow (f^1 (\text{fn } y \Rightarrow y^2)^3)^4)^5 \\ & \text{in } (g^6 (\text{fn } z \Rightarrow z^7)^8)^9)^{10} \end{aligned}$$

of Example 3.2 and see how the informal explanation of its semantics is captured in the formal semantics. First we introduce abbreviations for three closures:

$$\begin{aligned} f &= \text{close } (\text{fun } f \ x \Rightarrow (f^1 (\text{fn } y \Rightarrow y^2)^3)^4) \text{ in } [] \\ \text{id}_y &= \text{close } (\text{fn } y \Rightarrow y^2) \text{ in } [] \\ \text{id}_z &= \text{close } (\text{fn } z \Rightarrow z^7) \text{ in } [] \end{aligned}$$

Then we have the following derivation sequence

$$\begin{aligned} [] &\vdash \text{loop} \\ &\rightarrow (\text{let } g = f^5 \text{ in } (g^6 (\text{fn } z \Rightarrow z^7)^8)^9)^{10} \\ &\rightarrow (\text{bind } [g \mapsto f] \text{ in } (g^6 (\text{fn } z \Rightarrow z^7)^8)^9)^{10} \\ &\rightarrow (\text{bind } [g \mapsto f] \text{ in } (f^6 (\text{fn } z \Rightarrow z^7)^8)^9)^{10} \\ &\rightarrow (\text{bind } [g \mapsto f] \text{ in } (f^6 \text{id}_z^8)^9)^{10} \\ &\rightarrow (\text{bind } [g \mapsto f] \text{ in} \\ &\quad (\text{bind } [f \mapsto f][x \mapsto \text{id}_z] \text{ in } (f^1 (\text{fn } y \Rightarrow y^2)^3)^4)^9)^{10} \\ &\rightarrow^* (\text{bind } [g \mapsto f] \text{ in} \\ &\quad (\text{bind } [f \mapsto f][x \mapsto \text{id}_z] \text{ in} \\ &\quad\quad (\text{bind } [f \mapsto f][x \mapsto \text{id}_y] \text{ in } (f^1 (\text{fn } y \Rightarrow y^2)^3)^4)^9)^9)^{10} \\ &\rightarrow^* \dots \end{aligned}$$

showing that the program does indeed loop. ■

$[bind]$	$(\widehat{C}, \widehat{\rho}) \models (\text{bind } \rho \text{ in } it_0^{\ell_0})^{\ell}$ iff $(\widehat{C}, \widehat{\rho}) \models it_0^{\ell_0} \wedge \widehat{C}(\ell_0) \subseteq \widehat{C}(\ell) \wedge \rho \mathcal{R} \widehat{\rho}$
$[close]$	$(\widehat{C}, \widehat{\rho}) \models (\text{close } t_0 \text{ in } \rho)^{\ell}$ iff $\{t_0\} \subseteq \widehat{C}(\ell) \wedge \rho \mathcal{R} \widehat{\rho}$

Table 3.4: Abstract Control Flow Analysis for intermediate expressions.

3.2.2 Semantic Correctness

We shall formulate semantic correctness of the Control Flow Analysis as a *subject reduction result*; this is an approach borrowed from type theory and merely says that an acceptable result of the analysis remains acceptable under evaluation. However, in order to do that we need to extend the analysis to intermediate expressions.

Analysis of intermediate expressions. The clauses for the constructs $\text{bind } \rho \text{ in } ie$ and $\text{close } t_0 \text{ in } \rho$ are given in Table 3.4; the remaining clauses are as in Table 3.1 (with the obvious replacements of expressions with intermediate expressions).

The clause $[bind]$ reflects that its body will be executed and hence whatever it evaluates to will also be a possible value for the construct. Additionally, it expresses that there is a certain relationship \mathcal{R} between the local environment (of the semantics) and the abstract environment (of the analysis). The clause $[close]$ is similar in spirit to the clauses for function abstraction: the term of the closure is a possible value of the construct. Additionally, there has to be a relationship \mathcal{R} between the two environments.

Correctness relation. The purpose of the global abstract environment, $\widehat{\rho}$, is to model *all* of the local environments arising during evaluation. We formalise this by defining the *correctness relation*

$$\mathcal{R} : (\mathbf{Env} \times \widehat{\mathbf{Env}}) \rightarrow \{\text{true}, \text{false}\}$$

and demanding that $\rho \mathcal{R} \widehat{\rho}$ for all local environments, ρ , occurring in the intermediate expressions. We then define:

$$\rho \mathcal{R} \widehat{\rho} \text{ iff } \forall x \in \text{dom}(\rho) \subseteq \text{dom}(\widehat{\rho}) \forall t_x \forall \rho_x : \\ (\rho(x) = \text{close } t_x \text{ in } \rho_x) \Rightarrow (t_x \in \widehat{\rho}(x) \wedge \rho_x \mathcal{R} \widehat{\rho})$$

This clearly demands that the function abstraction, t_x , in $\rho(x)$ must be an element of $\widehat{\rho}(x)$. It also shows that all local environments reachable from ρ , e.g. ρ_x , must be modelled by $\widehat{\rho}$ as well. Note that the relation \mathcal{R} is well-defined because each recursive call is performed on a local environment that




Figure 3.3: Preservation of analysis result.

is *strictly smaller* than that of the call itself; thus a simple proof by well-founded induction (Appendix B) suffices for showing the well-definedness of \mathcal{R} .

Example 3.9 Suppose that:

$$\begin{aligned} \rho &= [x \mapsto \text{close } t_1 \text{ in } \rho_1][y \mapsto \text{close } t_2 \text{ in } \rho_2] \\ \rho_1 &= [] \\ \rho_2 &= [x \mapsto \text{close } t_3 \text{ in } \rho_3] \\ \rho_3 &= [] \end{aligned}$$

Then $\rho \mathcal{R} \hat{\rho}$ amounts to $\{t_1, t_3\} \subseteq \hat{\rho}(x) \wedge \{t_2\} \subseteq \hat{\rho}(y)$. ■

We shall sometimes find it helpful to split the definition of \mathcal{R} into two components. For this we make use of the auxiliary relation

$$\mathcal{V} : (\text{Val} \times (\widehat{\text{Env}} \times \widehat{\text{Val}})) \rightarrow \{\text{true}, \text{false}\}$$

and define \mathcal{V} and \mathcal{R} by mutual recursion:

$$\begin{aligned} v \mathcal{V} (\hat{\rho}, \hat{v}) &\text{ iff } \forall t \forall \rho : (v = \text{close } t \text{ in } \rho) \Rightarrow (t \in \hat{v} \wedge \rho \mathcal{R} \hat{\rho}) \\ \rho \mathcal{R} \hat{\rho} &\text{ iff } \forall x \in \text{dom}(\rho) \subseteq \text{dom}(\hat{\rho}) : \rho(x) \mathcal{V} (\hat{\rho}, \hat{\rho}(x)) \end{aligned}$$

Clearly the two definitions are equivalent.

Correctness result. The correctness result is now expressed by:

Theorem 3.10

If $\rho \mathcal{R} \hat{\rho}$ and $\rho \vdash ie \rightarrow ie'$ then $(\hat{C}, \hat{\rho}) \models ie$ implies $(\hat{C}, \hat{\rho}) \models ie'$.

This is illustrated in Figure 3.3 for a terminating evaluation sequence $\rho \vdash ie \rightarrow^* v^l$; note that the result is analogous to that of Corollary 2.17 for the Live Variables Analysis in Chapter 2.

The intuitive content of the result is as follows:

If there is a possible evaluation of the program such that the function at a call point evaluates to some abstraction, then this abstraction has to be in the set of possible abstractions computed by the analysis.

To see this assume that $\rho \vdash t^\ell \rightarrow^* (\text{close } t_0 \text{ in } \rho_0)^\ell$ and that $(\widehat{C}, \widehat{\rho}) \models t^\ell$ as well as $\rho \mathcal{R} \widehat{\rho}$. Then Theorem 3.10 (and an immediate numerical induction) gives that $(\widehat{C}, \widehat{\rho}) \models (\text{close } t_0 \text{ in } \rho_0)^\ell$. Now from the clause [close] of Table 3.4 we get that $t_0 \in \widehat{C}(\ell)$ as was claimed. It is worth noticing that if the program is closed, i.e. if it does not contain free variables, then ρ will be [] and the condition $\rho \mathcal{R} \widehat{\rho}$ is trivially fulfilled.

Note that the theorem expresses that *all* acceptable analysis results remain acceptable under evaluation. One advantage of this is that we do not need to rely on the existence of a least or “best” solution (to be proved in Subsection 3.2.3) in order to formulate the result. Indeed the result does not say that the “best” solution remains “best” – merely that it remains acceptable. More importantly, the result opens up the possibility that the efficient realisation of Sections 3.3 and 3.4 computes a more approximate solution than the least (perhaps using the techniques of Chapter 4). Finally, note that the formulation of the theorem crucially depends on having defined the analysis for all intermediate expressions rather than just all ordinary expressions.

We shall now turn to the proof of Theorem 3.10. We first state an important observation:

Fact 3.11 If $(\widehat{C}, \widehat{\rho}) \models it^{\ell_1}$ and $\widehat{C}(\ell_1) \subseteq \widehat{C}(\ell_2)$ then $(\widehat{C}, \widehat{\rho}) \models it^{\ell_2}$. ■

Proof By cases on the clauses for “ \models ”. ■

We then prove Theorem 3.10:

Proof We assume that $\rho \mathcal{R} \widehat{\rho}$ and $(\widehat{C}, \widehat{\rho}) \models ie$ and prove $(\widehat{C}, \widehat{\rho}) \models ie'$ by induction on the structure of the inference tree for $\rho \vdash ie \rightarrow ie'$. Most cases simply amount to inspecting the defining clause for $(\widehat{C}, \widehat{\rho}) \models ie$; note that this method of proof applies to all fixed points of a recursive definition and in particular also to the greatest fixed point. We only give the proofs for some of the more interesting cases.

The case [var]. Here $\rho \vdash ie \rightarrow ie'$ is:

$$\rho \vdash x^\ell \rightarrow v^\ell \text{ because } x \in \text{dom}(\rho) \text{ and } v = \rho(x)$$

If $v = c$ there is nothing to prove so suppose that $v = \text{close } t_0 \text{ in } \rho_0$. From $\rho \mathcal{R} \widehat{\rho}$ we get $v \vee (\widehat{\rho}, \widehat{\rho}(x))$ and hence $t_0 \in \widehat{\rho}(x)$ and $\rho_0 \mathcal{R} \widehat{\rho}$. From $(\widehat{C}, \widehat{\rho}) \models ie$ we get $\widehat{\rho}(x) \subseteq \widehat{C}(\ell)$, and hence $t_0 \in \widehat{C}(\ell)$. Since $t_0 \in \widehat{C}(\ell)$ and $\rho_0 \mathcal{R} \widehat{\rho}$ we have established $(\widehat{C}, \widehat{\rho}) \models ie'$.

The case [fn]. Here $\rho \vdash ie \rightarrow ie'$ is:

$$\rho \vdash (\text{fn } x \Rightarrow e_0)^{\ell} \rightarrow (\text{close } (\text{fn } x \Rightarrow e_0) \text{ in } \rho_0)^{\ell}$$

where $\rho_0 = \rho \upharpoonright FV(\text{fn } x \Rightarrow e_0)$

From $(\widehat{C}, \widehat{\rho}) \models ie$ we get $(\text{fn } x \Rightarrow e_0) \in \widehat{C}(\ell)$; from $\rho \mathcal{R} \widehat{\rho}$ it is immediate to get $\rho_0 \mathcal{R} \widehat{\rho}$; this then establishes $(\widehat{C}, \widehat{\rho}) \models ie'$.

The case [app₁]. Here $\rho \vdash ie \rightarrow ie'$ is:

$$\rho \vdash (ie_1 \ ie_2)^{\ell} \rightarrow (ie'_1 \ ie_2)^{\ell} \text{ because } \rho \vdash ie_1 \rightarrow ie'_1$$

The defining clauses of $(\widehat{C}, \widehat{\rho}) \models ie$ and $(\widehat{C}, \widehat{\rho}) \models ie'$ are equal except that the former has $(\widehat{C}, \widehat{\rho}) \models ie_1$ where the latter has $(\widehat{C}, \widehat{\rho}) \models ie'_1$. From the induction hypothesis applied to

$$(\widehat{C}, \widehat{\rho}) \models ie_1, \quad \rho \mathcal{R} \widehat{\rho}, \quad \text{and} \quad \rho \vdash ie_1 \rightarrow ie'_1$$

we get $(\widehat{C}, \widehat{\rho}) \models ie'_1$ and the desired result then follows.

The case [app_{fn}]. Here $\rho \vdash ie \rightarrow ie'$ is:

$$\rho \vdash ((\text{close } (\text{fn } x \Rightarrow t_0^{\ell_0}) \text{ in } \rho_1)^{\ell_1} \ v_2^{\ell_2})^{\ell} \rightarrow (\text{bind } \rho_1[x \mapsto v_2] \text{ in } t_0^{\ell_0})^{\ell}$$

From $(\widehat{C}, \widehat{\rho}) \models ie$ we have $(\widehat{C}, \widehat{\rho}) \models (\text{close } (\text{fn } x \Rightarrow t_0^{\ell_0}) \text{ in } \rho_1)^{\ell_1}$ which yields:

$$(\text{fn } x \Rightarrow t_0^{\ell_0}) \in \widehat{C}(\ell_1) \quad \text{and} \quad \rho_1 \mathcal{R} \widehat{\rho}$$

Further we have $(\widehat{C}, \widehat{\rho}) \models v_2^{\ell_2}$; in the case where $v_2 = c$, it is immediate that

$$v_2 \vee (\widehat{\rho}, \widehat{C}(\ell_2))$$

and in the case where $v_2 = \text{close } t_2 \text{ in } \rho_2$ it follows from the definition of $(\widehat{C}, \widehat{\rho}) \models v_2^{\ell_2}$. Finally, the first universally quantified formula of the definition of $(\widehat{C}, \widehat{\rho}) \models ie$ gives:

$$(\widehat{C}, \widehat{\rho}) \models t_0^{\ell_0}, \quad \widehat{C}(\ell_2) \subseteq \widehat{\rho}(x), \quad \text{and} \quad \widehat{C}(\ell_0) \subseteq \widehat{C}(\ell)$$

Now observe that $v_2 \vee (\widehat{\rho}, \widehat{\rho}(x))$ since $\widehat{C}(\ell_2) \subseteq \widehat{\rho}(x)$ follows from the clause (app_{fn}). Since $\rho_1 \mathcal{R} \widehat{\rho}$ we now have

$$(\widehat{C}, \widehat{\rho}) \models t_0^{\ell_0}, \quad \widehat{C}(\ell_0) \subseteq \widehat{C}(\ell), \quad \text{and} \quad (\rho_1[x \mapsto v_2]) \mathcal{R} \widehat{\rho}$$

and this establishes the desired $(\widehat{C}, \widehat{\rho}) \models ie'$.

The case [bind₂]. Here $\rho \vdash ie \rightarrow ie'$ is:

$$\rho \vdash (\text{bind } \rho_1 \text{ in } v_1^{\ell_1})^{\ell} \rightarrow v_1^{\ell}$$

From $(\widehat{C}, \widehat{\rho}) \models ie$ we have $(\widehat{C}, \widehat{\rho}) \models v_1^{\ell_1}$ as well as $\widehat{C}(\ell_1) \subseteq \widehat{C}(\ell)$ and the desired $(\widehat{C}, \widehat{\rho}) \models v_1^{\ell}$ follows from Fact 3.11.

This completes the proof. ■

Example 3.12 From Example 3.7 we have:

$$[] \vdash ((\text{fn } x \Rightarrow x^1)^2 (\text{fn } y \Rightarrow y^3)^4)^5 \rightarrow^* (\text{close } (\text{fn } y \Rightarrow y^3) \text{ in } [])^5$$

Next let $(\widehat{C}_e, \widehat{\rho}_e)$ be as in Example 3.3. Clearly $[] \mathcal{R} \widehat{\rho}_e$ and from Example 3.5 we have:

$$(\widehat{C}_e, \widehat{\rho}_e) \models ((\text{fn } x \Rightarrow x^1)^2 (\text{fn } y \Rightarrow y^3)^4)^5$$

According to Theorem 3.10 we can now conclude:

$$(\widehat{C}_e, \widehat{\rho}_e) \models (\text{close } (\text{fn } y \Rightarrow y^3) \text{ in } [])^5$$

Using Table 3.4 it is easy to check that this is indeed the case. ■

3.2.3 Existence of Solutions

Having defined the analysis in Table 3.1 it is natural to ask the following question: Does each expression e admit a Control Flow Analysis, i.e. does there exist $(\widehat{C}, \widehat{\rho})$ such that $(\widehat{C}, \widehat{\rho}) \models e$? We shall show that the answer to this question is *yes*.

However, this does not exclude the possibility of having many different analyses for the same expression so an additional question is: Does each expression e have a “least” Control Flow Analysis, i.e. does there exist $(\widehat{C}_0, \widehat{\rho}_0)$ such that $(\widehat{C}_0, \widehat{\rho}_0) \models e$ and such that whenever $(\widehat{C}, \widehat{\rho}) \models e$ then $(\widehat{C}_0, \widehat{\rho}_0)$ is “less than” $(\widehat{C}, \widehat{\rho})$? Again, the answer will be *yes*.

Here “least” is with respect to the partial order defined by:

$$(\widehat{C}_1, \widehat{\rho}_1) \sqsubseteq (\widehat{C}_2, \widehat{\rho}_2) \quad \text{iff} \quad (\forall \ell \in \text{Lab} : \widehat{C}_1(\ell) \subseteq \widehat{C}_2(\ell)) \wedge \\ (\forall x \in \text{Var} : \widehat{\rho}_1(x) \subseteq \widehat{\rho}_2(x))$$

It will be the topic of Sections 3.3 and 3.4 (and Mini Project 3.1) to show that the least solution can be computed efficiently for all expressions. However, it may be instructive to give a general proof for the existence of least solutions also for intermediate expressions. To this end we recall the notion of a Moore family (see Appendix A and Exercise 2.7):

A subset Y of a complete lattice $L = (L, \sqsubseteq)$ is a *Moore family* if and only if $(\bigcap Y') \in Y$ for all $Y' \subseteq Y$.

This property is also called the *model intersection property* because whenever we take the “intersection” of a number of “models” we still get a “model”.

Proposition 3.13

For all $ie \in \mathbf{IExp}$ the set $\{(\widehat{C}, \widehat{\rho}) \mid (\widehat{C}, \widehat{\rho}) \models ie\}$ is a Moore family.

It is an immediate corollary that all intermediate expressions ie admit a Control Flow Analysis: Let Y' be the empty set; then $\prod Y'$ is an element of $\{(\widehat{C}, \widehat{\rho}) \mid (\widehat{C}, \widehat{\rho}) \models ie\}$ showing that there exists at least one analysis of ie .

It is also an immediate corollary that all intermediate expressions have a least Control Flow Analysis: Let Y' be the set $\{(\widehat{C}, \widehat{\rho}) \mid (\widehat{C}, \widehat{\rho}) \models ie\}$; then $\prod Y'$ is an element of $\{(\widehat{C}, \widehat{\rho}) \mid (\widehat{C}, \widehat{\rho}) \models ie\}$ so it will also be an analysis of ie . Clearly $\prod Y' \sqsubseteq (\widehat{C}, \widehat{\rho})$ for all other analyses $(\widehat{C}, \widehat{\rho})$ of ie so it is the least analysis result.

In preparation for the proof of Proposition 3.13 we shall first establish an auxiliary result for \mathcal{R} and \mathcal{V} :

Lemma 3.14

- (i) For all $\rho \in \mathbf{Env}$ the set $\{\widehat{\rho} \mid \rho \mathcal{R} \widehat{\rho}\}$ is a Moore family.
- (ii) For all $v \in \mathbf{Val}$ the set $\{(\widehat{\rho}, \widehat{v}) \mid v \mathcal{V} (\widehat{\rho}, \widehat{v})\}$ is a Moore family. ■

Proof To prove (i) we proceed by well-founded induction on ρ (which is also the manner in which the existence of the predicate was proved). Now assume that

$$\forall i \in I : \rho \mathcal{R} \widehat{\rho}_i$$

for some index set I and let us show that $\rho \mathcal{R} (\prod_i \widehat{\rho}_i)$. For this consider x, t_x , and ρ_x such that:

$$\rho(x) = \text{close } t_x \text{ in } \rho_x$$

We then know

$$\forall i \in I : t_x \in \widehat{\rho}_i(x) \wedge \rho_x \mathcal{R} \widehat{\rho}_i$$

and using the induction hypothesis it follows that

$$t_x \in (\prod_i \widehat{\rho}_i)(x) \wedge \rho_x \mathcal{R} (\prod_i \widehat{\rho}_i)$$

(taking care when $I = \emptyset$).

To prove (ii) we simply expand the definition of \mathcal{V} and note that the result then follows from (i). ■

We now prove Proposition 3.13 using coinduction (see Appendix B):

Proof The ternary relation \models of Tables 3.1 and 3.4 is the greatest fixed point of a function \mathcal{Q} as explained in Section 3.1. Now assume that

$$\forall i \in I : (\widehat{C}_i, \widehat{\rho}_i) \models ie$$

and let us prove that $\prod_i (\widehat{C}_i, \widehat{\rho}_i) \models ie$. We shall proceed by coinduction (see Appendix B) so we start by defining the ternary relation \mathcal{Q}' by:

$$(\widehat{C}', \widehat{\rho}') \mathcal{Q}' ie' \text{ iff } (\widehat{C}', \widehat{\rho}') = \prod_i (\widehat{C}_i, \widehat{\rho}_i) \wedge \forall i \in I : (\widehat{C}_i, \widehat{\rho}_i) \models ie'$$

It is then immediate that we have:

$$\prod_i (\widehat{C}_i, \widehat{\rho}_i) Q' ie'$$

The coinduction proof principle requires that we prove

$$Q' \sqsubseteq Q(Q')$$

and this amounts to assuming $(\widehat{C}', \widehat{\rho}') Q' ie'$ and proving that $(\widehat{C}, \widehat{\rho}) (Q(Q')) ie'$. So let us assume that

$$\forall i \in I : (\widehat{C}_i, \widehat{\rho}_i) \models ie'$$

and let us show that:

$$\prod_i (\widehat{C}_i, \widehat{\rho}_i) (Q(Q')) ie'$$

For this we consider each of the clauses for ie' in turn.

Here we shall only deal with the more complicated choice $ie' = (it_1^{t_1} it_2^{t_2})^t$. From

$$\forall i \in I : (\widehat{C}_i, \widehat{\rho}_i) \models (it_1^{t_1} it_2^{t_2})^t$$

we get $\forall i \in I : (\widehat{C}_i, \widehat{\rho}_i) \models it_1^{t_1}$ and hence:

$$\prod_i (\widehat{C}_i, \widehat{\rho}_i) Q' it_1^{t_1}$$

Similarly we get:

$$\prod_i (\widehat{C}_i, \widehat{\rho}_i) Q' it_2^{t_2}$$

Next consider $(\text{fn } x \Rightarrow t_0^{t_0}) \in \prod_i (\widehat{C}_i(\ell_1))$ and let us prove that:

$$\prod_i (\widehat{C}_i(\ell_2)) \subseteq \prod_i (\widehat{\rho}_i(x)), \quad \prod_i (\widehat{C}_i(\ell_0)) \subseteq \prod_i (\widehat{C}_i(\ell)), \quad \prod_i (\widehat{C}_i, \widehat{\rho}_i) Q' t_0^{t_0} \quad (3.1)$$

For all $i \in I$ we have that $(\widehat{C}_i, \widehat{\rho}_i) \models ie'$ and since $(\text{fn } x \Rightarrow t_0^{t_0}) \in \widehat{C}_i(\ell_1)$ we have

$$\widehat{C}_i(\ell_2) \subseteq \widehat{\rho}_i(x), \quad \widehat{C}_i(\ell_0) \subseteq \widehat{C}_i(\ell), \quad \text{and} \quad (\widehat{C}_i, \widehat{\rho}_i) \models t_0^{t_0}$$

and this then gives (3.1) as desired (taking care when $I = \emptyset$). The case of $(\text{fun } f \ x \Rightarrow t_0^{t_0}) \in \prod_i (\widehat{C}_i(\ell_1))$ is similar. This completes the proof. ■

Example 3.15 Let us return to Example 3.5 and consider the following potential analysis results for $((\text{fn } x \Rightarrow x^1)^2 (\text{fn } y \Rightarrow y^3)^4)^5$:

	$(\widehat{C}_e, \widehat{\rho}_e)$	$(\widehat{C}'_e, \widehat{\rho}'_e)$	$(\widehat{C}''_e, \widehat{\rho}''_e)$
1	{fn y => y ³ }	{fn y => y ³ }	{fn y => y ³ }
2	{fn x => x ¹ }	{fn x => x ¹ }	{fn x => x ¹ }
3	∅	{fn x => x ¹ }	{fn y => y ³ }
4	{fn y => y ³ }	{fn y => y ³ }	{fn y => y ³ }
5	{fn y => y ³ }	{fn y => y ³ }	{fn y => y ³ }
x	{fn y => y ³ }	{fn y => y ³ }	{fn y => y ³ }
y	∅	{fn x => x ¹ }	{fn y => y ³ }

It is straightforward to verify that

$$\begin{aligned}(\widehat{C}_e, \widehat{\rho}_e) &\models ((\text{fn } x \Rightarrow x^1)^2 (\text{fn } y \Rightarrow y^3)^4)^5 \\(\widehat{C}_e'', \widehat{\rho}_e'') &\models ((\text{fn } x \Rightarrow x^1)^2 (\text{fn } y \Rightarrow y^3)^4)^5\end{aligned}$$

Now Proposition 3.13 ensures that also:

$$(\widehat{C}_e' \sqcap \widehat{C}_e'', \widehat{\rho}_e' \sqcap \widehat{\rho}_e'') \models ((\text{fn } x \Rightarrow x^1)^2 (\text{fn } y \Rightarrow y^3)^4)^5$$

Neither $(\widehat{C}_e', \widehat{\rho}_e')$ nor $(\widehat{C}_e'', \widehat{\rho}_e'')$ is a least solution. Their “intersection” $(\widehat{C}_e' \sqcap \widehat{C}_e'', \widehat{\rho}_e' \sqcap \widehat{\rho}_e'')$ is smaller and equals $(\widehat{C}_e, \widehat{\rho}_e)$ which turns out to be the least analysis result for the expression. ■

3.2.4 Coinduction versus Induction

One of the important aspects of the development of the abstract Control Flow Analysis in Table 3.1 is the *coinductive definition* of the acceptability relation:

\models as the *greatest* fixed point of a function Q

An alternative might be an *inductive definition* of an acceptability relation:

\models' as the *least* fixed point of the function Q .

However, in Example 3.6 we argued that this might be inappropriate and here we are going to demonstrate that an important part of the development of the previous subsection fails for the least fixed point of Q .

Proposition 3.16

There exists $e_* \in \mathbf{Exp}$ such that $\{(\widehat{C}, \widehat{\rho}) \mid (\widehat{C}, \widehat{\rho}) \models' e_*\}$ is *not* a Moore family.

Proof (sketch) This proof is rather demanding and is best omitted on a first reading. To make the proof tractable we consider

$$\begin{aligned}e_* &= t_*^t \\t_* &= (\text{fn } x \Rightarrow (x^t x^t)^t)^t (\text{fn } x \Rightarrow (x^t x^t)^t)^t\end{aligned}$$

and take:

$$\begin{aligned}\mathbf{Lab}_e &= \{\ell\} \\ \mathbf{Var}_e &= \{x\} \\ \mathbf{Term}_e &= \{t_*, \text{fn } x \Rightarrow (x^t x^t)^t, x^t x^t, x\} \\ \mathbf{IExp}_e &= \{t^t \mid t \in \mathbf{Term}_e\} \\ \widehat{\mathbf{Val}}_e &= \mathcal{P}(\{\text{fn } x \Rightarrow (x^t x^t)^t\}) = \{\emptyset, \{\text{fn } x \Rightarrow (x^t x^t)^t\}\}\end{aligned}$$

This is in line with Exercise 3.2 (and the development of Subsection 3.3.2) and as we shall see the proof of Proposition 3.13 is not invalidated.

Next let Q be the function defined by Table 3.1 and let Q be in the domain of Q . The condition

$$Q = Q(Q)$$

is equivalent to:

$$\forall t \in \text{Term}_e : \forall (\widehat{C}, \widehat{\rho}) : ((\widehat{C}, \widehat{\rho}) Q t^t \text{ iff } (\widehat{C}, \widehat{\rho}) Q(Q) t^t)$$

By considering the four possibilities of $t \in \text{Term}_e$ this is then equivalent to the conjunction of the following four conditions (where $(\widehat{C}, \widehat{\rho})$ is universally quantified):

$$\begin{aligned} (\widehat{C}, \widehat{\rho}) Q x^t &\text{ iff } \widehat{\rho}(x) \subseteq \widehat{C}(\ell) \\ (\widehat{C}, \widehat{\rho}) Q (\text{fn } x \Rightarrow (x^t x^t)^t) &\text{ iff } \{\text{fn } x \Rightarrow (x^t x^t)^t\} \subseteq \widehat{C}(\ell) \\ (\widehat{C}, \widehat{\rho}) Q (x^t x^t)^t &\text{ iff } (\widehat{C}, \widehat{\rho}) Q x^t \wedge \\ &\quad \widehat{C}(\ell) \neq \emptyset \Rightarrow ((\widehat{C}, \widehat{\rho}) Q (x^t x^t)^t \wedge \\ &\quad \widehat{C}(\ell) \subseteq \widehat{\rho}(x)) \\ (\widehat{C}, \widehat{\rho}) Q t_*^t &\text{ iff } (\widehat{C}, \widehat{\rho}) Q (\text{fn } x \Rightarrow (x^t x^t)^t) \wedge \\ &\quad \widehat{C}(\ell) \neq \emptyset \Rightarrow ((\widehat{C}, \widehat{\rho}) Q (x^t x^t)^t \wedge \\ &\quad \widehat{C}(\ell) \subseteq \widehat{\rho}(x)) \end{aligned}$$

Here we have used that $\widehat{C}(\ell) \neq \emptyset$ implies that $\widehat{C}(\ell) = \{\text{fn } x \Rightarrow (x^t x^t)^t\}$ as follows from the definition of $\widehat{\text{Val}}_e$ in the beginning of this proof.

The conjunction of the above four conditions implies the conjunction of the following four conditions:

$$\begin{aligned} (\widehat{C}, \widehat{\rho}) Q x^t &\text{ iff } \widehat{\rho}(x) \subseteq \widehat{C}(\ell) \\ (\widehat{C}, \widehat{\rho}) Q (\text{fn } x \Rightarrow (x^t x^t)^t) &\text{ iff } \{\text{fn } x \Rightarrow (x^t x^t)^t\} \subseteq \widehat{C}(\ell) \\ (\widehat{C}, \widehat{\rho}) Q (x^t x^t)^t &\text{ iff } \widehat{\rho}(x) \subseteq \widehat{C}(\ell) \wedge \\ &\quad (\widehat{C}(\ell) \neq \emptyset \Rightarrow (\widehat{C}, \widehat{\rho}) Q (x^t x^t)^t) \wedge \\ &\quad \widehat{C}(\ell) \subseteq \widehat{\rho}(x) \\ (\widehat{C}, \widehat{\rho}) Q t_*^t &\text{ iff } \{\text{fn } x \Rightarrow (x^t x^t)^t\} \subseteq \widehat{C}(\ell) \wedge \\ &\quad (\widehat{C}, \widehat{\rho}) Q (x^t x^t)^t \wedge \\ &\quad \widehat{C}(\ell) \subseteq \widehat{\rho}(x) \end{aligned}$$

This implication can be reversed and this shows that also the conjunct of the above four conditions is equivalent to $Q = Q(Q)$.

Using that $\widehat{\rho}(x)$ can only be \emptyset or $\{\text{fn } x \Rightarrow (x^t x^t)^t\}$, and similarly for $\widehat{C}(\ell)$, the above four conditions are equivalent to the following:

$$(\widehat{C}, \widehat{\rho}) Q x^t \text{ iff } \widehat{\rho}(x) \subseteq \widehat{C}(\ell)$$

$$\begin{aligned}
(\widehat{C}, \widehat{\rho}) Q (\text{fn } x \Rightarrow (x^l x^l)^l) &\text{ iff } \{\text{fn } x \Rightarrow (x^l x^l)^l\} = \widehat{C}(\ell) \\
(\widehat{C}, \widehat{\rho}) Q (x^l x^l)^l &\text{ iff } \widehat{\rho}(x) = \widehat{C}(\ell) \wedge \\
&\quad (\widehat{C}(\ell) \neq \emptyset \Rightarrow (\widehat{C}, \widehat{\rho}) Q (x^l x^l)^l) \\
(\widehat{C}, \widehat{\rho}) Q t_*^l &\text{ iff } \{\text{fn } x \Rightarrow (x^l x^l)^l\} = \widehat{C}(\ell) = \widehat{\rho}(x) \wedge \\
&\quad (\widehat{C}, \widehat{\rho}) Q (x^l x^l)^l
\end{aligned}$$

It follows that the conjunct of the above four conditions is once more equivalent to $Q = \mathcal{Q}(Q)$.

The crucial case in the definition of $(\widehat{C}, \widehat{\rho}) Q e$ is for $e = (x^l x^l)^l$ as this determines the truth or falsity of all other cases. We shall now try to get a handle on the candidates Q_1, \dots, Q_n for satisfying $Q_i = \mathcal{Q}(Q_i)$. Concentrating on the condition for $(x^l x^l)^l$ it follows that $(\widehat{C}, \widehat{\rho}) Q_i (x^l x^l)^l$ must demand that $\widehat{C}(\ell) = \widehat{\rho}(x)$. Since each of $\widehat{C}(\ell)$ and $\widehat{\rho}(x)$ can only be $\{\text{fn } x \Rightarrow (x^l x^l)^l\}$ or \emptyset there are at most the following four candidates for Q_i :

$$\begin{aligned}
(\widehat{C}, \widehat{\rho}) Q_1 (x^l x^l)^l &\text{ iff } \widehat{C}(\ell) = \widehat{\rho}(x) \\
(\widehat{C}, \widehat{\rho}) Q_2 (x^l x^l)^l &\text{ iff } \widehat{C}(\ell) = \widehat{\rho}(x) = \emptyset \\
(\widehat{C}, \widehat{\rho}) Q_3 (x^l x^l)^l &\text{ iff } \widehat{C}(\ell) = \widehat{\rho}(x) \neq \emptyset \\
(\widehat{C}, \widehat{\rho}) Q_4 (x^l x^l)^l &\text{ iff } \text{false}
\end{aligned}$$

Verifying the condition

$$\forall (\widehat{C}, \widehat{\rho}) : \left((\widehat{C}, \widehat{\rho}) Q_i (x^l x^l)^l \text{ iff } \widehat{\rho}(x) = \widehat{C}(\ell) \wedge (\widehat{C}(\ell) \neq \emptyset \Rightarrow (\widehat{C}, \widehat{\rho}) Q_i (x^l x^l)^l) \right)$$

for $i \in \{1, 2, 3, 4\}$ it follows that Q_1 and Q_2 satisfy the condition whereas Q_3 and Q_4 do not.

It is now straightforward to verify also the remaining three conditions and it follows that:

$$Q_i = \mathcal{Q}(Q_i) \text{ for } i = 1, 2$$

This means that Q_1 equals \models (the greatest fixed point of \mathcal{Q}) and that Q_2 equals \models' (the least fixed point of \mathcal{Q}). One can then calculate that

$$\begin{aligned}
(\widehat{C}, \widehat{\rho}) Q_1 t_*^l &\text{ iff } \widehat{C}(\ell) = \widehat{\rho}(x) \neq \emptyset \\
(\widehat{C}, \widehat{\rho}) Q_2 t_*^l &\text{ iff } \text{false}
\end{aligned}$$

and this shows that

$$\{(\widehat{C}, \widehat{\rho}) \mid (\widehat{C}, \widehat{\rho}) Q_1 e_*\} = \{(\widehat{C}, \widehat{\rho}) \mid \widehat{C}(\ell) = \widehat{\rho}(x) = \{\text{fn } x \Rightarrow (x^l x^l)^l\}\}$$

which is a singleton set and in fact a Moore family, whereas

$$\{(\widehat{C}, \widehat{\rho}) \mid (\widehat{C}, \widehat{\rho}) Q_2 e_*\} = \emptyset$$

which cannot be a Moore family (since a Moore family is never empty). This completes the proof. \blacksquare

3.3 Syntax Directed 0-CFA Analysis

We shall now show how to obtain efficient realisations of 0-CFA analyses. So assume throughout this section that $e_* \in \mathbf{Exp}$ is the expression of interest and that we want to find a “good” solution $(\widehat{C}, \widehat{\rho})$ satisfying $(\widehat{C}, \widehat{\rho}) \models e_*$. This entails finding a solution that is *as small as possible* with respect to the partial order \sqsubseteq defined in Section 3.2 by:

$$(\widehat{C}_1, \widehat{\rho}_1) \sqsubseteq (\widehat{C}_2, \widehat{\rho}_2) \text{ iff } (\forall \ell : \widehat{C}_1(\ell) \subseteq \widehat{C}_2(\ell)) \wedge (\forall x : \widehat{\rho}_1(x) \subseteq \widehat{\rho}_2(x))$$

Proposition 3.13 shows that a least solution does exist; however, the algorithm that is implicit in the proof does not have tractable (i.e. polynomial) complexity: it involves enumerating all candidate solutions, determining if they are indeed solutions, and if so taking the greatest lower bound with respect to the others found so far.

An alternative approach is somehow to obtain a finite set of constraints, say of the form $lhs \subseteq rhs$ (where lhs and rhs are much as described in Section 3.1), and then take the least solution to this system of constraints. The most obvious method is to expand the formula $(\widehat{C}, \widehat{\rho}) \models e_*$ by unfolding all “recursive calls”, using memorisation to keep track of all the expansions that have been performed so far, and stopping the expansion whenever a previously expanded call is re-encountered.

Three phases. We shall take a more direct route motivated by the above considerations; it has three phases:

- (i) The specification of Table 3.1 is reformulated in a syntax directed manner (Subsection 3.3.1).
- (ii) The syntax directed specification is turned into an algorithm for constructing a finite set of constraints (Subsection 3.4.1).
- (iii) The least solution of this set of constraints is computed (Subsection 3.4.2).

This is indeed a *common phenomenon*: a specification “ \models_A ” is reformulated into a specification “ \models_B ” ensuring that

$$(\widehat{C}, \widehat{\rho}) \models_A e_* \Leftarrow (\widehat{C}, \widehat{\rho}) \models_B e_*$$

so that “ \models_B ” is a *safe approximation* to “ \models_A ” and in particular the best (i.e. least) solution to “ $\models_B e_*$ ” is also a solution to “ $\models_A e_*$ ”. This also ensures that all solutions to “ \models_B ” are semantically correct (assuming that this has already been established for all solutions to “ \models_A ”).

If additionally

$$(\widehat{C}, \widehat{\rho}) \models_A e_* \Rightarrow (\widehat{C}, \widehat{\rho}) \models_B e_*$$

then we can be assured that *no solutions are lost* and hence the best (i.e. least) solution to " $\models_B e_*$ " will also be the best (i.e. least) solution to " $\models_A e_*$ ". However, as we shall see, it may be necessary to *restrict the attention* to only solutions $(\widehat{C}, \widehat{\rho})$ satisfying some additional properties (e.g. that only program fragments of e_* appear in the range of \widehat{C} and $\widehat{\rho}$).

3.3.1 Syntax Directed Specification

In reformulating the specification of " $\models e_*$ " into a more computationally oriented specification " $\models_s e_*$ " we shall ensure that each function body is analysed *at most once* rather than each time the function could be applied. One way to achieve this is to analyse each function body *exactly once* as is done in the *syntax directed O-CFA* analysis of Table 3.5; an alternative would be to analyse only reachable function bodies and we refer to Mini Project 3.1 for how to achieve this. In Table 3.5 each function body is therefore analysed in the relevant clause for function abstraction rather than in the clause for function application; thus we now risk analysing unreachable program fragments.

Since semantic correctness was dealt with in Section 3.2 there is no longer any need to consider intermediate expressions and consequently our specification of

$$(\widehat{C}, \widehat{\rho}) \models_s e$$

in Table 3.5 considers ordinary expressions only. We shall take " \models_s " to be the *largest* relation that satisfies the specification; however, given the syntax directed nature of the specification there is in fact only one relation that satisfies the specification (see Exercise 3.9). Hence it would be technically correct, but intuitively misleading, to claim that we take the least relation that satisfies the specification.

Example 3.17 Consider the expression loop

$$\begin{aligned} &(\text{let } g = (\text{fun } f \ x \Rightarrow (f^1 (fn \ y \Rightarrow y^2)^3)^4)^5 \\ &\text{in } (g^6 (fn \ z \Rightarrow z^7)^8)^9)^{10} \end{aligned}$$

of Example 3.4. We shall verify that $(\widehat{C}_{lp}, \widehat{\rho}_{lp}) \models_s \text{loop}$ where \widehat{C}_{lp} and $\widehat{\rho}_{lp}$ are as in Example 3.4. Using the clause [let], it is sufficient to show

$$(\widehat{C}_{lp}, \widehat{\rho}_{lp}) \models_s (\text{fun } f \ x \Rightarrow (f^1 (fn \ y \Rightarrow y^2)^3)^4)^5 \quad (3.2)$$

$$(\widehat{C}_{lp}, \widehat{\rho}_{lp}) \models_s (g^6 (fn \ z \Rightarrow z^7)^8)^9 \quad (3.3)$$

since we have $\widehat{C}_{lp}(5) \subseteq \widehat{\rho}_{lp}(g)$ and $\widehat{C}_{lp}(9) \subseteq \widehat{C}_{lp}(10)$. To show (3.2) we use the clause [fun] and it is sufficient to show

[con]	$(\widehat{C}, \widehat{\rho}) \models_s c^l$ always
[var]	$(\widehat{C}, \widehat{\rho}) \models_s x^l$ iff $\widehat{\rho}(x) \subseteq \widehat{C}(l)$
[fn]	$(\widehat{C}, \widehat{\rho}) \models_s (\text{fn } x \Rightarrow e_0)^l$ iff $\{\text{fn } x \Rightarrow e_0\} \subseteq \widehat{C}(l) \wedge$ $(\widehat{C}, \widehat{\rho}) \models_s e_0$
[fun]	$(\widehat{C}, \widehat{\rho}) \models_s (\text{fun } f \ x \Rightarrow e_0)^l$ iff $\{\text{fun } f \ x \Rightarrow e_0\} \subseteq \widehat{C}(l) \wedge$ $(\widehat{C}, \widehat{\rho}) \models_s e_0 \wedge \{\text{fun } f \ x \Rightarrow e_0\} \subseteq \widehat{\rho}(f)$
[app]	$(\widehat{C}, \widehat{\rho}) \models_s (t_1^{l_1} t_2^{l_2})^l$ iff $(\widehat{C}, \widehat{\rho}) \models_s t_1^{l_1} \wedge (\widehat{C}, \widehat{\rho}) \models_s t_2^{l_2} \wedge$ $(\forall (\text{fn } x \Rightarrow t_0^{l_0}) \in \widehat{C}(l_1) :$ $\quad \widehat{C}(l_2) \subseteq \widehat{\rho}(x) \wedge \widehat{C}(l_0) \subseteq \widehat{C}(l)) \wedge$ $(\forall (\text{fun } f \ x \Rightarrow t_0^{l_0}) \in \widehat{C}(l_1) :$ $\quad \widehat{C}(l_2) \subseteq \widehat{\rho}(x) \wedge \widehat{C}(l_0) \subseteq \widehat{C}(l))$
[if]	$(\widehat{C}, \widehat{\rho}) \models_s (\text{if } t_0^{l_0} \text{ then } t_1^{l_1} \text{ else } t_2^{l_2})^l$ iff $(\widehat{C}, \widehat{\rho}) \models_s t_0^{l_0} \wedge$ $(\widehat{C}, \widehat{\rho}) \models_s t_1^{l_1} \wedge (\widehat{C}, \widehat{\rho}) \models_s t_2^{l_2} \wedge$ $\widehat{C}(l_1) \subseteq \widehat{C}(l) \wedge \widehat{C}(l_2) \subseteq \widehat{C}(l)$
[let]	$(\widehat{C}, \widehat{\rho}) \models_s (\text{let } x = t_1^{l_1} \text{ in } t_2^{l_2})^l$ iff $(\widehat{C}, \widehat{\rho}) \models_s t_1^{l_1} \wedge (\widehat{C}, \widehat{\rho}) \models_s t_2^{l_2} \wedge$ $\widehat{C}(l_1) \subseteq \widehat{\rho}(x) \wedge \widehat{C}(l_2) \subseteq \widehat{C}(l)$
[op]	$(\widehat{C}, \widehat{\rho}) \models_s (t_1^{l_1} \text{ op } t_2^{l_2})^l$ iff $(\widehat{C}, \widehat{\rho}) \models_s t_1^{l_1} \wedge (\widehat{C}, \widehat{\rho}) \models_s t_2^{l_2}$

Table 3.5: Syntax directed Control Flow Analysis.

$$(\widehat{C}_{lp}, \widehat{\rho}_{lp}) \models_s (f^1 (\text{fn } y \Rightarrow y^2)^3)^4$$

since $f \in \widehat{C}_{lp}(5)$ and $f \in \widehat{\rho}_{lp}(f)$. Now $\widehat{C}_{lp}(1) = \{f\}$ so, according to clause [app] this follows from

$$\begin{aligned} (\widehat{C}_{lp}, \widehat{\rho}_{lp}) &\models_s f^1 \\ (\widehat{C}_{lp}, \widehat{\rho}_{lp}) &\models_s (\text{fn } y \Rightarrow y^2)^3 \end{aligned}$$

since $\widehat{C}_{lp}(3) \subseteq \widehat{\rho}_{lp}(x)$ and $\widehat{C}_{lp}(4) \subseteq \widehat{C}_{lp}(4)$. The first clause follows from [var] since $\widehat{\rho}_{lp}(f) \subseteq \widehat{C}_{lp}(1)$ and for the last clause we observe that $\text{id}_y \in \widehat{C}_{lp}(3)$ and $(\widehat{C}_{lp}, \widehat{\rho}_{lp}) \models_s y^2$ as follows from $\widehat{\rho}_{lp}(y) \subseteq \widehat{C}_{lp}(2)$.

To show (3.3) we observe that $\widehat{C}_{lp}(6) = \{f\}$ so using [app] it is sufficient to show

$$\begin{aligned} (\widehat{C}_{lp}, \widehat{\rho}_{lp}) &\models_s g^6 \\ (\widehat{C}_{lp}, \widehat{\rho}_{lp}) &\models_s (\text{fn } z \Rightarrow z^7)^8 \end{aligned}$$

since $\widehat{C}_{lp}(8) \subseteq \widehat{\rho}_{lp}(x)$ and $\widehat{C}_{lp}(4) \subseteq \widehat{C}_{lp}(8)$. This is straightforward except for the last clause where we observe that $\text{id}_z \in \widehat{C}_{lp}(8)$ and $(\widehat{C}_{lp}, \widehat{\rho}_{lp}) \models_s z^7$ as follows from $\widehat{\rho}_{lp}(y) \subseteq \widehat{C}_{lp}(7)$.

Note that because the analysis is syntax directed we have had *no need for coinduction*, unlike what was the case in Example 3.6. ■

3.3.2 Preservation of Solutions

The specification of the analysis in Table 3.5 uses potentially infinite value spaces although this is not really necessary (as Exercise 3.2 demonstrates for Table 3.1). We can easily restrict ourselves to entities occurring in the original expression and this forms the basis for relating the results of the analysis of Table 3.5 to those of the analysis of Table 3.1.

So let $\text{Lab}_* \subseteq \text{Lab}$ be the finite set of labels occurring in the program e_* of interest, let $\text{Var}_* \subseteq \text{Var}$ be the finite set of variables occurring in e_* and let Term_* be the finite set of subterms occurring in e_* . Define $(\widehat{C}_*^T, \widehat{\rho}_*^T)$ by:

$$\begin{aligned} \widehat{C}_*^T(\ell) &= \begin{cases} \emptyset & \text{if } \ell \notin \text{Lab}_* \\ \text{Term}_* & \text{if } \ell \in \text{Lab}_* \end{cases} \\ \widehat{\rho}_*^T(x) &= \begin{cases} \emptyset & \text{if } x \notin \text{Var}_* \\ \text{Term}_* & \text{if } x \in \text{Var}_* \end{cases} \end{aligned}$$

Then the claim

$$(\widehat{C}, \widehat{\rho}) \sqsubseteq (\widehat{C}_*^T, \widehat{\rho}_*^T)$$

intuitively expresses that $(\widehat{C}, \widehat{\rho})$ is concerned only with subterms occurring in the expression e_* ; obviously, we are primarily interested in analysis results with that property. Actually, this condition can be “reformulated” as the technically more manageable

$$(\widehat{C}, \widehat{\rho}) \in \widehat{\text{Cache}}_* \times \widehat{\text{Env}}_*$$

where we define $\widehat{\text{Cache}}_* = \text{Lab}_* \rightarrow \widehat{\text{Val}}_*$, $\widehat{\text{Env}}_* = \text{Var}_* \rightarrow \widehat{\text{Val}}_*$ and $\widehat{\text{Val}}_* = \mathcal{P}(\text{Term}_*)$.

We can now show that all the solutions to “ $\models_s e_*$ ” that are “less than” $(\widehat{C}_*^T, \widehat{\rho}_*^T)$ are solutions to “ $\models e_*$ ” as well:

Proposition 3.18

If $(\widehat{C}, \widehat{\rho}) \models_s e_*$ and $(\widehat{C}, \widehat{\rho}) \sqsubseteq (\widehat{C}_*^T, \widehat{\rho}_*^T)$ then $(\widehat{C}, \widehat{\rho}) \models e_*$.

Proof Assume that $(\widehat{C}, \widehat{\rho}) \models_s e_*$ and that $(\widehat{C}, \widehat{\rho}) \sqsubseteq (\widehat{C}_*^T, \widehat{\rho}_*^T)$. Furthermore let Exp_* be the set of expressions occurring in e_* and note that

$$\forall e \in \text{Exp}_* : (\widehat{C}, \widehat{\rho}) \models_s e \quad (3.4)$$

is an immediate consequence of the syntax directed nature of the definition of \models_s .

To show that $(\widehat{C}, \widehat{\rho}) \models e_*$ we proceed by coinduction. We know that " \models " is defined coinductively by the specification of Table 3.1, i.e. " $\models = \text{gfp}(\mathcal{Q})$ " where \mathcal{Q} is the function (implicitly) defined by Table 3.1. Similarly, we know that " $\models_s = \text{gfp}(\mathcal{Q}_s)$ " where \mathcal{Q}_s is the function (implicitly) defined by Table 3.5.

Next write $(\widehat{C}', \widehat{\rho}') \models^* e'$ for $(\widehat{C}', \widehat{\rho}') = (\widehat{C}, \widehat{\rho}) \wedge e' \in \text{Exp}_*$. It now suffices to show

$$(\mathcal{Q}_s(\models_s) \cap \models^*) \subseteq \mathcal{Q}(\models_s \cap \models^*) \quad (3.5)$$

because then " $(\models_s \cap \models^*) \subseteq \mathcal{Q}(\models_s \cap \models^*)$ " follows and hence by coinduction " $(\models_s \cap \models^*) \subseteq \models$ " and since $(\widehat{C}, \widehat{\rho}) \models_s e_*$ as well as $(\widehat{C}, \widehat{\rho}) \models^* e_*$ we then have the required $(\widehat{C}, \widehat{\rho}) \models e_*$.

The proof of (3.5) amounts to a comparison of the right hand sides of Table 3.5 and Table 3.1: for each clause we shall assume that the right hand side of Table 3.5 holds for $(\widehat{C}, \widehat{\rho}, e)$ and that $e \in \text{Exp}_*$ and we shall show that the corresponding right hand side of Table 3.1 holds when all occurrences of " \models " are replaced by " $\models_s \cap \models^*$ ".

The clauses *[con]*, *[var]*, *[if]*, *[let]* and *[op]* are trivial as the right hand sides of Tables 3.5 and 3.1 are similar and the subterms will all be in Exp_* . The clauses *[fn]* and *[fun]* are straightforward as the right hand sides of Table 3.5 imply the right-hand sides of Table 3.1. Finally, we consider the clause *[app]*. For $(\text{fn } x \Rightarrow t_0^{t_0}) \in \widehat{C}(\ell_1)$ we need to show that $(\widehat{C}, \widehat{\rho}) \models_s t_0^{t_0}$; but since $(\widehat{C}, \widehat{\rho}) \sqsubseteq (\widehat{C}_*^T, \widehat{\rho}_*^T)$ this follows from (3.4). For $(\text{fun } f \ x \Rightarrow t_0^{t_0}) \in \widehat{C}(\ell_1)$ we need to show that $(\widehat{C}, \widehat{\rho}) \models_s t_0^{t_0}$ and that $(\text{fun } f \ x \Rightarrow t_0^{t_0}) \in \widehat{\rho}(f)$; the first follows from (3.4) and the second is an immediate consequence of $(\widehat{C}, \widehat{\rho}) \models_s (\text{fun } f \ x \Rightarrow t_0^{t_0})^\ell$ (for some ℓ) that again follows from (3.4). ■

We can also show an analogue of Proposition 3.13 for the syntax directed analysis:

Proposition 3.19

$\{(\widehat{C}, \widehat{\rho}) \in \widehat{\text{Cache}}_* \times \widehat{\text{Env}}_* \mid (\widehat{C}, \widehat{\rho}) \models_s e_*\}$ is a Moore family.

This result has as immediate corollaries that:

- each expression e_* has a Control Flow Analysis that is “less than” $(\widehat{C}_*, \widehat{\rho}_*)$, and
- each expression e_* has a “least” Control Flow Analysis that is “less than” $(\widehat{C}_*, \widehat{\rho}_*)$.

This means that the properties obtained for the analysis of Table 3.1 in Subsection 3.2.3 also hold for the analysis of Table 3.5 with the additional restriction on the range of the analysis functions. In particular, any analysis result that is acceptable with respect to Table 3.5 (and properly restricted to $\widehat{\text{Cache}}_* \times \widehat{\text{Env}}_*$) is also an acceptable analysis result with respect to Table 3.1. The converse relationship is studied in Exercise 3.11 and Mini Project 3.1.

Proof We shall write $(\widehat{C}_*, \widehat{\rho}_*)$ also for the greatest element of $\widehat{\text{Cache}}_* \times \widehat{\text{Env}}_*$. It is immediate to show that

- (a) $(\widehat{C}_*, \widehat{\rho}_*) \models_s e$
- (b) if $(\widehat{C}_1, \widehat{\rho}_1) \models_s e$ and $(\widehat{C}_2, \widehat{\rho}_2) \models_s e$ then $((\widehat{C}_1, \widehat{\rho}_1) \sqcap (\widehat{C}_2, \widehat{\rho}_2)) \models_s e$.

for all subexpressions e of e_* by means of structural induction on e . This establishes (a) and (b) also for $e = e_*$. Next consider some

$$Y \subseteq \{(\widehat{C}, \widehat{\rho}) \in \widehat{\text{Cache}}_* \times \widehat{\text{Env}}_* \mid (\widehat{C}, \widehat{\rho}) \models_s e_*\}$$

and note that one can write $Y = \{(\widehat{C}_i, \widehat{\rho}_i) \mid i \in \{1, \dots, n\}\}$ for some $n \geq 0$ since $\widehat{\text{Cache}}_* \times \widehat{\text{Env}}_*$ is finite. That

$$\prod Y \in \{(\widehat{C}, \widehat{\rho}) \in \widehat{\text{Cache}}_* \times \widehat{\text{Env}}_* \mid (\widehat{C}, \widehat{\rho}) \models_s e_*\}$$

then follows from (a) and (b) because $\prod Y = (\widehat{C}_*, \widehat{\rho}_*) \sqcap (\widehat{C}_1, \widehat{\rho}_1) \sqcap \dots \sqcap (\widehat{C}_n, \widehat{\rho}_n)$. ■

3.4 Constraint Based 0-CFA Analysis

We are now ready to consider efficient ways of finding the least solution $(\widehat{C}, \widehat{\rho})$ such that $(\widehat{C}, \widehat{\rho}) \models_s e_*$. To do so we first construct a finite set $\mathcal{C}_*[e_*]$ of *constraints* and *conditional constraints* of the form

$$lhs \subseteq rhs \tag{3.6}$$

$$\{t\} \subseteq rhs' \Rightarrow lhs \subseteq rhs \tag{3.7}$$

where rhs is of the form $C(\ell)$ or $r(x)$, and lhs is of the form $C(\ell)$, $r(x)$, or $\{t\}$, and all occurrences of t are of the form $\text{fn } x \Rightarrow e_0$ or $\text{fun } f \ x \Rightarrow e_0$. To simplify the technical development we shall read (3.7) as

$$(\{t\} \subseteq rhs' \Rightarrow lhs) \subseteq rhs$$

[con]	$C_*[c^\ell] = \emptyset$
[var]	$C_*[x^\ell] = \{r(x) \subseteq C(\ell)\}$
[fn]	$C_*[(\text{fn } x \Rightarrow e_0)^\ell] = \{\{\text{fn } x \Rightarrow e_0\} \subseteq C(\ell)\} \cup C_*[e_0]$
[fun]	$C_*[(\text{fun } f \ x \Rightarrow e_0)^\ell] = \{\{\text{fun } f \ x \Rightarrow e_0\} \subseteq C(\ell)\} \cup C_*[e_0] \cup \{\{\text{fun } f \ x \Rightarrow e_0\} \subseteq r(f)\}$
[app]	$C_*[(t_1^{\ell_1} \ t_2^{\ell_2})^\ell] = C_*[t_1^{\ell_1}] \cup C_*[t_2^{\ell_2}] \cup \{\{t\} \subseteq C(\ell_1) \Rightarrow C(\ell_2) \subseteq r(x) \mid t = (\text{fn } x \Rightarrow t_0^{\ell_0}) \in \mathbf{Term}_*\}$ $\cup \{\{t\} \subseteq C(\ell_1) \Rightarrow C(\ell_0) \subseteq C(\ell) \mid t = (\text{fn } x \Rightarrow t_0^{\ell_0}) \in \mathbf{Term}_*\}$ $\cup \{\{t\} \subseteq C(\ell_1) \Rightarrow C(\ell_2) \subseteq r(x) \mid t = (\text{fun } f \ x \Rightarrow t_0^{\ell_0}) \in \mathbf{Term}_*\}$ $\cup \{\{t\} \subseteq C(\ell_1) \Rightarrow C(\ell_0) \subseteq C(\ell) \mid t = (\text{fun } f \ x \Rightarrow t_0^{\ell_0}) \in \mathbf{Term}_*\}$
[if]	$C_*[(\text{if } t_0^{\ell_0} \text{ then } t_1^{\ell_1} \text{ else } t_2^{\ell_2})^\ell] = C_*[t_0^{\ell_0}] \cup C_*[t_1^{\ell_1}] \cup C_*[t_2^{\ell_2}] \cup \{C(\ell_1) \subseteq C(\ell)\} \cup \{C(\ell_2) \subseteq C(\ell)\}$
[let]	$C_*[(\text{let } x = t_1^{\ell_1} \text{ in } t_2^{\ell_2})^\ell] = C_*[t_1^{\ell_1}] \cup C_*[t_2^{\ell_2}] \cup \{C(\ell_1) \subseteq r(x)\} \cup \{C(\ell_2) \subseteq C(\ell)\}$
[op]	$C_*[(t_1^{\ell_1} \ \text{op} \ t_2^{\ell_2})^\ell] = C_*[t_1^{\ell_1}] \cup C_*[t_2^{\ell_2}]$

Table 3.6: Constraint based Control Flow Analysis.

and we shall write ls for lhs as well as $\{t\} \subseteq rhs' \Rightarrow lhs$.

Informally, the constraints are obtained by expanding the clauses defining $(\widehat{C}, \widehat{\rho}) \models_s e_*$ into a finite set of constraints of the above form and then letting $C_*[e_*]$ be the set of individual conjuncts. One caveat is that all occurrences of “ \widehat{C} ” are changed into “ C ” and that all occurrences of “ $\widehat{\rho}$ ” are changed into “ r ” to avoid confusion: $\widehat{C}(\ell)$ will be a *set of terms* whereas $C(\ell)$ is *pure syntax* and similarly for $\widehat{\rho}(x)$ and $r(x)$.

Formally, the *constraint based 0-CFA* analysis is defined by the function C_* of Table 3.6: it makes use of the set \mathbf{Term}_* of subterms occurring in the expression e_* in order to generate only a finite number of constraints in the clause for application; this is justified by Propositions 3.18 and 3.19.

If the size of the expression e_* is n then it might seem that there could

be $O(n^2)$ constraints of the form (3.6) and $O(n^4)$ constraints of the form (3.7). However, inspection of the definition of C_* ensures that at most $O(n)$ constraints of the form (3.6) and $O(n^2)$ constraints of the form (3.7) are ever generated: each of the $O(n)$ constituents only generate $O(1)$ constraints of the form (3.6) and $O(n)$ constraints of the form (3.7).

Example 3.20 Consider the expression

$$((\text{fn } x \Rightarrow x^1)^2 (\text{fn } y \Rightarrow y^3)^4)^5$$

of Example 3.7. We generate the following set of constraints

$$\begin{aligned} C_*[((\text{fn } x \Rightarrow x^1)^2 (\text{fn } y \Rightarrow y^3)^4)^5] = \\ \{ \{ \text{fn } x \Rightarrow x^1 \} \subseteq C(2), \\ r(x) \subseteq C(1), \\ \{ \text{fn } y \Rightarrow y^3 \} \subseteq C(4), \\ r(y) \subseteq C(3), \\ \{ \text{fn } x \Rightarrow x^1 \} \subseteq C(2) \Rightarrow C(4) \subseteq r(x), \\ \{ \text{fn } x \Rightarrow x^1 \} \subseteq C(2) \Rightarrow C(1) \subseteq C(5), \\ \{ \text{fn } y \Rightarrow y^3 \} \subseteq C(2) \Rightarrow C(4) \subseteq r(y), \\ \{ \text{fn } y \Rightarrow y^3 \} \subseteq C(2) \Rightarrow C(3) \subseteq C(5) \} \end{aligned}$$

where we use that $\text{fn } x \Rightarrow x^1$ and $\text{fn } y \Rightarrow y^3$ are the only abstractions in Term_* . ■

3.4.1 Preservation of Solutions

It is important to stress that while $(\widehat{C}, \widehat{\rho}) \models_s e_*$ is a logical formula, $C_*[e_*]$ is a set of syntactic entities. To give meaning to the syntax we first translate the “C” and “r” symbols into the sets “ \widehat{C} ” and “ $\widehat{\rho}$ ”:

$$\begin{aligned} (\widehat{C}, \widehat{\rho})[C(\ell)] &= \widehat{C}(\ell) \\ (\widehat{C}, \widehat{\rho})[r(x)] &= \widehat{\rho}(x) \end{aligned}$$

To deal with the possible forms of ls we additionally take:

$$\begin{aligned} (\widehat{C}, \widehat{\rho})[\{t\}] &= \{t\} \\ (\widehat{C}, \widehat{\rho})[\{t\} \subseteq rhs' \Rightarrow lhs] &= \begin{cases} (\widehat{C}, \widehat{\rho})[lhs] & \text{if } \{t\} \subseteq (\widehat{C}, \widehat{\rho})[rhs'] \\ \emptyset & \text{otherwise} \end{cases} \end{aligned}$$

Next we define a satisfaction relation $(\widehat{C}, \widehat{\rho}) \models_c (ls \subseteq rhs)$ on the individual constraints:

$$(\widehat{C}, \widehat{\rho}) \models_c (ls \subseteq rhs) \quad \text{iff} \quad (\widehat{C}, \widehat{\rho})[ls] \subseteq (\widehat{C}, \widehat{\rho})[rhs]$$

This definition can be lifted to a set C of constraints by:

$$(\widehat{C}, \widehat{\rho}) \models_c C \text{ iff } \forall (ls \subseteq rhs) \in C : (\widehat{C}, \widehat{\rho}) \models_c (ls \subseteq rhs)$$

We then have the following result showing that all solutions to the set $C_*[e_*]$ of constraints also satisfy the syntax directed specification of the Control Flow Analysis and vice versa:

Proposition 3.21

$$(\widehat{C}, \widehat{\rho}) \models_s e_* \text{ if and only if } (\widehat{C}, \widehat{\rho}) \models_c C_*[e_*].$$

Thus the least solution $(\widehat{C}, \widehat{\rho})$ to $(\widehat{C}, \widehat{\rho}) \models_s e_*$ equals the least solution to $(\widehat{C}, \widehat{\rho}) \models_c C_*[e_*]$.

Proof A simple structural induction on e shows that

$$(\widehat{C}, \widehat{\rho}) \models_s e \text{ iff } (\widehat{C}, \widehat{\rho}) \models_c C_*[e]$$

for all subexpressions e of e_* . ■

3.4.2 Solving the Constraints

We shall present two approaches to solving the set of constraints $C_*[e_*]$. First we shall show that finding the least solution to $C_*[e_*]$ is equivalent to finding the least fixed point of a certain function; straightforward techniques allow us to compute that in time $O(n^5)$ when the size of the expression e_* is n . Improvements upon this are possible, but to obtain the best known result we shall consider a graph representation of the problem; this will give us a $O(n^3)$ algorithm. This is indeed a *common phenomenon* in program analysis: syntax directed specifications are appropriate for correctness considerations but often they need to be “massaged” in order to obtain efficient implementations.

Fixed point formulation. To show that finding the solution of the set $C_*[e_*]$ of constraints is a *fixed point problem* we shall define a function

$$F_* : \widehat{\text{Cache}}_* \times \widehat{\text{Env}}_* \rightarrow \widehat{\text{Cache}}_* \times \widehat{\text{Env}}_*$$

and show that it has a least fixed point $\text{lfp}(F_*)$ that is indeed the least solution whose existence is guaranteed by Propositions 3.18 and 3.21.

We define the function F_* by

$$F_*(\widehat{C}, \widehat{\rho}) = (F_1(\widehat{C}, \widehat{\rho}), F_2(\widehat{C}, \widehat{\rho}))$$

where:

$$F_1(\widehat{C}, \widehat{\rho})(\ell) = \bigcup \{(\widehat{C}, \widehat{\rho})[ls] \mid (ls \subseteq C(\ell)) \in C_*[e_*]\}$$

$$F_2(\widehat{C}, \widehat{\rho})(x) = \bigcup \{(\widehat{C}, \widehat{\rho})[ls] \mid (ls \subseteq r(x)) \in C_*[e_*]\}$$

To see that this defines a monotone function it suffices to consider a constraint

$$lhs' \subseteq rhs' \Rightarrow lhs \subseteq rhs$$

in $C_*[e_*]$ and to observe that lhs' is of the form $\{t\}$; this ensures that $(\widehat{C}_1, \widehat{\rho}_1) \subseteq (\widehat{C}_2, \widehat{\rho}_2)$ implies $F_i(\widehat{C}_1, \widehat{\rho}_1) \subseteq F_i(\widehat{C}_2, \widehat{\rho}_2)$ (for $i = 1, 2$) because if $\{t\} \subseteq (\widehat{C}_1, \widehat{\rho}_1)[rhs']$ then also $\{t\} \subseteq (\widehat{C}_2, \widehat{\rho}_2)[rhs']$. Since $\widehat{Cache}_* \times \widehat{Env}_*$ is a complete lattice this means that F_* has a least fixed point and it turns out to be the least solution also to the set $C_*[e_*]$ of constraints:

Proposition 3.22

$$lfp(F_*) = \bigcap \{(\widehat{C}, \widehat{\rho}) \mid (\widehat{C}, \widehat{\rho}) \models_c C_*[e_*]\}$$

Proof It is easy to verify that:

$$F_*(\widehat{C}, \widehat{\rho}) \subseteq (\widehat{C}, \widehat{\rho}) \text{ iff } (\widehat{C}, \widehat{\rho}) \models_c C_*[e_*]$$

Using the formula $lfp(f) = \bigcap \{x \mid f(x) \subseteq x\}$ (see Appendix A) the result then follows. ■

If the size of e_* is n then an element $(\widehat{C}, \widehat{\rho})$ of $\widehat{Cache}_* \times \widehat{Env}_*$ may be viewed as an $O(n)$ -tuple of values from \widehat{Val}_* . Since \widehat{Val}_* is a lattice of height $O(n)$ this means that $\widehat{Cache}_* \times \widehat{Val}_*$ has height $O(n^2)$ and hence the formula

$$lfp(F_*) = \bigsqcup_m F_*^m(\perp)$$

may be used to compute the least fixed point in at most $O(n^2)$ iterations. A naive approach will need to consider all $O(n^2)$ constraints to determine the value of each of the $O(n)$ components of the new iterant; this yields an overall $O(n^5)$ bound on the cost.

Graph formulation. An alternative method for computing the least solution to the set $C_*[e_*]$ of constraints is to use a *graph formulation of the constraints*. The graph will have nodes $C(\ell)$ and $r(x)$ for $\ell \in \mathbf{Lab}_*$ and $x \in \mathbf{Var}_*$. Associated with each node p we have a data field $D[p]$ that initially is given by:

$$D[p] = \{t \mid (\{t\} \subseteq p) \in C_*[e_*]\}$$

The graph will have edges for a subset of the constraints in $C_*[e_*]$; each edge will be decorated with the constraint that gives rise to it:

- a constraint $p_1 \subseteq p_2$ gives rise to an edge from p_1 to p_2 , and

p	$D[p]$	$E[p]$
C(1)	\emptyset	$[\text{id}_x \subseteq C(2) \Rightarrow C(1) \subseteq C(5)]$
C(2)	id_x	$[\text{id}_y \subseteq C(2) \Rightarrow C(3) \subseteq C(5), \text{id}_y \subseteq C(2) \Rightarrow C(4) \subseteq r(y),$ $\text{id}_x \subseteq C(2) \Rightarrow C(1) \subseteq C(5), \text{id}_x \subseteq C(2) \Rightarrow C(4) \subseteq r(x)]$
C(3)	\emptyset	$[\text{id}_y \subseteq C(2) \Rightarrow C(3) \subseteq C(5)]$
C(4)	id_y	$[\text{id}_y \subseteq C(2) \Rightarrow C(4) \subseteq r(y), \text{id}_x \subseteq C(2) \Rightarrow C(4) \subseteq r(x)]$
C(5)	\emptyset	$[\]$
$r(x)$	\emptyset	$[r(x) \subseteq C(1)]$
$r(y)$	\emptyset	$[r(y) \subseteq C(3)]$

Figure 3.4: Initialisation of data structure for example program.

- a constraint $\{t\} \subseteq p \Rightarrow p_1 \subseteq p_2$ gives rise to an edge from p_1 to p_2 and an edge from p to p_2 .

Having constructed the graph we now traverse all edges in order to propagate information from one $D[p_1]$ to another $D[p_2]$. We make certain only to traverse an edge from p_1 to p_2 when $D[p_1]$ is extended with a term not previously there (and this incorporates the situation where $D[p_1]$ is initially set to a non-empty set). Furthermore, an edge decorated with $\{t\} \subseteq p \Rightarrow p_1 \subseteq p_2$ is only traversed if in fact $t \in D[p]$.

To be more specific consider the algorithm of Figure 3.7. It takes as *input* a set $C_*[e_*]$ of constraints and produces as *output* a solution $(\widehat{C}, \widehat{\rho}) \in \widehat{\text{Cache}}_* \times \widehat{\text{Env}}_*$. It operates on the following main *data structures*:

- a *worklist* W i.e. a list of nodes whose outgoing edges should be traversed;
- a *data array* D that for each node gives an element of $\widehat{\text{Val}}_*$; and
- an *edge array* E that for each node gives a list of constraints from which a list of the successor nodes can be computed.

The set Nodes consists of $C(\ell)$ for all ℓ in Lab_* and $r(x)$ for all x in Var_* .

The first step of the algorithm is to initialise the data structures. The second step is to build the graph and to perform the initial assignments to the data fields. This is established using the procedure $\text{add}(q,d)$ that incorporates d into $D[q]$ and adds q to the worklist if d was not part of $D[q]$. The third step is to continue propagating contributions along edges as long as the worklist is non-empty. The fourth and final step is to record the solution in a more familiar form.

INPUT: $C_*[e_*]$
OUTPUT: $(\widehat{C}, \widehat{\rho})$
METHOD: Step 1: Initialisation
 $W := \text{nil};$
for q in Nodes do $D[q] := \emptyset;$
for q in Nodes do $E[q] := \text{nil};$

Step 2: Building the graph
for cc in $C_*[e_*]$ do
case cc of
 $\{t\} \subseteq p$: $\text{add}(p, \{t\});$
 $p_1 \subseteq p_2$: $E[p_1] := \text{cons}(cc, E[p_1]);$
 $\{t\} \subseteq p \Rightarrow p_1 \subseteq p_2$:
 $E[p_1] := \text{cons}(cc, E[p_1]);$
 $E[p] := \text{cons}(cc, E[p]);$

Step 3: Iteration
while $W \neq \text{nil}$ do
 $q := \text{head}(W); W := \text{tail}(W);$
for cc in $E[q]$ do
case cc of
 $p_1 \subseteq p_2$: $\text{add}(p_2, D[p_1]);$
 $\{t\} \subseteq p \Rightarrow p_1 \subseteq p_2$:
if $t \in D[p]$ then $\text{add}(p_2, D[p_1]);$

Step 4: Recording the solution
for ℓ in Lab_* do $\widehat{C}(\ell) := D[C(\ell)];$
for x in Var_* do $\widehat{\rho}(x) := D[r(x)];$

USING: procedure $\text{add}(q, d)$ is
if $\neg (d \subseteq D[q])$
then $D[q] := D[q] \cup d;$
 $W := \text{cons}(q, W);$

Table 3.7: Algorithm for solving constraints.

Example 3.23 Let us consider how the algorithm operates on the expression $((\text{fn } x \Rightarrow x^1)^2 (\text{fn } y \Rightarrow y^3)^4)^5$ of Example 3.20. After step 2 the data structure W has been initialised to

$$W = [C(4), C(2)],$$

and the data structures D and E have been initialised as in Figure 3.4 where we have written id_x for $\{\text{fn } x \Rightarrow x^1\}$ and id_y for $\{\text{fn } y \Rightarrow y^3\}$. The algo-

W	[C(4),C(2)]	[r(x),C(2)]	[C(1),C(2)]	[C(5),C(2)]	[C(2)]	[]
p	D[p]	D[p]	D[p]	D[p]	D[p]	D[p]
C(1)	\emptyset	\emptyset	id_y	id_y	id_y	id_y
C(2)	id_x	id_x	id_x	id_x	id_x	id_x
C(3)	\emptyset	\emptyset	\emptyset	\emptyset	\emptyset	\emptyset
C(4)	id_y	id_y	id_y	id_y	id_y	id_y
C(5)	\emptyset	\emptyset	\emptyset	id_y	id_y	id_y
r(x)	\emptyset	id_y	id_y	id_y	id_y	id_y
r(y)	\emptyset	\emptyset	\emptyset	\emptyset	\emptyset	\emptyset

Figure 3.5: Iteration steps of example program.

rithm will now iterate through the worklist and update the data structures W and D as described by step 3. The various intermediate stages are recorded in Figure 3.5. The algorithm computes the solution in the last column and this agrees with the solution presented in Example 3.5. ■

The following result shows that the algorithm of Table 3.7 does indeed compute the solution we want:

Proposition 3.24

Given input $C_*[e_*]$ the algorithm of Table 3.7 terminates and the result $(\widehat{C}, \widehat{\rho})$ produced by the algorithm satisfies

$$(\widehat{C}, \widehat{\rho}) = \bigcap \{ (\widehat{C}', \widehat{\rho}') \mid (\widehat{C}', \widehat{\rho}') \models_c C_*[e_*] \}$$

and hence it is the least solution to $C_*[e_*]$.

Proof It is immediate that steps 1, 2 and 4 terminate, and this leaves us with step 3. It is immediate that the values of $D[q]$ never decrease and that they can be increased at most a finite number of times. It is also immediate that a node q is added to the worklist only if some value of $D[q]$ actually increased. For each node placed on the worklist only a finite amount of calculation (bounded by the number of outgoing edges) needs to be performed in order to remove the node from the worklist. This guarantees termination.

Next let $(\widehat{C}', \widehat{\rho}')$ be a solution to $(\widehat{C}', \widehat{\rho}') \models_c C_*[e_*]$. It is possible to show that the following invariant

$$\begin{aligned} \forall \ell \in \text{Lab}_* : D[C(\ell)] &\subseteq \widehat{C}'(\ell) \\ \forall x \in \text{Var}_* : D[r(x)] &\subseteq \widehat{\rho}'(x) \end{aligned}$$

is maintained at all points after step 1. It follows that $(\widehat{C}, \widehat{\rho}) \sqsubseteq (\widehat{C}', \widehat{\rho}')$ upon completion of the algorithm.

We prove that $(\widehat{C}, \widehat{\rho}) \models_c C_*[e_*]$ by contradiction. So suppose there exists $cc \in C_*[e_*]$ such that $(\widehat{C}, \widehat{\rho}) \models_c cc$ does *not* hold. If cc is $\{t\} \subseteq p$ then step 2 ensures that $\{t\} \subseteq D[p]$ and this is maintained throughout the algorithm; hence cc cannot have this form. If cc is $p_1 \subseteq p_2$ it must be the case that the final value of D satisfies $D[p_1] \neq \emptyset$ since otherwise $(\widehat{C}, \widehat{\rho}) \models_c cc$ would hold; now consider the last time $D[p_1]$ was modified and note that p_1 was placed on the worklist at that time (by the procedure add); since the final worklist is empty we must have considered the constraint cc (which is in $E[p_1]$) and updated $D[p_2]$ accordingly; hence cc cannot have this form. If cc is $\{t\} \subseteq p \Rightarrow p_1 \subseteq p_2$ it must be the case that the final value of D satisfies $D[p] \neq \emptyset$ as well as $D[p_1] \neq \emptyset$; now consider the last time one of $D[p]$ and $D[p_1]$ was modified and note that p or p_1 was placed on the worklist at that time; since the final worklist is empty we must have considered the constraint cc and updated $D[p_2]$ accordingly; hence cc cannot have this form. Thus $(\widehat{C}, \widehat{\rho}) \models_c cc$ for all $cc \in C_*[e_*]$.

We have now shown that $(\widehat{C}, \widehat{\rho}) \models_c C_*[e_*]$ and that $(\widehat{C}, \widehat{\rho}) \sqsubseteq (\widehat{C}', \widehat{\rho}')$ whenever $(\widehat{C}', \widehat{\rho}') \models_c C_*[e_*]$. It now follows that

$$(\widehat{C}, \widehat{\rho}) = \prod \{(\widehat{C}', \widehat{\rho}') \mid (\widehat{C}', \widehat{\rho}') \models_c C_*[e_*]\}$$

as required. ■

The proof showing that the algorithm terminates can be refined to show that it takes at most $O(n^3)$ steps if the original expression e_* has size n . To see this recall that $C_*[e_*]$ contains at most $O(n)$ constraints of the form $\{t\} \subseteq p$ or $p_1 \subseteq p_2$, and at most $O(n^2)$ constraints of the form $\{t\} \subseteq p \Rightarrow p_1 \subseteq p_2$. We therefore know that the graph has at most $O(n)$ nodes and $O(n^2)$ edges and that each data field can be enlarged at most $O(n)$ times. Assuming that the operations upon $D[p]$ take unit time we can perform the following calculations: step 1 takes time $O(n)$, step 2 takes time $O(n^2)$, and step 4 takes time $O(n)$; step 3 traverses each of the $O(n^2)$ edges at most $O(n)$ times and hence takes time $O(n^3)$; it follows that the overall algorithm takes no more than $O(n^3)$ basic steps.

Combining the three phases. From Proposition 3.24 we get that the pair $(\widehat{C}, \widehat{\rho})$ computed by the algorithm of Table 3.7 is the least solution to $C[e_*]$, so in particular $(\widehat{C}, \widehat{\rho}) \models_c C[e_*]$. Proposition 3.21 shows that a solution to the constraints will also be an acceptable analysis result for the syntax directed specification, hence $(\widehat{C}, \widehat{\rho}) \models_e e_*$. Proposition 3.18 shows that a solution that only involves program fragments of e_* and that is acceptable for the syntax directed specification, also is acceptable for the abstract specification, and therefore $(\widehat{C}, \widehat{\rho}) \models_{e_*} e_*$. Thus we have the following important corollary:

Corollary 3.25 Assume that $(\widehat{C}, \widehat{\rho})$ is the solution to the constraints $C[e_*]$ computed by the algorithm of Table 3.7; then $(\widehat{C}, \widehat{\rho}) \models_{e_*} e_*$. ■

It is not the case that any $(\widehat{C}, \widehat{\rho})$ satisfying $(\widehat{C}, \widehat{\rho}) \models e_*$ can be obtained using the above approach – see Exercise 3.11 and Mini Project 3.1.

For many applications it is the ability to compute the least $(\widehat{C}, \widehat{\rho})$ satisfying $(\widehat{C}, \widehat{\rho}) \models e_*$ that is of primary interest, rather than the ability to check $(\widehat{C}, \widehat{\rho}) \models e_*$ for a proposed guess $(\widehat{C}, \widehat{\rho})$. However, the ability to check $(\widehat{C}, \widehat{\rho}) \models e_*$ is indispensable for *open systems* where the environment e.g. provides a library to be used with e_* . When analysing and optimising e_* it is $(\widehat{C}, \widehat{\rho})$ that expresses the assumptions about the environment; indeed if an existing library e satisfies $(\widehat{C}, \widehat{\rho}) \models e$ for the $(\widehat{C}, \widehat{\rho}) \models e_*$ used to optimise e_* , then one can exchange the library e with any other e' as long as $(\widehat{C}, \widehat{\rho}) \models e'$ continues to hold, and the optimisation made in e_* will continue to hold.

3.5 Adding Data Flow Analysis

In Section 3.1 we indicated that our Control Flow Analysis could be extended with Data Flow Analysis components. Basically, this amounts to extending the set $\widehat{\text{Val}}$ to contain other abstract values than just abstractions. We shall first see how this can be done when the data flow component is a powerset and next we shall see how it can be generalised to complete lattices. We shall present the two approaches as abstract specifications only (in the manner of Section 3.1).

3.5.1 Abstract Values as Powersets

Abstract domains. There are several ways to extend the value domain $\widehat{\text{Val}}$ so as to specify both Control Flow Analysis and Data Flow Analysis. A particularly simple approach is to use a set Data of *abstract data values* (i.e. abstract properties of booleans and integers) since this allows us to define:

$$\widehat{v} \in \widehat{\text{Val}}_d = \mathcal{P}(\text{Term} \cup \text{Data}) \quad \text{abstract values}$$

For each constant $c \in \text{Const}$ we need an element $d_c \in \text{Data}$ specifying the abstract property of c . Similarly, for each operator $op \in \text{Op}$ we need a total function

$$\widehat{op} : \widehat{\text{Val}}_d \times \widehat{\text{Val}}_d \rightarrow \widehat{\text{Val}}_d$$

telling how op operates on abstract properties. Typically, \widehat{op} will have a definition of the form

$$\widehat{v}_1 \widehat{op} \widehat{v}_2 = \bigcup \{d_{op}(d_1, d_2) \mid d_1 \in \widehat{v}_1 \cap \text{Data}, d_2 \in \widehat{v}_2 \cap \text{Data}\}$$

for some function $d_{op} : \text{Data} \times \text{Data} \rightarrow \mathcal{P}(\text{Data})$ specifying how the operator computes with the abstract properties of integers and booleans.

Example 3.26 For a *Detection of Signs Analysis* we take $\text{Data}_{\text{sign}} = \{\text{tt}, \text{ff}, -, 0, +\}$ where tt and ff stand for the two truth values and -, 0, and + for the negative numbers, the number 0, and the positive numbers, respectively. It is then natural to define $d_{\text{true}} = \text{tt}$ and $d_7 = +$ and similarly for the other constants. Taking $\hat{+}$ as an example, it is natural to base its definition on the following table

d_+	tt	ff	-	0	+
tt	\emptyset	\emptyset	\emptyset	\emptyset	\emptyset
ff	\emptyset	\emptyset	\emptyset	\emptyset	\emptyset
-	\emptyset	\emptyset	$\{-\}$	$\{-\}$	$\{-, 0, +\}$
0	\emptyset	\emptyset	$\{-\}$	$\{0\}$	$\{+\}$
+	\emptyset	\emptyset	$\{-, 0, +\}$	$\{+\}$	$\{+\}$

and similarly for the other operators. ■

Acceptability relation. The acceptability relation of the combined analysis has the form

$$(\hat{C}, \hat{\rho}) \models_d e$$

and is presented in Table 3.8. Compared with the analysis of Table 3.1 the clause $[con]$ now records that d_c is a possible value of c and the clause $[op]$ makes use of the function \widehat{op} described above. In the case of $[if]$ we have made sure only to analyse those branches of the conditional that the analysis of the condition indicates the need for; hence we can be more precise than in the pure Control Flow Analysis – the Data Flow Analysis component of the analysis can influence the outcome of the Control Flow Analysis. In the manner of Exercise 3.3 similar improvements can be made to many of the clauses (see Exercise 3.14) and thereby the specification becomes more flow-sensitive.

