available at www.sciencedirect.com

ScienceDirect

www.compseconline.com/publications/prodinf.htm

Information Security Technical Report

Connectivity models of interdependency in mixed-type critical infrastructure networks

Nils Kalstad Svendsen[a], Stephen D. Wolthusen[a,b,*]

[a]Norwegian Information Security Laboratory, Gjøvik University College, P.O. Box 191, N-2802 Gjøvik, Norway
[b]Information Security Group, Department of Mathematics, Royal Holloway, University of London, Egham Hill, Egham TW20 0EX, UK

ARTICLE INFO

Published online 12 March 2007

ABSTRACT

Determining interdependencies and cascading failure modes in critical infrastructures is a complex problem that is exacerbated further by the diverging characteristics of the interconnected infrastructure types. Services in some types of infrastructure such as telecommunications or the electric grid are provided and consumed instantly. Others, notably oil and gas but also other infrastructures built on physical resources, however, exhibit buffering characteristics. In this paper we describe a model for the abstract representation of both types of infrastructure networks and their interdependencies. The model is then validated and demonstrated using characteristic topologies and interconnections.

1. Introduction

Critical infrastructures, including primarily the energy, financial services, health care, public services, and transportation sectors (Marsh, 1997; Brömmelhörster et al., 2004), are interconnected and interdependent on multiple levels. This leads to a number of questions which must be answered satisfactorily to protect the well-being of the population, functioning of government, and economic capabilities. Questions may include what cascading effects a regional failure of one critical infrastructure (such as the recent November 2006 failure of the electric power grid throughout much of continental Europe (E.ON Netz GmbH, 2006) and earlier failures in this infrastructure such as the August 2003 power outages in the northeastern U.S. and Canada (Hilt, 2004)) may have on other infrastructure components, or to elaborate how adding small and hence cost-effective amounts of redundancy can significantly enhance the overall robustness of this interconnected network of infrastructure services.

Based on a scalable multigraph-based model we have described earlier (Svendsen and Wolthusen, 2007), in this paper we are describing a model capable of not only representing the types of instantaneous interactions between infrastructure components as may arise in networks such as the electric power grid and in the telecommunications sector but can now also cope with so-called buffered resources, wherein reservoirs of resources may be retained and thereby allow the provision of services even if the required resources for their composition or manufacture have become (e.g. temporarily) unavailable.

While the buffered characteristic is exhibited by a number of fungible resources (including food supplies or fuels such as coal), we are concentrating on the special case of fluid or gaseous resources transported over pressurized pipeline networks since other types of fungible resources can obviously be trivially reduced to this case. The model, however, retains its overall simplicity and focuses on efficient computability over large-scale networks and is therefore not aiming for accurate representation of all relevant physical effects. While elaborate models, also incorporating physical characteristics and effects and with predictive capabilities exist for many of the individual critical infrastructure services such as for

* Corresponding author.
E-mail address: stephen.wolthusen@rhul.ac.uk (S.D. Wolthusen).
1363-4127/$ – see front matter
doi:10.1016/j.istr.2007.02.005

electrical power grids at national and transnational levels, it is clearly desirable to also investigate larger-scale interactions among multiple infrastructure sectors.

Moreover, from the perspective of critical infrastructure protection it is not of primary interest to retain maximum efficiency and effectiveness for each individual infrastructure element or sector. Instead, the primary focus is on the provision of adequate levels of service particularly if those services are being degraded, e.g. as a result of a large-scale natural disaster or coordinated acts of terrorism. Issues to be explored in a model such as the one presented in this paper therefore include questions of the prioritization of services (e.g. to maintain communication and power supplies for first responders and medical services) and the required interval in which some infrastructure services (e.g. gas pipelines and electric grid) must be restored or alternative services established before cascading effects (widespread power outages owing to grid partitioning and inability to provide fuel to a gas-fired power plant in the example above) become unacceptable. The models required for this type of exploratory research must provide both acceptable computational complexity and at the same time yield adequate levels of detail and accuracy in their unavoidable simplification. Even though sector-specific models can and do provide much greater levels of accuracy than simple connectivity-based mechanisms for determining dependencies both are valuable tools, it is only a combination of both approaches that can provide a quantitative and qualitative overview of infrastructure interdependencies and ways to address them and which can then be used to refine further investigations.

We therefore provide a refinement of the unbuffered model first described in Svendsen and Wolthusen (2007) by adding buffered resource models into the model, first reviewing the overall model for this purpose in Section 2. We subsequently describe a number of complex interdependencies, including transitive and cyclical interdependencies among heterogeneous infrastructure types along with examples of possible disruptions (or attacks) on such networks in Section 3. Analogously, Section 4 discusses infrastructure interdependencies arising when incorporating buffered resource types as exemplified in the case of natural gas pipelines. The results of the preceding model instances are then discussed in Section 5, while Section 6 briefly reviews related work before our conclusions and notes on ongoing research in Section 7.

2. Model overview

Interactions among infrastructure components and infrastructure users are modeled in the form of directed multigraphs, which can be further augmented by response functions defining interactions between components. In the model, the vertices $\mathcal{V} = \{v_1, \ldots, v_k\}$ are interpreted as producers and consumers of m different types of services. A single node can act both as a producer and a consumer at the same time. If a node is not able to generate a needed type, the node is said to be dependent on some other node delivering this service. Such a dependency is named dependency type d_j, and is chosen from the set $\mathcal{D} = \{d_1, \ldots, d_m\}$ (Fig. 1).

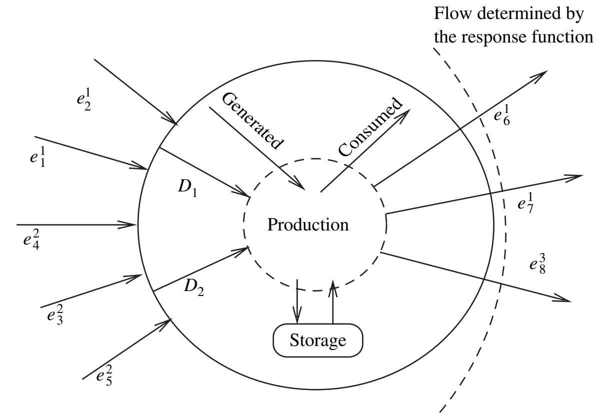

Fig. 1 – The parameters that define the functionality of a node, and its outputs.

