Simulation of heterogeneous and interdependent critical infrastructures

S. De Porcellinis and R. Setola*

Complex Systems & Security Lab.
University CAMPUS Bio-Medico
via E. Longoni 83, 00155, Rome, Italy
E-mail: s.deporcellinis@unicampus.it
E-mail: r.setola@unicampus.it
*Corresponding author

S. Panzieri and G. Ulivi

Dip. Informatica e Automazione
University 'Roma Tre', Italy
E-mail: panzieri@uniroma3.it
E-mail: ulivi@uniroma3.it

Abstract: In this paper, a simulation tool specifically designed for the analysis of heterogeneous and (inter)dependent infrastructures is proposed. The simulator, named Critical Infrastructure Simulation by Interdependent Agents (CISIA), adopts a modular and sufficiently abstract representation of the different infrastructures' components to allow consistent descriptions, starting from the incomplete and generic data acquirable from stakeholders. An important part of the modelling effort was reserved for the representation of the dependencies and interdependencies, these being the cause of the complex behaviours we are interested in. Each component interacts with the others via a multitude of mechanisms that codify different concepts of proximity. The simulator has been used to analyse, in a simplified scenario, crisis evolution in the urban area of Rome, in the presence of a failure in the electric power system.

Keywords: interdependent infrastructures; complex systems simulation; network analysis; heterogeneous systems.

Reference to this paper should be made as follows: De Porcellinis, S., Setola, R., Panzieri, S. and Ulivi, G. (2008) 'Simulation of heterogeneous and interdependent critical infrastructures', *Int. J. Critical Infrastructures*, Vol. 4, Nos. 1/2, pp.110–128.

Biographical notes: Stefano De Porcellinis received his Laurea degree in Computer Science Engineering from the University 'Roma Tre' in Rome in 2005. He is a PhD student at University CAMPUS Bio-Medico of Roma. His research interests include critical infrastructures modelling, simulation environments and fuzzy control techniques for biomedical devices.

Roberto Setola received his Laurea degree in Electronic Engineering (1992) and his PhD in Electronic Engineering and Computer Science (1996) from the University of Naples. From 1999 to 2004, he worked at the Italian Prime Minister's Office. At present, he is an Assistant Professor of Automatic

Copyright © 2008 Inderscience Enterprises Ltd.

Control at University CAMPUS Bio-Medico of Roma. He is the Technical Representative of the Italian Government Working Group on Critical Infrastructure Protection, and a member of the G8 Senior Experts' group for Critical Information Infrastructure Protection. His research interests include critical infrastructures modelling and control, control of complex systems, nonlinear estimation, mobile robot cooperation and development of biomedical systems. He is the author of more than 90 scientific papers.

Stefano Panzieri received his Laurea degree in Electronic Engineering in 1989 and his PhD in Systems Engineering in 1994, both from the University of Roma 'La Sapienza'. Since February 1996, he has been with the 'Dipartimento di Informatica e Automazione' of the University of 'Roma Tre', where he is currently an Associate Professor. His research interests are in the field of industrial control systems. He has published several papers concerning the study of iterative learning control applied to robots with elastic elements and to nonholonomic systems. In particular, in the area of mobile robots, he studied the problem of localisation in structured and unstructured environments, with special attention to the problem of sensor-based navigation.

Giovanni Ulivi received his Laurea degree in Electrical Engineering from the University of Rome 'La Sapienza' in 1974. Starting that year, he was with the Department of Computer and System Science of the same university. In 1992, he moved to the University of Rome 'RomaTre' and is now head of the Department of Computer Science and Industrial Automation. He teaches courses in automatic control and fuzzy control and is tutor to several PhD students. His research interests, begun with the control of electric motors, now include robotics (in particular autonomous vehicles control) and complex systems. He is the author of more than 100 scientific papers. He is a member of IEEE and the International Federation of Automatic Control (IFAC).

1 Introduction

The welfare of large segments of the population in developed countries depends on many technological infrastructures, such as energy production, transportation and distribution; telecommunications networks; water management and supply networks; transportation (air, rail, marine, surface); banking and financial services (Dunn and Wigert, 2006; US Government, 2003).

Owing to their relevance, they are generally indicated as *Critical Infrastructures*, because they, "if disrupted or destroyed, would have a serious impact on the health, safety, security or economic well-being of Citizens or the effective functioning of governments" (EU Commission Communication 576, 2005).

In the very last years, for a lot of economical, social, political and technological reasons, we observed a rapid change in the organisational, operational and technological structures of these infrastructures. Indeed, to reduce costs, improve efficiency and provide innovative services, infrastructures have become more and more interoperable and have extensively adopted Information and Communication Technologies (ICTs). Moreover, to effectively and profitably operate in the global market in the absence of monopolistic privileges, infrastructures' stakeholders have started to focalise on their core business and then to use outsourcing strategies.

Although designed as logically separate systems delivering different services, these infrastructures have become highly dependent and interdependent, each relying (directly or indirectly) on the services provided by the others. Those connections contribute to create a very huge and complex system of systems that appears prone to large cascade failures. Indeed, owing to the presence of interdependencies, a failure in any subsystem may easily propagate to the others, with the result of affecting a large, unpredictable and geographically distributed sets of users.

As an example, on December 1998 the failure of the telecommunications satellite Galaxy IV in geo-stationary orbit on the US West Coast (Rosenbush, 1998), besides creating communication problems (almost 90% of US pagers were affected), led to significant difficulties in the transportation system: numerous flights were delayed owing to the absence of high-altitude weather information, while refueling on highways became difficult as gas stations could not process credit card transactions (for more information on interdependency-related incidents, see Bologna and Setola (2005)).

Another example about negative effects of interdependencies can be discovered by analysing what happened in Italy during the blackout of 2003. Specifically, there was a considerable delay in power recovery mechanisms, owing to the cascade failures of the telecommunications systems: SCADA operators were not able to telecontrol the generator plants and had to resort to manual procedures for restarting (waiting also for the time required by the technicians to reach the installations) (PIC – Italian Government Working Group on Critical Information Infrastructure Protection, 2004).

A last citation, regarding the failure that happened, on January 2004, to the air conditioning system of an important telecom node near Rome (PIC – Italian Government Working Group on Critical Information Infrastructure Protection, 2004). This failure caused a widespread blackout of land and wireless telecommunications (affecting almost all the service providers), the quitting of financial transactions in 5.000 banks and in 3.000 postal offices, and also difficulties at the international airport, where about 70% of check-in desks were forced to use manual procedures.