Example 3.27 Consider the expression:

$$\begin{aligned} &(\text{let } f = (\text{fn } x \Rightarrow (\text{if } (x^1 > 0^2)^3 \text{ then } (\text{fn } y \Rightarrow y^4)^5 \\ &\quad \text{else } (\text{fn } z \Rightarrow 25^6)^7)^8)^9 \\ &\text{in } ((f^{10} \ 3^{11})^{12} \ 0^{13})^{14})^{15} \end{aligned}$$

A pure 0-CFA analysis will not be able to discover that the else-branch of the conditional will never be executed so it will conclude that the subterm with label 12 may evaluate to $\text{fn } y \Rightarrow y^4$ as well as $\text{fn } z \Rightarrow 25^6$ as shown in the first column of Figure 3.6. The second column of Figure 3.6 shows that when we combine the analysis with a Detection of Signs Analysis (outlined in Example 3.26) then the analysis can determine that only $\text{fn } y \Rightarrow y^4$ is a possible abstraction at label 12. Note that the Detection of Signs Analysis (correctly) determines that the expression will evaluate to a value with property $\{0\}$. ■

[con]	$(\widehat{C}, \widehat{\rho}) \models_d c^\ell$ iff $\{d_c\} \subseteq \widehat{C}(\ell)$
[var]	$(\widehat{C}, \widehat{\rho}) \models_d x^\ell$ iff $\widehat{\rho}(x) \subseteq \widehat{C}(\ell)$
[fn]	$(\widehat{C}, \widehat{\rho}) \models_d (\text{fn } x \Rightarrow e_0)^\ell$ iff $\{\text{fn } x \Rightarrow e_0\} \subseteq \widehat{C}(\ell)$
[fun]	$(\widehat{C}, \widehat{\rho}) \models_d (\text{fun } f \ x \Rightarrow e_0)^\ell$ iff $\{\text{fun } f \ x \Rightarrow e_0\} \subseteq \widehat{C}(\ell)$
[app]	$(\widehat{C}, \widehat{\rho}) \models_d (t_1^{t_1'} t_2^{t_2'})^\ell$ iff $(\widehat{C}, \widehat{\rho}) \models_d t_1^{t_1'} \wedge (\widehat{C}, \widehat{\rho}) \models_d t_2^{t_2'} \wedge$ $(\forall (\text{fn } x \Rightarrow t_0^{t_0'}) \in \widehat{C}(\ell_1) :$ $\quad (\widehat{C}, \widehat{\rho}) \models_d t_0^{t_0'} \wedge$ $\quad \widehat{C}(\ell_2) \subseteq \widehat{\rho}(x) \wedge \widehat{C}(\ell_0) \subseteq \widehat{C}(\ell)) \wedge$ $(\forall (\text{fun } f \ x \Rightarrow t_0^{t_0'}) \in \widehat{C}(\ell_1) :$ $\quad (\widehat{C}, \widehat{\rho}) \models_d t_0^{t_0'} \wedge$ $\quad \widehat{C}(\ell_2) \subseteq \widehat{\rho}(x) \wedge \widehat{C}(\ell_0) \subseteq \widehat{C}(\ell) \wedge$ $\quad \{\text{fun } f \ x \Rightarrow t_0^{t_0'}\} \subseteq \widehat{\rho}(f))$
[if]	$(\widehat{C}, \widehat{\rho}) \models_d (\text{if } t_0^{t_0'} \text{ then } t_1^{t_1'} \text{ else } t_2^{t_2'})^\ell$ iff $(\widehat{C}, \widehat{\rho}) \models_d t_0^{t_0'} \wedge$ $(d_{\text{true}} \in \widehat{C}(\ell_0) \Rightarrow ((\widehat{C}, \widehat{\rho}) \models_d t_1^{t_1'} \wedge \widehat{C}(\ell_1) \subseteq \widehat{C}(\ell))) \wedge$ $(d_{\text{false}} \in \widehat{C}(\ell_0) \Rightarrow ((\widehat{C}, \widehat{\rho}) \models_d t_2^{t_2'} \wedge \widehat{C}(\ell_2) \subseteq \widehat{C}(\ell)))$
[let]	$(\widehat{C}, \widehat{\rho}) \models_d (\text{let } x = t_1^{t_1'} \text{ in } t_2^{t_2'})^\ell$ iff $(\widehat{C}, \widehat{\rho}) \models_d t_1^{t_1'} \wedge (\widehat{C}, \widehat{\rho}) \models_d t_2^{t_2'} \wedge$ $\widehat{C}(\ell_1) \subseteq \widehat{\rho}(x) \wedge \widehat{C}(\ell_2) \subseteq \widehat{C}(\ell)$
[op]	$(\widehat{C}, \widehat{\rho}) \models_d (t_1^{t_1'} \text{ op } t_2^{t_2'})^\ell$ iff $(\widehat{C}, \widehat{\rho}) \models_d t_1^{t_1'} \wedge (\widehat{C}, \widehat{\rho}) \models_d t_2^{t_2'} \wedge$ $\widehat{C}(\ell_1) \widehat{\text{op}} \widehat{C}(\ell_2) \subseteq \widehat{C}(\ell)$

Table 3.8: Abstract values as powersets.

The proof techniques introduced in Section 3.2 should suffice for proving the correctness of the analysis with respect to the operational semantics. A slight extension of the algorithmic techniques presented in Sections 3.3 and 3.4 (and in Mini Project 3.1) suffices for obtaining an implementation of the analysis provided that the set **Data** is finite.

Finally, we should stress that a solution to the analysis of Table 3.8 does *not* immediately give a solution to the analysis of Table 3.1. More precisely, $(\widehat{C}, \widehat{\rho}) \models_d e$ does *not* guarantee that $(\widehat{C}', \widehat{\rho}') \models e$ where $\forall \ell : \widehat{C}'(\ell) = \widehat{C}(\ell) \cap \mathbf{Term}$ and $\forall x : \widehat{\rho}'(x) = \widehat{\rho}(x) \cap \mathbf{Term}$. The reason is that the Control Flow

	Section 3.1	Subsection 3.5.1	Subsection 3.5.2	
	$(\widehat{C}, \widehat{\rho})$	$(\widehat{C}, \widehat{\rho})$	$(\widehat{C}, \widehat{\rho})$	$(\widehat{D}, \widehat{\delta})$
1	\emptyset	$\{+\}$	\emptyset	$\{+\}$
2	\emptyset	$\{0\}$	\emptyset	$\{0\}$
3	\emptyset	$\{tt\}$	\emptyset	$\{tt\}$
4	\emptyset	$\{0\}$	\emptyset	$\{0\}$
5	$\{fn\ y \Rightarrow y^4\}$	$\{fn\ y \Rightarrow y^4\}$	$\{fn\ y \Rightarrow y^4\}$	\emptyset
6	\emptyset	\emptyset	\emptyset	\emptyset
7	$\{fn\ z \Rightarrow 25^6\}$	\emptyset	\emptyset	\emptyset
8	$\{fn\ y \Rightarrow y^4,$ $fn\ z \Rightarrow 25^6\}$	$\{fn\ y \Rightarrow y^4\}$	$\{fn\ y \Rightarrow y^4\}$	\emptyset
9	$\{fn\ x \Rightarrow (\dots)^8\}$	$\{fn\ x \Rightarrow (\dots)^8\}$	$\{fn\ x \Rightarrow (\dots)^8\}$	\emptyset
10	$\{fn\ x \Rightarrow (\dots)^8\}$	$\{fn\ x \Rightarrow (\dots)^8\}$	$\{fn\ x \Rightarrow (\dots)^8\}$	\emptyset
11	\emptyset	$\{+\}$	\emptyset	$\{+\}$
12	$\{fn\ y \Rightarrow y^4,$ $fn\ z \Rightarrow 25^6\}$	$\{fn\ y \Rightarrow y^4\}$	$\{fn\ y \Rightarrow y^4\}$	\emptyset
13	\emptyset	$\{0\}$	\emptyset	$\{0\}$
14	\emptyset	$\{0\}$	\emptyset	$\{0\}$
15	\emptyset	$\{0\}$	\emptyset	$\{0\}$
f	$\{fn\ x \Rightarrow (\dots)^8\}$	$\{fn\ x \Rightarrow (\dots)^8\}$	$\{fn\ x \Rightarrow (\dots)^8\}$	\emptyset
x	\emptyset	$\{+\}$	\emptyset	$\{+\}$
y	\emptyset	$\{0\}$	\emptyset	$\{0\}$
z	\emptyset	\emptyset	\emptyset	\emptyset

Figure 3.6: Control Flow and Data Flow Analysis for example program.

Analysis part of Table 3.8 is influenced by the Data Flow Analysis part in the clause $[if]$: if for example the abstract value of the condition does not include d_{true} , then the then-branch will not be analysed.

3.5.2 Abstract Values as Complete Lattices

Abstract domains. Clearly $\widehat{Val}_d = \mathcal{P}(\text{Term} \cup \text{Data})$ is isomorphic to $\mathcal{P}(\text{Term}) \times \mathcal{P}(\text{Data})$. This suggests that the abstract cache $\widehat{C} : \text{Lab} \rightarrow \widehat{Val}_d$ could be split into a term component and a data component and similarly for the abstract environment $\widehat{\rho} : \text{Var} \rightarrow \widehat{Val}_d$.

Having decoupled $\mathcal{P}(\text{Term})$ and $\mathcal{P}(\text{Data})$ we can now consider replacing $\mathcal{P}(\text{Data})$ by a more general collection of properties. An obvious possibility

is to replace $\mathcal{P}(\mathbf{Data})$ by a complete lattice L and perform a development closely related to that of the (forward) Monotone Frameworks of Chapter 2.

So let us define a *monotone structure* to consist of:

- a complete lattice L , and
- a set \mathcal{F} of monotone functions of $L \times L \rightarrow L$.

An *instance* of a monotone structure then consists of the structure (L, \mathcal{F}) and

- a mapping ι from the constants $c \in \mathbf{Const}$ to values in L , and
- a mapping f from the binary operators $op \in \mathbf{Op}$ to functions of \mathcal{F} .

Compared with the instances of the Monotone Frameworks of Section 2.3 we omit the flow component since it will be the responsibility of the Control Flow Analysis to determine this. The component ι has been replaced by the mapping ι giving the *extremal value* for all the constants and the component f mapping labels to transfer functions has been replaced by a mapping of the binary operators to their interpretation.

Example 3.28 A monotone structure corresponding to the development of Subsection 3.5.1 will have L to be $\mathcal{P}(\mathbf{Data})$ and \mathcal{F} to be the monotone functions of $\mathcal{P}(\mathbf{Data}) \times \mathcal{P}(\mathbf{Data}) \rightarrow \mathcal{P}(\mathbf{Data})$.

An instance of the monotone structure is then obtained by taking

$$\iota_c = \{d_c\}$$

for all constants c (and with $d_c \in \mathbf{Data}$ as above) and

$$f_{op}(l_1, l_2) = \bigcup \{d_{op}(d_1, d_2) \mid d_1 \in l_1, d_2 \in l_2\}$$

for all binary operators op (and where $d_{op} : \mathbf{Data} \times \mathbf{Data} \rightarrow \mathcal{P}(\mathbf{Data})$ is as above). ■

Example 3.29 A monotone structure for *Constant Propagation Analysis* will have L to be $\mathbf{Z}_{\perp}^{\top} \times \mathcal{P}(\{\text{tt}, \text{ff}\})$ and \mathcal{F} to be the monotone functions of $L \times L \rightarrow L$.

An instance of the monotone structure is obtained by taking e.g. $\iota_{\top} = (\top, \emptyset)$ and $\iota_{\text{true}} = (\perp, \{\text{tt}\})$. For a binary operator such as $+$ we can take:

$$f_{+}(l_1, l_2) = \begin{cases} (z_1 + z_2, \emptyset) & \text{if } l_1 = (z_1, \dots), l_2 = (z_2, \dots), \\ & \text{and } z_1, z_2 \in \mathbf{Z} \\ (\perp, \emptyset) & \text{if } l_1 = (z_1, \dots), l_2 = (z_2, \dots), \\ & \text{and } z_1 = \perp \text{ or } z_2 = \perp \\ (\top, \emptyset) & \text{otherwise} \end{cases} \quad \blacksquare$$

We can now define the following abstract domains

$$\begin{aligned} \widehat{v} \in \widehat{\mathbf{Val}} &= \mathcal{P}(\mathbf{Term}) && \text{abstract values} \\ \widehat{\rho} \in \widehat{\mathbf{Env}} &= \mathbf{Var} \rightarrow \widehat{\mathbf{Val}} && \text{abstract environments} \\ \widehat{C} \in \widehat{\mathbf{Cache}} &= \mathbf{Lab} \rightarrow \widehat{\mathbf{Val}} && \text{abstract caches} \end{aligned}$$

to take care of the Control Flow Analysis and furthermore

$$\begin{aligned} \widehat{d} \in \widehat{\mathbf{Data}} &= L && \text{abstract data values} \\ \widehat{\delta} \in \widehat{\mathbf{DEnv}} &= \mathbf{Var} \rightarrow \widehat{\mathbf{Data}} && \text{abstract data environments} \\ \widehat{D} \in \widehat{\mathbf{DCache}} &= \mathbf{Lab} \rightarrow \widehat{\mathbf{Data}} && \text{abstract data caches} \end{aligned}$$

to take care of the Data Flow Analysis.

Acceptability relation. The acceptability relation now has the form

$$(\widehat{C}, \widehat{D}, \widehat{\rho}, \widehat{\delta}) \models_D e$$

and it is defined by the clauses of Table 3.9. In the clause $[con]$ we see that the ι . component of the instance is used to restrict the value of the \widehat{D} component of the analysis and in the clause $[op]$ we see how the f . component is used. The clause $[if]$ has explicit tests for the two branches as in the previous approach thereby allowing the Control Flow Analysis to benefit from results obtained by the Data Flow Analysis component. As in the previous subsection, similar improvements can be made to many of the other clauses so as to produce a more flow-sensitive analysis.

Example 3.30 Returning to the expression of Example 3.27 and the Detection of Signs Analysis we now get the analysis result of the last column of Figure 3.6. So we see that the result is as before. ■

The proof techniques introduced in Section 3.2 should suffice for proving the correctness of the analysis with respect to the operational semantics. A slight extension of the algorithmic techniques presented in Sections 3.3 and 3.4 (and in Mini Project 3.1) suffices for obtaining an implementation of the analysis provided that L satisfies the Ascending Chain Condition (as is the case for Monotone Frameworks).

Staging the specification. Let us briefly consider the following alternative clause for $[if]$ where the data flow component *cannot* influence the control flow component because we always make sure that that the analysis result is acceptable for both branches:

$$\begin{aligned} (\widehat{C}, \widehat{D}, \widehat{\rho}, \widehat{\delta}) \models'_D (\text{if } t_0^{\iota_0} \text{ then } t_1^{\iota_1} \text{ else } t_2^{\iota_2})^{\iota} \\ \text{iff } (\widehat{C}, \widehat{D}, \widehat{\rho}, \widehat{\delta}) \models'_D t_0^{\iota_0} \wedge \\ (\widehat{C}, \widehat{D}, \widehat{\rho}, \widehat{\delta}) \models'_D t_1^{\iota_1} \wedge \widehat{C}(\ell_1) \subseteq \widehat{C}(\ell) \wedge \widehat{D}(\ell_1) \sqsubseteq \widehat{D}(\ell) \wedge \\ (\widehat{C}, \widehat{D}, \widehat{\rho}, \widehat{\delta}) \models'_D t_2^{\iota_2} \wedge \widehat{C}(\ell_2) \subseteq \widehat{C}(\ell) \wedge \widehat{D}(\ell_2) \sqsubseteq \widehat{D}(\ell) \end{aligned}$$

[con]	$(\widehat{C}, \widehat{D}, \widehat{\rho}, \widehat{\delta}) \models_D c^l$ iff $\iota_c \sqsubseteq \widehat{D}(l)$
[var]	$(\widehat{C}, \widehat{D}, \widehat{\rho}, \widehat{\delta}) \models_D x^l$ iff $\widehat{\rho}(x) \subseteq \widehat{C}(l) \wedge \widehat{\delta}(x) \sqsubseteq \widehat{D}(l)$
[fn]	$(\widehat{C}, \widehat{D}, \widehat{\rho}, \widehat{\delta}) \models_D (\text{fn } x \Rightarrow e_0)^l$ iff $\{\text{fn } x \Rightarrow e_0\} \subseteq \widehat{C}(l)$
[fun]	$(\widehat{C}, \widehat{D}, \widehat{\rho}, \widehat{\delta}) \models_D (\text{fun } f \ x \Rightarrow e_0)^l$ iff $\{\text{fun } f \ x \Rightarrow e_0\} \subseteq \widehat{C}(l)$
[app]	$(\widehat{C}, \widehat{D}, \widehat{\rho}, \widehat{\delta}) \models_D (t_1^{l_1} \ t_2^{l_2})^l$ iff $(\widehat{C}, \widehat{D}, \widehat{\rho}, \widehat{\delta}) \models_D t_1^{l_1} \wedge (\widehat{C}, \widehat{D}, \widehat{\rho}, \widehat{\delta}) \models_D t_2^{l_2} \wedge$ $(\forall (\text{fn } x \Rightarrow t_0^{l_0}) \in \widehat{C}(l_1) :$ $\quad (\widehat{C}, \widehat{D}, \widehat{\rho}, \widehat{\delta}) \models_D t_0^{l_0} \wedge$ $\quad \widehat{C}(l_2) \subseteq \widehat{\rho}(x) \wedge \widehat{D}(l_2) \sqsubseteq \widehat{\delta}(x) \wedge$ $\quad \widehat{C}(l_0) \subseteq \widehat{C}(l) \wedge \widehat{D}(l_0) \sqsubseteq \widehat{D}(l)) \wedge$ $(\forall (\text{fun } f \ x \Rightarrow t_0^{l_0}) \in \widehat{C}(l_1) :$ $\quad (\widehat{C}, \widehat{D}, \widehat{\rho}, \widehat{\delta}) \models_D t_0^{l_0} \wedge$ $\quad \widehat{C}(l_2) \subseteq \widehat{\rho}(x) \wedge \widehat{D}(l_2) \sqsubseteq \widehat{\delta}(x) \wedge$ $\quad \widehat{C}(l_0) \subseteq \widehat{C}(l) \wedge \widehat{D}(l_0) \sqsubseteq \widehat{D}(l) \wedge$ $\quad \{\text{fun } f \ x \Rightarrow t_0^{l_0}\} \subseteq \widehat{\rho}(f))$
[if]	$(\widehat{C}, \widehat{D}, \widehat{\rho}, \widehat{\delta}) \models_D (\text{if } t_0^{l_0} \text{ then } t_1^{l_1} \text{ else } t_2^{l_2})^l$ iff $(\widehat{C}, \widehat{D}, \widehat{\rho}, \widehat{\delta}) \models_D t_0^{l_0} \wedge$ $(\iota_{\text{true}} \sqsubseteq \widehat{D}(l_0) \Rightarrow (\widehat{C}, \widehat{D}, \widehat{\rho}, \widehat{\delta}) \models_D t_1^{l_1} \wedge$ $\quad \widehat{C}(l_1) \subseteq \widehat{C}(l) \wedge$ $\quad \widehat{D}(l_1) \sqsubseteq \widehat{D}(l)) \wedge$ $(\iota_{\text{false}} \sqsubseteq \widehat{D}(l_0) \Rightarrow (\widehat{C}, \widehat{D}, \widehat{\rho}, \widehat{\delta}) \models_D t_2^{l_2} \wedge$ $\quad \widehat{C}(l_2) \subseteq \widehat{C}(l) \wedge$ $\quad \widehat{D}(l_2) \sqsubseteq \widehat{D}(l))$
[let]	$(\widehat{C}, \widehat{D}, \widehat{\rho}, \widehat{\delta}) \models_D (\text{let } x = t_1^{l_1} \text{ in } t_2^{l_2})^l$ iff $(\widehat{C}, \widehat{D}, \widehat{\rho}, \widehat{\delta}) \models_D t_1^{l_1} \wedge (\widehat{C}, \widehat{D}, \widehat{\rho}, \widehat{\delta}) \models_D t_2^{l_2} \wedge$ $\widehat{C}(l_1) \subseteq \widehat{\rho}(x) \wedge \widehat{D}(l_1) \sqsubseteq \widehat{\delta}(x) \wedge$ $\widehat{C}(l_2) \subseteq \widehat{C}(l) \wedge \widehat{D}(l_2) \sqsubseteq \widehat{D}(l)$
[op]	$(\widehat{C}, \widehat{D}, \widehat{\rho}, \widehat{\delta}) \models_D (t_1^{l_1} \text{ op } t_2^{l_2})^l$ iff $(\widehat{C}, \widehat{D}, \widehat{\rho}, \widehat{\delta}) \models_D t_1^{l_1} \wedge (\widehat{C}, \widehat{D}, \widehat{\rho}, \widehat{\delta}) \models_D t_2^{l_2} \wedge$ $f_{\text{op}}(\widehat{D}(l_1), \widehat{D}(l_2)) \sqsubseteq \widehat{D}(l)$

Table 3.9: Abstract values as complete lattices.

Unlike what was the case for the analyses of Tables 3.8 and 3.9, a solution to the analysis modified in this way *does* give rise to a solution to the analysis of Table 3.1; to be precise $(\widehat{C}, \widehat{D}, \widehat{\rho}, \widehat{\delta}) \models'_D e$ guarantees $(\widehat{C}, \widehat{\rho}) \models e$.

In terms of implementation this modification means that the constraints for the control flow component (\widehat{C} and $\widehat{\rho}$) can be solved first, and based on this the constraints for the data flow component (\widehat{D} and $\widehat{\delta}$) can be solved next. If both sets of constraints are solved for their least solution this will still yield the least solution of the combined set of constraints.

Example 3.31 Let us return to Example 3.30. If we modify the clause for *if* as discussed above then the resulting analysis will have \widehat{C} and $\widehat{\rho}$ as in the column for the pure analysis from Section 3.1 and \widehat{D} and $\widehat{\delta}$ will associate slightly larger sets with some of the labels and variables:

$$\begin{aligned}\widehat{D}(6) &= \{+\} \\ \widehat{D}(14) &= \{0, +\} \\ \widehat{D}(15) &= \{0, +\} \\ \widehat{\delta}(z) &= \{0\}\end{aligned}$$

This analysis is less precise than those of Tables 3.8 and 3.9: it will only determine that the expression will evaluate to a value with the property $\{0, +\}$. ■

3.6 Adding Context Information

The Control Flow Analyses presented so far are imprecise in that they cannot distinguish the various instances of function calls from one another. In the terminology of Section 2.5 the 0-CFA analysis is *context-insensitive* and in the terminology of Control Flow Analysis it is *mono-variant*.

Example 3.32 Consider the expression:

$$\begin{aligned} &(\text{let } f = (\text{fn } x \Rightarrow x^1)^2 \\ &\text{in } ((f^3 f^4)^5 (\text{fn } y \Rightarrow y^6)^7)^8)^9 \end{aligned}$$

The least 0-CFA analysis is given by $(\widehat{C}_{id}, \widehat{\rho}_{id})$:

$$\begin{aligned}\widehat{C}_{id}(1) &= \{\text{fn } x \Rightarrow x^1, \text{fn } y \Rightarrow y^6\} \\ \widehat{C}_{id}(2) &= \{\text{fn } x \Rightarrow x^1\} \\ \widehat{C}_{id}(3) &= \{\text{fn } x \Rightarrow x^1\} \\ \widehat{C}_{id}(4) &= \{\text{fn } x \Rightarrow x^1\}\end{aligned}$$

$$\begin{aligned}
\widehat{C}_{id}(5) &= \{\text{fn } x \Rightarrow x^1, \text{fn } y \Rightarrow y^6\} \\
\widehat{C}_{id}(6) &= \{\text{fn } y \Rightarrow y^6\} \\
\widehat{C}_{id}(7) &= \{\text{fn } y \Rightarrow y^6\} \\
\widehat{C}_{id}(8) &= \{\text{fn } x \Rightarrow x^1, \text{fn } y \Rightarrow y^6\} \\
\widehat{C}_{id}(9) &= \{\text{fn } x \Rightarrow x^1, \text{fn } y \Rightarrow y^6\} \\
\widehat{\rho}_{id}(f) &= \{\text{fn } x \Rightarrow x^1\} \\
\widehat{\rho}_{id}(x) &= \{\text{fn } x \Rightarrow x^1, \text{fn } y \Rightarrow y^6\} \\
\widehat{\rho}_{id}(y) &= \{\text{fn } y \Rightarrow y^6\}
\end{aligned}$$

So we see that x can be bound to $\text{fn } x \Rightarrow x^1$ as well as $\text{fn } y \Rightarrow y^6$ and hence the overall expression (label 9) may evaluate to either of these two abstractions. However, it is easy to see that in fact only $\text{fn } y \Rightarrow y^6$ is a possible result. ■

To get a more precise analysis it is useful to introduce a mechanism that distinguishes different dynamic instances of variables and labels from one another. This results in a *context-sensitive* analysis and in the terminology of Control Flow Analysis the term *poly-variant* is used. There are several approaches to how this can be done. One simple possibility is to expand the program such that the problem does not arise.

Example 3.33 For the expression of Example 3.32 we could for example consider

```

let f1 = (fn x1 => x1)
in let f2 = (fn x2 => x2)
   in (f1 f2) (fn y => y)

```

and then analyse the expanded expression: the 0-CFA analysis is now able to deduce that $x1$ can only be bound to $\text{fn } x2 \Rightarrow x2$ and that $x2$ can only be bound to $\text{fn } y \Rightarrow y$ so the overall expression will evaluate to $\text{fn } y \Rightarrow y$ only. ■

However, a more satisfactory solution to the problem is to extend the analysis with *context information* allowing it to distinguish between the various instances of variables and program points and still analyse the original expression. Examples of such analyses include k -CFA analyses, uniform k -CFA analyses and polynomial k -CFA analyses (for $k \geq 0$). This approach is clearly related to the use of call-strings in interprocedural analysis as studied in Section 2.5.

[con]	$(\widehat{C}, \widehat{\rho}) \models_{\delta}^{ce} c^{\ell}$ always
[var]	$(\widehat{C}, \widehat{\rho}) \models_{\delta}^{ce} x^{\ell}$ iff $\widehat{\rho}(x, ce(x)) \subseteq \widehat{C}(\ell, \delta)$
[fn]	$(\widehat{C}, \widehat{\rho}) \models_{\delta}^{ce} (\text{fn } x \Rightarrow e_0)^{\ell}$ iff $\{(\text{fn } x \Rightarrow e_0, ce_0)\} \subseteq \widehat{C}(\ell, \delta)$ where $ce_0 = ce \mid FV(\text{fn } x \Rightarrow e_0)$
[fun]	$(\widehat{C}, \widehat{\rho}) \models_{\delta}^{ce} (\text{fun } f \text{ } x \Rightarrow e_0)^{\ell}$ iff $\{(\text{fun } f \text{ } x \Rightarrow e_0, ce_0)\} \subseteq \widehat{C}(\ell, \delta)$ where $ce_0 = ce \mid FV(\text{fun } f \text{ } x \Rightarrow e_0)$
[app]	$(\widehat{C}, \widehat{\rho}) \models_{\delta}^{ce} (t_1^{\ell_1} t_2^{\ell_2})^{\ell}$ iff $(\widehat{C}, \widehat{\rho}) \models_{\delta}^{ce} t_1^{\ell_1} \wedge (\widehat{C}, \widehat{\rho}) \models_{\delta}^{ce} t_2^{\ell_2} \wedge$ $(\forall (\text{fn } x \Rightarrow t_0^{\ell_0}, ce_0) \in \widehat{C}(\ell_1, \delta) :$ $\quad (\widehat{C}, \widehat{\rho}) \models_{\delta_0}^{ce_0} t_0^{\ell_0} \wedge$ $\quad \widehat{C}(\ell_2, \delta) \subseteq \widehat{\rho}(x, \delta_0) \wedge \widehat{C}(\ell_0, \delta_0) \subseteq \widehat{C}(\ell, \delta)$ $\quad \text{where } \delta_0 = [\delta, \ell]_k$ $\quad \text{and } ce_0 = ce_0[x \mapsto \delta_0]) \wedge$ $(\forall (\text{fun } f \text{ } x \Rightarrow t_0^{\ell_0}, ce_0) \in \widehat{C}(\ell_1, \delta) :$ $\quad (\widehat{C}, \widehat{\rho}) \models_{\delta_0}^{ce_0} t_0^{\ell_0} \wedge$ $\quad \widehat{C}(\ell_2, \delta) \subseteq \widehat{\rho}(x, \delta_0) \wedge \widehat{C}(\ell_0, \delta_0) \subseteq \widehat{C}(\ell, \delta) \wedge$ $\quad \{(\text{fun } f \text{ } x \Rightarrow t_0^{\ell_0}, ce_0)\} \subseteq \widehat{\rho}(f, \delta_0)$ $\quad \text{where } \delta_0 = [\delta, \ell]_k$ $\quad \text{and } ce_0 = ce_0[f \mapsto \delta_0, x \mapsto \delta_0])$
[if]	$(\widehat{C}, \widehat{\rho}) \models_{\delta}^{ce} (\text{if } t_0^{\ell_0} \text{ then } t_1^{\ell_1} \text{ else } t_2^{\ell_2})^{\ell}$ iff $(\widehat{C}, \widehat{\rho}) \models_{\delta}^{ce} t_0^{\ell_0} \wedge (\widehat{C}, \widehat{\rho}) \models_{\delta}^{ce} t_1^{\ell_1} \wedge (\widehat{C}, \widehat{\rho}) \models_{\delta}^{ce} t_2^{\ell_2} \wedge$ $\widehat{C}(\ell_1, \delta) \subseteq \widehat{C}(\ell, \delta) \wedge \widehat{C}(\ell_2, \delta) \subseteq \widehat{C}(\ell, \delta)$
[let]	$(\widehat{C}, \widehat{\rho}) \models_{\delta}^{ce} (\text{let } x = t_1^{\ell_1} \text{ in } t_2^{\ell_2})^{\ell}$ iff $(\widehat{C}, \widehat{\rho}) \models_{\delta}^{ce} t_1^{\ell_1} \wedge (\widehat{C}, \widehat{\rho}) \models_{\delta}^{ce'} t_2^{\ell_2} \wedge$ $\widehat{C}(\ell_1, \delta) \subseteq \widehat{\rho}(x, \delta) \wedge \widehat{C}(\ell_2, \delta) \subseteq \widehat{C}(\ell, \delta)$ where $ce' = ce[x \mapsto \delta]$
[op]	$(\widehat{C}, \widehat{\rho}) \models_{\delta}^{ce} (t_1^{\ell_1} \text{ op } t_2^{\ell_2})^{\ell}$ iff $(\widehat{C}, \widehat{\rho}) \models_{\delta}^{ce} t_1^{\ell_1} \wedge (\widehat{C}, \widehat{\rho}) \models_{\delta}^{ce} t_2^{\ell_2}$

Table 3.10: Control Flow Analysis with Contexts.

3.6.1 Uniform k -CFA Analysis

Abstract domains. A key idea is to introduce context to distinguish between the various dynamic instances of variables and program points. There are many choices concerning how to model contexts and how they

can be modified in the course of the analysis. In a *uniform k-CFA* analysis (as well as in a *k-CFA* analysis) a context δ records the last k dynamic call points; hence in this case contexts will be sequences of labels of length at most k and they will be updated whenever a function application is analysed. This is modelled by taking:

$$\delta \in \Delta = \text{Lab}^{\leq k} \quad \text{context information}$$

Since the contexts will be used to distinguish between the various instances of the variables we will need a *context environment* to determine the context associated with the current instance of a variable:

$$ce \in \text{CEnv} = \text{Var} \rightarrow \Delta \quad \text{context environments}$$

The context environment will play a rôle similar to the environment of the semantics; in particular, this means that we shall extend the *abstract values* to contain a context environment:

$$\hat{v} \in \widehat{\text{Val}} = \mathcal{P}(\text{Term} \times \text{CEnv}) \quad \text{abstract values}$$

So in addition to recording the abstractions ($\text{fn } x \Rightarrow e$ and $\text{fun } f \ x \Rightarrow e$) we will also record the context environment at the *definition point* for the free variables of the term. This should be compared with the Structural Operational Semantics of Section 3.2 where the closures contain information about the abstraction as well as the environment determining the values of the free variables at the definition point.

The *abstract environment* $\hat{\rho}$ will now map a variable and a context to an abstract value:

$$\hat{\rho} \in \widehat{\text{Env}} = (\text{Var} \times \Delta) \rightarrow \widehat{\text{Val}} \quad \text{abstract environments}$$

Typically we will use a context environment to find the context associated with the variable of interest and then use it together with the variable to access the abstract environment. This means that indirectly we get the effect of having local abstract environments in the abstract values although $\hat{\rho}$ is still a global entity as in the previous sections.

The uniform *k-CFA* analysis differs from the *k-CFA* analysis in performing a similar development for the *abstract cache* that now maps a label and a context to an abstract value:

$$\hat{c} \in \widehat{\text{Cache}} = (\text{Lab} \times \Delta) \rightarrow \widehat{\text{Val}} \quad \text{abstract caches}$$

Given information about the context of interest we can determine the abstract value associated with a label. Again we indirectly get the effect of having a cache for each possible context although it is still a global entity. (In *k-CFA* one has $\widehat{\text{Cache}} = (\text{Lab} \times \text{CEnv}) \rightarrow \widehat{\text{Val}}$.)

Acceptability relation. The acceptability relation for *uniform k-CFA* is presented in Table 3.10. It is defined by formulae of the form

$$(\widehat{C}, \widehat{\rho}) \models_{\delta}^{ce} e$$

where ce is the current context environment and δ is the current context. The formula expresses that $(\widehat{C}, \widehat{\rho})$ is an acceptable analysis of e in the *context* specified by ce and δ . The clauses for the various constructs are very much as those in Table 3.1 and will be explained below.

In the clause *[var]* we use the current context environment ce to determine the context $ce(x)$ of the current instance of the variable x and then the abstract value of the variable is given by $\widehat{\rho}(x, ce(x))$. The current context is δ so we have to ensure that $\widehat{\rho}(x, ce(x)) \subseteq \widehat{C}(\ell, \delta)$.

In the clause *[fn]* we record the current context environment as part of the abstract value and (as in the Structural Operational Semantics of Table 3.2) we restrict the context environment to the set of variables of interest for the abstraction. The clause *[fun]* is similar.

In the clause *[app]* we analyse the two subexpressions using the same context and context environment as the composite expression. When we find a potential abstract value, say $(fn\ x \Rightarrow t_0^{l_0}, ce_0)$, that the operator may evaluate to, it will contain a local context environment ce_0 that was created at its definition point. When analysing $t_0^{l_0}$ we will have passed through the application point ℓ , so the current context will be updated to include ℓ and this will also be the context associated with the variable x in the updated version of the context environment ce_0 used for the analysis of $t_0^{l_0}$. The new context is $[\delta, \ell]_k$ which (as in Section 2.5) denotes the sequence $[\delta, \ell]$ but possibly truncated (by omitting elements on the left) so as have length at most k . In the case where the operator has the form $(fun\ f\ x \Rightarrow t_0^{l_0}, ce_0)$ we proceed in a similar way and note that f as well as x will be associated with the new context in the analysis of the body of the function.

The clauses for *[if]*, *[let]* and *[op]* are fairly straightforward modifications of those of Table 3.1; however, note that the context of the bound variable of the *let*-construct is the current context (as no application point is passed). – We shall dispense with proving the correctness of the analysis and with showing how it can be implemented.

Example 3.34 We shall now specify a uniform 1-CFA analysis for the expression of Example 3.33:

$$(\text{let } f = (\text{fn } x \Rightarrow x^1)^2 \text{ in } ((f^3\ f^4)^5 (\text{fn } y \Rightarrow y^6)^7)^8)^9$$

The initial context will be Λ , the empty sequence of labels. In the course of the analysis the current context will be modified at the two application points with labels 5 and 8; since we only records call strings of length at most

one the only contexts of interest will therefore be Λ , 5 and 8. There are four context environments of interest:

$ce_0 = []$	the initial (empty) context environment,
$ce_1 = ce_0[f \mapsto \Lambda]$	the context environment for the analysis of the body of the let -construct,
$ce_2 = ce_0[x \mapsto 5]$	the context environment used for the analysis of the body of f initiated at the application point 5, and
$ce_3 = ce_0[x \mapsto 8]$	the context environment used for the analysis of the body of f initiated at the application point 8.

Let us take \widehat{C}'_{id} and $\widetilde{\rho}'_{id}$ to be:

$$\begin{aligned}
\widehat{C}'_{id}(1, 5) &= \{(fn\ x \Rightarrow x^1, ce_0)\} & \widehat{C}'_{id}(1, 8) &= \{(fn\ y \Rightarrow y^6, ce_0)\} \\
\widehat{C}'_{id}(2, \Lambda) &= \{(fn\ x \Rightarrow x^1, ce_0)\} & \widehat{C}'_{id}(3, \Lambda) &= \{(fn\ x \Rightarrow x^1, ce_0)\} \\
\widehat{C}'_{id}(4, \Lambda) &= \{(fn\ x \Rightarrow x^1, ce_0)\} & \widehat{C}'_{id}(5, \Lambda) &= \{(fn\ x \Rightarrow x^1, ce_0)\} \\
\widehat{C}'_{id}(7, \Lambda) &= \{(fn\ y \Rightarrow y^6, ce_0)\} & \widehat{C}'_{id}(8, \Lambda) &= \{(fn\ y \Rightarrow y^6, ce_0)\} \\
\widehat{C}'_{id}(9, \Lambda) &= \{(fn\ y \Rightarrow y^6, ce_0)\} & & \\
\widetilde{\rho}'_{id}(f, \Lambda) &= \{(fn\ x \Rightarrow x^1, ce_0)\} & & \\
\widetilde{\rho}'_{id}(x, 5) &= \{(fn\ x \Rightarrow x^1, ce_0)\} & \widetilde{\rho}'_{id}(x, 8) &= \{(fn\ y \Rightarrow y^6, ce_0)\}
\end{aligned}$$

We shall now show that this is an acceptable analysis result for the example expressions:

$$(\widehat{C}'_{id}, \widetilde{\rho}'_{id}) \models_{\Lambda}^{ce_0} (\text{let } f = (fn\ x \Rightarrow x^1)^2 \text{ in } ((f^3\ f^4)^5 (fn\ y \Rightarrow y^6)^7)^8)^9$$

According to clause [let], it is sufficient to verify that

$$\begin{aligned}
(\widehat{C}'_{id}, \widetilde{\rho}'_{id}) &\models_{\Lambda}^{ce_0} (fn\ x \Rightarrow x^1)^2 \\
(\widehat{C}'_{id}, \widetilde{\rho}'_{id}) &\models_{\Lambda}^{ce_1} ((f^3\ f^4)^5 (fn\ y \Rightarrow y^6)^7)^8
\end{aligned}$$

because $\widehat{C}'_{id}(2, \Lambda) \subseteq \widetilde{\rho}'_{id}(f, \Lambda)$ and $\widehat{C}'_{id}(8, \Lambda) \subseteq \widehat{C}'_{id}(9, \Lambda)$. This is straightforward except for the last clause. Since $\widehat{C}'_{id}(5, \Lambda) = \{(fn\ x \Rightarrow x^1, ce_0)\}$ it is, according to [app], sufficient to verify that

$$\begin{aligned}
(\widehat{C}'_{id}, \widetilde{\rho}'_{id}) &\models_{\Lambda}^{ce_1} (f^3\ f^4)^5 \\
(\widehat{C}'_{id}, \widetilde{\rho}'_{id}) &\models_{\Lambda}^{ce_1} (fn\ y \Rightarrow y^6)^7 \\
(\widehat{C}'_{id}, \widetilde{\rho}'_{id}) &\models_8^{ce_3} x^1
\end{aligned}$$

because $\widehat{C}'_{id}(7, \Lambda) \subseteq \widetilde{\rho}'_{id}(x, 8)$ and $\widehat{C}'_{id}(1, 8) \subseteq \widehat{C}'_{id}(8, \Lambda)$. This is straightforward except for the first clause. Proceeding as above we see that $\widehat{C}'_{id}(3, \Lambda) = \{(fn\ x \Rightarrow x^1, ce_0)\}$ and it is sufficient to verify

$$(\widehat{C}'_{id}, \widehat{\rho}'_{id}) \models_{\Lambda}^{ce_1} f^3$$

$$(\widehat{C}'_{id}, \widehat{\rho}'_{id}) \models_{\Lambda}^{ce_1} f^4$$

$$(\widehat{C}'_{id}, \widehat{\rho}'_{id}) \models_5^{ce_2} x^1$$

because $\widehat{C}'_{id}(4, \Lambda) \subseteq \widehat{\rho}'_{id}(x, 5)$ and $\widehat{C}'_{id}(1, 5) \subseteq \widehat{C}'_{id}(5, \Lambda)$. This is straightforward.

The importance of this example is that it shows that the uniform 1-CFA analysis is strong enough to determine that $\text{fn } y \Rightarrow y^6$ is the *only* result possible for the overall expression unlike what was the case for the 0-CFA analysis in Example 3.32. We can also see that, since $\widehat{\rho}'_{id}(y, \delta) = \emptyset$ for all $\delta \in \{\Lambda, 5, 8\}$ it follows that $\text{fn } y \Rightarrow y^6$ is never called upon a function. ■

The resulting analysis will have exponential worst case complexity even for the case where $k = 1$. To see this assume that the expression has size n and that it has p different variables. Then Δ has $O(n)$ elements and hence there will be $O(p \cdot n)$ different pairs (x, δ) and $O(n^2)$ different pairs (ℓ, δ) . This means that $(\widehat{C}, \widehat{\rho})$ can be seen as an $O(n^2)$ tuple of values from $\widehat{\text{Val}}$. Since $\widehat{\text{Val}}$ itself is a powerset of pairs of the form (t, ce) and there are $O(n \cdot n^p)$ such pairs it follows that $\widehat{\text{Val}}$ has height $O(n \cdot n^p)$. Since $p = O(n)$ we have the exponential worst case complexity claimed above.

This should be contrasted with the 0-CFA analysis developed in the previous sections. It corresponds to letting Δ be a singleton. Repeating the above calculations we can see $(\widehat{C}, \widehat{\rho})$ as an $O(p+n)$ tuple of values from $\widehat{\text{Val}}$, and $\widehat{\text{Val}}$ will be a lattice of height $O(n)$. In total this gives us a polynomial analysis as we already saw in Section 3.4.

The worst case complexity of the uniform k -CFA analysis (as well as the k -CFA analysis) can be improved in different ways. One possibility is to reduce the height of the lattice $\widehat{\text{Val}}$ using the techniques of Chapter 4. Another possibility is to replace all context environments with contexts, i.e. to have $\widehat{\text{Val}} = \mathcal{P}(\text{Term} \times \Delta)$; clearly this will give a lattice of polynomial height. This idea is closely related to the so-called *polynomial k -CFA* analysis where the analogues of context environments are forced to be constant functions, i.e. to map all variables to the same context. In the case of polynomial 1-CFA the analysis is of complexity $O(n^6)$.

3.6.2 Interprocedural Analysis Revisited

Let us conclude this section by comparing the above development with that of Section 2.5 where we considered *interprocedural analysis* for a simple imperative procedure language.

Recall that the abstract domain of interest in Section 2.5 has the form

$$\Delta \rightarrow L$$

where Δ is the context information and L is the complete lattice of abstract values of interest. For each label ℓ the analysis will determine two elements $A_o(\ell)$ and $A_a(\ell)$ of $\Delta \rightarrow L$ describing the situation before and after the elementary block labelled ℓ is executed. So we have

$$A_o, A_a : \text{Lab} \rightarrow (\Delta \rightarrow L)$$

and in the terminology of the present chapter we may regard these functions as abstract caches. There is no analogue of the abstract environment in Section 2.5 – the reason is that the procedure language is so simple that it is not needed: the abstract environment records the context of the free variables and since all free variables in the procedures are global variables there is no need for this component.

We can now reformulate the development of this section as follows. We can take the abstract domain of interest to be

$$\Delta \rightarrow \mathcal{P}(\text{Term} \times \text{CEnv})$$

and reformulate the abstract cache and the abstract environment as having the functionalities:

$$\begin{aligned} \hat{C} : \text{Lab} &\rightarrow \Delta \rightarrow \mathcal{P}(\text{Term} \times \text{CEnv}) \\ \hat{\rho} : \text{Var} &\rightarrow \Delta \rightarrow \mathcal{P}(\text{Term} \times \text{CEnv}) \end{aligned}$$

Thus the abstract caches of the interprocedural analysis and the uniform k -CFA analysis have the same overall functionality.

Concluding Remarks

Control Flow Analysis for functional languages. Many of the key ideas for Control Flow Analysis have been developed within the context of functional languages. The concept of k -CFA analysis seems due to Shivers [121, 122, 123] although the main focus was on 0-CFA. Other works on 0-CFA-like analyses include [128, 104, 41, 40]. The ideas behind k -CFA and *polynomial k -CFA* analysis were further clarified in [59] that also established the exponential complexity of k -CFA analysis for $k > 0$; it also related a form of *Set Based Analysis* [53] to 0-CFA. The *uniform k -CFA* analyses were introduced in [94] as a simplification of the k -CFA analyses; an obvious variation over this is to record the set of the last k distinct call points. Yet another variation over the same theme is *Closure Analysis*; an early and often neglected development may be found in [118].

To the extent these developments go beyond 0-CFA they establish additional context (called mementoes, tokens or contours) for representing information concerning the *dynamic* call chain; this clearly links back to the use of call

strings in Section 2.5. The formulation of the analyses (as well as the one presented in Table 3.1) would often seem to be more appropriate for a dynamically scoped than for a statically scoped language: The 0-CFA analysis coalesces information about variables having several defining occurrences even if they differ in their scope; clearly the analysis can easily be modified so that it more directly models static scope (see Exercise 3.7) rather than relying on no variable having more than one defining occurrence.

Another way to establish context is to represent the *static* call chain. This seems first to be described by [60] as part of their so-called “polymorphic splitting” analysis. A more general set-up was formulated in [94] that also argued for the need to base abstract specifications on coinductive methods – bearing in mind that coinductive and inductive methods may coincide as in the case of syntax directed specifications.

Most of the papers cited above directly formulate a syntax directed specification, perhaps proving it semantically sound, and perhaps showing how to obtain constraints so as to obtain an efficient implementation. The use of abstract specifications first appeared in [60, 94] and has the advantage of being more directly applicable to *open systems* (that allow to interface with the library routines provided by the environment) and also to the ideas of Abstract Interpretation of Chapter 4. In particular, the notion of reachability suggests itself rather naturally [14, 44], it becomes clearer how to integrate ideas from Abstract Interpretation into Control Flow Analysis, and one does not inadvertently restrict oneself to *closed systems* only. (The notion of reachability is considered in Mini Project 3.1 which is based on [44].)

Only few papers [96] discuss the interplay between the choice of specification style for the analysis and the choice of semantics. We have used a small-step Structural Operational Semantics rather than a big-step semantics in order to express the semantic correctness also of looping programs. We have used an environment based semantics in order to ensure that we do not “modify” the bodies of functions before they are called, so that function abstractions can meaningfully be used in the value domains of our analysis [96]. As a consequence we have had to introduce intermediate expressions (closures and bindings) and have had to specify the abstract analysis also for intermediate expressions; for the syntax directed specification and the constraint based analysis this was not necessary given that semantic correctness had already been dealt with. Alternative choices are clearly possible but are likely to sacrifice at least some of the generality offered by the present approach.

Control Flow Analysis for other language paradigms. Another main application of Control Flow Analysis has been for *object-oriented languages*: one simply tracks objects rather than functions [1, 95, 105]. As a reminder of the close links between Data Flow Analysis and Control Flow Analysis we should also mention that some approaches [139, 107] are closer to the presentation of Chapter 2. A common theme among the more advanced

studies is the incorporation of context (related to k -CFA) and an abstract store [102] (to deal with imperative aspects like method update). To increase the precision, local versions of the abstract store need to exist at all program points, and abstract reference counts are needed to incorporate a “kill” component (in the manner of Chapter 2). We refer to the above literature for further details of how to formulate such analyses and how to choose a proper balance between precision and cost.

Control Flow Analysis for *concurrent languages* has received relatively little attention [15]. However, variations of the techniques presented in this chapter have been used to analyse functional languages extended with concurrency primitives allowing processes to be created dynamically and to communicate via shared locations or channels [58, 41, 44].

In this book we do not consider *logic languages*. However, we should point out that Control Flow Analysis also has applications for logic languages and that set based analysis was first developed for this class of languages [55, 56].

Set-Constraint Based Analysis. Control Flow Analysis is just one approach to program analysis where the use of constraints pays off. In this chapter we have taken the following approach: (i) first we have given an abstract specification of when a proposed solution is acceptable, (ii) then we have developed an algorithm for generating a set of constraints expressing that a proposed solution is acceptable, and (iii) finally we have solved the set of constraints for the least solution. For the solution of set constraints in step (iii), it is unimportant how the constraints were in fact obtained. For this reason, it is often said that set constraints allow to separate the specification of an analysis from its implementation and that set constraints are able to deal with forward analyses as well as backward analyses and indeed mixtures of these.

Set constraints [55, 9] have a long history [110, 64]. They allow to express general inclusions of the form

$$S_1 \subseteq S_2$$

where set expressions, S , set variables, V , and set constructors, C , may be built as follows:

$$S ::= V \mid \emptyset \mid S_1 \cup S_2 \mid S_1 \cap S_2 \mid C(S_1, \dots, S_n) \\ \mid (S_1 \subseteq S_2) \Rightarrow S_2 \mid (S_1 \neq \emptyset) \Rightarrow S_2 \mid C^{-i}(S) \mid \neg S \mid \dots$$

$$V ::= X \mid Y \mid \dots$$

$$C ::= \text{true} \mid \text{false} \mid 0 \mid \dots \mid \text{cons} \mid \text{nil} \mid \dots$$

Set constraints allow to consider constructors that are not just nullary and this allows to record also the shape of data structures, e.g. $\text{cons}(S_1, S_2)\text{Unil}$ expresses possibly empty lists whose heads come from S_1 and whose tails come from S_2 . The associated projection selects those terms (if any) having

the required shape, e.g. $\text{cons}^{-1}(S)$ produces the heads that may be present in S . We have seen conditional constraints before and it turns out that projection is so powerful that it can be used to code conditional constraints. Finally, it is sometimes possible to explicitly take the complement of a set but this adds to the complexity of the development. (It means that solutions can no longer be guaranteed using Tarski's Theorem and sometimes a version of Banach's Theorem can be used instead.)

The complexity of solving a system of set constraints depends rather dramatically on the set forming operations allowed and therefore many versions have been considered in the literature. We refer to [3, 108] for an overview of what is known in this area; here we just mention [19] for a general result and [53, 5] for some cubic time fragments.

However, it is worth pointing out that many of these results are worst-case results; benchmark results of Jaganathan and Wright [60] shows e.g. that in practice a 1-CFA analysis may be faster than a 0-CFA analysis despite the fact that the former has exponential worst-case complexity and the latter polynomial worst-case complexity. The reason seems to be that the 0-CFA analysis explores most of its polynomial sized state space whereas the 1-CFA analysis is so precise that it only explores a fraction of its exponentially sized state space.

The basic idea behind many of the solution procedures for set constraints is roughly as follows [4]:

1. Dynamically expand conditional constraints, based on the condition being fulfilled, until no more expansion is possible.
2. Remove all conditional constraints and combine the remaining constraints to obtain the least solution.

This is not quite the algorithm used in Section 3.4 where we were only interested in solving a rather limited class of constraints for 0-CFA analysis. When generating the constraints in Table 3.6 we were able to "guess a universe" Term_* that was sufficiently large and this allowed us to generate explicit versions of the conditional constraints; in fact a superset of all those to be considered in step 1 of the above algorithm. Therefore our subsequent constraint solving algorithm in Table 3.7 merely needed to check the already existing constraints and to determine whether or not they could contribute to the solution. In practice, the above "lazy" algorithm is likely to perform better than the "eager" algorithm of Tables 3.6 and 3.7.

A final note, to be expanded upon in Chapter 6, is that also state-of-the-art algorithms for Data Flow Analysis work by internally regarding data flow equations as presented in a constraint formulation. This is yet more evidence of the close connections between Data Flow Analysis, Control Flow Analysis and Set-Constraint Based Analysis.

Mini Projects

Mini Project 3.1 Reachability Analysis

The syntax directed analysis of Table 3.5 analyses each subexpression of e_* “exactly once” rather than “at most once” as really called for. In this mini project we shall study one way to amend this.

The idea is to introduce an *abstract reachability component*

$$\widehat{R} \in \mathbf{Reach} = \mathbf{Lab} \rightarrow \mathcal{P}(\{\mathbf{on}\})$$

and to modify the syntax directed analysis to have a relation of the form

$$(\widehat{C}, \widehat{\rho}, \widehat{R}) \models'_s e$$

The idea is that $\text{fn } x \Rightarrow t_0^{\ell_0}$ has $\{\mathbf{on}\} \subseteq \widehat{R}(\ell_0)$ if and only if the function is indeed applied somewhere, and that the “recursive call” $(\widehat{C}, \widehat{\rho}, \widehat{R}) \models_s t_0^{\ell_0}$ is performed if and only if $\{\mathbf{on}\} \subseteq \widehat{R}(\ell_0)$.

1. Modify Table 3.5 to incorporate this idea.
2. Show the following analogue of Proposition 3.18: If $(\widehat{C}, \widehat{\rho}, \widehat{R}) \models'_s t_*^{\ell_*}$, $\{\mathbf{on}\} \subseteq \widehat{R}(\ell_*)$ and $(\widehat{C}, \widehat{\rho}) \sqsubseteq (\widehat{C}_*^T, \widehat{\rho}_*^T)$ then $(\widehat{C}, \widehat{\rho}) \models t_*^{\ell_*}$.
3. Determine whether or not the statements
 - if $(\widehat{C}, \widehat{\rho}) \models e_*$ then $(\widehat{C}, \widehat{\rho}, \widehat{R}) \models'_s e_*$ for some \widehat{R}
 - if $(\widehat{C}, \widehat{\rho}) \models e_*$ and $(\widehat{C}, \widehat{\rho}) \sqsubseteq (\widehat{C}_*^T, \widehat{\rho}_*^T)$ then $(\widehat{C}, \widehat{\rho}, \widehat{R}) \models'_s e_*$ for some \widehat{R}
 hold in general. ■

Mini Project 3.2 Data Structures

The language considered so far only includes simple data like integers and booleans. In this mini project we shall extend the language with more general data structures:

$$e ::= \dots \mid C(e_1, \dots, e_n)^{\ell} \mid (\text{case } e_0 \text{ of } C(x_1, \dots, x_n) \Rightarrow e_1 \text{ or } x \Rightarrow e_2)^{\ell}$$

Here $C \in \mathbf{Constr}$ denotes an n -ary data constructor. A data element is constructed by $C(e_1, \dots, e_n)$: it has tag C and its components are the values of e_1, \dots, e_n . The case construct will first determine the value v_0 of e_0 , if v_0 has the tag C then x_1, \dots, x_n will be bound to the components of v_0 and e_1 is evaluated. If v_0 does not have tag C then x is bound to v_0 and e_2 is evaluated.

As an example we may have $\mathbf{Constr} = \{\text{cons}, \text{nil}\}$ so we have the following expression (omitting labels) for reversing a list:


```

let append = fun app xs => fn ys =>
    case xs of cons(z,zs) => cons(z,app zs ys)
    or xs => ys
in fun rev xs => case xs of cons(y,ys) =>
    append (rev ys) (cons(y,nil()))
    or xs => nil()

```

To specify a 0-CFA analysis for this language we shall take

$$\widehat{\mathbf{Val}} = \mathcal{P}(\mathbf{Term} \cup \{C(\ell_1, \dots, \ell_n) \mid C \in \mathbf{Constr}, \ell_1, \dots, \ell_n \in \mathbf{Lab}\})$$

As before the terms of interest are $\text{fn } x \Rightarrow e_0$ and $\text{fun } f \ x \Rightarrow e_0$ for recording the abstractions. The new contribution is a number of elements of the form $C(\ell_1, \dots, \ell_n)$ denoting a data element $C(v_1, \dots, v_n)$ whose i 'th component might have been created by the expression at program point ℓ_i .

1. Develop a syntax directed analogue of the analysis in Table 3.5.
2. Modify the constraint generation algorithm of Table 3.6 to handle the new constructs and make the necessary changes to the constraint solving algorithm of Table 3.7.

For the more ambitious: are there any difficulties in developing an abstract analogue of the analysis in Table 3.1? ■

Mini Project 3.3 A Prototype Implementation

In this mini project we shall implement the pure 0-CFA analysis considered in Section 3.3. As implementation language we shall choose a functional language such as Standard ML or Haskell. We can then define a suitable data type for FUN expressions as follows:

```

type var      = string
type label    = int
datatype const = Num of int | True | False
datatype exp  = Label of term * label
and term      = Const of const | Var of var
              | Fn of var * exp | Fun of var * var * exp
              | App of exp * exp | If of exp * exp * exp
              | Let of var * exp * exp | Op of string * exp * exp

```

Now proceed as follows:

1. Implement the constraint based control flow analysis of Section 3.4; this includes defining an appropriate data structure *constraints* for (conditional) constraints.
2. Implement the graph based algorithm of Section 3.4 for solving constraints; this involves choosing appropriate data structures for the worklist and the two arrays used by the algorithm.

For the more ambitious: generalise your program to perform some of the more advanced analyses, e.g. by incorporating data flow information or context information. ■

Exercises

Exercise 3.1 Consider the following expression (omitting labels):

```
let f = fn x => x 1
in let g = fn y => y+2
   in let h = fn z => z+3
      in (f g) + (f h)
```

Add labels to the program and guess an analysis result. Use Table 3.1 to verify that it is indeed an acceptable guess. ■

Exercise 3.2 The specification of the Control Flow Analysis in Table 3.1 uses potentially infinite value spaces and this is not really necessary. To see this choose some expression $e_* \in \mathbf{Exp}$ that is to be analysed. Let $\mathbf{Var}_* \subseteq \mathbf{Var}$ be the finite set of variables occurring in e_* , let $\mathbf{Lab}_* \subseteq \mathbf{Lab}$ be the finite set of labels occurring in e_* , and let \mathbf{Term}_* be the finite set of subterms of e_* . Next define

$$\begin{aligned} \widehat{v} &\in \widehat{\mathbf{Val}}_* &= \mathcal{P}(\mathbf{Term}_*) \\ \widehat{\rho} &\in \widehat{\mathbf{Env}}_* &= \mathbf{Var}_* \rightarrow \widehat{\mathbf{Val}}_* \\ \widehat{C} &\in \widehat{\mathbf{Cache}}_* &= \mathbf{Lab}_* \rightarrow \widehat{\mathbf{Val}}_* \end{aligned}$$

and note that these value spaces are finite. Show that the specification of the analysis in Table 3.1 still makes sense when $(\widehat{C}, \widehat{\rho})$ is restricted to be in $\widehat{\mathbf{Cache}}_* \times \widehat{\mathbf{Env}}_*$. ■

Exercise 3.3 Modify the Control Flow Analysis of Table 3.1 to take account of the left to right evaluation order imposed by a call-by-value semantics: in the clause $[app]$ there is no need to analyse the operand if the operator cannot produce any closures. Try to find a program where the modified analysis accepts analysis results $(\widehat{C}, \widehat{\rho})$ rejected by Table 3.1. ■

Exercise 3.4 So far we have defined “ $(\widehat{C}, \widehat{\rho})$ is an acceptable solution for e ” to mean that

$$(\widehat{C}, \widehat{\rho}) \models e \quad (3.8)$$

but an alternative condition is that

$$\exists (\widehat{C}', \widehat{\rho}') : (\widehat{C}', \widehat{\rho}') \models e \wedge (\widehat{C}', \widehat{\rho}') \sqsubseteq (\widehat{C}, \widehat{\rho}) \quad (3.9)$$

Show that (3.8) implies (3.9) but not vice versa. Discuss which of (3.8) or (3.9) is the preferable definition. ■

Exercise 3.5 Consider an alternative specification of the analysis in Table 3.1 where the condition

$$(\text{fun } f \ x \Rightarrow t_0^{\ell_0}) \in \widehat{\rho}(f)$$

in [app] is replaced by

$$\widehat{C}(\ell_1) \subseteq \widehat{\rho}(f)$$

also in [app]. Show that the proof of Theorem 3.10 can be modified accordingly. Discuss the relative precision of the two analyses. ■

Exercise 3.6 Reconsider our decision to use $\widehat{\text{Val}} = \mathcal{P}(\text{Term})$ and consider using $\widehat{\text{Val}} = \mathcal{P}(\text{Exp})$ instead. Show that the specification of the Control Flow Analysis may be modified accordingly but that then Fact 3.11 (and hence the correctness result) would fail. ■

Exercise 3.7 The operational semantics allow us to rename bound variables without changing the semantics; this is in accord with the language being statically scoped (or lexically scoped) rather than dynamically scoped. As an example

$$((\text{fn } x \Rightarrow x^1)^2 (\text{fn } y \Rightarrow y^3)^4)^5 =_{\alpha} ((\text{fn } x \Rightarrow x^1)^2 (\text{fn } x \Rightarrow x^3)^4)^5$$

and clearly the two programs have the same semantics.

However, renaming bound variables changes the acceptability of a solution as well as influences the precision of the analysis specified in Table 3.1. Develop an abstract specification of a 0-CFA analysis that is more faithful to the static scoping than that of Table 3.1; it should agree with the specification of Table 3.1 for expressions that do not have multiple defining occurrences. ■

Exercise 3.8 In Section 3.2 we equipped FUN with a call-by-value semantics. An alternative would be to use a *call-by-name* semantics. It can be obtained as a simple modification of the semantics of Tables 3.2 and 3.3 by

allowing the environments $\rho \in \mathbf{Env}$ to map variables to intermediate terms (and not just values), by deleting the rules $[app_2]$ and $[let_2]$ and then make some obvious modifications to the axioms $[var]$, $[app_{fn}]$, $[app_{fun}]$ and $[let_1]$; in the case of $[var]$ we will take:

$$\rho \vdash x^l \rightarrow it^l \quad \text{if } x \in \text{dom}(\rho) \text{ and } it = \rho(x)$$

Complete the specification of the call-by-name semantics of FUN and show that the correctness result (Theorem 3.10) still holds for the analysis of Table 3.1.

What does that tell us about the precision of the analysis? ■

Exercise 3.9 Let \models'_s and \models''_s be two relations satisfying the specification of Table 3.5. Show that

$$(\widehat{C}, \widehat{\rho}) \models'_s e \text{ iff } (\widehat{C}, \widehat{\rho}) \models''_s e$$

by structural induction on e . ■

Exercise 3.10 Consider Proposition 3.18 and determine whether or not the statement

$$\text{if } (\widehat{C}, \widehat{\rho}) \models_s e_* \text{ then } (\widehat{C}, \widehat{\rho}) \models e_*$$

holds in general. ■

Exercise 3.11 Give an example showing that both of the statements

$$\text{if } (\widehat{C}, \widehat{\rho}) \models e_* \text{ then } (\widehat{C}, \widehat{\rho}) \models_s e_*$$

$$\text{if } (\widehat{C}, \widehat{\rho}) \models e_* \text{ and } (\widehat{C}, \widehat{\rho}) \sqsubseteq (\widehat{C}_*, \widehat{\rho}_*) \text{ then } (\widehat{C}, \widehat{\rho}) \models_s e_*$$

fail in general. ■

Exercise 3.12 Give a direct proof of the correctness of the syntax directed analysis of Table 3.5, i.e. establish an analogue of Theorem 3.10. This involves first extending the syntax directed analysis to the bind- and close-constructs and next proving that if $\rho \mathcal{R} \widehat{\rho}$, $\rho \vdash ie \rightarrow ie'$ and $(\widehat{C}, \widehat{\rho}) \models_s ie$ then also $(\widehat{C}, \widehat{\rho}) \models_s ie'$. ■

Exercise 3.13 Consider the system $\mathcal{C}_*^-[e_*]$ that contains a constraint

$$ls_1 \cup \dots \cup ls_n = rhs$$

whenever $\mathcal{C}_*[e_*]$ contains the $n \geq 1$ constraints

$$ls_i \subseteq rhs$$

Show that

$$(\widehat{C}, \widehat{\rho}) \models_c C_*^-[e_*] \text{ implies } (\widehat{C}, \widehat{\rho}) \models_c C_*[e_*]$$

where $(\widehat{C}, \widehat{\rho}) \models_c (ls = rhs)$ is defined in the obvious way. Also show that

$$(\widehat{C}, \widehat{\rho}) \models_c C_*[e_*] \text{ implies } (\widehat{C}, \widehat{\rho}) \models_c C_*^-[e_*]$$

holds in the special case where $(\widehat{C}, \widehat{\rho})$ is *least* such that $(\widehat{C}, \widehat{\rho}) \models_c C_*[e_*]$. ■

Exercise 3.14 Use the ideas of Exercise 3.3 to develop an improvement of Table 3.8 where expressions are only analysed when absolutely needed. Next develop a syntax directed analysis using the same ideas. Discuss the relationship between the two specifications: are they more closely related than is the case for \models and \models_s of Table 3.1 and 3.5 (see Exercises 3.10 and 3.11)? ■

Exercise 3.15 Modify the abstract specification of the uniform k -CFA analysis so that it does not record the last k function calls but the last k *changes* function calls: if the calling sequence is $[1,2,2,1,1]$ then 2-CFA records $[1,1]$ but the modified analysis records $[2,1]$. Discuss which of the two analyses (say for $k = 2$) is likely to be most useful in practice. ■

Exercise 3.16 Let us consider a language of first-order recursion equation schemes: the programs have the form

define D_* in e_*

where D_* is a sequence of function definitions of the form:

$$f(x) = e$$

Here f is a function name, x is the formal parameter and e is the body of the function; the functions defined in D may be mutually recursive and the parameter mechanism is call-by-value. The expressions are given by

$$\begin{aligned} e &::= t^t \\ t &::= c \mid x \mid f e \mid \text{if } e_0 \text{ then } e_1 \text{ else } e_2 \mid e_1 \text{ op } e_2 \end{aligned}$$

where $c \in \mathbf{Const}$ and $op \in \mathbf{Op}$ as before; we shall assume that f and x belong to distinct syntactic categories. As an example we may define the Fibonacci function by the following expression (omitting labels):

```
define fib(z) = if z<3 then 0
                else fib (z-1) + fib (z-2)
in      fib x
```

Define a uniform k -CFA analysis for this language. For $k = 0$ and $k = 1$ compare the development with that for the procedure language in Section 2.5. ■

Exercise* 3.17 In Section 3.6 we showed how to extend the Control Flow Analysis with contest similar to the call strings of Section 2.5. Investigate the possibility of performing a similar development based on the assumption sets of Section 2.5. ■

Chapter 4

Abstract Interpretation

The purpose of this chapter is to convey some of the essential ideas of Abstract Interpretation. We shall mainly do so in a programming language independent way and thus focus on the design of the property spaces, the functions and computations upon them, and the relationships between them.

We first formulate a notion of correctness for a restricted class of analyses as this will allow us to motivate better some of the key definitions in the development. Then we cover the widening and narrowing techniques that can be used to obtain approximations of the least fixed point and for limiting the number of computation steps needed. Next we consider Galois connections and Galois insertions that allow a costly space of properties to be replaced with one that is less costly. Galois connections can be constructed in a systematic way and can be used to induce one specification of an analysis from another.

4.1 A Mundane Approach to Correctness

To set the scene, imagine some programming language. Its *semantics* identifies some set V of values (like states, closures, double precision reals) and specifies how a program p transforms one value v_1 to another v_2 ; we may write

$$p \vdash v_1 \rightsquigarrow v_2 \quad (4.1)$$

for this without committing ourselves to the details of the semantics and without necessarily imposing determinacy (that $p \vdash v_1 \rightsquigarrow v_2$ and $p \vdash v_1 \rightsquigarrow v_3$ imply $v_2 = v_3$).

In a similar way, a *program analysis* identifies the set L of properties (like shapes of states, abstract closures, lower and upper bounds for reals) and

specifies how a program p transforms one property l_1 to another l_2 ; we may write

$$p \vdash l_1 \triangleright l_2 \quad (4.2)$$

for this without committing ourselves to the method used for specification of the analysis. However, unlike what is the case for the semantics, it is customary to require \triangleright to be deterministic and thereby define a function; this will allow us to write $f_p(l_1) = l_2$ to mean $p \vdash l_1 \triangleright l_2$.

In the rest of this section we shall show how to relate the semantics to the analysis. We shall present two approaches based on correctness relations and representation functions, respectively. In both cases we shall define a notion of correctness of the analysis with respect to the semantics and we shall show that the two notions are equivalent.

This is a mundane approach in the sense that it only applies to analyses where properties directly describe sets of values. This is the case for the Constant Propagation Analysis of Section 2.3, the Shape Analysis of Section 2.6 and the Control Flow Analysis of Chapter 3 but it is not the case for the Live Variable Analysis of Section 2.1 where properties are related to relations between values. In the literature, the terms *first-order analyses* versus *second-order analyses* have been used to differentiate between these classes of analyses. It is important to stress that the development of Sections 4.2 to 4.5 apply equally well to both classes.

We begin by showing how the development of Chapters 2 and 3 can be rephrased in the style of (4.1) and (4.2).

Example 4.1 Consider the WHILE language of Chapter 2. Recall that the semantics is a Structural Operational Semantics with transitions of the forms $\langle S, \sigma \rangle \rightarrow \langle S', \sigma' \rangle$ and $\langle S, \sigma \rangle \rightarrow \sigma'$, where S and S' are statements of **Stmt** and σ and σ' are states of **State** = $\text{Var} \rightarrow \mathbf{Z}$. With S_* being the program of interest we shall now write

$$S_* \vdash \sigma_1 \rightsquigarrow \sigma_2$$

for the reflexive transitive closure of the transition relation, i.e. for:

$$\langle S_*, \sigma_1 \rangle \rightarrow^* \sigma_2$$

Note that the set V of values is the set **State**.

We shall now consider the *Constant Propagation Analysis* of Section 2.3. Recall that the analysis of S_* gives rise to a set of equations CP^\perp formulated in terms of an instance of a Monotone Framework: the properties L of interest are given by $\widehat{\text{State}}_{\text{CP}} = (\text{Var}_* \rightarrow \mathbf{Z}^\top)_\perp$, E is $\{\text{init}(S_*)\}$, F is $\text{flow}(S_*)$, and ι is $\lambda x. \top$. Further recall that a solution to the equations is a pair $(\text{CP}_\circ, \text{CP}_\bullet)$

of mappings $CP_o : \mathbf{Lab}_* \rightarrow \widehat{\mathbf{State}}_{CP}$ and $CP_e : \mathbf{Lab}_* \rightarrow \widehat{\mathbf{State}}_{CP}$ satisfying the equations. Given a solution (CP_o, CP_e) to $CP^=$ we take

$$S_* \vdash \widehat{\sigma}_1 \triangleright \widehat{\sigma}_2$$

to mean that:

$$\iota = \widehat{\sigma}_1 \wedge \widehat{\sigma}_2 = \bigsqcup \{CP_e(\ell) \mid \ell \in \mathit{final}(S_*)\}$$

Thus for a program S_* with isolated entries, $\widehat{\sigma}_1$ is the abstract state associated with the entry point of S_* , and $\widehat{\sigma}_2$ is the abstract state associated with the exit points; we use the least upper bound operation (\bigsqcup) on the complete lattice $\widehat{\mathbf{State}}_{CP}$ to combine the contributions from the (possibly several) exit points of S_* . ■

Example 4.2 Consider the FUN language of Chapter 3. Recall that the semantics is given by a Structural Operational Semantics with transitions of the form $\rho \vdash ie \rightarrow ie'$ where ρ is an environment of $\mathbf{Env} = \mathbf{Var} \rightarrow_{\text{fin}} \mathbf{Val}$ and ie and ie' are intermediate expressions of \mathbf{IExp} . Let now e_* be the closed expression of interest. We shall write

$$e_* \vdash v_1 \rightsquigarrow v_2$$

to mean that e_* when given the argument v_1 will evaluate to the value v_2 , i.e. that

$$[] \vdash (e_* v_1^{\ell_1})^{\ell_2} \rightarrow^* v_2^{\ell_2}$$

where ℓ_1 and ℓ_2 are fresh labels. Note that the set V of values now is the set \mathbf{Val} .

We shall next consider the pure *Control Flow Analysis* of Section 3.1. Recall that the result of analysing the expression e_* is a pair $(\widehat{C}, \widehat{\rho})$ satisfying $(\widehat{C}, \widehat{\rho}) \models e_*$ as defined in Table 3.1. Here \widehat{C} is an element of $\widehat{\mathbf{Cache}} = \mathbf{Lab}_* \rightarrow \mathbf{Val}$ and $\widehat{\rho}$ is an element of $\widehat{\mathbf{Env}} = \mathbf{Var}_* \rightarrow \mathbf{Val}$ where $\mathbf{Val} = \mathcal{P}(\mathbf{Term}_*)$. For this analysis we shall take the properties L of interest to be pairs $(\widehat{\rho}, \widehat{v})$ of $\widehat{\mathbf{Env}} \times \mathbf{Val}$ and assume that $(\widehat{C}, \widehat{\rho}) \models (e_* c^{\ell_1})^{\ell_2}$ for some constant c . Then we define

$$e_* \vdash (\widehat{\rho}_1, \widehat{v}_1) \triangleright (\widehat{\rho}_2, \widehat{v}_2)$$

to mean that when e_* is given an argument with property $(\widehat{\rho}_1, \widehat{v}_1)$ then the result of the application will have property $(\widehat{\rho}_2, \widehat{v}_2)$:

$$\widehat{C}(\ell_1) = \widehat{v}_1 \wedge \widehat{C}(\ell_2) = \widehat{v}_2 \wedge \widehat{\rho}_1 = \widehat{\rho}_2 = \widehat{\rho}$$

Note that the “dummy” constant c used as an argument to e_* is used as a place holder for all potential arguments being described by \widehat{v}_1 ; for this idea to work it is important that the analysis of c puts no constraints on $(\widehat{C}, \widehat{\rho})$ as in indeed the case for the specification in Table 3.1. ■

4.1.1 Correctness Relations

Every program analysis should be correct with respect to the semantics. For a class of (so-called first-order) program analyses this is established by directly relating properties to values using a *correctness relation*:

$$R : V \times L \rightarrow \{\text{true}, \text{false}\}$$

The intention is that $v R l$ formalises our claim that the value v is described by the property l .

Correctness formulation. To be useful one has to prove that the correctness relation R is preserved under computation: if the relation holds between the initial value and the initial property then it also holds between the final value and the final property. This may be formulated as the implication

$$v_1 R l_1 \wedge p \vdash v_1 \rightsquigarrow v_2 \wedge p \vdash l_1 \triangleright l_2 \Rightarrow v_2 R l_2 \quad (4.3)$$

and is also expressed by the following diagram:

$$\begin{array}{ccccc} p \vdash & v_1 & \rightsquigarrow & v_2 & \\ & \vdots & & \vdots & \\ & R & \Rightarrow & R & \\ & \vdots & & \vdots & \\ p \vdash & l_1 & \triangleright & l_2 & \end{array}$$

A relation R satisfying a condition like this is often called a *logical relation* and the implication is sometimes written $(p \vdash \cdot \rightsquigarrow \cdot)(R \rightarrow R)(p \vdash \cdot \triangleright \cdot)$.

The theory of Abstract Interpretation comes to life when we augment the set of properties L with a preorder structure and relate this to the correctness relation R . The most common scenario is when $L = (L, \sqsubseteq, \sqcup, \sqcap, \perp, \top)$ is a complete lattice with partial ordering \sqsubseteq (see Appendix A). We then impose the following relationship between R and L :

$$v R l_1 \wedge l_1 \sqsubseteq l_2 \Rightarrow v R l_2 \quad (4.4)$$

$$(\forall l \in L' \subseteq L : v R l) \Rightarrow v R (\bigsqcap L') \quad (4.5)$$

Condition (4.4) says that the smaller the property is with respect to the partial ordering, the better (i.e. more precise) it is. This is an "arbitrary" decision in the sense that we could instead have decided that the larger the property is, the better it is, as is indeed the case in much of the literature on

Data Flow Analysis; luckily the principle of duality from lattice theory (see the Concluding Remarks) tells us that this difference is only a cosmetic one.

Condition (4.5) says that there is always a best property for describing a value. This is important for having to perform only one analysis (using the best property, i.e. the greatest lower bound of the candidates) instead of several analyses (one for each of the candidates). Recall from Appendix A that a subset Y of L is a *Moore family* if and only if $(\bigcap Y') \in Y$ for all subsets Y' of Y . We can then see that condition (4.5) is equivalent to the demand that $\{l \mid v R l\}$ is a Moore family.

Condition (4.5) has two immediate consequences:

$$\begin{aligned} v R \top \\ v R l_1 \wedge v R l_2 \Rightarrow v R (l_1 \sqcap l_2) \end{aligned}$$

The first formula says that \top describes any value and the second formula says that if we have two descriptions of a value then their greatest lower bound is also a description of the value.

Example 4.3 Returning to the Constant Propagation Analysis of Example 4.1 we can now specify the relation

$$R_{CP} : \text{State} \times \widehat{\text{State}}_{CP} \rightarrow \{\text{true}, \text{false}\}$$

between the values (i.e. the states) and the properties (i.e. the abstract states):

$$\sigma R_{CP} \hat{\sigma} \text{ iff } \forall x \in FV(S_*) : (\hat{\sigma}(x) = \top \vee \sigma(x) = \hat{\sigma}(x))$$

Thus $\hat{\sigma}$ may map some variables to \top but if $\hat{\sigma}$ maps a variable x to an element in \mathbf{Z} then this must also be the value of $\sigma(x)$.

Let us observe that the conditions (4.4) and (4.5) are fulfilled by the Constant Propagation Analysis. Recall from Section 2.3 that $(\widehat{\text{State}}_{CP}, \sqsubseteq_{CP})$ is a complete lattice with the ordering \sqsubseteq_{CP} . It is then straightforward to verify that (4.4) and (4.5) do indeed hold. (Also compare with Exercise 2.7.) ■

Example 4.4 For the Control Flow Analysis mentioned in Example 4.2 we shall define

$$R_{CFA} : \text{Val} \times (\widehat{\text{Env}} \times \widehat{\text{Val}}) \rightarrow \{\text{true}, \text{false}\}$$

to be the relation \mathcal{V} of Section 3.2:

$$v R_{CFA} (\hat{\rho}, \hat{v}) \text{ iff } v \mathcal{V} (\hat{\rho}, \hat{v})$$

Recall that we have two kinds of values $v \in \mathbf{Val}$, constants c and closures $\text{close } t$ in ρ , and that \mathcal{V} is given by:

$$v \mathcal{V} (\widehat{\rho}, \widehat{v}) \text{ iff } \begin{cases} \text{true} & \text{if } v = c \\ t \in \widehat{v} \wedge \forall x \in \text{dom}(\rho) : \rho(x) \mathcal{V} (\widehat{\rho}, \widehat{\rho}(x)) & \text{if } v = \text{close } t \text{ in } \rho \end{cases}$$

The correctness condition (4.3) can be reformulated as

$$(v_1 \mathcal{V} (\widehat{\rho}, \widehat{v}_1) \wedge [] \vdash (e_* v_1^{l_1})^{l_2} \rightarrow^* v_2^{l_2} \wedge (\widehat{C}, \widehat{\rho}) \models (e_* c^{l_1})^{l_2} \wedge \widehat{C}(l_1) = \widehat{v}_1 \wedge \widehat{C}(l_2) = \widehat{v}_2) \Rightarrow v_2 \mathcal{V} (\widehat{\rho}, \widehat{v}_2)$$

and it follows from the correctness result established by Theorem 3.10 in Section 3.2 (see Exercise 4.3).

Finally, let us observe that the Control Flow Analysis also satisfies the conditions (4.4) and (4.5). For this we shall equip $\widehat{\mathbf{Env}} \times \widehat{\mathbf{Val}}$ with the partial ordering \sqsubseteq_{CFA} defined by:

$$(\widehat{\rho}_1, \widehat{v}_1) \sqsubseteq_{\text{CFA}} (\widehat{\rho}_2, \widehat{v}_2) \text{ iff } \widehat{v}_1 \subseteq \widehat{v}_2 \wedge \forall x : \widehat{\rho}_1(x) \subseteq \widehat{\rho}_2(x)$$

This will turn $\widehat{\mathbf{Env}} \times \widehat{\mathbf{Val}}$ into a complete lattice. By induction on $v \in \mathbf{Val}$ one can then easily prove that (4.4) and (4.5) are fulfilled. ■

4.1.2 Representation Functions

An alternative approach to the use of a correctness relation $R : V \times L \rightarrow \{\text{true}, \text{false}\}$ between values and properties is to use a *representation function*:

$$\beta : V \rightarrow L$$

The idea is that β maps a value to the *best* property describing it. The correctness criterion for the analysis will then be formulated as follows:

$$\beta(v_1) \sqsubseteq l_1 \wedge p \vdash v_1 \rightsquigarrow v_2 \wedge p \vdash l_1 \triangleright l_2 \Rightarrow \beta(v_2) \sqsubseteq l_2 \quad (4.6)$$

This is also expressed by the following diagram:

$$\begin{array}{ccccc} p \vdash & v_1 & \rightsquigarrow & v_2 & \\ & \beta \downarrow & \Rightarrow & \downarrow \beta & \\ & \sqcap \mid & & \sqcap \mid & \\ p \vdash & l_1 & \triangleright & l_2 & \end{array}$$

Thus the idea is that if the initial value v_1 is safely described by l_1 then the final value v_2 will be safely described by the result l_2 of the analysis.

Equivalence of correctness formulations. Lemma 4.5 below shows that the formulations (4.3) and (4.6) of the correctness of the analysis are indeed equivalent (when R and β are suitably related). To establish this we shall first show how to define a correctness relation R_β from a given representation function β :

$$v R_\beta l \text{ iff } \beta(v) \sqsubseteq l$$

Next we show how to define a representation function β_R from a correctness relation R :

$$\beta_R(v) = \bigcap \{l \mid v R l\}$$

Lemma 4.5

- (i) Given $\beta : V \rightarrow L$, then the relation $R_\beta : V \times L \rightarrow \{\text{true}, \text{false}\}$ satisfies conditions (4.4) and (4.5), and furthermore $\beta_{R_\beta} = \beta$.
- (ii) Given $R : V \times L \rightarrow \{\text{true}, \text{false}\}$ satisfying conditions (4.4) and (4.5), then β_R is well-defined and $R_{\beta_R} = R$.

Hence the two formulations (4.3) and (4.6) of correctness are equivalent. ■

Proof To prove (i) we first observe that condition (4.4) is immediate since \sqsubseteq is transitive. Condition (4.5) is immediate because when $\beta(v)$ is a lower bound for L' we have $\beta(v) \sqsubseteq \bigcap L'$. The calculation $\beta_{R_\beta}(v) = \bigcap \{l \mid v R_\beta l\} = \bigcap \{l \mid \beta(v) \sqsubseteq l\} = \beta(v)$ then concludes the proof of (i).

To prove (ii) we observe that from $v R l$ we get $\beta_R(v) \sqsubseteq l$ and hence $v R_{\beta_R} l$. Conversely, from $v R_{\beta_R} l$ we get $\beta_R(v) \sqsubseteq l$; writing $L' = \{l \mid v R l\}$ it is clear that (4.5) gives $v R (\bigcap L')$ and this amounts to $v R (\beta_R(v))$; we then get the desired result $v R l$ by (4.4). ■

Motivated by these results we shall say that the relation R is *generated by* the representation function β whenever $v R l$ is equivalent to $\beta(v) \sqsubseteq l$. This relationship is illustrated in Figure 4.1: The relation R expresses that v is described by all the properties above $\beta(v)$ and β expresses that among all the properties that describe v , $\beta(v)$ is the best.

Example 4.6 For the Constant Propagation Analysis studied in Examples 4.1 and 4.3 we shall define

$$\beta_{\text{CP}} : \text{State} \rightarrow \widehat{\text{State}}_{\text{CP}}$$

as the injection of State into $\widehat{\text{State}}_{\text{CP}}$: $\beta_{\text{CP}}(\sigma) = \lambda x. \sigma(x)$. It is straightforward to verify that R_{CP} is *generated by* β_{CP} , i.e.

$$\sigma R_{\text{CP}} \hat{\sigma} \Leftrightarrow \beta_{\text{CP}}(\sigma) \sqsubseteq_{\text{CP}} \hat{\sigma}$$

using the definition of the ordering \sqsubseteq_{CP} on $\widehat{\text{State}}_{\text{CP}}$. ■




Figure 4.1: Correctness relation R generated by representation function β .