Dependency types are classified into storable and non-storable or ephemeral types. It is assumed that all nodes v_a have a buffer of volume $\mathcal{V}_a^j$ (indicating a scalar resource; this may represent both physical and logical resources and, moreover, may be subject to further constraints such as integral values) for all dependency types d_j, and we define $N_{\mathrm{Max}}(v_a, d_j)$ to be the capacity of the storage in terms of amount of resource d_j. In the case that d_j is a non-storable dependency type, we have $\mathcal{V}_a^j = 0$ for all nodes v_a, and it follows that $N_{\mathrm{Max}}(v_a, d_j) = 0$. If d_j is storable but incompressible (as may be for most physical resources) we have that $N_{\mathrm{Max}}(v_a, d_j) = \rho\mathcal{V}_a$, where ρ is the density of the resource. Finally in the case of storable and compressible dependency types we find $N_{\mathrm{Max}}(v_a, d_j) = P_{\mathrm{Max}}(v_a, d_j)\mathcal{V}_a$, where $P_{\mathrm{Max}}(v_a, d_j)$ is the maximum pressure supported in the storage of resource d_j in the node v_a. Further refinements such as multiple storage stages (e.g. requiring staging of resources from long-term storage to operational status) and logistical aspects are not covered at the abstraction level of the model described here.

Pairwise dependencies between nodes are represented with directed edges, where the head node is dependent on the tail node. The edges of a given infrastructure are defined by a subset $\mathcal{E}$ of $\mathcal{E} = \left\{e_1^1, e_2^1, \ldots, e_{n_1}^1, e_1^2, \ldots, e_{n_m}^m\right\}$, where $n_1, \ldots, n_m$, respectively, are the numbers of dependencies of type $d_1, \ldots, d_m$ and e_i^j is the edge number i of dependency type j in the network. A given dependency between two nodes v_a and v_b, is uniquely determined by $e_i^j(v_a, v_b)$.

In addition to the type, two predicates $C_{\mathrm{Max}}(e_i^j(v_a, v_b)) \in \mathbb{N}_0$ and $C_{\mathrm{Min}}(e_i^j(v_a, v_b)) \in \mathbb{N}_0$ are defined for each edge. These values, respectively, represent the maximum capacity of the edge $e_i^j(v_a, v_b)$ and the lower threshold for flow through the edge. Hence, two $k \times m$ matrices C_{Max} and C_{Min} are sufficient to summarize this information.

Let $r_a^j(t)$ be the amount of a resource of dependency type j produced in node v_a at time t. We define $\mathcal{D}(t)$ to be a $k \times m$ matrix over $\mathbb{Z}$ describing the amount of resources of dependency type j available at the node v_a at time t. It follows that the initial state of $\mathcal{D}$ is given by

$$\mathcal{D}_{aj}(0) = r_a^j(0). \tag{1}$$

For every edge in $\mathcal{E}$ a response function $R_i^j(v_a, v_b)$:

$$\mathcal{D}_{aj} \times \mathcal{V}_a^j \times N_a^j \times N_{\mathrm{Max}}(v_a, j) \times C_{\mathrm{Max}} \times C_{\mathrm{Min}} \to \mathbb{N}_0 \tag{2}$$

that determines the i-th flow of type j between the nodes v_a and v_b is defined (see Sections 3 and 4 for some simple examples of response functions). The function $R_i^j(v_a, v_b)$ w.l.o.g. is defined as a linear function, and may contain some prioritizing scheme over i and v_b. By constraining the response function to a linear function and discrete values for both time steps and resources, linear programming approaches can be employed for optimization of the relevant parameters; interior point methods for this type of problem such as given in Karmarkar (1984) and Schrijver (2003) can achieve computational complexity on the order of $O(n^{3.5})$, making the analysis of large graphs feasible.

Given the responses at time t, the amount of resource j available in any node v_a at time $t+1$ is given by

$$\mathcal{D}_{aj}(t+1) = r_a^j(t) + N_a^j(t) + \sum_{i,s|e_i^j(v_s, v_a) \in \mathcal{E}} R_i^j(v_s, v_a, t). \tag{3}$$

A node v_a is said to be functional at time t if it receives or generates the resources needed to satisfy its internal needs, that is $\mathcal{D}_{aj}(t) > 0$ for all dependency types j which are such that $e_i^j(v_b, v_a) \in \mathcal{E}$, where $b \in \{1,\ldots,a-1,\ a+1,\ldots k\}$. If this is the case for only some of the dependency types the node is said to be partially functional. Finally, if no requirements of the node are satisfied, it is said to be dysfunctional.

As seen from Eq. (3) a single step model with one state memory has been chosen. This is a natural choice, as we are currently not concerned with long-term feedback, although the model naturally extends to longer-term state retention. The presented model can be used to represent any topology given a set of infrastructures and their interconnections. The model cannot achieve the level of accuracy found, e.g. in dedicated network simulators as described in Section 6; however, it has the advantage of being able to estimate the consequences of cascading failures through large-scale interconnected infrastructures.

The implemented model is primarily intended to assist in the manual and particularly algorithmic investigation of how high-level network effects (functionality of nodes) and interrelations (connectivity of nodes) in interconnected infrastructures react to different attack scenarios as well as criteria and mechanisms for enhancing the robustness of the resulting interdependency multigraphs. This provides a natural progression from the initial studies of large complex networks which concentrated on evaluating the robustness of attacks towards the infrastructure based on static failures (Cohen et al., 2000; Callaway et al., 2000), i.e. removing a certain percentage of the nodes in the network and estimating how the performance or connectivity of network is affected by the induced failure. In dependency networks, as in the case of electric power distribution networks and the telephony transport network used in subsequent (purely illustrative) examples, the breakdown or partial degradation of a node may cause cascading failures and have other time-dependent dynamic effects through the network detectable only through a dynamic approach to the networks, which the present model aims to integrate.

3. Interdependencies between infrastructures with no buffering

As a first example of the mechanisms provided by the model, the mutual influence of interconnected infrastructures depending on non-storable resources shall be considered. This analysis must be based on several abstractions and represents an approximation to actual network topologies for the purposes of this example. The selection of networks used in the following example was based on public availability of topology information for several network instances, the role of the network in the society, as well as a direct and obvious interdependency.

These criteria have led to the selection of the electric power distribution grid and the telephony transport layer, along with an additional key rationale. These two infrastructures are key components in modern industrial societies at both direct and indirect levels (the latter is not considered in this example); for instance, the BAS study (Hæsken et al., 1997), carried out by the Norwegian Defense Research Establishment in 1997, established core aspects of the criticality of the electric power grid and telecommunication networks in the Norwegian society. These networks are thus natural subjects of analysis and have been the subject of separate investigations (see Section 6). In addition the networks are interesting candidates for model verification because of a fundamental difference in how service deliveries flow through the networks.

3.1. *Non-storable resources*

In networks with edges representing a non-storable resource d_j, outbound edges are immediately impacted by inbound edges. Using the model described in Section 2, we have that $\mathcal{V}_a^j = N_{\mathrm{Max}}(v_a, d_j) = 0$. The response function is thus a function $R_i^j(v_a, v_b)$:

$$\mathcal{D}_{aj} \times \mathcal{V}_a^j \times C_{\mathrm{Max}} \times C_{\mathrm{Min}} \to \mathbb{N}_0, \tag{4}$$

and given the available resources and constraints in a node v_a at time t the available resources at time $t+1$ are given by

$$\mathcal{D}_{aj}(t+1) = r_a^j(t) + \sum_{i,s|e_i^j(v_s, v_a) \in \mathcal{E}} R_i^j(v_s, v_a, t). \tag{5}$$

The function $R_i^j(v_a, v_b)$ depends on the mechanism governing flows and conservation laws in the different networks.

