Smart Sensor Network for Space Vehicle Monitoring

Mark Hedley
CSIRO Australia
PO Box 76, Epping NSW 1710
Australia
+61-2-9372 4136
Mark.Hedley@csiro.au

Mark Johnson
CSIRO Australia
PO Box 76, Epping NSW 1710
Australia
+61-2-9372 4628
Mark.Johnson@csiro.au

Chris Lewis
CSIRO Australia
PO Box 76, Epping NSW 1710
Australia
+61-2-9413 7129
Chris.Lewis@csiro.au

David Carpenter
CSIRO Australia
PO Box 76, Epping NSW 1710
Australia
+61-2-9413 7604
David.Carpenter@csiro.au

Howard Lovatt
CSIRO Australia
PO Box 76, Epping NSW 1710
Australia
+61-2-9413 7412
Howard.Lovatt@csiro.au

Don Price
CSIRO Australia
PO Box 76, Epping NSW 1710
Australia
+61-2-9413 7447
Don.Price@csiro.au

ABSTRACT

A multi-disciplinary team at CSIRO has begun developing concepts for Integrated Vehicle Health Monitoring (IVHM) systems for space vehicles as part of the NASA Robust Aerospace Vehicle Program (RAV). This paper describes the architectural design of the 'concept demonstrator', which is a modularized smart sensor network for detecting the location, energy and flux of micro-meteoroid impacts on the skin of a space vehicle. This paper will concentrate on design of the signal processing and communications hardware and protocols for the concept demonstrator.

Categories and Subject Descriptors

C.2.1 Network Architecture and Design

C.2.2 Network Protocols

C.2.4 Distributed Systems

General Terms

Algorithms, Measurement, Documentation, Performance, Design, Reliability, Experimentation, Verification.

Keywords

Vehicle health monitoring, sensor arrays, digital signal processing, systolic array, communications protocols, electronics, agent technology.

Conference '00, Month 1-2, 2000, City, State.

1. INTRODUCTION

A multi-disciplinary team at Australia's Commonwealth Scientific and Industrial Research Organisation (CSIRO) has been developing concepts for Integrated Vehicle Health Monitoring (IVHM) systems for space vehicles as part of the NASA Robust Aerospace Vehicle Program (RAV). RAV is a long-term program with an ultimate goal of the development of vehicles capable of structural self-assessment and repair, and structural reconfiguration. The purpose of the IVHM system is to detect and measure quantities such as structural fatigue and impact damage, and to use the information to make intelligent, forward-looking decisions and initiate actions. A practical IVHM system will be a network of thousands, or perhaps millions, of heterogeneous sensors, and the system needs to be sufficiently robust that it can continue to operate effectively with significant damage to the vehicle and to the IVHM system itself.

CSIRO's initial work was to develop ideas for an IVHM system, and the reports on this phase of the work have been released [1,2]. Part of the current work that CSIRO is undertaking for NASA is to develop a smart sensor network, called the Concept Demonstrator (CD), that will be used as a test bed and demonstrator of IVHM technology. The Concept Demonstrator is a structure containing a large array (over one thousand) of sensors (initially for detecting impact damage), along with signal acquisition hardware, a robust communications network and distributed processing of the signals. The CD has been initially designed to detect impact damage on an aerospace 'skin', and to measure the location and approximate energy of each impact. This test bed allows experimentation in the sensing technology, communications topologies and protocols, and in particular in the agent technology and distributed processing that implements the system intelligence. We are particularly interested in testing the robustness of the IVHM system to continue monitoring the structure while portions of the IVHM system are induced to fail.

The overall architecture of the CD will be described in the following section, and then the rest of the paper will be devoted to describing the signal processing and communications aspects of the design.

2. CD ARCHITECTURE

2.1 Overview

The Concept Demonstrator (CD) is a laboratory prototype and test-bed for detecting and measuring impacts on the aluminium skin of a space vehicle, where the sensing is performed using an array of acoustic emission sensors mounted on the inside of the skin. The system must be able to handle a large number of impact events, and it must be robust to damage to the sensors and monitoring system, which in turn requires redundancy and a degree of reconfigurability to work around damage in the system.

The ultimate purpose of a space vehicle IVHM system is to detect damage and to determine an appropriate response to the damage in real-time. In some cases the response need only be local to the damaged area, but in other cases a more global response may be required (e.g. move the space vehicle or deploy repair robots). Although the CD will not (at this stage) have the ability to undertake actions to repair damage, it will be a test-bed to develop systems and algorithms to not only detect damage but also to assess the response required and ensure the information is communicated to the appropriate levels. Agent-based software technology will be used to develop and implement these algorithms. The CD will be a platform on which heterogeneous communicating agents may be deployed, but there will be no further discussion of these agents in this paper.