Besides being vulnerable to accidental failures, critical infrastructures are more exposed to natural disasters, especially owing to the extremisation of climate events. Moreover, because of their increased relevance, they may become targets of terroristic and criminal actions (OCIPEP, 2003). Indeed, attacks could be carried out against infrastructures to create damage, panic or mistrust, or even to increase the effects of acts against more visible targets, *e.g.*, slowing down emergency services and delaying rescue operations (US General Accounting Office, 2003).

Such conditions imposed on governments and international organisations to improve the security, robustness, resilience and 'plasticity' of these infrastructures (UN 58th Generally Assembly, 2003; Dunn and Wigert, 2006). To this end, the actual scenario also imposes the need to consider, besides intrasectorial (vertical) strategies for security, the definition of coordinated (horizontal) intersectorial strategies, with the aim of integrating into a single framework security requirements and constraints about critical infrastructures as a whole. These strategies are usually referred to as Critical Infrastructure Protection (CIP), and Critical Information Infrastructure Protection (CIIP) when the focus is on the ICT component.

Obviously, to develop such strategies we need methods and tools to understand and foresee the global behaviour of these systems, in particular when they are forced to operate in critical situations. This is a very hard challenge (US Government, 2005). Any critical infrastructure is a complex, highly nonlinear, geographically dispersed cluster of

systems, interacting with their human owners, operators and users. The complexity of these systems has grown to a point where, as stressed in Amin (2002), there is no hope of applying the existing standard methodologies, owing to the presence of many dependency and interdependency links with other infrastructures. Moreover, many of these dependencies are very often hidden or poorly understood.

In this paper, we illustrate the Critical Infrastructure Simulation by Interdependent Agents (CISIA) project, with the aim to define an approach specifically devoted to analyse heterogeneous and interdependent infrastructures, oriented to capture the most important phenomena related to failure propagation and performance degradations.

The paper is organised as follows: In the next section, some of the most interesting approaches proposed in the literature for modelling, analysis and simulation of interdependent infrastructures are briefly reviewed. In Section 3, we summarise some key features of our project, and improve the definition of its goals, scope and limitations. Section 4 is devoted to describe the proposed simulator and Section 5 presents the analysis of a failure in the electric power system in the urban area of Rome. Finally, some conclusive remarks are collected in Section 6.

2 Independent infrastructures modelling

Modelling procedures and simulation techniques for individual infrastructures represent a rather well-developed field. Numerous products are commercially available to analyse each single infrastructure at different abstraction degrees, on multiple time scales and with a prescribed level of detail. On the other hand, the modelling and simulation of multiple, interdependent and heterogeneous infrastructures are still immature techniques, but a number of approaches are under development. Some approaches are very qualitative and mainly used to identify the infrastructures that might be considered 'critical', but not to single out their role in the global framework. For example, the Quick-Scan project, supported by the Dutch government (Luiijf *et al.*, 2003), is aimed at obtaining answers to the following questions: What are the sectors, products and services critical to The Netherlands? What are the underlying processes? What are the (inter)dependencies?

The same questions have been investigated in Moteff *et al.* (2002), where the authors emphasise that none of the definitions given over the years about what constitutes a critical infrastructure may be considered exhaustive: they appear too ambiguous and open to interpretation. Moreover, we have to consider that any infrastructure is completely isolated from the others, and also that, even in the same country, the notion of what is critical changes over time. They stress that, in any case, it is mandatory to identify the actual critical elements inside the different infrastructures.

Considering each infrastructure as a unicum represents a very crude simplification that does not take into account its geographical extension and its structure. Indeed, as noted in Rinaldi *et al.* (2001), interactions among different components produce the emergence of behaviours that are not predictable from the knowledge of any single isolated part. This suggests the adoption of a bottom-up approach, as largely done when we have to deal with scarce or ill-defined macroscale information, like in the bio-complexity researches.

Following this bottom-up philosophy, in Setola and Ulivi (2003) the authors proposed to decompose each infrastructure into its macrocomponents and to analyse the failure propagation among them. To manage the uncertainties about the partially unknown interdependency phenomenon, the level of inoperability associated with each macrocomponent is represented by fuzzy values.

In MacDonald and Bologna (2001), it is emphasised that, to correctly understand the behaviour of these infrastructures, it is mandatory to adopt a three-layer model:

1 Physical layer – the physical component of the infrastructure, *e.g.*, the grid for the electrical network.

2 Cyber layer – hardware and software components of the system devoted to control and manage the infrastructure, *e.g.*, SCADA and DCS.

3 Organisational layer – procedures and functions used to define activities of human operators and to support cooperation among infrastructures.

Here, the authors emphasise that each component of an infrastructure largely interacts with elements in the same layers of other infrastructures, besides elements belonging to the same infrastructure, by means of *interdependency* links. The increasing presence of these links creates many functional dependencies among infrastructures. Moreover, the authors emphasise that, with respect to ten years ago, the importance of the cyber layer is largely increased, becoming one of the most important sources of interdependencies. Notice that a similar kind of decomposition was also used to analyse the 2003 blackout in the USA and Canada (US and Canada Power System Outage Task Force, 2004). As a matter of fact, to explain the multitude of causes that produced that episode, the USA and Canada government commission described the event in terms of grid (physical), computer and human layers. Only by considering all the layers together is it possible to correctly understand what really led to the blackout.

Going further into detail, Rinaldi *et al.* (2001) emphasise how interdependencies should be analysed with respect to different dimensions. In particular, they catalogue interdependencies into four not mutually exclusive classes:

1 Physical interdependency – Two infrastructures are physically interdependent if the operations of one infrastructure depend on the physical output(s) of the other.

2 Cyber interdependency – An infrastructure presents a cyber interdependency if its state depends on information transmitted through the information infrastructure.

3 Geographical interdependency – A geographical interdependency occurs when elements of multiple infrastructures are in close spatial proximity. In this case, particular events, such as an explosion or a fire in an element of an infrastructure, may create failures in one or more nearby infrastructures.

4 Logical interdependency – Two infrastructures are logically interdependent if the state of each one depends on the state of the other via control, regulatory or other mechanisms that cannot be considered physical, geographical or cyber.