3.2. *Electric grid network*

In the power distribution network the generated power originates from a small number of power plants or generators. A transportation network, which may well interconnect several power plants, delivers the power to a large number of transformers, which serve the low voltage distribution network, potentially through several intermediate sub-distribution networks. As a consequence the resulting graph is a directed network where multiple edges of different orientations between two nodes are rare occurrences.

Recently, extensive studies of the power distribution grid have been published in the open literature. One of the early analysis was of the Western States Power Grid in the U.S.

carried out by Watts and Strogatz in 1998. The degree distribution of the network was found to be exponential-like, but the clustering coefficients identified were too large for the network to be a classical random graph. The observed network consisted of approximately 3500 nodes, a number which might be too small for being conclusive regarding the categorization of the network (Dorogovtsev and Mendes, 2003). For the purposes of the present study, however, an exact representation of the power distribution is not necessary as we are primarily interested in topological characteristics. To this end, a network topology generator was implemented based on the following assumptions:

(i) The number of source nodes is small compared to the number of transport and sink nodes in the network.
(ii) Power generating nodes are not directly interconnected.
(iii) The network is constructed in order to cover a topological area as efficiently as possible.
(iv) Some redundant links are forced on the network in order to interconnect distribution networks and create redundancy.

Based on this, a tree-like model for the power distribution network provides an acceptable first-order approximation, although pure binary or k-trees are too regular to represent the topology. The basic Barabási-Albert (BA) model (Albert and Barabási, 2002) with some modifications provides a tree-like structure together with the level of irregularity found in real networks. The original BA model is initiated with a connected graph. In the power distribution network case, the source nodes are not directly interconnected. This is solved by simply providing the originating nodes with an initial degree $k_{\mathrm{Init}} \geq 1$ which does not represent any real edges, just the centrality of the node in the network.

Given that one node is added at each time step in the BA-model, as many disconnected trees as there are initial nodes in the network will be generated. In order to connect lower level nodes with each other a sparse random graph is placed on the top of the scale free networks. Given that this is a very sparse network it will not affect the statistical properties of the network, but has a major influence on the connectivity of the network and may also generate feedback loops.

The following simplified procedure is used to generate the power distribution network topology:

(i) *Growth*: At every time step a new node is added to the network. This node defines the head of an edge connecting it to an already existing node.
(ii) *Preferential attachment*: The tail of the edge is selected among the existing nodes with probability proportional to the degree of the node.
(iii) *Redundant connection*: After the final time step a sparse random graph is placed on the top of the network.

As the network grows large, the influence of the sparse random graph will be small, and the probability of a node having k edges will follow a power law with exponent $\gamma = 3$ (Dorogovtsev and Mendes, 2003).

Finally the response function for each edge is defined. In the case of quantitative analysis of service delivery this function should be an implementation of Kirchhoff's first rule, ensuring that all the flow into a node together with the flow generated by a node equals the output and the consumption of the node for the given resource type. In order to explore the presented model, however, such a detailed approach is not necessary, as the model instance under discussion focuses primarily on the functionality of the node. The principal issue in this case is that the electricity is consumed as it propagates through the networks and cannot, e.g. be stored in subgraph cycles. Thus the response function as described here only illustrates an abstract resource which is being consumed as it flows through the network. Introducing a threshold function

$$T(x, c) = \delta(x - c)x, \tag{6}$$

where

$$\delta(x) = \begin{cases} 0, & x < 0 \\ 1, & x \geq 0. \end{cases} \tag{7}$$

the implemented response function is of the form

$$R_i(v_a, v_b, t) = T\left(\frac{1}{2}\mathcal{D}_a(t), C_{\mathrm{Min}}(e_i(v_a, v_b))\right), \tag{8}$$

where $\mathcal{D}_a$ is the current available in the node a at time t. Eq. (8) indicates that two units of input current to the node are required to produce one unit of output current along an outgoing edge. As there is only one dependency in the network, the dependency type is not specified. Moreover, for the purposes of this example we also assume that there exists only one power dependency between two nodes and no prioritization scheme is defined over the outgoing edges.

A node in the power distribution network is defined to be functional if it has incoming current or generates current internally. The given response function can provide information on whether a node is functional or not, but does not provide any physical representation of the level of functionality of a given node in the network, which provides a sufficient level of details for the purposes of the present study.

3.3. *Telephony transport layer*

Traditionally the telecommunication layer has been a hierarchical network, as described, e.g. in Freeman (1999). Although there has been a decided trend away from this owing to progress in transmission and switching technology particularly since the early 1990s, we chose to use this model since it is representative of much of the currently deployed telecommunication infrastructure. The telecommunication transport layer could be described as an onion structure with a very low diameter. The signal always starts from the outer layer, then depending on the range of the connection, it goes through the core of the network before returning to a local switch in the outer layer of the network. As the flow through these edges is bidirectional, all connected nodes will be connected by an edge in each direction.

Compared to the electric grid, the telephony transport layer has received much less attention in the scientific community modeling critical infrastructures. Whereas the Internet and autonomous system networks have been modeled extensively (Pastor-Satorras and Vespignani, 2004), modeling the telephony transport layers' functionality and design has been left to operators for the most part, reducing the availability of open literature. As mentioned in the introduction to this section, we assume a telephony transport layer having a traditional hierarchical network structure. This is a network which is optimized locally for full connectivity, and globally to reduce the number of switches included in an average connection circuit.[1] In order to be functional a switch needs to be connected to other switches and to power supply, which is the focus of the following analysis.

The network is modeled as a number of disconnected trees which are connected to a fully connected transportation network through their root nodes. The simplified response function of the telephony network depends on whether the node has power as input or not. If there is no power available, circuit switching cannot take place and no communication is possible. In the example model discussed here, the response function for edges in the telephony transport layer is thus a threshold function given by

$$R_i(\upsilon_a, \upsilon_b, t) = \delta(\mathcal{D}_a(t) - C_{\mathrm{Min}}(e_i(\upsilon_a, \upsilon_b))), \tag{9}$$

where $\mathcal{D}_a$ is the current available in the node a at time t and δ is as defined in Eq. (7). It follows from the definition of Eq. (2) that a directed edge between the nodes υ_a and υ_b, is defined if there is power available in node υ_a. Again, no redundant links are defined between two nodes and no prioritization scheme is defined over the edges.

As mentioned earlier, each connection in the telephony transport layer is bidirectional, as one way voice communications are typically of limited interest. The functionality of a node thus depends on if the node itself and the node it is connected to has an effective power supply. Only in this case can the switch deliver the two way service it is meant to.