The CD will be constructed with over one thousand acoustic emission sensors, and processing the large amount of data produced requires a large number of processors. Furthermore, to ensure that the system is robust requires the use of multiple processors in such a way that the failure of some of these processors results in only a gradual degradation of the performance of the IVHM system. To satisfy these requirements the hardware consists of a distributed system of communicating processors, with each processor communicating with multiple neighbours. The physical organisation of the processors is purely distributed, not hierarchical, so that the failure of one processor will not compromise the use of any working processor, and the failure of a communications link may lead to the rerouting of communications, but will not isolate any processors. This does not prevent a logical hierarchy of processors being used for the distributed processing, but ensures that should a processor higher in the hierarchy fail then another processor can be promoted to replace the failed processor's role.

2.2 Modules and Layers

The CD is being designed as a network of sensors with distributed processing. The electronics needs to be conformal to the surface it is monitoring, it needs to be able to be flexibly deployed, and it should be designed for ease of manufacture, testing and repair. From these considerations it was decided to use a modular construction where a small number of sensors and processors are packaged into a physical module that can communicate with neighbouring modules.

There is a logical progression of layers in the hardware:

1. Physical outer skin – this is protective and possibly structural.
2. Sensors – initially these are piezoelectric PVDF acoustic emission sensors attached to the skin, but other types of sensors could also be used.
3. Conditioning Electronics – this filters and converts the signals from the sensors to a form suitable for digitizing.
4. Sampling and data pre-processing – this includes digitization, calibration correction and time-stamping of the signals, and possibly pre-processing for data reduction or other purposes.
5. Data analysis – at this level data is processed not only from the local sensors, but also using information from neighbouring modules. This is the layer at which agent based software will run, which provide the 'smart' processing for the smart sensor network.
6. Inter-module communications stack – this provides the communications between modules.
7. Physical inter-module communications link – this provides the hardware for NAL communications.

These layers are grouped into a Data Acquisition Layer (DAL) which consists of layers 3 and 4, and a Network Application Layer (NAL) which consists of layers 5, 6, 7. Separate processors will be used in the DAL and the NAL, with the respective processors being called the Data Acquisition Processor (DAP) and Network Application Processor (NAP). By using separate processors on separate sub-modules, it is possible use another type of sensor while keeping the same analysis and communications infrastructure. The two layers, and the links between them, are shown in Figure 1 for an array of three by three modules.

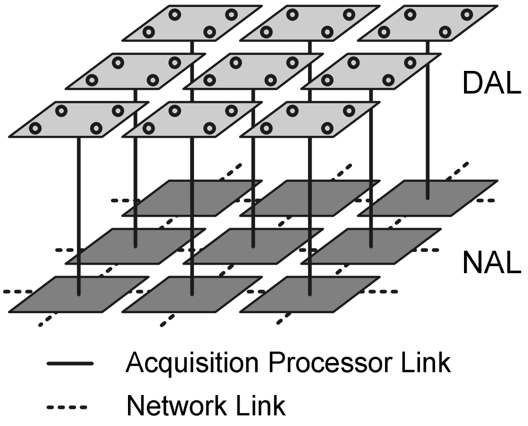

Figure 1: Layer structure and communications.

This paper will not provide any further information on layers 1, 2 or 3, nor will it discuss the mechanical or electrical interconnection aspects of the CD. Rather, this paper will focus on the design of the processing systems and communications architecture and protocols.

2.3 Module Organisation

Our current estimates for the required module capabilities are:

1. Following signal conditioning, data acquisition from the acoustic emission sensors requires roughly one million samples per second for each sensor.
2. Processing roughly 100 MIPS per sensor.
3. Communications roughly 1 Mbit/sec per sensor.

Two sizes of acoustic emission sensors will be used, and there is a natural grouping of five sensors. Given the required capabilities this is a good match to using a single DAP in the DAL and a single NAP in the NAL. Thus a module will consist of two sub-modules:

- DAL sub-module containing five sensors, conditioning electronics, one DAP and a link to a NAP.
- NAL sub-module containing a link to a single DAP, one NAP and network links.