Example 4.7 For the Control Flow Analysis studied in Examples 4.2 and 4.4 we shall define

$$\beta_{\text{CFA}} : \mathbf{Val} \rightarrow \widehat{\mathbf{Env}} \times \widehat{\mathbf{Val}}$$

inductively on the structure of values $v \in \mathbf{Val}$:

$$\beta_{\text{CFA}}(v) = \begin{cases} (\lambda x. \emptyset, \emptyset) & \text{if } v = c \\ (\beta_{\text{CFA}}^E(\rho), \{t\}) & \text{if } v = \text{close } t \text{ in } \rho \end{cases}$$

The first clause reflects that we do not collect constants in a pure 0-CFA analysis. In the second clause we only have one closure so the abstract value will be a singleton set and we construct the associated “minimal” abstract environment by extending β_{CFA} to operate on environments. To do that we shall “merge” all the abstract environments occurring in $\bigcup\{\beta_{\text{CFA}}(\rho(x)) \mid x \in \mathbf{Var}_*\}$; this reflects that the 0-CFA analysis uses one global abstract environment to describe all the possible local environments of the semantics. So we define $\beta_{\text{CFA}}^E : \widehat{\mathbf{Env}} \rightarrow \widehat{\mathbf{Env}}$ by:

$$\beta_{\text{CFA}}^E(\rho)(x) = \bigcup \{ \widehat{\rho}_y(x) \mid \beta_{\text{CFA}}(\rho(y)) = (\widehat{\rho}_y, \widehat{v}_y) \text{ and } y \in \text{dom}(\rho) \} \\ \cup \begin{cases} \{ \widehat{v}_x \} & \text{if } x \in \text{dom}(\rho) \text{ and } \beta_{\text{CFA}}(\rho(x)) = (\widehat{\rho}_x, \widehat{v}_x) \\ \emptyset & \text{otherwise} \end{cases}$$

To show that R_{CFA} is generated by β_{CFA} we have to show that:

$$v R_{\text{CFA}} (\widehat{\rho}, \widehat{v}) \Leftrightarrow \beta_{\text{CFA}}(v) \sqsubseteq_{\text{CFA}} (\widehat{\rho}, \widehat{v})$$

This follows by induction on $v \in \mathbf{Val}$ and we leave the details to Exercise 4.4. ■

4.1.3 A Modest Generalisation

We shall conclude by performing a modest generalisation of the development performed so far. A program p specifies how one value v_1 is transformed into another value v_2 :

$$p \vdash v_1 \rightsquigarrow v_2$$

Here $v_1 \in V_1$ and $v_2 \in V_2$ and we shall subsequently refrain from imposing the condition that $V_1 = V_2$; thus we shall allow the programs to have different "argument" and "result" types – this will e.g. be the case for most functional programs. The analysis of p specifies how a property l_1 is transformed into a property l_2 :

$$p \vdash l_1 \triangleright l_2$$

Here $l_1 \in L_1$ and $l_2 \in L_2$ and again we shall refrain from imposing the restriction that $L_1 = L_2$. As previously argued it is natural to demand that $p \vdash l_1 \triangleright l_2$ specifies a function

$$f_p : L_1 \rightarrow L_2$$

given by $f_p(l_1) = l_2$ iff $p \vdash l_1 \triangleright l_2$.

Turning to the correctness conditions we shall now assume that we have two correctness relations, one for V_1 and L_1 and one for V_2 and L_2 :

$$R_1 : V_1 \times L_1 \rightarrow \{\text{true}, \text{false}\} \text{ generated by } \beta_1 : V_1 \rightarrow L_1$$

$$R_2 : V_2 \times L_2 \rightarrow \{\text{true}, \text{false}\} \text{ generated by } \beta_2 : V_2 \rightarrow L_2$$

Correctness of f_p now amounts to

$$v_1 R_1 l_1 \wedge p \vdash v_1 \rightsquigarrow v_2 \Rightarrow v_2 R_2 f_p(l_1)$$

for all $v_1 \in V_1$, $v_2 \in V_2$ and $l_1 \in L_1$. Using the concept of *logical relations* (briefly mentioned above) this can be written as:

$$(p \vdash \cdot \rightsquigarrow \cdot) (R_1 \twoheadrightarrow R_2) f_p$$

To be precise, $\rightsquigarrow (R_1 \twoheadrightarrow R_2) f$ means that:

$$\forall v_1, v_2, l_1 : v_1 \rightsquigarrow v_2 \wedge v_1 R_1 l_1 \Rightarrow v_2 R_2 f(l_1)$$

Higher-order formulation. We can now ask whether the relation $R_1 \twoheadrightarrow R_2$ defined above is a correctness relation. Lemma 4.8 below shows that this is indeed the case and furthermore that we can find a representation function β such that $R_1 \twoheadrightarrow R_2$ is *generated by* β . The representation function β can be defined from the representation functions β_1 and β_2 and it will be denoted $\beta_1 \twoheadrightarrow \beta_2$:

$$(\beta_1 \twoheadrightarrow \beta_2)(\rightsquigarrow) = \lambda l_1. \bigsqcup \{ \beta_2(v_2) \mid \beta_1(v_1) \sqsubseteq l_1 \wedge v_1 \rightsquigarrow v_2 \}$$

Lemma 4.8 If R_i is a correctness relation for V_i and L_i that is *generated* by the representation function $\beta_i : V_i \rightarrow L_i$ (for $i = 1, 2$) then $R_1 \rightarrow R_2$ is a correctness relation and it is *generated by* the representation function $\beta_1 \rightarrow \beta_2$. ■

Proof We shall prove $\sim (R_1 \rightarrow R_2) f \Leftrightarrow (\beta_1 \rightarrow \beta_2)(\sim) \sqsubseteq f$. We calculate:

$$\begin{aligned} (\beta_1 \rightarrow \beta_2)(\sim) \sqsubseteq f &\Leftrightarrow \forall l_1 : \bigsqcup \{ \beta_2(v_2) \mid \beta_1(v_1) \sqsubseteq l_1 \wedge v_1 \sim v_2 \} \sqsubseteq f(l_1) \\ &\Leftrightarrow \forall l_1, v_1, v_2 : (\beta_1(v_1) \sqsubseteq l_1 \wedge v_1 \sim v_2 \Rightarrow \beta_2(v_2) \sqsubseteq f(l_1)) \\ &\Leftrightarrow \forall l_1, v_1, v_2 : (v_1 R_1 l_1 \wedge v_1 \sim v_2 \Rightarrow v_2 R_2 f(l_1)) \\ &\Leftrightarrow \sim (R_1 \rightarrow R_2) f \end{aligned}$$

Note that it now follows (from Lemma 4.5) that if each R_i satisfies conditions (4.4) and (4.5) then so does $R_1 \rightarrow R_2$. ■

Example 4.9 Consider the program `plus` with the semantics given by

$$\text{plus} \vdash (z_1, z_2) \sim z_1 + z_2$$

where $z_1, z_2 \in \mathbf{Z}$. A very precise analysis might use the complete lattices $(\mathcal{P}(\mathbf{Z}), \subseteq)$ and $(\mathcal{P}(\mathbf{Z} \times \mathbf{Z}), \subseteq)$ as follows:

$$f_{\text{plus}}(ZZ) = \{z_1 + z_2 \mid (z_1, z_2) \in ZZ\}$$

where $ZZ \subseteq \mathbf{Z} \times \mathbf{Z}$. Consider now the correctness relations R_Z and $R_{Z \times Z}$ *generated by* the representation functions:

$$\begin{aligned} \beta_Z(z) &= \{z\} \\ \beta_{Z \times Z}(z_1, z_2) &= \{(z_1, z_2)\} \end{aligned}$$

The correctness of the analysis of `plus` can now be expressed by

$$\forall z_1, z_2, z, ZZ : \text{plus} \vdash (z_1, z_2) \sim z \wedge (z_1, z_2) R_{Z \times Z} ZZ \Rightarrow z R_Z f_{\text{plus}}(ZZ)$$

or more succinctly

$$(\text{plus} \vdash \cdot \sim \cdot) (R_{Z \times Z} \rightarrow R_Z) f_{\text{plus}}$$

The representation function $\beta_{Z \times Z} \rightarrow \beta_Z$ satisfies

$$(\beta_{Z \times Z} \rightarrow \beta_Z)(p \vdash \cdot \sim \cdot) = \lambda ZZ. \{z \mid (z_1, z_2) \in ZZ \wedge p \vdash (z_1, z_2) \sim z\}$$

so the correctness can also be expressed as $(\beta_{Z \times Z} \rightarrow \beta_Z)(\text{plus} \vdash \cdot \sim \cdot) \sqsubseteq f_{\text{plus}}$. ■

The above example illustrates how the abstract concepts can give a more succinct formulation of the correctness of the analysis. In the following we shall see several cases where we move freely between what we may call a “concrete” formulation of a property and an “abstract” formulation of the same property. And we shall see that the latter often will allow us to reuse general results so that we do not have to redevelop parts of the theory for each application considered.




Figure 4.2: The complete lattice $\text{Interval} = (\text{Interval}, \subseteq)$.

4.2 Approximation of Fixed Points

It should be clear by now that complete lattices play a crucial role in program analysis and in the remainder of this chapter we shall tacitly *assume* that property spaces such as L and M are indeed complete lattices. We refer to Appendix A for the basic notions of complete lattices and monotone functions. Here we present an interesting complete lattice that forms the basis of many analyses over the integers.

Example 4.10 We shall now present a complete lattice that may be used for *Array Bound Analysis*, i.e. for determining if an array index is always within the bounds of the array – if this is the case then a number of run-time checks can be eliminated.

The lattice $(\text{Interval}, \subseteq)$ of intervals over \mathbf{Z} may be described as follows. The elements are

$$\text{Interval} = \{\perp\} \cup \{[z_1, z_2] \mid z_1 \leq z_2, z_1 \in \mathbf{Z} \cup \{-\infty\}, z_2 \in \mathbf{Z} \cup \{\infty\}\}$$

where the ordering \leq on \mathbf{Z} is extended to an ordering on $\mathbf{Z}' = \mathbf{Z} \cup \{-\infty, \infty\}$ by setting $-\infty \leq z$, $z \leq \infty$, and $-\infty \leq \infty$ (for all $z \in \mathbf{Z}$). Intuitively, \perp denotes the empty interval and $[z_1, z_2]$ is the interval from z_1 to z_2 including the end points if they are in \mathbf{Z} . We shall use *int* to range over elements of Interval .

The partial ordering \sqsubseteq on **Interval** is depicted in Figure 4.2; the idea is that $int_1 \sqsubseteq int_2$ corresponds to $\{z \mid z \text{ is in } int_1\} \subseteq \{z \mid z \text{ is in } int_2\}$ where the meaning of "is in" should be immediate. To give a succinct definition of the partial ordering we define the infimum and supremum operations on intervals as follows:

$$\begin{aligned} \text{inf}(int) &= \begin{cases} \infty & \text{if } int = \perp \\ z_1 & \text{if } int = [z_1, z_2] \end{cases} \\ \text{sup}(int) &= \begin{cases} -\infty & \text{if } int = \perp \\ z_2 & \text{if } int = [z_1, z_2] \end{cases} \end{aligned}$$

This allows us to define:

$$int_1 \sqsubseteq int_2 \quad \text{iff} \quad \text{inf}(int_2) \leq \text{inf}(int_1) \wedge \text{sup}(int_1) \leq \text{sup}(int_2)$$

We claim that $(\mathbf{Interval}, \sqsubseteq)$ is indeed a complete lattice. We shall prove this by showing that each subset of **Interval** has a least upper bound and then refer to Lemma A.2 of Appendix A to get that $(\mathbf{Interval}, \sqsubseteq)$ is a complete lattice. So let Y be a subset of **Interval**. The idea is that each interval int of Y should be "contained in" the interval $\bigsqcup Y$ defined by:

$$\bigsqcup Y = \begin{cases} \perp & \text{if } Y \subseteq \{\perp\} \\ [\text{inf}'\{\text{inf}(int) \mid int \in Y\}, \text{sup}'\{\text{sup}(int) \mid int \in Y\}] & \text{otherwise} \end{cases}$$

where inf' and sup' are the infimum and supremum operators on \mathbf{Z}' corresponding to the ordering \leq on \mathbf{Z}' ; they are given by $\text{inf}'(\emptyset) = \infty$, $\text{inf}'(Z) = z'$ if $z' \in Z$ is the least element of Z , and $\text{inf}'(Z) = -\infty$ otherwise; and similarly $\text{sup}'(\emptyset) = -\infty$, $\text{sup}'(Z) = z'$ if $z' \in Z$ is the greatest element of Z , and $\text{sup}'(Z) = \infty$ otherwise. It is now straightforward to show that $\bigsqcup Y$ is indeed the least upper bound of Y . ■

Given a complete lattice $L = (L, \sqsubseteq, \sqcup, \sqcap, \perp, \top)$ the effect of a program, p , in transforming one property, l_1 , into another, l_2 , i.e. $p \vdash l_1 \triangleright l_2$, is normally given by an equation

$$f(l_1) = l_2$$

for a monotone function $f : L \rightarrow L$ dependent on the program p . Note that the demand that f is monotone is very natural for program analysis; it merely says that if l'_1 describes at least the values that l_1 does then also $f(l'_1)$ describes at least the values that $f(l_1)$ does.

For recursive or iterative program constructs we ideally want to obtain the least fixed point, $\text{lfp}(f)$, as the result of a *finite* iterative process. However, the iterative sequence $(f^n(\perp))_n$ need *not* eventually stabilise nor need its least upper bound necessarily equal $\text{lfp}(f)$. This might suggest considering

the iterative sequence $(f^n(\top))_n$ and even when it does not eventually stabilise we can always terminate the iteration at an arbitrary point in time. While this is safe (thanks to condition (4.4) of Section 4.1) it turns out to be grossly imprecise in practice.

Fixed points. We shall begin by recalling some of the properties of fixed points of monotone functions over complete lattices; we refer to Appendix A for the details of this development. So consider a monotone function $f : L \rightarrow L$ on a complete lattice $L = (L, \sqsubseteq, \sqcup, \sqcap, \perp, \top)$. A *fixed point* of f is an element $l \in L$ such that $f(l) = l$ and we write

$$\text{Fix}(f) = \{l \mid f(l) = l\}$$

for the set of fixed points. The function f is *reductive* at l if and only if $f(l) \sqsubseteq l$ and we write

$$\text{Red}(f) = \{l \mid f(l) \sqsubseteq l\}$$

for the set of elements upon which f is reductive; we shall say that f itself is reductive if $\text{Red}(f) = L$. Similarly, the function f is *extensive* at l if and only if $f(l) \sqsupseteq l$ and we write

$$\text{Ext}(f) = \{l \mid f(l) \sqsupseteq l\}$$

for the set of elements upon which f is extensive; we shall say that f itself is extensive if $\text{Ext}(f) = L$.

Since L is a complete lattice it is always the case that the set $\text{Fix}(f)$ will have a greatest lower bound in L and we denote it by $\text{lfp}(f)$; this is actually the *least fixed point* of f because Tarski's Theorem (Proposition A.10) ensures that:

$$\text{lfp}(f) = \bigsqcap \text{Fix}(f) = \bigsqcap \text{Red}(f) \in \text{Fix}(f) \subseteq \text{Red}(f)$$

Similarly, the set $\text{Fix}(f)$ will have a least upper bound in L and we denote it by $\text{gfp}(f)$; this is actually the *greatest fixed point* of f because Tarski's Theorem ensures that

$$\text{gfp}(f) = \bigsqcup \text{Fix}(f) = \bigsqcup \text{Ext}(f) \in \text{Fix}(f) \subseteq \text{Ext}(f)$$

In *Denotational Semantics* it is customary to iterate to the least fixed point by taking the least upper bound of the sequence $(f^n(\perp))_n$. However, we have not imposed any continuity requirements on f (e.g. that $f(\bigsqcup_n l_n) = \bigsqcup_n (f(l_n))$ for all ascending chains $(l_n)_n$) and consequently we cannot be sure to actually reach the fixed point. In a similar way one could consider the greatest lower bound of the sequence $(f^n(\top))_n$. One can show that

$$f^n(\perp) \sqsubseteq \bigsqcup_n f^n(\perp) \sqsubseteq \text{lfp}(f) \sqsubseteq \text{gfp}(f) \sqsubseteq \bigsqcap_n f^n(\top) \sqsubseteq f^n(\top)$$

as is illustrated in Figure 4.3; indeed all inequalities (i.e. \sqsubseteq) can be strict (i.e. \neq).




Figure 4.3: Fixed points of f .

4.2.1 Widening Operators

Since we cannot guarantee that the iterative sequence $(f^n(\perp))_n$ eventually stabilises nor that its least upper bound necessarily equals $lfp(f)$, we must consider another way of approximating $lfp(f)$. The idea is now to replace it by a new sequence $(f^n_{\nabla})_n$ that is known to eventually stabilise and to do so with a value that is a safe (upper) approximation of the least fixed point. The construction of the new sequence is parameterised on the operator ∇ , called a *widening operator*; the precision of the approximated fixed point as well as the cost of computing it depends on the actual choice of widening operator.

Upper bound operators. In preparation for the development, an operator $\boxplus : L \times L \rightarrow L$ on a complete lattice $L = (L, \sqsubseteq)$ is called an *upper bound operator* if

$$l_1 \sqsubseteq l_1 \boxplus l_2 \sqsupseteq l_2$$

for all $l_1, l_2 \in L$, i.e. it always returns an element larger than both its arguments. Note that we do *not* require \boxplus to be monotone, commutative, associative, nor absorptive (i.e. that $l \boxplus l = l$).

Let $(l_n)_n$ be a sequence of elements of L and let $\phi : L \times L \rightarrow L$ be a total function on L . We shall now use ϕ to construct a new sequence $(l_n^\phi)_n$ defined

by:

$$l_n^\phi = \begin{cases} l_n & \text{if } n = 0 \\ l_{n-1}^\phi \phi l_n & \text{if } n > 0 \end{cases}$$

The following result expresses that any sequence can be turned into an ascending chain by an upper bound operator:

Fact 4.11 If $(l_n)_n$ is a sequence and $\check{\cup}$ is an upper bound operator then $(l_n^{\check{\cup}})_n$ is an ascending chain; furthermore $l_n^{\check{\cup}} \supseteq \bigsqcup\{l_0, l_1, \dots, l_n\}$ for all n . ■

Proof To prove that $(l_n^{\check{\cup}})_n$ is an ascending chain it is sufficient to prove that $l_n^{\check{\cup}} \subseteq l_{n+1}^{\check{\cup}}$ for all n . If $n = 0$ we calculate $l_0^{\check{\cup}} = l_0 \subseteq l_0 \check{\cup} l_1 = l_1^{\check{\cup}}$. For the inductive step we have $l_n^{\check{\cup}} \subseteq l_n^{\check{\cup}} \check{\cup} l_{n+1} = l_{n+1}^{\check{\cup}}$ as required.

To prove that $l_n^{\check{\cup}} \supseteq \bigsqcup\{l_0, l_1, \dots, l_n\}$ we first observe that it holds trivially for $n = 0$. For the inductive step we have $l_{n+1}^{\check{\cup}} = l_n^{\check{\cup}} \check{\cup} l_{n+1} \supseteq \bigsqcup\{l_0, l_1, \dots, l_n\} \sqcup l_{n+1} = \bigsqcup\{l_0, l_1, \dots, l_n, l_{n+1}\}$ and the result follows. ■

Example 4.12 Consider the complete lattice $(\mathbf{Interval}, \sqsubseteq)$ of Figure 4.2 and let int be an arbitrary but fixed element of $\mathbf{Interval}$. Consider the following operator $\check{\cup}^{int}$ defined on $\mathbf{Interval}$:

$$int_1 \check{\cup}^{int} int_2 = \begin{cases} int_1 \sqcup int_2 & \text{if } int_1 \sqsubseteq int \vee int_2 \sqsubseteq int_1 \\ [-\infty, \infty] & \text{otherwise} \end{cases}$$

Note that the operation is not symmetric: for $int = [0, 2]$ we e.g. have $[1, 2] \check{\cup}^{int} [2, 3] = [1, 3]$ whereas $[2, 3] \check{\cup}^{int} [1, 2] = [-\infty, \infty]$.

It is immediate that $\check{\cup}^{int}$ is an upper bound operator. Consider now the sequence:

$$[0, 0], [1, 1], [2, 2], [3, 3], [4, 4], [5, 5], \dots$$

If $int = [0, \infty]$, then the upper bound operator will transform the above sequence into the ascending chain:

$$[0, 0], [0, 1], [0, 2], [0, 3], [0, 4], [0, 5], \dots$$

However, if $int = [0, 2]$, then we will get the following ascending chain

$$[0, 0], [0, 1], [0, 2], [0, 3], [-\infty, \infty], [-\infty, \infty], \dots$$

which eventually stabilises. ■

Widening operators. We can now introduce a special class of upper bound operators that will help us to approximate the least fixed points: An operator $\nabla : L \times L \rightarrow L$ is a *widening operator* if and only if:

- it is an upper bound operator, and




Figure 4.4: The widening operator ∇ applied to f .

- for all ascending chains $(l_n)_n$ the ascending chain $(l_n^\nabla)_n$ eventually stabilises.

Note that it follows from Fact 4.11 that $(l_n^\nabla)_n$ is indeed an ascending chain. The idea is as follows: Given a monotone function $f : L \rightarrow L$ on a complete lattice L and given a widening operator ∇ on L , we shall calculate the sequence $(f_\nabla^n)_n$ defined by

$$f_\nabla^n = \begin{cases} \perp & \text{if } n = 0 \\ f_\nabla^{n-1} & \text{if } n > 0 \wedge f(f_\nabla^{n-1}) \sqsubseteq f_\nabla^{n-1} \\ f_\nabla^{n-1} \nabla f(f_\nabla^{n-1}) & \text{otherwise} \end{cases}$$

As in Fact 4.11 it follows that this is an ascending chain and Proposition 4.13 below will ensure that this sequence eventually stabilises. From Fact 4.14 below we shall see that this means that we will eventually have $f(f_\nabla^m) \sqsubseteq f_\nabla^m$ for some value of m (corresponding to the second clause in the definition of f_∇^n). This means that f is reductive at f_∇^m and from Tarski's Theorem (Proposition A.10) we then know that $f_\nabla^m \supseteq lfp(f)$ must be the case; hence we take

$$lfp_\nabla(f) = f_\nabla^m$$

as the desired safe approximation of $lfp(f)$. This is illustrated in Figure 4.4. We shall now establish the necessary results.

Proposition 4.13

If ∇ is a widening operator then the ascending chain $(f_{\nabla}^n)_n$ eventually stabilises.

In preparation for the proof of Proposition 4.13 we shall first show the following technical result:

Fact 4.14 If ∇ is a widening operator then:

- (i) the sequence $(f_{\nabla}^n)_n$ is an ascending chain;
- (ii) if $f(f_{\nabla}^m) \sqsubseteq f_{\nabla}^m$ for some m then the sequence $(f_{\nabla}^n)_n$ eventually stabilises and furthermore $\forall n > m : f_{\nabla}^n = f_{\nabla}^m$ and $\bigsqcup_n f_{\nabla}^n = f_{\nabla}^m$;
- (iii) if $(f_{\nabla}^n)_n$ eventually stabilises then there exists an m such that $f(f_{\nabla}^m) \sqsubseteq f_{\nabla}^m$; and
- (iv) if $(f_{\nabla}^n)_n$ eventually stabilises then $\bigsqcup_n f_{\nabla}^n \sqsupseteq \text{lfp}(f)$.

These claims also hold if ∇ is just an upper bound operator. ■

Proof The proof of (i) is analogous to that of Fact 4.11 so we omit the details.

To prove (ii) we assume that $f(f_{\nabla}^m) \sqsubseteq f_{\nabla}^m$ for some m . By induction on $n > m$ we prove that $f_{\nabla}^n = f_{\nabla}^m$: for $n = m + 1$ it follows from the assumption and for the inductive step we note that $f(f_{\nabla}^n) \sqsubseteq f_{\nabla}^n$ will be the case so $f_{\nabla}^{n+1} = f_{\nabla}^n$. Thus $(f_{\nabla}^n)_n$ eventually stabilises and it follows that $\bigsqcup_n f_{\nabla}^n = f_{\nabla}^m$.

To prove (iii) we assume that $(f_{\nabla}^n)_n$ eventually stabilises. This means that there exists m such that $\forall n > m : f_{\nabla}^n = f_{\nabla}^m$. By way of contradiction assume that $f(f_{\nabla}^m) \sqsubseteq f_{\nabla}^m$ does not hold; then $f_{\nabla}^m = f_{\nabla}^{m+1} = f_{\nabla}^m \nabla f(f_{\nabla}^m) \sqsupseteq f(f_{\nabla}^m)$ and we have the desired contradiction.

To prove (iv) we observe that (ii) and (iii) give $\bigsqcup_n f_{\nabla}^n = f_{\nabla}^m$ for some m where $f(f_{\nabla}^m) \sqsubseteq f_{\nabla}^m$. Hence $f_{\nabla}^m \in \text{Red}(f)$ and by Tarski's Theorem (Proposition A.10) this shows $f_{\nabla}^m \sqsupseteq \text{lfp}(f)$. ■

We now turn to the proof of Proposition 4.13:

Proof By way of contradiction we shall assume that the ascending chain $(f_{\nabla}^n)_n$ never stabilises; i.e.:

$$\forall n_0 : \exists n \geq n_0 : f_{\nabla}^n \neq f_{\nabla}^{n_0}$$

It follows that $f(f_{\nabla}^{n-1}) \sqsubseteq f_{\nabla}^{n-1}$ never holds for any $n > 0$; because if it did, then Fact 4.14 gives that $(f_{\nabla}^n)_n$ eventually stabilises and our hypothesis would be contradicted. This means that the definition of $(f_{\nabla}^n)_n$ specialises to:

$$f_{\nabla}^n = \begin{cases} \perp & \text{if } n = 0 \\ f_{\nabla}^{n-1} \nabla f(f_{\nabla}^{n-1}) & \text{otherwise} \end{cases}$$

Now define the sequence $(l_n)_n$ by

$$l_n = \begin{cases} \perp & \text{if } n = 0 \\ f(f_{\nabla}^{n-1}) & \text{if } n > 0 \end{cases}$$

and note that $(l_n)_n$ is an ascending chain because $(f_{\nabla}^n)_n$ is an ascending chain (Fact 4.14) and f is monotone. We shall now prove that

$$\forall n : l_n^{\nabla} = f_{\nabla}^n$$

by induction on n : for $n = 0$ it is immediate and for $n > 0$ we have $l_n^{\nabla} = l_{n-1}^{\nabla} \nabla l_n = f_{\nabla}^{n-1} \nabla f(f_{\nabla}^{n-1}) = f_{\nabla}^n$. Since $(l_n)_n$ is an ascending chain and ∇ is a widening operator it follows that the sequence $(l_n^{\nabla})_n$ eventually stabilises, i.e. $(f_{\nabla}^n)_n$ eventually stabilises. This provides the desired contradiction and proves the result. ■

Example 4.15 Consider the complete lattice $(\text{Interval}, \sqsubseteq)$ of Figure 4.2. Let K be a *finite* set of integers, e.g. the set of integers explicitly mentioned in a given program. We shall now define a widening operator ∇_K based on K . The idea is that $[z_1, z_2] \nabla_K [z_3, z_4]$ is something like

$$[\text{LB}(z_1, z_3), \text{UB}(z_2, z_4)]$$

where $\text{LB}(z_1, z_3) \in \{z_1\} \cup K \cup \{-\infty\}$ is the best possible lower bound and $\text{UB}(z_2, z_4) \in \{z_2\} \cup K \cup \{\infty\}$ is the best possible upper bound. In this way a change in any of the bounds of the interval $[z_1, z_2]$ can only take place in a finite number of steps (corresponding to the elements of K).

For the precise definition we let $z_i \in \mathbf{Z}' = \mathbf{Z} \cup \{-\infty, \infty\}$ and write:

$$\text{LB}_K(z_1, z_3) = \begin{cases} z_1 & \text{if } z_1 \leq z_3 \\ k & \text{if } z_3 < z_1 \wedge k = \max\{k \in K \mid k \leq z_3\} \\ -\infty & \text{if } z_3 < z_1 \wedge \forall k \in K : z_3 < k \end{cases}$$

$$\text{UB}_K(z_2, z_4) = \begin{cases} z_2 & \text{if } z_4 \leq z_2 \\ k & \text{if } z_2 < z_4 \wedge k = \min\{k \in K \mid z_4 \leq k\} \\ \infty & \text{if } z_2 < z_4 \wedge \forall k \in K : k < z_4 \end{cases}$$

We can now define $\nabla = \nabla_K$ by:

$$\text{int}_1 \nabla \text{int}_2 = \begin{cases} \perp & \text{if } \text{int}_1 = \text{int}_2 = \perp \\ [\text{LB}_K(\text{inf}(\text{int}_1), \text{inf}(\text{int}_2)), \text{UB}_K(\text{sup}(\text{int}_1), \text{sup}(\text{int}_2))] & \\ \text{otherwise} & \end{cases}$$

As an example consider the ascending chain $(\text{int}_n)_n$:

$$[0, 1], [0, 2], [0, 3], [0, 4], [0, 5], [0, 6], [0, 7], \dots$$

and assume that $K = \{3, 5\}$. Then $(\text{int}_n^{\nabla})_n$ will be the chain

$$[0, 1], [0, 3], [0, 3], [0, 5], [0, 5], [0, \infty], [0, \infty], \dots$$

It is straightforward to show that ∇ is indeed an upper bound operator. To show that it is a widening operator we consider an ascending chain $(int_n)_n$ and must show that the ascending chain $(int_n^\nabla)_n$ eventually stabilises. By way of contradiction suppose that $(int_n^\nabla)_n$ does not eventually stabilise. Then at least one of the following properties will hold:

$$\begin{aligned} (\forall n : \inf(int_n^\nabla) > -\infty) \quad \wedge \quad \inf(\bigsqcup_n int_n^\nabla) = -\infty \\ (\forall n : \sup(int_n^\nabla) < \infty) \quad \wedge \quad \sup(\bigsqcup_n int_n^\nabla) = \infty \end{aligned}$$

Without loss of generality we can assume that the second property holds. Hence there must exist an infinite sequence $n_1 < n_2 < \dots$ such that

$$\forall i : \infty > \sup(int_{n_i+1}^\nabla) > \sup(int_{n_i}^\nabla)$$

and by finiteness of K there must be some j such that

$$\forall i \geq j : \infty > \sup(int_{n_i+1}^\nabla) > \sup(int_{n_i}^\nabla) > \max(K)$$

where for the purpose of this definition we set $\max(\emptyset) = -\infty$. Since we also have

$$int_{n_j+1}^\nabla = int_{n_j}^\nabla \nabla int_{n_j+1}$$

it must be the case that $\sup(int_{n_j+1}^\nabla) > \sup(int_{n_j}^\nabla)$ as otherwise $\sup(int_{n_j+1}^\nabla) = \sup(int_{n_j}^\nabla)$. But then the construction yields

$$\sup(int_{n_j+1}^\nabla) = \infty$$

and we have the desired contradiction. This shows that ∇ is a widening operator on the complete lattice of intervals. ■

4.2.2 Narrowing Operators

Using the technique of widening we managed to arrive at an upper approximation f_{∇}^m of the least fixed point of f . However, we have $f(f_{\nabla}^m) \sqsubseteq f_{\nabla}^m$ so f is reductive at f_{∇}^m and this immediately suggests a way of improving the approximation by considering the iterative sequence $(f^n(f_{\nabla}^m))_n$. Since $f_{\nabla}^m \in Red(f)$ this will be a descending chain with $f^n(f_{\nabla}^m) \in Red(f)$ and hence $f^n(f_{\nabla}^m) \sqsupseteq lfp(f)$ for all n . Once again we have no reason to believe that this descending chain eventually stabilises although it is of course safe to stop at an arbitrary point. This scenario is not quite the one we saw above but inspired by the notion of widening we can define the notion of narrowing that encapsulates a termination criterion. This development can safely be omitted on a first reading.

An operator $\Delta : L \times L \rightarrow L$ is a *narrowing operator* if:




Figure 4.5: The narrowing operator Δ applied to f .

- $l_2 \sqsubseteq l_1 \Rightarrow l_2 \sqsubseteq (l_1 \Delta l_2) \sqsubseteq l_1$ for all $l_1, l_2 \in L$, and
- for all descending chains $(l_n)_n$ the sequence $(l_n^\Delta)_n$ eventually stabilises.

Note that we do *not* require Δ to be monotone, commutative, associative or absorptive. One can show that $(l_n^\Delta)_n$ is a descending chain when $(l_n)_n$ is a descending chain (Exercise 4.10).

The idea is as follows: For f_{∇}^m satisfying $f(f_{\nabla}^m) \sqsubseteq f_{\nabla}^m$, i.e. $lfp_{\nabla}(f) = f_{\nabla}^m$, we now construct the sequence $([f]_{\Delta}^n)_n$ by

$$[f]_{\Delta}^n = \begin{cases} f_{\nabla}^m & \text{if } n = 0 \\ [f]_{\Delta}^{n-1} \Delta f([f]_{\Delta}^{n-1}) & \text{if } n > 0 \end{cases}$$

Lemma 4.16 below guarantees that this is a descending chain where all elements satisfy $lfp(f) \sqsubseteq [f]_{\Delta}^n$. Proposition 4.17 below tells us that this chain eventually stabilises so $[f]_{\Delta}^{m'} = [f]_{\Delta}^{m'+1}$ for some value m' . We shall therefore take

$$lfp_{\nabla}^{\Delta}(f) = [f]_{\Delta}^{m'}$$

as the desired approximation of $lfp(f)$. The complete development is illustrated in Figures 4.4 and 4.5. We shall now establish the required results.

Lemma 4.16 If Δ is a narrowing operator and $f(f_{\nabla}^m) \subseteq f_{\nabla}^m$ then $([f]_{\Delta}^n)_n$ is a descending chain in $\text{Red}(f)$ and

$$[f]_{\Delta}^n \supseteq f^n(f_{\nabla}^m) \supseteq \text{Ifp}(f)$$

for all n . ■

Proof By induction on n we prove:

$$f^{n+1}(f_{\nabla}^m) \subseteq f([f]_{\Delta}^n) \subseteq [f]_{\Delta}^{n+1} \subseteq [f]_{\Delta}^n \quad (4.7)$$

For the basis ($n = 0$) it is immediate that

$$f^{n+1}(f_{\nabla}^m) \subseteq f([f]_{\Delta}^n) \subseteq [f]_{\Delta}^n$$

using that $f(f_{\nabla}^m) \subseteq f_{\nabla}^m$. By construction of $[f]_{\Delta}^{n+1}$ we then get

$$f([f]_{\Delta}^n) \subseteq [f]_{\Delta}^{n+1} \subseteq [f]_{\Delta}^n$$

and together these two results establish the basis of the induction proof.

For the inductive step we may apply f to the induction hypothesis (4.7) and get

$$f^{n+2}(f_{\nabla}^m) \subseteq f^2([f]_{\Delta}^n) \subseteq f([f]_{\Delta}^{n+1}) \subseteq f([f]_{\Delta}^n)$$

since f is assumed to be monotone. Using the induction hypothesis we also have $f([f]_{\Delta}^n) \subseteq [f]_{\Delta}^{n+1}$ so we obtain:

$$f^{n+2}(f_{\nabla}^m) \subseteq f([f]_{\Delta}^{n+1}) \subseteq [f]_{\Delta}^{n+1}$$

By construction of $[f]_{\Delta}^{n+2}$ we get

$$f([f]_{\Delta}^{n+1}) \subseteq [f]_{\Delta}^{n+2} \subseteq [f]_{\Delta}^{n+1}$$

and together these two results complete the proof of (4.7).

From (4.7) it now follows that $([f]_{\Delta}^n)_n$ is a descending chain in $\text{Red}(f)$. It also follows that $f^n(f_{\nabla}^m) \subseteq [f]_{\Delta}^n$ holds for $n > 0$ and it clearly also holds for $n = 0$. From the assumption $f(f_{\nabla}^m) \subseteq f_{\nabla}^m$ it is immediate that $f(f^n(f_{\nabla}^m)) \subseteq f^n(f_{\nabla}^m)$ for $n \geq 0$ and hence $f^n(f_{\nabla}^m) \in \text{Red}(f)$. But then $f^n(f_{\nabla}^m) \supseteq \text{Ifp}(f)$. This completes the proof. ■

Proposition 4.17

If Δ is a narrowing operator and $f(f_{\nabla}^m) \subseteq f_{\nabla}^m$ then the descending chain $([f]_{\Delta}^n)_n$ eventually stabilises.

Proof Define the sequence $(l_n)_n$ by

$$l_n = \begin{cases} f_{\nabla}^m & \text{if } n = 0 \\ f([f]_{\Delta}^{n-1}) & \text{if } n > 0 \end{cases}$$

and note that this defines a descending chain because $([f]_{\Delta}^n)_n$ is a descending chain and because $f([f]_{\Delta}^0) \sqsubseteq f_{\nabla}^m$. Therefore the sequence $(l_n^{\Delta})_n$ eventually stabilises. We now prove by induction on n that:

$$l_n^{\Delta} = [f]_{\Delta}^n$$

The basis ($n = 0$) is immediate. For the inductive step we calculate $l_{n+1}^{\Delta} = l_n^{\Delta} \Delta l_{n+1} = [f]_{\Delta}^n \Delta f([f]_{\Delta}^n) = [f]_{\Delta}^{n+1}$. It thus follows that $([f]_{\Delta}^n)_n$ eventually stabilises. ■

It is important to stress that narrowing operators are *not* the dual concept of widening operators. In particular, the sequence $(f_{\nabla}^n)_n$ may step outside $Ext(f)$ in order to end in $Red(f)$, whereas the sequence $([f]_{\Delta}^n)_n$ stays in $Red(f)$ all the time.

Example 4.18 Consider the complete lattice $(\mathbf{Interval}, \sqsubseteq)$ of Figure 4.2. Basically there are two kinds of infinite descending chains in $\mathbf{Interval}$: those with elements of the form $[-\infty, z]$ and those with elements of the form $[z, \infty]$ where $z \in \mathbf{Z}$. Consider an infinite sequence of the latter form; it will have elements

$$[z_1, \infty], [z_2, \infty], [z_3, \infty], \dots$$

where $z_1 < z_2 < z_3 < \dots$. The idea is now to define a narrowing operator Δ_N that will force the sequence to stabilise when $z_i \geq N$ for some fixed non-negative integer N . Similarly, for a descending chain with elements of the form $[-\infty, z_i]$ the narrowing operator will force it to stabilise when $z_i \leq -N$.

Formally, we shall define $\Delta = \Delta_N$ by

$$int_1 \Delta int_2 = \begin{cases} \perp & \text{if } int_1 = \perp \vee int_2 = \perp \\ [z_1, z_2] & \text{otherwise} \end{cases}$$

where

$$z_1 = \begin{cases} \inf(int_1) & \text{if } N < \inf(int_2) \wedge \sup(int_2) = \infty \\ \inf(int_2) & \text{otherwise} \end{cases}$$

$$z_2 = \begin{cases} \sup(int_1) & \text{if } \inf(int_2) = -\infty \wedge \sup(int_2) < -N \\ \sup(int_2) & \text{otherwise} \end{cases}$$

So consider e.g. the infinite descending chain $([n, \infty])_n$

$$[0, \infty], [1, \infty], [2, \infty], [3, \infty], [4, \infty], [5, \infty], \dots$$

and assume that $N = 3$. Then the operator will give the sequence $([n, \infty]^{\Delta})_n$:

$$[0, \infty], [1, \infty], [2, \infty], [3, \infty], [3, \infty], [3, \infty], \dots$$

Let us show that Δ is indeed a narrowing operator. It is immediate to verify that

$$int_2 \sqsubseteq int_1 \text{ implies } int_2 \sqsubseteq int_1 \Delta int_2 \sqsubseteq int_1$$

by cases on whether $\text{int}_2 = \perp$ or $\text{int}_2 \neq \perp$. We shall then show that if $(\text{int}_n)_n$ is a descending chain then the sequence $(\text{int}_n^\Delta)_n$ eventually stabilises. So assume that $(\text{int}_n)_n$ is a descending chain. One can then show that $(\text{int}_n^\Delta)_n$ is a descending chain (Exercise 4.9). Next suppose by way of contradiction that $(\text{int}_n^\Delta)_n$ never eventually stabilises. It follows that there exists $n_1 \geq 0$ such that:

$$\text{int}_n^\Delta \neq [-\infty, \infty] \text{ for all } n \geq n_1$$

It furthermore follows for all $n \geq n_1$ that int_n^Δ must have:

$$\sup(\text{int}_n^\Delta) = \infty \text{ or } \inf(\text{int}_n^\Delta) = -\infty$$

Without loss of generality let us assume that all $n \geq n_1$ have $\sup(\text{int}_n^\Delta) = \infty$ and $\inf(\text{int}_n^\Delta) \in \mathbf{Z}$. Hence there exists $n_2 \geq n_1$ such that:

$$\inf(\text{int}_n^\Delta) > N \text{ for all } n \geq n_2$$

But then $\text{int}_n^\Delta = \text{int}_{n_2}^\Delta$ for all $n \geq n_2$ and we have the desired contradiction. This shows that Δ is a narrowing operator. ■

4.3 Galois Connections

Sometimes calculations on a complete lattice L may be too costly or even uncomputable and this may motivate replacing L by a simpler lattice M . An example is when L is the powerset of integers and M is a lattice of intervals. So rather than performing the analysis $p \vdash l_1 \triangleright l_2$ in L , the idea will be to find a description of the elements of L in M and to perform the analysis $p \vdash m_1 \triangleright m_2$ in M . To express the relationship between L and M it is customary to use an *abstraction function*

$$\alpha : L \rightarrow M$$

giving a representation of the elements of L as elements of M and a *concretisation function*

$$\gamma : M \rightarrow L$$

that expresses the meaning of elements of M in terms of elements of L . We shall write

$$(L, \alpha, \gamma, M)$$

or

$$L \begin{array}{c} \xleftarrow{\gamma} \\ \xrightarrow{\alpha} \end{array} M$$




Figure 4.6: The Galois connection (L, α, γ, M) .

for this setup and would expect that α and γ should be somehow related. We shall study this relationship in the present section and then return to the connection between $p \vdash l_1 \triangleright l_2$ and $p \vdash m_1 \triangleright m_2$ in Section 4.5; in Section 4.4 we shall study the systematic construction of such relationships.

We define (L, α, γ, M) to be a *Galois connection* between the complete lattices (L, \sqsubseteq) and (M, \sqsubseteq) if and only if

$$\alpha : L \rightarrow M \text{ and } \gamma : M \rightarrow L \text{ are monotone functions}$$

that satisfy:

$$\gamma \circ \alpha \sqsupseteq \lambda l.l \quad (4.8)$$

$$\alpha \circ \gamma \sqsubseteq \lambda m.m \quad (4.9)$$

Conditions (4.8) and (4.9) express that we do not lose safety by going back and forth between the two lattices although we may lose precision. In the case of (4.8) this ensures that if we start with an element $l \in L$ we can first find a description $\alpha(l)$ of it in M and next determine which element $\gamma(\alpha(l))$ of L that describes $\alpha(l)$; this need not be l but it will be a safe approximation to l , i.e. $l \sqsubseteq \gamma(\alpha(l))$. This is illustrated in Figure 4.6.

Example 4.19 Let $\mathcal{P}(\mathbf{Z}) = (\mathcal{P}(\mathbf{Z}), \sqsubseteq)$ be the complete lattice of sets of integers and let $\mathbf{Interval} = (\mathbf{Interval}, \sqsubseteq)$ be the complete lattice of Figure 4.2. We shall now define a Galois connection

$$(\mathcal{P}(\mathbf{Z}), \alpha_{\mathbf{ZI}}, \gamma_{\mathbf{ZI}}, \mathbf{Interval})$$

between $\mathcal{P}(\mathbf{Z})$ and $\mathbf{Interval}$.

The concretisation function $\gamma_{ZI} : \mathbf{Interval} \rightarrow \mathcal{P}(\mathbf{Z})$ is defined by

$$\gamma_{ZI}(int) = \{z \in \mathbf{Z} \mid \inf(int) \leq z \leq \sup(int)\}$$

where \inf and \sup are as in Example 4.10. Thus γ_{ZI} will extract the set of elements described by the interval, e.g. $\gamma_{ZI}([0, 3]) = \{0, 1, 2, 3\}$ and $\gamma_{ZI}([0, \infty]) = \{z \in \mathbf{Z} \mid z \geq 0\}$.

The abstraction function $\alpha_{ZI} : \mathcal{P}(\mathbf{Z}) \rightarrow \mathbf{Interval}$ is defined by

$$\alpha_{ZI}(Z) = \begin{cases} \perp & \text{if } Z = \emptyset \\ [\inf'(Z), \sup'(Z)] & \text{otherwise} \end{cases}$$

where \inf' and \sup' are as in Example 4.10. Thus α_{ZI} will determine the smallest interval that includes all the elements of the set, e.g. $\alpha_{ZI}(\{0, 1, 3\}) = [0, 3]$ and $\alpha_{ZI}(\{2 * z \mid z > 0\}) = [2, \infty]$.

Let us verify that $(\mathcal{P}(\mathbf{Z}), \alpha_{ZI}, \gamma_{ZI}, \mathbf{Interval})$ is indeed a Galois connection. It is easy to see that α_{ZI} and γ_{ZI} are monotone functions. We shall next prove that (4.8) holds, i.e. that $\gamma_{ZI} \circ \alpha_{ZI} \supseteq \lambda Z.Z$. If $Z \neq \emptyset$ we have:

$$\begin{aligned} \gamma_{ZI}(\alpha_{ZI}(Z)) &= \gamma_{ZI}([\inf'(Z), \sup'(Z)]) \\ &= \{z \in \mathbf{Z} \mid \inf'(Z) \leq z \leq \sup'(Z)\} \\ &\supseteq Z \end{aligned}$$

For $Z = \emptyset$ we trivially have $\gamma_{ZI}(\alpha_{ZI}(\emptyset)) = \gamma_{ZI}(\perp) = \emptyset$ so we have proved (4.8). Intuitively, this condition expresses that if we start with a subset of \mathbf{Z} , find the smallest interval containing it, and next determine the corresponding subset of \mathbf{Z} , then we will get a (possibly) larger subset of \mathbf{Z} than the one we started with.

Finally, we shall prove (4.9), i.e. that $\alpha_{ZI} \circ \gamma_{ZI} \sqsubseteq \lambda int.int$. Consider first $int = [z_1, z_2]$ where we have:

$$\begin{aligned} \alpha_{ZI}(\gamma_{ZI}([z_1, z_2])) &= \alpha_{ZI}(\{z \in \mathbf{Z} \mid z_1 \leq z \leq z_2\}) \\ &= [z_1, z_2] \end{aligned}$$

For $int = \perp$ we trivially have $\alpha_{ZI}(\gamma_{ZI}(\perp)) = \alpha_{ZI}(\emptyset) = \perp$ so we have proved (4.9). Intuitively, the condition expresses that if we start with an interval, determine the corresponding subset of \mathbf{Z} , and next find the smallest interval containing this set, then the resulting interval will include the interval we started with; actually we showed that the two intervals are equal. ■

Adjunctions. There is an alternative formulation of the Galois connection (L, α, γ, M) that is frequently easier to work with. We define (L, α, γ, M) to be an *adjunction* between complete lattices $L = (L, \sqsubseteq)$ and $M = (M, \sqsubseteq)$ if and only if

$$\alpha : L \rightarrow M \text{ and } \gamma : M \rightarrow L \text{ are total functions}$$

that satisfy

$$\alpha(l) \sqsubseteq m \Leftrightarrow l \sqsubseteq \gamma(m) \quad (4.10)$$

for all $l \in L$ and $m \in M$.

Condition (4.10) expresses that α and γ "respect" the orderings of the two lattices: If an element $l \in L$ is safely described by the element $m \in M$, i.e. $\alpha(l) \sqsubseteq m$, then it is also the case that the element described by m is safe with respect to l , i.e. $l \sqsubseteq \gamma(m)$.

Proposition 4.20

(L, α, γ, M) is an adjunction if and only if (L, α, γ, M) is a Galois connection.

Proof First assume that (L, α, γ, M) is a Galois connection and let us show that it is an adjunction. So assume first that $\alpha(l) \sqsubseteq m$; since γ is monotone we get $\gamma(\alpha(l)) \sqsubseteq \gamma(m)$; using that $\gamma \circ \alpha \sqsupseteq \lambda l.l$ we then get $l \sqsubseteq \gamma(\alpha(l)) \sqsubseteq \gamma(m)$ as required. The proof showing that $l \sqsubseteq \gamma(m)$ implies $\alpha(l) \sqsubseteq m$ is analogous.

Next assume that (L, α, γ, M) is an adjunction and let us prove that it is a Galois connection. First we prove that $\gamma \circ \alpha \sqsupseteq \lambda l.l$: for $l \in L$ we trivially have $\alpha(l) \sqsubseteq \alpha(l)$ and using that $\alpha(l) \sqsubseteq m \Rightarrow l \sqsubseteq \gamma(m)$ we get $l \sqsubseteq \gamma(\alpha(l))$ as required. The proof showing that $\alpha \circ \gamma \sqsubseteq \lambda m.m$ is analogous. To complete the proof we also have to show that α and γ are monotone. To see that α is monotone suppose that $l_1 \sqsubseteq l_2$; we have already proved that $\gamma \circ \alpha \sqsupseteq \lambda l.l$ so we have $l_1 \sqsubseteq l_2 \sqsubseteq \gamma(\alpha(l_2))$; using that $l \sqsubseteq \gamma(m) \Rightarrow \alpha(l) \sqsubseteq m$ we then get $\alpha(l_1) \sqsubseteq \alpha(l_2)$. The proof showing that γ is monotone is analogous. ■

Galois connections defined by extraction functions. We shall now see that representation functions (introduced in Section 4.1) can be used to define Galois connections. So consider once again the representation function $\beta : V \rightarrow L$ mapping the values of V to the properties of the complete lattice L . It gives rise to a Galois connection

$$(\mathcal{P}(V), \alpha, \gamma, L)$$

between $\mathcal{P}(V)$ and L where the abstraction and concretisation functions are defined by

$$\begin{aligned} \alpha(V') &= \bigsqcup\{\beta(v) \mid v \in V'\} \\ \gamma(l) &= \{v \in V \mid \beta(v) \sqsubseteq l\} \end{aligned}$$

for $V' \subseteq V$ and $l \in L$. Let us pause for a minute to see that this indeed defines an adjunction:

$$\begin{aligned} \alpha(V') \sqsubseteq l &\Leftrightarrow \bigsqcup\{\beta(v) \mid v \in V'\} \sqsubseteq l \\ &\Leftrightarrow \forall v \in V' : \beta(v) \sqsubseteq l \\ &\Leftrightarrow V' \subseteq \gamma(l) \end{aligned}$$

It follows from Proposition 4.20 that we also have a Galois connection.

A special case of the above construction that is frequently useful is when $L = (\mathcal{P}(D), \subseteq)$ for some set D and we have an *extraction function*

$$\eta : V \rightarrow D$$

mapping the values of V to their descriptions in D . We will then define the representation function $\beta_\eta : V \rightarrow \mathcal{P}(D)$ by $\beta_\eta(v) = \{\eta(v)\}$ and the Galois connection between $\mathcal{P}(V)$ and $\mathcal{P}(D)$ will now be written

$$(\mathcal{P}(V), \alpha_\eta, \gamma_\eta, \mathcal{P}(D))$$

where

$$\begin{aligned} \alpha_\eta(V') &= \bigcup \{\beta_\eta(v) \mid v \in V'\} &= \{\eta(v) \mid v \in V'\} \\ \gamma_\eta(D') &= \{v \in V \mid \beta_\eta(v) \subseteq D'\} &= \{v \mid \eta(v) \in D'\} \end{aligned}$$

for $V' \subseteq V$ and $D' \subseteq D$.

Example 4.21 Let us consider the two complete lattices $(\mathcal{P}(\mathbf{Z}), \subseteq)$ and $(\mathcal{P}(\mathbf{Sign}), \subseteq)$ where $\mathbf{Sign} = \{-, 0, +\}$; see Figure 4.7. The extraction function

$$\text{sign} : \mathbf{Z} \rightarrow \mathbf{Sign}$$

simply defines the signs of the integers and is specified by:

$$\text{sign}(z) = \begin{cases} - & \text{if } z < 0 \\ 0 & \text{if } z = 0 \\ + & \text{if } z > 0 \end{cases}$$

The above construction then gives us a Galois connection

$$(\mathcal{P}(\mathbf{Z}), \alpha_{\text{sign}}, \gamma_{\text{sign}}, \mathcal{P}(\mathbf{Sign}))$$

with

$$\begin{aligned} \alpha_{\text{sign}}(Z) &= \{\text{sign}(z) \mid z \in Z\} \\ \gamma_{\text{sign}}(S) &= \{z \in \mathbf{Z} \mid \text{sign}(z) \in S\} \end{aligned}$$

where $Z \subseteq \mathbf{Z}$ and $S \subseteq \mathbf{Sign}$. ■

4.3.1 Properties of Galois Connections

We shall present three interesting results. The first result says that a Galois connection is fully determined by either one of the abstraction and concretisation functions:




Figure 4.7: The complete lattice $\mathcal{P}(\text{Sign}) = (\mathcal{P}(\text{Sign}), \subseteq)$.

Lemma 4.22 If (L, α, γ, M) is a Galois connection then:

- (i) α uniquely determines γ by $\gamma(m) = \bigsqcup\{l \mid \alpha(l) \subseteq m\}$ and γ uniquely determines α by $\alpha(l) = \prod\{m \mid l \subseteq \gamma(m)\}$.
- (ii) α is completely additive and γ is completely multiplicative.

In particular $\alpha(\perp) = \perp$ and $\gamma(\top) = \top$. ■

Proof To show (i) we shall first show that γ is determined by α . Since (L, α, γ, M) is an adjunction by Proposition 4.20 we have $\gamma(m) = \bigsqcup\{l \mid l \subseteq \gamma(m)\} = \bigsqcup\{l \mid \alpha(l) \subseteq m\}$. This shows that α uniquely determines γ : if both (L, α, γ_1, M) and (L, α, γ_2, M) are Galois connection then $\gamma_1(m) = \bigsqcup\{l \mid \alpha(l) \subseteq m\} = \gamma_2(m)$ for all m and hence $\gamma_1 = \gamma_2$.

Similarly, we have $\alpha(l) = \prod\{m \mid \alpha(l) \subseteq m\} = \prod\{m \mid l \subseteq \gamma(m)\}$ and this shows that γ uniquely determines α .

To show (ii) consider $L' \subseteq L$; using Proposition 4.20 we then have

$$\begin{aligned} \alpha(\bigsqcup L') \subseteq m &\Leftrightarrow \bigsqcup L' \subseteq \gamma(m) \\ &\Leftrightarrow \forall l \in L' : l \subseteq \gamma(m) \\ &\Leftrightarrow \forall l \in L' : \alpha(l) \subseteq m \\ &\Leftrightarrow \bigsqcup\{\alpha(l) \mid l \in L'\} \subseteq m \end{aligned}$$

and it follows that $\alpha(\bigsqcup L') = \bigsqcup\{\alpha(l) \mid l \in L'\}$.

The proof that $\gamma(\prod M') = \prod\{\gamma(m) \mid m \in M'\}$ is analogous. ■

Motivated by Lemma 4.22 we shall say that if (L, α, γ, M) is a Galois connection then α is the *lower adjoint* (or left adjoint) of γ and that γ is the *upper adjoint* (or right adjoint) of α .

The next result shows that it suffices to specify either a completely additive abstraction function or a completely multiplicative concretisation function in order to obtain a Galois connection:

Lemma 4.23 If $\alpha : L \rightarrow M$ is completely additive then there exists $\gamma : M \rightarrow L$ such that (L, α, γ, M) is a Galois connection. Similarly, if $\gamma : M \rightarrow L$ is completely multiplicative then there exists $\alpha : L \rightarrow M$ such that (L, α, γ, M) is a Galois connection. ■

Proof Consider the claim for α and define γ by:

$$\gamma(m) = \bigsqcup \{l' \mid \alpha(l') \sqsubseteq m\}$$

Then we have $\alpha(l) \sqsubseteq m \Rightarrow l \in \{l' \mid \alpha(l') \sqsubseteq m\} \Rightarrow l \sqsubseteq \gamma(m)$ where the last implication follows from the definition of γ . For the other direction we first observe that $l \sqsubseteq \gamma(m) \Rightarrow \alpha(l) \sqsubseteq \alpha(\gamma(m))$ because α is completely additive and hence monotone. Now

$$\begin{aligned} \alpha(\gamma(m)) &= \alpha(\bigsqcup \{l' \mid \alpha(l') \sqsubseteq m\}) \\ &= \bigsqcup \{\alpha(l') \mid \alpha(l') \sqsubseteq m\} \\ &\sqsubseteq m \end{aligned}$$

so $l \sqsubseteq \gamma(m) \Rightarrow \alpha(l) \sqsubseteq m$. It follows that (L, α, γ, M) is an adjunction and hence a Galois connection by Proposition 4.20.

The proof of the claim for γ is analogous. ■

The final result shows that we neither lose nor gain precision by iterating abstraction and concretisation:

Fact 4.24 If (L, α, γ, M) is a Galois connection then $\alpha \circ \gamma \circ \alpha = \alpha$ and $\gamma \circ \alpha \circ \gamma = \gamma$. ■

Proof We have $\lambda l. l \sqsubseteq \gamma \circ \alpha$ and since α is monotone we get $\alpha \sqsubseteq \alpha \circ (\gamma \circ \alpha)$. Similarly $(\alpha \circ \gamma) \circ \alpha \sqsubseteq \alpha$ follows from $\alpha \circ \gamma \sqsubseteq \lambda m. m$. Thus $\alpha \circ \gamma \circ \alpha = \alpha$.

The proof of $\gamma \circ \alpha \circ \gamma = \gamma$ is analogous. ■

Example 4.25 As a somewhat more complex example consider the complete lattices **Interval** = $(\mathbf{Interval}, \sqsubseteq)$ and $\mathcal{P}(\mathbf{Sign}) = (\mathcal{P}(\mathbf{Sign}), \sqsubseteq)$ of Figures 4.2 and 4.7, respectively.

Let us define a concretisation function $\gamma_{\text{IS}} : \mathcal{P}(\mathbf{Sign}) \rightarrow \mathbf{Interval}$ by:

$$\begin{array}{ll} \gamma_{\text{IS}}(\{-, 0, +\}) = [-\infty, \infty] & \gamma_{\text{IS}}(\{-, 0\}) = [-\infty, 0] \\ \gamma_{\text{IS}}(\{-, +\}) = [-\infty, \infty] & \gamma_{\text{IS}}(\{0, +\}) = [0, \infty] \\ \gamma_{\text{IS}}(\{-\}) = [-\infty, -1] & \gamma_{\text{IS}}(\{0\}) = [0, 0] \\ \gamma_{\text{IS}}(\{+\}) = [1, \infty] & \gamma_{\text{IS}}(\emptyset) = \perp \end{array}$$

To determine whether or not there exists an abstraction function

$$\alpha_{\text{IS}} : \mathbf{Interval} \rightarrow \mathcal{P}(\mathbf{Sign})$$

such that $(\mathbf{Interval}, \alpha_{\mathbf{IS}}, \gamma_{\mathbf{IS}}, \mathcal{P}(\mathbf{Sign}))$ is a Galois connection, we shall simply determine whether or not $\gamma_{\mathbf{IS}}$ is completely multiplicative: If it is, then Lemma 4.23 guarantees that there does indeed exist a Galois connection and Lemma 4.22 tells us how to construct the abstraction function. If $\gamma_{\mathbf{IS}}$ is not completely multiplicative then Lemma 4.22 guarantees that there cannot exist a Galois connection. In order to determine whether or not $\gamma_{\mathbf{IS}}$ is completely multiplicative we shall use Lemma A.4 of Appendix A: It is immediate to verify that $\mathcal{P}(\mathbf{Sign})$ is finite and that $\gamma_{\mathbf{IS}}$ satisfies conditions (i) and (ii) of Lemma A.4. To verify condition (iii) we need to compare $\gamma_{\mathbf{IS}}(S_1 \cap S_2)$ to $\gamma_{\mathbf{IS}}(S_1) \sqcap \gamma_{\mathbf{IS}}(S_2)$ for all pairs $(S_1, S_2) \in \mathcal{P}(\mathbf{Sign}) \times \mathcal{P}(\mathbf{Sign})$ of incomparable sets of signs, i.e. all pairs in the following list:

$$\begin{array}{lll} (\{-, 0\}, \{-, +\}), & (\{-, 0\}, \{0, +\}), & (\{-, 0\}, \{+\}), \\ (\{-, +\}, \{0, +\}), & (\{-, +\}, \{0\}), & (\{0, +\}, \{-\}), \\ (\{-\}, \{0\}), & (\{-\}, \{+\}), & (\{0\}, \{+\}) \end{array}$$

In checking the pair $(\{-, 0\}, \{-, +\})$ we calculate

$$\begin{aligned} \gamma_{\mathbf{IS}}(\{-, 0\} \cap \{-, +\}) &= \gamma_{\mathbf{IS}}(\{-\}) = [-\infty, -1] \\ \gamma_{\mathbf{IS}}(\{-, 0\}) \sqcap \gamma_{\mathbf{IS}}(\{-, +\}) &= [-\infty, 0] \sqcap [-\infty, \infty] = [-\infty, 0] \end{aligned}$$

and we deduce that $\gamma_{\mathbf{IS}}$ is *not* completely multiplicative. Hence according to Lemma 4.22 we cannot find any $\alpha_{\mathbf{IS}} : \mathbf{Interval} \rightarrow \mathcal{P}(\mathbf{Sign})$ such that $(\mathbf{Interval}, \alpha_{\mathbf{IS}}, \gamma_{\mathbf{IS}}, \mathcal{P}(\mathbf{Sign}))$ is a Galois connection. ■

The mundane approach to correctness. We shall now return to further motivating that the connection between L and M should be a Galois connection. Recall from Section 4.1 that the semantic correctness of an analysis may be expressed by a *correctness relation* R between the values of V and the properties of L or it may be expressed by a representation function β mapping the values of V to their description in L . When we replace L with some other complete lattice M we would obviously like the correctness results still to hold. We shall now see that if there is a Galois connection (L, α, γ, M) between L and M then we can construct a correctness relation between V and M and a representation function from V to M .

Let us first focus on the correctness relations. So let $R : V \times L \rightarrow \{\text{true}, \text{false}\}$ be a correctness relation that satisfies the conditions (4.4) and (4.5) of Section 4.1. Further let (L, α, γ, M) be a Galois connection between the complete lattices L and M . It is then natural to define $S : V \times M \rightarrow \{\text{true}, \text{false}\}$ by

$$v S m \text{ iff } v R (\gamma(m))$$

and we shall now argue that this is a correctness relation between V and M fulfilling conditions corresponding to (4.4) and (4.5). We have

$$\begin{aligned} (v S m_1) \wedge m_1 \sqsubseteq m_2 &\Rightarrow v R (\gamma(m_1)) \wedge \gamma(m_1) \sqsubseteq \gamma(m_2) \\ &\Rightarrow v R (\gamma(m_2)) \\ &\Rightarrow v S m_2 \end{aligned}$$

using that γ is monotone and that R satisfies condition (4.4); this shows that S also satisfies condition (4.4). Also for all $M' \subseteq M$ we have

$$\begin{aligned} (\forall m \in M' : v S m) &\Rightarrow (\forall m \in M' : v R (\gamma(m))) \\ &\Rightarrow v R (\bigsqcap \{\gamma(m) \mid m \in M'\}) \\ &\Rightarrow v R (\gamma(\bigsqcap M')) \\ &\Rightarrow v S (\bigsqcap M') \end{aligned}$$

using that γ is completely multiplicative (Lemma 4.22) and that R satisfies condition (4.5); this shows that S also satisfies condition (4.5). Hence S defines a correctness relation between V and M .

Continuing the above line of reasoning assume now that R is *generated by the representation function* $\beta : V \rightarrow L$, i.e. $v R l \Leftrightarrow \beta(v) \sqsubseteq l$. Since (L, α, γ, M) is a Galois connection and hence an adjunction (Proposition 4.20) we may calculate

$$\begin{aligned} v S m &\Leftrightarrow v R (\gamma(m)) \\ &\Leftrightarrow \beta(v) \sqsubseteq \gamma(m) \\ &\Leftrightarrow (\alpha \circ \beta)(v) \sqsubseteq m \end{aligned}$$

showing that S is *generated by* $\alpha \circ \beta : V \rightarrow M$. This shows how the Galois connection facilitates the definition of correctness relations and representation functions and concludes our motivation for why it is natural to require that the connection between L and M is specified by a Galois connection.

4.3.2 Galois Insertions

For a Galois connection (L, α, γ, M) there may be several elements of M that describe the same element of L , i.e. γ need not be injective, and this means that M may contain elements that are not relevant for the approximation of L .

The concept of Galois insertion is intended to rectify this: (L, α, γ, M) is a *Galois insertion* between the complete lattices $L = (L, \sqsubseteq)$ and $M = (M, \sqsubseteq)$ if and only if

$$\alpha : L \rightarrow M \text{ and } \gamma : M \rightarrow L \text{ are monotone functions}$$

that satisfy:

$$\begin{aligned} \gamma \circ \alpha &\sqsupseteq \lambda l.l \\ \alpha \circ \gamma &= \lambda m.m \end{aligned}$$

Thus we now require that we do not lose precision by first doing a concretisation and next an abstraction. As a consequence M cannot contain elements




Figure 4.8: The Galois insertion (L, α, γ, M) .

that do not describe elements of L , i.e. M does not contain superfluous elements.

Example 4.26 The calculations of Example 4.19 show that

$$(\mathcal{P}(\mathbf{Z}), \alpha_{\mathbf{ZI}}, \gamma_{\mathbf{ZI}}, \mathbf{Interval})$$

is indeed a Galois insertion: we start with an interval, use $\gamma_{\mathbf{ZI}}$ to determine the set of integers it describes and next use $\alpha_{\mathbf{ZI}}$ to determine the smallest interval containing this set and we get exactly the same interval as we started with. ■

The concept of a Galois insertion is illustrated in Figure 4.8 and is further clarified by:

Lemma 4.27 For a Galois connection (L, α, γ, M) the following claims are equivalent:

- (i) (L, α, γ, M) is a Galois insertion;
- (ii) α is surjective: $\forall m \in M : \exists l \in L : \alpha(l) = m$;
- (iii) γ is injective: $\forall m_1, m_2 \in M : \gamma(m_1) = \gamma(m_2) \Rightarrow m_1 = m_2$; and
- (iv) γ is an order-similarity: $\forall m_1, m_2 \in M : \gamma(m_1) \sqsubseteq \gamma(m_2) \Leftrightarrow m_1 \sqsubseteq m_2$. ■

Proof First we prove that (i) \Rightarrow (iii). For this we calculate

$$\gamma(m_1) = \gamma(m_2) \Rightarrow \alpha(\gamma(m_1)) = \alpha(\gamma(m_2)) \Rightarrow m_1 = m_2$$

where the last step follows from (i).

Next we prove that (iii) \Rightarrow (ii). For $m \in M$ we have $\gamma(\alpha(\gamma(m))) = \gamma(m)$ by Fact 4.24 and hence $\alpha(\gamma(m)) = m$ by (iii).

Now we prove that (ii) \Rightarrow (iv). That $m_1 \sqsubseteq m_2 \Rightarrow \gamma(m_1) \sqsubseteq \gamma(m_2)$ is immediate since γ is monotone. Next suppose that $\gamma(m_1) \sqsubseteq \gamma(m_2)$; by monotonicity of α this gives $\alpha(\gamma(m_1)) \sqsubseteq \alpha(\gamma(m_2))$; using (ii) we can write $m_1 = \alpha(l_1)$ and $m_2 = \alpha(l_2)$ for some $l_1, l_2 \in L$; using Fact 4.24 this then gives $m_1 \sqsubseteq m_2$ as desired.

Finally we prove that (iv) \Rightarrow (i). For this we calculate

$$\alpha(\gamma(m_1)) \sqsubseteq m_2 \Leftrightarrow \gamma(m_1) \sqsubseteq \gamma(m_2) \Leftrightarrow m_1 \sqsubseteq m_2$$

where the first step is using Proposition 4.20 and the last step is using (iv). This shows that $\alpha(\gamma(m_1))$ and m_1 have the same upper bounds and hence are equal. ■

Lemma 4.27 has an interesting consequence for a Galois connection given by an extraction function $\eta : V \rightarrow D$: it is a Galois insertion if and only if η is surjective.

Example 4.28 Consider the complete lattices $(\mathcal{P}(\mathbf{Z}), \subseteq)$ and $(\mathcal{P}(\mathbf{Sign} \times \mathbf{Parity}), \subseteq)$ where $\mathbf{Sign} = \{-, 0, +\}$ as before and $\mathbf{Parity} = \{\text{odd}, \text{even}\}$. Define the extraction function $\text{signparity} : \mathbf{Z} \rightarrow \mathbf{Sign} \times \mathbf{Parity}$ by:

$$\text{signparity}(z) = \begin{cases} (\text{sign}(z), \text{odd}) & \text{if } z \text{ is odd} \\ (\text{sign}(z), \text{even}) & \text{if } z \text{ is even} \end{cases}$$

This gives rise to a Galois connection $(\mathcal{P}(\mathbf{Z}), \alpha_{\text{signparity}}, \gamma_{\text{signparity}}, \mathcal{P}(\mathbf{Sign} \times \mathbf{Parity}))$. The property $(0, \text{odd})$ describes no integers so clearly signparity is not surjective and we have an example of a Galois connection that is not a Galois insertion. ■

Construction of Galois insertions. Given a Galois connection (L, α, γ, M) it is always possible to obtain a Galois insertion by enforcing that the concretisation function γ is injective. Basically, this amounts to removing elements from the complete lattice M using a *reduction operator*

$$\varsigma : M \rightarrow M$$

defined from the Galois connection. We have the following result:

Proposition 4.29

Let (L, α, γ, M) be a Galois connection and define the reduction operator $\varsigma : M \rightarrow M$ by

$$\varsigma(m) = \bigsqcap \{m' \mid \gamma(m) = \gamma(m')\}$$

Then $\varsigma[M] = (\{\varsigma(m) \mid m \in M\}, \sqsubseteq_M)$ is a complete lattice and $(L, \alpha, \gamma, \varsigma[M])$ is a Galois insertion.




Figure 4.9: The reduction operator $\varsigma : M \rightarrow M$.

This is illustrated in Figure 4.9: The idea is that two elements of M are identified by ς if they map to the same value by γ ; in particular, m and $\alpha(\gamma(m))$ will be identified.

In preparation for the proof we shall first show the following two results:

Fact 4.30 Let (L, α, γ, M) be a Galois connection. Then

$$\alpha[L] = (\{\alpha(l) \mid l \in L\}, \subseteq_M)$$

is a complete lattice. ■

Proof Clearly $\alpha[L] \subseteq M$ is partially ordered by the ordering \subseteq_M of $M = (M, \subseteq_M)$. We now want to show for $M' \subseteq \alpha[L] \subseteq M$ that

$$\bigsqcup_{\alpha[L]} M' = \bigsqcup_M M'$$

We first show that $\bigsqcup_M M' \in \alpha[L]$ which means that there exists $l \in L$ such that $\alpha(l) = \bigsqcup_M M'$. For this take $l = \bigsqcup_L \gamma[M']$ which clearly exists in L . Since $\alpha \circ \gamma \circ \alpha = \alpha$ (Fact 4.24) and α is completely additive (Lemma 4.22) we then have

$$\begin{aligned} \alpha(l) &= \alpha(\bigsqcup_L \gamma[M']) \\ &= \bigsqcup_M \alpha[\gamma[M']] \\ &= \bigsqcup_M M' \end{aligned}$$

thus showing $\bigsqcup_M M' \in \alpha[L]$. Since $\bigsqcup_M M'$ is the least upper bound of M' in M it follows that it also is in $\alpha[L]$. By Lemma A.2 we then have the result. ■

Fact 4.31 If (L, α, γ, M) is a Galois connection and

$$\varsigma(m) = \prod \{m' \mid \gamma(m') = \gamma(m)\}$$

then

$$\zeta(m) = \alpha(\gamma(m))$$

and hence $\alpha[L] = \zeta[M]$. ■

Proof We have $\zeta(m) \sqsubseteq \alpha(\gamma(m))$ because $\gamma(m) = \gamma(\alpha(\gamma(m)))$ using Fact 4.24. That $\alpha(\gamma(m)) \sqsubseteq \zeta(m)$ is equivalent to $\gamma(m) \sqsubseteq \gamma(\zeta(m))$ using Proposition 4.20 and $\gamma(\zeta(m)) = \bigcap \{\gamma(m') \mid \gamma(m') = \gamma(m)\} = \gamma(m)$ follows from Lemma 4.22.

Next consider $\alpha(l)$ for $l \in L$; by Fact 4.24 we have $\alpha(l) = \alpha(\gamma(\alpha(l)))$ and hence $\alpha(l) = \zeta(\alpha(l))$; this shows $\alpha[L] \subseteq \zeta[M]$. Finally consider $\zeta(m)$ for $m \in M$; then $\zeta(m) = \alpha(\gamma(m))$; this shows $\zeta[M] \subseteq \alpha[M]$. ■

We now turn to the proof of Proposition 4.29:

Proof Facts 4.30 and 4.31 give that $\zeta[M] = \alpha[L]$ is a complete lattice. Since (L, α, γ, M) is an adjunction (Proposition 4.20) it follows that also $(L, \alpha, \gamma, \alpha[L])$ is. Since α is surjective on $\alpha[L]$ it follows from Lemma 4.27 that we have a Galois insertion. ■

Reduction operators defined by extraction functions. We shall now specialise the construction of Proposition 4.29 to the setting where a Galois connection $(\mathcal{P}(V), \alpha_\eta, \gamma_\eta, \mathcal{P}(D))$ is given by an *extraction function* $\eta : V \rightarrow D$. Then the reduction operator ζ_η is given by

$$\zeta_\eta(D') = D' \cap \eta[V]$$

where $\eta[V] = \{\eta(v) \mid v \in V\}$ denotes the image of η and furthermore $\zeta_\eta[\mathcal{P}(D)]$ is isomorphic (see Appendix A) to $\mathcal{P}(\eta[V])$ (see Exercise 4.14). The resulting Galois insertion will therefore be isomorphic to $(\mathcal{P}(V), \alpha_\eta, \gamma_\eta, \mathcal{P}(\eta[V]))$: Formally, the Galois connection $(L_1, \alpha_1, \gamma_1, M_1)$ is *isomorphic* to the Galois connection $(L_2, \alpha_2, \gamma_2, M_2)$ when there are isomorphisms $\theta_L : L_1 \rightarrow L_2$ and $\theta_M : M_1 \rightarrow M_2$ (see Appendix A) such that $\alpha_2 = \theta_M \circ \alpha_1 \circ \theta_L^{-1}$ and $\gamma_2 = \theta_L \circ \gamma_1 \circ \theta_M^{-1}$.

Example 4.32 Returning to Example 4.28 we can use the above technique to construct a Galois insertion. Now

$$\text{signparity}[Z] = \{(-, \text{odd}), (-, \text{even}), (0, \text{even}), (+, \text{odd}), (+, \text{even})\}$$

showing that $(\mathcal{P}(Z), \alpha_{\text{signparity}}, \gamma_{\text{signparity}}, \mathcal{P}(\text{signparity}[Z]))$ is the resulting Galois insertion. ■

4.4 Systematic Design of Galois Connections

Sequential composition. When developing a program analysis it is often useful to do so in stages: The starting point will typically be a complete lattice (L_0, \sqsubseteq) fairly closely related to the semantics; an example is

$(\mathcal{P}(V), \sqsubseteq)$. We may then decide to use a more approximate set of properties and introduce the complete lattice (L_1, \sqsubseteq) related to L_0 by a Galois connection $(L_0, \alpha_1, \gamma_1, L_1)$. This step can then be repeated any number of times: We replace one complete lattice L_i of properties with a more approximate complete lattice (L_{i+1}, \sqsubseteq) related to L_i by a Galois connection $(L_i, \alpha_{i+1}, \gamma_{i+1}, L_{i+1})$. This process will stop when we have an analysis with the required computational properties. So the situation can be depicted as follows:

$$\begin{array}{ccccccc}
 & \xleftarrow{\gamma_1} & & \xleftarrow{\gamma_2} & & \xleftarrow{\gamma_3} & \cdots & \xleftarrow{\gamma_k} & \\
 L_0 & & L_1 & & L_2 & & \cdots & & L_k \\
 & \xrightarrow{\alpha_1} & & \xrightarrow{\alpha_2} & & \xrightarrow{\alpha_3} & & \xrightarrow{\alpha_k} &
 \end{array}$$

The above sequence of approximations of the analysis could as well have been done in one step, i.e. the “functional composition” of two Galois connections is also a Galois connection. Formally, if $(L_0, \alpha_1, \gamma_1, L_1)$ and $(L_1, \alpha_2, \gamma_2, L_2)$ are Galois connections, then also

$$(L_0, \alpha_2 \circ \alpha_1, \gamma_1 \circ \gamma_2, L_2)$$

is a Galois connection. To verify this we simply observe that $\alpha_2(\alpha_1(l_0)) \sqsubseteq l_2 \Leftrightarrow \alpha_1(l_0) \sqsubseteq \gamma_2(l_2) \Leftrightarrow l_0 \sqsubseteq \gamma_1(\gamma_2(l_2))$ and, using Proposition 4.20, this shows the result. A similar result holds for Galois insertions because the functional composition of surjective functions yields a surjective function.

Each of the Galois connections $(L_i, \alpha_{i+1}, \gamma_{i+1}, L_{i+1})$ may have been obtained by combining other Galois connections and we shall shortly introduce a number of techniques for doing so. We shall illustrate these techniques in a sequence of examples that in the end will give us a finite complete lattice that has turned out to be very useful in practice for performing an *Array Bound Analysis*.

Example 4.33 One of the components in the Array Bound Analysis is concerned with approximating the difference in magnitude between two numbers (typically the bound and the index). We shall proceed in two stages: First we shall approximate pairs (z_1, z_2) of integers by their difference in magnitude $|z_1| - |z_2|$ and next we shall further approximate this difference using a finite lattice. The two Galois connections will be defined by extraction functions and they will then be combined by taking their functional composition.

The first stage is specified by the Galois connection

$$(\mathcal{P}(\mathbf{Z} \times \mathbf{Z}), \alpha_{\text{diff}}, \gamma_{\text{diff}}, \mathcal{P}(\mathbf{Z}))$$

where $\text{diff} : \mathbf{Z} \times \mathbf{Z} \rightarrow \mathbf{Z}$ is the extraction function calculating the difference in magnitude:

$$\text{diff}(z_1, z_2) = |z_1| - |z_2|$$

The abstraction and concretisation functions α_{diff} and γ_{diff} will then be

$$\begin{aligned}\alpha_{\text{diff}}(ZZ) &= \{|z_1| - |z_2| \mid (z_1, z_2) \in ZZ\} \\ \gamma_{\text{diff}}(Z) &= \{(z_1, z_2) \mid |z_1| - |z_2| \in Z\}\end{aligned}$$

for $ZZ \subseteq \mathbf{Z} \times \mathbf{Z}$ and $Z \subseteq \mathbf{Z}$.

The second stage is specified by the Galois connection

$$(\mathcal{P}(\mathbf{Z}), \alpha_{\text{range}}, \gamma_{\text{range}}, \mathcal{P}(\mathbf{Range}))$$

where $\mathbf{Range} = \{<-1, -1, 0, +1, >+1\}$. The extraction function $\text{range} : \mathbf{Z} \rightarrow \mathbf{Range}$ clarifies the meaning of the elements of \mathbf{Range} :

$$\text{range}(z) = \begin{cases} <-1 & \text{if } z < -1 \\ -1 & \text{if } z = -1 \\ 0 & \text{if } z = 0 \\ +1 & \text{if } z = 1 \\ >+1 & \text{if } z > 1 \end{cases}$$

The abstraction and concretisation functions α_{range} and γ_{range} will then be

$$\begin{aligned}\alpha_{\text{range}}(Z) &= \{\text{range}(z) \mid z \in Z\} \\ \gamma_{\text{range}}(R) &= \{z \mid \text{range}(z) \in R\}\end{aligned}$$

for $Z \subseteq \mathbf{Z}$ and $R \subseteq \mathbf{Range}$.

We then have that the functional composition

$$(\mathcal{P}(\mathbf{Z} \times \mathbf{Z}), \alpha_{\mathbf{R}}, \gamma_{\mathbf{R}}, \mathcal{P}(\mathbf{Range}))$$

where $\alpha_{\mathbf{R}} = \alpha_{\text{range}} \circ \alpha_{\text{diff}}$ and $\gamma_{\mathbf{R}} = \gamma_{\text{diff}} \circ \gamma_{\text{range}}$, is a Galois connection. We obtain the following formulae for the abstraction and concretisation functions:

$$\begin{aligned}\alpha_{\mathbf{R}}(ZZ) &= \{\text{range}(|z_1| - |z_2|) \mid (z_1, z_2) \in ZZ\} \\ \gamma_{\mathbf{R}}(R) &= \{(z_1, z_2) \mid \text{range}(|z_1| - |z_2|) \in R\}\end{aligned}$$

Using Exercise 4.15 we see that this is the Galois connection specified by the extraction function $\text{range} \circ \text{diff} : \mathbf{Z} \times \mathbf{Z} \rightarrow \mathbf{Range}$. ■

A catalogue of combination techniques. We have seen that the “functional (or sequential) composition” of Galois connections gives rise to a new Galois connection. It is also useful to be able to do “parallel combinations” and this will be the topic for the remainder of this section. We may have analyses for the *individual* components of a composite structure and may want to combine them into an analysis for the whole structure; we shall see two techniques for that: the independent attribute method and the relational method. We shall also show how Galois connections can be used

to approximate the total function spaces and the monotone function spaces. Alternatively, we may have several analyses of the *same* structure and may want to combine them into one analysis and here the direct product and the direct tensor product allow us to do that. Using the notion of Galois insertions this leads to a study of the reduced product and the reduced tensor product.

The benefit of having such a catalogue of techniques is that a relative small set of "basic" analyses, whose correctness have been established, can be used to construct rather sophisticated analyses and, from an implementation point of view, opens up the possibility of reusing existing implementations. Finally, it is worth stressing that the complete lattices used in program analysis obviously can be constructed without using these techniques but that it often provides additional insights to view them as combinations of simpler Galois connection.

4.4.1 Component-wise Combinations

The first techniques we shall consider are applicable when we have several analyses of *individual* components of a structure and we want to combine them into a single analysis.

Independent attribute method. Let $(L_1, \alpha_1, \gamma_1, M_1)$ and $(L_2, \alpha_2, \gamma_2, M_2)$ be Galois connections. The *independent attribute method* will then give rise to a Galois connection

$$(L_1 \times L_2, \alpha, \gamma, M_1 \times M_2)$$

where:

$$\begin{aligned}\alpha(l_1, l_2) &= (\alpha_1(l_1), \alpha_2(l_2)) \\ \gamma(m_1, m_2) &= (\gamma_1(m_1), \gamma_2(m_2))\end{aligned}$$

To see that this indeed does define a Galois connection we simply calculate

$$\begin{aligned}\alpha(l_1, l_2) \sqsubseteq (m_1, m_2) &\Leftrightarrow (\alpha_1(l_1), \alpha_2(l_2)) \sqsubseteq (m_1, m_2) \\ &\Leftrightarrow \alpha_1(l_1) \sqsubseteq m_1 \wedge \alpha_2(l_2) \sqsubseteq m_2 \\ &\Leftrightarrow l_1 \sqsubseteq \gamma_1(m_1) \wedge l_2 \sqsubseteq \gamma_2(m_2) \\ &\Leftrightarrow (l_1, l_2) \sqsubseteq (\gamma_1(m_1), \gamma_2(m_2)) \\ &\Leftrightarrow (l_1, l_2) \sqsubseteq \gamma(m_1, m_2)\end{aligned}$$

and use Proposition 4.20. A similar result holds for Galois insertions (Exercise 4.17).

Example 4.34 The Array Bound Analysis will contain a component that performs a Detection of Signs Analysis on pairs of integers. As a starting point, we take the Galois connection

$$(\mathcal{P}(\mathbf{Z}), \alpha_{\text{sign}}, \gamma_{\text{sign}}, \mathcal{P}(\mathbf{Sign}))$$

specified by the extraction function `sign` in Example 4.21. It can be used to analyse both components of a pair of integers so using the independent attribute method we will get a Galois connection

$$(\mathcal{P}(\mathbf{Z}) \times \mathcal{P}(\mathbf{Z}), \alpha_{\text{SS}}, \gamma_{\text{SS}}, \mathcal{P}(\mathbf{Sign}) \times \mathcal{P}(\mathbf{Sign}))$$

where α_{SS} and γ_{SS} are given by

$$\begin{aligned} \alpha_{\text{SS}}(Z_1, Z_2) &= (\{\text{sign}(z) \mid z \in Z_1\}, \{\text{sign}(z) \mid z \in Z_2\}) \\ \gamma_{\text{SS}}(S_1, S_2) &= (\{z \mid \text{sign}(z) \in S_1\}, \{z \mid \text{sign}(z) \in S_2\}) \end{aligned}$$

where $Z_i \subseteq \mathbf{Z}$ and $S_i \subseteq \mathbf{Sign}$.

This Galois connection cannot be described using an extraction function because neither $\mathcal{P}(\mathbf{Z}) \times \mathcal{P}(\mathbf{Z})$ nor $\mathcal{P}(\mathbf{Sign}) \times \mathcal{P}(\mathbf{Sign})$ is a powerset. However, they are both isomorphic to powersets:

$$\begin{aligned} \mathcal{P}(\mathbf{Z}) \times \mathcal{P}(\mathbf{Z}) &\cong \mathcal{P}(\{1, 2\} \times \mathbf{Z}) \\ \mathcal{P}(\mathbf{Sign}) \times \mathcal{P}(\mathbf{Sign}) &\cong \mathcal{P}(\{1, 2\} \times \mathbf{Sign}) \end{aligned}$$

By defining the extraction function `twosigns` : $\{1, 2\} \times \mathbf{Z} \rightarrow \{1, 2\} \times \mathbf{Sign}$ using the formula `twosigns(i, z) = (i, sign(z))` we obtain a Galois connection

$$(\mathcal{P}(\{1, 2\} \times \mathbf{Z}), \alpha_{\text{twosigns}}, \gamma_{\text{twosigns}}, \mathcal{P}(\{1, 2\} \times \mathbf{Sign}))$$

that is isomorphic to $(\mathcal{P}(\mathbf{Z}) \times \mathcal{P}(\mathbf{Z}), \alpha_{\text{SS}}, \gamma_{\text{SS}}, \mathcal{P}(\mathbf{Sign}) \times \mathcal{P}(\mathbf{Sign}))$.