3.4. *Attack scenario*

The studies of complex networks frequently conclude that many man-made and natural networks possess scale free characteristics, and thus are exposed to the well-known Achilles heel of robustness against random breakdown while at the same time being vulnerable to targeted attacks (Albert et al., 2000). There is obviously a plethora of scenarios that may cause the failure of a set of nodes in some infrastructure. Causes may be intentionally or unintentionally created by humans or they may be the consequence of some changes in the environment of the network, e.g. flooding and temperature changes, or be induced by technical errors leading to failures. Analyzing the presented network models we have the following three attack scenarios in mind:

(i) Single node removal: This can be the consequence of a targeted terrorist attack or single technical failure.
(ii) Removal of small connected component: Representing non-localized failures such as flooding or other natural disasters.
(iii) Removal of disconnected components: This could, e.g. be the result of a coordinated terrorist attack.

In order to visually illustrate the properties of the presented model, artificially small topologies were generated. A power distribution topology based on two power sources and 28 power distribution nodes were connected to a telephony transport network with three core switches and a total of 21 switches. The switches were connected to randomly selected lower level power distribution nodes (meaning that no power generating nodes were connected directly to the telephony transport layer). None of the nodes of the telephony transport layer were assumed to have an independent power supply. Moreover, some of the power distribution nodes were connected to the telephony transport layer.

For all the presented scenarios two attacks are considered. Random node removal follows the procedure

(i) Remove a node from the network.
(ii) Run the response function until the number of functional nodes in the network stabilizes.
(iii) Count the number of functional nodes in the network.
(iv) Reinsert the node.

The procedure is repeated for all nodes in the network. Pairwise removal of nodes follows the same procedures, but in this case two nodes are removed from the graph at the time. Results of the attacks are presented as fraction of functional remaining nodes after the removal of one or two nodes from the network. The visualization of the results can be done with histograms as in Fig. 2, where we see the fraction of functional power nodes in the case of one (Fig. 2(a)) and two (Fig. 2(b)) node removal.

The presented results are deduced from one topology generated as described in Section 5. A single topology is not sufficient to draw any general conclusions on the properties of the proposed topologies, but illustrates the ability and flexibility of the presented model.

3.5. *Multi-domain one-way dependencies*

The dependency between the power distribution network and the telephony transport layer is first assumed to be one way as a simplification. This implies that the modeled power distribution network can be fully functional even if no switches in the telephony transport layer are functional. Conversely the flow along an edge in the telecommunication network will halt if either the head node or tail node loses its power supply. The connection of the telephony transport layer to the power grid is randomized in the present model (i.e. does not take into account geospatial proximity or other factors that would otherwise result in functional clustering). However, for the purposes of the present analysis, this is considered adequate.

The nodes of the telephony transportation layer now have two inputs, current and information, and give output in form of information. At every time step the response functions for power distribution and telephony transportation edges can be computed over their respective domains given the network state in the previous time step. From this the functionality of the telephony transport layer follows directly.

[1] This is primarily a vestige of analog switching system limitations required to consider attenuation and noise levels but is still underlying the existing network topology.

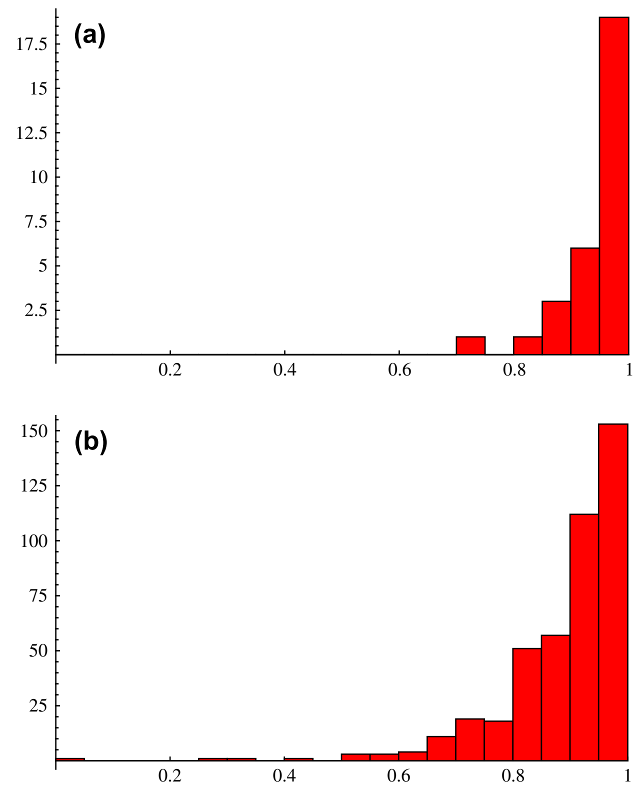

Fig. 2 – The consequences of one and two node removals from a scale free network with added redundancy. (a) Remaining fraction of functional power nodes after random removal of one power node (51 runs of the algorithm). (b) Remaining fraction of functional power nodes after random removal of two power nodes (1275 runs of the algorithm).

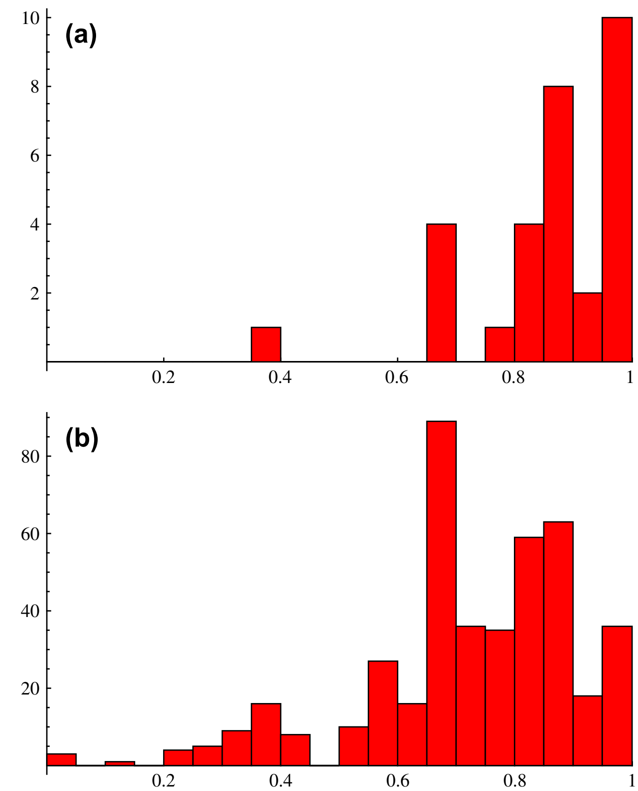

Fig. 3 – The consequences of one and two node removals in the power distribution network on telephony transport layer. (a) Remaining fraction of functional telecommunication nodes after random removal of one power node (51 runs of the algorithm). (b) Remaining fraction of functional telecommunication nodes after random removal of two power nodes (1275 runs of the algorithm).