Cyber interdependency is a relatively new phenomenon, strictly related to the pervasiveness of ICT and the integration of computerised control systems (*e.g.*, SCADA) with other information systems, further stressed by the use of public communication networks. Rinaldi *et al.* (2001) stress that cyber interdependency tends to become

an absolute and global characteristic of all the infrastructures, while other types of interdependencies are more local. Cyber interdependency potentially couples an infrastructure with every other infrastructure that uses the cyberspace, in spite of their nature, type or geographical location. Incidentally, note that considerations of this kind suggested to many governments to pay great attention to cyber threats, specifically to the vulnerability induced in critical plants by SCADA systems exposed to cyberspace (US Government, 2003).

A different aspect of the problem is connected with the topology of the networks and their robustness. The pioneer works by Watts and Strogatz (1998) and by Albert *et al.* (2000) have emphasised some peculiarities common to many networks, and among them to technological infrastructures, never pointed out before. Specifically, the presence of hubs (*i.e.*, nodes connected with a large number of other nodes) increases the robustness with respect to accidental failure, but makes the network prone to deliberate attacks (Albert and Barabasi, 2002). The studies on the role played by the topological structure in the framework of interdependent infrastructures (see Newman *et al.*, 2005) are only preliminary, but their relevance to failure analysis can be predicted.

3 Guidelines for the simulator design

A visionary project about the simulation of critical infrastructures is under development at the National Infrastructures Simulation and Analysis Centre (NISAC). This centre, established by Los Alamos and Sandia Laboratories, aims to model and to simulate all the infrastructures that are critical to the USA, with their interdependencies. To this end, NISAC is developing different suites composed of interoperable modules able to support very detailed analysis.

For example, in the Urban Infrastructure Suite (UIS), they integrate seven interoperable modules that employ advanced modelling and simulation methodologies to represent urban infrastructures and populations. Each one of these modules adopts very detailed models; for example, the Transportation Analysis Simulation System module simulates the daily activities and movements of a cluster of individuals in an urban region statistically representing the actual population. During their movements, people impose demands on transportation, telecommunications, *etc.*, and in the presence of an epidemiological event (*e.g.*, a biological attack), the people's movement will affect the epidemic spread.

To realise this simulator, NISAC and the US government are trying to collect huge quantities of information about each critical infrastructure. Unfortunately, more details about NISAC activities are not available.

It is evident that setting up simulators, such as those proposed by NISAC, is a very hard task, both for the technological challenges and the extreme difficulty to acquire detailed information from stakeholders. Moreover, it is mandatory to guarantee continuous updating of these data, in order to avoid incoherent views with catastrophic effects.

From the previous considerations, we derived some desirable characteristics for the simulator, taking also into account that only a limited amount of resources is available, both in terms of computing power and software complexity, and in terms of available or to be elicited information:

- Each infrastructure should be modelled starting from its macrocomponents, *i.e.*, objects with specific and easily recognisable roles. The global behaviour should be figured out from interactions.
- To reduce the need for detailed information, each element should be defined with a sufficiently high level of abstraction that permits consistent descriptions, starting from the incomplete and generic data acquirable from stakeholders and open documents. This is also instrumental for the next point.
- The external representation of all the elements (of the same or different infrastructures) should be kept as uniform as possible to ease the descriptions of the networks and the coding of the interfaces.
- Fuzzy numbers should be used to code parameters and values. In this way, not only is it possible to represent vague statements such as 'component B depends very much on component A', but the results of simulations can also be analysed in terms of their reliability.
- The description of the macrocomponents should depend only on internal parameters and the values explicitly exchanged with other blocks.
- The simulator should not impose any limitation on the representable behaviours, to leave the experts free to use the most suitable descriptions for the macrocomponent.

These considerations have inspired an approach to capture the most significant behaviours related with CIP topics, starting from fragmented, nonhomogeneous and ambiguous information. To this end, we decompose each infrastructure into its macrocomponents in order to generate a 'lumped' model composed of n elements, where n depends on spatial/temporal scale and on the required detail level.

Each of these macrocomponents is characterised by of its capability to correctly perform its own task (*i.e.*, to produce a given amount of goods or services) and its level of failure (actually, we consider that each macrocomponent may be affected by different types of failures, each one with an arbitrary level of severity).

Because we are interested in analysing the system during a crisis scenario, we ignore any internal physical dynamic (that we suppose to be adequate in normal situations) and adopt an abstract description for the behaviours of the macrocomponents based on information related to resource availability and failures level. In this framework, we can easily model interactions among macrocomponents through the exchange of a small and common set of values, representing the different types of resources and failures considered. Notice that we have to consider multiple types of failures because, as explained in the next section, failure type greatly influences the mechanism of spreading and the impact on the different macrocomponents.

The scenario is then described by means of n macrocomponents, which exchange p types of resources and m types of failures. For each type of resource or failure, we have to consider peculiar mechanisms for their diffusion, considering different concepts of proximity and then different sets of neighbours.

In our model, variables and parameters are expressed by Fuzzy Numbers (FNs), which can be seen as the most natural way to introduce model and data uncertainty in technical talk. Consider the simple statement we could obtain from an expert: 'The emergency supply will last for about two days.' According to this highly qualitative

information, we may represent the value 'about two days' with one of the diagrams in Figure 1. Here the so-called triangular fuzzy numbers are assumed. In particular, the first plot (continuous curve) represents 'I am maximally confident that it can last two days and I am sure it cannot last less than one day or more than three', while the second plot (dotted curve) shows less uncertainty.

Figure 1 Triangular Fuzzy Numbers (FNs) representation

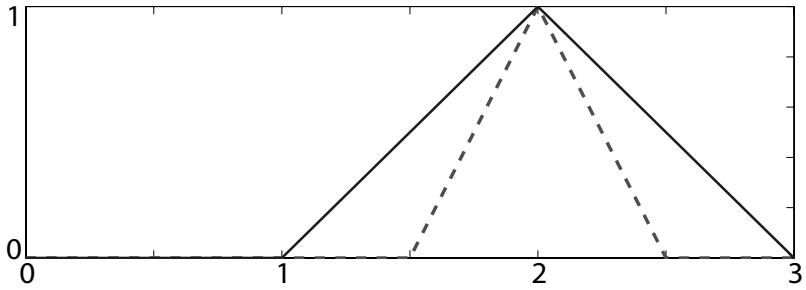

Here, fuzzy numbers are used mainly as an extension of interval arithmetic. Actually, they represent our belief in a given assertion and are more flexible to use than the *a priori* probability. A complete analysis of the relation between probability and fuzzy measures can be found in Dubois and Prade (1998) or in Appendix A of Ross (2004).