In general the independent attribute method often leads to imprecision. An expression like $(x, -x)$ in the source language may have a value in $\{(z, -z) \mid z \in \mathbf{Z}\}$ but in the present setting where we use $\mathcal{P}(\mathbf{Z}) \times \mathcal{P}(\mathbf{Z})$ to represent sets of pairs of integers we cannot do better than representing $\{(z, -z) \mid z \in \mathbf{Z}\}$ by (\mathbf{Z}, \mathbf{Z}) and hence the best property describing it in the analysis of Example 4.34 will be $\alpha_{\text{SS}}(\mathbf{Z}, \mathbf{Z}) = (\{-, 0, +\}, \{-, 0, +\})$. Thus we lose all information about the relative signs of the two components. ■

Relational method. In the independent attribute method there is absolutely no interplay between the two pairs of abstraction and concretisation functions. It is possible to do better by allowing the two components of the analysis to interact with one another so as to get more precise descriptions.

Let $(\mathcal{P}(V_1), \alpha_1, \gamma_1, \mathcal{P}(D_1))$ and $(\mathcal{P}(V_2), \alpha_2, \gamma_2, \mathcal{P}(D_2))$ be Galois connections. The *relational method* will give rise to the Galois connection

$$(\mathcal{P}(V_1 \times V_2), \alpha, \gamma, \mathcal{P}(D_1 \times D_2))$$

where

$$\begin{aligned}\alpha(VV) &= \bigcup \{ \alpha_1(\{v_1\}) \times \alpha_2(\{v_2\}) \mid (v_1, v_2) \in VV \} \\ \gamma(DD) &= \{ (v_1, v_2) \mid \alpha_1(\{v_1\}) \times \alpha_2(\{v_2\}) \subseteq DD \}\end{aligned}$$

where $VV \subseteq V_1 \times V_2$ and $DD \subseteq D_1 \times D_2$. Let us check that this does indeed define a Galois connection. From the definition it is immediate that α is completely additive and hence that there exists a Galois connection (Lemma 4.23). It remains to show that γ (as defined above) is the upper adjoint of α . For this we can use Lemma 4.22 and calculate

$$\begin{aligned}\gamma(DD) &= \{ (v_1, v_2) \mid \alpha_1(\{v_1\}) \times \alpha_2(\{v_2\}) \subseteq DD \} \\ &= \{ (v_1, v_2) \mid \alpha(\{(v_1, v_2)\}) \subseteq DD \} \\ &= \bigcup \{ VV \mid \alpha(VV) \subseteq DD \}\end{aligned}$$

where we have used that α is completely additive. This shows the required result.

It is instructive to see how the relational method is simplified if the Galois connections $(\mathcal{P}(V_i), \alpha_i, \gamma_i, \mathcal{P}(D_i))$ are given by extraction functions $\eta_i : V_i \rightarrow D_i$, i.e. if $\alpha_i(V'_i) = \{ \eta_i(v_i) \mid v_i \in V'_i \}$ and $\gamma_i(D'_i) = \{ v_i \mid \eta_i(v_i) \in D'_i \}$. We then have

$$\begin{aligned}\alpha(VV) &= \{ (\eta_1(v_1), \eta_2(v_2)) \mid (v_1, v_2) \in VV \} \\ \gamma(DD) &= \{ (v_1, v_2) \mid (\eta_1(v_1), \eta_2(v_2)) \in DD \}\end{aligned}$$

which also can be obtained directly from the extraction function $\eta : V_1 \times V_2 \rightarrow D_1 \times D_2$ defined by $\eta(v_1, v_2) = (\eta_1(v_1), \eta_2(v_2))$.

Example 4.35 Let us return to Example 4.34 and show how the relational method can be used to construct a more precise analysis. We will now get a Galois connection

$$(\mathcal{P}(\mathbf{Z} \times \mathbf{Z}), \alpha_{SS'}, \gamma_{SS'}, \mathcal{P}(\mathbf{Sign} \times \mathbf{Sign}))$$

where $\alpha_{SS'}$ and $\gamma_{SS'}$ are given by

$$\begin{aligned}\alpha_{SS'}(ZZ) &= \{ (\text{sign}(z_1), \text{sign}(z_2)) \mid (z_1, z_2) \in ZZ \} \\ \gamma_{SS'}(SS) &= \{ (z_1, z_2) \mid (\text{sign}(z_1), \text{sign}(z_2)) \in SS \}\end{aligned}$$

where $ZZ \subseteq \mathbf{Z} \times \mathbf{Z}$ and $SS \subseteq \mathbf{Sign} \times \mathbf{Sign}$. This corresponds to using an extraction function $\text{twosigns}' : \mathbf{Z} \times \mathbf{Z} \rightarrow \mathbf{Sign} \times \mathbf{Sign}$ defined by $\text{twosigns}'(z_1, z_2) = (\text{sign}(z_1), \text{sign}(z_2))$.

Once again consider the expression $(x, -x)$ in the source language that has a value in $\{(z, -z) \mid z \in \mathbf{Z}\}$. In the present setting $\{(z, -z) \mid z \in \mathbf{Z}\}$ is an

element of $\mathcal{P}(\mathbf{Z} \times \mathbf{Z})$ and it is described by the set $\alpha_{SS'}(\{(z, -z) \mid z \in \mathbf{Z}\}) = \{(-, +), (0, 0), (+, -)\}$ of $\mathcal{P}(\mathbf{Sign} \times \mathbf{Sign})$. Hence the information about the relative signs of the two components is preserved. This will be the Galois connection that we will use in our further development of the Array Bound Analysis. ■

The above treatment of the relational method can be extended in a very general way. Let us write $\mathcal{P}(V_1) \otimes \mathcal{P}(V_2)$ for $\mathcal{P}(V_1 \times V_2)$ and similarly $\mathcal{P}(D_1) \otimes \mathcal{P}(D_2)$ for $\mathcal{P}(D_1 \times D_2)$. It is possible to perform a more general development using a notion of *tensor product* for which $L_1 \otimes L_2$ exists even when the complete lattices L_1 and L_2 are not powersets. We refer to the Concluding Remarks for information about this.

Total function space. In Appendix A it is established that if L is a complete lattice then so is the *total function space* $S \rightarrow L$ for S being a set. We have a similar result for Galois connections:

Let (L, α, γ, M) be a Galois connection and let S be a set. Then we obtain a Galois connection

$$(S \rightarrow L, \alpha', \gamma', S \rightarrow M)$$

by taking

$$\begin{aligned}\alpha'(f) &= \alpha \circ f \\ \gamma'(g) &= \gamma \circ g\end{aligned}$$

To see this we first observe that α' and γ' are monotone functions because α and γ are; furthermore

$$\begin{aligned}\gamma'(\alpha'(f)) &= \gamma \circ \alpha \circ f \sqsupseteq f \\ \alpha'(\gamma'(g)) &= \alpha \circ \gamma \circ g \sqsubseteq g\end{aligned}$$

follow since (L, α, γ, M) is a Galois connection. A similar result holds for Galois insertions.

Example 4.36 Assume that we have some analysis mapping the program variables to properties from the complete lattice L , i.e. operating on the abstract states $\mathbf{Var} \rightarrow L$. Given a Galois connection (L, α, γ, M) the above construction will show us how the abstract states of $\mathbf{Var} \rightarrow L$ are approximated by the abstract states of $\mathbf{Var} \rightarrow M$. ■

Monotone function space. In Appendix A it is established that the *monotone function space* between two complete lattices is a complete lattice. We have a similar result for Galois connections:

Let $(L_1, \alpha_1, \gamma_1, M_1)$ and $(L_2, \alpha_2, \gamma_2, M_2)$ be Galois connections. Then we obtain the Galois connection

$$(L_1 \rightarrow L_2, \alpha, \gamma, M_1 \rightarrow M_2)$$

by taking

$$\alpha(f) = \alpha_2 \circ f \circ \gamma_1$$

$$\gamma(g) = \gamma_2 \circ g \circ \alpha_1$$


To check this we first observe that the functions α and γ are monotone because α_2 and γ_2 are; next we calculate

$$\gamma(\alpha(f)) = (\gamma_2 \circ \alpha_2) \circ f \circ (\gamma_1 \circ \alpha_1) \sqsupseteq f$$

$$\alpha(\gamma(g)) = (\alpha_2 \circ \gamma_2) \circ g \circ (\alpha_1 \circ \gamma_1) \sqsubseteq g$$

using the monotonicity of $f : L_1 \rightarrow L_2$ and $g : M_1 \rightarrow M_2$ together with (4.8) and (4.9). A similar result holds for Galois insertions (Exercise 4.17).

This construction is illustrated by the following commuting diagrams:



4.4.2 Other Combinations

So far our constructions have shown how to combine Galois connections dealing with individual components of the data into Galois connections dealing with composite data. We shall now show how two analyses dealing with the *same* data can be combined into one analysis; this amounts to performing two analyses in parallel. We shall consider two variants of this analysis, one “corresponding” to the independent attribute method and one “corresponding” to the relational method.

Direct product. Let $(L, \alpha_1, \gamma_1, M_1)$ and $(L, \alpha_2, \gamma_2, M_2)$ be Galois connections. The *direct product* of the two Galois connection will be the Galois connection

$$(L, \alpha, \gamma, M_1 \times M_2)$$

where α and γ are given by:

$$\begin{aligned} \alpha(l) &= (\alpha_1(l), \alpha_2(l)) \\ \gamma(m_1, m_2) &= \gamma_1(m_1) \sqcap \gamma_2(m_2) \end{aligned}$$

To see that this indeed defines a Galois connection we calculate

$$\begin{aligned} \alpha(l) \sqsubseteq (m_1, m_2) &\Leftrightarrow \alpha_1(l) \sqsubseteq m_1 \wedge \alpha_2(l) \sqsubseteq m_2 \\ &\Leftrightarrow l \sqsubseteq \gamma_1(m_1) \wedge l \sqsubseteq \gamma_2(m_2) \\ &\Leftrightarrow l \sqsubseteq \gamma(m_1, m_2) \end{aligned}$$

and then use Proposition 4.20 to get the result.

Example 4.37 Let us consider how this construction can be used to combine the detection of signs analysis for pairs of integers given in Example 4.35 with the analysis of difference in magnitude given in Example 4.33. We get the Galois connection

$$(\mathcal{P}(\mathbf{Z} \times \mathbf{Z}), \alpha_{\text{SSR}}, \gamma_{\text{SSR}}, \mathcal{P}(\mathbf{Sign} \times \mathbf{Sign}) \times \mathcal{P}(\mathbf{Range}))$$

where α_{SSR} and γ_{SSR} are given by:

$$\begin{aligned} \alpha_{\text{SSR}}(ZZ) &= (\{\text{sign}(z_1), \text{sign}(z_2) \mid (z_1, z_2) \in ZZ\}, \\ &\quad \{\text{range}(|z_1| - |z_2|) \mid (z_1, z_2) \in ZZ\}) \\ \gamma_{\text{SSR}}(SS, R) &= \{(z_1, z_2) \mid (\text{sign}(z_1), \text{sign}(z_2)) \in SS\} \\ &\quad \cap \{(z_1, z_2) \mid \text{range}(|z_1| - |z_2|) \in R\} \end{aligned}$$

Note that the expression $(x, 3 * x)$ in the source language has a value in $\{(z, 3 * z) \mid z \in \mathbf{Z}\}$ which is described by $\alpha_{\text{SSR}}(\{(z, 3 * z) \mid z \in \mathbf{Z}\}) = (\{(-, -), (0, 0), (+, +)\}, \{0, <-1\})$. Thus we do not exploit the fact that if the pair is described by $(0, 0)$ then the difference in magnitude will indeed be described by 0 whereas if the pair is described by $(-, -)$ or $(+, +)$ then the difference in magnitude will indeed be described by <-1 . ■

Direct tensor product. In the direct product there is no interplay between the two abstraction functions and as we saw above this gives rise to the same loss of precision as in the independent attribute method. It is possible to do better by letting the two components interact with one another. Again we shall only consider the simple case of powersets so let $(\mathcal{P}(V), \alpha_i, \gamma_i, \mathcal{P}(D_i))$ be Galois connections. Then the *direct tensor product* is the Galois connection

$$(\mathcal{P}(V), \alpha, \gamma, \mathcal{P}(D_1 \times D_2))$$

where α and γ are defined by:

$$\begin{aligned} \alpha(V') &= \bigcup \{\alpha_1(\{v\}) \times \alpha_2(\{v\}) \mid v \in V'\} \\ \gamma(DD) &= \{v \mid \alpha_1(\{v\}) \times \alpha_2(\{v\}) \subseteq DD\} \end{aligned}$$

where $V' \subseteq V$ and $DD \subseteq D_1 \times D_2$. To verify that this defines a Galois connection we calculate

$$\begin{aligned} \alpha(V') \subseteq DD &\Leftrightarrow \forall v \in V' : \alpha_1(\{v\}) \times \alpha_2(\{v\}) \subseteq DD \\ &\Leftrightarrow \forall v \in V' : v \in \gamma(DD) \\ &\Leftrightarrow V' \subseteq \gamma(DD) \end{aligned}$$

and then use Proposition 4.20.

The construction can be simplified if the two Galois connections $(\mathcal{P}(V), \alpha_i, \gamma_i, \mathcal{P}(D_i))$ are given by extraction functions $\eta_i : V \rightarrow D_i$, i.e. if $\alpha_i(V') = \{\eta_i(v) \mid v \in V'\}$ and $\gamma_i(D'_i) = \{v \mid \eta_i(v) \in D'_i\}$. Then we have

$$\begin{aligned} \alpha(V') &= \{(\eta_1(v), \eta_2(v)) \mid v \in V'\} \\ \gamma(DD) &= \{v \mid (\eta_1(v), \eta_2(v)) \in DD\} \end{aligned}$$

which also can be obtained directly from the extraction function $\eta : V \rightarrow D_1 \times D_2$ defined by $\eta(v) = (\eta_1(v), \eta_2(v))$.

Example 4.38 Let us return to Example 4.37 and show how the direct tensor product gives a more precise analysis. We will now get a Galois connection

$$(\mathcal{P}(\mathbf{Z} \times \mathbf{Z}), \alpha_{SSR'}, \gamma_{SSR'}, \mathcal{P}(\mathbf{Sign} \times \mathbf{Sign} \times \mathbf{Range}))$$

where

$$\begin{aligned} \alpha_{SSR'}(ZZ) &= \{(\text{sign}(z_1), \text{sign}(z_2), \text{range}(|z_1| - |z_2|)) \mid (z_1, z_2) \in ZZ\} \\ \gamma_{SSR'}(SSR) &= \{(z_1, z_2) \mid (\text{sign}(z_1), \text{sign}(z_2), \text{range}(|z_1| - |z_2|)) \in SSR\} \end{aligned}$$

for $ZZ \subseteq \mathbf{Z} \times \mathbf{Z}$ and $SSR \subseteq \mathbf{Sign} \times \mathbf{Sign} \times \mathbf{Range}$.

It is worth pointing out that this Galois connection is also obtainable from the extraction function

$$\text{twosignsrange} : \mathbf{Z} \times \mathbf{Z} \rightarrow \mathbf{Sign} \times \mathbf{Sign} \times \mathbf{Range}$$

defined by $\text{twosignsrange}(z_1, z_2) = (\text{sign}(z_1), \text{sign}(z_2), \text{range}(|z_1| - |z_2|))$.

Returning to the precision of the analysis we will now have $\alpha_{SSR'}(\{(z, 3 * z) \mid z \in \mathbf{Z}\}) = \{(-, -, <-1), (0, 0, 0), (+, +, <-1)\}$ and hence have a more precise description than in Example 4.37.

However, it is worth noticing that the above Galois connection is not a Galois insertion. To see this consider the two elements \emptyset and $\{(0, 0, <-1)\}$ of $\mathcal{P}(\mathbf{Sign} \times \mathbf{Sign} \times \mathbf{Range})$ and observe that

$$\gamma_{SSR'}(\emptyset) = \emptyset = \gamma_{SSR'}(\{(0, 0, <-1)\})$$

Thus $\gamma_{SSR'}$ is not injective and hence Lemma 4.27 shows that we do not have a Galois insertion. ■

Reduced product and reduced tensor product. The construction of Proposition 4.29 gives us a general method for turning a Galois connection into a Galois insertion. This technique can now be combined with the other techniques for combining Galois connections and this is of particular interest for the direct product and the direct tensor product.

Let $(L, \alpha_1, \gamma_1, M_1)$ and $(L, \alpha_2, \gamma_2, M_2)$ be Galois connections. Then the *reduced product* is the Galois insertion

$$(L, \alpha, \gamma, \varsigma[M_1 \times M_2])$$

where

$$\alpha(l) = (\alpha_1(l), \alpha_2(l))$$

$$\gamma(m_1, m_2) = \gamma_1(m_1) \sqcap \gamma_2(m_2)$$

$$\varsigma(m_1, m_2) = \bigcap \{(m'_1, m'_2) \mid \gamma_1(m_1) \sqcap \gamma_2(m_2) = \gamma_1(m'_1) \sqcap \gamma_2(m'_2)\}$$

To see that this is indeed a Galois insertion recall that we already know that the direct product $(L, \alpha, \gamma, M_1 \times M_2)$ is a Galois connection and that Proposition 4.29 then shows that $(L, \alpha, \gamma, \varsigma[M_1 \times M_2])$ is a Galois insertion.

Next let $(\mathcal{P}(V), \alpha_i, \gamma_i, \mathcal{P}(D_i))$ be Galois connections. Then the *reduced tensor product* is the Galois insertion

$$(\mathcal{P}(V), \alpha, \gamma, \varsigma[\mathcal{P}(D_1 \times D_2)])$$

where

$$\alpha(V') = \bigcup \{\alpha_1(\{v\}) \times \alpha_2(\{v\}) \mid v \in V'\}$$

$$\gamma(DD) = \{v \mid \alpha_1(\{v\}) \times \alpha_2(\{v\}) \subseteq DD\}$$

$$\varsigma(DD) = \bigcap \{DD' \mid \gamma(DD) = \gamma(DD')\}$$

Again it follows from Proposition 4.29 that this is indeed a Galois insertion.

Example 4.39 Let us return to Example 4.38 where we noted that the complete lattice $\mathcal{P}(\mathbf{Sign} \times \mathbf{Sign} \times \mathbf{Range})$ contains more than one element that describes the empty set of $\mathcal{P}(\mathbf{Z} \times \mathbf{Z})$. The superfluous elements will be removed by the construction of Proposition 4.29. The function $\varsigma_{SSR'}$ will amount to

$$\varsigma_{SSR'}(SSR) = \bigcap \{SSR' \mid \gamma_{SSR'}(SSR) = \gamma_{SSR'}(SSR')\}$$

where $SSR, SSR' \subseteq \mathbf{Sign} \times \mathbf{Sign} \times \mathbf{Range}$. In particular, $\varsigma_{SSR'}$ will map the singleton sets constructed from the 16 elements

$$\begin{array}{lll} (-, 0, <-1), & (-, 0, -1), & (-, 0, 0), \\ (0, -, 0), & (0, -, +1), & (0, -, >+1), \\ (0, 0, <-1), & (0, 0, -1), & (0, 0, +1), & (0, 0, >+1), \\ (0, +, 0), & (0, +, +1), & (0, +, >+1), \\ (+, 0, <-1), & (+, 0, -1), & (+, 0, 0) \end{array}$$

to the empty set. The remaining 29 elements of $\mathbf{Sign} \times \mathbf{Sign} \times \mathbf{Range}$ are

$$\begin{array}{cccccc} (-, -, <-1), & (-, -, -1), & (-, -, 0), & (-, -, +1), & (-, -, >+1), \\ (-, 0, +1), & (-, 0, >+1), & & & \\ (-, +, <-1), & (-, +, -1), & (-, +, 0), & (-, +, +1), & (-, +, >+1), \\ (0, -, <-1), & (0, -, -1), & (0, 0, 0), & (0, +, <-1), & (0, +, -1), \\ (+, -, <-1), & (+, -, -1), & (+, -, 0), & (+, -, +1), & (+, -, >+1), \\ (+, 0, +1), & (+, 0, >+1), & & & \\ (+, +, <-1), & (+, +, -1), & (+, +, 0), & (+, +, +1), & (+, +, >+1) \end{array}$$

and they describe disjoint subsets of $\mathbf{Z} \times \mathbf{Z}$. Let us call the above set of 29 elements for **AB** (for Array Bound); then $\alpha_{\text{SSR}'}[\mathcal{P}(\mathbf{Sign} \times \mathbf{Sign} \times \mathbf{Range})]$ is isomorphic to $\mathcal{P}(\mathbf{AB})$.

To conclude the development of the complete lattice and the associated Galois connection for the Array Bound Analysis we shall simply construct the reduced tensor product of the Galois connections of Examples 4.35 and 4.33. This will yield a *Galois insertion* isomorphic to

$$(\mathcal{P}(\mathbf{Z} \times \mathbf{Z}), \alpha_{\text{SSR}'}, \gamma_{\text{SSR}'}, \mathcal{P}(\mathbf{AB}))$$

Note that from an implementation point of view the last step of the construction has paid off: if we had stopped with the direct tensor product in Example 4.38 then the properties would need 45 bits for their representation whereas now 29 bits suffice.

Summary. The Array Bound Analysis has been designed from three simple Galois connections specified by extraction functions:

- (i) an analysis approximating integers by their sign (Example 4.21),
- (ii) an analysis approximating pairs of integers by their difference in magnitude (Example 4.33), and
- (iii) an analysis approximating integers by their closeness to 0, 1 and -1 (Example 4.33).

We have illustrated different ways of combining these analyses:

- (iv) the relational product of analysis (i) with itself,
- (v) the functional composition of analysis (ii) and (iii), and
- (vi) the reduced tensor product of analysis (iv) and (v).

It is worth noting that because the resulting complete lattice $\mathcal{P}(\mathbf{AB})$ is a powerset then it is indeed possible to obtain the very same Galois insertion using an extraction function $\text{twosignsrange}' : \mathbf{Z} \times \mathbf{Z} \rightarrow \mathbf{AB}$. ■

4.5 Induced Operations

We shall now show that Galois connections are indeed useful for transforming computations into more approximate computations that have better time-, space-, or termination behaviour. We can do so in two different ways. In both cases we assume that we have an analysis using the complete lattice L and that we have a Galois connection (L, α, γ, M) .

One possibility is to replace the analysis using L with an analysis using M . In Subsection 4.5.1 we shall show that if the analysis using M is an upper approximation to the analysis *induced* from L then the correctness properties are preserved. We shall illustrate this approach in Subsection 4.5.2 for the Monotone Frameworks considered in Chapter 2.

An alternative is only to use the complete lattice M for approximating the fixed point computations in L . So rather than performing all computations on the more approximate lattice M the idea is only to use M to ensure convergence of fixed point computations and not needlessly reduce the precision of all other operations. We shall illustrate this in Subsection 4.5.3.

4.5.1 Inducing along the Abstraction Function

Now suppose that we have Galois connections $(L_i, \alpha_i, \gamma_i, M_i)$ such that each M_i is a more approximate version of L_i (for $i = 1, 2$). One way to make use of this is to replace an existing analysis $f_p : L_1 \rightarrow L_2$ with a new and more approximate analysis $g_p : M_1 \rightarrow M_2$. We already saw in Section 4.4 that

$$\alpha_2 \circ f_p \circ \gamma_1 \text{ is a candidate for } g_p$$

(just as $\gamma_2 \circ g_p \circ \alpha_1$ would be a candidate for f_p). The analysis $\alpha_2 \circ f_p \circ \gamma_1$ is said to be *induced* by f_p and the two Galois connections. This is illustrated by the diagram:

$$\begin{array}{ccc}
 L_1 & \xrightarrow{f_p} & L_2 \\
 \gamma_1 \uparrow & & \downarrow \alpha_2 \\
 M_1 & \xrightarrow{\alpha_2 \circ f_p \circ \gamma_1} & M_2
 \end{array}$$

Example 4.40 Let us return to Example 4.9 where we studied the simple program `plus` and specified the very precise analysis

$$f_{\text{plus}}(ZZ) = \{z_1 + z_2 \mid (z_1, z_2) \in ZZ\}$$

using the complete lattices $(\mathcal{P}(\mathbf{Z}), \subseteq)$ and $(\mathcal{P}(\mathbf{Z} \times \mathbf{Z}), \subseteq)$. In Example 4.21 we introduced the Galois connection

$$(\mathcal{P}(\mathbf{Z}), \alpha_{\text{sign}}, \gamma_{\text{sign}}, \mathcal{P}(\mathbf{Sign}))$$

for approximating sets of integers by sets of signs. In Example 4.35 we used the relational method to get the Galois connection

$$(\mathcal{P}(\mathbf{Z} \times \mathbf{Z}), \alpha_{SS'}, \gamma_{SS'}, \mathcal{P}(\mathbf{Sign} \times \mathbf{Sign}))$$

operating on pairs of integers. We now want to induce a more approximate analysis for the plus program

$$g_{\text{plus}} : \mathcal{P}(\mathbf{Sign} \times \mathbf{Sign}) \rightarrow \mathcal{P}(\mathbf{Sign})$$

from the existing analysis f_{plus} . To do so we take

$$g_{\text{plus}} = \alpha_{\text{sign}} \circ f_{\text{plus}} \circ \gamma_{SS'}$$

and simply calculate (for $SS \subseteq \mathbf{Sign} \times \mathbf{Sign}$)

$$\begin{aligned} g_{\text{plus}}(SS) &= \alpha_{\text{sign}}(f_{\text{plus}}(\gamma_{SS'}(SS))) \\ &= \alpha_{\text{sign}}(f_{\text{plus}}(\{(z_1, z_2) \in \mathbf{Z} \times \mathbf{Z} \mid (\text{sign}(z_1), \text{sign}(z_2)) \in SS\})) \\ &= \alpha_{\text{sign}}(\{z_1 + z_2 \mid z_1, z_2 \in \mathbf{Z}, (\text{sign}(z_1), \text{sign}(z_2)) \in SS\}) \\ &= \{\text{sign}(z_1 + z_2) \mid z_1, z_2 \in \mathbf{Z}, (\text{sign}(z_1), \text{sign}(z_2)) \in SS\} \\ &= \bigcup \{s_1 \oplus s_2 \mid (s_1, s_2) \in SS\} \end{aligned}$$

where $\oplus : \mathbf{Sign} \times \mathbf{Sign} \rightarrow \mathcal{P}(\mathbf{Sign})$ is the "addition" operator on signs (so e.g. $+\oplus+ = \{+\}$ and $+\oplus- = \{-, 0, +\}$). ■

The mundane approach to correctness. We shall now follow the approach of Section 4.1 and show that correctness of f_p carries over to g_p . For this assume that:

$$R_i : V_i \times L_i \rightarrow \{\text{true}, \text{false}\} \text{ is generated by } \beta_i : V_i \rightarrow L_i$$

The correctness of the analysis $f_p : L_1 \rightarrow L_2$ is then expressed by

$$(p \vdash \cdot \rightsquigarrow \cdot) (R_1 \twoheadrightarrow R_2) f_p$$

where $R_1 \twoheadrightarrow R_2$ is generated by $\beta_1 \twoheadrightarrow \beta_2$ (Lemma 4.8). As argued in Section 4.3 we get a correctness relation S_i for V_i and M_i by letting

$$S_i : V_i \times M_i \rightarrow \{\text{true}, \text{false}\} \text{ be generated by } \alpha_i \circ \beta_i : V_i \rightarrow M_i$$

which is equivalent to saying that $v_i S_i m_i \Leftrightarrow v_i R_i (\gamma_i(m_i))$. The correctness relation for the analysis using M_1 and M_2 will be $S_1 \twoheadrightarrow S_2$ which will be generated by $(\alpha_1 \circ \beta_1) \twoheadrightarrow (\alpha_2 \circ \beta_2)$ (Lemma 4.8). We now have the following useful result:

Lemma 4.41 If $(L_i, \alpha_i, \gamma_i, M_i)$ are Galois connections, and $\beta_i : V_i \rightarrow L_i$ are representation functions then

$$((\alpha_1 \circ \beta_1) \rightarrow (\alpha_2 \circ \beta_2))(\sim) = \alpha_2 \circ ((\beta_1 \rightarrow \beta_2)(\sim)) \circ \gamma_1$$

holds for all \sim . ■

Proof To see this we simply calculate

$$\begin{aligned} ((\alpha_1 \circ \beta_1) \rightarrow (\alpha_2 \circ \beta_2))(\sim)(m_1) &= \bigsqcup \{ \alpha_2(\beta_2(v_2)) \mid \alpha_1(\beta_1(v_1)) \sqsubseteq m_1 \wedge v_1 \sim v_2 \} \\ &= \alpha_2(\bigsqcup \{ \beta_2(v_2) \mid \beta_1(v_1) \sqsubseteq \gamma_1(m_1) \wedge v_1 \sim v_2 \}) \\ &= \alpha_2((\beta_1 \rightarrow \beta_2)(\sim)(\gamma_1(m_1))) \end{aligned}$$

and the result follows. ■


We shall now show that Lemma 4.41 yields:

$$(p \vdash \cdot \sim \cdot)(R_1 \rightarrow R_2) f_p \wedge \alpha_2 \circ f_p \circ \gamma_1 \sqsubseteq g_p \Rightarrow (p \vdash \cdot \sim \cdot)(S_1 \rightarrow S_2) g_p$$

This just means that if f_p is correct and if g_p is an upper approximation to the induced analysis then also g_p is correct. So suppose that

$$(p \vdash \cdot \sim \cdot)(R_1 \rightarrow R_2) f_p$$

and that $\alpha_2 \circ f_p \circ \gamma_1 \sqsubseteq g_p$. Since $(L_i, \alpha_i, \gamma_i, M_i)$ are Galois connections and f_p and g_p are monotone we get $f_p \sqsubseteq \gamma_2 \circ g_p \circ \alpha_1$ as illustrated in the diagrams:



It follows that $(\beta_1 \rightarrow \beta_2)(p \vdash \cdot \sim \cdot) \sqsubseteq \gamma_2 \circ g_p \circ \alpha_1$ and hence

$$\alpha_2 \circ (\beta_1 \rightarrow \beta_2)(p \vdash \cdot \sim \cdot) \circ \gamma_1 \sqsubseteq g_p$$

By Lemma 4.41 this is the desired result.

We shall say that a function $f_p : L_1 \rightarrow L_2$ is *optimal* for the program p if and only if correctness of a function $f' : L_1 \rightarrow L_2$ amounts to $f_p \sqsubseteq f'$. An equivalent formulation is that f_p is optimal if and only if

$$(\beta_1 \rightarrow \beta_2)(p \vdash \cdot \sim \cdot) = f_p$$

Lemma 4.41 may then be read as saying that if $f_p : L_1 \rightarrow L_2$ is optimal then so is $\alpha_2 \circ f_p \circ \gamma_1 : M_1 \rightarrow M_2$.

Fixed points in the induced analysis. Let us next consider the situation where the analysis $f_p : L_1 \rightarrow L_2$ requires the computation of the *least fixed point* of a monotone function $F : (L_1 \rightarrow L_2) \rightarrow (L_1 \rightarrow L_2)$ so that $f_p = \text{lfp}(F)$. The Galois connections $(L_i, \alpha_i, \gamma_i, M_i)$ give rise to a Galois connection $(L_1 \rightarrow L_2, \alpha, \gamma, M_1 \rightarrow M_2)$ between the monotone function spaces as shown in Section 4.4. We can now apply our technique of inducing and let $G : (M_1 \rightarrow M_2) \rightarrow (M_1 \rightarrow M_2)$ be an upper approximation to $\alpha \circ F \circ \gamma$. It will be natural to take $g_p : M_1 \rightarrow M_2$ to be $g_p = \text{lfp}(G)$. That correctness of f_p carries over to g_p follows from the following general result:

Lemma 4.42 Assume that (L, α, γ, M) is a Galois connection and let $f : L \rightarrow L$ and $g : M \rightarrow M$ be monotone functions satisfying that g is an upper approximation to the function induced by f , i.e.:

$$\alpha \circ f \circ \gamma \sqsubseteq g$$

Then for all $m \in M$:

$$g(m) \sqsubseteq m \Rightarrow f(\gamma(m)) \sqsubseteq \gamma(m)$$

and furthermore $\text{lfp}(f) \sqsubseteq \gamma(\text{lfp}(g))$ and $\alpha(\text{lfp}(f)) \sqsubseteq \text{lfp}(g)$. ■

Proof First assume $g(m) \sqsubseteq m$. The assumption $\alpha \circ f \circ \gamma \sqsubseteq g$ gives $\alpha(f(\gamma(m))) \sqsubseteq g(m)$ and hence $\alpha(f(\gamma(m))) \sqsubseteq m$. Using that a Galois connection is an adjunction (Proposition 4.20) we get $f(\gamma(m)) \sqsubseteq \gamma(m)$ as required.

To prove the second result we observe that $\{\gamma(m) \mid g(m) \sqsubseteq m\} \subseteq \{l \mid f(l) \sqsubseteq l\}$ follows from the previous result. Hence we get (using Lemma 4.22):

$$\gamma(\bigsqcap \{m \mid g(m) \sqsubseteq m\}) = \bigsqcap \{\gamma(m) \mid g(m) \sqsubseteq m\} \supseteq \bigsqcap \{l \mid f(l) \sqsubseteq l\}$$

Using Tarski's Theorem (Proposition A.10) twice we have $\text{lfp}(g) = \bigsqcap \text{Red}(g) = \bigsqcap \{m \mid g(m) \sqsubseteq m\}$ and $\text{lfp}(f) = \bigsqcap \text{Red}(f) = \bigsqcap \{l \mid f(l) \sqsubseteq l\}$ so it follows that $\gamma(\text{lfp}(g)) \supseteq \text{lfp}(f)$ as required. Then $\alpha(\text{lfp}(f)) \sqsubseteq \text{lfp}(g)$ follows because a Galois connection is an adjunction. ■

4.5.2 Application to Data Flow Analysis

Generalised Monotone Frameworks. To illustrate how these techniques can be applied we shall now consider a generalisation of the Monotone Frameworks of Section 2.3. So let a *generalised Monotone Framework* consist of:

- a complete lattice $L = (L, \sqsubseteq)$.

Here we do not demand that L satisfies the Ascending Chain Condition and we do not specify the space \mathcal{F} of transfer functions as we shall be taking

\mathcal{F} to be the entire space of monotone functions from L to L (which clearly contains the identity function and is closed under composition of functions).

An *instance* A of a generalised Monotone Framework then consists of:

- the complete lattice, L , of the framework;
- a finite flow, $F \subseteq \mathbf{Lab} \times \mathbf{Lab}$;
- a finite set of extremal labels, $E \subseteq \mathbf{Lab}$;
- an extremal value, $\iota \in L$; and
- a mapping, f , from the labels \mathbf{Lab} of F and E to monotone transfer functions from L to L .

As in Chapter 2 this gives rise to a set A^\exists of constraints

$$A_o(\ell) \supseteq \bigsqcup \{A_\bullet(\ell') \mid (\ell', \ell) \in F\} \sqcup \iota_E^\ell \quad \text{where } \iota_E^\ell = \begin{cases} \iota & \text{if } \ell \in E \\ \perp & \text{if } \ell \notin E \end{cases}$$

$$A_\bullet(\ell) \supseteq f_\ell(A_o(\ell))$$

where ℓ ranges over the labels \mathbf{Lab} of F and E . We write $(A_o, A_\bullet) \models A^\exists$ whenever $A_o, A_\bullet : \mathbf{Lab} \rightarrow L$ is a solution to the constraints A^\exists . It is useful to consider the associated monotone function

$$\tilde{f} : (\mathbf{Lab} \rightarrow L) \times (\mathbf{Lab} \rightarrow L) \rightarrow (\mathbf{Lab} \rightarrow L) \times (\mathbf{Lab} \rightarrow L)$$

defined by:

$$\tilde{f}(A_o, A_\bullet) = (\lambda \ell. \bigsqcup \{A_\bullet(\ell') \mid (\ell', \ell) \in F\} \sqcup \iota_E^\ell, \lambda \ell. f_\ell(A_o(\ell)))$$

We then have the following important result (in the manner of Section 4.4):

$$(A_o, A_\bullet) \supseteq \tilde{f}(A_o, A_\bullet) \text{ is equivalent to } (A_o, A_\bullet) \models A^\exists$$

Galois connections and Monotone Frameworks. Let now (L, α, γ, M) be a Galois connection and consider an instance B of the generalised Monotone Framework M that satisfies

- the mapping g from the labels \mathbf{Lab} of F and E to monotone transfer functions of $M \rightarrow M$ satisfies $g_\ell \supseteq \alpha \circ f_\ell \circ \gamma$ for all ℓ ; and
- the extremal value j satisfies $j \supseteq \alpha(\iota)$

and otherwise B is as A , i.e. has the same F and E .

As above we get a set of constraints B^\exists and we write $(B_o, B_\bullet) \models B^\exists$ whenever $B_o, B_\bullet : \mathbf{Lab} \rightarrow M$ is a solution to the constraints. The alternative

formulation is $(B_o, B_*) \sqsupseteq \bar{g}(B_o, B_*)$ where $\bar{g}: (\mathbf{Lab} \rightarrow M) \times (\mathbf{Lab} \rightarrow M) \rightarrow (\mathbf{Lab} \rightarrow M) \times (\mathbf{Lab} \rightarrow M)$ is the monotone function associated with the constraints.

We shall now see that whenever we have a solution to the constraints obtained from B then we also have a solution to the constraints obtained from A. This can be expressed by:

$$(B_o, B_*) \models B^\sqsupseteq \text{ implies } (\gamma \circ B_o, \gamma \circ B_*) \models A^\sqsupseteq$$

We can give a direct ("concrete") proof of this result but it is instructive to see how it follows from the general ("abstract") results established earlier. The idea is to "lift" the Galois connection (L, α, γ, M) to a Galois connection

$$((\mathbf{Lab} \rightarrow L) \times (\mathbf{Lab} \rightarrow L), \alpha', \gamma', (\mathbf{Lab} \rightarrow M) \times (\mathbf{Lab} \rightarrow M))$$

using the techniques for total function spaces and the independent attribute method presented in Section 4.4. The assumptions $g_\ell \sqsupseteq \alpha \circ f_\ell \circ \gamma$ (for all ℓ) and $j \sqsupseteq \alpha(\iota)$ can then be used to establish

$$\bar{g} \sqsupseteq \alpha' \circ \bar{f} \circ \gamma'$$

as the following calculations show

$$\begin{aligned} (\alpha' \circ \bar{f} \circ \gamma')(B_o, B_*) &= (\lambda \ell. \bigsqcup \{ \alpha(\gamma(B_o(\ell'))) \mid (\ell', \ell) \in F \} \sqcup \alpha(\iota_E), \\ &\quad \lambda \ell. \alpha(f_\ell(\gamma(B_o(\ell)))) \\ &\sqsubseteq \bar{g}(B_o, B_*) \end{aligned}$$

where we have used Lemma 4.22. We can now use Lemma 4.42 to obtain

$$\bar{g}(B_o, B_*) \sqsubseteq (B_o, B_*) \text{ implies } \bar{f}(\gamma'(B_o, B_*)) \sqsubseteq \gamma'(B_o, B_*)$$

and it follows that if $(B_o, B_*) \models B^\sqsupseteq$ then $(\gamma \circ B_o, \gamma \circ B_*) \models A^\sqsupseteq$ as stated above.

The mundane approach to correctness. The above result shows that any solution to the constraints obtained for B also is a solution to the constraints obtained from A. We shall now show that semantic correctness of A implies semantic correctness of B.

Let us reconsider the approach taken to semantic correctness in Section 4.1; here $F = \text{flow}(S_*)$ and $E = \{\text{init}(S_*)\}$. For the analysis A this calls for using a *representation function*

$$\beta: \text{State} \rightarrow L$$

and the correctness of all solutions to A^\sqsupseteq then amounts to the claim:

$$\begin{aligned} \text{Assume that } (A_o, A_*) \models A^\sqsupseteq \text{ and } \langle S_*, \sigma_1 \rangle \rightarrow^* \sigma_2; \\ \text{then } \beta(\sigma_1) \sqsubseteq \iota \text{ implies } \beta(\sigma_2) \sqsubseteq \bigsqcup \{ A_*(\ell) \mid \ell \in \text{final}(S_*) \}. \end{aligned} \quad (4.11)$$

For the analysis B it follows from Section 4.3 that it is natural to use the representation function

$$\alpha \circ \beta : \text{State} \rightarrow M$$

and the correctness of all solutions to B^{\sqsupseteq} then amounts to the claim:

$$\begin{aligned} & \text{Assume that } (B_o, B_*) \models B^{\sqsupseteq} \text{ and } \langle S_*, \sigma_1 \rangle \rightarrow^* \sigma_2; \\ & \text{then } (\alpha \circ \beta)(\sigma_1) \sqsubseteq j \text{ implies } (\alpha \circ \beta)(\sigma_2) \sqsubseteq \bigsqcup \{B_*(\ell) \mid \ell \in \text{final}(S_*)\}. \end{aligned} \quad (4.12)$$

We know that B is an upper approximation of the analysis induced from A and shall now prove that (4.11) implies (4.12). To do so we shall need to strengthen the relationship between the extremal values of A and B and we

assume that j satisfies $\gamma(j) = \iota$

from which $j \sqsupseteq \alpha(\iota)$ readily follows. For the proof that (4.11) implies (4.12) suppose that:

$$(B_o, B_*) \models B^{\sqsupseteq}, \langle S_*, \sigma_1 \rangle \rightarrow^* \sigma_2 \text{ and } (\alpha \circ \beta)(\sigma_1) \sqsubseteq j$$

It follows that:

$$(\gamma \circ B_o, \gamma \circ B_*) \models A^{\sqsupseteq}, \langle S_*, \sigma_1 \rangle \rightarrow^* \sigma_2 \text{ and } \beta(\sigma_1) \sqsubseteq \gamma(j) \sqsubseteq \iota$$

From (4.11) we get $\beta(\sigma_2) \sqsubseteq \bigsqcup \{\gamma \circ B_*(\ell) \mid \ell \in \text{final}(S_*)\}$ and hence $\beta(\sigma_2) \sqsubseteq \gamma(\bigsqcup \{B_*(\ell) \mid \ell \in \text{final}(S_*)\})$ showing the desired $(\alpha \circ \beta)(\sigma_2) \sqsubseteq \bigsqcup \{B_*(\ell) \mid \ell \in \text{final}(S_*)\}$.

A Worked Example

As a concrete example we shall now consider an analysis SS for the `WHILE` language that approximates how *sets* of states are transformed into *sets* of states. First we prove that it is correct. Then we show that the Constant Propagation Analysis of Section 2.3 is an upper approximation of an analysis induced from SS . This will then establish the correctness of the *Constant Propagation Analysis*.

Sets of states analysis. The analysis SS approximating the sets of states will be a generalised Monotone Framework with:

- the complete lattice $(\mathcal{P}(\text{State}), \sqsubseteq)$.

Given a label consistent statement S_* in `Stmt` we can now specify the instance as follows:

- the flow F is $\text{flow}(S_*)$;

- the set E of extremal labels is $\{init(S_*)\}$;
- the extremal value ι is **State**; and
- the transfer functions are given by f_t^{SS} :

$$\begin{aligned} f_t^{SS}(\Sigma) &= \{\sigma[x \mapsto \mathcal{A}[a]\sigma] \mid \sigma \in \Sigma\} && \text{if } [x := a]^\ell \text{ is in } S_* \\ f_t^{SS}(\Sigma) &= \Sigma && \text{if } [\text{skip}]^\ell \text{ is in } S_* \\ f_t^{SS}(\Sigma) &= \Sigma && \text{if } [b]^\ell \text{ is in } S_* \end{aligned}$$

where $\Sigma \subseteq \text{State}$.

Correctness. The following result shows that this analysis is correct in the sense explained above:

Lemma 4.43 Assume that $(SS_0, SS_*) \models SS^\supseteq$ and $\langle S_*, \sigma_1 \rangle \rightarrow^* \sigma_2$; then $\sigma_1 \in \text{State}$ implies $\sigma_2 \in \bigcup \{SS_*(\ell) \mid \ell \in \text{final}(S_*)\}$. ■

Proof From Section 2.2 we have:

$$\langle S, \sigma \rangle \rightarrow \langle S', \sigma' \rangle \text{ implies } \text{final}(S) \supseteq \text{final}(S') \wedge \text{flow}(S) \supseteq \text{flow}(S')$$

and as in Chapter 2 it is immediate that

$$\text{flow}(S) \supseteq \text{flow}(S') \wedge (SS_0, SS_*) \models SS^\supseteq(S) \text{ implies } (SS_0, SS_*) \models SS^\supseteq(S')$$

It then suffices to show

$$\begin{aligned} (SS_0, SS_*) \models SS^\supseteq(S) \wedge \langle S, \sigma \rangle \rightarrow \sigma' \wedge \sigma \in SS_0(\text{init}(S)) \\ \text{implies } \sigma' \in \bigcup \{SS_*(\ell) \mid \ell \in \text{final}(S)\} \\ (SS_0, SS_*) \models SS^\supseteq(S) \wedge \langle S, \sigma \rangle \rightarrow \langle S', \sigma' \rangle \wedge \sigma \in SS_0(\text{init}(S)) \\ \text{implies } \sigma' \in SS_0(\text{init}(S')) \end{aligned}$$

since then an induction on the length of the derivation sequence $\langle S_*, \sigma_1 \rangle \rightarrow^* \sigma_2$ will give the result. The proof proceeds by induction on the inference in the semantics. We only consider a few of the interesting cases.

The case $\langle [x := a]^\ell, \sigma \rangle \rightarrow \sigma[x \mapsto \mathcal{A}[a]\sigma]$. Then $SS^\supseteq(S)$ will contain the equation

$$SS_*(\ell) \supseteq \{\sigma[x \mapsto \mathcal{A}[a]\sigma] \mid \sigma \in SS_0(\ell)\}$$

and since $\text{init}([x := a]^\ell) = \ell$ and $\text{final}([x := a]^\ell) = \{\ell\}$ we see that that the required relationship holds: if $\sigma \in SS_0(\ell)$ then $\sigma[x \mapsto \mathcal{A}[a]\sigma] \in SS_*(\ell)$.

The case $\langle S_1; S_2, \sigma \rangle \rightarrow \langle S'_1; S_2, \sigma' \rangle$ because $\langle S_1, \sigma \rangle \rightarrow \langle S'_1, \sigma' \rangle$. From the assumption $\sigma \in SS_0(\text{init}(S_1; S_2))$ we get $\sigma \in SS_0(\text{init}(S_1))$ and the induction hypothesis gives $\sigma' \in SS_0(\text{init}(S'_1))$. But then $\sigma' \in SS_0(\text{init}(S'_1; S_2))$ as required.

The case $\langle S_1; S_2, \sigma \rangle \rightarrow \langle S_2, \sigma' \rangle$ because $\langle S_1, \sigma \rangle \rightarrow \sigma'$. From the assumption $\sigma \in SS_0(\text{init}(S_1; S_2))$ we get $\sigma \in SS_0(\text{init}(S_1))$ and the induction hypothesis gives $\sigma' \in \bigcup \{SS_*(\ell) \mid \ell \in \text{final}(S_1)\}$. We have

$$\{(\ell, \text{init}(S_2)) \mid \ell \in \text{final}(S_1)\} \subseteq \text{flow}(S_1; S_2)$$

and since we have the constraints

$$SS_0(\ell) \supseteq \bigcup \{SS_0(\ell') \mid (\ell', \ell) \in \text{flow}(S_1; S_2)\}$$

for all ℓ we get

$$SS_0(\text{init}(S_2)) \supseteq \bigcup \{SS_0(\ell) \mid \ell \in \text{final}(S_1)\}$$

and hence $\sigma' \in SS_0(\text{init}(S_2))$ as required.

The remaining cases follow the same pattern and are omitted. ■

Remark. The SS analysis is unlikely to be optimal and hence is unlikely to equal the collecting semantics (see Exercise 4.5). This may be demonstrated by exhibiting an example where $\bigcup \{SS_0(\ell) \mid \ell \in \text{final}(S_*)\}$ is strictly larger than $\{\sigma' \mid \langle S_*, \sigma \rangle \rightarrow^* \sigma' \wedge \sigma \in \text{State}\}$ and it is fairly easy to do so. To obtain a specification of the collecting semantics we should let transfer functions be associated with edges rather than nodes as this would allow us to record the outcome of tests (see Exercise 2.11). ■

Constant Propagation Analysis. The analysis of Section 2.3 is specified by a generalised Monotone Framework consisting of

- the complete lattice $\widehat{\text{State}}_{\text{CP}} = ((\text{Var} \rightarrow \mathbf{Z}^T)_{\perp}, \sqsubseteq)$.

The instance for the statement S_* is determined by

- the flow F is $\text{flow}(S_*)$;
- the set E of extremal labels is $\{\text{init}(S_*)\}$;
- the extremal value ι is $\lambda x. \top$; and
- the transfer functions are given by the mapping f^{CP} defined in Table 2.7.

Galois connection. The relationship between the two analyses is established by defining the representation function

$$\beta_{\text{CP}} : \text{State} \rightarrow \widehat{\text{State}}_{\text{CP}}$$

by $\beta_{\text{CP}}(\sigma) = \sigma$ (as in Example 4.7). As in Section 4.3 this gives rise to a Galois connection $(\mathcal{P}(\text{State}), \alpha_{\text{CP}}, \gamma_{\text{CP}}, \widehat{\text{State}}_{\text{CP}})$ where $\alpha_{\text{CP}}(\Sigma) = \bigsqcup \{\beta_{\text{CP}}(\sigma) \mid \sigma \in \Sigma\}$ and $\gamma_{\text{CP}}(\hat{\sigma}) = \{\sigma \mid \beta_{\text{CP}}(\sigma) \sqsubseteq \hat{\sigma}\}$. One can now show that for all labels ℓ

$$f_{\ell}^{\text{CP}} \supseteq \alpha_{\text{CP}} \circ f_{\ell}^{\text{SS}} \circ \gamma_{\text{CP}}$$

as well as $\gamma_{\text{CP}}(\lambda x. \top) = \text{State}$. Let us only consider the case where $[x := a]^{\ell}$ occurs in S_* and calculate

$$\alpha_{\text{CP}}(f_{\ell}^{\text{SS}}(\gamma_{\text{CP}}(\hat{\sigma}))) = \alpha_{\text{CP}}(f_{\ell}^{\text{SS}}(\{\sigma \mid \sigma \sqsubseteq \hat{\sigma}\}))$$

$$\begin{aligned}
&= \alpha_{\text{CP}}(\{\sigma[x \mapsto \mathcal{A}[a]\sigma] \mid \sigma \sqsubseteq \hat{\sigma}\}) \\
&= \bigsqcup \{\sigma[x \mapsto \mathcal{A}[a]\sigma] \mid \sigma \sqsubseteq \hat{\sigma}\} \\
&\sqsubseteq \hat{\sigma}[x \mapsto \bigsqcup \{\mathcal{A}[a]\sigma \mid \sigma \sqsubseteq \hat{\sigma}\}] \\
&\sqsubseteq f_t^{\text{CP}}(\hat{\sigma})
\end{aligned}$$

and where the last step follows from $\bigsqcup \{\mathcal{A}[a]\sigma \mid \sigma \sqsubseteq \hat{\sigma}\} \sqsubseteq \mathcal{A}_{\text{CP}}[a]\hat{\sigma}$ which can be proved by a straightforward structural induction on a .

Thus we conclude that CP is an upper approximation to the analysis induced from SS by the Galois connection and hence it is correct.

4.5.3 Inducing along the Concretisation Function

Widening operator induced by Galois connection. Suppose that we have a Galois connection (L, α, γ, M) between the complete lattices L and M , and also a monotone function $f : L \rightarrow L$. Often the motivation for approximating f arises because a fixed point of f is desired, and the ascending chain $(f^n(\perp))_n$ does not eventually stabilise (or may do so in too many iterations). Instead of using $\alpha \circ f \circ \gamma : M \rightarrow M$ to remedy this situation it is often possible to consider a *widening operator* $\nabla_M : M \times M \rightarrow M$ and use it to define $\nabla_L : L \times L \rightarrow L$ by the formula:

$$l_1 \nabla_L l_2 = \gamma(\alpha(l_1) \nabla_M \alpha(l_2))$$

If ∇_L turns out to be a widening operator we then know how to approximate the least fixed point of $f : L \rightarrow L$ while calculating over L . This has the advantage that the coarser structure of M is only used to ensure convergence and does not needlessly reduce the precision of all other operations. The following result gives sufficient criteria for this approach to work:

Proposition 4.44

Let (L, α, γ, M) be a Galois connection and let $\nabla_M : M \times M \rightarrow M$ be an upper bound operator. Then the formula

$$l_1 \nabla_L l_2 = \gamma(\alpha(l_1) \nabla_M \alpha(l_2))$$

defines an upper bound operator $\nabla_L : L \times L \rightarrow L$. It defines a widening operator if one of the following two conditions are fulfilled:

- (i) M satisfies the Ascending Chain Condition, or
- (ii) (L, α, γ, M) is a Galois insertion and $\nabla_M : M \times M \rightarrow M$ is a widening operator.

Proof First we prove that ∇_L is an upper bound operator. Since ∇_M is an upper bound operator we have $\alpha(l_i) \sqsubseteq \alpha(l_1) \nabla_M \alpha(l_2)$. Using that a Galois connection is an adjunction we get $l_i \sqsubseteq \gamma(\alpha(l_1) \nabla_M \alpha(l_2))$, i.e. $l_i \sqsubseteq l_1 \nabla_L l_2$.

Assume now that condition (i) is fulfilled and consider an ascending chain $(l_n)_n$ in L . We know that also $(l_n^{\nabla_L})_n$ is an ascending chain and that $l_n^{\nabla_L} \in \gamma[M]$ for $n > 0$. Hence $(\alpha(l_n^{\nabla_L}))_n$ is an ascending chain and since M satisfies the Ascending Chain Condition there exists $n_0 \geq 1$ such that $\alpha(l_n^{\nabla_L}) = \alpha(l_{n_0}^{\nabla_L})$ for all $n \geq n_0$. So $\gamma(\alpha(l_n^{\nabla_L})) = \gamma(\alpha(l_{n_0}^{\nabla_L}))$ for all $n \geq n_0$ and using that $\gamma \circ \alpha \circ \gamma = \gamma$ (Fact 4.24) we get $l_n^{\nabla_L} = l_{n_0}^{\nabla_L}$ for all $n \geq n_0$. This completes the proof.

Assume next that condition (ii) is fulfilled and consider again an ascending chain $(l_n)_n$ in L . Since α is monotone it follows that $(\alpha(l_n))_n$ is an ascending chain in M . Now, ∇_M is a widening operator on M so there exists n_0 such that $(\alpha(l_n))^{\nabla_M} = (\alpha(l_{n_0}))^{\nabla_M}$ for $n \geq n_0$. We shall now prove that

$$(\alpha(l_n))^{\nabla_M} = \alpha(l_n^{\nabla_L}) \quad (4.13)$$

for all $n \geq 0$. The case $n = 0$ is immediate since $(\alpha(l_0))^{\nabla_M} = \alpha(l_0) = \alpha(l_0^{\nabla_L})$. For the induction step we assume that $(\alpha(l_n))^{\nabla_M} = \alpha(l_n^{\nabla_L})$. Then

$$\begin{aligned} (\alpha(l_{n+1}))^{\nabla_M} &= (\alpha(l_n))^{\nabla_M} \nabla_M \alpha(l_{n+1}) \\ &= \alpha(l_n^{\nabla_L}) \nabla_M \alpha(l_{n+1}) \end{aligned}$$

and

$$\begin{aligned} \alpha(l_{n+1}^{\nabla_L}) &= \alpha(l_n^{\nabla_L} \nabla_L l_{n+1}) \\ &= \alpha(\gamma(\alpha(l_n^{\nabla_L}) \nabla_M \alpha(l_{n+1}))) \\ &= \alpha(l_n^{\nabla_L}) \nabla_M \alpha(l_{n+1}) \end{aligned}$$

since (L, α, γ, M) is a Galois insertion.

Using (4.13) we thus have that there exists n_0 such that $\alpha(l_n^{\nabla_L}) = \alpha(l_{n_0}^{\nabla_L})$ for all $n \geq n_0$. But then $\gamma(\alpha(l_n^{\nabla_L})) = \gamma(\alpha(l_{n_0}^{\nabla_L}))$ and hence $l_n^{\nabla_L} = l_{n_0}^{\nabla_L}$ for all $n \geq n_0$ because (L, α, γ, M) is a Galois insertion. This completes the proof. ■

Precision of induced widening operator. The following result compares the precision of using the widening operator ∇_L with the precision of using the widening operator ∇_M .

Lemma 4.45 If (L, α, γ, M) is a Galois insertion such that $\gamma(\perp_M) = \perp_L$, and if $\nabla_M : M \times M \rightarrow M$ is a widening operator, then the widening operator $\nabla_L : L \times L \rightarrow L$ defined by $l_1 \nabla_L l_2 = \gamma(\alpha(l_1) \nabla_M \alpha(l_2))$ satisfies

$$\text{Ifp}_{\nabla_L}(f) = \gamma(\text{Ifp}_{\nabla_M}(\alpha \circ f \circ \gamma))$$

for all monotone functions $f : L \rightarrow L$. ■

Proof By Proposition 4.44 we already know that ∇_L is a widening operator. Hence there exists $n_f \geq 0$ such that $\text{Ifp}_{\nabla_L}(f) = f_{\nabla_L}^{n_f} = f_{\nabla_L}^n$ for all $n \geq n_f$. Next write

$g = \alpha \circ f \circ \gamma$ and recall that ∇_M is a widening operator. Hence there exists $n_g \geq 0$ such that $\text{fp}_{\nabla_M}(g) = g_{\nabla_M}^{n_g} = g_{\nabla_M}^n$ for all $n \geq n_g$. To obtain the desired result it therefore suffices to prove

$$f_{\nabla_L}^n = \gamma(g_{\nabla_M}^n) \quad (4.14)$$

by induction on n . The base case ($n = 0$) is immediate since $f_{\nabla_L}^0 = \perp_L$ and $g_{\nabla_M}^0 = \perp_M$ and we assumed that $\perp_L = \gamma(\perp_M)$.

To prepare for the induction step we prove that (4.14) implies that:

$$f(f_{\nabla_L}^n) \sqsubseteq f_{\nabla_L}^n \Leftrightarrow g(g_{\nabla_M}^n) \sqsubseteq g_{\nabla_M}^n \quad (4.15)$$

For " \Rightarrow " we calculate (using (4.14) and that (L, α, γ, M) is a Galois connection):

$$\begin{aligned} f(f_{\nabla_L}^n) \sqsubseteq f_{\nabla_L}^n &\Rightarrow \alpha(f(f_{\nabla_L}^n)) \sqsubseteq \alpha(f_{\nabla_L}^n) \\ &\Rightarrow \alpha(f(\gamma(g_{\nabla_M}^n))) \sqsubseteq \alpha(\gamma(g_{\nabla_M}^n)) \\ &\Rightarrow g(g_{\nabla_M}^n) \sqsubseteq \alpha(\gamma(g_{\nabla_M}^n)) \\ &\Rightarrow g(g_{\nabla_L}^n) \sqsubseteq g_{\nabla_M}^n \end{aligned}$$

For " \Leftarrow " we calculate (using (4.14) and that (L, α, γ, M) is a Galois connection):

$$\begin{aligned} g(g_{\nabla_M}^n) \sqsubseteq g_{\nabla_M}^n &\Rightarrow \gamma(g(g_{\nabla_M}^n)) \sqsubseteq \gamma(g_{\nabla_M}^n) \\ &\Rightarrow \gamma(\alpha(f(\gamma(g_{\nabla_M}^n)))) \sqsubseteq \gamma(g_{\nabla_M}^n) \\ &\Rightarrow \gamma(\alpha(f(f_{\nabla_L}^n))) \sqsubseteq f_{\nabla_L}^n \\ &\Rightarrow f(f_{\nabla_L}^n) \sqsubseteq f_{\nabla_L}^n \end{aligned}$$

Returning to the induction step ($n > 0$) of the proof of (4.14) we calculate:

$$\begin{aligned} f_{\nabla_L}^n &= \begin{cases} f_{\nabla_L}^{n-1} & \text{if } f(f_{\nabla_L}^{n-1}) \sqsubseteq f_{\nabla_L}^{n-1} \\ f_{\nabla_L}^{n-1} \nabla_L f(f_{\nabla_L}^{n-1}) & \text{otherwise} \end{cases} \\ &= \begin{cases} f_{\nabla_L}^{n-1} & \text{if } g(g_{\nabla_M}^{n-1}) \sqsubseteq g_{\nabla_M}^{n-1} \\ f_{\nabla_L}^{n-1} \nabla_L f(f_{\nabla_L}^{n-1}) & \text{otherwise} \end{cases} \\ &= \begin{cases} \gamma(g_{\nabla_M}^{n-1}) & \text{if } g(g_{\nabla_M}^{n-1}) \sqsubseteq g_{\nabla_M}^{n-1} \\ \gamma(\alpha(\gamma(g_{\nabla_M}^{n-1}))) \nabla_M \alpha(f(\gamma(g_{\nabla_M}^{n-1}))) & \text{otherwise} \end{cases} \\ &= \gamma \left(\begin{cases} g_{\nabla_M}^{n-1} & \text{if } g(g_{\nabla_M}^{n-1}) \sqsubseteq g_{\nabla_M}^{n-1} \\ \alpha(\gamma(g_{\nabla_M}^{n-1})) \nabla_M g(g_{\nabla_M}^{n-1}) & \text{otherwise} \end{cases} \right) \\ &= \gamma \left(\begin{cases} g_{\nabla_M}^{n-1} & \text{if } g(g_{\nabla_M}^{n-1}) \sqsubseteq g_{\nabla_M}^{n-1} \\ g_{\nabla_M}^{n-1} \nabla_M g(g_{\nabla_M}^{n-1}) & \text{otherwise} \end{cases} \right) \\ &= \gamma(g_{\nabla_M}^n) \end{aligned}$$

In this calculation we have used that (4.14) and (4.15) hold for $n - 1$, the definition of ∇_L , and that (L, α, γ, M) is a Galois insertion. ■

This result then provides the formal justification for the motivating remarks in the beginning of the subsection. To be specific let M be of finite height,

let (L, α, γ, M) be a Galois insertion satisfying $\gamma(\perp_M) = \perp_L$, and let ∇_M be the least upper bound operator \sqcup_M . Then the above lemma shows that

$$\text{lfp}_{\nabla_L}(f) = \gamma(\text{lfp}(\alpha \circ f \circ \gamma))$$

which means that $\text{lfp}_{\nabla_L}(f)$ equals the result we would have obtained if we decided to work with $\alpha \circ f \circ \gamma : M \rightarrow M$ instead of the given $f : L \rightarrow L$; furthermore the number of iterations needed turn out to be the same. Since the greater precision of L over M is available for all other operations, this suggests that the use of widening operators is often preferable to the approach of Subsection 4.5.1.

Concluding Remarks

In this chapter we have only been able to touch upon a few of the central concepts of Abstract Interpretation: this has mainly been based on [27, 29, 25] and [30] while the series of examples leading up to Example 4.39 are based on [140]. Much more can be said both about the development of the theory and about its applications. In this section we briefly discuss some of the more important concepts that have been omitted so far.

Upper closure operators. An *upper closure operator* $\rho : L \rightarrow L$ is a monotone function that is extensive (i.e. satisfies $\rho \sqsupseteq \lambda.l$) and idempotent (i.e. $\rho \circ \rho = \rho$). Such operators arise naturally in Abstract Interpretation [29] because whenever (L, α, γ, M) is a Galois connection the function $\gamma \circ \alpha : L \rightarrow L$ is easily seen to be an upper closure operator. Furthermore, if $\rho : L \rightarrow L$ is an upper closure operator then the image $\rho[L] = \{\rho(l) \mid l \in L\}$ of L under ρ equals the set $\text{Fix}(\rho) = \{l \in L \mid l = \rho(l)\}$ of fixed points of ρ and is a complete lattice under the partial ordering of L ; in fact $(L, \rho, \lambda.l, \rho[L])$ is a Galois insertion.

It follows that upper closure operators may be used to represent Galois connections by simply demanding that the more approximate space M is actually a subset of L and that no essential features are lost by doing this. This then opens up for directly comparing the precision of various Galois connections over the same complete lattice L by simply relating the closure operators. The relation $\rho_1 \sqsubseteq \rho_2$ is naturally defined by $\forall l \in L : \rho_1(l) \sqsubseteq \rho_2(l)$ but turns out to be equivalent to the condition that $\rho_2(L) \subseteq \rho_1(L)$ and represents the fact that ρ_2 is more approximate than ρ_1 .

Having defined an ordering on the set of upper closure operators on L one can next show that it is itself a complete lattice: the least element is given by the upper closure operator $\lambda.l$ and the greatest element by $\lambda.\top$. The binary greatest lower bound operator \sqcap is of special interest: it gives rise to the construction of the reduced product [29] as we saw it in Section 4.3.

A number of additional constructs can be explained by means of upper closure operators: we just mention reduced cardinal power [29] and disjunctive completion [29]. By "inverting" some of these constructions it may then be possible to find the "optimal bases" with respect to a given combination: for reduced product the notion of pseudo-complementation has been used to find the minimal factors [23, 24], and for disjunctive completion one can identify the basic irreducible properties that are inherently necessary for the analysis [47, 48].

Stronger properties on the complete lattices. Being a complete lattice is a rather weak notion compared to being a powerset. By considering more structure on the complete lattices, say distributivity, and identifying the elements that correspond to singletons, e.g. atoms or join irreducible elements, it is frequently possible to lift some of the stronger results that hold for powersets to a larger class of complete lattices.

Since powersets are isomorphic to bit vectors this gives a way of finding more general conditions on analyses for when they are as efficient as the *Bit Vector Frameworks*. This is of particular interest in the case of fixed points, where one can use properties of distributive lattices and distributive analyses to give rather low bounds on the number of iterations needed for the ascending chain $(f^n(\perp))_n$ to stabilise [99, 90].

Another line of work concerns the development of the *tensor product* for complete lattices that are not also powersets [83, 85, 87]. Several notions of tensor product have been studied in lattice theory but the development of tensor products suitable for program analysis was first done in [83].

Concrete analyses. In this chapter we have concentrated on introducing some of the key notions in the theory of Abstract Interpretation and only occasionally have we hinted at concrete applications.

One of the main applications of Abstract Interpretation has been in the area of logic programming. To implement a program efficiently it is important to have precise information about the substitutions that may reach the various program points; a central question to be asked for a substitution is whether or not it is ground. A number of analyses have been designed for this and most of these build upon the framework of Abstract Interpretation. This includes the design of iteration strategies based on widening, and the decomposition of base domains using the techniques mentioned above under upper closure operators.

Another main application of Abstract Interpretation has been to approximate subsets of n -dimensional vector spaces over integers or rationals. For the purpose of this discussion we shall limit ourselves to at most two dimensions (the line and the plane). In the case of one dimension there are two main techniques. One we already illustrated: the lattice of intervals, and it may be generalised to consider (possibly finite) unions of intervals.

The other technique records sets of numbers modulo some base value, e.g. $\{x \mid x \bmod k_1 = k_2\}$. Clearly these two analyses can be combined. In the case of two (or more) dimensions it is straightforward to perform an independent attribute analysis where the techniques above are applied component-wise for each dimension.

A large effort has been conducted to develop more interesting relational analyses for two (or more) dimensions where the choice of axes is of less importance for the ability to approximate subsets of vector space. An early method was the affine subspaces of Karr [71] where sets of the form $\{(x, y) \mid k_1x + k_2y = k_3\}$ can be described. The generalisation from equality to inequality, and allowing to take intersections of such subsets, was considered by Cousot and Halb wachs [31] and resulted in a study of convex polygons. Generalisations and combinations of these ideas have been developed by Granger [50, 51] and by Masdupuy [79].

An interesting line of work pioneered by Deutsch [34, 35] is to change the problem of describing regular sets of words over a finite alphabet to the problem of describing sets of integer vectors. This is by no means trivial but once it has been achieved it opens up for using all of the above techniques to represent also regular sets of words. This is very important for the analysis of higher-order and concurrent programs, as shown by Deutsch and Colby [21, 22], since it can describe the shape of activation records and communication patterns in much greater precision than other comparative techniques.

We should also mention techniques for building the abstract space of properties "dynamically" [16] and for using widening and narrowing to improve the performance of chaotic iteration [17].

Duality. The dual \sqsubseteq^d of a partial ordering \sqsubseteq is obtained by defining $l_1 \sqsubseteq^d l_2$ if and only if $l_2 \sqsubseteq l_1$; thus we could write \sqsubseteq^d as \sqsupseteq . Any concept defined in terms of partial orderings can be dualised by replacing all partial orderings by their dual. In this way the dual least element is the greatest element, and the dual least upper bound is the greatest lower bound etc. The principle of lattice duality of Lattice Theory says that if any statement about partially ordered sets is true then so is the dual statement. (Interestingly the concept of monotonicity is its own dual.) However, we should like to point out that the dual of a complete lattice may of course differ from the complete lattice itself; pictorially we represent this by drawing the complete lattice "up-side down".

The principle of lattice duality is important for program analysis because it gives an easy way of relating the literature on Abstract Interpretation to the "classical" literature on Data Flow Analysis: simply dualise the complete lattices. So in Abstract Interpretation the greatest element is trivially safe and conveys no information whereas in "classical" Data Flow Analysis it is the least element that has this role. Similarly, in Abstract Interpretation we

are interested in least fixed points whereas in “classical” Data Flow Analysis we are interested in greatest fixed points.

Staying within the basic approach of Abstract Interpretation, that going up in the complete lattice means losing information, it is still possible to dualise much of the development: in particular we can define the notion of dual Galois connections. To see why this may be worthwhile consider the following scenario. In program analysis we aim at establishing an element $l_\ell \in L$ for describing the set of values that may reach a given program point ℓ . In program transformation it is frequently the case that a certain transformation Ξ is valid only if the set of values that reach a certain point have certain properties; we may formulate this as the condition $l_\ell \sqsubseteq l_\Xi$. Now if we want to be more approximate we approximate l_ℓ to l'_ℓ and l_Ξ to l'_Ξ and formulate the approximate condition $l'_\ell \sqsubseteq l'_\Xi$. To ensure that $l'_\ell \sqsubseteq l'_\Xi$ implies $l_\ell \sqsubseteq l_\Xi$ we demand that $l_\ell \sqsubseteq l'_\ell$ and that $l'_\Xi \sqsubseteq l_\Xi$. Thus properties of program points are approximated by going *up* in the complete lattice, for which Galois connections are useful, whereas enabling conditions for program transformations are approximated by going *down* in the complete lattice, and for this the concept of dual Galois connections is useful.

A final word of advice concerns the interplay between Abstract Interpretation and *Denotational Semantics*. In Denotational Semantics the least element conveys absolutely no information, and we learn more when things get larger according to the partial order; had there been a greatest element it would have denoted conflicting information. This is quite opposite to the situation in Abstract Interpretation where the greatest element conveys absolutely no information and we learn more when things get smaller according to the partial order; the least element often denotes non-reachability. Hence it would be dangerous to simply apply too many of the intuitions from Denotational Semantics when performing Abstract Interpretation not least because both formalisms ask for least fixed points and therefore are *not duals* of one another.

Mini Projects

Mini Project 4.1 A Galois Connection for Lists

In a series of examples leading up to Example 4.39 we constructed a Galois insertion for recording the relationship between pairs of integers; it was given by

$$(\mathcal{P}(\mathbf{Z} \times \mathbf{Z}), \alpha_{\text{SSR}'}, \gamma_{\text{SSR}'}, \mathcal{P}(\mathbf{AB}))$$

where $\mathbf{AB} \subseteq \mathbf{Sign} \times \mathbf{Sign} \times \mathbf{Range}$ contained only 29 elements (out of the 45 possibilities).

In this mini project we are going to construct a Galois insertion for recording the relationship between pairs of lists. Let V be the domain of lists of finite length over some simple data type. We write $x = [x_1, \dots, x_n]$ for a list with n elements whose first element is x_1 ; when $n = 0$ we write $x = []$. Next let $x = [x_1, \dots, x_n]$ and $y = [y_1, \dots, y_m]$ be two lists. They have the same head if and only if $n > 0$, $m > 0$ and $x_1 = y_1$. The list x is a suffix of y if and only if there exists $k \geq 0$ such that $n + k = m$ and $x_i = y_{i+k}$ for $i \in \{1, \dots, n\}$. Finally we write $\text{length}(x) = n$ and $\text{length}(y) = m$.

The Galois insertion should have the form

$$(\mathcal{P}(V \times V), \alpha, \gamma, \mathcal{P}(\mathbf{LR}))$$

where $\mathbf{LR} \subseteq \mathcal{P}(\{\mathbf{H}, \mathbf{S}\}) \times \mathbf{Range}$. Here \mathbf{H} means that the lists have the same head, \mathbf{S} means that the *shorter* list is a suffix of the other, and the range components describe $\text{length}(x) - \text{length}(y)$ where x is the first list and y the second list.

Complete the details of the specification. ■

Mini Project 4.2 Correctness of the Shape Analysis

We shall now return to the *Shape Analysis* of Section 2.6 and show how it gives rise to a Galois connection. Recall that the semantics uses configurations with a state $\sigma \in \mathbf{State}$ and a heap component $\mathcal{H} \in \mathbf{Heap}$ and that the analysis works on shape graphs consisting of an abstract state \mathbf{S} , an abstract heap \mathbf{H} and a sharing component \mathbf{is} .

We shall begin by defining a function vars that given a location and a state will determine the associated abstract location:

$$\text{vars}(\xi)(\sigma) = n_X \text{ where } X = \{x \mid \sigma(x) = \xi\}$$

Proceed as follows:

1. Define a representation function

$$\beta_{\mathbf{SA}} : \mathbf{State} \times \mathbf{Heap} \rightarrow \mathcal{P}(\mathbf{SG})$$

that to each state and heap associates a singleton set with a compatible shape graph (as defined in Section 2.6) and construct the associated Galois connection

$$(\mathcal{P}(\mathbf{State} \times \mathbf{Heap}), \alpha_{\mathbf{SA}}, \gamma_{\mathbf{SA}}, \mathcal{P}(\mathbf{SG}))$$

Is it a Galois insertion?

To establish the correctness of the analysis we shall follow the approach of Section 4.5.2:

2. Specify an analysis SH approximating the sets of pairs of states and heaps as a generalised Monotone Framework over the complete lattice $(\mathcal{P}(\mathbf{State} \times \mathbf{Heap}), \subseteq)$; write f_l^{SH} for the associated transfer functions. Prove the correctness of the analysis SH (i.e. prove an analogue of Lemma 4.43).
3. Show that $f_l^{\text{SA}} \sqsupseteq \alpha_{\text{SA}} \circ f_l^{\text{SH}} \circ \gamma_{\text{SA}}$ for all transfer functions and conclude that the Shape Analysis is correct. Determine whether or not $f_l^{\text{SA}} = \alpha_{\text{SA}} \circ f_l^{\text{SH}} \circ \gamma_{\text{SA}}$ holds for all transfer functions. (See Exercise 2.22.) ■

Mini Project 4.3 Application to Control Flow Analysis

In this mini project we shall perform an analogue of the development of Subsection 4.5.2 for the Control and Data Flow Analysis of Section 3.5.

1. Specify a “sets of values” analysis

$$(\widehat{\mathbf{C}}_{\text{sv}}, \widehat{\rho}_{\text{sv}}) \models_{\text{sv}} e$$

in the manner of Section 3.5.1 (by taking $\mathbf{Data} = \mathbf{Val}$ where \mathbf{Val} is as in Section 3.2). Formulate and prove a semantic correctness result in the manner of Example 4.40 and Theorem 3.10.

2. Let a *monotone structure* (L, \mathcal{F}) be given as in Section 3.5.2 and consider a Galois connection $(\mathcal{P}(\mathbf{Val}), \alpha, \gamma, L)$. Motivated by the judgements of the acceptability relation $(\widehat{\mathbf{C}}, \widehat{\mathbf{D}}, \widehat{\rho}, \widehat{d}) \models_D e$ construct a Galois connection:

$$(\{(\widehat{\mathbf{C}}_{\text{sv}}, \widehat{\rho}_{\text{sv}}) \mid \dots\}, \alpha', \gamma', \{(\widehat{\mathbf{C}}, \widehat{\mathbf{D}}, \widehat{\rho}, \widehat{d}) \mid \dots\})$$

Formulate and prove a result intended to establish

$$(\widehat{\mathbf{C}}, \widehat{\mathbf{D}}, \widehat{\rho}, \widehat{d}) \models_D e \Rightarrow \gamma'((\widehat{\mathbf{C}}, \widehat{\mathbf{D}}, \widehat{\rho}, \widehat{d})) \models_{\text{sv}} e$$

and argue that this shows the semantic correctness of the Control and Data Flow Analysis. ■

Exercises

Exercise 4.1 For the complete lattice $(\mathbf{Sign}', \sqsubseteq)$ of Figure 4.10 define a correctness relation $R_{\mathbf{ZS}'} : \mathbf{Z} \times \mathbf{Sign}' \rightarrow \{\text{true}, \text{false}\}$. Verify that it indeed has the properties (4.4) and (4.5). Next define a representation function $\beta_{\mathbf{ZS}'} : \mathbf{Z} \rightarrow \mathbf{Sign}'$ and show that the $R_{\mathbf{ZS}'}$ constructed above is indeed generated by $\beta_{\mathbf{ZS}'}$. ■




Figure 4.10: The complete lattice $(\text{Sign}', \sqsubseteq)$.

Exercise 4.2 Show that if (4.4) and (4.5) hold for R and L then we also have:

$$v R l_1 \wedge v R l_2 \Rightarrow v R (l_1 \sqcup l_2)$$

and more generally:

$$(\forall l \in L' \neq \emptyset : v R l) \Rightarrow v R (\bigsqcup L')$$

Give an example showing that $v R \perp$ fails even though (4.4) and (4.5) are fulfilled. ■

Exercise 4.3 Show that the Control Flow Analysis of Chapter 3 is indeed correct in the sense of condition (4.3) as claimed in Example 4.4. To do so first show that

$$(\widehat{C}, \widehat{\rho}) \models v^\ell \text{ iff } v \vee (\widehat{\rho}, \widehat{C}(\ell))$$

whenever v is a value (i.e. a closure or a constant). Next show that

$$[] \vdash (e_* v_1^{\ell_1})^\ell \rightarrow^* v^\ell \wedge (\widehat{C}, \widehat{\rho}) \models (e_* v_1^{\ell_1})^\ell \Rightarrow (\widehat{C}, \widehat{\rho}) \models v^\ell$$

is a corollary of Theorem 3.7. Finally assume the premise of (4.3) and use the above results to obtain the desired result. ■

Exercise 4.4 Show that the relation R_{CFA} defined in Example 4.4 is generated by the representation function β_{CFA} also defined in Example 4.7. To do so prove that

$$v R_{\text{CFA}} (\widehat{\rho}, \widehat{v}) \text{ iff } \beta_{\text{CFA}}(v) \sqsubseteq_{\text{CFA}} (\widehat{\rho}, \widehat{v})$$

by induction on the size of v ; only the case where v is close t in ρ is non-trivial. ■

Exercise 4.5 Define $L_i = (\mathcal{P}(V_i), \subseteq)$ (for $i = 1, 2$) and define $f_p : L_1 \rightarrow L_2$ by

$$f_p(l_1) = \{v_2 \in V_2 \mid \exists v_1 \in l_1 : p \vdash v_1 \rightsquigarrow v_2\}$$

Show that f_p is monotone. Next show that $(p \vdash \cdot \rightsquigarrow \cdot)(R_1 \rightarrow R_2) f_p$ where $v_i R_i l_i$ is defined by $v_i \in l_i$. Also, show that for $f' : L_1 \rightarrow L_2$ we have $(p \vdash \cdot \rightsquigarrow \cdot)(R_1 \rightarrow R_2) f'$ if and only if $f_p \subseteq f'$. A semantics that associates a program p with f_p as defined here is sometimes called a *collecting semantics*. Finally, note that R_i is generated by β_i defined by $\beta_i(v_i) = \{v_i\}$; show that $f_p = (\beta_1 \rightarrow \beta_2)(p \vdash \cdot \rightsquigarrow \cdot)$. ■

Exercise 4.6 Show that all of

- \sqcup
- $\lambda(l_1, l_2). \top$
- $\lambda(l_1, l_2). \begin{cases} l_1 & \text{if } l_2 \subseteq l_1 \\ \top & \text{otherwise} \end{cases}$
- $\lambda(l_1, l_2). \begin{cases} l_2 & \text{if } l_1 = \perp \\ l_1 & \text{if } l_2 \subseteq l_1 \wedge l_1 \neq \perp \\ \top & \text{otherwise} \end{cases}$
- $\lambda(l_1, l_2). \begin{cases} l_1 \sqcup l_2 & \text{if } l_1 \subseteq l' \wedge l_2 \subseteq l_1 \\ \top & \text{otherwise} \end{cases}$

are upper bound operators (where l' is some element of L). Determine which of them that are also widening operators. Try to find sufficient conditions on l' such that the operator involving l' is a widening operator. ■

Exercise 4.7 Show that if L satisfies the Ascending Chain Condition then an operator on L is a widening operator if and only if it is an upper bound operator. Conclude that if L satisfies the Ascending Chain Condition then the least upper bound operator $\sqcup : L \times L \rightarrow L$ is a widening operator. ■

Exercise 4.8 Consider changing the definition of f_{∇}^0 from $f_{\nabla}^0 = \perp$ to $f_{\nabla}^0 = l_0$ for some $l_0 \in L$. Possible assumptions on l_0 are:

- $l_0 = f(\perp)$;
- $l_0 = f^{2^7}(\perp)$;
- $l_0 \in \text{Ext}(f)$;
- l_0 arbitrary.