Fig. 3 shows the fraction of fully functional telecommunication nodes as one or two nodes are being removed from the power distribution network. The histograms clearly illustrate the error propagation from the power distribution network to the telephony transport layer. In particular, we note the peaks observed as one of the two power generating nodes are being removed in Fig. 3(a), and the obvious observation that there is no functional switch when both power generating nodes are removed in Fig. 3(b).

3.6. *Multi-domain two way dependencies*

The dependency between the power distribution network and the telephony transport layer is now assumed to be bidirectional. Flow along an edge in the telecommunication network will halt if either the head or tail node loses its power supply. In the power distribution network, some of the nodes are dependent on the functionality of the telephony transport layer. This can, e.g. be remote switches. The interconnection of the telephony transport layer and the power grid is randomized in the present model (i.e. does not take into account geospatial proximity or other factors that would otherwise result in functional clustering). However, for the purposes of the present analysis, this is considered adequate.

The nodes of the telephony transportation layer, and some of the nodes in the power distribution network now have two inputs, current and information, and give output, respectively, in the form of information and current. At every time step the response functions for power distribution and telephony transportation edges can be computed over their respective domains given the network state in the previous time step. From this the functionality of the networks follows directly.

The cascading effects of two way cyclic dependencies between infrastructures are illustrated in Fig. 4. Fig. 4(a) illustrates the consequence of removal of one particular node from the electricity grid when there is no cyclic dependency. The level of functionality of the electric grid and telephony transport layer drops, respectively, to 0.8 and 0.6. Fig. 4(a) shows the consequences of removing the same node, but now with a couple of two way cyclic dependencies introduced. In this case we see a dramatic drop of functionality to 0.4 and 0.05, respectively, for the electric grid and the telephony transport layer.

Fig. 5 shows the effect of the removal of any pair of power distribution nodes on the considered networks. A comparison with Figs. 2(b) and 3(b) clearly show the increased vulnerability of the system introduced by the two way cyclic dependencies.

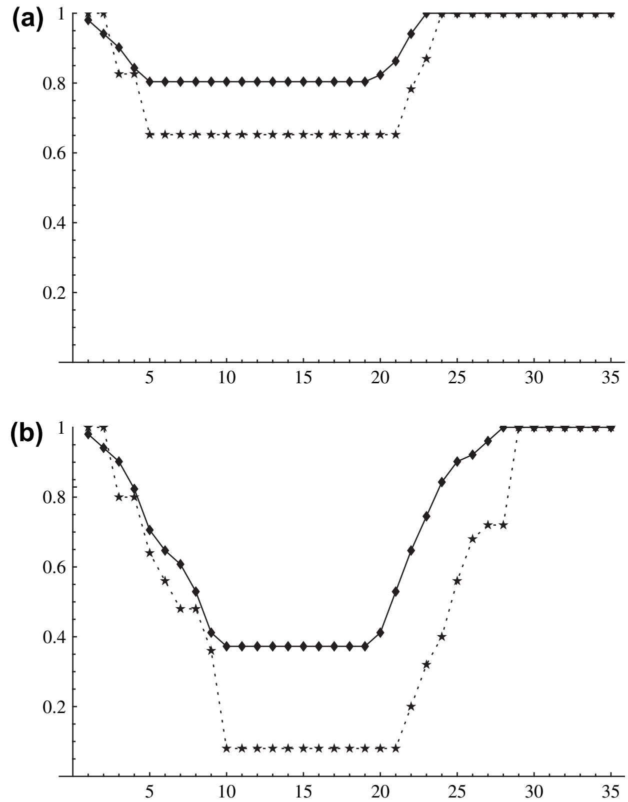

Fig. 4 – The level of functionality of the electric grid network (continuous line) and telephony transport layer (dashed line) as a particular node (not source node) is removed from the electric grid network. The node is removed in iteration 1 and reinserted in iteration 20. (a) No cyclic dependency between the electric grid and the telephony transport layer. (b) Two cyclic dependencies between the electric grid and the telephony transport layer.

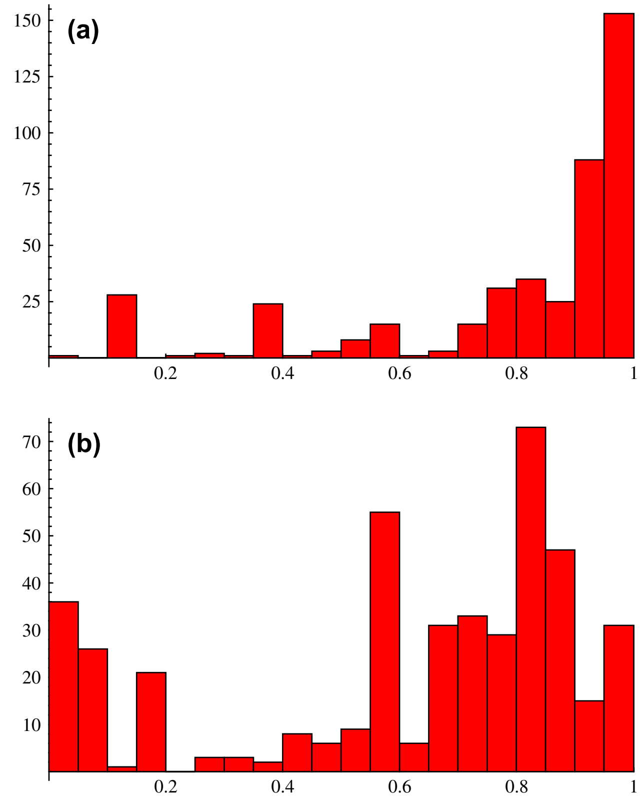

Fig. 5 – The consequences of two node removals in the power distribution network on the power distribution network and telephony transport layer with two multi-domain two way dependencies. All unordered pairs of power nodes has been removed over 1275 runs of the algorithm. (a) Remaining fraction of functional power nodes. (b) Remaining fraction of functional telecommunication nodes.

Multi-domain two way dependencies can create several challenges for infrastructure managers. A particular problem is the initial startup of the networks and may for example not be done without the use of additional measures such as backup power and other ancillary requirements (which, however, is part of the core contingency cases for such infrastructure components that are validated regularly). To visualize this two of the power distribution nodes of the networks of Section 3.5 are made dependent on two of the telecommunication switches.

4. Buffered infrastructure interdependencies

In numerous infrastructures the service delivery resource can be stored and accumulated in the nodes over time. This section proposes two extensions of our graph model, allowing simulations involving the depletion of both compressible and incompressible buffered resources. The section elaborates on a simplified model of a gas pipeline infrastructure as example of a buffered resource with the additional complexity of compressibility.