4 CISIA

Using the previous consideration, we have set up a simulation framework, named CISIA, to analyse failure propagation and performance degradation in a system composed of different, heterogeneous and interdependent infrastructures.

In CISIA, each macrocomponent is modelled with a high level of abstraction and, to reach an acceptable level of modularity and scalability, it is self-contained. Then its input-output behaviour is independent of the state of its neighbours, and is based only on the quantities exchanged with them. In other words, each element has no (explicit or implicit) information about which are its neighbours or about their state or internal model.

4.1 Macrocomponent dynamic

Each macrocomponent is modelled as an autonomous block (called 'entity' in the following), whose behaviour is described, at least, by the following quantities:

- Operative Level (*OL*) – represents the capability of the entity to perform its job. It represents only the potential capability, in the sense that an *OL* of 100% does not mean that the system is working at its maximum capacity, but that it could if required. This quantity dynamically varies over time in the range [0,1], and any abnormal situation is identified by the OL's reduction from the unit. It is represented by a FN.

- Failures (F) – is a structured variable that enumerates the types of internal failure that can affect the entity. It also stores the associated levels of severity. While the types of failure are static properties of the entity, the associated levels of severity can dynamically increase (they cannot decrease as we do not take into account fixing or rescue procedures).

Failure severity is normalised in the range [0,1], where 0 represents the absence of failure and 1 its maximum possible amount. Some failures are binary quantities, but others are represented via FN values that allow discriminating between *very small* to *disastrous* failures.

Each entity is characterised also by a set of constant quantities used to better specify its behaviour. Among others, we cite *Requirements* (REQ) and *Nominal Production Level* (NPL), which specify types, units of measurements and nominal value of *Resources* needed to be reached or produced when $OL = 100\%$.

Entities interact only via the exchange of the following quantities:

- Resources (R) – goods and services produced (or used) by the entity, expressed in terms of their types, units of measurements, nominal values and actual levels. Notice that, as for the OL, *Resources* do not represent the effective amount of goods produced (used), but the maximum affordable quantity that the entity could produce (use) if required. It is assumed that the i-th entity produces and uses, respectively, p_i and r_i different types of resources (where the number and types depend on the characteristic of the entity).

- Consumption (C) – resources effectively used by the entity, expressed in terms of their types, units of measurements, nominal values and amount. These quantities are fed back to producers, where they are used to estimate the overall consumption in order to check for saturation effects or overload conditions.

- Failures (F) – failures propagated from (or to) the entity, expressed in terms of their types and the associated levels of severity. It is assumed that the i-th entity is affected by m_i different types of failures (where number and types depend on the characteristic of the entity).

Specifically, an entity's input and output are:

R_{IN} = resources used by the entity
C_{IN} = amount of resources actually required by downstream elements
F_{IN} = failures propagated to the entity
R_{OUT} = maximum resources that the entity can produce if needed
C_{OUT} = actual amount of resources consumed by the entity
F_{OUT} = failures propagated from the entity to other elements.

We base entities' interactions on exchange of *Resources*, instead of actual production, to improve simulation efficiency by reducing the presence of loops. Indeed, even if we adopt a producer-consumer paradigm, such components are partially decoupled because the consumer operates on potential capability of the producer, which, in turn, is independent by the consumer's requests, unless they do not induce overload.

The *OL* of each macrocomponent is set to 100% if there is no internal failure (*F*) and the available resources (R_{IN}) are greater or at least equal to the entity requirements (*REQ*). Otherwise, the entity's *OL* is progressively reduced to zero. When $OL = 0$, the entity is not able to supply any resource, but even in this case, it may still generate and/or transmit failures.

Specifically:

$$OL_i = \prod_{k=1}^{m_i} \left(\Lambda_k \left(F(k) \right) \right) * \prod_{j=1}^{r_i} \Theta_j \left(\frac{REQ(j)}{R_{IN}(j)} \right) \tag{1}$$

where Λ and Θ are functions which assume values in the range [0,1]. Λ is a decreasing function with $\Lambda_k(0) = 1$, while Θ_j is an increasing one with $\Theta_j(1) = 1$.

Inside CISIA, to allow the modelling of the most complete as possible range of behaviours with a common framework, we used functions belonging only to the classes illustrated in Figure 2, or to linear combinations of them. Notice that, linear (a) and threshold (b) functions are generally easier to configure, even if a better fit of experimental data are obtained with logistic curves (c).

Figure 2 Functions used inside the entity models

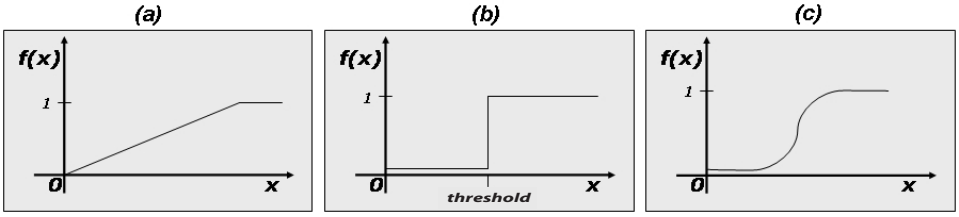

The same classes of curves are adopted for all of the functions used in this section.

The resources produced by the entity is assumed to be proportional to the *OL* via its *NPL*:

$$R_{OUT} = OL * NPL. \tag{2}$$

On the other hand, the entity's consumptions depend on the *OL* and on all the requests coming from other elements.

$$C_{OUT} = \Omega\left(OL,\ C_{IN},\ NSL\right) \tag{3}$$

where *NSL* specifies the *Normal Suppliable Level*, a reference value that represents the amount of resource that the entity can supply with an *OL* of 100% in a steady state. The previous expression for an entity with a single output may be specialised as:

$$C_{OUT} = OL * \frac{C_{IN}}{R_{OUT}} * NSL. \tag{4}$$

Upon the occurrence that the entity receives a failure of type *X* that is 'supported' by the entity, it updates the corresponding internal state by means of:

$$F(X) = \Psi_X \left\{ F(X),\ \Delta\left(F_{IN}(X)\right),\ C_{IN} \right\} \tag{5}$$

where Ψ_X is a suitable function that models how internal failure is influenced by external failures, by overload or other entity-specific causes; Δ represents a suitable function that scales the severity of the input failure in accordance with the entity protection profile.