Which of these suffice for proving Fact 4.14 and Proposition 4.13? ■

Exercise 4.9 Let ∇_K be as in Example 4.15 and define

$$int_1 \nabla int_2 = \begin{cases} int_1 \sqcup int_2 & \text{if } int_1 \sqsubseteq int' \vee int_2 \sqsubseteq int_1 \\ int_1 \nabla_K int_2 & \text{otherwise} \end{cases}$$

where int' is an interval satisfying $\inf(int') > -\infty$ and $\sup(int') < \infty$. Show that

$$\forall int_1, int_2 : int_1 \nabla int_2 \sqsubseteq int_1 \nabla_K int_2$$

and that the inequality may be strict. Show that ∇ is an upper bound operator. Determine whether or not ∇ is a widening operator. ■

Exercise 4.10 Let $(l_n)_n$ be a descending chain and let $\Delta : L \times L \rightarrow L$ be a total function that satisfies $l'_2 \sqsubseteq l'_1 \Rightarrow l'_2 \sqsubseteq (l'_1 \Delta l'_2) \sqsubseteq l'_1$ for all $l'_1, l'_2 \in L$. Show that the sequence $(l_n^\Delta)_n$ is a descending chain and that $l_n^\Delta \sqsupseteq l_n$ for all n . ■

Exercise 4.11 Consider the following alternative strategy to narrowing for improving the approximation $f_\nabla^m \in Red(f)$ to the fixed point $lfp(f)$ of $f : L \rightarrow L$. A *descending chain truncator* is a function Υ that maps descending chains $(l_n)_n$ to a non-negative number such that

$$\text{if } (l_n)_n \text{ and } (l'_n)_n \text{ are descending chains and } \forall n \leq \Upsilon((l_n)_n) : l_n = l'_n \\ \text{then } \Upsilon((l_n)_n) = \Upsilon((l'_n)_n).$$

This ensures that Υ is finitely calculatable. The truncated descending chain then is

$$(f_\nabla^m, \dots, f^n(f_\nabla^m), \dots, f^{m'}(f_\nabla^m))$$

where $m' = \Upsilon((f^n(f_\nabla^m))_n)$ and the desired approximation to $lfp(f)$ is

$$lfp_\nabla^\Upsilon = f^{m'}(f_\nabla^m)$$

Prove that this development can be used as an alternative to narrowing and try to determine the relationship between the two concepts. ■

Exercise 4.12 Show that if L satisfies the Descending Chain Condition then the binary greatest lower bound operator $\sqcap : L \times L \rightarrow L$ is a narrowing operator. ■

Exercise 4.13 Consider the complete lattice **Interval** of Figure 4.2 and the complete lattice **Sign'** of Figure 4.10 and define $\gamma_{IS'}$ by

$$\begin{array}{ll} \gamma_{IS'}(\top) &= [-\infty, \infty] & \gamma_{IS'}(-0) &= [-\infty, 0] \\ \gamma_{IS'}(0+) &= [0, \infty] & \gamma_{IS'}(-) &= [-\infty, -1] \\ \gamma_{IS'}(0) &= [0, 0] & \gamma_{IS'}(+) &= [1, \infty] \\ \gamma_{IS'}(\perp) &= \perp \end{array}$$

Show that there exists a Galois connection between **Interval** and **Sign'** with $\gamma_{IS'}$ as the upper adjoint. ■

Exercise 4.14 Let $(\mathcal{P}(V), \alpha_\eta, \gamma_\eta, \mathcal{P}(D))$ be a Galois connection that is given by an extraction function $\eta : V \rightarrow D$. Show that α_η is surjective if and only if η is. Conclude that $(\mathcal{P}(V), \alpha_\eta, \gamma_\eta, \mathcal{P}(D))$ is a Galois insertion if and only if η is surjective. Next show that $\zeta_\eta[\mathcal{P}(D)] = \alpha_\eta[\mathcal{P}(V)]$ is isomorphic (see Appendix A) to $\mathcal{P}(\eta[V])$ where $\eta[V]$ is the image of η and finally verify that $\zeta_\eta(V') = V' \cap \eta[V]$. ■

Exercise 4.15 Assume that the Galois connections $(\mathcal{P}(D_i), \alpha_{\eta_{i+1}}, \gamma_{\eta_{i+1}}, \mathcal{P}(D_{i+1}))$ are given by the extraction functions $\eta_{i+1} : D_i \rightarrow D_{i+1}$. Show that the composition of the Galois connections, $(\mathcal{P}(D_0), \alpha, \gamma, \mathcal{P}(D_2))$, will have $\alpha = \alpha_{\eta_2} \circ \alpha_{\eta_1} = \alpha_{\eta_2 \circ \eta_1}$ and $\gamma = \gamma_{\eta_1} \circ \gamma_{\eta_2} = \gamma_{\eta_2 \circ \eta_1}$; i.e. the Galois connection is given by the extraction function $\eta_2 \circ \eta_1$. ■

Exercise 4.16 Let $(\mathcal{P}(V_1), \alpha_1, \gamma_1, \mathcal{P}(D_1))$ and $(\mathcal{P}(V_2), \alpha_2, \gamma_2, \mathcal{P}(D_2))$ be Galois connections given by the extraction functions $\eta_1 : V_1 \rightarrow D_1$ and $\eta_2 : V_2 \rightarrow D_2$. Furthermore assume that V_1 and V_2 are disjoint and similarly for D_1 and D_2 . Define an extraction function

$$\eta : V_1 \cup V_2 \rightarrow D_1 \cup D_2$$

by

$$\eta(v) = \begin{cases} \eta_1(v) & \text{if } v \in V_1 \\ \eta_2(v) & \text{if } v \in V_2 \end{cases}$$

Show that this defines a Galois connection

$$(\mathcal{P}(V_1 \cup V_2), \alpha_\eta, \gamma_\eta, \mathcal{P}(D_1 \cup D_2))$$

and reformulate it as an isomorphic Galois connection

$$(\mathcal{P}(V_1) \times \mathcal{P}(V_2), \alpha_\eta, \gamma_\eta, \mathcal{P}(D_1) \times \mathcal{P}(D_2))$$

in the manner of the independent attribute method. How important is it that V_1 and V_2 are disjoint and that D_1 and D_2 are disjoint? ■

Exercise 4.17 Let $(L_1, \alpha_1, \gamma_1, M_1)$ and $(L_2, \alpha_2, \gamma_2, M_2)$ be Galois insertions. First define

$$\begin{aligned}\alpha(l_1, l_2) &= (\alpha_1(l_1), \alpha_2(l_2)) \\ \gamma(m_1, m_2) &= (\gamma_1(m_1), \gamma_2(m_2))\end{aligned}$$

and show that $(L_1 \times L_2, \alpha, \gamma, M_1 \times M_2)$ is a Galois insertion. Then define

$$\begin{aligned}\alpha(f) &= \alpha_2 \circ f \circ \gamma_1 \\ \gamma(g) &= \gamma_2 \circ g \circ \alpha_1\end{aligned}$$

and show that $(L_1 \rightarrow L_2, \alpha, \gamma, M_1 \rightarrow M_2)$ is a Galois insertion. ■

Exercise 4.18 Let $(\mathcal{P}(V_1), \alpha_1, \gamma_1, \mathcal{P}(D_1))$ and $(\mathcal{P}(V_2), \alpha_2, \gamma_2, \mathcal{P}(D_2))$ be Galois insertions. Define

$$\begin{aligned}\alpha(VV) &= \bigcup \{ \alpha_1(\{v_1\}) \times \alpha_2(\{v_2\}) \mid (v_1, v_2) \in VV \} \\ \gamma(DD) &= \{ (v_1, v_2) \mid \alpha_1(\{v_1\}) \times \alpha_2(\{v_2\}) \subseteq DD \}\end{aligned}$$

and determine whether or not $(\mathcal{P}(V_1 \times V_2), \alpha, \gamma, \mathcal{P}(D_1 \times D_2))$ is a Galois insertion. ■

Exercise 4.19 Let $(L, \alpha_1, \gamma_1, M_1)$ and $(L, \alpha_2, \gamma_2, M_2)$ be Galois insertions. Define

$$\begin{aligned}\alpha(l) &= (\alpha_1(l), \alpha_2(l)) \\ \gamma(m_1, m_2) &= \gamma_1(m_1) \sqcap \gamma_2(m_2)\end{aligned}$$

and determine whether or not $(L, \alpha, \gamma, M_1 \times M_2)$ is a Galois insertion. ■

Exercise 4.20 Let $(L_1, \alpha_1, \gamma_1, M_1)$ be a Galois connection and define

$$\begin{aligned}\alpha(f) &= \alpha_1 \circ f \circ \gamma_1 \\ \gamma(g) &= \gamma_1 \circ g \circ \alpha_1\end{aligned}$$

(in the manner of Section 4.4). Do any of the following equalities

$$\begin{aligned}\alpha(\lambda l) &= \lambda m.m & \gamma(\lambda m.m) &= \lambda l.l \\ \alpha(f_1 \circ f_2) &= \alpha(f_1) \circ \alpha(f_2) & \gamma(g_1 \circ g_2) &= \gamma(g_1) \circ \gamma(g_2)\end{aligned}$$

necessarily hold? Which of the equalities hold when $(L_1, \alpha_1, \gamma_1, M_1)$ is known to be a Galois insertion? ■

Exercise 4.21 Consider the Galois insertion

$$(\mathcal{P}(\mathbf{Z} \times \mathbf{Z}), \alpha_{\text{SSR}'}, \gamma_{\text{SSR}'}, \mathcal{P}(\mathbf{AB}))$$

developed in Example 4.39. Determine for each of the sets

$$\begin{aligned} &\{(x, y) \mid x = y\} \\ &\{(x, y) \mid x = -y\} \\ &\{(x, y) \mid x = y + 1\} \\ &\{(x, y) \mid x = y + 3\} \\ &\{(x, y) \mid x \geq y\} \\ &\{(x, y) \mid x \geq y + 1\} \\ &\{(x, y) \mid x^2 + y^2 \leq 100\} \end{aligned}$$

the best description in $\mathcal{P}(\mathbf{AB})$. ■

Exercise 4.22 Let $(L_i, \alpha_i, \gamma_i, M_i)$ be Galois connections for $i = 1, 2, 3$. Use the approach of Section 4.4 to define

$$\begin{aligned} \alpha(f) &= \dots \\ \gamma(g) &= \dots \end{aligned}$$

such that $((L_1 \times L_2) \rightarrow L_3, \alpha, \gamma, (M_1 \times M_2) \rightarrow M_3)$ is a Galois connection.

Next let all of $(L_i, \alpha_i, \gamma_i, M_i)$ be the Galois connection

$$(\mathcal{P}(\mathbf{Z}), \alpha_{\text{ZI}}, \gamma_{\text{ZI}}, \mathbf{Interval})$$

of Example 4.19 that relates the set of integers to the intervals. Let $\text{plus} : \mathcal{P}(\mathbf{Z}) \times \mathcal{P}(\mathbf{Z}) \rightarrow \mathcal{P}(\mathbf{Z})$ be the "pointwise" application of addition defined by:

$$\lambda(Z_1, Z_2). \{z_1 + z_2 \mid z_1 \in Z_1 \wedge z_2 \in Z_2\}$$

Next define

$$\alpha(\text{plus}) = \lambda(\text{int}_1, \text{int}_2). \dots$$

and supply the details of the definition. ■

Exercise 4.23 Let L be the complete lattice of sets of integers, let M be the complete lattice of intervals of Example 4.10, let (L, α, γ, M) be the Galois insertion of Example 4.19, and let $\nabla_M : M \times M \rightarrow M$ be the widening operator of Example 4.15. Observe that the formula

$$l_1 \nabla_L l_2 = \gamma(\alpha(l_1) \nabla_M \alpha(l_2))$$

defines a widening operator $\nabla_L : L \times L \rightarrow L$ and develop a formula for it (in the manner of the formula for ∇_M of Example 4.15). ■

Exercise 4.24 Suppose that (L, α, γ, M) is a Galois connection or a Galois insertion and that possibly M satisfies the Descending Chain Condition. Let $\Delta_M : M \times M \rightarrow M$ be a narrowing operator and try to determine if the formula

$$l_1 \Delta_L l_2 = \gamma(\alpha(l_1) \Delta_M \alpha(l_2))$$

defines a narrowing operator $\Delta_L : L \times L \rightarrow L$. ■

Chapter 5

Type and Effect Systems

So far our techniques have applied equally well to typed and untyped programming languages. This flexibility does not apply to the development to be performed in this chapter: here we demand that our programming language is typed because we will use the syntax for types in order to express the program analysis properties of interest (as was already illustrated in Section 1.6).

We shall first present an Annotated Type System for *Control Flow Analysis* in Section 5.1, demonstrate its semantic soundness and other theoretical properties in Section 5.2, and then in Section 5.3 show how to obtain an algorithm for computing the annotated types (and prove that it is sound and complete). In Sections 5.4 and 5.5 we give examples of other analyses specified by Type and Effect Systems. In Section 5.4 we study Type and Effect Systems with rules for subtyping, polymorphism and polymorphic recursion and illustrate their use in an analysis for tracking *Side Effects*, an *Exception Analysis* and an analysis for *Region Inference*. Finally, in Section 5.5 we show that the annotations can be given more structure and we illustrate this for a *Communication Analysis*.

5.1 Control Flow Analysis

Syntax of the FUN language. To illustrate the approach we shall make use of the functional language FUN also considered in Chapter 3; that the approach also applies to the imperative language of Chapter 2 was briefly sketched in Section 1.6. However, in this chapter we shall use a slightly different labelling scheme from the one in Chapter 3; the syntactic category of interest is

$e \in \text{Exp}$ expressions

and it is defined by:

$$e ::= c \mid x \mid \text{fn}_{\pi} x \Rightarrow e_0 \mid \text{fun}_{\pi} f x \Rightarrow e_0 \mid e_1 e_2 \\ \mid \text{if } e_0 \text{ then } e_1 \text{ else } e_2 \mid \text{let } x = e_1 \text{ in } e_2 \mid e_1 \text{ op } e_2$$

The *program points*, $\pi \in \mathbf{Pnt}$, are used to name the function abstractions in the program; this could also be done using the notion of labelled terms from Chapter 3 but for the present purposes we do not need the full generality of this machinery – the reason is that now we will use the types to record information that was previously associated with labels. Hence our syntax just makes use of expressions and dispenses with terms.

As in the previous chapters we shall assume that a countable set of variables is given and that constants (including the truth values), binary operators (including the natural arithmetic, boolean and relational operators) and program points are left unspecified:

$$\begin{array}{ll} c \in \mathbf{Const} & \text{constants} \\ op \in \mathbf{Op} & \text{binary operators} \\ f, x \in \mathbf{Var} & \text{variables} \\ \pi \in \mathbf{Pnt} & \text{program points} \end{array}$$

Example 5.1 The functional program $(\text{fn } x \Rightarrow x) (\text{fn } y \Rightarrow y)$ considered in Chapters 1 and 3 is now written as

$$(\text{fn}_x x \Rightarrow x) (\text{fn}_y y \Rightarrow y)$$

just as we did in Example 1.5. ■

Example 5.2 The expression loop of Example 3.2 is now written:

$$\text{let } g = (\text{fun}_f f x \Rightarrow f (\text{fn}_y y \Rightarrow y)) \\ \text{in } g (\text{fn}_z z \Rightarrow z)$$

Recall that this is a looping program: g is first applied to the identity function $\text{fn}_z z \Rightarrow z$ but it ignores its argument and calls itself recursively with the function $\text{fn}_y y \Rightarrow y$. ■

5.1.1 The Underlying Type System

The analyses will be specified as extensions of the ordinary type system in order to record the program properties of interest. For this reason the ordinary type system is sometimes called the *underlying type system* and we shall start by specifying it.

[con]	$\Gamma \vdash_{\text{UL}} c : \tau_c$
[var]	$\Gamma \vdash_{\text{UL}} x : \tau \quad \text{if } \Gamma(x) = \tau$
[fn]	$\frac{\Gamma[x \mapsto \tau_x] \vdash_{\text{UL}} e_0 : \tau_0}{\Gamma \vdash_{\text{UL}} \text{fn}_\pi x \Rightarrow e_0 : \tau_x \rightarrow \tau_0}$
[fun]	$\frac{\Gamma[f \mapsto \tau_x \rightarrow \tau_0][x \mapsto \tau_x] \vdash_{\text{UL}} e_0 : \tau_0}{\Gamma \vdash_{\text{UL}} \text{fun}_\pi f x \Rightarrow e_0 : \tau_x \rightarrow \tau_0}$
[app]	$\frac{\Gamma \vdash_{\text{UL}} e_1 : \tau_2 \rightarrow \tau_0 \quad \Gamma \vdash_{\text{UL}} e_2 : \tau_2}{\Gamma \vdash_{\text{UL}} e_1 e_2 : \tau_0}$
[if]	$\frac{\Gamma \vdash_{\text{UL}} e_0 : \text{bool} \quad \Gamma \vdash_{\text{UL}} e_1 : \tau \quad \Gamma \vdash_{\text{UL}} e_2 : \tau}{\Gamma \vdash_{\text{UL}} \text{if } e_0 \text{ then } e_1 \text{ else } e_2 : \tau}$
[let]	$\frac{\Gamma \vdash_{\text{UL}} e_1 : \tau_1 \quad \Gamma[x \mapsto \tau_1] \vdash_{\text{UL}} e_2 : \tau_2}{\Gamma \vdash_{\text{UL}} \text{let } x = e_1 \text{ in } e_2 : \tau_2}$
[op]	$\frac{\Gamma \vdash_{\text{UL}} e_1 : \tau_{op}^1 \quad \Gamma \vdash_{\text{UL}} e_2 : \tau_{op}^2}{\Gamma \vdash_{\text{UL}} e_1 \text{ op } e_2 : \tau_{op}}$

Table 5.1: The Underlying Type System.

Types. Let us first introduce the notions of *types* and *type environments*:

$\tau \in \mathbf{Type}$ types

$\Gamma \in \mathbf{TEnv}$ type environments

We shall assume that the types are given by

$$\tau ::= \text{int} \mid \text{bool} \mid \tau_1 \rightarrow \tau_2$$

where `int` and `bool` are the only two kinds of base types and as usual we use arrows for function types. Each constant $c \in \mathbf{Const}$ has a type that we shall denote τ_c so e.g. `true` has type $\tau_{\text{true}} = \text{bool}$ and `7` has type $\tau_7 = \text{int}$. Each binary operator op will expect two arguments of type τ_{op}^1 and τ_{op}^2 , respectively, and give a result of type τ_{op} – an example is the relational operation \leq that expects two arguments of type `int` and gives a result of type `bool`. For the sake of simplicity we shall *assume* that all the constants have base types and that all binary operators expect arguments of base types and return values of base types.

The type environments are given by:

$$\Gamma ::= [] \mid \Gamma[x \mapsto \tau]$$

Formally, Γ is a list but nevertheless we shall feel free to regard it as a finite mapping: we write $\text{dom}(\Gamma)$ for $\{x \mid \Gamma \text{ contains } [x \mapsto \dots]\}$; we write $\Gamma(x) = \tau$ if $x \in \text{dom}(\Gamma)$ and the *rightmost* occurrence of $[x \mapsto \dots]$ in Γ is $[x \mapsto \tau]$, and we write $\Gamma|X$ for the type environment obtained from Γ by removing all occurrences of $[x \mapsto \dots]$ with $x \notin X$. For the sake of readability we shall write $[x \mapsto \tau]$ for $[\] [x \mapsto \tau]$.

Typing judgements. The general form of a typing is given by

$$\Gamma \vdash_{\text{UL}} e : \tau$$

that says that the expression e has type τ assuming that any free variable has type given by Γ . The axioms and rules for the judgements are listed in Table 5.1 and are explained below.

The axioms $[\text{con}]$ and $[\text{var}]$ are straightforward: the first uses the predefined type for the constant and the second consults the type environment. In the rules $[\text{fn}]$ and $[\text{fun}]$ we *guess* a type for the bound variables and determine the types of the bodies under the additional assumptions; the rule $[\text{fun}]$ has the implicit requirement that the guess of the type for f matches that of the resulting function. As a consequence of these two rules the type system is *nondeterministic* in the sense that $\Gamma \vdash_{\text{UL}} e : \tau_1$ and $\Gamma \vdash_{\text{UL}} e : \tau_2$ does not necessarily imply that $\tau_1 = \tau_2$.

The rule $[\text{app}]$ requires that the operator and the operand of the application can be typed and implicitly it requires that the type of the operator is a function type where the type before the arrow equals that of the operand – in this way we express that the types of the formal and actual parameter must be equal.

The rules $[\text{if}]$, $[\text{let}]$ and $[\text{op}]$ are straightforward. In particular, the let-construct $\text{let } x = e_1 \text{ in } e_2$ admits exactly the same typings as the application $(\text{fn}_\pi x \Rightarrow e_2) e_1$ and regardless of the choice of π . In Sections 5.4 and 5.5 we shall consider a polymorphic let-construct where $\text{let } x = e_1 \text{ in } e_2$ may admit more typings than $(\text{fn}_\pi x \Rightarrow e_2) e_1$.

Example 5.3 Let us show that the expression loop

```
let g = (funF f x => f (fnY y => y))
in g (fnZ z => z)
```

of Example 5.2 has type $\tau \rightarrow \tau$ for each type τ . We shall first consider the expression $\text{fun}_F f x \Rightarrow f (\text{fn}_Y y \Rightarrow y)$ where we write Γ_{f_x} for the type environment $[f \mapsto (\tau \rightarrow \tau) \rightarrow (\tau \rightarrow \tau)][x \mapsto \tau \rightarrow \tau]$. Then we get

$$\begin{aligned} \Gamma_{f_x} \vdash_{\text{UL}} f : (\tau \rightarrow \tau) \rightarrow (\tau \rightarrow \tau) \\ \Gamma_{f_x} \vdash_{\text{UL}} \text{fn}_Y y \Rightarrow y : \tau \rightarrow \tau \end{aligned}$$

using the axiom $[var]$ and the rule $[fn]$. The rule $[app]$ then gives

$$\Gamma_{fx} \vdash_{UL} f (fn_Y y \Rightarrow y) : \tau \rightarrow \tau$$

and we have a judgement matching the premise of the rule $[fun]$. Hence we get

$$[] \vdash_{UL} fun_F f x \Rightarrow f (fn_Y y \Rightarrow y) : (\tau \rightarrow \tau) \rightarrow (\tau \rightarrow \tau)$$

Taking Γ_g to be the type environment $[g \mapsto (\tau \rightarrow \tau) \rightarrow (\tau \rightarrow \tau)]$ we get

$$\Gamma_g \vdash_{UL} g (fn_Z z \Rightarrow z) : \tau \rightarrow \tau$$

using the axiom $[var]$ and the rules $[fn]$ and $[app]$. The rule $[let]$ can now be used to show that the expression `loop` has type $\tau \rightarrow \tau$. ■

5.1.2 The Analysis

Annotated types. That a function has type $\tau_1 \rightarrow \tau_2$ means that given an argument of type τ_1 it will return a value of type τ_2 in case it terminates. To get a *Control Flow Analysis* we shall annotate the type with information, $\varphi \in \mathbf{Ann}$, about which function it might be. The annotations are given by

$$\varphi \in \mathbf{Ann} \text{ annotations}$$

where:

$$\varphi ::= \{\pi\} \mid \varphi_1 \cup \varphi_2 \mid \emptyset$$

So φ will be a *set* of function names – describing the set of function definitions that can result in a function of a given type; as will be discussed below, we shall feel free to write $\{\pi_1, \dots, \pi_n\}$ for $\{\pi_1\} \cup \dots \cup \{\pi_n\}$. We now take

$$\hat{\tau} \in \widehat{\mathbf{Type}} \text{ annotated types}$$

$$\hat{\Gamma} \in \widehat{\mathbf{TEnv}} \text{ annotated type environments}$$

and define:

$$\hat{\tau} ::= \text{int} \mid \text{bool} \mid \hat{\tau}_1 \xrightarrow{\varphi} \hat{\tau}_2$$

$$\hat{\Gamma} ::= [] \mid \hat{\Gamma}[x \mapsto \hat{\tau}]$$

We shall write $[\hat{\tau}]$ for the *underlying type* corresponding to the annotated type $\hat{\tau}$; it is defined as follows:

$$\begin{aligned} [\text{int}] &= \text{int} \\ [\text{bool}] &= \text{bool} \\ [\hat{\tau}_1 \xrightarrow{\varphi} \hat{\tau}_2] &= [\hat{\tau}_1] \rightarrow [\hat{\tau}_2] \end{aligned}$$

As an example we have $[\text{int} \xrightarrow{\{x\}} \text{int}] = (\text{int} \rightarrow \text{int})$. Furthermore, we extend the notation to operate on type environments so $[\hat{\Gamma}](x) = [\hat{\Gamma}(x)]$ for all x .

$[con]$	$\hat{\Gamma} \vdash_{CFA} c : \tau_c$
$[var]$	$\hat{\Gamma} \vdash_{CFA} x : \hat{\tau} \quad \text{if } \hat{\Gamma}(x) = \hat{\tau}$
$[fn]$	$\frac{\hat{\Gamma}[x \mapsto \hat{\tau}_x] \vdash_{CFA} e_0 : \hat{\tau}_0}{\hat{\Gamma} \vdash_{CFA} fn_{\pi} x \Rightarrow e_0 : \hat{\tau}_x \quad \{\pi\} \cup \varphi \rightarrow \tau_0}$
$[fun]$	$\frac{\hat{\Gamma}[f \mapsto \hat{\tau}_x \quad \{\pi\} \cup \varphi \rightarrow \hat{\tau}_0][x \mapsto \hat{\tau}_x] \vdash_{CFA} e_0 : \hat{\tau}_0}{\hat{\Gamma} \vdash_{CFA} fun_{\pi} f x \Rightarrow e_0 : \hat{\tau}_x \quad \{\pi\} \cup \varphi \rightarrow \hat{\tau}_0}$
$[app]$	$\frac{\hat{\Gamma} \vdash_{CFA} e_1 : \hat{\tau}_2 \xrightarrow{\varphi} \hat{\tau}_0 \quad \hat{\Gamma} \vdash_{CFA} e_2 : \hat{\tau}_2}{\hat{\Gamma} \vdash_{CFA} e_1 e_2 : \hat{\tau}_0}$
$[if]$	$\frac{\hat{\Gamma} \vdash_{CFA} e_0 : bool \quad \hat{\Gamma} \vdash_{CFA} e_1 : \hat{\tau} \quad \hat{\Gamma} \vdash_{CFA} e_2 : \hat{\tau}}{\hat{\Gamma} \vdash_{CFA} \text{if } e_0 \text{ then } e_1 \text{ else } e_2 : \hat{\tau}}$
$[let]$	$\frac{\hat{\Gamma} \vdash_{CFA} e_1 : \hat{\tau}_1 \quad \hat{\Gamma}[x \mapsto \hat{\tau}_1] \vdash_{CFA} e_2 : \hat{\tau}_2}{\hat{\Gamma} \vdash_{CFA} \text{let } x = e_1 \text{ in } e_2 : \hat{\tau}_2}$
$[op]$	$\frac{\hat{\Gamma} \vdash_{CFA} e_1 : \tau_{op}^1 \quad \hat{\Gamma} \vdash_{CFA} e_2 : \tau_{op}^2}{\hat{\Gamma} \vdash_{CFA} e_1 \cdot op \ e_2 : \tau_{op}}$

Table 5.2: Control Flow Analysis.

Judgements. The judgements of the *Control Flow Analysis* have the form

$$\hat{\Gamma} \vdash_{CFA} e : \hat{\tau}$$

and are defined by the axioms and rules of Table 5.2. The clauses $[fn]$ and $[fun]$ annotate the arrow of the resulting function type with the information that the abstraction named π should be included in the set of possible functions; the use of $\{\pi\} \cup \varphi$ indicates that we may want to include other names as well and we shall say that the Type and Effect System allows *subeffecting* (see the Concluding Remarks). In Example 5.5 below we shall give an example where subeffecting is indeed needed to analyse the expression. The remaining clauses of Table 5.2 are straightforward modifications of the similar clauses of Table 5.1.

Example 5.4 Let us return to the expression

$$(fn_X x \Rightarrow x) (fn_Y y \Rightarrow y)$$

of Example 5.1. Writing $\hat{\tau}_Y$ for $\text{int} \xrightarrow{\{Y\}} \text{int}$ we have the following inference tree:

$$\frac{\frac{[x \mapsto \hat{\tau}_Y] \vdash_{\text{CFA}} x : \hat{\tau}_Y}{[] \vdash_{\text{CFA}} \text{fn}_X x \Rightarrow x : \hat{\tau}_Y \xrightarrow{\{X\}} \hat{\tau}_Y} \quad \frac{[y \mapsto \text{int}] \vdash_{\text{CFA}} y : \text{int}}{[] \vdash_{\text{CFA}} \text{fn}_Y y \Rightarrow y : \hat{\tau}_Y}}{[] \vdash_{\text{CFA}} (\text{fn}_X x \Rightarrow x) (\text{fn}_Y y \Rightarrow y) : \hat{\tau}_Y}$$

Note that the whole inference tree is needed to get full information about the control flow properties of the expression. If we label all the subexpressions

$$((\text{fn}_X x \Rightarrow x^1)^2 (\text{fn}_Y y \Rightarrow y^3)^4)^5$$

as in Chapter 3 then we can list the types of the subexpressions as follows

ℓ	1	2	3	4	5
$\hat{C}(\ell)$	$\hat{\tau}_Y$	$\hat{\tau}_Y \xrightarrow{\{X\}} \hat{\tau}_Y$	int	$\hat{\tau}_Y$	$\hat{\tau}_Y$

and we are close to the information supplied by \hat{C} in Chapter 3. The information corresponding to $\hat{\rho}$ can be obtained by “merging” information from the various type environments of the inference tree (see Exercise 5.4). ■

Example 5.5 Consider once again the expression loop

```
let g = (funF f x => f (fnY y => y))
in g (fnZ z => z)
```

and let us write $\hat{\Gamma}_{f_x}$ for $[f \mapsto (\hat{\tau} \xrightarrow{\{Y,Z\}} \hat{\tau}) \xrightarrow{\{F\}} (\hat{\tau} \xrightarrow{\emptyset} \hat{\tau})][x \mapsto \hat{\tau} \xrightarrow{\{Y,Z\}} \hat{\tau}]$. Using the clause $[fn]$ we have

$$\hat{\Gamma}_{f_x} \vdash_{\text{CFA}} \text{fn}_Y y \Rightarrow y : \hat{\tau} \xrightarrow{\{Y,Z\}} \hat{\tau}$$

where we exploit that subeffecting allows us to enlarge the annotation from the minimal $\{Y\}$ to $\{Y, Z\}$. Using this we can construct an inference tree for:

$$[] \vdash_{\text{CFA}} \text{fun}_F f x \Rightarrow f (\text{fn}_Y y \Rightarrow y) : (\hat{\tau} \xrightarrow{\{Y,Z\}} \hat{\tau}) \xrightarrow{\{F\}} (\hat{\tau} \xrightarrow{\emptyset} \hat{\tau})$$

Next let $\hat{\Gamma}_g$ be $[g \mapsto (\hat{\tau} \xrightarrow{\{Y,Z\}} \hat{\tau}) \xrightarrow{\{F\}} (\hat{\tau} \xrightarrow{\emptyset} \hat{\tau})]$. We now exploit that the annotation $\{Z\}$ can be enlarged to $\{Z, Y\}$ in the clause $[fn]$ so that:

$$\hat{\Gamma}_g \vdash_{\text{CFA}} \text{fn}_Z z \Rightarrow z : \hat{\tau} \xrightarrow{\{Z,Y\}} \hat{\tau}$$

Since $\{Z, Y\} = \{Y, Z\}$ we can use the clause $[app]$ and get

$$\hat{\Gamma}_g \vdash_{\text{CFA}} g (\text{fn}_Z z \Rightarrow z) : \hat{\tau} \xrightarrow{\emptyset} \hat{\tau}$$

$[unit]$	$\varphi = \varphi \cup \emptyset$	$[idem]$	$\varphi = \varphi \cup \varphi$
$[com]$	$\varphi_1 \cup \varphi_2 = \varphi_2 \cup \varphi_1$	$[ass]$	$\varphi_1 \cup (\varphi_2 \cup \varphi_3) = (\varphi_1 \cup \varphi_2) \cup \varphi_3$
$[trans]$	$\varphi = \varphi$		
$[ref]$	$\frac{\varphi_1 = \varphi_2 \quad \varphi_2 = \varphi_3}{\varphi_1 = \varphi_3}$	$[cong]$	$\frac{\varphi_1 = \varphi'_1 \quad \varphi_2 = \varphi'_2}{\varphi_1 \cup \varphi_2 = \varphi'_1 \cup \varphi'_2}$

Table 5.3: Equality of Annotations.

and eventually $[] \vdash_{CFA} \text{loop} : \hat{\tau} \xrightarrow{\emptyset} \hat{\tau}$. This can be interpreted as saying that the expression `loop` does not terminate: the type is $\hat{\tau} \xrightarrow{\emptyset} \hat{\tau}$ but the annotation \emptyset indicates that there will be no function abstractions with the given type.

Actually we can show

$$[] \vdash_{CFA} \text{fun}_f f x \Rightarrow f (f_{NY} y \Rightarrow y) : (\hat{\tau} \xrightarrow{\{Y,Z\}} \hat{\tau}) \xrightarrow{\{F\}} (\hat{\tau} \xrightarrow{\varphi} \hat{\tau})$$

for every annotation φ and hence we have $[] \vdash_{CFA} \text{loop} : \hat{\tau} \xrightarrow{\varphi} \hat{\tau}$ for every annotation φ ; clearly the judgement with $\varphi = \emptyset$ is the most informative. ■

Equivalence of annotations. There are a few subtleties involved in using this simple-minded system for Control Flow Analysis. One is the implicit decision to regard a type like $\hat{\tau} \xrightarrow{\{Y,Z\}} \hat{\tau}$ as being equal to $\hat{\tau} \xrightarrow{\{Z,Y\}} \hat{\tau}$.

Concerning this subtlety, we already explained that we feel free to write $\{\pi_1, \dots, \pi_n\}$ for $\{\pi_1\} \cup \dots \cup \{\pi_n\}$. To be utterly formal we should really say that we write $\{\pi_1, \dots, \pi_n\}$ for $((\emptyset \cup \{\pi_1\}) \cup \dots) \cup \{\pi_n\}$.

Next we allow to replace $\tau_1 \xrightarrow{\varphi_1} \tau_2$ by $\tau_1 \xrightarrow{\varphi_2} \tau_2$ whenever φ_1 and φ_2 are “equal as sets”. To be utterly formal this can be axiomatised by the axioms and rules of Table 5.3: the axioms $[unit]$, $[idem]$, $[com]$ and $[ass]$ express that set union has a unit and is idempotent, commutative and associative, and the axioms and rules $[trans]$, $[ref]$ and $[cong]$ ensure that equality is an equivalence relation as well as a congruence.

Finally, we allow to replace $\hat{\tau}_1$ by $\hat{\tau}_2$ if they have the same underlying types and all annotations on corresponding function arrows are “equal as sets”. To be utterly formal we could axiomatise this by:

$$\hat{\tau} = \hat{\tau} \quad \frac{\hat{\tau}_1 = \hat{\tau}'_1 \quad \hat{\tau}_2 = \hat{\tau}'_2 \quad \varphi = \varphi'}{(\hat{\tau}_1 \xrightarrow{\varphi} \hat{\tau}_2) = (\hat{\tau}'_1 \xrightarrow{\varphi'} \hat{\tau}'_2)}$$

It is customary to be informal about these fine technical points. But to avoid confusion one should at the very least point out that annotations are

considered equal modulo the existence of a unit, commutativity, associativity, and idempotence; the abbreviation *UCAI* is often used for this.

Conservative extension. Another subtlety is the ability to give the abstraction $\text{fn}_Y y \Rightarrow y$ the type $\hat{\tau} \xrightarrow{\{Y,Z\}} \hat{\tau}$. Suppose for a moment that the two rules for function abstraction did *not* have a $\{\pi\} \cup \varphi$ annotation on the function arrows but only a $\{\pi\}$. Then $\text{fn}_Y y \Rightarrow y$ would have type $\hat{\tau} \xrightarrow{\{Y\}} \hat{\tau}$ but not $\hat{\tau} \xrightarrow{\{Y,Z\}} \hat{\tau}$; consequently the program from Example 5.5 would have no type in the Annotated Type System for Control Flow Analysis! This is a very undesirable property: we ought to be able to analyse all programs.

To ensure that our system does not have the above deficiency we shall formally prove that the Control Flow Analysis of Table 5.2 is a *conservative extension* of the underlying type system of Table 5.1. This is expressed by:

Fact 5.6

- (i) If $\hat{\Gamma} \vdash_{\text{CFA}} e : \hat{\tau}$ then $[\hat{\Gamma}] \vdash_{\text{UL}} e : [\hat{\tau}]$.
- (ii) If $\Gamma \vdash_{\text{UL}} e : \tau$ then there exists $\hat{\Gamma}$ and $\hat{\tau}$ such that $\hat{\Gamma} \vdash_{\text{CFA}} e : \hat{\tau}$, $[\hat{\Gamma}] = \Gamma$ and $[\hat{\tau}] = \tau$. ■

Proof The proof of (i) is straightforward. For the proof of (ii) one annotates all arrows with the set of all program points in e . ■

This result paves the way for extending $[\cdot]$ to operate on entire typings: applied to the typing $\hat{\Gamma} \vdash_{\text{CFA}} e : \hat{\tau}$ it produces a typing of the form $[\hat{\Gamma}] \vdash_{\text{UL}} e : [\hat{\tau}]$. As an example, the typing of Example 5.5 is transformed into the typing of Example 5.3 (assuming that $[\hat{\tau}] = \tau$).

In Section 5.4 we shall study explicit inference rules for *subeffecting*: this is a related technique for ensuring that the analysis is a conservative extension of the underlying type system.

5.2 Theoretical Properties

Having specified the analysis we shall now ensure that it is semantically correct. Furthermore, the fact that the analysis is a conservative extension of the underlying type system motivates the following result: whenever we have a typing in the underlying type system then the set of typings of the Control Flow Analysis constitutes a Moore family. So as in Section 3.2 (and Exercise 2.7) every acceptable expression can be analysed and it has a best analysis.

As in Sections 2.2 and 3.2, the material of this section may be skimmed through on a first reading; however, we re-iterate that it is frequently when

$[con]$	$\vdash c \rightarrow c$
$[fn]$	$\vdash (fn_{\pi} x \Rightarrow e_0) \rightarrow (fn_{\pi} x \Rightarrow e_0)$
$[fun]$	$\vdash (fun_{\pi} f x \Rightarrow e_0) \rightarrow fn_{\pi} x \Rightarrow (e_0[f \mapsto fun_{\pi} f x \Rightarrow e_0])$
$[app]$	$\frac{\vdash e_1 \rightarrow (fn_{\pi} x \Rightarrow e_0) \quad \vdash e_2 \rightarrow v_2 \quad \vdash e_0[x \mapsto v_2] \rightarrow v_0}{\vdash e_1 e_2 \rightarrow v_0}$
$[if_1]$	$\frac{\vdash e_0 \rightarrow true \quad \vdash e_1 \rightarrow v_1}{\vdash \text{if } e_0 \text{ then } e_1 \text{ else } e_2 \rightarrow v_1}$
$[if_2]$	$\frac{\vdash e_0 \rightarrow false \quad \vdash e_2 \rightarrow v_2}{\vdash \text{if } e_0 \text{ then } e_1 \text{ else } e_2 \rightarrow v_2}$
$[let]$	$\frac{\vdash e_1 \rightarrow v_1 \quad \vdash e_2[x \mapsto v_1] \rightarrow v_2}{\vdash \text{let } x = e_1 \text{ in } e_2 \rightarrow v_2}$
$[op]$	$\frac{\vdash e_1 \rightarrow v_1 \quad \vdash e_2 \rightarrow v_2}{\vdash e_1 \text{ op } e_2 \rightarrow v} \quad \text{if } v_1 \text{ op } v_2 = v$

Table 5.4: Natural Semantics for FUN.

conducting the correctness proof that the final and subtle errors in the analysis are found and corrected!

5.2.1 Natural Semantics

To prove the semantic correctness of the analysis we need to define the semantics. Many kinds of semantics would be appropriate but among the operational semantics we shall now prefer a *Natural Semantics* (i.e. big-step operational semantics) without environments because this makes semantic correctness somewhat easier to establish. This is related to the discussion in the Concluding Remarks of Chapter 3 about the difference between using a big-step operational semantics without environments and a small-step operational semantics with environments; thus our correctness statement will be somewhat weak in case of looping programs.

Transitions. The Natural Semantics will have transitions of the form

$$\vdash e \rightarrow v$$

meaning that the expression e evaluates to the value v . We shall *assume* that $e \in \mathbf{Exp}$ is a *closed* expression, i.e. $FV(e) = \emptyset$, meaning that e does not have

any free variables. The *values*

$$v \in \mathbf{Val} \text{ values}$$

will be a subset of the expressions given by the syntax

$$v ::= c \mid \mathbf{fn}_\pi x \Rightarrow e_0 \quad \text{provided that } FV(\mathbf{fn}_\pi x \Rightarrow e_0) = \emptyset$$

where we demand that values are closed expressions. Compared to the Structural Operational Semantics of Section 3.2, it is not necessary to introduce environments since the bound variables will be syntactically replaced by their value as soon as they become free. Hence there is no need for a `close-construct` nor for a `bind-construct`.

As usual we shall write $e_1[x \mapsto e_2]$ for substituting e_2 for all free occurrences of x in e_1 . It will be the case that e_2 is closed whenever we use this notation and therefore there is no risk of variable capture and hence no need to rename bound variables. Throughout we shall *assume* that $\mathbf{fn}_\pi f x \Rightarrow e_0$ uses distinct variables for f and x . The semantics is given by Table 5.4 and is explained below.

The axioms $[con]$ and $[fn]$ express that the constant and the function abstraction, respectively, evaluates to themselves. For recursive function abstractions we shall unfold the recursion one level as expressed by the axiom $[fun]$; note that the function abstraction being created inherits the program point of the recursive function abstraction. The rule $[app]$ for application expresses that first we evaluate the operator, then the operand, and next we substitute the actual parameter for the formal parameter in the body of the abstraction and evaluate the body. We have only one rule for application because the axiom $[fun]$ ensures that all function abstractions will be of the form $\mathbf{fn}_\pi x \Rightarrow e_0$. The rules for the conditional, the `let-construct` and the binary operators should be straightforward.

Example 5.7 Consider the expression $(\mathbf{fn}_x x \Rightarrow x) (\mathbf{fn}_y y \Rightarrow y)$ of Example 5.1. Using the axiom $[fn]$ we have

$$\begin{aligned} \vdash \mathbf{fn}_x x \Rightarrow x &\longrightarrow \mathbf{fn}_x x \Rightarrow x \\ \vdash \mathbf{fn}_y y \Rightarrow y &\longrightarrow \mathbf{fn}_y y \Rightarrow y \\ \vdash x[x \mapsto \mathbf{fn}_y y \Rightarrow y] &\longrightarrow \mathbf{fn}_y y \Rightarrow y \end{aligned}$$

and we can apply the rule $[app]$ to get:

$$\vdash (\mathbf{fn}_x x \Rightarrow x) (\mathbf{fn}_y y \Rightarrow y) \longrightarrow \mathbf{fn}_y y \Rightarrow y$$

In Example 3.7 we showed how the Structural Operational Semantics deals with this expression. ■

Example 5.8 Next consider the expression `loop`

```
let g = (funF f x => f (fnY y => y))
in g (fnZ z => z)
```

and let us see how this looping program is modelled in the Natural Semantics. First we observe that the axiom *[fun]* gives

$$\begin{aligned} \vdash \text{fun}_F f x => f (\text{fn}_Y y => y) \longrightarrow \\ \text{fn}_F x => ((\text{fun}_F f x => f (\text{fn}_Y y => y)) (\text{fn}_Y y => y)) \end{aligned}$$

so we have replaced the recursive call of `f` with the recursive function definition itself. Turning to the body of the `let`-construct we have to replace the occurrence of `g` with the abstraction:

$$\text{fn}_F x => ((\text{fun}_F f x => f (\text{fn}_Y y => y)) (\text{fn}_Y y => y))$$

The operator will now evaluate to this value and the operand `fnZ z => z` will evaluate to itself. So the next step is to determine a value v such that we have an inference tree for

$$\vdash (\text{fun}_F f x => f (\text{fn}_Y y => y)) (\text{fn}_Y y => y) \longrightarrow v \quad (5.1)$$

and after that we are in a position to use the rule for application. The evaluation of the operator in (5.1) proceeds as before and so does the evaluation of the operand and once more we are left with the obligation to construct an inference tree for the judgement (5.1). Thus we have encountered a *circularity* and we see that the *looping* of the program is modelled in the semantics by the *absence* of an inference tree. In Example 3.8 we showed how the Structural Operational Semantics deals with the expression loop. ■

It is immediate to verify that if e is a closed expression and $\vdash e \longrightarrow v$ then all $\vdash e' \longrightarrow v'$ occurring in the corresponding inference tree (in particular $\vdash e \longrightarrow v$ itself) will have both e' and v' to be closed.

5.2.2 Semantic Correctness

To be able to express the semantic correctness of the analysis we need to assume that the types of the binary operators *op* and their semantics are suitably related. Recall that for the underlying type system we assume that *op* takes two arguments with the base types τ_{op}^1 and τ_{op}^2 and gives a result of type τ_{op} and since base types do not have any annotations we shall now *assume* that:

$$\begin{aligned} \text{If } [\] \vdash_{\text{CFA}} v_1 : \tau_{op}^1 \text{ and } [\] \vdash_{\text{CFA}} v_2 : \tau_{op}^2 \text{ then } [\] \vdash_{\text{CFA}} v : \tau_{op} \\ \text{where } v = v_1 \text{ op } v_2. \end{aligned}$$

This ensures that when given arguments of appropriate types the operator will return a value of the expected type.

The *semantic correctness* of the analysis now expresses that the annotated type foreseen by the analysis will also be the annotated type of the value obtained by evaluating the expression:

Theorem 5.9

If $[\] \vdash_{\text{CFA}} e : \hat{\tau}$, and $\vdash e \rightarrow v$ then $[\] \vdash_{\text{CFA}} v : \hat{\tau}$.

It follows that if $[\] \vdash e : \hat{\tau}_1 \xrightarrow{\varphi_0} \hat{\tau}_2$ and $\vdash e \rightarrow \text{fn}_\pi x \Rightarrow e_0$ then $\pi \in \varphi_0$; hence the analysis correctly tracks the closures that can result from a given expression. Also note that if $[\] \vdash e : \hat{\tau}_1 \xrightarrow{\emptyset} \hat{\tau}_2$ then e cannot terminate.

In preparation for the proof of the theorem we need a few technical results of the sort that are standard for type systems. The first result expresses that the type environment may be extended with information that does not influence the analysis result:

Fact 5.10 If $\hat{\Gamma}_1 \vdash_{\text{CFA}} e : \hat{\tau}$ and $\hat{\Gamma}_1(x) = \hat{\Gamma}_2(x)$ for all $x \in FV(e)$ then $\hat{\Gamma}_2 \vdash_{\text{CFA}} e : \hat{\tau}$. ■

Proof The proof is by induction on the inference tree for $\hat{\Gamma}_1 \vdash_{\text{CFA}} e : \hat{\tau}$ using that it is the rightmost occurrence of a variable in a type environment that determines its type. ■

The next result shows that we can safely substitute an expression of the correct annotated type for a variable:

Lemma 5.11 Assume $[\] \vdash_{\text{CFA}} e_0 : \hat{\tau}_0$ and $\hat{\Gamma}[x \mapsto \hat{\tau}_0] \vdash_{\text{CFA}} e : \hat{\tau}$. Then $\hat{\Gamma} \vdash_{\text{CFA}} e[x \mapsto e_0] : \hat{\tau}$. ■

Proof The proof is by structural induction on e . Most cases are straightforward so we shall only present the cases of variables and function abstractions.

The case y . We assume that

$$\hat{\Gamma}[x \mapsto \hat{\tau}_0] \vdash_{\text{CFA}} y : \hat{\tau}$$

so from the axiom $\{var\}$ of Table 5.2 we have $(\hat{\Gamma}[x \mapsto \hat{\tau}_0])(y) = \hat{\tau}$. If $x = y$ then $y[x \mapsto e_0] = e_0$ and $\hat{\tau} = \hat{\tau}_0$, and from $[\] \vdash_{\text{CFA}} e_0 : \hat{\tau}_0$ and Fact 5.10 we get the required $\hat{\Gamma} \vdash_{\text{CFA}} e_0 : \hat{\tau}$. If $x \neq y$ then $y[x \mapsto e_0] = y$ and it is easy to see that $\hat{\Gamma} \vdash_{\text{CFA}} y : \hat{\tau}$.

The case $\text{fn}_\pi y \Rightarrow e$. We assume that

$$\hat{\Gamma}[x \mapsto \hat{\tau}_0] \vdash_{\text{CFA}} \text{fn}_\pi y \Rightarrow e : \hat{\tau}$$

so, according to rule $[fn]$ of Table 5.2, it must be the case that $\hat{\tau} = \hat{\tau}_y \xrightarrow{\{\pi\} \cup \varphi} \hat{\tau}'$ and:

$$\hat{\Gamma}[x \mapsto \hat{\tau}_0][y \mapsto \hat{\tau}_y] \vdash_{\text{CFA}} e : \hat{\tau}'$$

If $x = y$ then it follows from Fact 5.10 that $\hat{\Gamma}[y \mapsto \hat{\tau}_y] \vdash_{\text{CFA}} e : \hat{\tau}'$ and since $(fn_{\pi} y \Rightarrow e)[x \mapsto e_0] = fn_{\pi} y \Rightarrow e$ the result follows. So assume that $x \neq y$; then $(fn_{\pi} y \Rightarrow e)[x \mapsto e_0] = fn_{\pi} y \Rightarrow (e[x \mapsto e_0])$. It follows from Fact 5.10 that $\hat{\Gamma}[y \mapsto \hat{\tau}_y][x \mapsto \hat{\tau}_0] \vdash_{\text{CFA}} e : \hat{\tau}'$ and then we get $\hat{\Gamma}[y \mapsto \hat{\tau}_y] \vdash_{\text{CFA}} e[x \mapsto e_0] : \hat{\tau}'$ from the induction hypothesis. Now the result follows using the rule $[fn]$. ■

We now turn to the proof of Theorem 5.9:

Proof The proof proceeds by induction on the inference tree for $\vdash e \longrightarrow v$.

The cases $[con]$ and $[fn]$ are immediate.

The case $[fun]$. We assume that

$$\vdash fun_{\pi} f x \Rightarrow e_0 \longrightarrow fn_{\pi} x \Rightarrow e_0[f \mapsto fun_{\pi} f x \Rightarrow e_0]$$

where f and x are distinct variables. Also we have

$$[] \vdash_{\text{CFA}} fun_{\pi} f x \Rightarrow e_0 : \hat{\tau}_x \xrightarrow{\{\pi\} \cup \varphi_0} \hat{\tau}_0$$

and according to Table 5.2 this is because $[f \mapsto \hat{\tau}_x \xrightarrow{\{\pi\} \cup \varphi_0} \hat{\tau}_0][x \mapsto \hat{\tau}_x] \vdash_{\text{CFA}} e_0 : \hat{\tau}_0$. Since $[f \mapsto \hat{\tau}_x \xrightarrow{\{\pi\} \cup \varphi_0} \hat{\tau}_0][x \mapsto \hat{\tau}_x]$ equals $[x \mapsto \hat{\tau}_x][f \mapsto \hat{\tau}_x \xrightarrow{\{\pi\} \cup \varphi} \hat{\tau}_0]$ (because we assumed that f and x are distinct) it follows from Lemma 5.11 that

$$[x \mapsto \hat{\tau}_x] \vdash_{\text{CFA}} e_0[f \mapsto fun_{\pi} f x \Rightarrow e_0] : \hat{\tau}_0$$

and hence $[] \vdash_{\text{CFA}} fn_{\pi} x \Rightarrow e_0[f \mapsto fun_{\pi} f x \Rightarrow e_0] : \hat{\tau}_x \xrightarrow{\{\pi\} \cup \varphi_0} \hat{\tau}_0$ which is the desired result.

The case $[app]$. We assume that

$$\vdash e_1 e_2 \longrightarrow v_0$$

because $\vdash e_1 \longrightarrow fn_{\pi} x \Rightarrow e_0$, $\vdash e_2 \longrightarrow v_2$ and $\vdash e_0[x \mapsto v_2] \longrightarrow v_0$. Also we have

$$[] \vdash_{\text{CFA}} e_1 e_2 : \hat{\tau}_0$$

and according to Table 5.2 this is because $[] \vdash_{\text{CFA}} e_1 : \hat{\tau}_2 \xrightarrow{\varphi} \hat{\tau}_0$ and $[] \vdash_{\text{CFA}} e_2 : \hat{\tau}_2$. The induction hypothesis applied to the inference tree for e_1 gives:

$$[] \vdash_{\text{CFA}} fn_{\pi} x \Rightarrow e_0 : \hat{\tau}_2 \xrightarrow{\varphi} \hat{\tau}_0$$

According to Table 5.2 this can only be the case if $\pi \in \varphi$ and $[x \mapsto \hat{\tau}_2] \vdash_{\text{CFA}} e_0 : \hat{\tau}_0$. The induction hypothesis applied to the inference tree for e_2 gives

$$[] \vdash_{\text{CFA}} v_2 : \hat{\tau}_2$$

and by Lemma 5.11 we now get $[] \vdash_{\text{CFA}} e_0[x \mapsto v_2] : \hat{\tau}_0$. The induction hypothesis applied to the inference tree for $e_0[x \mapsto v_2]$ now gives

$$[] \vdash_{\text{CFA}} v_0 : \hat{\tau}_0$$

and this is the desired result.

The case [if₁]. We assume that

$$\vdash \text{if } e_0 \text{ then } e_1 \text{ else } e_2 \longrightarrow v_1$$

because $\vdash e_0 \longrightarrow \text{true}$ and $\vdash e_1 \longrightarrow v_1$. Also we have

$$[] \vdash_{\text{CFA}} \text{if } e_0 \text{ then } e_1 \text{ else } e_2 : \hat{\tau}$$

and from Table 5.2 we see that this is because $[] \vdash_{\text{CFA}} e_0 : \text{bool}$, $[] \vdash_{\text{CFA}} e_1 : \hat{\tau}$ and $[] \vdash_{\text{CFA}} e_2 : \hat{\tau}$. The induction hypothesis gives

$$[] \vdash_{\text{CFA}} v_1 : \hat{\tau}$$

and this is the desired result.

The case [if₂] is analogous.

The case [let]. We assume that

$$\vdash \text{let } x = e_1 \text{ in } e_2 \longrightarrow v_2$$

because $\vdash e_1 \longrightarrow v_1$ and $\vdash e_2[x \mapsto v_1] \longrightarrow v_2$. Also we have

$$[] \vdash_{\text{CFA}} \text{let } x = e_1 \text{ in } e_2 : \hat{\tau}_2$$

and this is because $[] \vdash_{\text{CFA}} e_1 : \hat{\tau}_1$ and $[x \mapsto \hat{\tau}_1] \vdash_{\text{CFA}} e_2 : \hat{\tau}_2$. The induction hypothesis now gives:

$$[] \vdash_{\text{CFA}} v_1 : \hat{\tau}_1$$

From Lemma 5.11 we then get $[] \vdash_{\text{CFA}} e_2[x \mapsto v_1] : \hat{\tau}_2$ and the induction hypothesis gives

$$[] \vdash_{\text{CFA}} v_2 : \hat{\tau}_2$$

as required.

The case [op]. We assume that

$$\vdash e_1 \text{ op } e_2 \longrightarrow v$$

because $\vdash e_1 \longrightarrow v_1$, $\vdash e_2 \longrightarrow v_2$ and $v_1 \text{ op } v_2 = v$. Also we assume

$$[] \vdash_{\text{CFA}} e_1 \text{ op } e_2 : \tau_{op}$$

and this can only be because $[] \vdash_{\text{CFA}} e_1 : \tau_{op}^1$ and $[] \vdash_{\text{CFA}} e_2 : \tau_{op}^2$. The induction hypothesis gives $[] \vdash_{\text{CFA}} v_1 : \tau_{op}^1$ and $[] \vdash_{\text{CFA}} v_2 : \tau_{op}^2$ and the desired result that $[] \vdash_{\text{CFA}} v : \tau_{op}$ then follows from the stated assumptions about the relationship between the semantics and the types of the binary operators. ■

5.2.3 Existence of Solutions

In Chapter 3 we showed that the set of solutions to the Control Flow Analysis constituted a Moore family: the greatest lower bound of a set of solutions is also a solution. (Also see Exercise 2.7.) From that result it then follows that all programs can be analysed and that there is a best analysis.

Complete lattice of annotations. A similar development is possible here except that special care has to be taken concerning the nature of annotations. We would like to be able to regard the set **Ann** of annotations as a complete lattice, and this necessitates a partial ordering

$$\varphi_1 \sqsubseteq \varphi_2 \quad \text{or} \quad \varphi_1 \subseteq \varphi_2$$

that intuitively means that the set φ_1 of program points is included in the set φ_2 of program points. One way to formalise this is to say

$$\exists \varphi' : \varphi_1 \cup \varphi' = \varphi_2$$

where we rely on the axiomatisation of “equal as sets” presented previously. Another way is by means of an explicit system of rules and axioms for inferring judgements of the form $\varphi_1 \subseteq \varphi_2$; we shall dispense with these details (but see Exercise 5.3).

One way to ensure that $(\mathbf{Ann}, \sqsubseteq)$ is a complete lattice is to demand that all annotations are subsets of a given finite set; for this one might replace **Pnt** by the finite set **Pnt**_{*} of program points occurring in the expression of interest. Another possibility will be to change the syntax of annotations to allow expressing an arbitrary subset of **Pnt**. Either approach works, since all that suffices for the subsequent development is to *assume* that:

$$(\mathbf{Ann}, \sqsubseteq) \text{ is a complete lattice isomorphic to } (\mathcal{P}(\mathbf{Pnt}), \subseteq).$$

The partial ordering on **Ann** will sometimes be written as \sqsubseteq and sometimes as \subseteq .

Complete lattice of annotated types. We can now extend the partial ordering on annotations to operate on annotated types that have the same underlying type. For this let $\tau \in \mathbf{Type}$ be a type and write

$$\widehat{\mathbf{Type}}[\tau]$$

for the set of annotated types $\hat{\tau}$ with underlying type τ , i.e. such that $[\hat{\tau}] = \tau$. Next define $\hat{\tau}_1 \sqsubseteq \hat{\tau}_2$ for $\hat{\tau}_1, \hat{\tau}_2 \in \widehat{\mathbf{Type}}[\tau]$ to mean that whenever $\hat{\tau}_i$ has the annotation φ_i in a given position then $\varphi_1 \subseteq \varphi_2$. Formally this is achieved by:

$$\hat{\tau} \sqsubseteq \hat{\tau} \quad \frac{\hat{\tau}_1 \sqsubseteq \hat{\tau}'_1 \quad \varphi \subseteq \varphi' \quad \hat{\tau}_2 \sqsubseteq \hat{\tau}'_2}{\hat{\tau}_1 \xrightarrow{\varphi} \hat{\tau}_2 \sqsubseteq \hat{\tau}'_1 \xrightarrow{\varphi'} \hat{\tau}'_2}$$

As an example, $(\text{int} \xrightarrow{\varphi_1} \text{int}) \xrightarrow{\varphi_2} \text{int} \sqsubseteq (\text{int} \xrightarrow{\varphi_3} \text{int}) \xrightarrow{\varphi_4} \text{int}$ will be the case if and only if $\varphi_1 \subseteq \varphi_3$ and $\varphi_2 \subseteq \varphi_4$. (Note that this ordering is *not* contravariant unlike what will be the case for the subtyping to be introduced in Section 5.4.) Clearly the least element in $\widehat{\mathbf{Type}}[\tau]$ will have \emptyset on all function arrows occurring in τ and the greatest element will have **Pnt** on all function arrows.

It is straightforward to prove that each $(\widehat{\mathbf{Type}}[\tau], \sqsubseteq)$ is a complete lattice. In a similar way the partial ordering can be extended to operate also on type environments having the same underlying structure. Finally, suppose that $\Gamma \vdash_{\text{UL}} e : \tau$ and write

$$\text{JUDG}_{\text{CFA}}[\Gamma \vdash_{\text{UL}} e : \tau]$$

for the set of typings $\widehat{\Gamma} \vdash_{\text{CFA}} e : \widehat{\tau}$ such that $[\cdot]$ maps $\widehat{\Gamma} \vdash_{\text{CFA}} e : \widehat{\tau}$ to $\Gamma \vdash_{\text{UL}} e : \tau$ (and hence $[\widehat{\Gamma}] = \Gamma$ and $[\widehat{\tau}] = \tau$). This facilitates extending the partial ordering \sqsubseteq to operate on the set $\text{JUDG}_{\text{CFA}}[\Gamma \vdash_{\text{UL}} e : \tau]$ of typings.

Moore family result. We are now ready to state the result about Moore families:

Proposition 5.12

$\text{JUDG}_{\text{CFA}}[\Gamma \vdash_{\text{UL}} e : \tau]$ is a Moore family whenever $\Gamma \vdash_{\text{UL}} e : \tau$.

Proof We assume that $\Gamma \vdash_{\text{UL}} e : \tau$ and prove that $\text{JUDG}_{\text{CFA}}[\Gamma \vdash_{\text{UL}} e : \tau]$ is a Moore family. For this we proceed by induction on the shape of the inference tree for $\Gamma \vdash_{\text{UL}} e : \tau$. We shall only give the proof for the cases $[\text{var}]$, $[\text{fn}]$ and $[\text{app}]$; the other cases follow the same overall pattern. In all cases let

$$Y = \{(\widehat{\Gamma}^i \vdash_{\text{CFA}} e : \widehat{\tau}^i) \mid i \in I\}$$

be a subset of $\text{JUDG}_{\text{CFA}}[\Gamma \vdash_{\text{UL}} e : \tau]$. Using that $(\widehat{\mathbf{Type}}[\tau'], \sqsubseteq)$ is a complete lattice for all choices of $\tau' \in \mathbf{Type}$ we get that $\prod Y$ exists and that it is defined in a pointwise manner. (Note that if $I = \emptyset$ then $\prod Y$ is obtained from $\Gamma \vdash_{\text{UL}} e : \tau$ by placing $\prod \emptyset = \text{Pnt}$ as annotation on all function arrows.) It remains to show that $\prod Y \in \text{JUDG}_{\text{CFA}}[\Gamma \vdash_{\text{UL}} e : \tau]$.

The case $[\text{var}]$. The result follows from $(\prod_i \widehat{\Gamma}^i)(x) = \prod_i (\widehat{\Gamma}^i(x))$.

The case $[\text{fn}]$. We have $\Gamma \vdash_{\text{UL}} \text{fn}_\pi x \Rightarrow e_0 : \tau_x \rightarrow \tau_0$ because $\Gamma[x \mapsto \tau_x] \vdash_{\text{UL}} e_0 : \tau_0$. For $i \in I$ we have $\widehat{\Gamma}^i \vdash_{\text{CFA}} \text{fn}_\pi x \Rightarrow e_0 : \widehat{\tau}_x \xrightarrow{\{\pi\} \cup \varphi^i} \widehat{\tau}_0$ because of

$$\widehat{\Gamma}^i[x \mapsto \widehat{\tau}_x] \vdash_{\text{CFA}} e_0 : \widehat{\tau}_0$$

and clearly $\widehat{\Gamma}^i[x \mapsto \widehat{\tau}_x] \vdash_{\text{CFA}} e_0 : \widehat{\tau}_0$ is an element of $\text{JUDG}_{\text{CFA}}[\Gamma[x \mapsto \tau_x] \vdash_{\text{UL}} e_0 : \tau_0]$. By the induction hypothesis we get $(\prod_i \widehat{\Gamma}^i)[x \mapsto \prod_i \widehat{\tau}_x] \vdash_{\text{CFA}} e_0 : \prod_i \widehat{\tau}_0$ and hence

$$\prod_i \widehat{\Gamma}^i \vdash_{\text{CFA}} \text{fn}_\pi x \Rightarrow e_0 : \prod_i \widehat{\tau}_x \xrightarrow{\{\pi\} \cup \varphi} \prod_i \widehat{\tau}_0$$

where $\varphi = \prod_i \varphi^i$.

The case $[\text{app}]$. We now have $\Gamma \vdash_{\text{UL}} e_1 e_2 : \tau_0$ because $\Gamma \vdash_{\text{UL}} e_1 : \tau_2 \rightarrow \tau_0$ and $\Gamma \vdash_{\text{UL}} e_2 : \tau_2$. For all $i \in I$ we have $\widehat{\Gamma}^i \vdash_{\text{CFA}} e_1 e_2 : \widehat{\tau}_0$ because

$$\widehat{\Gamma}^i \vdash_{\text{CFA}} e_1 : \widehat{\tau}_2 \xrightarrow{\varphi^i} \widehat{\tau}_0 \quad \text{and} \quad \widehat{\Gamma}^i \vdash_{\text{CFA}} e_2 : \widehat{\tau}_2$$

and clearly $\widehat{\Gamma}^i \vdash_{\text{CFA}} e_1 : \widehat{\tau}_2 \xrightarrow{\varphi^i} \widehat{\tau}_0^i$ is an element of $\text{JUDG}_{\text{CFA}}[\Gamma \vdash_{\text{UL}} e_1 : \tau_2 \rightarrow \tau_0]$ and $\widehat{\Gamma}^i \vdash_{\text{CFA}} e_2 : \widehat{\tau}_2^i$ is an element of $\text{JUDG}_{\text{CFA}}[\Gamma \vdash_{\text{UL}} e_2 : \tau_2]$. By the induction hypothesis we get $\prod_i \widehat{\Gamma}^i \vdash e_1 : \prod_i \widehat{\tau}_2^i \xrightarrow{\varphi} \prod_i \widehat{\tau}_0^i$ and $\prod_i \widehat{\Gamma}^i \vdash e_2 : \prod_i \widehat{\tau}_2^i$ where $\varphi = \prod_i \varphi^i$ and hence

$$\prod_i \widehat{\Gamma}^i \vdash e_1 e_2 : \prod_i \widehat{\tau}_0^i$$

which is the desired result. ■

Example 5.13 Consider the expression e :

$$f (fn_x x \Rightarrow x+1) + f (fn_y y \Rightarrow y+2) + (fn_z z \Rightarrow z+3)(4)$$

In the underlying type system we have:

$$[f \mapsto (\text{int} \rightarrow \text{int}) \rightarrow \text{int}] \vdash e : \text{int}$$

In the Control Flow Analysis we have

$$[f \mapsto (\text{int} \xrightarrow{\varphi_1} \text{int}) \xrightarrow{\varphi_2} \text{int}] \vdash e : \text{int}$$

whenever $\{X, Y\} \subseteq \varphi_1$. The least solution therefore has $\varphi_1 = \{X, Y\}$ and $\varphi_2 = \emptyset$. This clearly tells us that f is only applied to $(fn_x x \Rightarrow x+1)$ and $(fn_y y \Rightarrow y+2)$ and not to $(fn_z z \Rightarrow z+3)$. A larger solution like $\varphi_1 = \{X, Y, Z\}$ and $\varphi_2 = \{V\}$ would not convey this information. Hence it seems sensible to ask for the least solution with respect to \sqsubseteq and the existence of this solution is guaranteed by Proposition 5.12. ■

5.3 Inference Algorithms

The main difference between an analysis expressed in the form of an inference system (as in Table 5.2) and in the form of an algorithm is that the user of the inference system is expected to have sufficient foresight to be able to *guess* the right types and annotations whereas the implementation of the algorithm will make use of a mechanism for making *tentative* guesses that are later *refined*. Let us first consider the simple case corresponding to the underlying type system of Table 5.1.

5.3.1 An Algorithm for the Underlying Type System

Augmented types. The algorithm corresponding to the type system of Table 5.1 will work on *augmented types* that allow the use of type variables to reflect that the details of a type are not fully determined yet:

$$\begin{aligned} \tau &\in \mathbf{AType} && \text{augmented types} \\ \alpha &\in \mathbf{TVar} && \text{type variables} \end{aligned}$$

We shall take:

$$\begin{aligned}\tau &::= \text{int} \mid \text{bool} \mid \tau_1 \rightarrow \tau_2 \mid \alpha \\ \alpha &::= 'a \mid 'b \mid 'c \mid \dots\end{aligned}$$

Substitutions. A *substitution* is a finite and partial mapping from type variables to augmented types, we write

$$\theta : \text{TVar} \rightarrow_{\text{fin}} \text{AType}$$

and note that the domain $\text{dom}(\theta) = \{\alpha \mid \theta \text{ is defined on } \alpha\}$ is finite. We shall allow to view a substitution as a total function from type variables to augmented types, setting $\theta \alpha = \alpha$ whenever $\alpha \notin \text{dom}(\theta)$. We shall say that a substitution θ is *defined* on α if and only if $\alpha \in \text{dom}(\theta)$.

The substitution θ is called a *ground substitution* if and only if it maps all type variables in its domain to ordinary types, i.e. if $\forall \alpha \in \text{dom}(\theta) : \theta \alpha \in \text{Type}$. The substitution θ is said to *cover* Γ , respectively τ , if and only if it is defined on all the type variables in Γ , respectively τ . Substitutions can be applied to augmented types in a pointwise manner:

$$\begin{aligned}\theta \text{ int} &= \text{int} \\ \theta \text{ bool} &= \text{bool} \\ \theta(\tau_1 \rightarrow \tau_2) &= (\theta \tau_1) \rightarrow (\theta \tau_2) \\ \theta \alpha &= \tau \text{ if } \theta \alpha = \tau\end{aligned}$$

We shall write $\theta_1 \circ \theta_2$ for the composition of θ_1 and θ_2 , i.e. $(\theta_1 \circ \theta_2)\tau = \theta_1(\theta_2 \tau)$ for all augmented types τ .

The idea. The *type reconstruction algorithm*, called \mathcal{W}_{UL} , is given two arguments: an augmented type environment Γ (mapping program variables to augmented types) and an expression e . If it succeeds in finding a typing for the expression then it will return its augmented type τ and a substitution θ telling how the type environment has to be refined in order to obtain a typing. As an example we will have

$$\mathcal{W}_{\text{UL}}([x \mapsto 'a], 1 + (x \ 2)) = (\text{int}, ['a \mapsto \text{int} \rightarrow \text{int}])$$

because during the inspection of the expression $1 + (x \ 2)$ it becomes clear that x must be a function from integers to integers if the expression is to be correctly typed. So the idea is that if

$$\mathcal{W}_{\text{UL}}(\Gamma, e) = (\tau, \theta) \text{ then } \theta_G(\theta \Gamma) \vdash_{\text{UL}} e : \theta_G \tau$$

for every ground substitution θ_G that covers $\theta \Gamma$ and τ , i.e. whenever we replace all the type variables with ordinary types in a consistent way. When

this property holds we shall say that the algorithm is *syntactically sound*. In order for the algorithm to be *syntactically complete* it is also required that all typings obtainable in the inference system can be reconstructed from the results of the algorithm. We shall discuss these properties at length in Subsection 5.3.3.

The algorithm. The algorithm \mathcal{W}_{UL} is specified in Table 5.5 and is explained below. The algorithm asks for *fresh type variables* at several places. By this is meant type variables that do not occur in the argument to \mathcal{W}_{UL} and that have not been generated as fresh variables elsewhere; this could be formalised by supplying \mathcal{W}_{UL} with yet another parameter for tracking the type variables that remain fresh but it is customary not to do so. There is a small amount of nondeterminism in \mathcal{W}_{UL} in that there may be many choices for the fresh variables; this could be made precise by assuming they are all numbered and by always supplying the candidate with the smallest number but again it is customary not to do so.

In the clause for variables we simply note that the type of c is τ_c and there is no need to adjust our assumptions Γ so we return the identity substitution id : we demand that $id\ \alpha = \alpha$ for all α and could take id to be the empty mapping. The clause for variables is similar except that now we consult the type environment Γ to determine the augmented type of x .

For function abstraction we assume that the formal parameter has type α_x for α_x being a fresh type variable – so far we have no constraints on the type of the formal parameter. Then we call \mathcal{W}_{UL} recursively on the function body to determine its type under the assumption that x has type α_x . The resulting type τ_0 and substitution θ_0 are then used to construct the overall type; in particular, θ_0 is used to replace the type variable α_x with a more refined type since the analysis of the function body may have provided additional information about the type of the formal parameter.

The clause for recursive function definition is somewhat more complicated. It starts out in a way similar to the previous clause and requires fresh type variables α_x and α_0 so that we can supply augmented types $\alpha_x \rightarrow \alpha_0$ and α_x for the occurrences of f and x , respectively, in the analysis of the function body. However, this will result in two possible types for the function body: one is the type variable α_0 (modified by the substitution θ_0 obtained by analysing the function body) and the other is the type τ_0 obtained from the analysis of the function body. These two types have to be equal according to the rule $[fn]$ of Table 5.1 and to ensure this we shall use a *unification procedure* \mathcal{U}_{UL} ; its definition will be discussed in detail below but the idea is that given two augmented types τ_1 and τ_2 , $\mathcal{U}_{UL}(\tau_1, \tau_2)$ will return a substitution θ that makes them equal i.e. such that $\theta\ \tau_1 = \theta\ \tau_2$. In the clause for recursive function definition we get that $\theta_1(\tau_0) = \theta_1(\theta_0\ \alpha_0)$ so the overall type will be $\theta_1(\theta_0\ \alpha_x) \rightarrow \theta_1\ \tau_0$. Also we record that the assumptions of Γ have to be modified by θ_0 as well as θ_1 .

$$\mathcal{W}_{UL}(\Gamma, c) = (\tau_c, id)$$

$$\mathcal{W}_{UL}(\Gamma, x) = (\Gamma(x), id)$$

$$\mathcal{W}_{UL}(\Gamma, \text{fn}_\pi x \Rightarrow e_0) = \text{let } \alpha_x \text{ be fresh}$$

$$\quad (\tau_0, \theta_0) = \mathcal{W}_{UL}(\Gamma[x \mapsto \alpha_x], e_0)$$

$$\quad \text{in } ((\theta_0 \alpha_x) \rightarrow \tau_0, \theta_0)$$

$$\mathcal{W}_{UL}(\Gamma, \text{fun}_\pi f x \Rightarrow e_0) =$$

$$\quad \text{let } \alpha_x, \alpha_0 \text{ be fresh}$$

$$\quad (\tau_0, \theta_0) = \mathcal{W}_{UL}(\Gamma[f \mapsto \alpha_x \rightarrow \alpha_0][x \mapsto \alpha_x], e_0)$$

$$\quad \theta_1 = \mathcal{U}_{UL}(\tau_0, \theta_0 \alpha_0)$$

$$\quad \text{in } (\theta_1(\theta_0 \alpha_x) \rightarrow \theta_1 \tau_0, \theta_1 \circ \theta_0)$$

$$\mathcal{W}_{UL}(\Gamma, e_1 e_2) = \text{let } (\tau_1, \theta_1) = \mathcal{W}_{UL}(\Gamma, e_1)$$

$$\quad (\tau_2, \theta_2) = \mathcal{W}_{UL}(\theta_1 \Gamma, e_2)$$

$$\quad \alpha \text{ be fresh}$$

$$\quad \theta_3 = \mathcal{U}_{UL}(\theta_2 \tau_1, \tau_2 \rightarrow \alpha)$$

$$\quad \text{in } (\theta_3 \alpha, \theta_3 \circ \theta_2 \circ \theta_1)$$

$$\mathcal{W}_{UL}(\Gamma, \text{if } e_0 \text{ then } e_1 \text{ else } e_2) = \text{let } (\tau_0, \theta_0) = \mathcal{W}_{UL}(\Gamma, e_0)$$

$$\quad (\tau_1, \theta_1) = \mathcal{W}_{UL}(\theta_0 \Gamma, e_1)$$

$$\quad (\tau_2, \theta_2) = \mathcal{W}_{UL}(\theta_1(\theta_0 \Gamma), e_2)$$

$$\quad \theta_3 = \mathcal{U}_{UL}(\theta_2(\theta_1 \tau_0), \text{bool})$$

$$\quad \theta_4 = \mathcal{U}_{UL}(\theta_3 \tau_2, \theta_3(\theta_2 \tau_1))$$

$$\quad \text{in } (\theta_4(\theta_3 \tau_2), \theta_4 \circ \theta_3 \circ \theta_2 \circ \theta_1)$$

$$\mathcal{W}_{UL}(\Gamma, \text{let } x = e_1 \text{ in } e_2) = \text{let } (\tau_1, \theta_1) = \mathcal{W}_{UL}(\Gamma, e_1)$$

$$\quad (\tau_2, \theta_2) = \mathcal{W}_{UL}((\theta_1 \Gamma)[x \mapsto \tau_1], e_2)$$

$$\quad \text{in } (\tau_2, \theta_2 \circ \theta_1)$$

$$\mathcal{W}_{UL}(\Gamma, e_1 \text{ op } e_2) = \text{let } (\tau_1, \theta_1) = \mathcal{W}_{UL}(\Gamma, e_1)$$

$$\quad (\tau_2, \theta_2) = \mathcal{W}_{UL}(\theta_1 \Gamma, e_2)$$

$$\quad \theta_3 = \mathcal{U}_{UL}(\theta_2 \tau_1, \tau_{op}^1)$$

$$\quad \theta_4 = \mathcal{U}_{UL}(\theta_3 \tau_2, \tau_{op}^2)$$

$$\quad \text{in } (\tau_{op}, \theta_4 \circ \theta_3 \circ \theta_2 \circ \theta_1)$$

Table 5.5: Algorithm \mathcal{W}_{UL} for the underlying type system.

Also the clause for function application relies on the unification procedure. In this clause we call \mathcal{W}_{UL} recursively on the operator and the operand and we use unification to ensure that the type of the operator τ_1 (modified by θ_2) is a function type with an argument type that equals the type τ_2 of the operand: $\theta_2 \tau_1$ has to have the form $\tau_2 \rightarrow \alpha$. Again we have to record that the assumptions of Γ have to be modified by all three substitutions that have

been constructed.

By now the clause for conditional is straightforward and similarly the clauses for the `let`-construct and the binary operator. (Note that the `let`-construct is *not* polymorphic so no special action is needed to determine the type of the bound variable; we shall return to the problem of polymorphism later.)

Example 5.14 Consider the expression $(\text{fn}_X x \Rightarrow x) (\text{fn}_Y y \Rightarrow y)$ of Example 5.1. The call

$$\mathcal{W}_{UL}([], (\text{fn}_X x \Rightarrow x) (\text{fn}_Y y \Rightarrow y))$$

gives rise to the call

$$\mathcal{W}_{UL}([], \text{fn}_X x \Rightarrow x)$$

which will create the fresh type variable 'a and return the pair $(a \rightarrow 'a, id)$. We now have the call

$$\mathcal{W}_{UL}([], \text{fn}_Y y \Rightarrow y)$$

that creates the fresh type variable 'b and returns $(b \rightarrow 'b, id)$. Thus we get the following call of \mathcal{U}_{UL}

$$\mathcal{U}_{UL}(a \rightarrow 'a, (b \rightarrow 'b) \rightarrow 'c)$$

where 'c is a fresh type variable. As we shall see in Example 5.15 below this gives rise to the substitution $[a \mapsto 'b \rightarrow 'b][c \mapsto 'b \rightarrow 'b]$ and the initial call of \mathcal{W}_{UL} will return $b \rightarrow 'b$ and $[a \mapsto 'b \rightarrow 'b][c \mapsto 'b \rightarrow 'b]$. ■

Analogy. The placement of substitutions in Table 5.5 may seem ad hoc at first sight. As an aid to the intuition we shall therefore offer the following analogy. Consider a society where a number of laws have been passed. One such law might stipulate that when the owner of a personal computer sells it, the owner is obliged to comply with the law and in particular provide the buyer with the original disks for all software that remains on the computer. And let us focus our attention on a particular owner preparing to sell a computer and who is determining an acceptable selling price. Then some day parliament passes a law stating that whenever original software disks are passed on from one private person to another, the previous owner has to pay a fee of ten percent of the original buying price to a government agency combating software piracy. From this day on the owner of the computer needs to reconsider all the considerations made in order to determine the acceptable selling price. In short: whenever a new law is passed all existing considerations need to be reconsidered in order to remain valid. — Coming back to Table 5.5, the *typings* determined by \mathcal{W}_{UL} correspond to the considerations of the owner, and the *substitutions* produced by \mathcal{U}_{UL} to the new laws being passed by parliament: whenever a new substitution (“law”) is constructed all existing typings (“considerations”) must be suitably modified so as to remain valid; this is exactly what is achieved by the carefully chosen use of substitutions in Table 5.5. □

$$\begin{aligned}
\mathcal{U}_{UL}(\text{int}, \text{int}) &= \text{id} \\
\mathcal{U}_{UL}(\text{bool}, \text{bool}) &= \text{id} \\
\mathcal{U}_{UL}(\tau_1 \rightarrow \tau_2, \tau'_1 \rightarrow \tau'_2) &= \text{let } \theta_1 = \mathcal{U}_{UL}(\tau_1, \tau'_1) \\
&\quad \theta_2 = \mathcal{U}_{UL}(\theta_1 \tau_2, \theta_1 \tau'_2) \\
&\quad \text{in } \theta_2 \circ \theta_1 \\
\mathcal{U}_{UL}(\tau, \alpha) &= \begin{cases} [\alpha \mapsto \tau] & \text{if } \alpha \text{ does not occur in } \tau \\ & \text{or if } \alpha \text{ equals } \tau \\ \text{fail} & \text{otherwise} \end{cases} \\
\mathcal{U}_{UL}(\alpha, \tau) &= \begin{cases} [\alpha \mapsto \tau] & \text{if } \alpha \text{ does not occur in } \tau \\ & \text{or if } \alpha \text{ equals } \tau \\ \text{fail} & \text{otherwise} \end{cases} \\
\mathcal{U}_{UL}(\tau_1, \tau_2) &= \text{fail} \quad \text{in all other cases}
\end{aligned}$$

Table 5.6: Unification of underlying types.

Unification. The algorithm \mathcal{U}_{UL} for unifying two augmented types is shown in Table 5.6. It takes two augmented types τ_1 and τ_2 as arguments and if it succeeds it returns a substitution θ such that $\theta \tau_1 = \theta \tau_2$.

In the clause for unifying the two function types $\tau_1 \rightarrow \tau_2$ and $\tau'_1 \rightarrow \tau'_2$ the desired result is obtained in two stages: θ_1 ensures that $\theta_1 \tau_1 = \theta_1 \tau'_1$ and hence $(\theta_2 \circ \theta_1) \tau_1 = (\theta_2 \circ \theta_1) \tau'_1$ and θ_2 ensures that $\theta_2(\theta_1 \tau_2) = \theta_2(\theta_1 \tau'_2)$; it follows that $\theta_2 \circ \theta_1$ succeeds in unifying the two function types. Note that the algorithm only fails at *top-level* if:

- two types with different *top-level* constructors (by which we mean `int`, `bool`, or `→`) are to be unified, or
- a type variable is to be unified with a function type containing that type variable.

Example 5.15 In Example 5.14 we had the following call of the unification procedure:

$$\mathcal{U}_{UL}('a \rightarrow 'a, ('b \rightarrow 'b) \rightarrow 'c)$$

It will first give rise to the call $\mathcal{U}_{UL}('a, 'b \rightarrow 'b)$ which returns the substitution $['a \mapsto 'b \rightarrow 'b]$. Then we will have the call $\mathcal{U}_{UL}('b \rightarrow 'b, 'c)$ and the substitution $['c \mapsto 'b \rightarrow 'b]$ is returned. Thus the overall result will be the substitution $['a \mapsto 'b \rightarrow 'b]['c \mapsto 'b \rightarrow 'b]$ as already used in Example 5.14. ■

5.3.2 An Algorithm for Control Flow Analysis

In applying the above ideas to the inference system for Control Flow Analysis presented in Table 5.2 we are going to face a difficulty. In the underlying type system two types are equal if and only if their syntactic representations are the same; we say that types constitute a *free algebra*. For annotated types, two annotated types *may* be equal even when their syntactic representations are different: $\text{int } \frac{\{\pi_1\} \cup \{\pi_2\}}{\rightarrow} \text{int}$ equals $\text{int } \frac{\{\pi_2\} \cup \{\pi_1\}}{\rightarrow} \text{int}$ because annotations and types are considered equal modulo *UCAI* as discussed in Subsection 5.1.2 – we say that annotated types constitute a *non-free algebra*.

The difficulty we encounter when transferring the development of the previous subsection to the Control Flow Analysis is that the algorithm \mathcal{W}_{UL} relies on the procedure \mathcal{U}_{UL} for unifying two types, and that this algorithm only applies to types in a free algebra. One way out of this difficulty would be to use instead an algorithm for unifying modulo *UCAI*; such algorithms exist but their properties are not so nice as those of \mathcal{U}_{UL} . Another way out of the difficulty, and the approach we shall take, is to arrange it such that a variant of \mathcal{U}_{UL} can still be used by introducing additional mechanisms for dealing with the annotations: one is the notion of *simple* types and annotations and the other is the use of *constraints*.

Simple types and annotations. The first step will be to restrict the form of the annotated types so that only annotation variables are allowed on the function arrows; later we shall combine this with a set of constraints restricting the values of the annotation variables.

A *simple type* is an augmented annotated type where the only annotations allowed on function arrows are *annotation variables* and where *type variables* are allowed in types, and a *simple annotation* is an annotation where also annotation variables are allowed:

$$\begin{aligned} \hat{\tau} &\in \mathbf{SType} && \text{simple types} \\ \alpha &\in \mathbf{TVar} && \text{type variables} \\ \beta &\in \mathbf{AVar} && \text{annotation variables} \\ \varphi &\in \mathbf{SAnn} && \text{simple annotations} \end{aligned}$$

Formally, the syntactic categories are given by:

$$\begin{aligned} \hat{\tau} &::= \text{int} \mid \text{bool} \mid \hat{\tau}_1 \xrightarrow{\beta} \hat{\tau}_2 \mid \alpha \\ \alpha &::= 'a \mid 'b \mid 'c \mid \dots \\ \beta &::= '1 \mid '2 \mid '3 \mid \dots \\ \varphi &::= \{\pi\} \mid \varphi_1 \cup \varphi_2 \mid \emptyset \mid \beta \end{aligned}$$

A *simple type environment* $\hat{\Gamma}$ then is a mapping from variables to simple

$\mathcal{U}_{\text{CFA}}(\text{int}, \text{int})$	$=$	id
$\mathcal{U}_{\text{CFA}}(\text{bool}, \text{bool})$	$=$	id
$\mathcal{U}_{\text{CFA}}(\hat{\tau}_1 \xrightarrow{\beta} \hat{\tau}_2, \hat{\tau}'_1 \xrightarrow{\beta'} \hat{\tau}'_2)$	$=$	let $\theta_0 = [\beta' \mapsto \beta]$ $\theta_1 = \mathcal{U}_{\text{CFA}}(\theta_0 \hat{\tau}_1, \theta_0 \hat{\tau}'_1)$ $\theta_2 = \mathcal{U}_{\text{CFA}}(\theta_1 (\theta_0 \hat{\tau}_2), \theta_1 (\theta_0 \hat{\tau}'_2))$ in $\theta_2 \circ \theta_1 \circ \theta_0$
$\mathcal{U}_{\text{CFA}}(\hat{\tau}, \alpha)$	$=$	$\begin{cases} [\alpha \mapsto \hat{\tau}] & \text{if } \alpha \text{ does not occur in } \hat{\tau} \\ & \text{or if } \alpha \text{ equals } \hat{\tau} \\ \text{fail} & \text{otherwise} \end{cases}$
$\mathcal{U}_{\text{CFA}}(\alpha, \hat{\tau})$	$=$	$\begin{cases} [\alpha \mapsto \hat{\tau}] & \text{if } \alpha \text{ does not occur in } \hat{\tau} \\ & \text{or if } \alpha \text{ equals } \hat{\tau} \\ \text{fail} & \text{otherwise} \end{cases}$
$\mathcal{U}_{\text{CFA}}(\hat{\tau}_1, \hat{\tau}_2)$	$=$	fail in all other cases

Table 5.7: Unification of simple types.

types.

Unification of simple types. Simple types constitute a free algebra and so we can apply the same technique as in Subsection 5.3.1 to define a function \mathcal{U}_{CFA} of the following form:

$$\mathcal{U}_{\text{CFA}}(\hat{\tau}_1, \hat{\tau}_2) = \theta$$

Here $\hat{\tau}_1$ and $\hat{\tau}_2$ are simple types and θ will be a simple substitution: A *simple substitution* is a substitution that maps type variables to simple types, and that maps annotation variables to annotation *variables* only. Thus a simple substitution applied to a simple type still gives a simple type. As for \mathcal{U}_{UL} the intention is that θ unifies $\hat{\tau}_1$ and $\hat{\tau}_2$, i.e. $\theta \tau_1 = \theta \tau_2$. In case this is not possible the outcome of $\mathcal{U}_{\text{CFA}}(\hat{\tau}_1, \hat{\tau}_2)$ will be a failure. The unification function \mathcal{U}_{CFA} is defined in Table 5.7 and is explained below.

As before id is the identity substitution and $\theta' \circ \theta''$ is the composition of two substitutions. In the clause for unifying the two function types $\hat{\tau}_1 \xrightarrow{\beta} \hat{\tau}_2$ and $\hat{\tau}'_1 \xrightarrow{\beta'} \hat{\tau}'_2$ the desired result is obtained by first ensuring that the two annotation variables are equal and then proceeding as in the unification algorithm for the underlying types. Note that $\mathcal{U}_{\text{CFA}}(\hat{\tau}_1, \hat{\tau}_2)$ will fail if and only if $\mathcal{U}_{\text{UL}}([\hat{\tau}_1], [\hat{\tau}_2])$ fails.