In networks with edges representing service delivery of a storable resource d_j, outbound edges are not immediately impacted by inbound edges. The amount of resource buffered in the node may increase in periods with an excess of input and may well decrease during periods while the amount of incoming resources are low or cut off completely. This ensures the operation of the node over a certain time in case of no input but may also well be the source of fluctuations in the system. Assuming that a node represents some facility capable of receiving, storing, and distributing a resource, the response function and the amount of resource d_j available in any node are given by Eqs. (3) and (2).

4.1. *Gas pipelines*

Pipelines are used for efficient and effective transportation of fluids over long distances from producing regions to consuming regions. Fluids frequently transported in pipelines are crude oil and natural gas, however, there are a variety of oil and gas products that can be transported in this way. There is nothing preventing liquid and gas phases from existing in the same pipeline, so called multiphase pipelines. Overviews of pipeline components and operating regimes can be found in several places, our description is based on the one given in Aalto (2005).

Natural gas pipeline systems transport natural gas from sources to users (sinks) through a system of interconnected pipeline segments. The difference in pressure at different points of the pipeline is the force driving the gas through the pipes. This difference in pressure is introduced by compressors at regular intervals along the pipeline. Besides the compressors, a pipeline also consists of metering stations, valves (discrete or continuous) and gas storages. All these components depend on power supply and in some cases also on telecommunication infrastructure enabling distant monitoring and control. In many cases electric power is provided from turbines driven with fuel from the pipeline, however, external power supply is increasingly common (Aalto, 2005).

A gas pipeline can be divided into three main parts: the gathering system, the transportation system, and finally the distribution system. The gathering system consists of low pressure, low diameter pipelines that transport raw natural gas from the wellhead to the processing plant. Transmission pipelines move gas in large quantities over long distances with few or no major supplies or off-takes between the end points of the pipeline. Distribution systems have a large number of off-takes and may be significantly branched. Evidently there are sliding boundaries between the different systems. In order to introduce or remove gas from a transmission pipeline special infrastructure such as pressure reduction stations and blending stations are required. These are omitted from our model.

Section 6 gives an overview of modeling efforts related to buffered resources, and in particular gas pipeline modeling. As our model is to be interconnected to other complex infrastructures we choose an approximation to the network model. The model does not aim to describe realistic behavior of fluids in pipelines, but to illustrate an approximation of the buffered behaviors in the pipeline and the interdependencies the mechanisms of the pipeline has with other infrastructures, e.g. power distribution and telecommunication networks.

We start by defining a network of interconnected pipes, the pipeline. The pipeline is divided into pipeline segments. As shown in Fig. 6, a pipeline segment is the piece of pipeline between two branchings. A branching can be a supply or an off-take, and is represented with an edge in our model. For simplicity, our sample networks are constructed without loss of generality such that compressors, metering stations, and valves only are present in branchings. Additional gas storages are also represented as pipe segments. We choose to model the gathering system as a number of in-branchings (Bang-Jensen and Gutin, 2006). The root nodes represent the network's transportation system. Meters, valves and compressor stations can be inserted as several nodes in line with the root node. The distribution network is, from the root nodes, viewed as out-branchings (Bang-Jensen and Gutin, 2006). In reality these spanning oriented trees are interconnected for redundancy purposes. The simplified case of one root node in the transportation network is shown in Fig. 6.

Every pipe segment P_a is represented as a node υ_a. In accordance with the general model presented in Section 2, each node has a buffer of volume $\mathcal{V}_a$, and a pressure limitation $P_{\mathrm{Max}}(\upsilon,a)$. The relation between pressure and volume for gases is given by the ideal gas law $PV = nRT$, where n is the number of moles, R is the universal gas constant, and T is the temperature (Zumdal, 1998). As an approximation we assume that the temperature remains constant in the pipeline, and normalizing the relations such that $RT = 1$, we get the relationship

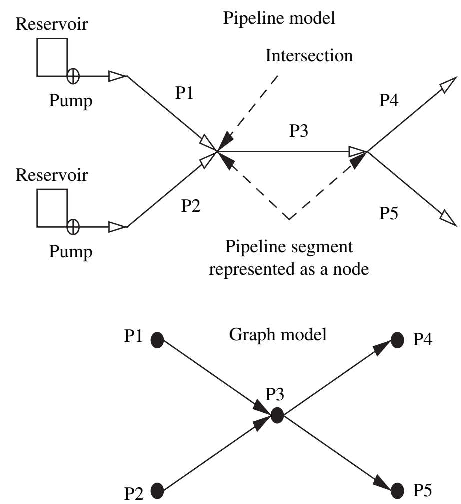

Fig. 6 – The principal components of a gas pipeline, and how they are transformed to a graph model. In the graph model P1 and P2 represent the gathering system, P3 is the transportation node, and P4 and P5 are distribution nodes.

$$PV = N, \tag{10}$$

where N is the amount of some arbitrary unit of gas.

Methods for calculating flows and losses in complicated pipe networks have much in common with methods for analyzing electrical networks, and the following rules are the basis for any calculation procedure (Gerhart et al., 1992):

– The net flow into any junction must equal the net flow out of the junction.
– The sum of head (or pressure) increases and losses around any loop must be zero.
– All losses must satisfy the pipe friction equations or the local loss equations. All pumps must operate at a point on their pump curve.

Our model follows these rules to satisfy the conservation laws.

The difference in pressure between different segments is the force driving gas from one node to another, and is therefore the input to our response function in the model. The velocity of the flow, or response, depends on several parameters, among others the length of the pipe segment, the diameter of the pipe, and the friction between the pipe and the gas. This means that the velocity of the gas will reach an upper bound, and may enter a regime of turbulence which again can cause a drop in velocity. Given the pressure P_a and P_b of two connected nodes υ_a and υ_b, we name the pressure difference ΔP_{ab}. The response function $R_i(\Delta P_{ab}, t)$ shall define the flow, in terms of amount of particles from v_a to v_b at time t. This is obviously not a linear relationship. Roughly it can be described as low for small ΔP_{ab}, then growing proportionally with ΔP_{ab} until the growth decays and the velocity tends to some threshold. This behaviour is approximately expressed by the logistic equation (Edwards and Penny, 1995). Bounded growth as a function of a variable x is of the form

$$\frac{A}{1+e^{(B-Cx)}},$$

where A is the upper bound, B defines the translation of the curve along the x axis, and C the maximal derivative of the curve. To represent that the ΔP_{ab}, e.g. due to friction, must be larger than some threshold to flow through a pipe a threshold is introduced. We can now define the response function as

$$R_i(\Delta P_{ab},t)=\begin{cases}0 & \text{if } \Delta P_{ab}(t)\leq\frac{C_{\text{Max}}}{10},\\ \frac{C_{\text{Max}}}{1+e^{(2.5-10\Delta P_{ab}(t))/C_{\text{Max}}}} & \text{if } \Delta P_{ab}(t)\geq\frac{C_{\text{Max}}}{10},\end{cases} \tag{11}$$

where $\Delta P_{ab}(t)$ is the pressure difference at time t, and C_{Max} is the capacity of the pipe. The values of B and C are chosen arbitrarily in order to get a suitable growth shape on the curve. An instance of $R_i(\Delta P_{ab}, t)$ with $C_{\text{Max}}=10$ is shown in Fig. 7.