Internal failures may generate output failure (F_{OUT}), which spreads to neighbouring entities, identified in accordance with the specific concept of proximity.

The vectors R_{OUT}, C_{OUT} and F_{OUT} are passed to the *Transmission Sub System* (TSS), which, as explained in the next section, decomposes them in messages that are sent to the different neighbours.

4.2 Dependencies modelling

An important part of the modelling effort was reserved for the representation of the dependencies and interdependencies, these being the cause of the complex behaviours we are interested in. Each macrocomponent interacts with the others via a multitude of mechanisms. Some of them, related to functional dependencies, may be assumed to be quite well known, because they have been conceived and voluntarily built to directly contribute to the normal operations.

Even if some kinds of failures may be propagated along such paths, we preferred to use logically distinct links for these kind of connections. Typically, over these functional links, the macrocomponents distribute the products or the services produced to the downstream components. Other connections exist, referred to as *indirect links* in Benoit (2004), which, not being specifically constructed or planned, are very often originated by some kind of proximity. Sometimes they do not exist at all in normal conditions and they appear as a consequence of specific failures, or emerge after some changes in the context in which the infrastructures operate.

In the following, to stress their different nature, we will distinguish between *Resource propagation links*, related to quantities directly used by each macrocomponent to perform its activities, and *Failure propagation links*, used to describe paths along which failures might be propagated.

In both categories, we have to consider a multitude of different propagation mechanisms, each one characterised by its own properties and related to different topologies or metrics. This implies that, for each macrocomponent, we need to define several sets of different neighbours.

Then, in CISIA, interactions among macrocomponents are described in terms of sending/receiving messages exploiting the presence of a multitude of *oriented weighted incidence matrices*, each one representing neighbours associated with a given phenomena, *e.g.*, geographic proximity or electric propagation.

The use of different kinds of matrices, beyond the emphasis on different characteristics of each dependency, simplifies interdependencies' discovery. Indeed, considering the taxonomy introduced in Rinaldi *et al.* (2001), physical interdependencies are generally well known to infrastructure experts and could be read from functional schemes. On the other hand, geographical interdependencies are less understood by experts, but they can be discovered by comparing infrastructure maps.

Each entry in these incidence matrices is characterised by two coefficients: one represents attenuation from the source to the destination, and the other represents the time delay needed for its propagation (in the actual version, this is the only 'crisp' quantity).

Moreover, for what concerns failure propagation, our approach distinguishes between spreading and diffusion of a failure. More specifically, any failure is spread to all the members of the corresponding set of neighbours, but this represents only a necessary condition for the propagation of the fault, *i.e.*, its diffusion. Indeed, we also consider other specific characteristics of the macrocomponent reached by the fault (*e.g.*, actual state, internal level of failure, presence of adequate protection profiles) that can reduce the severity or even block the propagation of such failure.

The simulation is primarily composed of the entities and the adjacency matrices. Entities and matrices are collected in two main structures, Entity Pool (EP) and TSS, depicted in Figure 3.

Figure 3 The simulator is composed of two main structures: the *Entity Pool* devoted to handle entities' evolution and the *Transmission Sub System* (TSS), which manages communication among entities.

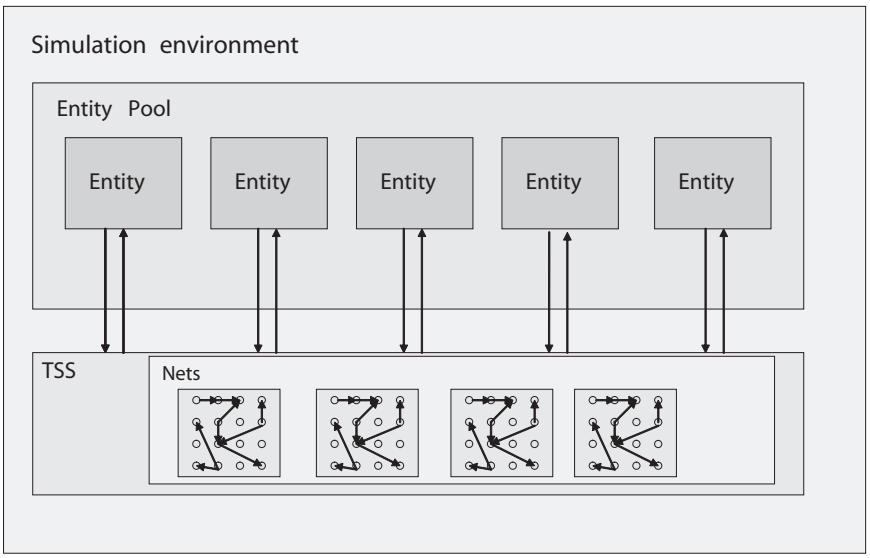

The TSS is devoted to managing the communication between the entities. Entities communicate via message exchange, where each message contains data about the type and the denormalised quantity of carried resource (or fault), the normalising factor, the unit of measurement and the sender port ID. The TSS collects the outgoing messages from all the entities and delivers each message to the neighbours of the sender entity, in accordance with the adjacencies described in the matrix associated with the type of the carried quantity. If a link between two adjacent entities is characterised by attenuation or delay factors, TSS provides to delay the delivery of the messages routed over that link and to suitably scale the carried quantities.

5 Case study

We tested CISIA by analysing the possible consequences of a failure inside one of the electric power plants that supply electricity to the area of Rome.

With respect to our goals, this area can be modelled in a very crude approximation, for many aspects not completely exact, considering the presence of five logically autonomous macrocomponents:

- *Power plant* – represents an electricity production plant located close to the urban area of Rome. This entity models a set of power plants geographically co-located in the same area (which share common resources, such as fuel storage and water cooling system).
- *380 kV substation* – this entity represents the point of contact from the local electric network and the national grid. Even if the area of Rome actually has several interconnection points, for our analysis we assumed the presence of a single interconnection point.
- *Energy transportation network* – composed of a single grid with the associated telecommunication and control network, able to dispatch electricity to urban areas.
- *Urban areas* – specifically we consider two close urban areas of different sizes. The *Big City* represents the town, while the *Small City* is the suburban area where the centre that supervises and controls the national power grid and telecontrols the different power plants is located. Obviously, the actual architecture used to supervise the national power grid is more complex and foresees the presence of different local centres that are coordinated by a national node (this latter is replicated in different geographic locations to guarantee high availability and dependability).