Example 5.16 Consider the following call of the unification procedure:

$$\mathcal{U}_{\text{CFA}}('a \xrightarrow{'1} 'a, ('b \xrightarrow{'2} 'b) \xrightarrow{'3} 'c)$$

We construct the substitution $['3 \mapsto '1]$ and then we perform the call

$$\mathcal{U}_{\text{CFA}}('a, 'b \xrightarrow{'2} 'b)$$

which returns the substitution $['a \mapsto 'b \xrightarrow{'2} 'b]$. Then we make the call

$$\mathcal{U}_{\text{CFA}}('b \xrightarrow{'2} 'b, 'c)$$

and the substitution $['c \mapsto 'b \xrightarrow{'2} 'b]$ is returned. Thus the overall result will be $['3 \mapsto '1] ['a \mapsto 'b \xrightarrow{'2} 'b] ['c \mapsto 'b \xrightarrow{'2} 'b]$. ■

The following fact, whose proof we leave to Exercise 5.7, expresses that the algorithm is correct. The first part of the result says that the algorithm is *syntactically sound*: if it succeeds then it produces the desired result. The second part says that the algorithm is *syntactically complete*: if there is some way of unifying the two simple types then the algorithm will succeed in doing so.

Fact 5.17 Let $\hat{\tau}_1$ and $\hat{\tau}_2$ be two simple types.

- If $\mathcal{U}_{\text{CFA}}(\hat{\tau}_1, \hat{\tau}_2) = \theta$ then θ is a simple substitution such that $\theta \hat{\tau}_1 = \theta \hat{\tau}_2$.
- If there exists a substitution θ'' such that $\theta'' \hat{\tau}_1 = \theta'' \hat{\tau}_2$ then there exists substitutions θ and θ' such that $\mathcal{U}_{\text{CFA}}(\hat{\tau}_1, \hat{\tau}_2) = \theta$ and $\theta'' = \theta' \circ \theta$. ■

Constraints. Annotated types can contain arbitrary annotations and simple types can only contain annotation variables. To make up for this deficiency we shall introduce constraints on the annotation variables. A *constraint* is an inclusion of the form

$$\beta \supseteq \varphi$$

where β is an annotation variable and φ is a simple annotation. A constraint set C is a finite set of such constraints.

We shall write θC for the set of inclusions obtained by applying the substitution θ to all the individual constraints of C : if $\beta \supseteq \varphi$ is in C then $\theta \beta \supseteq \theta \varphi$ is in θC . If C is a constraint set and θ is a simple substitution then also θC is a constraint set.

A *type substitution* is a substitution that is defined on type variables only and that maps them to types in **Type** (i.e. to annotated types without type

and annotation variables); we shall say that it *covers* $\hat{\Gamma}$, respectively $\hat{\tau}$, if it is defined on all type variables in $\hat{\Gamma}$, respectively $\hat{\tau}$. Similarly, an *annotation substitution* is a substitution that is defined on annotation variables only and that maps them to annotations in **Ann** (i.e. to annotations without annotation variables); it covers $\hat{\Gamma}$, respectively $\hat{\tau}$ or C , if it is defined on all annotation variables in $\hat{\Gamma}$, respectively $\hat{\tau}$ or C . An *annotation substitution* θ_A *solves* a constraint set C , written

$$\theta_A \models C$$

if and only if it covers C and for each $\beta \supseteq \varphi$ in C it is the case that $\theta_A \beta$ is a superset of, or is equal to, $\theta_A \varphi$.

A *ground substitution* is an annotation substitution on annotation variables and a type substitution on type variables. A substitution θ is a *ground validation* of $(\hat{\Gamma}, \hat{\tau}, C)$ if and only if it is a ground substitution that covers $\hat{\Gamma}, \hat{\tau}$ and C and such that $\theta \models C$ (or more precisely, $\theta_A \models C$ where θ_A is the restriction of θ to the annotation variables).

The algorithm. We are now ready to define an analogue of the type reconstruction algorithm \mathcal{W}_{UL} . It has the form

$$\mathcal{W}_{CFA}(\hat{\Gamma}, e) = (\hat{\tau}, \theta, C)$$

where we demand that $\hat{\Gamma}$ is a simple type environment (i.e. that it maps variables to simple types), and it will be the case that $\hat{\tau}$ is a simple type, θ is a simple substitution, and C is a constraint set of a very special form: it only contains constraints of the form $\beta \supseteq \{\pi\}$. Since \mathcal{W}_{CFA} will make use of \mathcal{U}_{CFA} there also is the possibility that \mathcal{W}_{CFA} fails. The algorithm \mathcal{W}_{CFA} is defined by the clauses of Table 5.8 and is explained below.

In the clauses for constants and variables we proceed as before and do not impose any constraints. In the case of function abstraction we shall additionally use a fresh annotation variable β_0 ; it will be the top-level annotation of the overall type and we shall add a constraint $\beta_0 \supseteq \{\pi\}$ requiring that it includes the label π of the function definition – this corresponds to the annotation $\{\pi\} \cup \varphi$ used in the rule $[fn]$ of Table 5.2.

The clause for recursive function definition is modified in a similar way. Here the annotation variable β will be used to annotate the function type of f and the call of \mathcal{W}_{CFA} on the body of the function may cause it to be modified to $\theta_0 \beta_0$. Next the call of the unification procedure may cause it to be further modified to $\theta_1 (\theta_0 \beta_0)$. Since both θ_0 and θ_1 are simple substitutions we know that $\theta_1 (\theta_0 \beta_0)$ is an annotation variable so the resulting type will still be simple. The recursive call of \mathcal{W}_{CFA} gives rise to a constraint set C_0 that has to be modified using the substitution θ_1 and as in the clause for ordinary function abstraction we have to add a new constraint expressing that the

$$\begin{aligned}
\mathcal{W}_{\text{CFA}}(\widehat{\Gamma}, c) &= (\tau_c, \text{id}, \emptyset) \\
\mathcal{W}_{\text{CFA}}(\widehat{\Gamma}, x) &= (\widehat{\Gamma}(x), \text{id}, \emptyset) \\
\mathcal{W}_{\text{CFA}}(\widehat{\Gamma}, \text{fn}_{\pi} x \Rightarrow e_0) &= \text{let } \alpha_x \text{ be fresh} \\
&\quad (\widehat{\tau}_0, \theta_0, C_0) = \mathcal{W}_{\text{CFA}}(\widehat{\Gamma}[x \mapsto \alpha_x], e_0) \\
&\quad \beta_0 \text{ be fresh} \\
&\quad \text{in } ((\theta_0 \alpha_x) \xrightarrow{\beta_0} \widehat{\tau}_0, \theta_0, C_0 \cup \{\beta_0 \supseteq \{\pi\}\}) \\
\mathcal{W}_{\text{CFA}}(\widehat{\Gamma}, \text{full}_{\pi} f x \Rightarrow e_0) &= \\
&\quad \text{let } \alpha_x, \alpha_0, \beta_0 \text{ be fresh} \\
&\quad (\widehat{\tau}_0, \theta_0, C_0) = \mathcal{W}_{\text{CFA}}(\widehat{\Gamma}[f \mapsto \alpha_x \xrightarrow{\beta_0} \alpha_0][x \mapsto \alpha_x], e_0) \\
&\quad \theta_1 = \mathcal{U}_{\text{CFA}}(\widehat{\tau}_0, \theta_0 \alpha_0) \\
&\quad \text{in } (\theta_1(\theta_0 \alpha_x) \xrightarrow{\theta_1(\theta_0 \beta_0)} \theta_1 \widehat{\tau}_0, \theta_1 \circ \theta_0, \\
&\quad \quad (\theta_1 C_0) \cup \{\theta_1(\theta_0 \beta_0) \supseteq \{\pi\}\}) \\
\mathcal{W}_{\text{CFA}}(\widehat{\Gamma}, e_1 e_2) &= \text{let } (\widehat{\tau}_1, \theta_1, C_1) = \mathcal{W}_{\text{CFA}}(\widehat{\Gamma}, e_1) \\
&\quad (\widehat{\tau}_2, \theta_2, C_2) = \mathcal{W}_{\text{CFA}}(\theta_1 \widehat{\Gamma}, e_2) \\
&\quad \alpha, \beta \text{ be fresh} \\
&\quad \theta_3 = \mathcal{U}_{\text{CFA}}(\theta_2 \widehat{\tau}_1, \widehat{\tau}_2 \xrightarrow{\beta} \alpha) \\
&\quad \text{in } (\theta_3 \alpha, \theta_3 \circ \theta_2 \circ \theta_1, \theta_3 (\theta_2 C_1) \cup \theta_3 C_2) \\
\mathcal{W}_{\text{CFA}}(\widehat{\Gamma}, \text{if } e_0 \text{ then } e_1 \text{ else } e_2) &= \\
&\quad \text{let } (\widehat{\tau}_0, \theta_0, C_0) = \mathcal{W}_{\text{CFA}}(\widehat{\Gamma}, e_0) \\
&\quad (\widehat{\tau}_1, \theta_1, C_1) = \mathcal{W}_{\text{CFA}}(\theta_0 \widehat{\Gamma}, e_1) \\
&\quad (\widehat{\tau}_2, \theta_2, C_2) = \mathcal{W}_{\text{CFA}}(\theta_1 (\theta_0 \widehat{\Gamma}), e_2) \\
&\quad \theta_3 = \mathcal{U}_{\text{CFA}}(\theta_2 (\theta_1 \widehat{\tau}_0), \text{bool}) \\
&\quad \theta_4 = \mathcal{U}_{\text{CFA}}(\theta_3 \widehat{\tau}_2, \theta_3 (\theta_2 \tau_1)) \\
&\quad \text{in } (\theta_4 (\theta_3 \widehat{\tau}_2), \theta_4 \circ \theta_3 \circ \theta_2 \circ \theta_1 \circ \theta_0, \\
&\quad \quad \theta_4 (\theta_3 (\theta_2 (\theta_1 C_0))) \cup \theta_4 (\theta_3 (\theta_2 C_1)) \cup \theta_4 (\theta_3 C_2)) \\
\mathcal{W}_{\text{CFA}}(\widehat{\Gamma}, \text{let } x = e_1 \text{ in } e_2) &= \\
&\quad \text{let } (\widehat{\tau}_1, \theta_1, C_1) = \mathcal{W}_{\text{CFA}}(\widehat{\Gamma}, e_1) \\
&\quad (\widehat{\tau}_2, \theta_2, C_2) = \mathcal{W}_{\text{CFA}}((\theta_1 \widehat{\Gamma})[x \mapsto \widehat{\tau}_1], e_2) \\
&\quad \text{in } (\widehat{\tau}_2, \theta_2 \circ \theta_1, (\theta_2 C_1) \cup C_2) \\
\mathcal{W}_{\text{CFA}}(\widehat{\Gamma}, e_1 \text{ op } e_2) &= \text{let } (\widehat{\tau}_1, \theta_1, C_1) = \mathcal{W}_{\text{CFA}}(\widehat{\Gamma}, e_1) \\
&\quad (\widehat{\tau}_2, \theta_2, C_2) = \mathcal{W}_{\text{CFA}}(\theta_1 \widehat{\Gamma}, e_2) \\
&\quad \theta_3 = \mathcal{U}_{\text{CFA}}(\theta_2 \widehat{\tau}_1, \tau_{\text{op}}^1) \\
&\quad \theta_4 = \mathcal{U}_{\text{CFA}}(\theta_3 \widehat{\tau}_2, \tau_{\text{op}}^2) \\
&\quad \text{in } (\tau_{\text{op}}, \theta_4 \circ \theta_3 \circ \theta_2 \circ \theta_1, \\
&\quad \quad \theta_4 (\theta_3 (\theta_2 C_1)) \cup \theta_4 (\theta_3 C_2))
\end{aligned}$$

Table 5.8: Algorithm \mathcal{W}_{CFA} for Control Flow Analysis.

annotation variable θ_1 ($\theta_0 \beta_0$) has to contain the label π of the function definition.

The clauses for the remaining constructs are fairly straightforward modifications of the similar clauses for \mathcal{W}_{UL} . Note that the substitutions are applied to the constraint sets as well as the types and type environments.

Example 5.18 Returning to the expression $(fn_X x \Rightarrow x) (fn_Y y \Rightarrow y)$ of Example 5.1 we shall now consider the call:

$$\mathcal{W}_{CFA}([], (fn_X x \Rightarrow x) (fn_Y y \Rightarrow y))$$

It gives rise to the call $\mathcal{W}_{CFA}([], fn_X x \Rightarrow x)$ which creates a fresh type variable 'a and a fresh annotation variable '1 and returns $(a \xrightarrow{'1} 'a, id, \{'1 \supseteq \{X\}\})$. Then we have the call $\mathcal{W}_{CFA}([], fn_Y y \Rightarrow y)$ that creates a fresh type variable 'b and a fresh annotation variable '2 and returns $(b \xrightarrow{'2} 'b, id, \{'2 \supseteq \{Y\}\})$. Thus we get the following call of \mathcal{U}_{CFA}

$$\mathcal{U}_{CFA}(a \xrightarrow{'1} 'a, (b \xrightarrow{'2} 'b) \xrightarrow{'3} 'c)$$

where 'c is a fresh type variable and '3 is a fresh annotation variable. This gives rise to $['3 \mapsto '1][a \mapsto 'b \xrightarrow{'2} 'b][c \mapsto 'b \xrightarrow{'2} 'b]$ as shown in Example 5.16 and the initial call of \mathcal{W}_{CFA} will return the type $b \xrightarrow{'2} 'b$, the substitution $['3 \mapsto '1][a \mapsto 'b \xrightarrow{'2} 'b][c \mapsto 'b \xrightarrow{'2} 'b]$ and the set $\{'1 \supseteq \{X\}, '2 \supseteq \{Y\}\}$ of constraints. This corresponds to the typing obtained in Example 5.4. ■

Example 5.19 Consider the program loop

```
let g = (fun_F f x => f (fn_Y y => y))
in g (fn_Z z => z)
```

of Example 5.2 and the call $\mathcal{W}_{CFA}([], \text{loop})$. This will first give rise to a call

$$\mathcal{W}_{CFA}([], \text{fun}_F f x \Rightarrow f (fn_Y y \Rightarrow y))$$

that returns the type $(a \xrightarrow{'2} 'a) \xrightarrow{'1} 'b$ and the set $\{'1 \supseteq \{F\}, '2 \supseteq \{Y\}\}$ of constraints. Then we have a call of \mathcal{W}_{CFA} on the body of the let-construct:

$$\mathcal{W}_{CFA}([g \mapsto (a \xrightarrow{'2} 'a) \xrightarrow{'1} 'b], g (fn_Z z \Rightarrow z))$$

The call of \mathcal{W}_{CFA} on $fn_Z z \Rightarrow z$ will give the type $c \xrightarrow{'3} 'c$ and the constraint set $\{'3 \supseteq \{Z\}\}$. The unification procedure will then be called with the types $(a \xrightarrow{'2} 'a) \xrightarrow{'1} 'b$ and $(c \xrightarrow{'3} 'c) \xrightarrow{'4} 'd$ and the resulting substitution is $['2 \mapsto '3, 'a \mapsto 'c, 'b \mapsto 'd]$. Thus the application will have type 'd and the constraint set will be $\{'3 \supseteq \{Z\}\}$. The initial call of \mathcal{W}_{CFA} on 'loop returns the type 'd and the constraint set $\{'1 \supseteq \{F\}, '3 \supseteq \{Y\}, '3 \supseteq \{Z\}\}$. This corresponds to the typing obtained in Example 5.5. ■

As illustrated in the above examples, when the *overall* call of $\mathcal{W}_{\text{CFA}}(\widehat{\Gamma}, e)$ gives $(\widehat{\tau}, \theta, C)$ we are only interested in the effect of θ on the type and annotation variables occurring in $\widehat{\Gamma}$.

5.3.3 Syntactic Soundness and Completeness

We are now ready to prove the correctness of the algorithm; this will take the form of a syntactic soundness result and a syntactic completeness result. Syntactic soundness is a rather straightforward result to prove. This normally also holds for Type and Effect Systems that are more complex than the Control Flow Analysis considered here, although the actual details of the Type and Effect Systems may of course require special treatment going beyond the techniques covered here.

By contrast, syntactic completeness is a somewhat harder result to prove. For a complex Type and Effect System it frequently involves establishing a result about *proof normalisation*: that the non syntax directed rules of the inference system need only be used at certain places (see Exercise 5.12). However, in the case of Control Flow Analysis the proof is going to be uncharacteristically simple because both the algorithm \mathcal{W}_{CFA} and the inference system for Control Flow Analysis are defined in syntax directed ways. Thus the present development does not indicate the breath of techniques needed for Type and Effect Systems in general and the the Concluding Remarks will contain references to the more general situation.

Syntactic soundness. The soundness result expresses that any information obtained from the algorithm is indeed correct with respect the the inference system:

Theorem 5.20

If $\mathcal{W}_{\text{CFA}}(\widehat{\Gamma}, e) = (\widehat{\tau}, \theta, C)$ and θ_G is a ground validation of $\theta \widehat{\Gamma}$, $\widehat{\tau}$ and C then $\theta_G(\theta \widehat{\Gamma}) \vdash_{\text{CFA}} e : \theta_G \widehat{\tau}$.

This theorem may be reformulated as follows: if $\mathcal{W}_{\text{CFA}}(\widehat{\Gamma}, e) = (\widehat{\tau}, \theta, C)$ and θ_T is a type substitution that covers $\theta \widehat{\Gamma}$ and $\widehat{\tau}$, and if θ_A is an annotation substitution that covers $\theta \widehat{\Gamma}$, $\widehat{\tau}$ and C and that satisfies $\theta_A \models C$, then $\theta_A(\theta_T(\theta \widehat{\Gamma})) \vdash_{\text{CFA}} e : \theta_A(\theta_T \widehat{\tau})$. To see this first note that $\theta_A \circ \theta_T = \theta_T \circ \theta_A$ is a ground substitution; next note that given a ground substitution θ_G we may obtain an annotation substitution θ_A by restricting θ_G to annotation variables and similarly we may obtain a type substitution θ_T by restricting θ_G to type variables and clearly $\theta_G = \theta_A \circ \theta_T = \theta_T \circ \theta_A$.

Proof The proof proceeds by structural induction on e (because \mathcal{W}_{CFA} is defined by structural induction on e).

The case c . We have $\mathcal{W}_{\text{CFA}}(\widehat{\Gamma}, c) = (\tau_c, id, \emptyset)$. Next let θ_G be a ground validation of $\widehat{\Gamma}$; clearly it also covers τ_c (because τ_c is a base type) and it also satisfies $\theta_G \models \emptyset$. From the axiom [con] of Table 5.2 it is immediate that

$$\theta_G(\Gamma) \vdash_{\text{CFA}} c : \tau_c$$

and since $\theta_G \tau_c = \tau_c$ this is the desired result.

The case x . We have $\mathcal{W}_{\text{CFA}}(\widehat{\Gamma}, x) = (\widehat{\Gamma}(x), id, \emptyset)$. Next let θ_G be a ground validation of $\widehat{\Gamma}$; clearly it also covers $\widehat{\Gamma}(x)$ and it satisfies $\theta_G \models \emptyset$. From the axiom [var] of Table 5.2 it is immediate that

$$\theta_G \widehat{\Gamma} \vdash_{\text{CFA}} x : \theta_G(\widehat{\Gamma}(x))$$

and this is the desired result.

The case $\text{fn}_\pi x \Rightarrow e_0$. We shall use the notation established in the clause for $\mathcal{W}_{\text{CFA}}(\widehat{\Gamma}, \text{fn}_\pi x \Rightarrow e_0)$. So let θ_G be a ground validation of $\theta_0 \widehat{\Gamma}$, $\theta_0 \alpha_x \xrightarrow{\beta_0} \widehat{\tau}_0$, and $C_0 \cup \{\beta_0 \supseteq \{\pi\}\}$. Then θ_G is a ground validation of $\theta_0(\widehat{\Gamma}[x \mapsto \alpha_x])$, $\widehat{\tau}_0$, and C_0 . Hence by the induction hypothesis we get:

$$\theta_G(\theta_0 \widehat{\Gamma})[x \mapsto \theta_G(\theta_0 \alpha_x)] \vdash_{\text{CFA}} e_0 : \theta_G \widehat{\tau}_0$$

Since $\theta_G \models C_0 \cup \{\beta_0 \supseteq \{\pi\}\}$ we have $\theta_G \beta_0 \supseteq \{\pi\}$ so we can apply the rule [fn] of Table 5.2 and get

$$\theta_G(\theta_0 \widehat{\Gamma}) \vdash_{\text{CFA}} \text{fn}_\pi x \Rightarrow e_0 : \theta_G(\theta_0 \alpha_x) \xrightarrow{\theta_G \beta_0} \theta_G \widehat{\tau}_0$$

which is the desired result.

The case $\text{fun}_\pi f x \Rightarrow e_0$. We shall use the notation already established in the clause for $\mathcal{W}_{\text{CFA}}(\widehat{\Gamma}, \text{fun}_\pi f x \Rightarrow e_0)$. So let θ_G be a ground validation of $\theta_1(\theta_0 \widehat{\Gamma})$, $\theta_1(\theta_0 \alpha_x) \xrightarrow{\theta_1(\theta_0 \beta_0)} \theta_1 \widehat{\tau}_0$, and $(\theta_1 C_0) \cup \{\theta_1(\theta_0 \beta_0) \supseteq \{\pi\}\}$. Then $\theta_G \circ \theta_1$ is a ground validation of $\theta_0 \widehat{\Gamma}$, $\theta_0 \alpha_x \xrightarrow{\theta_0 \beta_0} \widehat{\tau}_0$ and C_0 . Since $\theta_1 \widehat{\tau}_0 = \theta_1(\theta_0 \alpha_0)$ by Fact 5.17 we also have that $\theta_G \circ \theta_1$ is a ground validation of $\theta_0 \widehat{\Gamma}$, $\theta_0 \alpha_x \xrightarrow{\theta_0 \beta_0} \theta_0 \alpha_0$ and C_0 . Hence we can apply the induction hypothesis and get:

$$\theta_G(\theta_1(\theta_0 \widehat{\Gamma})[f \mapsto \alpha_x \xrightarrow{\beta_0} \alpha_0][x \mapsto \alpha_x])) \vdash_{\text{CFA}} e_0 : \theta_G(\theta_1 \widehat{\tau}_0)$$

Since $\theta_1 \widehat{\tau}_0 = \theta_1(\theta_0 \alpha_0)$ and $\theta_G \models (\theta_1 C_0) \cup \{\theta_1(\theta_0 \beta_0) \supseteq \{\pi\}\}$ we get $\theta_G(\theta_1(\theta_0 \beta_0)) \supseteq \{\pi\}$ so we can apply the rule [fun] of Table 5.2 and get

$$\theta_G(\theta_1(\theta_0 \widehat{\Gamma})) \vdash_{\text{CFA}} \text{fun}_\pi f x \Rightarrow e_0 : \theta_G(\theta_1(\theta_0 \alpha_x)) \xrightarrow{\theta_G(\theta_1(\theta_0 \beta_0))} \theta_G(\theta_1 \widehat{\tau}_0)$$

and this is the desired result.

The case $e_1 e_2$. We shall use the notation already established in the clause for $\mathcal{W}_{\text{CFA}}(\widehat{\Gamma}, e_1 e_2)$. So let θ_G be a ground validation of $\theta_3(\theta_2(\theta_1 \widehat{\Gamma}))$ and $\theta_3 \alpha$ and $\theta_3(\theta_2 C_1) \cup \theta_3 C_2$. Let θ'_G be a *ground extension* of θ_G upon $\theta_3(\theta_2 \widehat{\tau}_1)$ and $\theta_3 \widehat{\tau}_2$. Then $\theta'_G \circ \theta_3 \circ \theta_2$ is a ground validation of $\theta_1 \widehat{\Gamma}$, $\widehat{\tau}_1$ and C_1 . Hence we can apply the induction hypothesis to e_1 and get:

$$\theta'_G(\theta_3(\theta_2(\theta_1 \widehat{\Gamma}))) \vdash_{\text{CFA}} e_1 : \theta'_G(\theta_3(\theta_2 \widehat{\tau}_1))$$

Similarly $\theta'_G \circ \theta_3$ is a ground validation of $\theta_2(\theta_1 \widehat{\Gamma})$, $\widehat{\tau}_2$ and C_2 . Hence we can apply the induction hypothesis to e_2 and get:

$$\theta'_G(\theta_3(\theta_2(\theta_1 \widehat{\Gamma}))) \vdash_{\text{CFA}} e_2 : \theta'_G(\theta_3 \widehat{\tau}_2)$$

Since $\theta_3(\theta_2 \widehat{\tau}_1) = (\theta_3 \widehat{\tau}_2) \xrightarrow{\theta_3 \beta} (\theta_3 \alpha)$ follows from Fact 5.17 we can use the rule [app] of Table 5.2 and get

$$\theta'_G(\theta_3(\theta_2(\theta_1 \widehat{\Gamma}))) \vdash_{\text{CFA}} e_1 e_2 : \theta'_G(\theta_3 \alpha)$$

which is equivalent to $\theta_G(\theta_3(\theta_2(\theta_1 \widehat{\Gamma}))) \vdash_{\text{CFA}} e_1 e_2 : \theta_G(\theta_3 \alpha)$ and this is the desired result.

The cases if e_0 then e_1 else e_2 , let $x = e_1$ in e_2 and e_1 op e_2 are analogous. ■

Syntactic completeness. It is not enough to show that \mathcal{W}_{CFA} is syntactically sound: an algorithm that always fails will indeed be a syntactically sound implementation of our analysis. We shall therefore be interested in a result saying that any judgement of the Annotated Type System can in fact be obtained by the algorithm:

Theorem 5.21

Assume that $\widehat{\Gamma}$ is a simple type environment and $\theta' \widehat{\Gamma} \vdash_{\text{CFA}} e : \widehat{\tau}'$ holds for some ground substitution θ' that covers $\widehat{\Gamma}$. Then there exists $\widehat{\tau}$, θ , C and θ_G such that

- $\mathcal{W}_{\text{CFA}}(\widehat{\Gamma}, e) = (\widehat{\tau}, \theta, C)$,
- θ_G is a ground validation of $\theta \Gamma$, $\widehat{\tau}$ and C ,
- $\theta_G \circ \theta = \theta'$ except on fresh type and annotation variables (as created by $\mathcal{W}_{\text{CFA}}(\widehat{\Gamma}, e)$), and
- $\theta_G \widehat{\tau} = \widehat{\tau}'$.

Proof The proof is by induction on the shape of the inference tree; since the Annotated Type System of Table 5.2 is syntax directed this means that the proof follows the syntactic structure of e . Without loss of generality we may assume that θ' is not defined on type and annotation variables that are freshly generated in the call $\mathcal{W}_{\text{CFA}}(\widehat{\Gamma}, e)$.

The case c . We have $\theta' \widehat{\Gamma} \vdash_{\text{CFA}} c : \widehat{\tau}'$ and $\widehat{\tau}' = \tau_c$. Clearly $\mathcal{W}_{\text{CFA}}(\widehat{\Gamma}, c) = (\tau_c, id, \emptyset)$ so it suffices to set $\theta_G = \theta'$ and clearly $\theta_G \widehat{\tau}' = \widehat{\tau}'$.

The case x . We have $\theta' \widehat{\Gamma} \vdash_{\text{CFA}} x : \widehat{\tau}'$ because $\widehat{\tau}' = \theta'(\widehat{\Gamma}(x))$. Clearly $\mathcal{W}_{\text{CFA}}(\widehat{\Gamma}, x) = (\widehat{\Gamma}(x), id, \emptyset)$ so it suffices to set $\theta_G = \theta'$.

The case $\text{fn}_\pi x \Rightarrow e_0$. We have

$$\theta' \widehat{\Gamma} \vdash_{\text{CFA}} \text{fn}_\pi x \Rightarrow e_0 : \widehat{\tau}'_x \xrightarrow{\{\pi\} \cup \varphi'} \widehat{\tau}'_0$$

where φ' is an annotation (i.e. it does not contain any annotation variables) and from Table 5.2 we get that also:

$$(\theta' \widehat{\Gamma})[x \mapsto \widehat{\tau}_x] \vdash_{\text{CFA}} e_0 : \widehat{\tau}_0$$

Let now α_x be a fresh type variable and note that $\alpha_x \notin \text{dom}(\theta')$. Define θ'' by

$$\theta'' \zeta = \begin{cases} \widehat{\tau}_x & \text{if } \zeta = \alpha_x \\ \theta' \zeta & \text{otherwise} \end{cases}$$

where ζ can either be a type variable or an annotation variable. Then we also have:

$$\theta''(\widehat{\Gamma}[x \mapsto \alpha_x]) \vdash_{\text{CFA}} e_0 : \widehat{\tau}_0$$

By the induction hypothesis there exists $\widehat{\tau}_0$, θ_0 , C_0 and θ'_G such that:

$$\begin{aligned} \mathcal{W}_{\text{CFA}}(\widehat{\Gamma}[x \mapsto \alpha_x], e_0) &= (\widehat{\tau}_0, \theta_0, C_0), \\ \theta'_G &\text{ is a ground validation of } (\theta_0 \widehat{\Gamma})[x \mapsto \theta_0(\alpha_x)], \widehat{\tau}_0, \text{ and } C_0, \\ \theta'_G \circ \theta_0 &= \theta'' \text{ except on fresh type and annotation variables} \\ &\text{ created by } \mathcal{W}_{\text{CFA}}(\widehat{\Gamma}[x \mapsto \alpha_x], e_0) \\ \theta'_G \widehat{\tau}_0 &= \widehat{\tau}_0 \end{aligned}$$

Next let β_0 be a fresh annotation variable and define:

$$\theta_G \zeta = \begin{cases} \{\pi\} \cup \varphi' & \text{if } \zeta = \beta_0 \\ \theta'_G \zeta & \text{otherwise} \end{cases}$$

Then we get:

$$\begin{aligned} \theta_G &\text{ is a ground validation of } \theta_0 \widehat{\Gamma}, (\theta_0 \alpha_x) \xrightarrow{\beta_0} \widehat{\tau}_0 \text{ and } C_0 \cup \{\beta_0 \supseteq \{\pi\}\}, \\ \theta_G \circ \theta_0 &= \theta' \text{ except on fresh type and annotation variables} \\ &\text{ created by } \mathcal{W}_{\text{CFA}}(\widehat{\Gamma}, \text{fn}_\pi x \Rightarrow e_0), \\ \theta_G(\theta_0 \alpha_x \xrightarrow{\beta_0} \widehat{\tau}_0) &= \widehat{\tau}_x \xrightarrow{\{\pi\} \cup \varphi'} \widehat{\tau}_0 \end{aligned}$$

This is the desired result.

The case $\text{fun}_\pi f x \Rightarrow e_0$. We have

$$\theta' \widehat{\Gamma} \vdash_{\text{CFA}} \text{fun}_\pi f x \Rightarrow e_0 : \widehat{\tau}_x \xrightarrow{\{\pi\} \cup \varphi'} \widehat{\tau}_0$$

where φ' does not contain annotation variables and according to Table 5.2 this is because:

$$(\theta' \widehat{\Gamma})[f \mapsto \widehat{\tau}_x \xrightarrow{\{\pi\} \cup \varphi'} \widehat{\tau}_0][x \mapsto \widehat{\tau}_x] \vdash_{\text{CFA}} e_0 : \widehat{\tau}_0$$

Let α_x , α_0 and β_0 be fresh type and annotation variables and note that they are not in $\text{dom}(\theta')$. Define θ'' by:

$$\theta'' \zeta = \begin{cases} \widehat{\tau}_x & \text{if } \zeta = \alpha_x \\ \{\pi\} \cup \varphi' & \text{if } \zeta = \beta_0 \\ \widehat{\tau}_0 & \text{if } \zeta = \alpha_0 \\ \theta' \zeta & \text{otherwise} \end{cases}$$

Then we also have:

$$\theta''(\widehat{\Gamma}[f \mapsto \alpha_x \xrightarrow{\beta_0} \alpha_0][x \mapsto \alpha_x]) \vdash_{\text{CFA}} e_0 : \widehat{\tau}_0$$

By the induction hypothesis there exists $\widehat{\tau}_0$, θ_0 , C_0 and θ'_G such that:

$$\begin{aligned} \mathcal{W}_{\text{CFA}}(\widehat{\Gamma}[f \mapsto \alpha_x \xrightarrow{\beta_0} \alpha_0][x \mapsto \alpha_x], e_0) &= (\widehat{\tau}_0, \theta_0, C_0) \\ \theta'_G &\text{ is a ground validation of } \theta_0(\widehat{\Gamma}[f \mapsto \alpha_x \xrightarrow{\beta_0} \alpha_0][x \mapsto \alpha_x]), \widehat{\tau}_0, \text{ and } C_0 \\ \theta'_G \circ \theta_0 &= \theta'' \text{ except on fresh type and annotation variables} \\ &\text{ created by } \mathcal{W}_{\text{CFA}}(\dots, e_0) \\ \theta'_G \widehat{\tau}_0 &= \widehat{\tau}_0 \end{aligned}$$

Since $\theta'_G(\theta_0 \alpha_0) = \theta'' \alpha_0 = \widehat{\tau}_0 = \theta'_G \widehat{\tau}_0$ it follows from Fact 5.17 that there exists θ_1 and θ_G such that $\mathcal{U}_{\text{CFA}}(\theta_0 \alpha_0, \widehat{\tau}_0) = \theta_1$ and $\theta'_G = \theta_G \circ \theta_1$. Hence

$$\begin{aligned} \theta_G &\text{ is a ground validation of } \theta_1(\theta_0 \widehat{\Gamma}), \theta_1(\theta_0 \alpha_x) \xrightarrow{\theta_1(\theta_0 \beta_0)} \theta_1 \widehat{\tau}_0, \\ &\text{ and } (\theta_1 C_0) \cup \{\theta_1(\theta_0 \beta_0) \supseteq \{\pi\}\} \\ \theta_G \circ \theta_1 \circ \theta_0 &= \theta' \text{ except on fresh type and annotation variables} \\ &\text{ created by } \mathcal{W}_{\text{CFA}}(\dots, \text{fun}_\pi f x \Rightarrow e_0) \\ \theta_G(\theta_1(\theta_0 \alpha_x)) &\xrightarrow{\theta_1(\theta_0 \beta_0)} \theta_1(\theta_0 \alpha_0) = \widehat{\tau}_x \xrightarrow{\{\pi\} \cup \varphi'} \widehat{\tau}_0 \end{aligned}$$

and this is the desired result.

The case $e_1 e_2$. We have

$$\theta' \widehat{\Gamma} \vdash_{\text{CFA}} e_1 e_2 : \widehat{\tau}_0$$

and according to Table 5.2 this is because:

$$\begin{aligned} \theta' \widehat{\Gamma} \vdash_{\text{CFA}} e_1 : \widehat{\tau}_2 &\xrightarrow{\varphi'} \widehat{\tau}_0 \\ \theta' \widehat{\Gamma} \vdash_{\text{CFA}} e_2 : \widehat{\tau}_2 & \end{aligned}$$

By the induction hypothesis applied to e_1 there exists $\widehat{\tau}_1$, θ_1 , C_1 and θ_G^1 such that:

$$\begin{aligned} \mathcal{W}_{\text{CFA}}(\widehat{\Gamma}, e_1) &= (\widehat{\tau}_1, \theta_1, C_1) \\ \theta_G^1 &\text{ is a ground validation of } \theta_1 \widehat{\Gamma}, \widehat{\tau}_1, \text{ and } C_1 \\ \theta_G^1 \circ \theta_1 &= \theta' \text{ except on fresh type and annotation variables} \\ &\text{ created by } \mathcal{W}_{\text{CFA}}(\theta' \widehat{\Gamma}, e_1) \\ \theta_G^1 \widehat{\tau}_1 &= \widehat{\tau}_2 \xrightarrow{\varphi'} \widehat{\tau}_0 \end{aligned}$$

Then we have

$$\theta_G^1(\theta_1 \widehat{\Gamma}) \vdash_{\text{CFA}} e_2 : \widehat{\tau}_2$$

so by the induction hypothesis applied to e_2 there exists $\widehat{\tau}_2$, θ_2 , C_2 and θ_G^2 such that:

$$\begin{aligned} \mathcal{W}_{\text{CFA}}(\theta_1 \widehat{\Gamma}, e_2) &= (\widehat{\tau}_2, \theta_2, C_2) \\ \theta_G^2 &\text{ is a ground validation of } \theta_2(\theta_1 \widehat{\Gamma}), \widehat{\tau}_2 \text{ and } C_2 \\ \theta_G^2 \circ \theta_2 &= \theta_G^1 \text{ except on fresh type and annotation variables} \\ &\text{ created by } \mathcal{W}_{\text{CFA}}(\dots, e_2) \\ \theta_G^2 \widehat{\tau}_2 &= \widehat{\tau}_2 \end{aligned}$$

It then follows that:

$$\begin{aligned} \theta_G^2 &\text{ is a ground validation of } \theta_2(\theta_1 \widehat{\Gamma}), \theta_2 \widehat{\tau}_1, \theta_2 C_1, \widehat{\tau}_2, \text{ and } C_2 \\ \theta_G^2 \circ \theta_2 \circ \theta_1 &= \theta' \text{ except on fresh type and annotation variables} \\ &\quad \text{created by } \mathcal{W}_{\text{CFA}}(\dots, e_1) \text{ and } \mathcal{W}_{\text{CFA}}(\dots, e_2) \\ \theta_G^2(\theta_2 \widehat{\tau}_1) &= \widehat{\tau}_2 \xrightarrow{\varphi'} \widehat{\tau}_0 \\ \theta_G^2 \widehat{\tau}_2 &= \widehat{\tau}_2 \end{aligned}$$

Next let α and β be fresh and define:

$$\theta_G^3 \zeta = \begin{cases} \widehat{\tau}_0 & \text{if } \zeta = \alpha \\ \varphi' & \text{if } \zeta = \beta \\ \theta_G^2 \zeta & \text{otherwise} \end{cases}$$

Since $\theta_G^3(\theta_2 \widehat{\tau}_1) = \theta_G^2(\theta_2 \widehat{\tau}_1) = \widehat{\tau}_2 \xrightarrow{\varphi'} \widehat{\tau}_0 = \theta_G^2 \widehat{\tau}_2 \xrightarrow{\varphi'} \widehat{\tau}_0 = \theta_G^3(\widehat{\tau}_2 \xrightarrow{\beta} \alpha)$ it follows from Fact 5.17 that there exists θ_3 and θ_G such that $\mathcal{U}_{\text{CFA}}(\theta_2 \widehat{\tau}_1, \widehat{\tau}_2 \xrightarrow{\beta} \alpha) = \theta_3$ and $\theta_G^3 = \theta_G \circ \theta_3$. It follows that

$$\begin{aligned} \theta_G &\text{ is a ground validation of } \theta_3(\theta_2(\theta_1 \widehat{\Gamma})), \theta_3 \alpha, \text{ and } \theta_3(\theta_2 C_1) \cup \theta_3 C_2 \\ \theta_G \circ \theta_3 \circ \theta_2 \circ \theta_1 &= \theta' \text{ except on fresh type and annotation variables} \\ &\quad \text{created by } \mathcal{W}_{\text{CFA}}(\dots, e_1, e_2) \\ \theta_G(\theta_3 \alpha) &= \theta_G^3 \alpha = \widehat{\tau}_0 \end{aligned}$$

and this is the desired result.

The cases if e_0 then e_1 else e_2 , let $x = e_1$ in e_2 and e_1 op e_2 are analogous. ■

5.3.4 Existence of Solutions

Since \mathcal{W}_{CFA} generates a set of constraints the statement of syntactic soundness (Theorem 5.20) is a little weaker than usual: if the constraints cannot be solved then we cannot use the soundness result to guarantee that the result produced by \mathcal{W}_{CFA} can be inferred in the inference system.

This suggests showing that the constraints always have solutions; in line with previous developments in this book we shall prove a stronger result. For this let $AV(C)$ be the set of annotation variables in C .

Lemma 5.22 If $\mathcal{W}_{\text{CFA}}(\widehat{\Gamma}, e) = (\widehat{\tau}, \theta, C)$ and X is a finite set of annotation variables such that $X \supseteq AV(C)$, then

$$\{\theta_A \mid \theta_A \models C \wedge \text{dom}(\theta_A) = X \wedge \theta_A \text{ is an annotation substitution}\}$$

is a Moore family. ■

Proof Let C be a finite set of constraints of the form $\beta \supseteq \varphi$ where $\varphi \in \text{SAnn}$ and such that $X \supseteq AV(C)$ where X is a finite set of annotation variables; it will not be of importance that C is generated by \mathcal{W}_{CFA} . Let Y be a possibly empty

subset of the set displayed in the lemma. Each element of Y will be an annotation substitution with domain X and that satisfies C . By setting

$$\theta \beta = \bigcap_{\theta_A \in Y} (\theta_A \beta) \quad \text{for } \beta \in X$$

we define an annotation substitution θ with $\text{dom}(\theta) = X$. For each $\beta \supseteq \varphi$ in C and θ_A in Y we have

$$\theta_A \beta \supseteq \theta_A \varphi \supseteq \theta \varphi$$

(since φ is monotone in any free annotation variables) and hence:

$$\theta \beta = \bigcap_{\theta_A \in Y} \theta_A \beta \supseteq \bigcap_{\theta_A \in Y} \theta \varphi \supseteq \theta \varphi$$

This establishes the result. \blacksquare

Consider now the call $\mathcal{W}_{\text{CFA}}(\widehat{\Gamma}, e) = (\widehat{\tau}, \theta, C)$ and the problem of finding a ground substitution θ_G that covers $\theta \widehat{\Gamma}$, $\widehat{\tau}$ and C . After the statement of Theorem 5.20 we made it clear that θ_G always can be written as $\theta_A \circ \theta_T$ for an annotation substitution θ_A covering $\theta \widehat{\Gamma}$, $\widehat{\tau}$ and C and a type substitution θ_T covering $\theta \widehat{\Gamma}$ and $\widehat{\tau}$. The choice of the type substitution θ_T must be performed by the user of \mathcal{W}_{CFA} ; one possibility is to let $\theta_T \alpha = \text{int}$ for all type variables α in $\theta \widehat{\Gamma}$ and $\widehat{\tau}$. The existence of an annotation substitution θ_A now follows from Lemma 5.22.

Corollary 5.23 If $\mathcal{W}_{\text{CFA}}(\widehat{\Gamma}, e) = (\widehat{\tau}, \theta, C)$ then there exists a ground validation θ_G of $\theta \widehat{\Gamma}$, $\widehat{\tau}$ and C . \blacksquare

Proof Let θ_T be given by $\theta_T \alpha = \text{int}$ for all type variables α in $\theta \widehat{\Gamma}$ and $\widehat{\tau}$; clearly θ_T covers $\theta \widehat{\Gamma}$ and $\widehat{\tau}$. Next let X be the set of annotation variables in $\theta \widehat{\Gamma}$, $\widehat{\tau}$ and C ; then Lemma 5.22 guarantees the existence of an annotation substitution θ_A that covers $\theta \widehat{\Gamma}$, $\widehat{\tau}$ and C and such that $\theta_A \models C$. Taking $\theta_G = \theta_A \circ \theta_T$ we obtain a ground validation of $\theta \widehat{\Gamma}$, $\widehat{\tau}$ and C . \blacksquare

The result obtained by choosing a type substitution is only unique in case there are no type variables present in $\theta \widehat{\Gamma}$ and $\widehat{\tau}$. If there is at least one type variable present then the result of using θ_T displayed above will differ from the result of using θ'_T given by $\theta'_T \alpha = \text{bool}$ for all type variables α in $\theta \widehat{\Gamma}$ and $\widehat{\tau}$. In general a type substitution θ''_T may have $\theta''_T \alpha \in \widehat{\text{Type}}[\tau_\alpha]$ for an arbitrary underlying type τ_α . However, in case $\widehat{\text{Type}}[\tau_\alpha]$ has more than one element one is likely to prefer the least element since it has empty annotations on all function arrows (but see Exercise 5.8).

In a similar way one is likely to prefer the least annotation substitution guaranteed by Lemma 5.22. Keeping in mind that all constraints in C have the form $\beta \supseteq \{\pi\}$ for $\beta \in \text{AVar}$ and $\pi \in \text{Pnt}$, we simply set:

$$\theta_A \beta = \begin{cases} \{\pi \mid \beta \supseteq \{\pi\} \text{ is in } C\} & \text{if } \beta \in \text{AV}(C) \\ \text{undefined} & \text{otherwise} \end{cases}$$

It is immediate that $\text{dom}(\theta_A) = \text{AV}(C)$ and that $\theta_A \models C$. If also $\text{dom}(\theta) = \text{AV}(C)$ and $\theta \models C$ then it is immediate that $\forall \beta : \theta \beta \supseteq \theta_A \beta$ which we may write as $\theta \supseteq \theta_A$. This shows that θ_A as constructed above is indeed the least element of the Moore family displayed in Lemma 5.22.

5.4 Effects

The Type and Effect System for Control Flow Analysis is fairly simple: it is a syntax directed system using a form of subeffecting and the annotations are just sets. Much more powerful Type and Effect Systems can be constructed by allowing subtyping, let-polymorphism or polymorphic recursion; the resulting analyses will be more powerful and, not surprisingly, the techniques required for the implementation will be more demanding.

Subtyping and the various notions of polymorphism can be combined but for the sake of simplicity we shall present them one at a time. We shall first present a Side Effect Analysis for an extension of the FUN language with assignments; it will use subeffecting and subtyping. Then we shall present an Exception Analysis for an extension of FUN with exceptions; it will use subeffecting, subtyping as well as polymorphism. Finally, we shall present a Region Analysis for the FUN language; it will be based on polymorphic recursion.

5.4.1 Side Effect Analysis

Syntax. Let us consider an extension of the language FUN (Section 5.1) with imperative constructs for creating *reference variables* and for accessing and updating their values:

$$e ::= \dots \mid \text{new}_\pi r := e_1 \text{ in } e_2 \mid !r \mid r := e_0 \mid e_1 ; e_2$$

The idea is that $\text{new}_\pi r := e_1 \text{ in } e_2$ creates a new reference variable called r and initialises it to the value of e_1 ; the scope of the reference variable is e_2 but we shall want the creation of the reference variable to be visible also outside its scope so as to be able to determine whether or not functions may need to allocate additional memory. The value of the reference variable r can be obtained by writing $!r$ and it may be set to a new value by the assignment $r := e_0$. The sequencing construct $e_1 ; e_2$ first evaluates e_1 (for its side effects) and then e_2 .

Example 5.24 The following program computes the Fibonacci number of a positive number x and leaves the result in the reference variable r :

```

newR r:=0
in let fib = funf f z => if z<3 then r:=!r+1
                                else f(z-1); f(z-2)
in fib x; !r

```

The program creates a new reference variable r and initialises it to 0, then it defines the function fib , applies it to the value of x and returns the value of r . The statement $r:=!r+1$ in the body of the recursive function will increment the value of r each time it is executed and each call of $\text{fib } x$ will increase the value of r with the Fibonacci number of x . ■

The aim of the *Side Effect Analysis* is to record:

For each subexpression which locations have been created, accessed and assigned.

So for the function fib in Example 5.24 the analysis will record that it accesses and assigns the reference variable created at program point R .

Semantics. Before presenting the analysis let us briefly sketch the semantics of the language. To distinguish between the various incarnations of the new -construct we shall introduce *locations* (or references) and, as for the imperative languages of Chapter 2, the configurations will then contain a *store* component mapping locations to their values:

$$\varsigma \in \text{Store} = \text{Loc} \rightarrow_{\text{fin}} \text{Val}$$

The values of Val include the constants c , the (closed) function abstractions of the form $\text{fn}_\pi x \Rightarrow e$ and the locations $\xi \in \text{Loc}$. The semantic clauses of Table 5.4 are now modified to trace the store as for example in the following clause for the let -construct:

$$\frac{\vdash \langle e_1, \varsigma_1 \rangle \longrightarrow \langle v_1, \varsigma_2 \rangle \quad \vdash \langle e_2[x \mapsto v_1], \varsigma_2 \rangle \longrightarrow \langle v_2, \varsigma_3 \rangle}{\vdash \langle \text{let } x = e_1 \text{ in } e_2, \varsigma_1 \rangle \longrightarrow \langle v_2, \varsigma_3 \rangle}$$

For the new constructs we then have the following axioms and rules (explained below):

$$\frac{\vdash \langle e_1, \varsigma_1 \rangle \longrightarrow \langle v_1, \varsigma_2 \rangle \quad \vdash \langle e_2[r \mapsto \xi], \varsigma_2[\xi \mapsto v_1] \rangle \longrightarrow \langle v_2, \varsigma_3 \rangle}{\vdash \langle \text{new}_\pi r := e_1 \text{ in } e_2, \varsigma_1 \rangle \longrightarrow \langle v_2, \varsigma_3 \rangle}$$

where ξ does not occur in the domain of ς_2

$$\vdash \langle !\xi, \varsigma \rangle \longrightarrow \langle \varsigma(\xi), \varsigma \rangle$$

$$\frac{\vdash \langle e, \varsigma_1 \rangle \longrightarrow \langle v, \varsigma_2 \rangle}{\vdash \langle \xi := e, \varsigma_1 \rangle \longrightarrow \langle v, \varsigma_2[\xi \mapsto v] \rangle}$$

$$\frac{\vdash \langle e_1, \varsigma_1 \rangle \longrightarrow \langle v_1, \varsigma_2 \rangle \quad \vdash \langle e_2, \varsigma_2 \rangle \longrightarrow \langle v_2, \varsigma_3 \rangle}{\vdash \langle e_1; e_2, \varsigma_1 \rangle \longrightarrow \langle v_2, \varsigma_3 \rangle}$$

In the rule for `new` we evaluate e_1 , create a new location ξ which is initialised to the value of e_1 , and then we syntactically replace all occurrences of r with that location so that all subsequent references to r will be to ξ . We exploit this when defining the semantics for the constructs for accessing and updating the reference. Note that the value returned by the assignment will be the value being assigned.

Annotated types. In the Side Effect Analysis a location will be represented by the program point where it could be created. We shall therefore define the *annotations* (of effects) $\varphi \in \mathbf{Ann}_{SE}$ by:

$$\varphi ::= \{!\pi\} \mid \{\pi:=\} \mid \{\mathbf{new}\pi\} \mid \varphi_1 \cup \varphi_2 \mid \emptyset$$

The annotation $!\pi$ means that the value of a location created at π is accessed, $\pi:=$ means that a location created at π is assigned, and $\mathbf{new}\pi$ that a new location has been created at π . As for the Control Flow Analysis we shall consider annotations equal modulo UCAI.

The *annotated types* $\hat{\tau} \in \mathbf{Type}_{SE}$ are now given by:

$$\hat{\tau} ::= \text{int} \mid \text{bool} \mid \hat{\tau}_1 \xrightarrow{\varphi} \hat{\tau}_2 \mid \text{ref}_{\pi} \hat{\tau}$$

Here $\text{ref}_{\pi} \hat{\tau}$ is the type of a location created at the program point π that will contain values of the annotated type $\hat{\tau}$. As before the type environment $\hat{\Gamma}$ will map variables to annotated types.

Example 5.25 Consider the Fibonacci program

```
newR r:=0
in let fib = funF f z => if z<3 then r:=!r+1
                               else f(z-1); f(z-2)
   in fib x; !r
```

of Example 5.24. The variable `r` has the annotated type $\text{ref}_R \text{int}$ and the variable `fib` has type $\text{int} \xrightarrow{\{!R,R:=\}} \text{int}$ since it maps integers to integers and whenever executed it may, as a side effect, access and update a reference created at the program point `R`. ■

Typing judgements. The typing judgements for the Side Effect Analysis will be of the form:

$$\hat{\Gamma} \vdash_{SE} e : \hat{\tau} \ \& \ \varphi$$

The idea is that under the assumption $\hat{\Gamma}$, the expression e will evaluate into a value with the annotated type $\hat{\tau}$ and that during computation the side effects expressed by φ might take place. The analysis is specified by the axioms and rules of Table 5.9.

$[con]$	$\widehat{\Gamma} \vdash_{SE} c : \tau_c \ \& \ \emptyset$
$[var]$	$\widehat{\Gamma} \vdash_{SE} x : \widehat{\tau} \ \& \ \emptyset \quad \text{if } \widehat{\Gamma}(x) = \widehat{\tau}$
$[fn]$	$\frac{\widehat{\Gamma}[x \mapsto \widehat{\tau}_x] \vdash_{SE} e_0 : \widehat{\tau}_0 \ \& \ \varphi_0}{\widehat{\Gamma} \vdash_{SE} fn_{\pi} x \Rightarrow e_0 : \widehat{\tau}_x \xrightarrow{\varphi_0} \widehat{\tau}_0 \ \& \ \emptyset}$
$[fun]$	$\frac{\widehat{\Gamma}[f \mapsto \widehat{\tau}_x \xrightarrow{\varphi_0} \widehat{\tau}_0][x \mapsto \widehat{\tau}_x] \vdash_{SE} e_0 : \widehat{\tau}_0 \ \& \ \varphi_0}{\widehat{\Gamma} \vdash_{SE} fun_{\pi} f x \Rightarrow e_0 : \widehat{\tau}_x \xrightarrow{\varphi_0} \widehat{\tau}_0 \ \& \ \emptyset}$
$[app]$	$\frac{\widehat{\Gamma} \vdash_{SE} e_1 : \widehat{\tau}_2 \xrightarrow{\varphi_0} \widehat{\tau}_0 \ \& \ \varphi_1 \quad \widehat{\Gamma} \vdash_{SE} e_2 : \widehat{\tau}_2 \ \& \ \varphi_2}{\widehat{\Gamma} \vdash_{SE} e_1 e_2 : \widehat{\tau}_0 \ \& \ \varphi_1 \cup \varphi_2 \cup \varphi_0}$
$[if]$	$\frac{\widehat{\Gamma} \vdash_{SE} e_0 : bool \ \& \ \varphi_0 \quad \widehat{\Gamma} \vdash_{SE} e_1 : \widehat{\tau} \ \& \ \varphi_1 \quad \widehat{\Gamma} \vdash_{SE} e_2 : \widehat{\tau} \ \& \ \varphi_2}{\widehat{\Gamma} \vdash_{SE} \text{if } e_0 \text{ then } e_1 \text{ else } e_2 : \widehat{\tau} \ \& \ \varphi_0 \cup \varphi_1 \cup \varphi_2}$
$[let]$	$\frac{\widehat{\Gamma} \vdash_{SE} e_1 : \widehat{\tau}_1 \ \& \ \varphi_1 \quad \widehat{\Gamma}[x \mapsto \widehat{\tau}_1] \vdash_{SE} e_2 : \widehat{\tau}_2 \ \& \ \varphi_2}{\widehat{\Gamma} \vdash_{SE} \text{let } x = e_1 \text{ in } e_2 : \widehat{\tau}_2 \ \& \ \varphi_1 \cup \varphi_2}$
$[op]$	$\frac{\widehat{\Gamma} \vdash_{SE} e_1 : \tau_{op}^1 \ \& \ \varphi_1 \quad \widehat{\Gamma} \vdash_{SE} e_2 : \tau_{op}^2 \ \& \ \varphi_2}{\widehat{\Gamma} \vdash_{SE} e_1 \ op \ e_2 : \tau_{op} \ \& \ \varphi_1 \cup \varphi_2}$
$[deref]$	$\widehat{\Gamma} \vdash_{SE} !x : \widehat{\tau} \ \& \ \{\pi\} \quad \text{if } \widehat{\Gamma}(x) = \text{ref}_{\pi} \widehat{\tau}$
$[ass]$	$\frac{\widehat{\Gamma} \vdash_{SE} e : \widehat{\tau} \ \& \ \varphi}{\widehat{\Gamma} \vdash_{SE} x := e : \widehat{\tau} \ \& \ \varphi \cup \{\pi :=\}} \quad \text{if } \widehat{\Gamma}(x) = \text{ref}_{\pi} \widehat{\tau}$
$[new]$	$\frac{\widehat{\Gamma} \vdash_{SE} e_1 : \widehat{\tau}_1 \ \& \ \varphi_1 \quad \widehat{\Gamma}[x \mapsto \text{ref}_{\pi} \widehat{\tau}_1] \vdash_{SE} e_2 : \widehat{\tau}_2 \ \& \ \varphi_2}{\widehat{\Gamma} \vdash_{SE} \text{new}_{\pi} x := e_1 \text{ in } e_2 : \widehat{\tau}_2 \ \& \ (\varphi_1 \cup \varphi_2 \cup \{\text{new}_{\pi}\})}$
$[seq]$	$\frac{\widehat{\Gamma} \vdash_{SE} e_1 : \widehat{\tau}_1 \ \& \ \varphi_1 \quad \widehat{\Gamma} \vdash_{SE} e_2 : \widehat{\tau}_2 \ \& \ \varphi_2}{\widehat{\Gamma} \vdash_{SE} e_1 ; e_2 : \widehat{\tau}_2 \ \& \ \varphi_1 \cup \varphi_2}$
$[sub]$	$\frac{\widehat{\Gamma} \vdash_{SE} e : \widehat{\tau} \ \& \ \varphi}{\widehat{\Gamma} \vdash_{SE} e : \widehat{\tau}' \ \& \ \varphi'} \quad \text{if } \widehat{\tau} \leq \widehat{\tau}' \text{ and } \varphi \subseteq \varphi'$

Table 5.9: Side Effect Analysis.

In the clauses *[con]* and *[var]* we record that there are no side effects so we use \emptyset for the overall effect. The premise of the clause *[fn]* gives the effect of the body of the function and we use that to annotate the arrow of the function type whereas we use \emptyset as the overall effect of the function definition itself: no side effects can be observed by simply defining the function. A similar explanation holds for the recursive function definition, the only difference being that the assumption about *f* in the premise has to use the *same* effect as the one determined from the function body. In the rule *[app]* we see how the information comes together: the overall effect is what we can observe from evaluating the argument e_1 , what we can observe from evaluating the argument e_2 , and what we get from evaluating the body of the function called. The rules *[if]*, *[let]* and *[op]* should be straightforward.

Turning to the axioms and rules involving reference variables we make sure that we only assign values of the appropriate type to the variable. Also, in each of the cases we make sure to record that a location at the relevant program point has been created, referenced or assigned. The rule *[seq]* is straightforward and the final rule *[sub]* will be explained below.

Example 5.26 The following program

```
newA x:=1
in (newB y=!x in (x:=!y+1; !y+3))
  + (newC x:=!x in (x:=!x+1; !x+1))
```

evaluates to 8 because both summands evaluate to 4. The first summand has type and effect

$$\text{int} \ \& \ \{\text{newB}, !A, A:=, !B\}$$

and the second summand has type and effect

$$\text{int} \ \& \ \{\text{newC}, !A, C:=, !C\}$$

because the reference variable that is updated is the local one. Hence

$$\text{int} \ \& \ \{\text{newA}, A:=, !A, \text{newB}, !B, \text{newC}, C:=, !C\}$$

is the type and effect of the overall program. It follows from the effect that the variable being created at B (*y* in the program) is never reassigned after its creation. This might suggest transforming the *new_B*-construct into a *let*-construct (i.e. *let y=!x in (x:=y+1; y+3)*). ■

Subeffecting and subtyping. The purpose of the rule *[sub]* in Table 5.9 is to ensure that we obtain a conservative extension of the underlying type system. The rule is really a combination of a separate rule for *subeffecting*

$$\frac{\widehat{\Gamma} \vdash_{\text{SE}} e : \widehat{\tau} \ \& \ \varphi}{\widehat{\Gamma} \vdash_{\text{SE}} e : \widehat{\tau} \ \& \ \varphi'} \quad \text{if } \varphi \subseteq \varphi'$$

and a separate rule for *subtyping*:

$$\frac{\widehat{\Gamma} \vdash_{SE} e : \widehat{\tau} \ \& \ \varphi}{\widehat{\Gamma} \vdash_{SE} e : \widehat{\tau}' \ \& \ \varphi} \quad \text{if } \widehat{\tau} \leq \widehat{\tau}'$$

Here $\varphi \subseteq \varphi'$ means that φ is “a subset” of φ' (modulo UCAI) as discussed in Subsection 5.2.3 (and Exercise 5.3). The ordering $\widehat{\tau} \leq \widehat{\tau}'$ on annotated types is derived from the ordering on annotations as follows:

$$\widehat{\tau} \leq \widehat{\tau} \quad \frac{\widehat{\tau}'_1 \leq \widehat{\tau}_1 \quad \widehat{\tau}'_2 \leq \widehat{\tau}_2 \quad \varphi \subseteq \varphi'}{\widehat{\tau}_1 \xrightarrow{\varphi} \widehat{\tau}_2 \leq \widehat{\tau}'_1 \xrightarrow{\varphi'} \widehat{\tau}'_2} \quad \frac{\widehat{\tau} \leq \widehat{\tau}' \quad \widehat{\tau}' \leq \widehat{\tau}}{\text{ref } \widehat{\tau} \leq \text{ref } \widehat{\tau}'}$$

Note that the order of the comparison is reversed for arguments to functions; we say that $\widehat{\tau}_1 \xrightarrow{\varphi} \widehat{\tau}_2$ is *contravariant* in $\widehat{\tau}_1$ but *covariant* in φ and $\widehat{\tau}_2$. (To familiarise oneself with this idea just pretend that the types are really logical propositions and that $\xrightarrow{\varphi}$ as well as \leq mean logical implication; then it should be clear that the rule is the right one.) Also $\text{ref } \widehat{\tau}$ is both covariant in $\widehat{\tau}$ (when the reference variable is used for accessing its value as in $!r$) and contravariant in $\widehat{\tau}$ (when the reference variable is used for assignments as in $r := \dots$). This turns out to be essential for semantic correctness to hold.

This form of subtyping we shall call *shape conformant subtyping* because $\widehat{\tau}_1 \leq \widehat{\tau}_2$ implies that the two annotated types have the same underlying types, i.e. $[\widehat{\tau}_1] = [\widehat{\tau}_2]$ and hence that $\widehat{\tau}_1$ and $\widehat{\tau}_2$ have the same “shape”. (There are more permissive notions of subtyping than this, and we shall return to it in the Concluding Remarks.)

Example 5.27 Consider the following program

```
newA x:=1
in (fn f => f (fn y => !x) + f (fn z => (x:=z; z)))
    (fn g => g 1)
```

where we have omitted the labels on the function definitions. The program evaluates to 2 because each summand evaluates to 1.

The type and effect of the two arguments to f are

$$\text{int} \xrightarrow{\{!A\}} \text{int} \ \& \ \emptyset$$

$$\text{int} \xrightarrow{\{A:=\}} \text{int} \ \& \ \emptyset$$

and when f has the type

$$(\text{int} \xrightarrow{\{!A,A:=\}} \text{int}) \xrightarrow{\{!A,A:=\}} \text{int}$$

the application of f to the arguments will be well-typed: we have

$$\begin{aligned} (\text{int} \xrightarrow{\{!A\}} \text{int}) &\leq (\text{int} \xrightarrow{\{!A, A:=\}} \text{int}) \\ (\text{int} \xrightarrow{\{A:=\}} \text{int}) &\leq (\text{int} \xrightarrow{\{!A, A:=\}} \text{int}) \end{aligned}$$

and may use the rule for subtyping to change the types of the arguments to the type expected by f .

If instead we had only used the rule for subeffecting we would be obliged to let the two arguments to f have type and effect:

$$\begin{aligned} \text{int} \xrightarrow{\{A:=, !A\}} \text{int} \quad \& \quad \emptyset \\ \text{int} \xrightarrow{\{A:=, !A\}} \text{int} \quad \& \quad \emptyset \end{aligned}$$

This is indeed possible by using the rule for subeffecting just before the rule for function abstraction. ■

The combined rule $[sub]$ for subeffecting and subtyping gives rise to a conservative extension of the underlying type system. This would also have been the case if we had adopted either the rule for subeffecting or the rule for subtyping. However, had we incorporated no additional rule then this would not be the case.

Remark. One can make a distinction between “Annotated Type Systems” and “Effect Systems” but it is a subtle one and it is hardly fruitful to distinguish them in a formal way; yet intuitively there is a difference. The analysis presented in this subsection is truly an Effect System because the annotations relate neither to the input nor the output of functions; rather they relate to the internal steps of the computation. By contrast the analysis of Subsection 5.1 is an Annotated Type System because the annotations relate to intensional aspects of the semantic values: what function abstraction might be the result of evaluating the expression. ■

5.4.2 Exception Analysis

Syntax. As our next analysis we consider an *Exception Analysis*; the aim of this analysis is to determine:

For each expression, what exceptions might result from evaluating the expression.

These exceptions may be raised by primitive operators (like division by zero) or they may be explicitly raised in the program. Furthermore, there may be

the possibility of trapping an exception by means of executing an expression designed to overcome the source of the abnormal situation. To illustrate this scenario we extend the syntax of expressions in FUN (Section 5.1) as follows:

$$e ::= \dots \mid \text{raise } s \mid \text{handle } s \text{ as } e_1 \text{ in } e_2$$

The exception is raised by the `raise`-construct. If e_2 raises some exception s_1 then `handle s_2 as e_1 in e_2` will trap the exception in case $s_1 = s_2$ and this means that e_1 will be executed; if $s_1 \neq s_2$ we will continue propagating the exception s_1 . We take a simple-minded approach to exceptions and use strings to denote their identity.

Example 5.28 Consider the following program computing the combinatorial $\binom{x}{y}$ of the values x and y of the variables x and y :

```
let comb = fun f x => fn y =>
  if x<0 then raise x-out-of-range
  else if y<0 or y>x then raise y-out-of-range
  else if y=0 or y=x then 1
  else f (x-1) y + f (x-1) (y-1)
in handle x-out-of-range as 0 in comb x y
```

The program raises the exception `x-out-of-range` if x is negative and in the body of the `let`-construct this exception is trapped and the value 0 is returned. The program raises the exception `y-out-of-range` if the value of y is negative or larger than the value of x and this exception is not trapped in the program. ■

Semantics. Before presenting the analysis let us briefly sketch the semantics of the language. An expression can now give rise to an exception so we shall extend the set **Val** of values to include entities of the form `raise s` . The semantic clauses of Table 5.4 then have to be extended to take care of the new kind of values. For function application we shall for example add three rules

$$\frac{\vdash e_1 \rightarrow \text{raise } s}{\vdash e_1 e_2 \rightarrow \text{raise } s}$$

$$\frac{\vdash e_1 \rightarrow (\text{fn}_\pi x \Rightarrow e_0) \quad \vdash e_2 \rightarrow \text{raise } s}{\vdash e_1 e_2 \rightarrow \text{raise } s}$$

$$\frac{\vdash e_1 \rightarrow (\text{fn}_\pi x \Rightarrow e_0) \quad \vdash e_2 \rightarrow v_2 \quad \vdash e_0[x \mapsto v_2] \rightarrow \text{raise } s}{\vdash e_1 e_2 \rightarrow \text{raise } s}$$

reflecting that an exception can be raised in any of the subcomputations in the premise of rule [app]. Similar rules are added for the other constructs.

We can then specify the meaning of the new constructs by the following axioms and rules:

$$\begin{array}{l} \vdash \text{raise } s \longrightarrow \text{raise } s \\ \frac{\vdash e_2 \longrightarrow v_2}{\vdash \text{handle } s \text{ as } e_1 \text{ in } e_2 \longrightarrow v_2} \quad \text{if } v_2 \neq \text{raise } s \\ \frac{\vdash e_2 \longrightarrow \text{raise } s \quad \vdash e_1 \longrightarrow v_1}{\vdash \text{handle } s \text{ as } e_1 \text{ in } e_2 \longrightarrow v_1} \end{array}$$

Note that the expression e_2 in `handle s as e_1 in e_2` may also raise an exception other than s in which case it will be propagated out of the `handle`-construct. Similarly, the expression e_1 may raise an exception which will then be propagated out of the `handle`-construct.

Annotated types. The purpose of Exception Analysis is to determine which exceptions might be raised and not trapped in the program. We shall therefore take the annotations to be sets of exceptions. To get a more flexible type system and a more powerful analysis we shall use a polymorphic type system. This means that we shall allow the annotated types to contain type variables and also we shall allow the annotations to contain annotation variables. So the *annotations* (or effects) $\varphi \in \mathbf{Ann}_{ES}$ will be given by

$$\varphi ::= \{s\} \mid \varphi_1 \cup \varphi_2 \mid \emptyset \mid \beta$$

and the *annotated types* $\hat{\tau} \in \mathbf{Type}_{ES}$ will be given by:

$$\hat{\tau} ::= \text{int} \mid \text{bool} \mid \hat{\tau}_1 \xrightarrow{\varphi} \hat{\tau}_2 \mid \alpha$$

As usual effects will be considered equal modulo UCAI. To express the *polymorphism* concisely we introduce *type schemes*. They have the form

$$\hat{\sigma} ::= \forall(\zeta_1, \dots, \zeta_n). \hat{\tau}$$

where ζ_1, \dots, ζ_n is a (possible empty) list of type variables and annotation variables; if the list is empty we simply write $\hat{\tau}$ for $\forall().\hat{\tau}$.

Example 5.29 Consider the function $\text{fn}_F f \Rightarrow \text{fn}_X x \Rightarrow f x$. We can give it the type schema

$$\forall 'a, 'b, 'l. ('a \xrightarrow{\{l\}} 'b) \xrightarrow{\emptyset} ('a \xrightarrow{\{l\}} 'b)$$

which has instances like

$$(\text{int} \xrightarrow{\{x\text{-out-of-range}\}} \text{int}) \xrightarrow{\emptyset} (\text{int} \xrightarrow{\{x\text{-out-of-range}\}} \text{int})$$

and $(\text{int} \xrightarrow{\emptyset} \text{bool}) \xrightarrow{\emptyset} (\text{int} \xrightarrow{\emptyset} \text{bool})$. ■

[con]	$\widehat{\Gamma} \vdash_{\text{ES}} c : \tau_c \ \& \ \emptyset$
[var]	$\widehat{\Gamma} \vdash_{\text{ES}} x : \widehat{\sigma} \ \& \ \emptyset \quad \text{if } \widehat{\Gamma}(x) = \widehat{\sigma}$
[fn]	$\frac{\widehat{\Gamma}[x \mapsto \widehat{\tau}_x] \vdash_{\text{ES}} e_0 : \widehat{\tau}_0 \ \& \ \varphi_0}{\widehat{\Gamma} \vdash_{\text{ES}} \text{fn}_\pi x \Rightarrow e_0 : \widehat{\tau}_x \xrightarrow{\varphi_0} \widehat{\tau}_0 \ \& \ \emptyset}$
[fun]	$\frac{\widehat{\Gamma}[f \mapsto \widehat{\tau}_x \xrightarrow{\varphi_0} \widehat{\tau}_0][x \mapsto \widehat{\tau}_x] \vdash_{\text{ES}} e_0 : \widehat{\tau}_0 \ \& \ \varphi_0}{\widehat{\Gamma} \vdash_{\text{ES}} \text{fun}_\pi f x \Rightarrow e_0 : \widehat{\tau}_x \xrightarrow{\varphi_0} \widehat{\tau}_0 \ \& \ \emptyset}$
[app]	$\frac{\widehat{\Gamma} \vdash_{\text{ES}} e_1 : \widehat{\tau}_2 \xrightarrow{\varphi_0} \widehat{\tau}_0 \ \& \ \varphi_1 \quad \widehat{\Gamma} \vdash_{\text{ES}} e_2 : \widehat{\tau}_2 \ \& \ \varphi_2}{\widehat{\Gamma} \vdash_{\text{ES}} e_1 e_2 : \widehat{\tau}_0 \ \& \ \varphi_1 \cup \varphi_2 \cup \varphi_0}$
[if]	$\frac{\widehat{\Gamma} \vdash_{\text{ES}} e_0 : \text{bool} \ \& \ \varphi_0 \quad \widehat{\Gamma} \vdash_{\text{ES}} e_1 : \widehat{\tau} \ \& \ \varphi_1 \quad \widehat{\Gamma} \vdash_{\text{ES}} e_2 : \widehat{\tau} \ \& \ \varphi_2}{\widehat{\Gamma} \vdash_{\text{ES}} \text{if } e_0 \text{ then } e_1 \text{ else } e_2 : \widehat{\tau} \ \& \ \varphi_0 \cup \varphi_1 \cup \varphi_2}$
[let]	$\frac{\widehat{\Gamma} \vdash_{\text{ES}} e_1 : \widehat{\sigma}_1 \ \& \ \varphi_1 \quad \widehat{\Gamma}[x \mapsto \widehat{\sigma}_1] \vdash_{\text{ES}} e_2 : \widehat{\tau}_2 \ \& \ \varphi_2}{\widehat{\Gamma} \vdash_{\text{ES}} \text{let } x = e_1 \text{ in } e_2 : \widehat{\tau}_2 \ \& \ \varphi_1 \cup \varphi_2}$
[op]	$\frac{\widehat{\Gamma} \vdash_{\text{ES}} e_1 : \tau_{op}^1 \ \& \ \varphi_1 \quad \widehat{\Gamma} \vdash_{\text{ES}} e_2 : \tau_{op}^2 \ \& \ \varphi_2}{\widehat{\Gamma} \vdash_{\text{ES}} e_1 \text{ op } e_2 : \tau_{op} \ \& \ \varphi_1 \cup \varphi_2}$
[raise]	$\widehat{\Gamma} \vdash_{\text{ES}} \text{raise } s : \widehat{\tau} \ \& \ \{s\}$
[handle]	$\frac{\widehat{\Gamma} \vdash_{\text{ES}} e_1 : \widehat{\tau} \ \& \ \varphi_1 \quad \widehat{\Gamma} \vdash_{\text{ES}} e_2 : \widehat{\tau} \ \& \ \varphi_2}{\widehat{\Gamma} \vdash_{\text{ES}} \text{handle } s \text{ as } e_1 \text{ in } e_2 : \widehat{\tau} \ \& \ \varphi_1 \cup (\varphi_2 \setminus \{s\})}$
[sub]	$\frac{\widehat{\Gamma} \vdash_{\text{ES}} e : \widehat{\tau} \ \& \ \varphi}{\widehat{\Gamma} \vdash_{\text{ES}} e : \widehat{\tau}' \ \& \ \varphi'} \quad \text{if } \widehat{\tau} \leq \widehat{\tau}' \text{ and } \varphi \subseteq \varphi'$
[gen]	$\frac{\widehat{\Gamma} \vdash_{\text{ES}} e : \widehat{\tau} \ \& \ \varphi}{\widehat{\Gamma} \vdash_{\text{ES}} e : \forall(\zeta_1, \dots, \zeta_n). \widehat{\tau} \ \& \ \varphi}$ if ζ_1, \dots, ζ_n do not occur free in $\widehat{\Gamma}$ and φ
[ins]	$\frac{\widehat{\Gamma} \vdash_{\text{ES}} e : \forall(\zeta_1, \dots, \zeta_n). \widehat{\tau} \ \& \ \varphi}{\widehat{\Gamma} \vdash_{\text{ES}} e : (\theta \widehat{\tau}) \ \& \ \varphi}$ if θ has $\text{dom}(\theta) \subseteq \{\zeta_1, \dots, \zeta_n\}$

Table 5.10: Exception Analysis.

Typing judgements. The typing judgements for the Exception Analysis will be of the form

$$\widehat{\Gamma} \vdash_{\text{ES}} e : \widehat{\sigma} \ \& \ \varphi$$

where the type environment $\widehat{\Gamma}$ now maps variables to type schemes. As expected $\widehat{\sigma}$ denotes the type schema of e , and φ is the set of exceptions that may be raised during evaluation of e . The analysis is specified by the axioms and rules of Table 5.10. Most of the axioms and rules are straightforward modifications of those we have seen before; in the rule *[let]* the use of a type schema for the *let*-bound variable is going to give us the required polymorphism.

The axiom *[raise]* ensures that a raised exception can have any type and the effect records that the exception s might be raised. The rule *[handle]* then determines the effects of its two subexpressions and records that any exception raised by e_1 and any exception except s raised by e_2 is a possible exception for the construct. Formally, $\varphi \setminus \{s\}$ is defined as follows: $\{s\} \setminus \{s\} = \emptyset$, $\{s'\} \setminus \{s\} = \{s'\}$ if $s' \neq s$, $(\varphi \cup \varphi') \setminus \{s\} = (\varphi \setminus \{s\}) \cup (\varphi' \setminus \{s\})$, $\emptyset \setminus \{s\} = \emptyset$ and $\beta \setminus \{s\} = \beta$. (We shall consider alternative definitions shortly.)

The rule *[sub]* is the rule for *subeffecting* and *subtyping* and the ordering $\widehat{\tau} \leq \widehat{\tau}'$ on types is given by

$$\widehat{\tau} \leq \widehat{\tau} \qquad \frac{\widehat{\tau}'_1 \leq \widehat{\tau}_1 \quad \widehat{\tau}_2 \leq \widehat{\tau}'_2 \quad \varphi \subseteq \varphi'}{\widehat{\tau}_1 \xrightarrow{\varphi} \widehat{\tau}_2 \leq \widehat{\tau}'_1 \xrightarrow{\varphi'} \widehat{\tau}'_2}$$

as was also the case in Subsection 5.4.1.

The rules *[gen]* and *[ins]* are responsible for the *polymorphism*. The generalisation rule *[gen]* is used to construct type schemes: we can quantify over any type or annotation variable that does not occur free in the assumptions or in the effect of the construct; this rule is usually used before applications of the rule *[let]*. The instantiation rule *[ins]* can then be used to turn type schemes into annotated types: we just apply a substitution in order to replace the bound type and annotation variables with other types and annotations (possibly containing type variables and annotation variables); this rule is usually used after applications of the axiom *[var]*.

Example 5.30 The program

```
let f = fn g => fn x => g x
in f (fn y => if y < 0 then raise neg else y) (3-2)
  + f (fn z => if z > 0 then raise pos else 0-z) (2-3)
```

evaluates to 2 because each of the summands evaluates to 1.

We may analyse f so as to obtain the type schema

$$\forall 'a, 'b, '0. ('a \xrightarrow{'0} 'b) \xrightarrow{\emptyset} ('a \xrightarrow{'0} 'b)$$

and we may analyse the two functional arguments to f so as to obtain:

$$\begin{aligned} \text{int} \xrightarrow{\{\text{neg}\}} \text{int} \ \& \ \emptyset \\ \text{int} \xrightarrow{\{\text{pos}\}} \text{int} \ \& \ \emptyset \end{aligned}$$

We can now take two instances of the type schema for f : to match the first argument of f we use the substitution $[a \mapsto \text{int}; 'b \mapsto \text{int}; '0 \mapsto \{\text{neg}\}]$ and to match the second argument we use $[a \mapsto \text{int}; 'b \mapsto \text{int}; '0 \mapsto \{\text{pos}\}]$. Hence the type and effect of each summand is

$$\begin{aligned} \text{int} \ \& \ \{\text{neg}\} \\ \text{int} \ \& \ \{\text{pos}\} \end{aligned}$$

and therefore

$$\text{int} \ \& \ \{\text{neg}, \text{pos}\}$$

is the overall type and effect of the entire program.

If we did not avail ourselves of polymorphism we would have to rely on subeffecting and subtyping and let f have the type

$$(\text{int} \xrightarrow{\{\text{neg}, \text{pos}\}} \text{int}) \xrightarrow{\emptyset} (\text{int} \xrightarrow{\{\text{neg}, \text{pos}\}} \text{int})$$

which is more imprecise than the type scheme displayed above, although we would still be able to obtain $\text{int} \ \& \ \{\text{neg}, \text{pos}\}$ as the overall type and effect of the entire program. ■

Remark. The judgements of the Exception Analyses were of the form $\widehat{\Gamma} \vdash_{\text{ES}} e : \widehat{\sigma} \ \& \ \varphi$ but nonetheless most axioms and rules of Table 5.10 have had conclusions of the form $\widehat{\Gamma} \vdash_{\text{ES}} e : \widehat{\tau} \ \& \ \varphi$. By changing rules $[if]$, $[let]$, $[raise]$, $[handle]$ and possibly $[sub]$ one can obtain a more liberal system for Exception Analysis. We leave the details for Exercise 5.14.

In the rule $[handle]$ we make use of the notation $\varphi \setminus \{s\}$ and we next defined $\beta \setminus \{s\} = \beta$ even though β might later be instantiated to $\{s\}$. Since $\beta \setminus \{s\} \subseteq \beta$ should hold for all instantiations of β , this is semantically sound (corresponding to what happens in the rule $[sub]$). To define a less approximate system it would be natural to let

$$\varphi ::= \dots \mid \varphi \setminus \{s\}$$

and then to extend the axiomatisation of UCAI to deal with set difference. ■

5.4.3 Region Inference

Let us once again consider the language FUN as introduced in Section 5.1. The purpose of *Region Inference* is to facilitate implementing FUN in a *stack-based* as opposed to a heap-based regime as is usually the case; since the




Figure 5.1: The memory model for the stack-based implementation of FUN.

stack-based regime is quite efficient in reusing dead memory locations without relying on explicit garbage collection this might lead to an efficient implementation. But functions are statically scoped and can produce other functions as results, so it is by no means clear that a stack-based regime should be possible; the purpose of region inference is to analyse how far locally allocated data can be passed around so that the allocation of memory can be performed in an appropriate manner. This leads to a memory model for the stack-based regime where the memory is a stack of dynamic regions (r_1, r_2, r_3, \dots) and each dynamic region is an indexed list (or array) of values as illustrated in Figure 5.1.

Syntax. To make this feasible we shall introduce a notion of extended expressions that contain explicit information about regions. To this end we introduce *region names*, *region variables*, and *static regions*

$$\begin{aligned} rn &\in \mathbf{RName} && \text{region names} \\ \varrho &\in \mathbf{RVar} && \text{region variables} \\ r &\in \mathbf{Reg}_{\mathbf{Rl}} && \text{regions} \end{aligned}$$

as follows:

$$\begin{aligned} rn &::= r_1 \mid r_2 \mid r_3 \mid \dots \\ \varrho &::= "1" \mid "2" \mid "3" \mid \dots \\ r &::= \varrho \mid rn \end{aligned}$$

The syntax of *extended expressions*

$$ee \in \mathbf{EExp}$$

is given by:

$$ee ::= c \text{ at } r \mid x \mid \text{fn}_\pi x \Rightarrow ee_0 \text{ at } r \mid \text{fun}_\pi f [\vec{\varrho}] x \Rightarrow ee_0 \text{ at } r \mid ee_1 ee_2$$

$$\begin{array}{l}
 | \quad \text{if } ee_0 \text{ then } ee_1 \text{ else } ee_2 \mid \text{let } x = ee_1 \text{ in } ee_2 \mid ee_1 \text{ op } ee_2 \text{ at } r \\
 | \quad ee[\vec{r}] \text{ at } r \mid \text{letregion } \vec{\varrho} \text{ in } ee
 \end{array}$$

Here we write $\vec{\varrho}$ for a possibly empty sequence $\varrho_1, \dots, \varrho_k$ of region variables and similarly \vec{r} for a possible empty sequence r_1, \dots, r_k of regions. The main point of the extended expressions is that if a subexpression explicitly produces a new value then it has an explicit placement component, "at r ", that indicates the region in which the value is (to be) placed. The `letregion`-construct explicitly allows to deallocate the regions, $\vec{\varrho}$, that are no longer needed and the placement construct, $ee[\vec{r}]$ at r , allows us to explicitly place a "copy" of ee (typically a recursive function) in the region r . The construct for recursive function definitions explicitly takes a sequence of region variables as parameters; intuitively, this ensures that the various incarnations of the recursive function can have its local data in different regions.

Example 5.31 The relationship between expressions and extended expressions will be clarified when presenting the typing judgements below. For now we merely state that the expression e

$$(\text{let } x = 7 \text{ in fn } y \Rightarrow y+x) 9$$

will give rise to the extended expression ee :

$$\begin{array}{l}
 \text{letregion } \varrho_1, \varrho_3, \vec{\varrho}_4 \\
 \text{in } (\text{let } x = (7 \text{ at } \varrho_1) \\
 \quad \text{in } (\text{fn } y \Rightarrow (y+x) \text{ at } \varrho_2) \text{ at } \varrho_3) (9 \text{ at } \varrho_4)
 \end{array}$$

Here the value 7 is placed in the region ϱ_1 , the value of $x+y$ in region ϱ_2 , the function named Y in the region ϱ_3 and the argument 9 in region ϱ_4 . The final value (which is $7 + 9 = 16$) is therefore to be found in region ϱ_2 . All other regions (ϱ_1 , ϱ_3 , and ϱ_4) no longer serve any purpose and can be deallocated; this is made explicit by the `letregion`-construct. In the version of the program that is actually run we should replace the metavariables ϱ_1 , ϱ_3 and ϱ_4 by region variables "1", "2", and "3 and furthermore the free region variable ϱ_2 should be replaced by a region name such as `r1`. ■

Semantics. The Natural Semantics of expressions is given by Table 5.4 and we now devise a Natural Semantics for the extended expressions so as to make the rôle of regions clear. The transitions takes the form

$$\rho \vdash \langle ee, \varsigma \rangle \longrightarrow \langle v, \varsigma' \rangle$$

where ρ is an *environment*, ee is an extended expression, ς and ς' are stores (as in Figure 5.1) and v is an expressible value. Formally we shall use the

domains:

$$\begin{aligned}
 \rho &\in \text{Env} &= \text{Var}_* \rightarrow \text{EVal} \\
 v &\in \text{EVal} &= \text{RName} \times \text{Offset} \\
 o &\in \text{Offset} &= \mathbb{N} \\
 \zeta &\in \text{Store} &= \text{RName} \rightarrow_{\text{fin}} (\text{Offset} \rightarrow_{\text{fin}} \text{SVal}) \\
 w &\in \text{SVal} &
 \end{aligned}$$

Here Var_* is the finite set of variables in the extended expression ee_* of interest, a *store* is a stack of regions (where the stack is modelled as a finitary mapping from region names), a *region* is a list of storable values (where the list is modelled as a finitary mapping from indices called offsets), and a *storable value* is given by

$$w ::= c \mid \langle x, ee, \rho \rangle \mid \langle \vec{d}, x, ee, \rho \rangle$$

consisting of ordinary constants, closures, and so-called *region polymorphic closures*. In addition to the formal parameter, the body of the function, and the environment at the definition point, a region polymorphic closure also contains a list of formal region parameters that have to be instantiated whenever the function is called – in this way it is explicitly ensured that the local data can be properly placed in the store for each function call. To simplify the notation we shall allow to view a store as an element of $\text{RName} \times \text{Offset} \rightarrow_{\text{fin}} \text{SVal}$ and so write $\zeta(r, o)$ for $\zeta(r)(o)$ etc. All ordinary values are “boxed” which means that they are always placed in the store and hence an expressible value is always a pair consisting of a region name and an offset.