It must here be noted that in the case of gas flow the order of the considered edges influences the result. Gas is consumed in the intersection, thus the potential pressure difference diminished. To approximate continuous behavior the sequence of the edges is altered from time step to time step. Therefore, on average, the behavior of the intersection is correct. By ordering the edges in a list a static prioritizing scheme can be set, if gas delivery to certain parts of the network is considered critical.

Based on the response function defined in Eq. (11) the amount of gas available in a pipe segment is given by

$$\mathcal{D}_a(t+1)=r_a(t)+N_a(t)+\sum_{i,s|e_i(v_s,v_a)\in\mathcal{E}} R_i(v_s,v_a,t). \tag{12}$$

By altering the computation of R_i and $\mathcal{D}_a$ we can observe the gas flow through the network over time. One must also be aware of the pressure limitation of the pipeline components. Infinite pressure cannot be tolerated, thus in the case that $r_a(t)$ gives a positive contribution it must be possible to diminish or eliminate it. The same also holds true for incoming pressure that valves can be closed. This is implemented in our model.

The goal of a gas pipe is to deliver gas to customers. In our model the customers are located as leaf nodes in the distribution network, or as off-takes from intermediate nodes in the distribution network. We choose to measure the functionality of the network in terms of the fraction of the leaf nodes receiving sufficient amount of gas to cover their needs. This will permit us to compare the performance of different network configurations. Additional constraints can be added, such as prioritized gas flow to specific nodes in order to maintain functionality of the particular node.

More interesting than to consider attack scenarios to the constructed network, after all the strength and weaknesses of tree like networks are well known, is a visualization of the effect of the buffering ability of the nodes in the network. This can be seen in Fig. 8. First there is an initialization phase before the gas is spread in the entire network. A phase of steady state is kept for a relatively long time for the reservoirs to be filled. At iteration 150 all sources are turned of. We see that the network remains fully functional for another 50 iterations, before a relatively dramatic decay of functionality.

5. Analysis

This section explores the presented models' ability to capture the consequences of network dependencies. The sample networks from the previous sections are considered, but now with added dependencies to interconnect all three networks. The argumentation behind the constructed scenario is that some of the central nodes in the gas distribution network are dependent on the telecommunication infrastructure to transmit metering information. If no metering information is available the pipeline has to be shut down due to the risks related to overload. The gas pipeline distributes power to a gas power plant facility which again is a supply for the power distribution network.

Critical network connections from the studied sample networks are shown in Fig. 9. The first part of the setup is

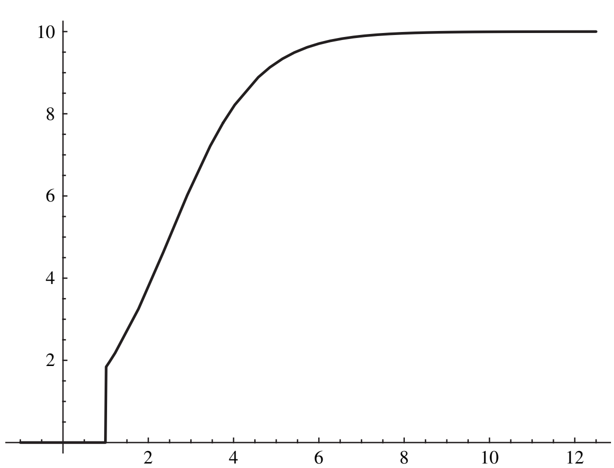

Fig. 7 – The transport between pipe segments as a function of the pressure difference between the segments with $C_{\text{Max}}=10$.

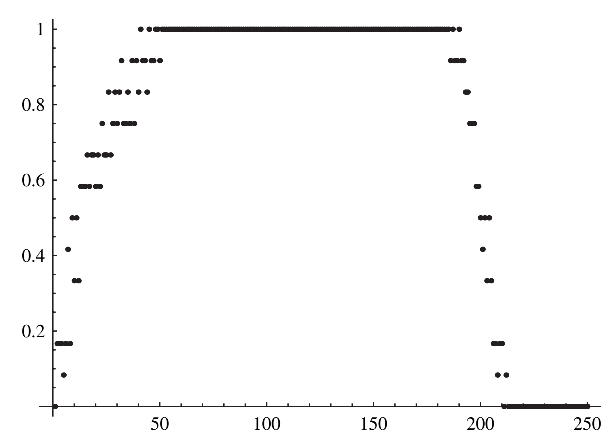

Fig. 8 – The functionality of a sample gas transportation network consisting of 23 nodes. At iteration 150 all sources are set to zero.

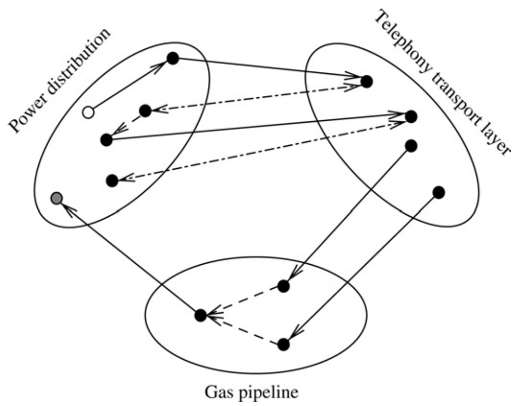

Fig. 9 – The interdependencies between the three networks. Continuous lines indicate directed dependencies, dashed lines indicates indirect dependencies, and mixed lines indicates bidirectional dependency of the telephony transport layer. The gray filled node is one of two power generating nodes in the power distribution network, and the node with no filling is the node were the error occurs in the following analysis.