In a very simplified manner, we assume that, in order to generate electricity, the power plant needs fuel from storage, water from the basin and control information from small city. Moreover, we assume that electricity consumptions are so high that local generation represents a mandatory component. On the other side, the transportation network receives electricity from the power plant and the 380 kV substation supplying to urban areas. In our scenario, the *OL* of the urban areas is strictly related to the presence of electricity: when a blackout persists for a certain amount of time, there is a degradation in basic services, and this induces social disorder and insecurity phenomena (*i.e.*, a social fault in the urban area). The electric grid adopts a dispatched policy to prioritise supplying electricity to the small city, where the control centre is located.

In our simulation, we assume that at the tick $T = 20$, a fire abruptly turns off the fuel storage that supplies the power plant (see Figure 5(a)).

Even if the 380 kV substation continues to supply electricity (see Figure 6), there is insufficient power to satisfy the requirements of both the urban areas. Then, in accordance with its dispatching policy, the electrical network reduces the amount of power supplied to the big city through the low-priority line (see Figure 7(a)). This shortage, after some time, induces problems in the big city owing to the absence of electricity in important services. Therefore, the *OL* of the town is reduced (see Figure 8(a)), as evident from the simulation results. Because we have adopted a triangular FN representation of the different quantities, the time histories of these quantities are composed of three curves (Figure 6), where the middle one (*i.e.*, the solid line) represents the most 'believable' behaviour, while the other lines represent, respectively, an estimation of the best and worst case. Moreover, we are able to characterise the evolution of the uncertainties, which, for each variable, is proportional to the difference between its maximum (line-dotted line) and minimum (dotted line) values.

Figure 4 The case study scenario: a simplified view of the urban area around Rome and its electricity supply system

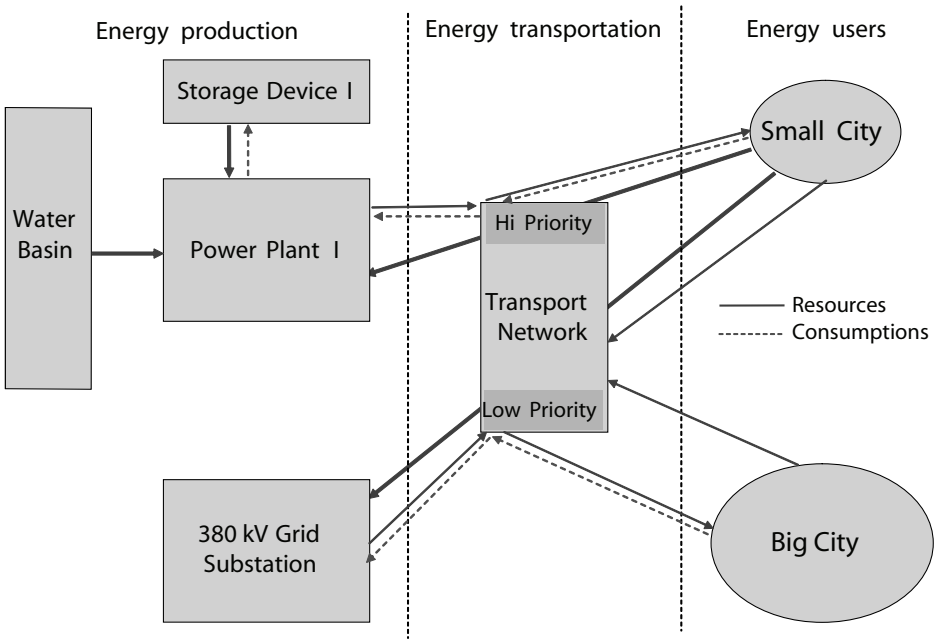

Figure 5 *OL* of the power plant (a) and 380 kV substation (b)

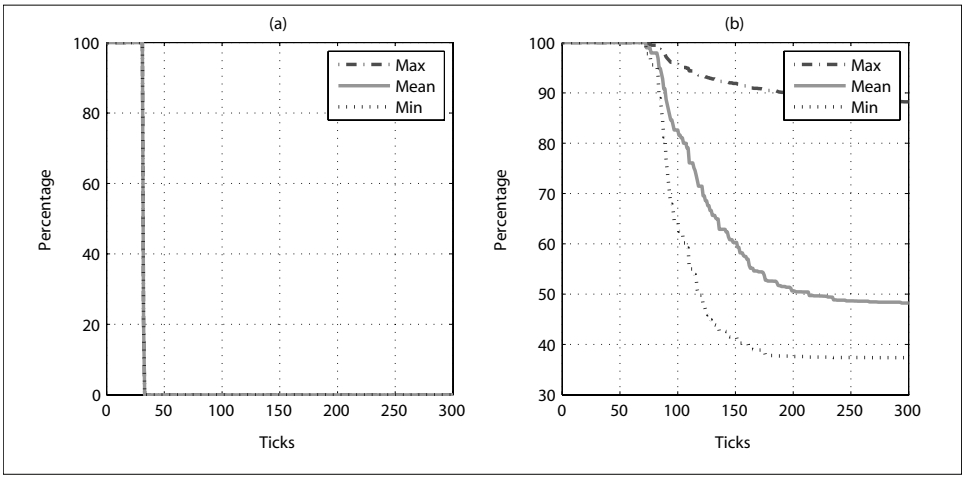

Note: All values are FNs.

Figure 6 Power supplied by the 380 kV substation

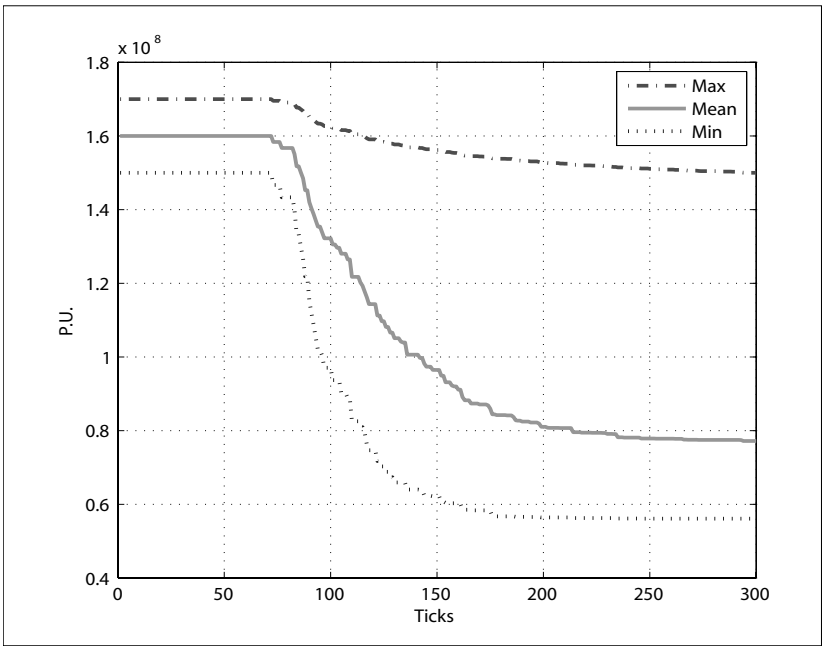

Note: All values are FNs.