An important decision for the module organization was the shape of the modules and the topology of the communications between them. Two shapes were considered, triangular and square modules. While triangles are better for conforming to three-dimensional surfaces, it was decided to use square modules for ease of fabrication. The most natural communications topology was then to have one communications port on each side of the square to communicate with a neighbour, which results in a mesh network. Thus, each NAL sub-module requires four communication ports. This is shown in Figure 1.

We are planning to produce 192 modules, which will contain a total of 960 sensors, 192 data acquisition processors and 192 network application processors. For the sensor spacing to be used this will monitor 1.92 m^2 of the space vehicle skin.

3. DATA ACQUISITION LAYER

The data acquisition layer needs to digitise signals from five sensors at a rate of at least one million samples per second (MSPS), thus must be capable of at least 5 MSPS. The processor needs to have the capability to perform at least simple filtering and pre-processing of the acquired data, thus needs to be able to execute tens of instructions per acquired sample, and hence requires a computational capability of the order of 100 MIPS (million instructions per second). The software that runs on this processor will be relatively simple, so large amounts of memory are not required, and having on-chip memory would minimise the external component count. The processor selected for the DAP was the Texas Instruments TMS320F2810, which has the following properties:

- 150 MHz, 150 MIPS 32-bit fixed point digital signal processor. This provides adequate computational performance.
- 16 multiplexed analogue input channels sharing two sample and hold circuits and a single 12-bit 16.6 MSPS pipelined analogue to digital converter (ADC).
- 36 kbytes RAM and 128 kbytes FLASH memory on-chip. This should be sufficient so that no external memory is required in the DAL.
- High-speed serial interface (McBSP) to allow rapid communication with the NAP. This supports data rates up to 75 Mbps.
- Other features, including timers and oscillators, that allow the sub-module to run with very little additional logic.
- Low power – 1.8 V core and dynamic clock control. Maximum power consumption is less than 1 W.
- Low cost development system available (TI F2812 eZdsp).

While other processors were considered, the F2810 had the best overall performance at a reasonable cost (about US$40 each). Should more memory be required, the TMS320F2812 is a similar processor with twice the FLASH and an external bus.

4. NETWORK APPLICATION LAYER

4.1 Layer Functions

The requirements of the NAL sub-module are:

- Communications to the DAP, with a standard hardware and software interface so that DAP modules with different functionality can be deployed.
- Communications to neighbouring NAP sub-modules. This includes not only point-to-point communications between modules but also routing through the mesh. This should be flexible so that the CD can be used to experiment with networking protocols.
- Data processing, particularly providing a platform for experimenting with agent-based processing algorithms. The processor requires processing capability of hundreds of MIPS (with 4 ports at 1 Mbps, 1000 instructions per byte requires 400 MIPS), and sufficient memory (at least hundreds of kilobytes of ROM and megabytes of RAM).

4.2 Processor Selection

The processor selected for the NAP was the Texas Instruments TMS320VC5509, which has the following properties:

- 200 MHz, 400 MIPS 32-bit fixed point digital signal processor.
- 256 kbytes RAM, *no user programmable ROM*, and glueless interface to external FLASH and SDRAM.
- McBSP port for communications with DAP.
- *Does not have four homogeneous serial ports.*
- Low power consumption, down to 0.25 mW/MIP, and with a 1.6V core.
- Each processor has 64-bit unique identifier.
- Two 20-bit timers, watchdog timer and real-time clock.

The selected processor is not ideal, but was the best compromise that could be found. The processor family has available a low cost development system (C5510 DSP Starter Kit), has a wide range of third-party software and IP, and each processor costs approximately US$25. The NAL sub-module will require external FLASH and SDRAM memory, but the glueless interface makes this simple to interface.

External hardware will also be required for the serial data links to the neighbouring sub-modules. For low-cost and simplicity of implementation (hardware and software) it was decided to use a four channel UART (universal asynchronous receiver transmitter) device, and the Texas Instruments TI16C754 was selected. This provides four ports, data rates up to 2 Mbps (at 3.3 V supply), and 64 byte transmit and receive buffers.

4.3 Communications Architecture

The CD not only implements an IVHM system as a laboratory prototype, but it is also a test-bed, and as such the communications required are not just for communication between the agents, but also for diagnostics, debugging and monitoring the operation of the system. Furthermore, downloading code into the large number of processors using the JTAG is slow, so the NAL will also be used to distribute code which can be loaded into RAM or FLASH on the NAL sub-modules. To further complicate the communications architecture, the software needs to have the flexibility to enable us to experiment with network protocols. To support these requirements the communications stack contains fixed code programmed into each NAL sub-module, and the ability to download and use new network protocols and applications. The software stack is illustrated in Figure 2.