The semantics is defined in Table 5.11 and is explained below. We intend that an extended expression does not contain free region variables when evaluated and hence many r 's have become rn 's. The axiom *[con]* for constants explicitly allocates a new offset in the relevant region and places the constant in that cell. The axiom *[var]* for variables does not involve the allocation of a new offset and merely performs a lookup in the environment.

The axiom *[fn]* for function abstraction allocates a new offset where it stores an ordinary closure consisting of the formal parameter, the body of the function, and the current environment. The axiom *[fun]* for recursive functions constructs a region polymorphic closure that records the list of formal region variables to be instantiated by means of the placement construct explained below; also note that recursion is handled by updating the current environment with a reference to the function itself.

The rules *[app]*, *[if₁]*, *[if₂]* and *[let]* should be straightforward and the rule *[op]* for binary operations allocates a new offset for the result.

The rule *[place]* takes care of the situation where an extended expression (such as a variable or a recursive function definition) evaluates to a region

[con]	$\rho \vdash \langle c \text{ at } rn, \varsigma \rangle \longrightarrow \langle (rn, o), \varsigma[(rn, o) \mapsto c] \rangle$ <p style="margin-left: 20px;">if $o \notin \text{dom}(\varsigma(rn))$</p>
[var]	$\rho \vdash \langle x, \varsigma \rangle \longrightarrow \langle \rho(x), \varsigma \rangle$
[fn]	$\rho \vdash \langle (\text{fn}_\pi x \Rightarrow ee_0) \text{ at } rn, \varsigma \rangle \longrightarrow$ $\langle (rn, o), \varsigma[(rn, o) \mapsto \langle x, ee_0, \rho \rangle] \rangle$ <p style="margin-left: 20px;">if $o \notin \text{dom}(\varsigma(rn))$</p>
[fun]	$\rho \vdash \langle (\text{fun}_\pi f[\vec{g}] x \Rightarrow ee_0) \text{ at } rn, \varsigma \rangle \longrightarrow$ $\langle (rn, o), \varsigma[(rn, o) \mapsto \langle \vec{g}, x, ee_0, \rho[f \mapsto (rn, o)] \rangle] \rangle$ <p style="margin-left: 20px;">if $o \notin \text{dom}(\varsigma(rn))$</p>
[app]	$\frac{\rho \vdash \langle ee_1, \varsigma_1 \rangle \longrightarrow \langle (rn_1, o_1), \varsigma_2 \rangle \quad \rho \vdash \langle ee_2, \varsigma_2 \rangle \longrightarrow \langle v_2, \varsigma_3 \rangle}{\rho_0[x \mapsto v_2] \vdash \langle ee_0, \varsigma_3 \rangle \longrightarrow \langle v_0, \varsigma_4 \rangle}$ $\rho \vdash \langle ee_1 ee_2, \varsigma_1 \rangle \longrightarrow \langle v_0, \varsigma_4 \rangle$ <p style="margin-left: 20px;">if $\varsigma_3(rn_1, o_1) = \langle x, ee_0, \rho_0 \rangle$</p>
[if ₁]	$\frac{\rho \vdash \langle ee_0, \varsigma_1 \rangle \longrightarrow \langle (rn, o), \varsigma_2 \rangle \quad \rho \vdash \langle ee_1, \varsigma_2 \rangle \longrightarrow \langle v_1, \varsigma_3 \rangle}{\rho \vdash \langle \text{if } ee_0 \text{ then } ee_1 \text{ else } ee_2, \varsigma_1 \rangle \longrightarrow \langle v_1, \varsigma_3 \rangle}$ <p style="margin-left: 20px;">if $\varsigma_2(rn, o) = \text{true}$</p>
[if ₂]	$\frac{\rho \vdash \langle ee_0, \varsigma_1 \rangle \longrightarrow \langle (r\tilde{n}_1, o), \varsigma_2 \rangle \quad \rho \vdash \langle ee_2, \varsigma_2 \rangle \longrightarrow \langle v_2, \varsigma_3 \rangle}{\rho \vdash \langle \text{if } ee_0 \text{ then } ee_1 \text{ else } ee_2, \varsigma_1 \rangle \longrightarrow \langle v_2, \varsigma_3 \rangle}$ <p style="margin-left: 20px;">if $\varsigma_2(rn, o) = \text{false}$</p>
[let]	$\frac{\rho \vdash \langle ee_1, \varsigma_1 \rangle \longrightarrow \langle v_1, \varsigma_2 \rangle \quad \rho[x \mapsto v_1] \vdash \langle ee_2, \varsigma_2 \rangle \longrightarrow \langle v_2, \varsigma_3 \rangle}{\rho \vdash \langle \text{let } x = ee_1 \text{ in } ee_2, \varsigma_1 \rangle \longrightarrow \langle v_2, \varsigma_3 \rangle}$
[op]	$\frac{\rho \vdash \langle ee_1, \varsigma_1 \rangle \longrightarrow \langle (rn_1, o_1), \varsigma_2 \rangle \quad \rho \vdash \langle ee_2, \varsigma_2 \rangle \longrightarrow \langle (rn_2, o_2), \varsigma_3 \rangle}{\rho \vdash \langle (ee_1 \text{ op } ee_2) \text{ at } rn, \varsigma_1 \rangle \longrightarrow \langle (rn, o), \varsigma_3[(rn, o) \mapsto v] \rangle}$ <p style="margin-left: 20px;">if $\varsigma_3(rn_1, o_1) \text{ op } \varsigma_3(rn_2, o_2) = v$ and $o \notin \text{dom}(\varsigma_3(rn))$</p>
[place]	$\frac{\rho \vdash \langle ee, \varsigma_1 \rangle \longrightarrow \langle (rn', o'), \varsigma_2 \rangle}{\rho \vdash \langle ee[r\tilde{n}] \text{ at } rn, \varsigma_1 \rangle$ $\longrightarrow \langle (rn, o), \varsigma_2[(rn, o) \mapsto \langle x, ee_0[\vec{g} \mapsto r\tilde{n}], \rho_0 \rangle] \rangle$ <p style="margin-left: 20px;">if $o \notin \text{dom}(\varsigma_2(rn))$ and $\varsigma_2(rn', o') = \langle \vec{g}, x, ee_0, \rho_0 \rangle$</p>
[region]	$\frac{\rho \vdash \langle ee[\vec{g} \mapsto r\tilde{n}], \varsigma_1[r\tilde{n} \mapsto [\]] \rangle \longrightarrow \langle v, \varsigma_2 \rangle}{\rho \vdash \langle \text{letregion } \vec{g} \text{ in } ee, \varsigma_1 \rangle \longrightarrow \langle v, \varsigma_2 \parallel r\tilde{n} \rangle}$ <p style="margin-left: 20px;">if $\{r\tilde{n}\} \cap \text{dom}(\varsigma) = \emptyset$</p>

Table 5.11: Natural Semantics for extended expressions.

polymorphic closure. The placement construct $ee[\vec{r}]$ at r then allocates a new cell in the region r and stores a copy of the region polymorphic closure in the cell except that it ensures that the list of formal region parameters is replaced by a list of the actual region names; this is an important feature for allowing each recursive call of a function to allocate its auxiliary data locally on the stack rather than being lumped together (in the heap) with data from other recursive calls.

Finally, the rule [region] for the `letregion`-construct allocates new unused region names to be used instead of the region variables, evaluates the enclosed extended expression, and finally deallocates the newly allocated region names; formally, $dom(\zeta \setminus r\vec{n}) = dom(\zeta) \setminus \{r\vec{n}\}$ and $\forall (rn, o) \in dom(\zeta \setminus r\vec{n}) : \zeta(rn, o) = (\zeta \setminus r\vec{n})(rn, o)$.

Annotated types. When analysing the extended expressions we shall want to keep track of those regions that may be affected during evaluation: in what regions do we place (or put) data and from what regions do we access (or get) data. This is somewhat analogous to the Side Effect Analysis of Subsection 5.4.1 and will be taken care of using the effects to be defined below. Another purpose of the analysis is to keep track of the regions in which the values reside. To this end we shall say that an *extended type* is a pair

$$\hat{\tau}@r$$

consisting of an annotated type and the region where the value resides. Formally, *annotations* (or effects), *annotated types* and *type schemes*

$$\begin{aligned} \varphi &\in \mathbf{Ann}_{RI} && \text{effects} \\ \hat{\tau} &\in \mathbf{Type}_{RI} && \text{annotated types} \\ \hat{\sigma} &\in \mathbf{Scheme}_{RI} && \text{type schemes} \end{aligned}$$

are given by:

$$\begin{aligned} \varphi &::= \{\text{put } r\} \mid \{\text{get } r\} \mid \varphi_1 \cup \varphi_2 \mid \emptyset \mid \beta \\ \hat{\tau} &::= \text{int} \mid \text{bool} \mid (\hat{\tau}_1@r_1) \xrightarrow{\beta.\varphi} (\hat{\tau}_2@r_2) \mid \alpha \\ \hat{\sigma} &::= \forall(\alpha_1, \dots, \alpha_n), (\beta_1, \dots, \beta_m), [\varrho_1, \dots, \varrho_k].\hat{\tau} \\ &\quad \mid \forall(\alpha_1, \dots, \alpha_n), (\beta_1, \dots, \beta_m).\hat{\tau} \end{aligned}$$

Here we distinguish between two kinds of type schemes: *compound* type schemes (with $[\varrho]$) to be used for recursive functions and *ordinary* type schemes (without $[\varrho]$). We shall write $\hat{\tau}$ for $\forall(), ().\hat{\tau}$ whereas we only allow to write $\forall[.]\hat{\tau}$ for $\forall(), (), [.]\hat{\tau}$; in this way the two kinds of type schemes can always be distinguished. The use of an annotation variable, β , in the annotation placed on the function arrow, $\beta.\varphi$, is related to the use of *simple types* in Section 5.3 and is mainly of interest for the inference algorithm;

for the present purposes $\beta.\varphi$ can be read as $\beta \cup \varphi$ although the inference algorithm will view it as the constraint $\beta \supseteq \beta \cup \varphi$.

Example 5.32 Returning to the extended expression ee

```

letregion  $\varrho_1, \varrho_3, \varrho_4$ 
in (let x = (7 at  $\varrho_1$ )
    in (fnv y => (y+x) at  $\varrho_2$ ) at  $\varrho_3$ ) (9 at  $\varrho_4$ )

```

of Example 5.31, the subexpression 7 at ϱ_1 will have the annotated type $\text{int}@_{\varrho_1}$, the function abstraction $(\text{fn}_v y \Rightarrow (y+x) \text{ at } \varrho_2) \text{ at } \varrho_3$ will have the annotated type

$$((\text{int}@_{\varrho_4}) \xrightarrow{\beta.\varphi} (\text{int}@_{\varrho_2}))@_{\varrho_3}$$

where $\varphi = \{\text{get } \varrho_4, \text{get } \varrho_1, \text{put } \varrho_2\}$, and the overall expression will have the extended type $\text{int}@_{\varrho_2}$. ■

Typing judgements. It would be quite feasible to define typing judgements for verifying that extended expressions are correctly typed. However, the main use of region inference is to facilitate the implementation of FUN and for this reason we shall touch upon Section 1.8 and let the typing judgements also describe how to translate expressions into extended expressions. This suggests using typing judgements of the form

$$\hat{\Gamma} \vdash_{\text{RI}} e \rightsquigarrow ee : \hat{r}@r \ \& \ \varphi$$

where $\hat{\Gamma}$ is a type environment mapping variables into extended type schemes that are pairs consisting of a type scheme and a region.

The analysis is defined in Tables 5.12 and 5.13 and is explained below. The axiom $[con]$ for constants inserts an explicit placement component and records the placement in the effect. The axiom $[var]$ for variables is straightforward as it involves no explicit placements.

The rule $[fn]$ for ordinary function abstraction transforms the body of the function and then inserts an explicit placement component for the function itself. The rule $[fun]$ for recursive functions involves a restricted form of *polymorphic recursion* (excluding type variables as this would be undecidable): polymorphic recursion means that when analysing the body of the function we are allowed to use the recursive function polymorphically. This generality gives much more precise type information and is essential for the success of region inference. For conciseness the rule has been presented in a non syntax directed manner where the rule for ordinary functions is needed to analyse the premise; clearly the rule could have been expanded so as to be syntax directed. The rules for application, conditional, local definitions and operators are straightforward.

[con]	$\widehat{\Gamma} \vdash_{\text{RI}} c \rightsquigarrow c \text{ at } r : (\tau_c @ r) \ \& \ \{\text{put } r\}$
[var]	$\widehat{\Gamma} \vdash_{\text{RI}} x \rightsquigarrow x : \widehat{\sigma} \ \& \ \emptyset \quad \text{if } \widehat{\Gamma}(x) = \widehat{\sigma}$
[fn]	$\frac{\widehat{\Gamma}[x \mapsto \widehat{\tau}_x @ r_x] \vdash_{\text{RI}} e_0 \rightsquigarrow ee_0 : (\widehat{\tau}_0 @ r_0) \ \& \ \varphi_0}{\widehat{\Gamma} \vdash_{\text{RI}} \text{fn}_\pi x \Rightarrow e_0 \rightsquigarrow (\text{fn}_\pi x \Rightarrow ee_0) \text{ at } r : ((\widehat{\tau}_x @ r_x \xrightarrow{\beta, \varphi_0} \widehat{\tau}_0 @ r_0) @ r) \ \& \ \{\text{put } r\}}$
[fun]	$\frac{\widehat{\Gamma}[f \mapsto (\forall \vec{\beta}, [\vec{\sigma}]. \widehat{\tau}) @ r] \vdash_{\text{RI}} \text{fn}_\pi x \Rightarrow e_0 \rightsquigarrow (\text{fn}_\pi x \Rightarrow ee_0) \text{ at } r : (\widehat{\tau} @ r) \ \& \ \varphi}{\widehat{\Gamma} \vdash_{\text{RI}} \text{fun}_\pi f x \Rightarrow e_0 \rightsquigarrow (\text{fun}_\pi f [\vec{\sigma}] x \Rightarrow ee_0) \text{ at } r : ((\forall \vec{\beta}, [\vec{\sigma}]. \widehat{\tau}) @ r) \ \& \ \varphi}$ <p style="margin-left: 20px;">if $\vec{\beta}$ and $\vec{\sigma}$ do not occur free in $\widehat{\Gamma}$ and φ</p>
[app]	$\frac{\widehat{\Gamma} \vdash_{\text{RI}} e_1 \rightsquigarrow ee_1 : ((\widehat{\tau}_2 @ r_2 \xrightarrow{\beta_0, \varphi_0} \widehat{\tau}_0 @ r_0) @ r_1) \ \& \ \varphi_1 \quad \widehat{\Gamma} \vdash_{\text{RI}} e_2 \rightsquigarrow ee_2 : (\widehat{\tau}_2 @ r_2) \ \& \ \varphi_2}{\widehat{\Gamma} \vdash_{\text{RI}} e_1 e_2 \rightsquigarrow ee_1 ee_2 : (\widehat{\tau}_0 @ r_0) \ \& \ \varphi_1 \cup \varphi_2 \cup \varphi_0 \cup \beta_0 \cup \{\text{get } r_1\}}$
[if]	$\frac{\widehat{\Gamma} \vdash_{\text{RI}} e_0 \rightsquigarrow ee_0 : (\text{bool} @ r_0) \ \& \ \varphi_0 \quad \widehat{\Gamma} \vdash_{\text{RI}} e_1 \rightsquigarrow ee_1 : (\widehat{\tau} @ r) \ \& \ \varphi_1 \quad \widehat{\Gamma} \vdash_{\text{RI}} e_2 \rightsquigarrow ee_2 : (\widehat{\tau} @ r) \ \& \ \varphi_2}{\widehat{\Gamma} \vdash_{\text{RI}} \text{if } e_0 \text{ then } e_1 \ \text{else } e_2 \rightsquigarrow \text{if } ee_0 \text{ then } ee_1 \ \text{else } ee_2 : (\widehat{\tau} @ r) \ \& \ \varphi_0 \cup \varphi_1 \cup \varphi_2 \cup \{\text{get } r_0\}}$
[let]	$\frac{\widehat{\Gamma} \vdash_{\text{RI}} e_1 \rightsquigarrow ee_1 : (\widehat{\sigma}_1 @ r_1) \ \& \ \varphi_1 \quad \widehat{\Gamma}[x \mapsto \widehat{\sigma}_1 @ r_1] \vdash_{\text{RI}} e_2 \rightsquigarrow ee_2 : (\widehat{\tau}_2 @ r_2) \ \& \ \varphi_2}{\widehat{\Gamma} \vdash_{\text{RI}} \text{let } x = e_1 \text{ in } e_2 \rightsquigarrow \text{let } x = ee_1 \text{ in } ee_2 : (\widehat{\tau}_2 @ r_2) \ \& \ \varphi_1 \cup \varphi_2}$
[op]	$\frac{\widehat{\Gamma} \vdash_{\text{RI}} e_1 \rightsquigarrow ee_1 : (\tau_{op}^1 @ r_1) \ \& \ \varphi_1 \quad \widehat{\Gamma} \vdash_{\text{RI}} e_2 \rightsquigarrow ee_2 : (\tau_{op}^2 @ r_2) \ \& \ \varphi_2}{\widehat{\Gamma} \vdash_{\text{RI}} e_1 \text{ op } e_2 \rightsquigarrow (ee_1 \text{ op } ee_2) \text{ at } r : (\tau_{op} @ r \ \& \ \varphi_1 \cup \varphi_2 \cup \{\text{get } r_1, \text{get } r_2, \text{put } r\}}$

Table 5.12: Region Inference Analysis and Translation (part 1).

Then we have a rule [sub] for *subeffecting* and *subtyping*; clearly one could decide to dispense with subtyping in which case subeffecting could be integrated with function abstraction as illustrated in Section 5.1.

There are two rules [gen₁] and [gen₂] for generalising over type variables. One

[sub]	$\frac{\widehat{\Gamma} \vdash_{\text{RI}} e \rightsquigarrow ee : (\widehat{\tau} @ r) \& \varphi}{\widehat{\Gamma} \vdash_{\text{RI}} e \rightsquigarrow ee : (\widehat{\tau}' @ r) \& \varphi'}$	if $\widehat{\tau} \leq \widehat{\tau}'$ and $\varphi \subseteq \varphi'$
[gen ₁]	$\frac{\widehat{\Gamma} \vdash_{\text{RI}} e \rightsquigarrow ee : (\widehat{\tau} @ r) \& \varphi}{\widehat{\Gamma} \vdash_{\text{RI}} e \rightsquigarrow ee : ((\forall \vec{\alpha}. \widehat{\tau}) @ r) \& \varphi}$	if $\vec{\alpha}$ do not occur free in $\widehat{\Gamma}$ and φ
[gen ₂]	$\frac{\widehat{\Gamma} \vdash_{\text{RI}} e \rightsquigarrow ee : ((\forall \vec{\beta}, [\vec{\rho}]. \widehat{\tau}) @ r) \& \varphi}{\widehat{\Gamma} \vdash_{\text{RI}} e \rightsquigarrow ee : ((\forall \vec{\alpha}, \vec{\beta}, [\vec{\rho}]. \widehat{\tau}) @ r) \& \varphi}$	if $\vec{\alpha}$ do not occur free in $\widehat{\Gamma}$ and φ
[ins ₁]	$\frac{\widehat{\Gamma} \vdash_{\text{RI}} e \rightsquigarrow ee : ((\forall \vec{\alpha}, \vec{\beta}. \widehat{\tau}) @ r) \& \varphi}{\widehat{\Gamma} \vdash_{\text{RI}} e \rightsquigarrow ee : ((\theta \widehat{\tau}) @ r) \& \varphi}$	if θ has $\text{dom}(\theta) \subseteq \{\vec{\alpha}, \vec{\beta}\}$
[ins ₂]	$\frac{\widehat{\Gamma} \vdash_{\text{RI}} e \rightsquigarrow ee : ((\forall \vec{\alpha}, \vec{\beta}, [\vec{\rho}]. \widehat{\tau}) @ r) \& \varphi}{\widehat{\Gamma} \vdash_{\text{RI}} e \rightsquigarrow ee[\theta \vec{\rho}] \text{ at } r' : ((\theta \widehat{\tau}) @ r') \& \varphi \cup \{\text{get } r, \text{put } r'\}}$	if θ has $\text{dom}(\theta) \subseteq \{\vec{\alpha}, \vec{\beta}, \vec{\rho}\}$
[region]	$\frac{\widehat{\Gamma} \vdash_{\text{RI}} e \rightsquigarrow ee : (\widehat{\tau} @ r) \& \varphi}{\widehat{\Gamma} \vdash_{\text{RI}} e \rightsquigarrow \text{letregion } \vec{\rho} \text{ in } ee : (\widehat{\tau} @ r) \& \varphi'}$	if $\varphi' = \text{Observe}(\widehat{\Gamma}, \widehat{\tau}, r)(\varphi)$ and $\vec{\rho}$ occurs in φ but not in φ'

Table 5.13: Region Inference Analysis and Translation (part 2).

applies to ordinary types and the other to compound type schemes. In both cases we could have adapted the rule so as also to generalise over (additional) region and effect variables; however, even when doing so, the type inference is quite separate from the region and effect inference.

There are two inference rules [ins₁] and [ins₂] for instantiating a type scheme. One is for ordinary types schemes and is invisible as far as the syntax of the extended expression is concerned. The other is for compound type schemes and is visible in the extended expression in that an explicit placement construct is introduced; additionally the effect records that the value has been accessed and placed again. Both rules would typically be used immediately after the axiom for variables.

The rule [region] for the letregion-construct uses an auxiliary function, *Observe*, to reduce the effect to what is visible from the outside; region variables that are no longer visible can then be encapsulated within the program.

The auxiliary function may be defined as follows:

$$\begin{aligned}
 \text{Observe}(\widehat{\Gamma}, \widehat{\tau}, r')(\{\text{put } r\}) &= \begin{cases} \{\text{put } r\} & \text{if } r \text{ occurs in } \widehat{\Gamma}, \widehat{\tau}, \text{ or } r' \\ \emptyset & \text{otherwise} \end{cases} \\
 \text{Observe}(\widehat{\Gamma}, \widehat{\tau}, r')(\{\text{get } r\}) &= \begin{cases} \{\text{get } r\} & \text{if } r \text{ occurs in } \widehat{\Gamma}, \widehat{\tau}, \text{ or } r' \\ \emptyset & \text{otherwise} \end{cases} \\
 \text{Observe}(\widehat{\Gamma}, \widehat{\tau}, r')(\varphi_1 \cup \varphi_2) &= \text{Observe}(\widehat{\Gamma}, \widehat{\tau}, r')(\varphi_1) \cup \text{Observe}(\widehat{\Gamma}, \widehat{\tau}, r')(\varphi_2) \\
 \text{Observe}(\widehat{\Gamma}, \widehat{\tau}, r')(\emptyset) &= \emptyset \\
 \text{Observe}(\widehat{\Gamma}, \widehat{\tau}, r')(\beta) &= \begin{cases} \beta & \text{if } \beta \text{ occurs in } \widehat{\Gamma}, \widehat{\tau}, \text{ or } r' \\ \emptyset & \text{otherwise} \end{cases}
 \end{aligned}$$

Example 5.33 In Example 5.31 we considered the expression e

$$(\text{let } x = 7 \text{ in } \text{fn}_V y \Rightarrow y+x) 9$$

and the extended expression ee :

$$\begin{aligned}
 &\text{letregion } \varrho_1, \varrho_3, \varrho_4 \\
 &\text{in } (\text{let } x = (7 \text{ at } \varrho_1) \\
 &\quad \text{in } (\text{fn}_V y \Rightarrow (y+x) \text{ at } \varrho_2) \text{ at } \varrho_3) (9 \text{ at } \varrho_4)
 \end{aligned}$$

One can now show that $[\] \vdash_{\text{RI}} e \rightsquigarrow ee : (\text{int@}\varrho_2) \ \& \ \{\text{put } \varrho_2\}$. ■

5.5 Behaviours

So far the effects have had a rather simple structure in that they merely denote *sets* of atomic actions like accessing a value or raising an exception. The effects have not attempted to capture the *temporal order* of these atomic actions. Often such information would be useful for example to check that a variable is not accessed until after it has been assigned a value. In this section we shall show how to devise a Type and Effect System where effects (called behaviours) are able to record such temporal ordering in the context of a *Communication Analysis* for a fragment of Concurrent ML.

5.5.1 Communication Analysis

Syntax. Let us consider an extension of the language FUN with constructs for generating new processes, for communicating between processes over typed channels, and for creating new channels. The syntactic category $e \in \text{Exp}$ of *expressions* is now given by:

$$e ::= \dots \mid \text{channel}_\pi \mid \text{spawn } e_0 \mid \text{send } e_1 \text{ on } e_2 \mid \text{receive } e_0 \mid e_1; e_2 \mid ch$$




Figure 5.2: The pipeline produced by pipe $[f_1, f_2]$ inp out.

Here `channelr` creates a new channel identifier (denoted ch above), `spawn e_0` generates a new parallel process that executes e_0 , and `send v on ch` sends the value v to another process ready to receive a value by means of `receive ch` ; sequential composition is as before. Channel identifiers

$ch \in \mathbf{Chan}$ channel identifiers

are created dynamically and are given by:

$ch ::= \text{chan1} \mid \text{chan2} \mid \dots$

Also we shall assume that the constants, $c \in \mathbf{Const}$, not only include integers and booleans but also a special value called unit and denoted by $()$; this is the value to be returned by the `spawn` and `send` constructs.

Example 5.34 In this example we shall imagine that the expressions are extended with operations on lists: `isnil e` tests whether or not the list e is empty, `hd e` selects the first element and `tl e` selects the remainder of the list. We now define a function `pipe` that takes a list of functions, an input channel and an output channel as arguments; it constructs a pipeline of processes that apply the functions to the data arriving at the input channel and returns the results on the output channel (see Figure 5.2). It makes use of the function `node` that takes a function, an input channel and an output channel as arguments and it is defined as follows:

```
let node = fnf f => fni inp => fno out =>
  spawn ((funH h d => let v = receive inp
                    in send (f v) on out;
                    h d) ())
in funp pipe fs => fni inp => fno out =>
  if isnil fs then node (fnx x =>x) inp out
  else let ch = channelc
        in (node (hd fs) inp ch; pipe (tl fs) ch out)
```

To deal with the empty list the function produces a process that just applies the identity function (denoted `id` in Figure 5.2). ■

$(\text{fn}_\pi x \Rightarrow e) v \rightarrow e[x \mapsto v]$
$\text{let } x = v \text{ in } e \rightarrow e[x \mapsto v]$
$v_1 \text{ op } v_2 \rightarrow v \text{ if } v_1 \text{ op } v_2 = v$
$\text{fun}_\pi f x \Rightarrow e \rightarrow (\text{fn}_\pi x \Rightarrow e)[f \mapsto (\text{fun}_\pi f x \Rightarrow e)]$
$\text{if true then } e_1 \text{ else } e_2 \rightarrow e_1$
$\text{if false then } e_1 \text{ else } e_2 \rightarrow e_2$
$v; e \rightarrow e$

Table 5.14: The sequential semantics.

Semantics. We shall begin by defining the operational semantics of the sequential fragment of the language and then show how to incorporate it into the concurrent fragment. This will make use of a notion of evaluation context in order to obtain a succinct specification of the semantics. The sequential fragment is evaluated in an eager left to right manner and does not need any notion of a store, a state or an environment. A fully evaluated expression gives rise to a value which is a constant, a channel identifier or a function abstraction. This may be modelled by

$$v \in \text{Val values}$$

defined by:

$$v ::= c \mid ch \mid \text{fn}_\pi x \Rightarrow e_0$$

There is no need to package the function abstraction with an environment because the semantics will treat a function application $(\text{fn}_\pi x \Rightarrow e_0) v$ by substituting v for x in e_0 yielding $e_0[x \mapsto v]$.

Part of the *sequential semantics* is specified in Table 5.14. In the manner of a Structural Operational Semantics it axiomatises the relation

$$e_1 \rightarrow e_2$$

for when the expression e_1 in one step evaluates to e_2 . However, it does not describe how e_1 may evaluate to e_2 as a result of a subcomponent e_{11} of e_1 evaluating to e_{12} unlike what has been the case for the Structural Operational Semantics used so far. As an example Table 5.14 cannot be used to show that $(1 + 2) + 4 \rightarrow 3 + 4$ although it can be used to show that $1 + 2 \rightarrow 3$.

To recover from this shortcoming we shall make use of *evaluation contexts*; these are expressions containing *one hole* which is written $[]$. Formally,

evaluation contexts E are given by:

$$\begin{aligned}
 E ::= & \ [] \mid E e \mid v E \mid \text{let } x = E \text{ in } e \\
 & \mid \text{if } E \text{ then } e_1 \text{ else } e_2 \mid E \text{ op } e \mid v \text{ op } E \\
 & \mid \text{send } E \text{ on } e \mid \text{send } v \text{ on } E \mid \text{receive } E \mid E; e
 \end{aligned}$$

Here the syntax ensures that E is an expression containing exactly one hole while e is an ordinary expression without any holes. The definition of E may be read as follows: you are allowed to evaluate an expression on its own, you may evaluate the function part of an application, you may evaluate the argument part of an application only after the function part has been fully evaluated etc. We shall write $E[e]$ for the expression obtained by replacing the hole of E with e so for example if E is $[] + 4$ then $E[e]$ is $e + 4$. We dispense with the detailed definition of $E[e]$ because it is straightforward: since the hole in E never occurs in the scope of a bound variable there is no risk of variable capture.

The basic idea then is to stipulate that e_1 evaluates to e_2 in one step ($e_1 \Rightarrow e_2$) if there exists E, e_{10} and e_{20} such that $e_1 = E[e_{10}]$, $e_{10} \rightarrow e_{20}$ and $e_2 = E[e_{20}]$. As an example, $(1+2)+4 \Rightarrow 3+4$ by taking E to be $[]+4$, e_{10} to be $1+2$ and e_{20} to be 3 . Note that having $E \text{ op } e$ as an evaluation context corresponds to having an inference rule

$$\frac{e_{10} \rightarrow e_{20}}{e_{10} \text{ op } e_2 \rightarrow e_{20} \text{ op } e_2}$$

as in Table 3.3. As already said, an advantage of using evaluation contexts is that the description of the semantics often becomes more succinct; we shall benefit from this below.

The *concurrent semantics* operates on a finite pool, PP , of processes and a finite set, CP , of channels. The set of channels keeps track of those channels that have been generated so far; this allows us to evaluate $\text{channel}_{1,\pi}$ by generating a new channel that has never been used before. The pool of processes both keeps track of the processes spawned so far and of the expression residing on each process; this allows us to allocate new processes when an expression is spawned and to let distinct processes communicate with one another. Formally

$$p \in \mathbf{Proc} \text{ processes}$$

is defined by

$$p ::= \text{proc1} \mid \text{proc2} \mid \dots$$

and we take

$$\begin{aligned}
 CP & \in \mathcal{P}_{\text{fin}}(\mathbf{Chan}) \\
 PP & \in \mathbf{Proc} \rightarrow_{\text{fin}} \mathbf{Exp}
 \end{aligned}$$

where it will be the case that each $PP(p)$ is a closed expression. We shall write $PP[p : e]$ for PP' given by $\text{dom}(PP') = \text{dom}(PP) \cup \{p\}$, $PP'(p) = e$ and $PP'(q) = PP(q)$ for $q \neq p$.

$[seq]$	$CP, PP[p : E[e_1]] \Rightarrow CP, PP[p : E[e_2]]$ if $e_1 \rightarrow e_2$
$[chan]$	$CP, PP[p : E[channel_\pi]] \Rightarrow CP \cup \{ch\}, PP[p : E[ch]]$ if $ch \notin CP$
$[spawn]$	$CP, PP[p : E[spawn e_0]] \Rightarrow CP, PP[p : E[()]] [p_0 : e_0]$ if $p_0 \notin dom(PP) \cup \{p\}$
$[comm]$	$CP, PP[p_1 : E_1[send v on ch]] [p_2 : E_2[receive ch]]$ $\Rightarrow CP, PP[p_1 : E_1[()]] [p_2 : E_2[v]]$ if $p_1 \neq p_2$

Table 5.15: The concurrent semantics.

The concurrent semantics is specified in Table 5.15. It is a Structural Operational Semantics that axiomatises the relation $CP_1, PP_1 \Rightarrow CP_2, PP_2$ for when one configuration CP_1, PP_1 in one step evolves into another configuration CP_2, PP_2 . Thanks to the use of evaluation contexts all clauses can be written succinctly. The clause $[seq]$ incorporates the sequential semantics into the concurrent semantics (and corresponds to the discussion of $e_1 \Rightarrow e_2$ above). The clause $[chan]$ takes care of the allocation of channels: $channel_\pi$ is replaced by a fresh channel identifier ch . The clause $[spawn]$ generates a new process, initialises it to the expression to be executed, and replaces the $spawn$ -construct by the unit value. Finally, the clause $[comm]$ allows synchronous communication between distinct processes: the $receive$ -construct is replaced by the value being sent and the $send$ -construct is replaced by the unit value.

Annotated types. The purpose of the *Communication Analysis* is for each expression to determine its communication behaviour: what channels will be allocated, what type of entities will be sent and received over channels, and what is the behaviour of the processes being generated: furthermore, we are interested in recording the temporal order (or “causality”) among these actions: what takes place before what. There are several ways to formalise this and we shall choose one where the inference system is not overly complicated (but where the inference algorithm presents some challenges).

To formalise our idea we introduce the following syntactic categories:

$\hat{\tau} \in$	Type _{CA}	types
$\varphi \in$	Ann _{CA}	annotations (or behaviours)
$r \in$	Reg _{CA}	regions
$\hat{\sigma} \in$	Scheme _{CA}	type schemes

Types are much as before but extended with a unit type (for the unit value) and a type for channels:

$$\hat{\tau} ::= \alpha \mid \text{bool} \mid \text{int} \mid \text{unit} \mid \hat{\tau}_1 \xrightarrow{\varphi} \hat{\tau}_2 \mid \hat{\tau} \text{ chan } r$$

Type variables, $\alpha \in \mathbf{TVar}$, are as before.

Behaviours differ from the annotations and effects used so far in that we shall not merely use union for combining them:

$$\begin{aligned} \varphi ::= & \beta \mid \Lambda \mid \varphi_1; \varphi_2 \mid \varphi_1 + \varphi_2 \mid \text{rec}\beta.\varphi \\ & \mid \hat{\tau} \text{ chan } r \mid \text{spawn } \varphi \mid r!\hat{\tau} \mid r?\hat{\tau} \end{aligned}$$

Behaviour variables, $\beta \in \mathbf{AVar}$, are as before and the behaviour Λ is used for atomic actions that do not involve communication; in a sense it corresponds to the empty set in previous annotations although it will be more intuitive to think of it as the empty string in regular expressions or as the silent action in process calculi. The behaviour $\varphi_1; \varphi_2$ says that φ_1 takes place before φ_2 whereas $\varphi_1 + \varphi_2$ indicates a choice between φ_1 and φ_2 (as will be the case for the conditional); this is reminiscent of constructs in regular expressions as well as in process algebras. The construct $\text{rec}\beta.\varphi$ indicates a recursive behaviour that acts as given by φ except that any occurrence of β stands for $\text{rec}\beta.\varphi$ itself.

The behaviour $\hat{\tau} \text{ chan } r$ indicates that a new channel has been allocated over which entities of type $\hat{\tau}$ can be communicated; the region r indicates the set of program points (see below) where the creation could have taken place. The behaviour $\text{spawn } \varphi$ indicates that a new process has been generated and that it operates as described by φ . Next $r!\hat{\tau}$ indicates that a value is sent over a channel of type $\hat{\tau} \text{ chan } r$ and $r?\hat{\tau}$ indicates that a value is received over a channel of that type; this is reminiscent of constructs in most process algebras (in particular CSP).

Regions have a bit more structure than in Subsection 5.4.3 because now we shall also be able to take the union of regions:

$$r ::= \{\pi\} \mid \varrho \mid r_1 \cup r_2 \mid \emptyset$$

Region variables, $\varrho \in \mathbf{RVar}$, are as before. In Subsection 5.4.3 the static regions were used to ensure that at run-time data would be allocated in the same dynamic regions; here regions are used to identify the set $\{\pi_1, \dots, \pi_n\}$ of program points where a channel might have been created. (So program points now play the rôle of region names.) However, these regions will not appear explicitly in the syntax of expressions, types or behaviours.

Type schemes have the form

$$\hat{\sigma} ::= \forall(\zeta_1, \dots, \zeta_n).\hat{\tau}$$

where ζ_1, \dots, ζ_n is a (possibly empty) list of type variables, behaviour variables and region variables; if the list is empty we simply write $\hat{\tau}$ for $\forall().\hat{\tau}$.

Example 5.35 Returning to the program of Example 5.34, the node function is intended to have the type schema

$$\forall a, 'b, '1, '2. (a \xrightarrow{1} 'b) \xrightarrow{\Lambda} (a \text{ chan } '1) \xrightarrow{\Lambda} ('b \text{ chan } '2) \xrightarrow{\varphi} \text{unit}$$

where $\varphi = \text{spawn}(\text{rec } '2. (('1?'a; '1; '2!'b; '2))$

corresponding to the function argument having type $a \xrightarrow{1} 'b$, the input channel having type $a \text{ chan } '1$ and the output channel having type $'b \text{ chan } '2$. When supplied with these arguments the node function will spawn a process that recursively will read on the input channel, execute the function supplied as its parameter, and write on the output channel – this is exactly what is expressed by φ .

The pipe function is intended to have the type schema

$$\forall a, '1, '2. ((a \xrightarrow{1} 'a) \text{ list } \xrightarrow{\Lambda} (a \text{ chan } ('1 \cup \{C\})) \xrightarrow{\Lambda} (a \text{ chan } '2) \xrightarrow{\varphi'} \text{unit})$$

where $\varphi' = \text{rec } '2. (\text{spawn}(\text{rec } '3. (('1 \cup \{C\})?'a; \Lambda; '2!'a; '3)) + a \text{ chan } C; \text{spawn}(\text{rec } '4. (('1 \cup \{C\})?'a; '1; C!'a; '4)); '2)$

where the first summand in the body of φ' corresponds to the `then`-branch where node is called with the identity function (which has behaviour Λ) and the second to the `else`-branch of the conditional. Here we see that a channel is created, a process is spawned and then the overall behaviour recurses. We shall return to these types schemes after presenting the typing rules. ■

Typing judgements. The typing judgements for the Communication Analysis will be of the form

$$\hat{\Gamma} \vdash_{\text{CA}} e : \hat{\sigma} \ \& \ \varphi$$

where the type environment $\hat{\Gamma}$ maps variables to type schemes (or types), $\hat{\sigma}$ is the type scheme (or type) for the expression e , and φ is the behaviour that may arise during evaluation of e . The analysis is specified by the axiom and rules of Tables 5.16 and 5.17 and have many points in common with those we have seen before; the differences are explained below.

The axioms $[con]$ and $[var]$ for constants and variables differ from the similar axioms in Table 5.10 in that Λ is used instead of \emptyset . A similar remark holds for the two rules $[fn]$ and $[fun]$ for functions. In the rule $[app]$ for function application we now use sequencing to express that we first evaluate the function part, then the argument and finally the body of the function. In the rule $[if]$ for conditional we additionally use choice to express that only one of the

[con]	$\widehat{\Gamma} \vdash_{CA} c : \tau_c \ \& \ \Lambda$
[var]	$\widehat{\Gamma} \vdash_{CA} x : \widehat{\sigma} \ \& \ \Lambda \quad \text{if } \widehat{\Gamma}(x) = \widehat{\sigma}$
[fn]	$\frac{\widehat{\Gamma}[x \mapsto \widehat{\tau}_x] \vdash_{CA} e_0 : \widehat{\tau}_0 \ \& \ \varphi_0}{\widehat{\Gamma} \vdash_{CA} \text{fn}_\pi x \Rightarrow e_0 : \widehat{\tau}_x \xrightarrow{\varphi_0} \widehat{\tau}_0 \ \& \ \Lambda}$
[fun]	$\frac{\widehat{\Gamma}[f \mapsto \widehat{\tau}_x \xrightarrow{\varphi_0} \widehat{\tau}_0][x \mapsto \widehat{\tau}_x] \vdash_{CA} e_0 : \widehat{\tau}_0 \ \& \ \varphi_0}{\widehat{\Gamma} \vdash_{CA} \text{fun}_\pi f x \Rightarrow e_0 : \widehat{\tau}_x \xrightarrow{\varphi_0} \widehat{\tau}_0 \ \& \ \Lambda}$
[app]	$\frac{\widehat{\Gamma} \vdash_{CA} e_1 : \widehat{\tau}_2 \xrightarrow{\varphi_0} \widehat{\tau}_0 \ \& \ \varphi_1 \quad \widehat{\Gamma} \vdash_{CA} e_2 : \widehat{\tau}_2 \ \& \ \varphi_2}{\widehat{\Gamma} \vdash_{CA} e_1 e_2 : \widehat{\tau}_0 \ \& \ \varphi_1; \varphi_2; \varphi_0}$
[if]	$\frac{\widehat{\Gamma} \vdash_{CA} e_0 : \text{bool} \ \& \ \varphi_0 \quad \widehat{\Gamma} \vdash_{CA} e_1 : \widehat{\tau} \ \& \ \varphi_1 \quad \widehat{\Gamma} \vdash_{CA} e_2 : \widehat{\tau} \ \& \ \varphi_2}{\widehat{\Gamma} \vdash_{CA} \text{if } e_0 \text{ then } e_1 \text{ else } e_2 : \widehat{\tau} \ \& \ \varphi_0; (\varphi_1 + \varphi_2)}$
[let]	$\frac{\widehat{\Gamma} \vdash_{CA} e_1 : \widehat{\sigma}_1 \ \& \ \varphi_1 \quad \widehat{\Gamma}[x \mapsto \widehat{\sigma}_1] \vdash_{CA} e_2 : \widehat{\tau}_2 \ \& \ \varphi_2}{\widehat{\Gamma} \vdash_{CA} \text{let } x = e_1 \text{ in } e_2 : \widehat{\tau}_2 \ \& \ \varphi_1; \varphi_2}$
[op]	$\frac{\widehat{\Gamma} \vdash_{CA} e_1 : \tau_{op}^1 \ \& \ \varphi_1 \quad \widehat{\Gamma} \vdash_{CA} e_2 : \tau_{op}^2 \ \& \ \varphi_2}{\widehat{\Gamma} \vdash_{CA} e_1 \text{ op } e_2 : \tau_{op} \ \& \ \varphi_1; \varphi_2; \Lambda}$

Table 5.16: Communication Analysis (part 1).

then- and else-branches are taken. Then the rules [let] and [op] should be straightforward.

The axiom [chan] for channel creation makes sure to record the program point in the type as well as the behaviour, the rule [spawn] encapsulates the behaviour of the spawned process in the behaviour of the construct itself and the rules [send] and [receive] for sending and receiving values over channels indicate the order in which the arguments are evaluated and then produce the behaviour for the axiom taken. The rules [seq] and [ch] are straightforward and the rules [gen] and [ins] are much as in Table 5.10.

The rule [sub] for *subeffecting* and *subtyping* looks like the one in Table 5.10 but involves a few subtleties. The ordering $\widehat{\tau} \leq \widehat{\tau}'$ on types is given by

$$\widehat{\tau} \leq \widehat{\tau}' \quad \frac{\widehat{\tau}_1 \leq \widehat{\tau}_1 \quad \widehat{\tau}_2 \leq \widehat{\tau}_2 \quad \varphi \sqsubseteq \varphi'}{\widehat{\tau}_1 \xrightarrow{\varphi} \widehat{\tau}_2 \leq \widehat{\tau}_1' \xrightarrow{\varphi'} \widehat{\tau}_2'} \quad \frac{\widehat{\tau} \leq \widehat{\tau}' \quad \widehat{\tau}' \leq \widehat{\tau} \quad r \sqsubseteq r'}{\widehat{\tau} \text{ chan } r \leq \widehat{\tau}' \text{ chan } r'}$$

and is similar to the definition in Subsection 5.4.1: $\widehat{\tau}_1 \xrightarrow{\varphi} \widehat{\tau}_2$ is contravariant in $\widehat{\tau}_1$ but covariant in φ and $\widehat{\tau}_2$, and $\widehat{\tau} \text{ chan } r$ is both covariant in $\widehat{\tau}$ (for when a value is sent) and contravariant in $\widehat{\tau}$ (for when a value is received)

[chan]	$\widehat{\Gamma} \vdash_{CA} \text{channel}_\pi : \widehat{\tau} \text{ chan } \{\pi\} \ \& \ \widehat{\tau} \text{ chan } \{\pi\}$
[spawn]	$\frac{\widehat{\Gamma} \vdash_{CA} e_0 : \widehat{\tau}_0 \ \& \ \varphi_0}{\widehat{\Gamma} \vdash_{CA} \text{spawn } e_0 : \text{unit} \ \& \ \text{spawn } \varphi_0}$
[send]	$\frac{\widehat{\Gamma} \vdash_{CA} e_1 : \widehat{\tau} \ \& \ \varphi_1 \quad \widehat{\Gamma} \vdash_{CA} e_2 : \widehat{\tau} \text{ chan } r_2 \ \& \ \varphi_2}{\widehat{\Gamma} \vdash_{CA} \text{send } e_1 \text{ on } e_2 : \text{unit} \ \& \ \varphi_1; \varphi_2; r_2! \widehat{\tau}}$
[receive]	$\frac{\widehat{\Gamma} \vdash_{CA} e_0 : \widehat{\tau} \text{ chan } r_0 \ \& \ \varphi_0}{\widehat{\Gamma} \vdash_{CA} \text{receive } e_0 : \widehat{\tau} \ \& \ \varphi_0; r_0? \widehat{\tau}}$
[seq]	$\frac{\widehat{\Gamma} \vdash_{CA} e_1 : \widehat{\tau}_1 \ \& \ \varphi_1 \quad \widehat{\Gamma} \vdash_{CA} e_2 : \widehat{\tau}_2 \ \& \ \varphi_2}{\widehat{\Gamma} \vdash_{CA} e_1; e_2 : \widehat{\tau}_2 \ \& \ \varphi_1; \varphi_2}$
[ch]	$\widehat{\Gamma} \vdash_{CA} ch : \widehat{\tau} \text{ chan } r \ \& \ \Lambda \quad \text{if } \widehat{\tau} \text{ chan } r = \widehat{\Gamma}(ch)$
[sub]	$\frac{\widehat{\Gamma} \vdash_{CA} e : \widehat{\tau} \ \& \ \varphi}{\widehat{\Gamma} \vdash_{CA} e : \widehat{\tau}' \ \& \ \varphi'} \quad \text{if } \widehat{\tau} \leq \widehat{\tau}' \text{ and } \varphi \sqsubseteq \varphi'$
[gen]	$\frac{\widehat{\Gamma} \vdash_{CA} e : \widehat{\tau} \ \& \ \varphi}{\widehat{\Gamma} \vdash_{CA} e : \forall(\zeta_1, \dots, \zeta_n). \widehat{\tau} \ \& \ \varphi}$ if ζ_1, \dots, ζ_n do not occur free in $\widehat{\Gamma}$ and φ
[ins]	$\frac{\widehat{\Gamma} \vdash_{CA} e : \forall(\zeta_1, \dots, \zeta_n). \widehat{\tau} \ \& \ \varphi}{\widehat{\Gamma} \vdash_{CA} e : (\theta \widehat{\tau}) \ \& \ \varphi}$ if θ has $\text{dom}(\theta) \subseteq \{\zeta_1, \dots, \zeta_n\}$

Table 5.17: Communication Analysis (part 2).

and it is covariant in r . The ordering $r \subseteq r'$ means that r is “a subset of” of r' (modulo UCAI) just as what was the case for the effects in Section 5.4. Finally, the ordering $\varphi \sqsubseteq \varphi'$ on behaviours is more complex than before because of the rich structure possessed by behaviours. The definition is given in Table 5.13 and will be explained below. Since the syntactic categories of types and behaviours are mutually recursive also the definitions of $\widehat{\tau} \leq \widehat{\tau}'$ and $\varphi \sqsubseteq \varphi'$ need to be interpreted recursively.

The axiomatisation of $\varphi \sqsubseteq \varphi'$ in Table 5.18 ensures that we obtain a preorder that is a congruence with respect to the operations for combining behaviours. Furthermore, sequencing is an associative operation with Λ as identity and we have a distributive law with respect to choice. It follows that choice is associative and commutative. Next the axioms for recursion allows us to

$\varphi \sqsubseteq \varphi$	$\frac{\varphi_1 \sqsubseteq \varphi_2 \quad \varphi_2 \sqsubseteq \varphi_3}{\varphi_1 \sqsubseteq \varphi_3}$
$\frac{\varphi_1 \sqsubseteq \varphi_2 \quad \varphi_3 \sqsubseteq \varphi_4}{\varphi_1; \varphi_3 \sqsubseteq \varphi_2; \varphi_4}$	$\frac{\varphi_1 \sqsubseteq \varphi_2 \quad \varphi_3 \sqsubseteq \varphi_4}{\varphi_1 + \varphi_3 \sqsubseteq \varphi_2 + \varphi_4}$
$\frac{\varphi_1 \sqsubseteq \varphi_2}{\text{spawn } \varphi_1 \sqsubseteq \text{spawn } \varphi_2}$	$\frac{\varphi_1 \sqsubseteq \varphi_2}{\text{rec}\beta.\varphi_1 \sqsubseteq \text{rec}\beta.\varphi_2}$
$\varphi_1; (\varphi_2; \varphi_3) \sqsubseteq (\varphi_1; \varphi_2); \varphi_3$	$(\varphi_1; \varphi_2); \varphi_3 \sqsubseteq \varphi_1; (\varphi_2; \varphi_3)$
$(\varphi_1 + \varphi_2); \varphi_3 \sqsubseteq (\varphi_1; \varphi_3) + (\varphi_2; \varphi_3)$	$(\varphi_1; \varphi_3) + (\varphi_2; \varphi_3) \sqsubseteq (\varphi_1 + \varphi_2); \varphi_3$
$\varphi \sqsubseteq \Lambda; \varphi \quad \Lambda; \varphi \sqsubseteq \varphi$	$\varphi \sqsubseteq \varphi; \Lambda \quad \varphi; \Lambda \sqsubseteq \varphi$
$\varphi_1 \sqsubseteq \varphi_1 + \varphi_2 \quad \varphi_2 \sqsubseteq \varphi_1 + \varphi_2$	$\varphi + \varphi \sqsubseteq \varphi$
$\text{rec}\beta.\varphi \sqsubseteq \varphi[\beta \mapsto \text{rec}\beta.\varphi]$	$\varphi[\beta \mapsto \text{rec}\beta.\varphi] \sqsubseteq \text{rec}\beta.\varphi$
$\frac{\hat{\tau} \leq \hat{\tau}' \quad \hat{\tau}' \leq \hat{\tau} \quad r \sqsubseteq r'}{\hat{\tau} \text{ chan } r \sqsubseteq \hat{\tau}' \text{ chan } r'}$	
$\frac{r_1 \sqsubseteq r_2 \quad \hat{\tau}_1 \leq \hat{\tau}_2}{r_1! \hat{\tau}_1 \sqsubseteq r_2! \hat{\tau}_2}$	$\frac{r_1 \sqsubseteq r_2 \quad \hat{\tau}_2 \leq \hat{\tau}_1}{r_1? \hat{\tau}_1 \sqsubseteq r_2? \hat{\tau}_2}$

Table 5.18: Ordering on behaviours.

unfold the *rec*-construct. The final three rules clarify how behaviours depend upon types and regions: $\hat{\tau} \text{ chan } r$ is both contravariant and covariant in $\hat{\tau}$ and is covariant in r (just as was the case for the type $\hat{\tau} \text{ chan } r$); $r! \hat{\tau}$ is covariant in both r and $\hat{\tau}$ (because a value is sent) whereas $r? \hat{\tau}$ is covariant in r and contravariant in $\hat{\tau}$ (because a value is received). There is no explicit law for renaming bound behaviour variables as we shall regard $\text{rec}\beta.\varphi$ as being equal to $\text{rec}\beta'.\varphi'$ whenever they are α -equivalent.

The Communication Analysis in Tables 5.16 and 5.17 differs from the Region Inference Analysis of Tables 5.12 and 5.13 in that there is no analogue of the rule [*region*] where an *Observe* function is used to reduce the annotation to what is visible from the outside. The reason is that the Communication Analysis is intended to record *all* the side effects (in the form of communications) that take place during computation.

Example 5.36 Returning to the program of Example 5.35 we can now

verify that the node function can be given the type:

$$('a \xrightarrow{i} 'b) \xrightarrow{\Lambda} ('a \text{ chan } '1) \xrightarrow{\Lambda} ('b \text{ chan } '2) \xrightarrow{\varphi} \text{unit}$$

where $\varphi = \text{spawn}(\text{rec } '2. (('1?'a; '1; '2!'b; '2))$

We use the rules $[fun]$ and $[sub]$ of Tables 5.16 and 5.17 together with the axiom $\varphi[\beta \mapsto \text{rec}\beta.\varphi] \sqsubseteq \text{rec}\beta.\varphi$ of Table 5.18. Then the rule $[gen]$ allows us to obtain the type schema given in Example 5.35.

Turning to the pipe function we first note that it can be given a type of the form

$$(('a \xrightarrow{i} 'a) \text{ list}) \xrightarrow{\Lambda} ('a \text{ chan } ('1 \cup \{C\})) \xrightarrow{\Lambda} ('a \text{ chan } '2) \xrightarrow{\varphi'} \text{unit}$$

where the regions for the input and local channels are “merged” because they can both be used as input channels in a call to `pipe` whereas the region for the output channel is always kept separate. The behaviour φ' is of the form

$$\text{rec } '2. (\text{spawn}(\text{rec } '3. (('1 \cup \{C\})?'a; \Lambda; '2!'a; '3)) \\ + 'a \text{ chan } C; \text{spawn}(\text{rec } '4. (('1 \cup \{C\})?'a; '1; C!'a; '4)); '2)$$

because in the `then`-branch the input channel for node has type `'a chan ('1 \cup {C})` and the output channel has type `'a chan '2`, whereas in the `else`-branch the input channel for node has type `'a chan ('1 \cup {C})` and the output channel has type `'a chan {C}` (as well as `'a chan ('1 \cup {C})`). The rule $[gen]$ allows us to obtain the type schema displayed in Example 5.35. ■

Concluding Remarks

Control Flow Analysis. The literature [37, 38, 54, 11] contains many formulations of non-standard type systems aiming at performing *Control Flow Analysis*. The formulation presented in Section 5.1 uses a particularly simple set of techniques where types are annotated but there is no additional effect component nor is there any coverage of polymorphism or subtyping. Although there is no explicit clause for subeffecting in the manner of Section 5.4, we regard the formulation as a *subeffecting analysis* because the rules for function abstraction allow to increase the annotation on the function arrow in much the same way as is the case in subeffecting (and in much more restricted ways than holds for subtyping; see also Exercise 5.12). References for type systems with subtyping include [43, 42, 81] as well as the more advanced [63, 125, 126] that also deal with polymorphism. To allow a general treatment of subtyping, these papers generally demand constraints to be an explicit part of the inference system unlike what was done in Section 5.1. Indeed, the formulation of Section 5.1 allows only *shape conformant subtyping*,

where the underlying type system does *not* make use of any form of subtyping, and is thus somewhat simpler than *atomic subtyping*, where an ordering is imposed upon base types, and *general subtyping*, where an ordering may be imposed between arbitrary types.

The semantic correctness of the Control Flow Analysis established in Section 5.2 is expressed as a subject reduction result but formulated for a Natural Semantics (rather than a Structural Operational Semantics); this approach to semantic correctness has a rather long history [86, 88, 144]. The *Moore family* result about the set of typings is inspired by the ideas of [86] and is included so as to stress the fundamental rôle played by partial orders in all of the approaches to program analysis considered in this book; it also relates to the study of principal types in type systems.

The development of a syntactically sound and complete algorithm for the Control Flow Analysis in Section 5.3 is based on the ideas in [69, 143, 132, 133, 134]; Mini Project 5.1 is based on [134]. The basic idea is to ensure that the algorithm operates on a *free algebra* by restricting annotations to be annotation variables only (the concept of *simple types*) and by recording a *set of constraints* for the meaning of the annotation variables; in our case this is particularly straightforward because the Control Flow Analysis does not deal with polymorphism. Our development differs somewhat from that of [42, 81] that deal with the more advanced notions of atomic subtyping and general subtyping. A different approach, not studied here, would be to dispense with simple types and constraints and instead use techniques for unifying types in a non-free algebra [124].

Restricted effects. The Type and Effect Systems presented in Section 5.4 all share the important property (also holding for the type system in Section 5.1) that *no type information* is recorded in the effects and that the shape of the type information *cannot be influenced* by the effects. All systems included a proper effect component and thereby illustrated the diversity of effects; some pioneering papers in Type and Effect Systems are [67, 77, 68, 69]. At the same time we illustrated a number of design considerations to be taken into account when devising a Type and Effect System: whether or not to incorporate subeffecting, subtyping, polymorphism, polymorphic recursion, whether or not types are allowed to be influenced by effects (which is not the case in Sections 5.4 and 5.1), and whether or not constraints are an explicit part of the inference system (as is implicitly the case in Subsection 5.4.3). However, it would be incorrect to surmise that the selection of components are inherently linked to the example analysis where they were illustrated. Rather, the techniques needed for semantic correctness and for syntactic soundness and completeness depend heavily on the particular selection of components; some are straightforward to deal with whereas others are beyond state-of-the-art.

The Side Effect Analysis presented in Subsection 5.4.1 illustrated the use

of subeffecting and subtyping, but did not incorporate polymorphism, there were no constraints in the inference system, and the effects did not influence the types. This system is sufficiently simple that semantic soundness may be established using the techniques of Section 5.2. If the rule for subtyping was omitted then also the techniques developed in Section 5.3 would suffice for obtaining a sound and complete inference algorithm. The presence of the rule for subtyping naturally leads to a *two stage* implementation process: first the underlying types are inferred and next the constraints on effects (or annotations) are determined [134, 136]. This works because we restrict ourselves to *shape conformant subtyping* where effects *do not influence* the type information. However, adding polymorphism to this development would dramatically increase the complexity of the development (see below).

The Exception Analysis presented in Subsection 5.4.2 illustrated the use of subeffecting, subtyping and polymorphism, but there were no constraints in the inference system, and the effects did not influence the types. Semantic soundness is a bit more complex than in Section 5.2 because of the polymorphism but the techniques of [135, 132, 133, 8] suffice. For the development of a syntactically sound and complete inference algorithm one may take the two stage approach described above [134, 136]; as before, it works because we restrict ourselves to shape conformant subtyping where effects do not influence the type information. Alternatively, one may use more powerful techniques [143, 132, 133, 97] that even allow to include type information inside effects; this amounts to an extension of the approach of Section 5.3 and will be explained below.

The Region Inference analysis presented in Subsection 5.4.3 illustrated the use of polymorphic recursion as far as the effects are concerned, there were implicitly constraints in the inference system (via the dot notation on function arrows), but still the effects cannot influence the types. The presentation is mainly based on [137] but adapted to the FUN language and the style of presentation used elsewhere in this chapter. To obtain effects that are as small as possible the inference system uses "effect masking" (developed in [77, 132, 133]) for removing internal components of the effect: effect components that only deal with regions that are not externally visible. Semantic correctness of the inference system can be shown using the approach of [138]. For the development of a syntactically sound inference algorithm one may once more take the two stage approach described above; the first stage (ordinary type inference) is standard and the second stage is considered in [136] where algorithm \mathcal{S} generates effect and region variables and algorithm \mathcal{R} deals with the complications due to polymorphic recursion (for effects and regions only). The inference algorithm is proved syntactically sound but is known not to be syntactically complete; indeed, obtaining an algorithm that is syntactically sound as well as complete, seems beyond state-of-the-art.

General effects. *One way to make effects more expressive is to allow type information inside the effects so that the shape of the type information can be influenced by the effects.* This idea occurred already in [143, 132, 133] for an extended Side Effect Analysis making use of polymorphism and subeffecting (but not subtyping); this work attempted to “generalise” previous work based on the idea of expansive expressions and imperative versus applicative type variables [135, 131]. As already indicated, semantic soundness amounts to an extension of the techniques of Section 5.2 as presented in [135, 132, 133, 8].

The two stage approach no longer works for obtaining an inference algorithm because the effects are used to control the shape of the underlying types in the form of which type variables are included in a polymorphic type. This suggests extending the techniques of Section 5.3 in that special care needs to be taken when deciding the variables over which to generalise when constructing a polymorphic type. The main idea is that the algorithm needs to consult the constraints in order to determine a larger set of forbidden variables than those directly occurring in the type environment or the effect; this can be formulated as a *downwards closure* with respect to the constraint set [143, 97] or by taking a *principal solution* of the constraints into account [132, 133].

Adding subtyping to this development dramatically increases the complexity of the development. The integration of atomic subtyping (a generalisation of shape conformant subtyping), polymorphism and subeffecting is done in [97, 103, 8] that establishes semantic soundness and develops an inference algorithm that is proved syntactically sound; extensions of this development incorporate a syntactic completeness result (see the discussion of [6] below). This work went a long way towards integrating the techniques for polymorphism and subeffecting (but no subtyping) from Effect Systems [143, 132, 133] with the techniques for polymorphism and subtyping (but no effects) from Type Systems [63, 125, 126].

Another way to make effects more expressive is to let them contain information about the temporal order and causality of actions, rather than just being an unordered set of possibilities. In Section 5.5 we considered the task of extracting *behaviours* (reminiscent of terms in a process algebra) from program in Concurrent ML by means of a Type and Effect System; here effects (the behaviours) have structure, they may influence the type information, there are no explicit constraints in the inference system (although this is the case in more advanced developments [6]), and there are inference rules for subeffecting and (atomic) subtyping. These ideas first occurred in [91, 92] (not involving polymorphism) and in [100, 101, 7] (involving polymorphism); our presentation in Section 5.5 is mainly based on [100, 101] with ingredients from [91, 92]. We refer to [6] for a comprehensive account of a more ambitious development where the inference system is massaged so as to facilitate devel-

oping a syntactically sound and complete inference algorithm; this includes having explicit constraints in the inference system as is usually the case in type systems that make use of subtyping. An application to the validation of embedded systems is presented in [98].

Other developments. All of the formulations presented in this chapter have had a number of common features: to the extent that polymorphism has been incorporated it has been based on the Hindley/Milner polymorphism also found in Standard ML, there has been no subtyping involved in the underlying type system, and there has been no treatment of conjunction or disjunction types as in [12, 13, 61, 62]. Also all of the formulations have expressed *safety* properties: if a certain point is reached then certain information will hold; *liveness* properties in the form of adding annotations that indicate whether or not functions can be assumed to be total was considered in [93].

Finally, linking up with the development of Chapter 4 on Abstract Interpretation, it is possible to allow annotations to be elements of a complete lattice (that is possibly of finite height as in Monotone Frameworks), and it may be profitable to describe Type and Effect Systems using the framework of Abstract Interpretation [82, 26].

Mini Projects

Mini Project 5.1 A Call-Tracking Analysis

Consider a Type and Effect System for Call-Tracking Analysis: it has judgements of the form

$$\hat{\Gamma} \vdash_{CT} e : \hat{\tau} \ \& \ \varphi$$

where φ denotes the set of functions that may be called during the evaluation of e (and similarly for the annotations on function arrows).

1. Formulate an inference system with subeffecting; next add subtyping and finally add polymorphism.

Next consider the inference system with subeffecting only:

2. Modify the Natural Semantics of Table 5.4 such that semantic correctness can be stated for the analysis and prove that the result holds.
3. Devise an algorithm for Call-Tracking Analysis and prove that it is syntactically sound (and complete).

For the more ambitious: can you also deal with subtyping and/or polymorphism? ■

Mini Project 5.2 Data Structures

As in Mini Project 3.2 we shall now extend the language with more general data structures and consider how to modify the Control Flow Analysis (Table 5.2) so as to track the creation points.

Pairs. To accommodate pairs we extend the syntax as follows:

$$\begin{aligned}\hat{\tau} & ::= \dots \mid \hat{\tau}_1 \times^{\varphi} \hat{\tau}_2 \\ e & ::= \dots \mid \text{Pair}_{\pi}(e_1, e_2) \mid (\text{case } e_0 \text{ of Pair}(x_1, x_2) \Rightarrow e_1)\end{aligned}$$

Here `Pair` is a binary constructor and the corresponding `case`-expression does not need an `or`-component as in Mini Project 3.2. As an example, consider the following program for “sorting” a pair of integers:

```
let srt = fnX x => case x of Pair(y,z) =>
    if y<z then x else PairB(z,y)
in srt(PairA(n,m))
```

Here the pair returned will be constructed at `A` if the value of `n` is smaller than the value of `m` and at `B` otherwise. The overall type is $\text{int} \times^{\{A,B\}} \text{int}$.

1. Modify the Control Flow Analysis of Table 5.2 to track the creation points of pairs.
2. Extend the Natural Semantics of Table 5.4 and augment the proof of semantic correctness (Theorem 5.9).
3. Extend the algorithms \mathcal{W}_{CFA} and \mathcal{U}_{CFA} and augment the proof of syntactic soundness and completeness (Theorems 5.20 and 5.21).

Lists. To accommodate lists we extend the syntax as follows:

$$\begin{aligned}\hat{\tau} & ::= \dots \mid \hat{\tau} \text{ list}^{\varphi} \\ e & ::= \dots \mid \text{Cons}_{\pi}(e_1, e_2) \mid \text{Nil}_{\pi} \mid (\text{case } e_0 \text{ of Cons}(x_1, x_2) \Rightarrow e_1 \text{ or } e_2)\end{aligned}$$

Now perform a similar development as the one you performed for pairs.

For the more ambitious: can you give a more general treatment of algebraic types in the manner of Mini Project 3.2? ■

Mini Project 5.3 A Prototype Implementation

In this mini project we shall implement the Control Flow Analysis considered in Sections 5.1 and 5.3. As implementation language we shall choose a functional language as Standard ML or Haskell. We can then define a suitable data type for FUN expressions as follows:

```

type var      = string
type point    = int
datatype const = Num of int | True | False
datatype exp  = Const of const | Var of var | Fn of point * var * exp
              | Fun of point * var * var * exp | App of exp * exp
              | If of exp * exp * exp | Let of var * exp * exp
              | Op of string * exp * exp

```

Now proceed as follows:

1. Define data types for simple types and simple substitutions and implement the function \mathcal{U}_{CFA} of Table 5.7.
2. Define data types for simple type environments and constraint sets and implement the function \mathcal{W}_{CFA} of Table 5.8.
3. Define data types for types and type environments and implement a function that pretty prints the result in the manner of Subsection 5.3.4: type variables must get instantiated to int and annotation variables to the least solution to the constraints.

Test your implementation on selected examples. ■

Exercises

Exercise 5.1 Consider the following expression:

```

let f = fnx x => x 1;
    g = fny y => y+2;
    h = fnz z => z+3
in (f g) + (f h)

```

Use Table 5.1 (the underlying type system) to obtain a type for this expression; what types do you use for f , g and h ? Next use Table 5.2 (Control Flow Analysis) to obtain the annotated type of this expression; what annotated types do you use for f , g and h ? ■

Exercise 5.2 Consider the following variation of FUN where function definitions explicitly involve type information as in $\text{fn}_\pi x : \tau_x \Rightarrow e_0$ and $\text{fun}_\pi f : (\tau_x \rightarrow \tau_0) x \Rightarrow e_0$. Modify Table 5.1 accordingly and prove that the resulting type system is deterministic: $\Gamma \vdash_{\text{UL}} e : \tau_1$ and $\Gamma \vdash_{\text{UL}} e : \tau_2$ imply $\tau_1 = \tau_2$. ■

Exercise 5.3 Consider the inclusion $\varphi_1 \subseteq \varphi_2$ between effects that was discussed in Subsections 5.4.1 and 5.2.3. Give an axiomatisation of $\varphi_1 \subseteq \varphi_2$ such that $\varphi_1 \subseteq \varphi_2$ holds if and only if the set of elements mentioned in φ_1 is a subset of the set of elements mentioned in φ_2 . ■

Exercise* 5.4 In Example 5.4 we showed how to record the annotated type information in a form close to that considered in Chapter 3. To make this precise suppose that expressions of FUN are simultaneously labelled as in Chapters 3 and 5:

$$\begin{aligned} e &::= t^\ell \\ t &::= c \mid x \mid \text{fn}_\pi x \Rightarrow e_0 \mid \text{fun}_\pi f x \Rightarrow e_0 \mid e_1 e_2 \\ &\quad \mid \text{if } e_0 \text{ then } e_1 \text{ else } e_2 \mid \text{let } x = e_1 \text{ in } e_2 \mid e_1 \text{ op } e_2 \end{aligned}$$

Next modify Table 5.2 so as to define a judgement

$$\tilde{\rho}, \tilde{C}; \hat{\Gamma} \vdash_{\text{CFA}} e : \hat{\tau} \quad (5.2)$$

where the idea is that

- $\tilde{C}(\ell) = \hat{\tau}_\ell$ ensures that all judgements $\tilde{\rho}, \tilde{C}; \hat{\Gamma}' \vdash_{\text{CFA}} t^\ell : \hat{\tau}'$ in (5.2) have $\hat{\tau}' = \hat{\tau}_\ell$, and
- $\tilde{\rho}(x) = \hat{\tau}_x$ ensures that all judgements $\tilde{\rho}, \tilde{C}; \hat{\Gamma}' \vdash_{\text{CFA}} x^\ell : \hat{\tau}'$ in (5.2) have $\hat{\tau}' = \hat{\tau}_x$.

Check that (5.2) holds for the expression in Example 5.4 when we take $\tilde{C}(1) = \hat{\tau}_Y$, $\tilde{C}(2) = \hat{\tau}_Y \xrightarrow{\{X\}} \hat{\tau}_Y$, $\tilde{C}(3) = \text{int}$, $\tilde{C}(4) = \hat{\tau}_Y$, $\tilde{C}(5) = \hat{\tau}_Y$, $\tilde{\rho}(x) = \hat{\tau}_Y$ and $\tilde{\rho}(y) = \text{int}$. ■

Exercise 5.5 Consider adding the following inference rule to Table 5.2 (Control Flow Analysis)

$$\frac{\hat{\Gamma} \vdash_{\text{CFA}} e_1 : \hat{\tau}_2 \xrightarrow{\varphi} \hat{\tau}_0 \quad \hat{\Gamma} \vdash_{\text{CFA}} e_2 : \hat{\tau}_2}{\hat{\Gamma} \vdash_{\text{CFA}} e_1 e_2 : \perp_{\hat{\tau}_0}} \quad \text{if } \varphi = \emptyset$$

where $\perp_{\hat{\tau}_0}$ is the least element of $\widehat{\text{Type}}[\tau_0]$ and $\tau_0 = \lfloor \hat{\tau}_0 \rfloor$. Explain what this rule does and determine whether or not Theorem 5.9 (semantic correctness) continues to hold. ■

Exercise 5.6 In Section 5.2 we equipped FUN with a call-by-value semantics. An alternative would be to use a call-by-name semantics. It can be obtained as a simple modification of the semantics of Table 5.4 by changing rule [app] such that the argument is not evaluated before the substitution takes place and similarly changing rule [let]. Make these changes to Table 5.4 and show that the correctness result (Theorem 5.9) still holds for the analysis of Table 5.2.

What does that tell us about the precision of the analysis? ■

Exercise 5.7 Prove Fact 5.17 (the syntactic soundness and completeness of Robinson unification). ■

Exercise 5.8 Consider the partial ordering $\hat{\tau} \leq \hat{\tau}'$ on the annotated types of Section 5.1 that is defined by:

$$\hat{\tau} \leq \hat{\tau}' \quad \frac{\hat{\tau}'_1 \leq \hat{\tau}_1 \quad \hat{\tau}_2 \leq \hat{\tau}'_2 \quad \varphi \subseteq \varphi'}{\hat{\tau}_1 \xrightarrow{\varphi} \hat{\tau}_2 \leq \hat{\tau}'_1 \xrightarrow{\varphi'} \hat{\tau}'_2}$$

We shall say that this ordering treats $\hat{\tau}_1 \xrightarrow{\varphi} \hat{\tau}_2$ *covariantly* in φ and $\hat{\tau}_2$ but *contravariantly* in $\hat{\tau}_1$; this is in line with the subtype ordering considered in Section 5.4 but differs from the partial ordering $\hat{\tau} \sqsubseteq \hat{\tau}'$ considered in Section 5.2.

Show that $(\widehat{\text{Type}}[\tau], \leq)$ is a complete lattice for all choices of the underlying type $\tau \in \text{Type}$. Next investigate whether or not an analogue of Proposition 5.12 holds for this ordering.

Finally reconsider the decision to let $\theta_T''(\alpha)$ in Subsection 5.3.4 be the least element of $(\widehat{\text{Type}}[\tau], \sqsubseteq)$; would it be preferable to let $\theta_T''(\alpha)$ be the least element of $(\widehat{\text{Type}}[\tau], \leq)$ or $(\widehat{\text{Type}}[\tau], \geq)$? ■

Exercise* 5.9 Suppose that

$$\mathcal{W}_{\text{UL}}([\], \text{fun}_F f x \Rightarrow e_0) = (\alpha_x \rightarrow \alpha_0, \theta)$$

where α_x and α_0 are *distinct* type variables. Let e be an arbitrary correctly typed closed expression, i.e. $[\] \vdash_{\text{UL}} e : \tau$ for some τ and show that the call

$$(\text{fun}_F f x \Rightarrow e_0) e$$

cannot terminate. (Hint: use Fact 5.6, Theorems 5.9, 5.20 and 5.21 and that \mathcal{W}_{UL} is syntactically sound.) ■

Exercise 5.10 Formulate what it means for the Side Effect Analysis of Table 5.9 to be semantically correct; this involves modifying the Natural Semantics of Table 5.4 to deal with the store and to record the side effects. (Proving the result would require quite some work.) ■

Exercise 5.11 Suppose that the language of Subsection 5.4.1 has a call-by-name semantics rather than a call-by-value semantics. Modify the Side Effect Analysis of Table 5.9 accordingly. ■

Exercise 5.12 We shall now illustrate the concept of *proof normalisation* for the Side Effect Analysis of Subsection 5.4.1; for this we shall assume that Table 5.9 does *not* include the combined rule for subeffecting and subtyping but *only* the rule for subeffecting.

For this system one can dispense with an explicit rule for subeffecting by integrating its effects into all other rules; this can be done by adding a “ $\cup\varphi$ ” to all effects occurring in the conclusions of all axioms and rules. Do this and argue that exactly the same judgements are provable in the two systems.

A further variation is only to incorporate “ $\cup\varphi$ ” where it is really needed: in all axioms and in the rules for function abstraction. Do this and argue that once more exactly the same judgements are provable in the two systems.

What you have performed amounts to *proof normalisation*: whenever one has an inference system with a number of syntax directed rules and axioms and at least one rule that is not syntax directed, it is frequently possible to restrict the use of the non syntax directed rules. In this way the structure of the inference trees comes closer to the structure of the syntax trees and this is often helpful for proving semantic correctness and is a useful starting point for developing inference algorithms. ■

Exercise 5.13 Consider the rule [*handle*] in the Exception Analysis of Table 5.10. Would the analysis still be semantically correct if we replaced it by the following two rules:

$$\frac{\widehat{\Gamma} \vdash_{\text{ES}} e_1 : \widehat{\tau} \ \& \ \varphi_1 \quad \widehat{\Gamma} \vdash_{\text{ES}} e_2 : \widehat{\tau} \ \& \ \varphi_2}{\widehat{\Gamma} \vdash_{\text{ES}} \text{handle } s \text{ as } e_1 \text{ in } e_2 : \widehat{\tau} \ \& \ \varphi_2}$$

if $s \notin \varphi_2$ and $AV(\varphi_2) = \emptyset$

$$\frac{\widehat{\Gamma} \vdash_{\text{ES}} e_1 : \widehat{\tau} \ \& \ \varphi_1 \quad \widehat{\Gamma} \vdash_{\text{ES}} e_2 : \widehat{\tau} \ \& \ \varphi_2}{\widehat{\Gamma} \vdash_{\text{ES}} \text{handle } s \text{ as } e_1 \text{ in } e_2 : \widehat{\tau} \ \& \ \varphi_1 \cup (\varphi_2 \setminus \{s\})}$$

if $s \in \varphi_2$ or $AV(\varphi_2) \neq \emptyset$

Here $s \in \varphi_2$ means $\{s\} \subseteq \varphi_2$ and $AV(\varphi_2)$ is the set of annotation variables in φ_2 . ■

Exercise 5.14 Consider the Exception Analysis of Table 5.10 and change the rule [*let*] to

$$\frac{\widehat{\Gamma} \vdash_{\text{ES}} e_1 : \widehat{\sigma}_1 \ \& \ \varphi_1 \quad \widehat{\Gamma}[x \mapsto \widehat{\sigma}_1] \vdash_{\text{ES}} e_2 : \widehat{\sigma}_2 \ \& \ \varphi_2}{\widehat{\Gamma} \vdash_{\text{ES}} \text{let } x = e_1 \text{ in } e_2 : \widehat{\sigma}_2 \ \& \ \varphi_1 \cup \varphi_2}$$

and perform similar changes in *[if]*, *[raise]*, *[handle]* and *[sub]*; do not forget to define the meaning of $\hat{\tau} \leq \hat{\tau}'$. Clearly the new system is at least as powerful as Table 5.10 but is it more powerful? (Hint: Consider the places where *[gen]* and *[ins]* can be used in Table 5.10.) ■

Chapter 6

Algorithms

(This material remains to be integrated.)

Appendix A

Partially Ordered Sets

Partially ordered sets and complete lattices play a crucial role in program analysis and in this appendix we shall summarise some of their properties. We review the basic approaches for how to construct complete lattices from other complete lattices and state the central properties of partially ordered sets satisfying the Ascending Chain and Descending Chain Conditions. We then review the classical results about least and greatest fixed points.

A.1 Basic Definitions

Partially ordered set. A *partial ordering* is a relation $\sqsubseteq: L \times L \rightarrow \{\text{true}, \text{false}\}$ that is reflexive (i.e. $\forall l : l \sqsubseteq l$), transitive (i.e. $\forall l_1, l_2, l_3 : l_1 \sqsubseteq l_2 \wedge l_2 \sqsubseteq l_3 \Rightarrow l_1 \sqsubseteq l_3$), and anti-symmetric (i.e. $\forall l_1, l_2 : l_1 \sqsubseteq l_2 \wedge l_2 \sqsubseteq l_1 \Rightarrow l_1 = l_2$). A *partially ordered set* (L, \sqsubseteq) is a set L equipped with a partial ordering \sqsubseteq (sometimes written \sqsubseteq_L). We shall write $l_2 \supseteq l_1$ for $l_1 \sqsubseteq l_2$ and $l_1 \subset l_2$ for $l_1 \sqsubseteq l_2 \wedge l_1 \neq l_2$.

A subset Y of L has $l \in L$ as an *upper bound* if $\forall l' \in Y : l' \sqsubseteq l$ and as a *lower bound* if $\forall l' \in Y : l' \supseteq l$. A *least upper bound* l of Y is an upper bound of Y that satisfies $l \sqsubseteq l_0$ whenever l_0 is another upper bound of Y ; similarly, a *greatest lower bound* l of Y is a lower bound of Y that satisfies $l_0 \sqsubseteq l$ whenever l_0 is another lower bound of Y . Note that subsets Y of a partially ordered set L need not have least upper bounds nor greatest lower bounds but when they exist they are unique (since \sqsubseteq is anti-symmetric) and they are denoted $\bigsqcup Y$ and $\bigsqcap Y$, respectively. Sometimes \bigsqcup is called the *join operator* and \bigsqcap the *meet operator* and we shall write $l_1 \sqcup l_2$ for $\bigsqcup\{l_1, l_2\}$ and similarly $l_1 \sqcap l_2$ for $\bigsqcap\{l_1, l_2\}$.

Complete lattice. A *complete lattice* $L = (L, \sqsubseteq) = (L, \sqsubseteq, \bigsqcup, \bigsqcap, \perp, \top)$ is a partially ordered set (L, \sqsubseteq) such that all subsets have least upper bounds




Figure A.1: Two complete lattices.

as well as greatest lower bounds. Furthermore, $\perp = \bigsqcup \emptyset = \bigsqcap L$ is the *least element* and $\top = \bigsqcap \emptyset = \bigsqcup L$ is the *greatest element*.

Example A.1 If $L = (\mathcal{P}(S), \subseteq)$ for some set S then \sqsubseteq is \subseteq and $\bigsqcup Y = \bigcup Y$, $\bigsqcap Y = \bigcap Y$, $\perp = \emptyset$ and $\top = S$. If $L = (\mathcal{P}(S), \supseteq)$ then \sqsubseteq is \supseteq and $\bigsqcup Y = \bigcap Y$, $\bigsqcap Y = \bigcup Y$, $\perp = S$ and $\top = \emptyset$.

Hence $(\mathcal{P}(S), \subseteq)$ as well as $(\mathcal{P}(S), \supseteq)$ are complete lattices. In the case where $S = \{1, 2, 3\}$ the two complete lattices are shown on Figure A.1; these drawings are often called Hasse diagrams. Here a line “going upwards” from some l_1 to some l_2 means that $l_1 \subseteq l_2$; we do not draw lines that follow from reflexivity or transitivity of the partial ordering. ■

Lemma A.2 For a partially ordered set $L = (L, \sqsubseteq)$ the claims

- (i) L is a complete lattice,
- (ii) every subset of L has a least upper bound, and
- (iii) every subset of L has a greatest lower bound

are equivalent. ■

Proof Clearly (i) implies (ii) and (iii). To show that (ii) implies (i) let $Y \subseteq L$ and define

$$\bigsqcap Y = \bigsqcup \{l \in L \mid \forall l' \in Y : l \sqsubseteq l'\} \quad (\text{A.1})$$

and let us prove that this indeed defines a greatest lower bound. All the elements of the set on the right hand side of (A.1) are lower bounds of Y so clearly (A.1) defines a lower bound of Y . Since any lower bound of Y will be in the set it follows that (A.1) defines the greatest lower bound of Y . Thus (i) holds.

To show that (iii) implies (i) we define $\bigsqcup Y = \bigsqcap \{l \in L \mid \forall l' \in Y : l' \sqsubseteq l\}$. Arguments analogous to those above show that this defines a least upper bound and that (i) holds. ■

Moore family. A *Moore family* is a subset Y of a complete lattice $L = (L, \sqsubseteq)$ that is closed under greatest lower bounds: $\forall Y' \subseteq Y : \bigsqcap Y' \in Y$. It follows that a Moore family always contains a least element, $\bigsqcap Y$, and a greatest element, $\bigsqcap \emptyset$, which equals the greatest element, \top , from L ; in particular, a Moore family is never empty.

Example A.3 Consider the complete lattice $(\mathcal{P}(S), \subseteq)$ of Figure A.1 (a). The subsets

$$\{\{2\}, \{1, 2\}, \{2, 3\}, \{1, 2, 3\}\} \text{ and } \{\emptyset, \{1, 2, 3\}\}$$

are both Moore families, whereas none of

$$\{\{1\}, \{2\}\} \text{ and } \{\emptyset, \{1\}, \{2\}, \{1, 2\}\}$$

are. ■

Properties of functions. A function $f : L_1 \rightarrow L_2$ between partially ordered sets $L_1 = (L_1, \sqsubseteq_1)$ and $L_2 = (L_2, \sqsubseteq_2)$ is *surjective* (or *onto* or *epic*) if

$$\forall l_2 \in L_2 : \exists l_1 \in L_1 : f(l_1) = l_2$$

and it is *injective* (or 1-1 or *monic*) if

$$\forall l, l' \in L_1 : f(l) = f(l') \Rightarrow l = l'$$

The function f is *monotone* (or *isotone* or *order-preserving*) if

$$\forall l, l' \in L_1 : l \sqsubseteq_1 l' \Rightarrow f(l) \sqsubseteq_2 f(l')$$

It is an *additive function* (or a *join morphism*, sometimes called a *distributive function*) if

$$\forall l_1, l_2 \in L_1 : f(l_1 \sqcup l_2) = f(l_1) \sqcup f(l_2)$$

and it is called a *multiplicative function* (or a *meet morphism*) if

$$\forall l_1, l_2 \in L_1 : f(l_1 \sqcap l_2) = f(l_1) \sqcap f(l_2)$$

The function f is a *completely additive function* (or a *complete join morphism*) if for all $Y \subseteq L_1$:

$$f(\bigsqcup_1 Y) = \bigsqcup_2 \{f(l') \mid l' \in Y\} \text{ whenever } \bigsqcup_1 Y \text{ exists}$$

and it is *completely multiplicative* (or a *complete meet morphism*) if for all $Y \subseteq L_1$:

$$f(\prod_1 Y) = \prod_2 \{f(l') \mid l' \in Y\} \text{ whenever } \prod_1 Y \text{ exists}$$

Clearly $\prod_1 Y$ and $\prod_1 Y$ always exist when L_1 is a complete lattice; when L_2 is not a complete lattice the above statements also require the appropriate least upper bounds and greatest lower bounds to exist in L_2 . The function f is *affine* if for all *non-empty* $Y \subseteq L_1$

$$f(\bigsqcup_1 Y) = \bigsqcup_2 \{f(l') \mid l' \in Y\} \text{ whenever } \bigsqcup_1 Y \text{ exists (and } Y \neq \emptyset)$$

and it is *strict* if $f(\perp_1) = \perp_2$; note that a function is completely additive if and only if it is both affine and strict.