Figure 7 Amount of electricity available for the big city (a) and the small city (b)

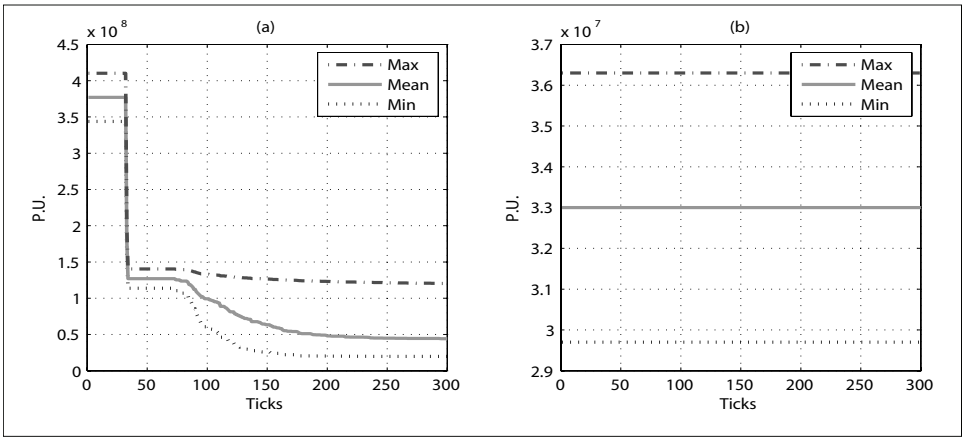

Note: All values are FNs.

Figure 8 *OL* of the big city (a) and of the small city (b)

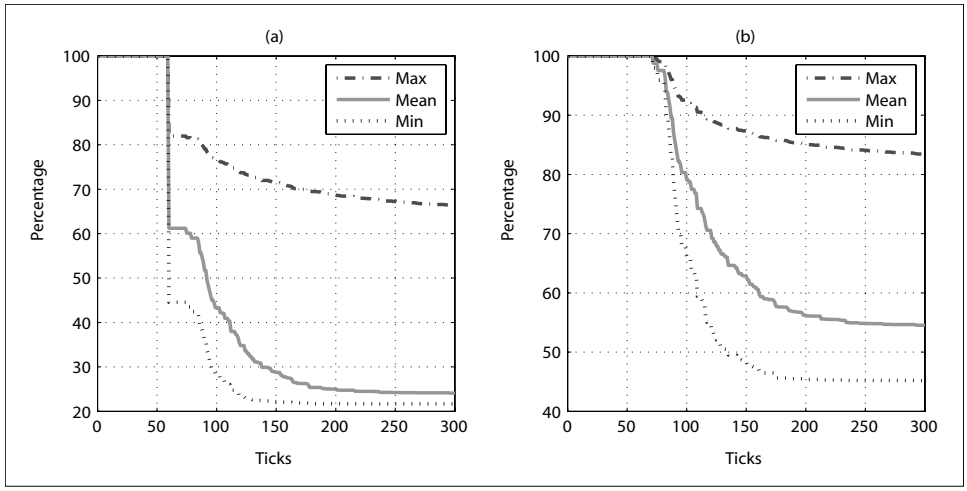

Note: All values are FNs.

Looking at Figure 8(b), we note that after some ticks, even if there is no reduction of the electric power supplied to the small city (see Figure 7(b)), we also record a degradation in the *OL* of this second urban area.

This behaviour is related to the propagation of the sociological failure generated inside the big city. Indeed, blackouts produce degradation in critical services (*e.g.*, stopping of railway transportation, traffic congestion, difficulties in communication, problems in healthcare and emergency services), and this induces social disorders, which in the model are supposed to spread also to the small city owing to their geographical proximity (see Figure 9(b)).

Figure 9 Sociological failure which affects the big city (a) and the small city (b)

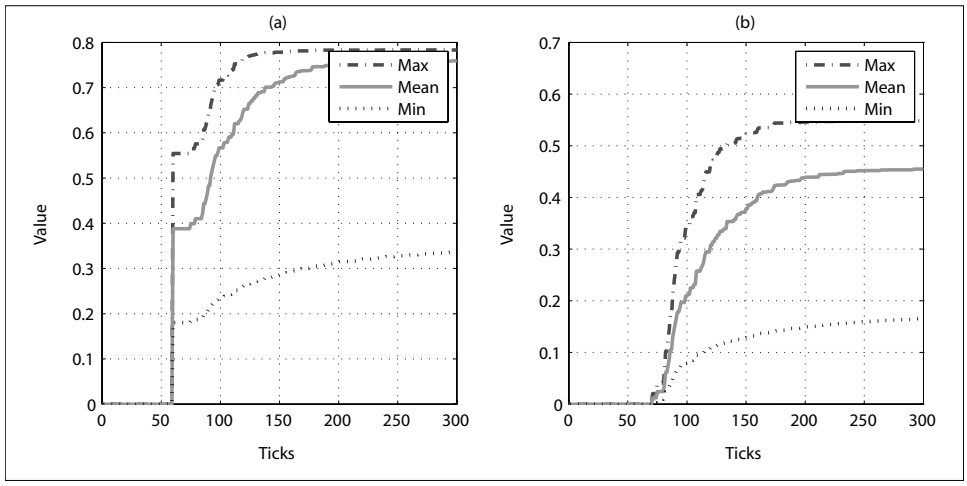

Note: All values are FNs.

The presence of these disorders has an impact also on the capability of the control centre to efficiently manage the national grid, and then on the 380 kV substation. This implies that the substation reduces the amount of electricity provided (see Figure 6). This exacerbates the blackout problems in the big city and, obviously, increases the level of social disorder. These disorders are then spread to the small city, where we register further reduction in its *OL*, and so on.