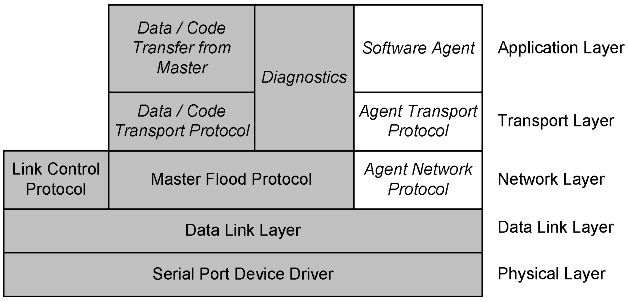

Figure 2: Concept demonstrator communications stack.

The lowest layer of the communications stack is the device driver to control the physical ports. Above this is the Data Link Layer (DLL), which is a point to point protocol between the two processors on a physical link. The protocol is similar to PPP [3,4], and supports multiple network protocols running simultaneously. The DLL has two network protocols fixed in each sub-module, but also allows other network protocols to be downloaded and registered for use. Each network protocol has a unique 8-bit identifier, and once a protocol is registered with the DLL received packets with the protocol identifier are passed to the registered network protocol handler.

The Link Control Protocol (LCP) is a special network protocol used to control the DLL. It is the LCP that determines when a processor is connected or disconnected to a port on the sub-module, negotiates the port speed to use when a connection is detected, and monitors the link.

The Master Flood Protocol is intended to provide simple but reliable communication between a master processor (probably a PC used to monitor and control the CD) and all processors in the NAL. It is not intended for communications between processors within the NAL. The master processor can send packets either to all processors or to single processors, and the other processors can send packets to the master processor. The routing of packets is based on flooding, and by being simple this software can be more readily implemented in a small memory footprint. This protocol is intended to be used for distributing software using the NAL, and to obtain diagnostic information about the operation of the CD. This protocol would not be used for agent communication.

The network protocol (or protocols) used for agent communication would not be programmed into each NAL, rather they would be downloaded (along with the agent software) and registered. This is done as these software components are expected to change frequently as the CD test-bed is used to experiment with agent algorithms.

5. CONCLUSIONS

This paper has described the hardware and communications architecture of a smart sensor network system that is currently being built to monitor space vehicles for impact damage.

6. ACKNOWLEDGEMENTS

Signal processing and communications are just part of the Concept Demonstrator project work, and other components include the development of the sensor and agent technologies, software simulators and mechanical design. We wish to acknowledge other members of the project, without whom the work presented in this paper would not have been possible, and the project members who are not authors of this paper are: D. Abbott, J. Ables, A. Batten, A. Collings, B. Doyle, J. Dunlop, G. Edwards, T. Farmer, P. Isaacs, G. James, K. James, B. Martin, G. Poulton, M. Prokopenko, T. Reda, D. Rees, A. Scott, S. Seneviratne, P. Valencia, D. Withnall, P. Wong and J. Winter.

7. REFERENCES

[1] D. Abbott, S. Cunningham, G. Daniels, B. Doyle, J. Dunlop, D. Economou, T. Farmer, D. Farrant, C. Foley, B. Fox, M. Hedley, J. Herrmann, C. Jacka, G. James, M. Johnson, B. Martin, G. Poulton, D. Price, T. Reda, G. Rosolen, A. Scott, V. Valencia, D. Ward, J. Winter and A. Young, "Development and Evaluation of Sensor Concepts for Ageless Aerospace Vehicles – Threats and Measurands," *NASA/CR-2002-211772*, July 2002, pp. 25.

[2] David Abbott, Briony Doyle, John Dunlop, Tony Farmer, Mark Hedley, Jan Herrmann, Geoff James, Mark Johnson, Bhautik Joshi, Geoff Poulton, Don Price, Mikhail Prokopenko, Torsten Reda, David Rees, Andrew Scott, Philip Valencia, Damon Ward and John Winter, "Development and Evaluation of Sensor Concepts for Ageless Aerospace Vehicles – Development of Concepts for an Intelligent Sensing System," *NASA/CR-2002-211773*, July 2002, pp. 138.

[3] W. Simpson, "The Point-to-Point Protocol (PPP)", RFC1661, http://www.faqs.org/rfcs/rfc1661.html.

[4] W. Simpson, "PPP in HDLC-like Framing", RFC1662, http://www.faqs.org/rfcs/rfc1662.html.