Lemma A.4 If $L = (L, \sqsubseteq, \sqcup, \sqcap, \perp, \top)$ and $M = (M, \sqsubseteq, \sqcup, \sqcap, \perp, \top)$ are complete lattices and M is finite then the three conditions

- (i) $\gamma : M \rightarrow L$ is monotone,
- (ii) $\gamma(\top) = \top$, and
- (iii) $\gamma(m_1 \sqcap m_2) = \gamma(m_1) \sqcap \gamma(m_2)$ whenever $m_1 \not\sqsubseteq m_2 \wedge m_2 \not\sqsubseteq m_1$

are jointly equivalent to $\gamma : M \rightarrow L$ being completely multiplicative. ■

Proof First note that if γ is completely multiplicative then (i), (ii) and (iii) hold. For the converse note that by monotonicity of γ we have $\gamma(m_1 \sqcap m_2) = \gamma(m_1) \sqcap \gamma(m_2)$ also when $m_1 \sqsubseteq m_2 \vee m_2 \sqsubseteq m_1$. We then prove by induction on the (finite) cardinality of $M' \subseteq M$ that:

$$\gamma(\prod M') = \prod \{\gamma(m) \mid m \in M'\} \quad (\text{A.2})$$

If the cardinality of M' is 0 then (A.2) follows from (ii). If the cardinality of M' is larger than 0 then we write $M' = M'' \cup \{m''\}$ where $m'' \notin M''$; this ensures that the cardinality of M'' is strictly less than that of M' ; hence:

$$\begin{aligned} \gamma(\prod M') &= \gamma((\prod M'') \sqcap m'') \\ &= \gamma(\prod M'') \sqcap \gamma(m'') \\ &= (\prod \{\gamma(m) \mid m \in M''\}) \sqcap \gamma(m'') \\ &= \prod \{\gamma(m) \mid m \in M'\} \end{aligned}$$

This proves the result. ■

Lemma A.5 A function $f : (\mathcal{P}(D), \subseteq) \rightarrow (\mathcal{P}(E), \subseteq)$ is *affine* if and only if there exists a function $\varphi : D \rightarrow \mathcal{P}(E)$ and an element $\varphi_\emptyset \in \mathcal{P}(E)$ such that

$$f(Y) = \bigcup \{\varphi(d) \mid d \in Y\} \cup \varphi_\emptyset$$

The function f is completely additive if and only if additionally $\varphi_\emptyset = \emptyset$. ■

Proof Suppose that f is of the form displayed and let \mathcal{Y} be a non-empty set; Then

$$\begin{aligned} \bigcup\{f(Y) \mid Y \in \mathcal{Y}\} &= \bigcup\{\bigcup\{\varphi(d) \mid d \in Y\} \cup \varphi_\emptyset \mid Y \in \mathcal{Y}\} \\ &= \bigcup\{\bigcup\{\varphi(d) \mid d \in Y\} \mid Y \in \mathcal{Y}\} \cup \varphi_\emptyset \\ &= \bigcup\{\varphi(d) \mid d \in \bigcup\mathcal{Y}\} \cup \varphi_\emptyset \\ &= f(\bigcup\mathcal{Y}) \end{aligned}$$

showing that f is affine.

Next suppose that f is affine and define $\varphi(d) = f(\{d\})$ and $\varphi_\emptyset = f(\emptyset)$. For $Y \in \mathcal{P}(D)$ let $\mathcal{Y} = \{\{d\} \mid d \in Y\} \cup \{\emptyset\}$ and note that $Y = \bigcup\mathcal{Y}$ and $\mathcal{Y} \neq \emptyset$. Then

$$\begin{aligned} f(Y) &= f(\bigcup\mathcal{Y}) \\ &= \bigcup\{f(\{d\}) \mid d \in Y\} \cup \{f(\emptyset)\} \\ &= \bigcup\{\varphi(d) \mid d \in Y\} \cup \{\varphi_\emptyset\} \\ &= \bigcup\{\varphi(d) \mid d \in Y\} \cup \varphi_\emptyset \end{aligned}$$

so f can be written in the required form. The additional statement about completely additivity is straightforward. ■

An *isomorphism* from a partially ordered set (L_1, \sqsubseteq_1) to a partially ordered set (L_2, \sqsubseteq_2) is a *monotone* function $\theta : L_1 \rightarrow L_2$ such that there exists a (necessarily unique) monotone function $\theta^{-1} : L_2 \rightarrow L_1$ with $\theta \circ \theta^{-1} = id_2$ and $\theta^{-1} \circ \theta = id_1$ (where id_i is the identity function over L_i , $i = 1, 2$).

A.2 Construction of Complete Lattices

Complete lattices can be combined to construct new complete lattices. We shall first see how to construct products and then two kinds of function spaces.

Cartesian product. Let $L_1 = (L_1, \sqsubseteq_1)$ and $L_2 = (L_2, \sqsubseteq_2)$ be partially ordered sets. Define $L = (L, \sqsubseteq)$ by

$$L = \{(l_1, l_2) \mid l_1 \in L_1 \wedge l_2 \in L_2\}$$

and

$$(l_{11}, l_{21}) \sqsubseteq (l_{12}, l_{22}) \text{ iff } l_{11} \sqsubseteq_1 l_{12} \wedge l_{21} \sqsubseteq_2 l_{22}$$

It is then straightforward to verify that L is a partially ordered set. If additionally each $L_i = (L_i, \sqsubseteq_i, \bigsqcup_i, \bigsqcap_i, \perp_i, \top_i)$ is a complete lattice then so is $L = (L, \sqsubseteq, \bigsqcup, \bigsqcap, \perp, \top)$ and furthermore

$$\bigsqcup Y = (\bigsqcup_1 \{l_1 \mid \exists l_2 : (l_1, l_2) \in Y\} , \bigsqcup_2 \{l_2 \mid \exists l_1 : (l_1, l_2) \in Y\})$$

and $\perp = (\perp_1, \perp_2)$ and similarly for $\prod Y$ and \top . We often write $L_1 \times L_2$ for L and call it the *cartesian product* of L_1 and L_2 .

A variant of the cartesian product called the *smash product* is obtained if we require that all the pairs (l_1, l_2) of the lattice satisfy $l_1 = \perp_1 \Leftrightarrow l_2 = \perp_2$.

Total function space. Let $L_1 = (L_1, \sqsubseteq_1)$ be a partially ordered set and let S be a set. Define $L = (L, \sqsubseteq)$ by

$$L = \{f : S \rightarrow L_1 \mid f \text{ is a total function}\}$$

and

$$f \sqsubseteq f' \text{ iff } \forall s \in S : f(s) \sqsubseteq_1 f'(s)$$

It is then straightforward to verify that L is a partially ordered set. If additionally $L_1 = (L_1, \sqsubseteq_1, \bigsqcup_1, \prod_1, \perp_1, \top_1)$ is a complete lattice then so is $L = (L, \sqsubseteq, \bigsqcup, \prod, \perp, \top)$ and furthermore

$$\bigsqcup Y = \lambda s. \bigsqcup_1 \{f(s) \mid f \in Y\}$$

and $\perp = \lambda s. \perp_1$ and similarly for $\prod Y$ and \top . We often write $S \rightarrow L_1$ for L and call it the *total function space* from S to L_1 .

Monotone function space. Again let $L_1 = (L_1, \sqsubseteq_1)$ and $L_2 = (L_2, \sqsubseteq_2)$ be partially ordered sets. Now define $L = (L, \sqsubseteq)$ by

$$L = \{f : L_1 \rightarrow L_2 \mid f \text{ is a monotone function}\}$$

and

$$f \sqsubseteq f' \text{ iff } \forall l_1 \in L_1 : f(l_1) \sqsubseteq_2 f'(l_1)$$

It is then straightforward to verify that L is a partially ordered set. If additionally each $L_i = (L_i, \sqsubseteq_i, \bigsqcup_i, \prod_i, \perp_i, \top_i)$ is a complete lattice then so is $L = (L, \sqsubseteq, \bigsqcup, \prod, \perp, \top)$ and furthermore

$$\bigsqcup Y = \lambda l_1. \bigsqcup_2 \{f(l_1) \mid f \in Y\}$$

and $\perp = \lambda l_1. \perp_2$ and similarly for $\prod Y$ and \top . We often write $L_1 \rightarrow L_2$ for L and call it the *monotone function space* from L_1 to L_2 .

A.3 Chains

The ordering \sqsubseteq on a complete lattice $L = (L, \sqsubseteq)$ expresses when one property is better (or more precise) than another property. When performing a program analysis we will typically construct a sequence of elements in L and it is the general properties of such sequences that we shall study now. In the next section we will be more explicit and consider the sequences obtained during a fixed point computation.

Chains. A subset $Y \subseteq L$ of a partially ordered set $L = (L, \subseteq)$ is a *chain* if

$$\forall l_1, l_2 \in Y : (l_1 \subseteq l_2) \vee (l_2 \subseteq l_1)$$

Thus a chain is a (possibly empty) subset of L that is totally ordered. We shall say that it is a *finite chain* if it is a finite subset of L .

A sequence $(l_n)_n = (l_n)_{n \in \mathbf{N}}$ of elements in L is an *ascending chain* if

$$n \leq m \Rightarrow l_n \subseteq l_m$$

Writing $(l_n)_n$ also for $\{l_n \mid n \in \mathbf{N}\}$ it is clear that an ascending chain also is a chain. Similarly, a sequence $(l_n)_n$ is a *descending chain* if

$$n \leq m \Rightarrow l_n \supseteq l_m$$

and clearly a descending chain is also a chain.

We shall say that a sequence $(l_n)_n$ *eventually stabilises* if and only if

$$\exists n_0 \in \mathbf{N} : \forall n \in \mathbf{N} : n \geq n_0 \Rightarrow l_n = l_{n_0}$$

For the sequence $(l_n)_n$ we write $\bigsqcup_n l_n$ for $\bigsqcup \{l_n \mid n \in \mathbf{N}\}$ and similarly we write $\bigsqcap_n l_n$ for $\bigsqcap \{l_n \mid n \in \mathbf{N}\}$.

Ascending Chain and Descending Chain Conditions. We shall say that a partially ordered set $L = (L, \subseteq)$ has *finite height* if and only if all chains are finite. It has finite height *at most* h if all chains contain at most $h + 1$ elements; it has finite height h if additionally there is a chain with $h + 1$ elements. The partially ordered set L satisfies the *Ascending Chain Condition* if and only if all ascending chains eventually stabilise. Similarly, it satisfies the *Descending Chain Condition* if and only if all descending chains eventually stabilise. These concepts are related as follows:

Lemma A.6 A partially ordered set $L = (L, \subseteq)$ has finite height if and only if it satisfies both the Ascending and Descending Chain Conditions. ■

Proof First assume that L has finite height. If $(l_n)_n$ is an ascending chain then it must be a finite chain and hence eventually stabilise; thus L satisfies the Ascending Chain Condition. In a similar way it is shown that L satisfies the Descending Chain Condition.

Next assume that L satisfies the Ascending Chain Condition as well as the Descending Chain Condition and consider a chain $Y \subseteq L$. We shall prove that Y is a finite chain. This is obvious if Y is empty so assume that it is not. Then also (Y, \subseteq) is a non-empty partially ordered set satisfying the Ascending and Descending Chain Conditions.

As an auxiliary result we shall now show that

$$\text{each non-empty subset } Y' \text{ of } Y \text{ contains a least element} \quad (\text{A.3})$$




Figure A.2: Two partially ordered sets.

To see this we shall construct a descending chain $(l'_n)_n$ in Y' as follows: first let l'_0 be an arbitrary element of Y' . For the inductive step let $l'_{n+1} = l'_n$ if l'_n is the least element of Y' ; otherwise we can find $l'_{n+1} \in Y'$ such that $l'_{n+1} \sqsubseteq l'_n \wedge l'_{n+1} \neq l'_n$. Clearly $(l'_n)_n$ is a descending chain in Y' ; since Y satisfies the Descending Chain Condition the chain will eventually stabilise, i.e. $\exists n'_0 : \forall n \geq n'_0 : l'_n = l'_{n'_0}$ and the construction is such that $l'_{n'_0}$ is the least element of Y' .

Returning to the main proof obligation we shall now construct an ascending chain $(l_n)_n$ in Y . Using (A.3) each l_n is chosen as the least element of the set $Y \setminus \{l_0, \dots, l_{n-1}\}$ as long as the latter set is non-empty, and this yields $l_{n-1} \sqsubseteq l_n \wedge l_{n-1} \neq l_n$; when $Y \setminus \{l_0, \dots, l_{n-1}\}$ is empty we set $l_n = l_{n-1}$, and since Y is non-empty we know that $n > 0$. Thus we have an ascending chain in Y and using the Ascending Chain Condition we have $\exists n_0 : \forall n \geq n_0 : l_n = l_{n_0}$. But this means that $Y \setminus \{l_0, \dots, l_{n_0}\} = \emptyset$ since this is the only way we can achieve that $l_{n_0+1} = l_{n_0}$. It follows that Y is finite. ■

Example A.7 The partially ordered set of Figure A.2 (a) satisfies the Ascending Chain Condition but does not have finite height; the partially ordered set of Figure A.2 (b) satisfies the Descending Chain Condition but does not have finite height. ■

One can show that each of the three conditions finite height, ascending chain, and descending chain, is preserved under the construction of cartesian product: if L_1 and L_2 satisfies one of the conditions then $L_1 \times L_2$ will also satisfy that condition. The construction of total function spaces $S \rightarrow L$ only preserves the conditions of L if S is finite and the construction of monotone function spaces $L_1 \rightarrow L_2$ does not in general preserve the conditions.

An alternative characterisation of complete lattices satisfying the Ascending Chain Condition is given by the following result:

Lemma A.8 For a partially ordered set $L = (L, \sqsubseteq)$ the conditions

- (i) L is a complete lattice satisfying the Ascending Chain Condition, and
- (ii) L has a least element, \perp , and binary least upper bounds and satisfies the Ascending Chain Condition

are equivalent. ■

Proof It is immediate that (i) implies (ii) so let us prove that (ii) implies (i). Using Lemma A.6 it suffices to prove that all subsets Y of L have a least upper bound $\bigsqcup Y$. If Y is empty clearly $\bigsqcup Y = \perp$. If Y is finite and non-empty then we can write $Y = \{y_1, \dots, y_n\}$ for $n \geq 1$ and it follows that $\bigsqcup Y = (\dots(y_1 \sqcup y_2) \sqcup \dots) \sqcup y_n$.

If Y is infinite then we construct a sequence $(l_n)_n$ of elements of L : let l_0 be an arbitrary element y_0 of Y and given l_n take $l_{n+1} = l_n$ in the case where $\forall y \in Y : y \sqsubseteq l_n$ and take $l_{n+1} = l_n \sqcup y_{n+1}$ in the case where some $y_{n+1} \in Y$ satisfies $y_{n+1} \not\sqsubseteq l_n$. Clearly this sequence is an ascending chain. Since L satisfies the Ascending Chain Condition it follows that the chain eventually stabilises, i.e. there exists n such that $l_n = l_{n+1} = \dots$. This means that $\forall y \in Y : y \sqsubseteq l_n$ because if $y \not\sqsubseteq l_n$ then $l_n \neq l_n \sqcup y$ and we have a contradiction. So we have constructed an upper bound for Y . Since it is actually the least upper bound of the subset $\{y_0, \dots, y_n\}$ of Y it follows that it is also the least upper bound of Y . ■

A related result is the following:

Lemma A.9 For a complete lattice $L = (L, \sqsubseteq)$ satisfying the Ascending Chain Condition and a total function $f : L \rightarrow L$, the conditions

- (i) f is additive, i.e. $\forall l_1, l_2 : f(l_1 \sqcup l_2) = f(l_1) \sqcup f(l_2)$, and
- (ii) f is affine, i.e. $\forall Y \subseteq L, Y \neq \emptyset : f(\bigsqcup Y) = \bigsqcup \{f(l) \mid l \in Y\}$

are equivalent and in both cases f is a monotone function. ■

Proof It is immediate that (ii) implies (i): take $Y = \{l_1, l_2\}$. It is also immediate that (i) implies that f is monotone since $l_1 \sqsubseteq l_2$ is equivalent to $l_1 \sqcup l_2 = l_2$.

Next suppose that f satisfies (i) and let us prove (ii). If Y is finite we can write $Y = \{y_1, \dots, y_n\}$ for $n \geq 1$ and

$$f(\bigsqcup Y) = f(y_1 \sqcup \dots \sqcup y_n) = f(y_1) \sqcup \dots \sqcup f(y_n) \sqsubseteq \bigsqcup \{f(l) \mid l \in Y\}$$

If Y is infinite then the construction of the proof of Lemma A.8 gives $\bigsqcup Y = l_n$ and $l_n = y_n \sqcup \dots \sqcup y_0$ for some $y_i \in Y$ and $0 \leq i \leq n$. We then have

$$f(\bigsqcup Y) = f(l_n) = f(y_n \sqcup \dots \sqcup y_0) = f(y_n) \sqcup \dots \sqcup f(y_0) \sqsubseteq \bigsqcup \{f(l) \mid l \in Y\}$$

Furthermore

$$f(\bigsqcup Y) \supseteq \bigsqcup \{f(l) \mid l \in Y\}$$

follows from the monotonicity of f . This completes the proof. ■

A.4 Fixed Points

Reductive and extensive functions. Consider a monotone function $f : L \rightarrow L$ on a complete lattice $L = (L, \sqsubseteq, \sqcup, \sqcap, \perp, \top)$. A *fixed point* of f is an element $l \in L$ such that $f(l) = l$ and we write

$$\text{Fix}(f) = \{l \mid f(l) = l\}$$

for the set of fixed points. The function f is *reductive at l* if and only if $f(l) \sqsubseteq l$ and we write

$$\text{Red}(f) = \{l \mid f(l) \sqsubseteq l\}$$

for the set of elements upon which f is reductive; we shall say that f itself is *reductive* if $\text{Red}(f) = L$. Similarly, the function f is *extensive at l* if and only if $f(l) \sqsupseteq l$ and we write

$$\text{Ext}(f) = \{l \mid f(l) \sqsupseteq l\}$$

for the set of elements upon which f is extensive; we shall say that f itself is *extensive* if $\text{Ext}(f) = L$.

Since L is a complete lattice it is always the case that the set $\text{Fix}(f)$ will have a greatest lower bound in L and we denote it by $\text{lfp}(f)$:

$$\text{lfp}(f) = \bigsqcap \text{Fix}(f)$$

Similarly, the set $\text{Fix}(f)$ will have a least upper bound in L and we denote it by $\text{gfp}(f)$:

$$\text{gfp}(f) = \bigsqcup \text{Fix}(f)$$

We then have the following result, known as *Tarski's Fixed Point Theorem*, showing that $\text{lfp}(f)$ is the *least fixed point* of f and that $\text{gfp}(f)$ is the *greatest fixed point* of f :


Proposition A.10

Let $L = (L, \sqsubseteq, \sqcup, \sqcap, \perp, \top)$ be a complete lattice. If $f : L \rightarrow L$ is a monotone function then $\text{lfp}(f)$ and $\text{gfp}(f)$ satisfy:

$$\begin{aligned} \text{lfp}(f) &= \bigsqcap \text{Red}(f) \in \text{Fix}(f) \\ \text{gfp}(f) &= \bigsqcup \text{Ext}(f) \in \text{Fix}(f) \end{aligned}$$

Proof To prove the claim for $\text{lfp}(f)$ we define $l_0 = \bigsqcap \text{Red}(f)$. We shall first show that $f(l_0) \sqsubseteq l_0$ so that $l_0 \in \text{Red}(f)$. Since $l_0 \sqsubseteq l$ for all $l \in \text{Red}(f)$ and f is monotone we have

$$f(l_0) \sqsubseteq f(l) \sqsubseteq l \text{ for all } l \in \text{Red}(f)$$

Figure A.3: Fixed points of f .

and hence $f(l_0) \sqsubseteq l_0$. To prove that $l_0 \sqsubseteq f(l_0)$ we observe that $f(f(l_0)) \sqsubseteq f(l_0)$ showing that $f(l_0) \in \text{Red}(f)$ and hence $l_0 \sqsubseteq f(l_0)$ by definition of l_0 . Together this shows that l_0 is a fixed point of f so $l_0 \in \text{Fix}(f)$. To see that l_0 is least in $\text{Fix}(f)$ simply note that $\text{Fix}(f) \subseteq \text{Red}(f)$. It follows that $\text{lfp}(f) = l_0$.

The claim for $\text{gfp}(f)$ is proved in a similar way. ■

In denotational semantics it is customary to iterate to the least fixed point by taking the least upper bound of the sequence $(f^n(\perp))_n$. However, we have not imposed any continuity requirements on f (e.g. that $f(\bigsqcup_n l_n) = \bigsqcup_n (f(l_n))$) for all ascending chains $(l_n)_n$ and consequently we cannot be sure to actually reach the fixed point. In a similar way one could consider the greatest lower bound of the sequence $(f^n(\top))_n$. One can show that

$$\begin{aligned} \perp \sqsubseteq f^n(\perp) \sqsubseteq \bigsqcup_n f^n(\perp) &\sqsubseteq \text{lfp}(f) \\ &\sqsubseteq \text{gfp}(f) \sqsubseteq \prod_n f^n(\top) \sqsubseteq f^n(\top) \sqsubseteq \top \end{aligned}$$

as is illustrated in Figure A.3; indeed all inequalities (i.e. \sqsubseteq) can be strict (i.e. \sqsubset). However, if L satisfies the Ascending Chain Condition then there exists n such that $f^n(\perp) = f^{n+1}(\perp)$ and hence $\text{lfp}(f) = f^n(\perp)$. (Indeed any monotone function f over a partially ordered set satisfying the Ascending Chain Condition is also continuous.) Similarly, if L satisfies the Descending Chain Condition then there exists n such that $f^n(\top) = f^{n+1}(\top)$ and hence

$$\text{gfp}(f) = f^n(\top).$$

Remark (for readers familiar with ordinal numbers). It is possible always to obtain $\text{lfp}(f)$ as the limit of an ascending (transfinite) sequence but one may have to iterate through the ordinals. To this effect define $f^{\uparrow\kappa} \in L$ for an ordinal κ by the equation

$$f^{\uparrow\kappa} = f\left(\bigsqcup_{\kappa' < \kappa} f^{\uparrow\kappa'}\right)$$

and note that for a natural number n we have $f^{\uparrow n} = f^{n+1}(\perp)$. Then $\text{lfp}(f) = f^{\uparrow\kappa}$ whenever κ is a cardinal number strictly greater than the cardinality of L , e.g. κ may be taken to be the cardinality of $\mathcal{P}(L)$. A similar construction allows to obtain $\text{gfp}(f)$ as the limit of a descending (transfinite) chain. ■

Concluding Remarks

For more information on partially ordered sets consult a text book (e.g. [33]).

Appendix B

Induction and Coinduction

We begin by reviewing a number of techniques for conducting inductive proofs. We then motivate the concept of coinduction and finally formulate a general proof principle for coinduction. This makes heavy use of Tarski's Fixed Point Theorem (Proposition A.10).

B.1 Proof by Induction

Mathematical induction. Perhaps the best known induction principle is that of *mathematical induction*. To prove that a property, $Q(n)$, holds for all natural numbers, n , we establish

$$Q(0) \\ \forall n : Q(n) \Rightarrow Q(n + 1)$$

and conclude

$$\forall n : Q(n)$$

Formally, the correctness of mathematical induction can be related to the fact that each natural number is either 0 or the successor, $n + 1$, of some other natural number, n . Thus the proof principle reflects the way the natural numbers are constructed.

Structural induction. Mathematical induction allows us to perform induction on the *size* of any structure for which a notion of size can be defined; this is just a mapping from the structure into the natural numbers. As an example consider an algebraic data type given by

$$d \in \mathbf{D} \\ d ::= \text{Base} \mid \text{Con}_1(d) \mid \text{Con}_2(d, d)$$

where Base is a base case, Con_1 is a unary constructor and Con_2 is a binary constructor. To prove that a certain property, $Q(d)$, holds for all elements, d , of \mathbf{D} we can define a size measure:

$$\begin{aligned} \text{size}(\text{Base}) &= 0 \\ \text{size}(\text{Con}_1(d)) &= 1 + \text{size}(d) \\ \text{size}(\text{Con}_2(d_1, d_2)) &= 1 + \text{size}(d_1) + \text{size}(d_2) \end{aligned}$$

and then proceed by mathematical induction on $\text{size}(d)$ to prove $Q(d)$.

Alternatively we can conceal the mathematical induction within a principle of *structural induction*: we must then show

$$\begin{aligned} Q(\text{Base}) \\ \forall d : Q(d) \Rightarrow Q(\text{Con}_1(d)) \\ \forall d_1, d_2 : Q(d_1) \wedge Q(d_2) \Rightarrow Q(\text{Con}_2(d_1, d_2)) \end{aligned}$$

from which we conclude

$$\forall d : Q(d)$$

Once again the proof principles reflects the way the data are constructed.

Induction on the shape. Now suppose that Base represents 0, that $\text{Con}_1(d)$ represents $d + 1$, and that $\text{Con}_2(d_1, d_2)$ represents $d_1 + d_2$. We can then define a Natural Semantics ..

$$d \rightarrow n$$

for evaluating d into the number, n , it represents:

$$\begin{aligned} [\text{base}] \quad & \text{Base} \rightarrow 0 \\ [\text{con}_1] \quad & \frac{d \rightarrow n}{\text{Con}_1(d) \rightarrow n + 1} \\ [\text{con}_2] \quad & \frac{d_1 \rightarrow n_1 \quad d_2 \rightarrow n_2}{\text{Con}_2(d_1, d_2) \rightarrow n_1 + n_2} \end{aligned}$$

This defines a notion of evaluation trees, $d \xrightarrow{\nabla} n$: there is one base case ($[\text{base}]$) and two constructors ($[\text{con}_1]$ and $[\text{con}_2]$). Again we can perform induction on the size of the evaluation trees but as above it is helpful to conceal the mathematical induction within a principle of *induction on the shape* of inference trees: we must show

$$\begin{aligned} Q(\text{Base} \rightarrow 0) \\ \forall (d \xrightarrow{\nabla} n) : Q(d \xrightarrow{\nabla} n) \Rightarrow Q\left(\frac{d \xrightarrow{\nabla} n}{\text{Con}_1(d) \rightarrow n + 1}\right) \end{aligned}$$

$$\forall(d_1 \Vdash n_1), (d_2 \Vdash n_2) : Q(d_1 \Vdash n_1) \wedge Q(d_2 \Vdash n_2) \Rightarrow$$

$$Q\left(\frac{d_1 \Vdash n_1 \quad d_2 \Vdash n_2}{\text{Con}_2(d_1, d_2) \rightarrow n_1 + n_2}\right)$$

from which we conclude

$$\forall(d \Vdash n) : Q(d \Vdash n)$$

As is to be expected, the proof principle once again reflects the way evaluation trees are constructed.

Course of values induction. All of the above induction principles have been constructive in the sense that we establish a predicate for the base cases and then show that it is maintained by all constructors. A variant of mathematical induction with a different flavour requires proving

$$\forall n : (\forall m < n : Q(m)) \Rightarrow Q(n)$$

from which we conclude

$$\forall n : Q(n)$$

Here the base case is dealt with in the same manner as the induction step. This induction principle is called *course of values induction*.

Well-founded induction. Course of values induction is an instance of a very powerful induction principle called *well-founded induction*. Given a partially ordered set (D, \preceq) , the partial ordering is a *well-founded ordering* if there is no infinite decreasing sequence

$$d_1 \succ d_2 \succ d_3 \succ \dots$$

where $d \succ d'$ means $d' \preceq d \wedge d \neq d'$ – this amounts to the Descending Chain Condition studied in Appendix A. The principle of well-founded induction then says: if we show

$$\forall d : (\forall d' \prec d : Q(d')) \Rightarrow Q(d)$$

we may then conclude

$$\forall d : Q(d)$$

(The proof of correctness of this principle is along the lines of the proof of (A.3) in Lemma A.6 and can be found also in the literature referenced below.)

B.2 Introducing Coinduction

To explain the difference between induction and coinduction, and to motivate the need for coinductive methods, let us consider a small example. Consider the program

if $j(27, m)$ then "something good" else "something bad"

where f is a function from pairs of natural numbers (i.e. pairs of non-negative integers) to truth values.

We want to ensure that the program never undertakes to do "something bad". Since the value of m is not known it is not feasible to prove that $f(27, m) \neq \text{false}$ by merely evaluating $f(27, m)$; we therefore need to perform some kind of proof. For this it is natural to define the predicate Q_f as follows

$$Q_f(n) \text{ iff } \forall m : f(n, m) \neq \text{false}$$

where it is implicit that $m, n \geq 0$.

Perhaps the most obvious approach is to use mathematical induction to prove $\forall n : Q_f(n)$. This amounts to proving

$$\begin{aligned} &Q_f(0) \\ &\forall n : Q_f(n) \Rightarrow Q_f(n+1) \end{aligned}$$

and then concluding

$$\forall n : Q_f(n)$$

from which the desired $Q_f(27)$ follows.

An alternative presentation of essentially the same idea is to establish the validity of the axiom and rule

$$Q_f(0) \quad \frac{Q_f(n)}{Q_f(n+1)}$$

and then deduce that

$$\forall n : Q_f(n)$$

Here the basic steps in the mathematical induction have been couched in terms of an *inductive definition* of the predicate Q_f .

The approach outlined above works nicely for the function f_0 defined by

$$\begin{aligned} f_0(0, m) &= \text{true} \\ f_0(n+1, m) &= f_0(n, m) \end{aligned}$$

but what about the functions f_1 , f_2 and f_3 defined by

$$\begin{aligned} f_1(0, m) &= f_1(0, m) & f_2(0, m) &= \text{true} \\ f_1(n+1, m) &= f_1(n, m) & f_2(n+1, m) &= f_2(n+1, m) \end{aligned}$$

$$\begin{aligned} f_3(0, m) &= f_3(0, m) \\ f_3(n+1, m) &= f_3(n+1, m) \end{aligned}$$

where $f_i(27, m)$ never terminates? Intuitively, they should be acceptable in the sense that “something bad” never happens. However, we cannot prove this by induction because we cannot establish the base case (for f_1 and f_3) and/or we cannot establish the inductive step (for f_2 and f_3).

An intuitive argument for why f_i is acceptable might go as follows: assume that all occurrences of f_i on the right hand sides of the above definitions satisfy Q_i ; then it follows that also the f_i on the left hand side does. Hence f_i satisfies Q_i , i.e. $\forall n : Q_i(n)$. This sounds very dangerous: we assume the desired result in order to prove it. However, with due care and given the proper definition of Q_i , this is a valid proof: it is a coinductive proof.

Obtaining a functional. Let us rewrite the defining clauses of f_i into clauses for Q_i so as to clarify the relationship between when Q_i holds on the left hand side and on the right hand side of the definitions of f_i :

$$\begin{array}{ll}
 Q_0(0) \text{ iff } true & Q_1(0) \text{ iff } Q_1(0) \\
 Q_0(n+1) \text{ iff } Q_0(n) & Q_1(n+1) \text{ iff } Q_1(n) \\
 \\
 Q_2(0) \text{ iff } true & Q_3(0) \text{ iff } Q_3(0) \\
 Q_2(n+1) \text{ iff } Q_2(n+1) & Q_3(n+1) \text{ iff } Q_3(n+1)
 \end{array} \tag{B.1}$$

Here the clauses for Q_0 looks just like our principle for mathematical induction whereas the others involve some amount of circularity. To make this evident let us rewrite the above as

$$Q_i = Q_i(Q_i) \tag{B.2}$$

where

$$\begin{array}{ll}
 Q_0(Q')(0) = true & Q_1(Q')(0) = Q'(0) \\
 Q_0(Q')(n+1) = Q'(n) & Q_1(Q')(n+1) = Q'(n) \\
 \\
 Q_2(Q')(0) = true & Q_3(Q')(0) = Q'(0) \\
 Q_2(Q')(n+1) = Q'(n+1) & Q_3(Q')(n+1) = Q'(n+1)
 \end{array} \tag{B.3}$$

Clearly Q_i satisfies (B.1) if and only if it satisfies (B.2) with Q_i as in (B.3).

It is immediate that each Q_i is a monotone function on the complete lattice

$$(\mathbb{N} \rightarrow \{true, false\}, \sqsubseteq)$$

of predicates where $Q_1 \sqsubseteq Q_2$ means that $\forall n : Q_1(n) \Rightarrow Q_2(n)$ and where the least element \perp is given by $\forall n : \perp(n) = false$ and the greatest element \top is given by $\forall n : \top(n) = true$. Using Tarski's Fixed Point Theorem (Proposition A.10) it follows that each Q_i has a least fixed point $lfp(Q_i)$ and a greatest fixed fixed point $gfp(Q_i)$; these are possibly different predicates in $(\mathbb{N} \rightarrow \{true, false\}, \sqsubseteq)$.

Least fixed point. Let us begin by looking at the least fixed points. It follows from Appendix A that

$$\bigsqcup_k Q_0^k(\perp) \sqsubseteq \text{lfp}(Q_0)$$

and given that the clauses for $Q_0(Q)$ only use a finite number of Q 's on the right hand sides (in fact zero or one), Q_0 satisfies a continuity property that ensures that

$$\bigsqcup_k Q_0^k(\perp) = \text{lfp}(Q_0)$$

This is good news because our previous proof by mathematical induction essentially defines the predicate $\bigsqcup_k Q_0^k(\perp)$: $Q_0^k(\perp)(n)$ holds if and only if at most k axioms and rules suffice for proving $Q_0(n)$. Thus it would seem that a proof by induction "corresponds" to taking the least fixed point of Q_0 .

Next let us look at Q_3 . Here

$$\text{lfp}(Q_3) = \perp$$

because $Q_3(\perp) = \perp$ so that \perp is a fixed point. This explains why we have $\text{lfp}(Q_3)(27) = \text{false}$ and why an inductive proof will not work for establishing $Q_3(27)$. Somewhat similar arguments can be given for Q_1 and Q_2 .

Greatest fixed point. Let us next look at the greatest fixed points. Here

$$\text{gfp}(Q_3) = \top$$

because $Q_3(\top) = \top$ so that \top is a fixed point. This explains why we have $\text{gfp}(Q_3)(27) = \text{true}$ and thus provides a formal underpinning for our belief that f_3 will not cause any harm in the example program. Somewhat similar arguments can be given for Q_1 and Q_2 .

Also for Q_0 it will be the case that $\text{gfp}(Q_0)(27) = \text{true}$. This is of course not surprising since $\text{lfp}(Q_0)(27) = \text{true}$ and $\text{lfp}(Q_0) \sqsubseteq \text{gfp}(Q_0)$. However, it is more interesting to note that for Q_0 there is no difference between the inductive and the coinductive approach (unlike what is the case for Q_1 , Q_2 and Q_3):

$$\text{lfp}(Q_0) = \text{gfp}(Q_0)$$

because mathematical induction on n suffices for proving that $\text{lfp}(Q_0)(n) = \text{gfp}(Q_0)(n)$.

Remark. To the mathematically inclined reader we should point out that the fact that $\text{lfp}(Q_0) = \text{gfp}(Q_0)$ is related to Banach's Fixed Point Theorem: a contractive operator on a metric space has a unique fixed point. Contractiveness of Q_0 (as opposed to Q_1 , Q_2 and Q_3) follows because the clause for $Q_0(Q)(n)$ only performs calls to Q on arguments smaller than n . ■

B.3 Proof by Coinduction

Consider once again the algebraic data type

$$d \in \mathbf{D}$$

$$d ::= \text{Base} \mid \text{Con}_1(d) \mid \text{Con}_2(d, d)$$

with one base case, one unary constructor and one binary constructor. Next consider the definition of a predicate

$$Q : \mathbf{D} \rightarrow \{\text{true}, \text{false}\}$$

by means of clauses of the form

$$\begin{aligned} Q(\text{Base}) & \text{ iff } \dots \\ Q(\text{Con}_1(d)) & \text{ iff } \dots Q(d') \dots \\ Q(\text{Con}_2(d_1, d_2)) & \text{ iff } \dots Q(d'_1) \dots Q(d'_2) \dots \end{aligned}$$

We can use this as the basis for defining a function \mathcal{Q} by cases as in

$$\begin{aligned} \mathcal{Q}(Q')(\text{Base}) & = \dots \\ \mathcal{Q}(Q')(\text{Con}_1(d)) & = \dots Q'(d') \dots \\ \mathcal{Q}(Q')(\text{Con}_2(d_1, d_2)) & = \dots Q'(d'_1) \dots Q'(d'_2) \dots \end{aligned}$$

We note that

$$(\mathbf{D} \rightarrow \{\text{true}, \text{false}\}, \sqsubseteq)$$

is a complete lattice under the ordering given by $Q_1 \sqsubseteq Q_2$ if and only if $\forall d : Q_1(d) \Rightarrow Q_2(d)$. We also

assume that \mathcal{Q} is monotone

and this means that e.g. a clause like “ $Q(\text{Con}_1(d))$ iff $\neg Q(d)$ ” will not be acceptable. From Proposition A.10 it follows that \mathcal{Q} has a least as well as a greatest fixed point.

Induction (or least fixed point). Consider first the inductive definition

$$Q = \text{lfp}(\mathcal{Q}) \tag{B.4}$$

This is more commonly written as

$$\frac{\dots}{Q(\text{Base})} \quad \frac{\dots Q(d') \dots}{Q(\text{Con}_1(d))} \quad \frac{\dots Q(d'_1) \dots Q(d'_2) \dots}{Q(\text{Con}_2(d_1, d_2))} \tag{B.5}$$

It is often the case that each rule only has a *finite* number of calls to Q and then the two definitions are equivalent: the predicate in (B.5) amounts

to $\bigsqcup_k Q^k(\perp)$ and by a continuity property as discussed above, this agrees with the predicate of (B.4). A *proof by induction* then simply amounts to establishing the validity of the axioms and rules in (B.5). Such a proof has a very constructive flavour: we take nothing for granted and only believe what can be demonstrated to hold. This proof strategy is often used to reason about semantics because the semantics of a program should not allow any spurious behaviour that is not forced by the semantics.

Coinduction (or greatest fixed point). Consider next the *coinductive definition*

$$Q = \text{gfp}(Q)$$

A *proof by coinduction* then amounts to using the proof rule

$$\frac{Q' \sqsubseteq Q(Q')}{Q' \sqsubseteq Q} \quad \left(\text{i.e. } \frac{Q' \sqsubseteq Q(Q')}{Q' \sqsubseteq \text{gfp}(Q)} \right)$$

as follows from the formula for $\text{gfp}(Q)$ given in Proposition A.10. So to prove $Q(d)$ one needs to

find some Q' such that

$$\begin{aligned} & Q'(d) \\ & \forall d' : Q'(d') \Rightarrow Q(Q')(d') \end{aligned}$$

Such a proof has a very optimistic flavour: we can assume everything we like as long as it cannot be demonstrated that we have violated any facts. It is commonly used for checking that a specification holds because the specification should not forbid some behaviour unless explicitly called for.

It sometimes saves a bit of work to use the derived proof rule

$$\frac{Q' \sqsubseteq Q(Q \sqcup Q')}{Q' \sqsubseteq Q}$$

To see that this is a valid proof rule suppose that $Q' \sqsubseteq Q(Q \sqcup Q')$. By definition of Q we also have $Q \sqsubseteq Q(Q)$ and by monotonicity of Q this gives $Q \sqsubseteq Q(Q \sqcup Q')$. Hence

$$Q \sqcup Q' \sqsubseteq Q(Q \sqcup Q')$$

and $Q \sqcup Q' \sqsubseteq Q$ follows by definition of Q . It is then immediate that $Q' \sqsubseteq Q$ as required.

Clearly this explanation can be generalised to algebraic data types with additional base cases and constructors, and from predicates to relations.

Example B.1 Consider the relations

$$\begin{aligned} R_1 &: D_{11} \times D_{12} \rightarrow \{\text{true}, \text{false}\} \\ R_2 &: D_{21} \times D_{22} \rightarrow \{\text{true}, \text{false}\} \end{aligned}$$

defined by

$$\begin{aligned} R_1 &= \mathcal{R}_1(R_1, R_2) \\ R_2 &= \mathcal{R}_2(R_1, R_2) \end{aligned}$$

This is intended to define

$$(R_1, R_2) = \text{gfp}(\mathcal{R})$$

where $\mathcal{R}(R'_1, R'_2) = (\mathcal{R}_1(R'_1, R'_2), \mathcal{R}_2(R'_1, R'_2))$ is assumed to be monotone.

Next write $R' \sqcup R''$ for the relation defined by

$$d_1 (R' \sqcup R'') d_2 \text{ iff } (d_1 R' d_2) \vee (d_1 R'' d_2)$$

and write $R' \sqsubseteq R''$ for the truth value defined by

$$R' \sqsubseteq R'' \text{ iff } \forall d_1, d_2 : d_1 R' d_2 \Rightarrow d_1 R'' d_2$$

We then have the following version of the coinduction principle: we establish

$$\begin{aligned} R'_1 &\sqsubseteq \mathcal{R}_1(R'_1, R'_2) \\ R'_2 &\sqsubseteq \mathcal{R}_2(R'_1, R'_2) \end{aligned}$$

and conclude

$$R'_1 \sqsubseteq R_1 \text{ and } R'_2 \sqsubseteq R_2$$

In analogy with before, it sometimes saves a bit of work only to show

$$\begin{aligned} R'_1 &\sqsubseteq \mathcal{R}_1(R_1 \sqcup R'_1, R_2 \sqcup R'_2) \\ R'_2 &\sqsubseteq \mathcal{R}_2(R_1 \sqcup R'_1, R_2 \sqcup R'_2) \end{aligned}$$

because we also have

$$\begin{aligned} R_1 &\sqsubseteq \mathcal{R}_1(R_1, R_2) \sqsubseteq \mathcal{R}_1(R_1 \sqcup R'_1, R_2 \sqcup R'_2) \\ R_2 &\sqsubseteq \mathcal{R}_2(R_1, R_2) \sqsubseteq \mathcal{R}_2(R_1 \sqcup R'_1, R_2 \sqcup R'_2) \end{aligned}$$

and this allows us to conclude

$$R'_1 \sqsubseteq R_1 \text{ and } R'_2 \sqsubseteq R_2$$

using the definition of $(R_1, R_2) = \text{gfp}(\mathcal{R})$. ■

Concluding Remarks

For more information on induction principles consult a text book (e.g. [10]). It is harder to find good introductory material on coinduction; one possibility is to study the work on strong bisimulation in CCS (e.g. Chapter 4 of [80]).

Appendix C

Graphs and Regular Expressions

(This material remains to be integrated.)

Index of Notation

- C , 302
 D_* , 80
 E , 67, 336
 P_* , 80
 $R \rightarrow R$, 208
 $R_1 \rightarrow R_2$, 213
 R_β , 211
 R , 208
 $S : \Sigma \xrightarrow{RD} \Sigma$, 21
 $S : \Sigma \rightarrow \Sigma$, 17
 $S : RD_1 \rightarrow RD_2$, 18
 S_* , 36
 S , 3
 $[B]^l$, 34
 $[f]_\Delta^2$, 224
 Δ , 189, 223
 $\Gamma \upharpoonright X$, 280
 $\Gamma \vdash e : \tau$, 22
 Γ , 22, 279
 Λ , 338
 Σ , 17
 α_η , 231
 α , 15, 294, 300
 $\beta \supseteq \varphi$, 302
 $\beta_1 \rightarrow \beta_2$, 213
 β_R , 211
 β_η , 231
 β , 210, 300
Proc, 336
 A , 53, 105
 B , 53, 105
 \mathcal{F} , 66
 \mathcal{N} , 53
 \mathcal{R} , 155
 \mathcal{V} , 156
 $[\delta, \ell]_k$, 190
 δ , 88, 189
 ℓ , 3, 140
 $\eta[V]$, 239
 η , 231
 γ_η , 231
 γ , 15
 op_a , 53, 71
 op_b , 53
 op_r , 53
 op_a , 4
 op_b , 4
 op_r , 4
 op , 140, 278
 sel , 101
 \varkappa , 104
 CS , 14
 DU , 52
 H , 108, 111
 $RD \vdash S \triangleright S'$, 26
 RD , 5
 SRD , 13
 S , 108, 111
 UD , 52
 inf , 216
 is , 110, 111
 sup , 216
 $ALoc(H)$, 108
 $ALoc(S)$, 108
 AV , 99, 311
 CP , 99
 Ext , 217, 366
 FV , 141
 Fix , 217, 366
 IAV , 99
 $Observe$, 332
 Red , 217, 366

- $blocks_*$, 82
 $blocks$, 34, 81
 $clear$, 50
 def , 50
 dom , 84, 151, 280, 295
 du , 50
 $final_*$, 82
 $final$, 34, 81
 $flow^R$, 35
 $flow_*$, 82
 $flow$, 35, 81
 gen_{AE} , 38
 gen_{LV} , 47
 gen_{RD} , 42
 gen_{VB} , 45
 gfp , 217, 366
 id , 296
 $init_*$, 82
 $init$, 33, 81
 $inter-flow_*$, 82
 $kill_{AE}$, 37
 $kill_{LV}$, 47
 $kill_{RD}$, 41
 $kill_{VB}$, 45
 $labels_*$, 82
 $labels$, 34, 81
 lfp_{Δ} , 224
 lfp_{∇} , 220
 lfp , 217, 366
 ran , 84
 ud , 50
 use , 50
 ι_E^i , 64, 67
 ι_* , 183
 ι , 67
 CP , 336
 PP , 336
 \bar{q} , 326
 \bar{r} , 326
 \bar{L} , 89
 $\hat{\Gamma} \vdash e : \hat{\tau} \ \& \ \varphi$, 23
 $\hat{\Gamma}$, 24, 281, 300, 315
 \widehat{Cache} , 142, 184, 189
 \widehat{DCache} , 184
 \widehat{DEnv} , 184
 \widehat{Data} , 184
 \widehat{Env} , 142, 184, 189
 \widehat{State}_{CP} , 70
 \widehat{TEnv} , 281
 $\widehat{Type}[\tau]$, 292
 \widehat{Type} , 281
 \widehat{Val}_d , 179
 \widehat{Val} , 142, 184, 189
 $\hat{\mathcal{F}}$, 89
 $\hat{\delta}$, 184
 \hat{i} , 93, 95, 97, 98
 $\hat{\rho}$, 11, 142, 184, 189
 \hat{C} , 11, 142, 184, 189
 \hat{D} , 184
 \hat{op} , 179
 $\hat{\sigma}$, 321, 329, 337
 $\hat{\tau} \leq \hat{\tau}'$, 318, 323, 340
 $\hat{\tau}_1 \sqsubseteq \hat{\tau}_2$, 292
 $\hat{\tau}$, 23, 24, 281, 300, 315, 321, 329, 337
 \hat{d} , 184
 $\hat{f}_{l_c}^1$, 90, 94, 95, 97, 98
 \hat{f}_{l_c, l_r}^2 , 91, 94, 95, 97, 98
 \hat{f}_l , 89
 \hat{f}_{l_n} , 90
 \hat{f}_{l_x} , 90
 \hat{v} , 142, 179, 184, 189
 ∇ , 219
 ϕ_l^{SA} , 115
 π , 278
 \diamond , 104
 $\rho | X$, 151
 $\rho \vdash_* \langle S, \varsigma \rangle \rightarrow \langle S', \varsigma' \rangle$, 84
 $\rho \vdash_* \langle S, \varsigma \rangle \rightarrow \varsigma'$, 84
 $\rho \vdash ie_1 \rightarrow ie_2$, 151
 ρ_* , 84
 ρ , 84, 150, 327
 W , 72
 $\sigma_1 \sim_V \sigma_2$, 60
 σ , 52, 104
 τ , 279, 294
 $\theta_A, \models C$, 303

- θ , 295
 $\varphi \sqsubseteq \varphi'$, 341
 $\varphi_1 \sqsubseteq \varphi_2$, 292
 $\varphi_1 \subseteq \varphi_2$, 292
 φ , 23, 281, 300, 315, 321, 329, 337
 e , 325
 $\varsigma[M]$, 237
 ς_η , 239
 ς , 84, 237, 314, 327
 \wp , 104
 ξ , 84, 103, 314
 ζ , 321
 a , 3
 b , 3
 ce , 189
 ch , 334
 c , 140, 278
 d_c , 179
 d_{op} , 179
 $e_1 \rightarrow e_2$, 335
 ee , 325
 e , 140, 277, 333
 $f_{\overline{\nu}}$, 220
 f_{ℓ}^{SA} , 112, 115
 f_{ℓ} , 183
 f_{ℓ} , 66
 f_p , 206, 213
 f_{ℓ}^{CP} , 71
 f , 140, 278
 ie , 151
 it , 151
 l_n^{ϕ} , 219
 n_X , 107
 n_{\emptyset} , 107
 n , 3
 o , 327
 $p \vdash l_1 \triangleright l_2$, 206, 213
 $p \vdash v_1 \rightsquigarrow v_2$, 205, 213
 p , 80, 102, 336
 rn , 325
 r , 325, 337
 t , 140
 v , 150, 287, 327, 335
 w , 327
 x , 3, 140, 278
 y , 3
 Δ , 88, 93, 95, 97, 98
AExp_{*}, 36
AExp, 3
AHeap, 108, 111
ALoc, 107
AState, 108, 111
AType, 294
AVar, 300
Ann_{CA}, 337
Ann_{ES}, 321
Ann_{RI}, 329
Ann_{SE}, 315
Ann, 281
BExp, 3
Blocks_{*}, 36
CEnv, 189
Chan, 334
Const, 140, 278
Data, 179
EExp, 325
EVal, 327
Env, 84, 150, 327
Exp, 140, 277, 333
Heap, 104
IExp, 151
ITerm, 151
Interval, 215
IsShared, 111
Lab_{*}, 168
Lab_{*}, 36
Lab, 3, 140
Loc, 84, 103, 314
Num, 3
Offset, 327
Op, 140, 278
Parity, 237
Pnt, 278
RName, 325
RVar, 325
Range, 241
Reg_{CA}, 337
Reg_{RI}, 325

- SAnn**, 300
SG, 112
SType, 300
SVal, 327
Scheme_{CA}, 337
Scheme_{RI}, 329
Sel, 101
Sign, 231
State, 52, 104
Stmt, 3
Store, 84, 314, 327
TEnv, 279
TVar, 294, 300
Term_{*}, 168
Term, 140
Trace, 13
Type_{CA}, 337
Type_{ES}, 321
Type_{RI}, 329
Type_{SE}, 315
Type, 279
Val, 150, 287, 335
Var_{*}, 36, 168
Var, 3, 140, 278
A_{CP}, 71
C_{}*, 171
F_{CP}, 71
U_{CFA}, 301
U_{UL}, 296, 299
W_{CFA}, 303
W_{UL}, 295
kill_x, 116
kill_{x.sel}, 123
tr, 13
AE_{entry}, 38
AE_{exit}, 38
Analysis⁼, 67
Analysis^E, 67
JUDG_{CFA}[$\Gamma \vdash_{UL} e : \tau$], 293
LV_{entry}, 47
LV_{exit}, 47
RD_{entry}, 42
RD_{exit}, 42
VB_{entry}, 45
VB_{exit}, 45
 $[\delta, \ell]$, 94
 $[\cdot]$, 281
 $[\delta, \ell]_k$, 95
LV⁼, 58
LV^E, 58
AExp(\cdot), 36
BExp(\cdot), 36
Analysis, 63
MFP, 72
MOP, 77
MVP, 88
Shape, 112
path, 77
vpath, 88

Index

- k*-CFA analysis, 189
- Abstract Interpretation, 13
- Annotated Type System, 17
- Array Bound Analysis, 215, 240
- Ascending Chain Condition, 65, 363
- Available Expressions Analysis, 37
- Bit Vector Framework, 135
- Call-tracking Analysis, 23
- Chaotic Iteration, 25
- Closure Analysis, 193
- Code Motion, 51
- Communication Analysis, 333, 337
- Constant Folding, 26, 70, 131
- Constant Propagation Analysis, 70, 131, 183, 206, 257, 259
- Control Flow Analysis, 10, 207, 281, 282
- Data Flow Analysis, 5
- Dead Code Elimination, 47, 51, 131
- Definition-Use chains, 143
- Definition-Use chain, 50
- Denotational Semantics, 217, 266
- Descending Chain Condition, 66, 363
- Detection of Signs Analysis, 89, 180
- Distributive Framework, 67
- Effect System, 17
- Exception Analysis, 319
- Galois connection, 15, 228
- Galois insertion, 235
- Live Variables Analysis, 47
- MFP solution, 72
- MOP solution, 76
- MVP solution, 86, 88
- Maximal Fixed Point, 72
- Meet Over all Paths, 76
- Meet over all Valid Paths, 86
- Modified Post Correspondence Problem, 77
- Monotone Framework, 66
- Moore family, 135, 159, 169, 209, 293, 359
- Natural Semantics, 286
- Reaching Definitions Analysis, 4, 41
- Region Inference, 324
- Set Based Analysis, 193
- Shape Analysis, 101, 267
- Side Effect Analysis, 314
- Structural Operational Semantics, 52, 150
- Tarski's Fixed Point Theorem, 366
- Type and Effect System, 17
- UCAI, 285, 300
- Use-Definition chain, 50
- Very Busy Expressions Analysis, 44
- abstract 0-CFA, 145
- abstract cache, 142, 189
- abstract data value, 179
- abstract environment, 142, 189
- abstract heap, 108
- abstract locations, 107
- abstract reachability component, 197
- abstract state, 107
- abstract summary location, 107
- abstract syntax, 4, 140
- abstract value, 142, 189
- abstraction function, 15, 227
- acceptable 0-CFA, 144
- additive, 359
- adjoint, 232

- adjunction, 15, 229
- affine, 360
- annotated type environment, 281
- annotated type, 281, 315, 321, 329
- annotation substitution, 303
- annotation variable, 300
- annotation, 281, 315, 321, 329
- ascending chain, 363
- assigned variables, 99
- assumption set, 97
- atomic subtyping, 344
- augmented type, 294
- backward analysis, 64, 83
- basic block, 134
- behaviour variable, 338
- behaviour, 338
- call string, 93
- call-by-name, 200
- call-by-result, 80
- call-by-value, 80
- called procedures, 99
- cartesian product, 362
- chain, 363
- channel identifier, 334
- channel pool, 336
- channel, 334
- closure, 150
- coinductive definition, 148, 162, 376
- collecting semantics, 13, 270
- combination operator, 65
- compatible shape graph, 111
- complete lattice, 357
- complete path, 86
- completely additive, 359
- completely multiplicative, 360
- concrete syntax, 4
- concretisation function, 15, 227
- conditional constraint, 12, 170
- conservative extension, 285
- constraint based O-CFA, 171
- constraint based approach, 8, 11
- constraint system, 58, 67
- constraint, 170, 302
- context environment, 189
- context information, 88, 187, 189
- context-insensitive, 93, 148, 186
- context-sensitive, 93, 187
- contravariant, 318
- correctness relation, 60, 155, 208, 234, 252
- course of values induction, 371
- covariant, 318
- cover, 295, 303
- data array, 175
- definition clear path, 50
- derivation sequence, 55
- descending chain, 363
- direct product, 246
- direct tensor product, 247
- duality, 265
- dynamic dispatch problem, 82, 140
- edge array, 175
- effect, 23, 315, 321, 329
- elementary block, 3, 34
- embellished monotone framework, 89
- environment, 84, 150, 326
- equality of annotations, 284
- equation system, 57, 63, 64, 67
- equational approach, 5
- evaluation context, 335
- eventually stabilises, 363
- expression, 140, 277, 333
- extended expression, 325
- extended type, 329
- extensive, 217, 366
- extraction function, 231, 239
- extremal label, 67
- extremal value, 67, 183
- faint variable, 134
- final label, 34
- finite chain, 363
- finite height, 363
- first-order analysis, 206
- fixed point, 217, 366
- flow graph, 35
- flow-insensitive, 98, 148
- flow-sensitive, 98

- flow, 35, 67
- forward analysis, 64, 83, 112
- free algebra, 300
- free variable, 141
- fresh type variable, 296
- functional composition, 240
- general subtyping, 344
- generalisation, 323, 331
- generalised Monotone Framework, 254
- generated by, 211
- graph formulation of the constraints, 174
- greatest element, 358
- greatest fixed point, 217, 366
- greatest lower bound, 357
- ground substitution, 295, 303
- ground validation, 303
- heap, 104
- hoisting, 44
- independent attribute analysis, 136, 242
- induced analysis, 16, 251
- induction on the shape, 370
- inductive definition, 162, 372
- initial label, 33
- injective, 359
- instance, 67, 183, 255
- instantiation, 323, 332
- instrumented semantics, 131
- intermediate expression, 151
- intermediate term, 151
- interprocedural analysis, 80, 192
- interprocedural flow, 82, 140
- intraprocedural analysis, 80
- isolated entries, 37
- isolated exits, 37
- isomorphism, 361
- join semi-lattice, 65
- label consistent, 37
- label, 34
- latent effect, 23
- least element, 358
- least fixed point, 217, 254, 366
- least upper bound, 357
- live variable, 47
- locations, 103
- location, 84, 314
- logical relation, 208, 213
- lower adjoint, 232
- lower bound, 357
- materialise, 120
- mathematical induction, 369
- may analysis, 64, 112
- model intersection property, 159
- mono-variant, 186
- monotone function space, 245, 362
- monotone structure, 183, 268
- monotone, 359
- multiplicative, 359
- must analysis, 64
- narrowing operator, 223
- non-free algebra, 300
- non-trivial expression, 36
- offset, 327
- open system, 148, 179
- optimal analysis, 17, 253
- partial ordering, 357
- partially ordered set, 357
- path, 76
- pointer expressions, 102
- poly-variant, 187
- polymorphic recursion, 330
- polymorphism, 321, 323
- polynomial k -CFA, 192, 193
- procedure call graph, 100
- process pool, 336
- process, 334
- program point, 278
- proof by coinduction, 376
- proof by induction, 376
- proof normalisation, 306, 352
- property space, 65
- reduced product, 249
- reduced tensor product, 249
- reduction operator, 237
- reductive, 217, 366
- reference variable, 313

- region name, 325
- region polymorphic closure, 327
- region variable, 325
- region, 325, 327, 338
- relational analysis, 136, 243
- representation function, 210, 230, 235, 252, 256
- reverse flows, 35
- safe approximation, 165
- second-order analysis, 206
- selector names, 101
- semantic correctness, 289
- semantically reaching definitions, 13
- semantics based, 3, 29
- semantics directed, 3, 29
- sequential composition, 240
- set constraints, 195
- sets of states analysis, 257
- shape conformant subtyping, 318, 343
- shape graph, 106, 111
- sharing information, 110
- simple annotation, 300
- simple substitution, 301
- simple type environment, 300
- simple type, 300
- smash product, 362
- states, 104
- state, 52
- static single assignment, 129
- storable value, 327
- store, 84, 314, 327
- strict, 360
- structural induction, 370
- subeffecting, 282, 317, 323, 331, 340
- subject reduction result, 155
- substitution, 295
- subsumption rule, 19, 21
- subtyping, 318, 323, 331, 340
- surjective, 359
- syntactic completeness, 302, 308
- syntactic soundness, 302, 306
- syntax directed 0-CFA, 166
- temporal order, 333
- tensor product, 245, 264
- term, 140
- total function space, 245, 362
- trace, 13, 132
- transfer function, 66
- transitions, 52
- type environment, 22, 279
- type reconstruction, 295, 303
- type schema, 321, 329, 338
- type substitution, 302
- type variable, 294, 300
- type, 22, 279, 338
- typing judgement, 280, 282, 315, 323, 330, 339
- underlying type system, 278
- underlying type, 281
- unification procedure, 296, 299, 301
- uniform k -CFA, 189, 190, 193
- upper adjoint, 232
- upper bound operator, 218
- upper bound, 357
- upper closure operator, 263
- valid path, 87
- values, 335
- value, 150, 287
- well-founded induction, 371
- well-founded ordering, 371
- widening operator, 219, 260
- worklist, 72, 175
- du -chain, 50, 130, 143
- ud -chain, 50, 130
- k -CFA, 193
- bind-construct, 84, 151
- close-construct, 150

Bibliography

- [1] O. Agesen, J. Palsberg, and M. Schwartzbach. Type inference of SELF: Analysis of objects with dynamic and multiple inheritance. In *Proc. ECOOP'93*, volume 707 of *Lecture Notes in Computer Science*. Springer, 1993.
- [2] A. V. Aho, R. Sethi, and J. D. Ullman. *Compilers: Principles, Techniques, Tools*. Addison Wesley, 1986.
- [3] A. Aiken. Set constraints: Results, applications, and future directions. In *Proc. Second Workshop on the Principles and Practices of Concurrent Programming*, volume 874 of *Lecture Notes in Computer Science*, pages 171–179. Springer, 1994.
- [4] A. Aiken and E. Wimmers. Type inclusion constraints and type inference. In *Proc. FPCA '93*, pages 31–41. ACM Press, 1993.
- [5] A. Aiken, E. Wimmers, and T.K. Lakshman. Soft typing with conditional types. In *Proc. POPL '94*, pages 163–173. ACM Press, 1994.
- [6] T. Amtoft, F. Nielson, and H. R. Nielson. Polymorphic subtyping for side effects. Technical Report DAIMI PB-529, Computer Science, Aarhus University, 1997. — [ANN97book] for the report but [ANN98book] for the final version.
- [7] T. Amtoft, F. Nielson, and H.R. Nielson. Type and behaviour reconstruction for higher-order concurrent programs. *Journal of Functional Programming*, 7(3):321–347, 1997.
- [8] T. Amtoft, F. Nielson, H.R. Nielson, and J. Ammann. Polymorphic subtyping for effect analysis: The dynamic semantics. In *Analysis and Verification of Multiple-Agent Languages*, volume 1192 of *Lecture Notes in Computer Science*, pages 172–206. Springer, 1997.
- [9] Anonymous. Details to be provided. ??, ??, 19??
- [10] A. Arnold and I. Guessarian. *Mathematics for Computer Science*. Prentice Hall International, 1996.

- [11] A. Banerjee. A modular, polyvariant, and type-based closure analysis. In *Proc. ICFP '97*, pages 1–10. ACM Press, 1997.
- [12] P. N. Benton. Strictness logic and polymorphic invariance. In *Proc. Second International Symposium on Logical Foundations of Computer Science*, volume 620 of *Lecture Notes in Computer Science*, pages 33–44. Springer, 1992.
- [13] P. N. Benton. Strictness properties of lazy algebraic datatypes. In *Proc. WSA '93*, volume 724 of *Lecture Notes in Computer Science*, pages 206–217. Springer, 1993.
- [14] S. K. Biswas. A demand-driven set-based analysis. In *Proc. POPL '97*, pages 372–385. ACM Press, 1997.
- [15] C. Bodei, P. Degano, F. Nielson, and H. R. Nielson. Control flow analysis for the π -calculus. In *Proc. CONCUR'98*, number 1466 in *Lecture Notes in Computer Science*, pages 84–98. Springer, 1998.
- [16] F. Bourdoncle. Abstract interpretation by dynamic partitioning. *Journal of Functional Programming*, 10:407–435, 1992.
- [17] F. Bourdoncle. Efficient chaotic iteration strategies with widenings. In *Proc. Formal Methods in Programming and Their Applications*, volume 735 of *Lecture Notes in Computer Science*, pages 128–141. Springer, 1993.
- [18] G. L. Burn, C. Hankin, and S. Abramsky. Strictness Analysis for Higher-Order Functions. *Science of Computer Programming*, 7:249–278, 1986.
- [19] W. Charatonik and L. Pacholski. Set constraints with projections are in NEXPTIME. In *Proc. FOCS '94*, pages 642–653, 1994.
- [20] D. Chase, M. Wegman, and F. Zadeck. Analysis of pointers and structures. In *Proc. PLDI '90*, pages 296–310. ACM Press, 1990.
- [21] C. Colby. Analyzing the communication topology of concurrent programs. In *Proc. PEPM '95*, pages 202–214. ACM Press, 1995.
- [22] C. Colby. Determining storage properties of sequential and concurrent programs with assignment and structured data. In *Proc. SAS '95*, volume 983 of *Lecture Notes in Computer Science*, pages 64–81. Springer, 1995.
- [23] A. Cortesi, G. Filé, R. Giacobazzi, C. Palamidessi, and F. Ranzato. Complementarity in abstract interpretation. In *Proc. SAS '95*, *Lecture Notes in Computer Science*, pages 100–117. Springer, 1995.

- [24] A. Cortesi, G. Filé, R. Giacobazzi, C. Palamidessi, and F. Ranzato. Complementation in abstract interpretation. *ACM TOPLAS*, 19(1):7–47, 1997.
- [25] P. Cousot. Semantics Foundation of Program Analysis. In S. S. Muchnick and N. D. Jones, editors, *Program Flow Analysis: Theory and Applications*, chapter 10, pages 303–342. Prentice Hall International, 1981.
- [26] P. Cousot. Types as abstract interpretations. In *Proc. POPL '97*, pages 316–331. ACM Press, 1997.
- [27] P. Cousot and R. Cousot. Abstract Interpretation: a Unified Lattice Model for Static Analysis of Programs by Construction or Approximation of Fixpoints. In *Proc. POPL '77*, pages 238–252. ACM Press, 1977.
- [28] P. Cousot and R. Cousot. Static determination of dynamic properties of generalised type unions. In *Conference on Language Design for Reliable Software*, volume 12(3) of *ACM SIGPLAN Notices*, pages 77–94, 1977.
- [29] P. Cousot and R. Cousot. Systematic Design of Program Analysis Frameworks. In *Proc. POPL '79*, pages 269–282, 1979.
- [30] P. Cousot and R. Cousot. Comparing the Galois Connection and Widening/Narrowing Approaches to Abstract Interpretation. In *Proc. PLILP '92*, volume 631 of *Lecture Notes in Computer Science*, pages 269–295. Springer, 1992.
- [31] P. Cousot and N. Halbwachs. Automatic Discovery of Linear Restraints Among Variables of a Program. In *Proc. POPL '78*, pages 84–97. ACM Press, 1978.
- [32] R. Cytron, J. Ferrante, B. K. Rosen, M. N. Wegman, and F. K. Zadek. Efficiently computing static single assignment form and the control dependence graph. *ACM TOPLAS*, 13(4):451–490, 1991.
- [33] B. A. Davey and H. A. Priestley. *Introduction to Lattices and Order*. Cambridge Mathematical Textbooks. Cambridge University Press, 1990.
- [34] A. Deutsch. On Determining Lifetime and Aliasing of Dynamically Allocated Data in Higher Order Functional Specifications. In *Proc. POPL '90*, pages 157–169. ACM Press, 1990.
- [35] A. Deutsch. Interprocedural may-alias analysis for pointers: Beyond k-limiting. In *Proc. PLDI '94*, pages 230–241. ACM Press, 1994.

- [36] M. Emami, R. Ghiya, and L. J. Hendren. Context-sensitive interprocedural points-to analysis in the presence of function pointers. In *Proc. PLDI '94*, pages 242–256. ACM Press, 1994.
- [37] K.-F. Faxén. Optimizing lazy functional programs using flow inference. In *Proc. SAS '95*, volume 983 of *Lecture Notes in Computer Science*, pages 136–153. Springer, 1995.
- [38] K.-F. Faxén. Polyvariance, polymorphism, and flow analysis. In *Proc. Analysis and Verification of Multiple-Agent Languages*, volume 1192 of *Lecture Notes in Computer Science*, pages 260–278. Springer, 1997.
- [39] C. N. Fischer and Jr. R. J. LeBlanc. *Crafting a Compiler*. Benjamin/Cummings, 1988.
- [40] C. Flanagan and M. Felleisen. Well-founded touch optimizations for futures. Technical Report Rice COMP TR94-239, Rice University, 1994.
- [41] C. Flanagan and M. Felleisen. The semantics of future and its use in program optimization. In *Proc. POPL '95*, pages 209–220. ACM Press, 1995.
- [42] Y.-C. Fuh and P. Mishra. Polymorphic subtype inference: Closing the theory-practice gap. In *Proc. TAPSOFT '89*, volume 352 of *Lecture Notes in Computer Science*, pages 167–183. Springer, 1989.
- [43] Y.-C. Fuh and P. Mishra. Type inference with subtypes. *Theoretical Computer Science*, 73:155–175, 1990.
- [44] K. L. S. Gasser, F. Nielson, and H. R. Nielson. Systematic realisation of control flow analyses for CML. In *Proc. ICFP '97*, pages 38–51. ACM Press, 1997.
- [45] R. Ghiya and L. Hendren. Connection analysis: a practical interprocedural analysis for C. In *Proc. of the eight workshop on languages and compilers for parallel computing*, 1995.
- [46] R. Ghiya and L. J. Hendren. Is it a tree, a dag, or a cyclic graph? a shape analysis for heap-directed pointers in c. In G. Kahn, editor, *Proc. POPL '96*, pages 1–15. ACM Press, 1996.
- [47] R. Giacobazzi and F. Ranzato. Compositional optimization of disjunctive abstract interpretations. In *Proc. ESOP '96*, volume 1058 of *Lecture Notes in Computer Science*, pages 141–155. Springer, 1996.
- [48] R. Giacobazzi and F. Ranzato. Optimal domains for disjunctive abstract interpretation. *Science of Computer Programming*, 32:177–210, 1998.

- [49] R. Giegerich, U. Möncke, and R. Wilhelm. Invariance of approximative semantics with respect to program transformations. In *Proc. GI - 11. Jahrestagung*, volume 50 of *Informatik Fachberichte*, pages 1-10. Springer, 1981.
- [50] P. Granger. Static analysis of arithmetical congruences. *International Journal of Computer Mathematics*, 30:165-190, 1989.
- [51] P. Granger. Static Analysis of Linear Congruence Equalities among Variables of a Program. In *Proc. TAPSOFT '91*, volume 493 of *Lecture Notes in Computer Science*, pages 169-192. Springer, 1991.
- [52] M. S. Hecht. *Flow Analysis of Computer Programs*. North Holland, 1977.
- [53] N. Heintze. Set-based analysis of ML programs. In *Proc. LFP '94*, pages 306-317, 1994.
- [54] N. Heintze. Control-flow analysis and type systems. In *Proc. SAS '95*, volume 983 of *Lecture Notes in Computer Science*, pages 189-206. Springer, 1995.
- [55] N. Heintze and J. Jaffar. A decision procedure for a class of herbrand set constraints. In *Proc. LICS '90*, pages 42-51, 1990.
- [56] N. Heintze and J. Jaffar. An engine for logic program analysis. In *Proc. LICS '92*, pages 318-328, 1992.
- [57] J. E. Hopcroft and J. D. Ullman. *Introduction to Automata Theory, Languages and Computation*. Addison Wesley, 1979.
- [58] S. Jagannathan and S. Weeks. Analyzing Stores and References in a Parallel Symbolic Language. In *Proc. LFP '94*, pages 294-305, 1994.
- [59] S. Jagannathan and S. Weeks. A unified treatment of flow analysis in higher-order languages. In *Proc. POPL '95*. ACM Press, 1995.
- [60] S. Jagannathan and A. Wright. Effective flow analysis for avoiding run-time checks. In *Proc. SAS '95*, volume 983 of *Lecture Notes in Computer Science*, pages 207-224. Springer, 1995.
- [61] T. P. Jensen. Strictness analysis in logical form. In *Proc. FPCA '91*, volume 523 of *Lecture Notes in Computer Science*, pages 352-366. Springer, 1991.
- [62] T. P. Jensen. Disjunctive strictness analysis. In *Proc. LICS '92*, pages 174-185, 1992.

- [63] M. P. Jones. A theory of qualified types. In *Proc. ESOP '92*, volume 582 of *Lecture Notes in Computer Science*, pages 287–306. Springer, 1992.
- [64] N. D. Jones and S. S. Muchnick. Flow analysis and optimization of Lisp-like structures. In S. S. Muchnick and N. D. Jones, editors, *Program Flow Analysis: Theory and Applications*, chapter 4, pages 102–131. Prentice Hall International, 1981.
- [65] N. D. Jones and S. S. Muchnick. A flexible approach to interprocedural data flow analysis and programs with recursive data structures. In *Proc. POPL '82*, pages 66–74. ACM Press, 1982.
- [66] N. D. Jones and F. Nielson. Abstract Interpretation: a Semantics-Based Tool for Program Analysis. In *Handbook of Logic in Computer Science volume 4*. Oxford University Press, 1995.
- [67] P. Jouvelot. Semantic Parallelization: a practical exercise in abstract interpretation. In *Proc. POPL '87*, pages 39–48, 1987.
- [68] P. Jouvelot and D. K. Gifford. Reasoning about continuations with control effects. In *Proc. PLDI '89*, ACM SIGPLAN Notices, pages 218–226. ACM Press, 1989.
- [69] P. Jouvelot and D. K. Gifford. Algebraic reconstruction of types and effects. In *Proc. POPL '91*, pages 303–310. ACM Press, 1990.
- [70] J. B. Kam and J. D. Ullman. Monotone data flow analysis frameworks. *Acta Informatica*, 7:305–317, 1977.
- [71] M. Karr. Affine Relationships among Variables of a Program. *Acta Informatica*, 6(2):133–151, 1976.
- [72] G. Kildall. A Unified Approach to Global Program Optimization. In *Proc. POPL '73*, pages 194–206. ACM Press, 1973.
- [73] K. Knobe and V. Sarkar. Array SSA form and its use in parallelization. In *Proc. POPL '98*, pages 107–120. ACM Press, 1998.
- [74] W. Landi and B. G. Ryder. Pointer-Induced Aliasing: A Problem Classification. In *Proc. POPL '91*, pages 93–103. ACM Press, 1991.
- [75] W. Landi and B. G. Ryder. A safe approximate algorithm for interprocedural pointer aliasing. In *Proc. PLDI '92*, pages 235–248. ACM Press, 1992.
- [76] J. Larus and P. Hilfinger. Detecting conflicts between structure accesses. In *Proc. PLDI '88*, pages 21–34. ACM Press, 1988.

- [77] J. M. Lucassen and D. K. Gifford. Polymorphic effect analysis. In *Proc. POPL '88*, pages 47–57. ACM Press, 1988.
- [78] T. J. Marlowe and B. G. Ryder. Properties of data flow frameworks - a unified model. *Acta Informatica*, 28(2):121–163, 1990.
- [79] F. Masdupuy. Using Abstract Interpretation to Detect Array Data Dependencies. In *Proc. International Symposium on Supercomputing*, pages 19–27, 1991.
- [80] R. Milner. *Communication and Concurrency*. Prentice Hall International, 1989.
- [81] J. Mitchell. Type inference with simple subtypes. *Journal of Functional Programming*, 1(3):245–285, 1991.
- [82] B. Monsuez. Polymorphic types and widening operators. In *Proc. Static Analysis (WSA '93)*, volume 724 of *Lecture Notes in Computer Science*, pages 267–281. Springer, 1993.
- [83] F. Nielson. *Abstract Interpretation using Domain Theory*. PhD thesis, University of Edinburgh, Scotland, 1984.
- [84] F. Nielson. Program Transformations in a denotational setting. *ACM TOPLAS*, 7:359–379, 1985.
- [85] F. Nielson. Tensor Products Generalize the Relational Data Flow Analysis Method. In *Proc. 4th Hungarian Computer Science Conference*, pages 211–225, 1985.
- [86] F. Nielson. A formal type system for comparing partial evaluators. In D. Bjørner, A. P. Ershov, and N. D. Jones, editors, *Proc. Partial Evaluation and Mixed Computation*, pages 349–384. North Holland, 1988.
- [87] F. Nielson. Two-Level Semantics and Abstract Interpretation. *Theoretical Computer Science — Fundamental Studies*, 69:117–242, 1989.
- [88] F. Nielson. The typed λ -calculus with first-class processes. In *Proc. PARLE'89*, volume 366 of *Lecture Notes in Computer Science*, pages 355–373. Springer, 1989.
- [89] F. Nielson. Semantics-directed program analysis: a tool-maker's perspective. In *Proc. Static Analysis Symposium (SAS)*, number 1145 in *Lecture Notes in Computer Science*, pages 2–21. Springer, 1996.
- [90] F. Nielson and H. R. Nielson. Finiteness Conditions for Fixed Point Iteration. In *Proc. LFP '92*, pages 96–108. ACM Press, 1992.

- [91] F. Nielson and H. R. Nielson. From CML to process algebras. In *Proc. CONCUR'93*, volume 715 of *Lecture Notes in Computer Science*, pages 493–508. Springer, 1993.
- [92] F. Nielson and H. R. Nielson. From CML to its process algebra. *Theoretical Computer Science*, 155:179–219, 1996.
- [93] F. Nielson and H. R. Nielson. Operational semantics of termination types. *Nordic Journal of Computing*, pages 144–187, 1996.
- [94] F. Nielson and H. R. Nielson. Infinitary Control Flow Analysis: a Collecting Semantics for Closure Analysis. In *Proc. POPL '97*. ACM Press, 1997.
- [95] F. Nielson and H. R. Nielson. The flow logic of imperative objects. In *Proc. MFCS'98*, number 1450 in *Lecture Notes in Computer Science*, pages 220–228. Springer, 1998.
- [96] F. Nielson and H. R. Nielson. Flow logics and operational semantics. *Electronic Notes of Theoretical Computer Science*, 10, 1998.
- [97] F. Nielson, H.R. Nielson, and T. Amtoft. Polymorphic subtyping for effect analysis: The algorithm. In *Analysis and Verification of Multiple-Agent Languages*, volume 1192 of *Lecture Notes in Computer Science*, pages 207–243. Springer, 1997.
- [98] H. R. Nielson, T. Amtoft, and F. Nielson. Behaviour analysis and safety conditions: a case study in CML. In *Proc. FASE '98*, number 1382 in *Lecture Notes in Computer Science*, pages 255–269. Springer, 1998.
- [99] H. R. Nielson and F. Nielson. Bounded fixed-point iteration. *Journal of Logic and Computation*, 2(4):441–464, 1992.
- [100] H. R. Nielson and F. Nielson. Higher-Order Concurrent Programs with Finite Communication Topology. In *Proc. POPL '94*. Springer, 1994.
- [101] H. R. Nielson and F. Nielson. Communication analysis for concurrent ML. In F. Nielson, editor, *ML with Concurrency*, Monographs in Computer Science, pages 185–235. Springer, 1997.
- [102] H. R. Nielson and F. Nielson. Flow logics for constraint based analysis. In *Proc. CC '98*, volume 1383 of *Lecture Notes in Computer Science*, pages 109–127. Springer, 1998.
- [103] H.R. Nielson, F. Nielson, and T. Amtoft. Polymorphic subtyping for effect analysis: The static semantics. In *Analysis and Verification of Multiple-Agent Languages*, volume 1192 of *Lecture Notes in Computer Science*, pages 141–171. Springer, 1997.

- [104] J. Palsberg. Closure analysis in constraint form. *ACM TOPLAS*, 17(1):47–62, 1995.
- [105] J. Palsberg and M. I. Schwartzbach. *Object-Oriented Type Systems*. Wiley, 1994.
- [106] H. D. Pande and B. G. Ryder. Data-flow-based virtual function resolution. In *Proc. SAS '96*, volume 1145 of *Lecture Notes in Computer Science*, pages 238–254. Springer, 1996.
- [107] J. Plevyak and A. A. Chien. Precise concrete type inference of object-oriented programs. In *Proc. OOPSLA '94*, 1994.
- [108] A. Podelski and L. Pacholski. Set constraints — a pearl in research on constraints. In *Proc. Third International Conference on the Principles and Practice of Constraint Programming*, volume 1330 of *Lecture Notes in Computer Science*, pages ??–?? Springer, 1997.
- [109] J. H. Reif and S. A. Smolka. Data Flow Analysis of Distributed Communicating Processes. *International Journal of Parallel Programming*, 19(1):1–30, 1990.
- [110] J. Reynolds. Automatic computation of data set definitions. In *Information Processing*, volume 68, pages 456–461. North Holland, 1969.
- [111] B. K. Rosen. Monoids for rapid data flow analysis. *SIAM J. Comp.*, 9:159–196, 1980.
- [112] E. Ruf. Context-insensitive alias analysis reconsidered. In *Proc. PLDI '95*, pages 13–22. ACM Press, 1995.
- [113] M. Sagiv, T. Reps, and S. Horwitz. Precise interprocedural dataflow analysis with applications to constant propagation. In *Proc. TAPSOFT '95*, volume 915 of *Lecture Notes in Computer Science*, pages 651–665, 1995.
- [114] M. Sagiv, T. Reps, and R. Wilhelm. Solving shape-analysis problems in languages with destructive updating. In *Proc. POPL '96*, pages 16–31. ACM Press, 1996.
- [115] M. Sagiv, T. Reps, and R. Wilhelm. Solving shape-analysis problems in languages with destructive updating. *ACM TOPLAS*, 20(1):1–50, 1998.
- [116] M. Sagiv, T. Reps, and R. Wilhelm. Parametric shape analysis via 3-valued logic. In *Proc. POPL '99*, pages ??–?? ACM Press, 1999.
- [117] D. Schmidt. Data flow analysis is model checking of abstract interpretations. In *Proc. POPL '98*, pages 38–48. ACM Press, 1998.

- [118] P. Sestoft. Replacing function parameters by global variables. Master's thesis, Department of Computer Science, University of Copenhagen, Denmark, 1988.
- [119] M. Shapiro and S. Horwitz. Fast and accurate flow-insensitive points-to analysis. In *Proc. POPL '97*, pages 1–14. ACM Press, 1997.
- [120] M. Sharir and A. Pnueli. Two approaches to interprocedural data flow analysis. In S. S. Muchnick and N. D. Jones, editors, *Program Flow Analysis*. Prentice Hall International, 1981.
- [121] O. Shivers. Control flow analysis in Scheme. In *Proc. PLDI '88*, volume 7 (1) of *ACM SIGPLAN Notices*, pages 164–174. ACM Press, 1988.
- [122] O. Shivers. Data-flow analysis and type recovery in Scheme. In P. Lee, editor, *In Topics in Advanced Language Implementation*, pages 47–87. MIT Press, 1991.
- [123] O. Shivers. The semantics of Scheme control-flow analysis. In *Proc. PEPM '91*, volume 26 (9) of *ACM SIGPLAN Notices*. ACM Press, 1991.
- [124] J. H. Siekmann. Unification theory. *Journal of Symbolic Computation*, 7:207–274, 1989.
- [125] G. S. Smith. Polymorphic inference with overloading and subtyping. In *Proc. TAPSOFT '93*, volume 668 of *Lecture Notes in Computer Science*, pages 671–685. Springer, 1993.
- [126] G. S. Smith. Polymorphic type schemes for functional programs with overloading and subtyping. *Science of Computer Programming*, 23:197–226, 1994.
- [127] B. Steensgaard. Points-to analysis in almost linear time. In *Proc. POPL '96*, pages 32–41. ACM Press, 1996.
- [128] D. Stefanescu and Y. Zhou. An equational framework for the flow analysis of higher order functional programs. In *Proc. LFP '94*, pages 318–327, 1994.
- [129] B. Steffen. Generating data flow analysis algorithms from modal specifications. *Science of Computer Programming*, 21:115–239, 1993.
- [130] J. Stransky. A lattice for abstract interpretation of dynamic (lisp-like) structures. *Information and Computation*, 1990.
- [131] J.-P. Talpin and P. Jouvelot. Polymorphic Type, Region and Effect Inference. *Journal of Functional Programming*, 2(3):245–271, 1992.

- [132] J.-P. Talpin and P. Jouvelot. The type and effect discipline. In *Proc. LICS '92*, pages 162–173, 1992.
- [133] J.-P. Talpin and P. Jouvelot. The type and effect discipline. *Information and Computation*, 111(2):245–296, 1994.
- [134] Y.-M. Tang. *Control-Flow Analysis by Effect Systems and Abstract Interpretation*. PhD thesis, Ecole des Mines de Paris, 1994.
- [135] M. Tofte. Type inference for polymorphic references. *Information and Computation*, 89:1–34, 1990.
- [136] M. Tofte and L. Birkedal. A region inference algorithm. *ACM TOPLAS*, 20(3):1–44, 1998.
- [137] M. Tofte and J.-P. Talpin. Implementing the call-by-value lambda-calculus using a stack of regions. In *Proc. POPL '94*, pages 188–201. ACM Press, 1994.
- [138] M. Tofte and J.-P. Talpin. Region-based memory management. *Information and Computation*, 132:109–176, 1997.
- [139] J. Vitek, R. N. Horspool, and J. S. Uhl. Compile-Time Analysis of Object-Oriented Programs. In *Proc. CC '92*, volume 641 of *Lecture Notes in Computer Science*, pages 236–250. Springer, 1992.
- [140] A. B. Webber. Program analysis using binary relations. In *Proc. PLDI '97*, volume 32 (5) of *ACM SIGPLAN Notices*, pages 249–260. ACM Press, 1997.
- [141] R. Wilhelm and D. Maurer. *Compiler Design*. Addison-Wesley, 1995. — [WiMa95].
- [142] R. P. Wilson and M. S. Lam. Efficient context-sensitive pointer analysis for C programs. In *Proc. PLDI '95*, pages 1–12. ACM Press, 1995.
- [143] A. K. Wright. Typing references by effect inference. In *Proc. ESOP '92*, volume 582 of *Lecture Notes in Computer Science*, pages 473–491. Springer, 1992.
- [144] A. K. Wright and M. Felleisen. A syntactic approach to type soundness. *Information and Computation*, 115:38–94, 1994.