In the steady state, we note that, even if the small city is continuously supplied, its *OL* is reduced to about 40% of its nominal level. This emphasises that, in this very simple scenario, the dispatched policy adopted appears to be partially effective, and suggests the opportunity to analyse different strategies.

Notice that the reported results have been unrealistically augmented and the time compressed to emphasise the consequences of cascade failures.

6 Conclusion

The paper describes an approach to model complex interdependent infrastructures that is useful in capturing some of the most important phenomena for impact analysis. The approach, implemented in a fully operative simulator, named CISIA, is specifically designed to take into account different interdependency mechanisms and handle uncertainties and vague information.

These latter elements appear very relevant in critical infrastructure scenarios because information about them and their interdependencies are often scarce, ambiguous and subjective. In the paper, we stressed the importance of aligning modelling details to the 'quality' of the available data.

To partially overcome this drawback, we propose the adoption of an abstract representation of the different elements in terms of their capability to produce goods and services on the base of the availability of external resources, and taking into account the presence (and severity) of different types of failures. The overall behaviour of the system is then obtained considering routing and spreading among macrocomponents of the 'potential' amount of resources and the effective severity of failures. Then, interactions among macrocomponents are modelled via the exchange of a small set of common data, each of them with its specific concept of proximity.

In this way, it is possible to adopt a standard framework to model a very large class of macrocomponents that is easily scalable, in order to better fit the granularity of available information.

Moreover, to better handle the large uncertainty that characterises this class of systems, the paper suggests representing the different quantities by means of FNs.

CISIA, besides providing simulation tools, also supports the performance of topological analysis over the different scenarios. Obviously, a plain topological analysis over a heavy multilayered environment, like a multiple-infrastructure framework, may not have a fulfilling result. In CISIA, we are able to discover both the effects of a given fault over the entire infrastructure network, and how a given fault may be propagated through the different interdependence layers; which nodes represent the weakest points with respect to a certain class of faults/attacks or how the combination of concurrent events may contribute to make a critical scenario worse.

Moreover, the possibility of describing different adjacency matrices permits the analyst to describe in a simpler way a complex scenario where several elements are tied together by multiple relationships, a hard task to carry out with a monolithic approach.

Further evolutions of this work will consist in building a block collection for common elements of important infrastructures, implementing a user-friendly interface to make easier the coding of the macrocomponent's behaviours, and testing it on different scenarios.

References

Albert, R. and Barabasi, A. (2002) 'Statistical mechanics of complex networks', *Reviews of Modern Physics*, Vol. 74, pp.48–97.

Albert, R., Jeong, H. and Barabasi, A. (2000) 'Error and attack tolerance of complex networks', *Nature*, Vol. 406, pp.378–382.

Amin, M. (2002) 'Modelling and control of complex interactive networks', *IEEE Control Syst. Mag.*, pp.22–27.

Benoit, R. (2004) 'A method for the study of cascading effects within lifeline networks', *Int. Journal of Critical Infrastructures*, Vol. 1, pp.86–99.

Bologna, S. and Setola, R. (2005) 'The need to improve local self-awareness in CIP/CIIP', *Proc. First IEEE International Workshop on Critical Infrastructure Protection (IWCIP 2005)*, Darmstadt, Germany, pp.84–89.

Dubois, D. and Prade, H. (1998) *Possibility Theory: An Approach to Computerized Processing of Uncertainty*, New York: Plenum Publishing Corporation.

Dunn, M. and Wigert, I. (2006) *International CIIP Handbook*, AETH, The Swiss Federal Institute of Technology, Zurich, Vol. 1.

EU Commission Communication 576 (2005) *Green Paper on a European Programme for Critical Infrastructure Protection COM(2005)576*, Brussels.

Luiijf, A., Burger, H. and Klaver, H. (2003) 'A critical (information) infrastructure protection in The Netherlands', *Proc. Critical Infrastructure Protection (CIP) Workshop*, Frankfurt, Germany.

MacDonald, R. and Bologna, S. (2001) 'Advanced modelling and simulation methods and tools for critical infrastructure protection', Technical Report, ACIP Project Report.

Moteff, J., Copeland, C. and Fischer, J. (2002) 'Critical infrastructures: what makes an infrastructure critical?', *Report for Congress RL31556*, The Library of Congress.

Newman, D., Nkei, B., Carreras, B., Dobson, I., Lynch, V. and Gradney, P. (2005) 'Risk assessment in complex interacting infrastructure systems', *Proc. of 38th Hawaii Int. Conf. on System Science*, Big Island, Hawaii.

OCIPEP (2003) *Threats to Canada's Critical Infrastructure*, Government of Canada Office of Critical Infrastructure Protection and Emergency Preparedness, TA03-001.

PIC – Italian Government Working Group on Critical Information Infrastructure Protection (2004) *La Protezione Delle Infrastrutture Critiche Informatizzate – La Realtá Italiana* [in Italian].

Rinaldi, S., Peerenboom, J. and Kelly, T. (2001) 'Identifying understanding and analyzing critical infrastructure interdependencies', *IEEE Control System Magazine*, pp.11–25.

Rosenbush, S. (1998) 'Satellite's death puts millions out of touch', *USA Today*.

Ross, T. (2004) *Fuzzy Logic With Engineering Applications*, Chap. 12, J. Wiley and Sons Ltd.

Setola, R. and Ulivi, G. (2003) 'Modelling interdependent critical infrastructure', in N.E. Mastorakis (Ed.) *Recent Advances in Intelligent Systems and Signal Processing*, Electrical and Computer Eng. Series, WSEAS Press, pp.366–372.

UN 58th Generally Assembly (2003) *Creation of a Global Culture of Cybersecurity and the Protection of Critical Information Infrastructures*, No. UN-58/159.

US and Canada Power System Outage Task Force (2004) *Final Report on the August 14, 2003 Blackout in the United States and Canada: Causes and Recommendations*.

US General Accounting Office (2003) *Critical Infrastructure Protection: Challenges for Selected Agencies and Industry Sectors*, GAO-03-233.

US Government (2003) *The National Strategy to Secure Cyberspace*, The White House, Washington, USA.

US Government (2005) *The National Plan for Research and Development in Support of Critical Infrastructure Protection*, The Executive Office of the President and the Office of Science and Technology Policy, Washington, USA.

Watts, D. and Strogatz, S. (1998) 'Collective dynamics of "small-world" networks', *Nature*, Vol. 393, pp.440–424.