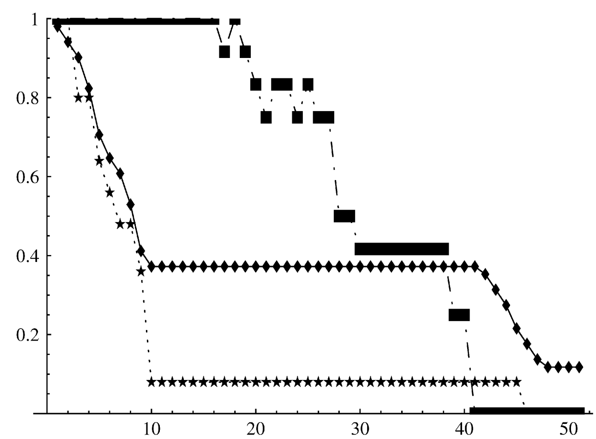

Fig. 10 – The level of functionality of the different network as failure propagates through the power distribution network (continuous line), telephony transport layer (dashed line) and gas distribution network (mixed dashes) as a result of a failure in one power node. All three networks are fully functional at the first observation.

identical to the one used to generate Fig. 5. Two power distribution nodes depend on two seemingly independent nodes in the telephony transport layer. The critical factor here is that the power supply to one of the telephony nodes indirectly depends on the functionality of the other telephony node. The removal of the power supply to the node with no filling starts the cascading failure observed in Section 3.6. In this analysis the two transportation nodes in the gas pipeline are dependent on two nodes in the telephony transport layer that seems to be independent of the dependencies described above, but which actually are taken out in the cascading failure. This leads to the supply to the majority of gas distribution nodes being shut off. Among these is the gas power plant in the power distribution network (the gray filled node).

Earlier sections introduced the level of functionality of the different networks as a metric to measure the effects of attacks or failures in a network. Fig. 9 shows how the level of functionality in the different networks evolves after the removal of the no fill node. For the power and telecommunication network we recognize the steep descent from Fig. 5. Most of the telecommunication network is taken out, while the functionality of the power distribution network remains at around 40%.

Due to the amount of gas buffered up in the gas pipeline there seems to be no immediate effect on this system. After about 20 iterations the first reservoirs are being depleted and a decay of functionality is initiated. The observed oscillations in Fig. 9 are due to the small size of our sample network, and to a rough resource prioritization scheme. We further observe a stabilization in the functionality of the gas consuming nodes around 40%. This, however, is just a temporary state since after about 40 iterations all the gas reserves are empty and the functionality of the gas consuming nodes goes to zero (Fig. 10).

The final consequence of this is the deprecation of the gas stock of the power plant represented by the gray filled node in Fig. 9. As a consequence only the nodes with supply from the second power which were not impacted by the first cascading failure remains functional (about 10%), whereas the functionality in the telecommunication network drops to zero.

Obviously these sample networks are of small size and constructed such that these effects could be observed. The result of the analysis of our model is that it is suitable to show the consequence of the removal of one or several nodes from some interconnected network. This provides incentives for research on algorithms providing an ability to detect critical interdependencies. Moreover, it also encourages future research on simulation tools capable of running what-if scenarios based on interconnections of critical infrastructure as may be required for decision support systems.

6. Related work

Related work tends to fall into two broad categories, namely abstract meta-level and therefore largely qualitative models for infrastructures and simulation systems at similar levels of detail on one hand and sector-specific models using detailed logical and physical constraints and models inherent of a given sector.

Examples of the former category include agent-based approaches such as the micro-simulation by Barton and Stamber (2000), the high-level simulation by North (2000), and the subsequent attempt to translate mechanisms for the modeling of complex adaptive systems by Thomas et al. (2003). Even more qualitative approaches also include system dynamics models such as that of Pasqualini and Witkowski (2005). Control systems approaches, e.g. the model proposed by Sullivan et al. (1999) can provide significant levels of detail but are constrained in their size and accuracy; the latter issue is at least in part addressed by the inclusion of hybrid control mechanisms as proposed by James and Mabry (2004).

Further qualitative efforts also include the results of the European Project ACIP (Schmitz, 2003) and related research on the Critical Infrastructure Modeling and Assessment Program (CIMAP) by Rahman et al. Similar approaches have also been described by Amin (2000) and Rinaldi (2004); these models vary considerably in their level of detail and range from simple binary dependency analyses to networks of models in which sub-aspects may be modeled by continuous physical submodels.

For the more constrained case of individual infrastructures such as electrical and pipeline grid environments, however, rich modeling and simulation environments already exist including the PSI control system and proprietary mechanisms employed by grid operators. Interconnections and interdependencies can only be modeled to a limited extent in such environments. Several properties are immediately derivable from interconnection characteristics alone as shown for power grid and Internet connectivity by Faloutsos et al. (1999), Broder et al. (2000), Yook et al. (2002), and Chen et al. (2002).

Frequently, the underlying structure of the networks can be identified as being wholly or partially scale-free; this was shown for network growth patterns by Dorogovtsev (Dorogovtsev and Mendes, 2001); further investigations of network properties by Casselman (2004) and Goh yielded insights into structural and spectral properties of scale-free networks (Goh et al., 2001) (for a survey of recent work on scale-free graphs see Newman, 2003). These results have severe repercussions for the assessment of vulnerabilities of interconnected and interdependent networks of critical infrastructure components to random failure as shown by Callaway et al. (2000) and Cohen et al. (2000), who also investigated the susceptibility of such graphs to targeted attacks (Cohen et al., 2001).

The more specific aspect of efficient modeling of buffered resources suitable for simulation and even real-time decision support as envisioned for the results of our paper, particularly in the case of pipelines for compressible resources as described in Section 4, has long been a subject of interest as made evident by the early modeling efforts by Kralik et al. (1984). Efficient pipeline models have been proposed among others by Zhu et al. (2001), while, e.g. Skvortsov and Sarychev (2002) provide models for pipeline sections. Other models tend, however, to remain proprietary and are not fully disclosed in the literature. The dissertation by Aalto (2005) does, however, provide a survey of the modeling and optimization issues for gas pipelines.

7. Conclusion

Based on a flexible framework for modeling infrastructures and their interdependencies we have first reported in Svendsen and Wolthusen (2007), we have described a graph-theoretical model augmented with a set of response functions that can model both unbuffered and particularly buffered resources along with their production and consumption in a network of infrastructure components. The model described allows the consideration of multiple concurrent types of interdependencies such as may arise in the provision of further infrastructure services (e.g. a hospital requiring electrical power, gas for heating, water, and telecommunications) along with simple prioritization mechanisms as may be necessary in case of some elements of the infrastructure network becoming unavailable or owing to a partitioning of the interdependency graph.

Based on this model we have demonstrated several types of multi-dependency structures for both linear and particularly cyclical dependencies among multiple infrastructure types of both unbuffered and buffered types. Moreover, we have also demonstrated simple attack scenarios over small fictional interconnected (and interdependent) infrastructure networks. The instantiations of the model presented in this paper are primarily intended for illustrative purposes; however, the topologies used for the respective infrastructure types and their interconnections are based on actual networks of the respective types.

Our ongoing research focuses on the identification of graph-theoretical and combinatorial optimization techniques (particularly as applicable to large-scale graphs) for both the identification of critical interdependencies and efficient mechanisms for increasing the robustness of such interdependent graphs. Future work includes further extensions of the model in which the response function can accommodate multiple resources being provided by each individual vertex in both discrete and continuous variables, resulting in a web of interdependencies. One such extension is the introduction of a more fine-grained time model capable of capturing certain effects such as ringing within dependency cycles and physical sub-models which the present coarse model cannot capture adequately.

REFERENCES

Aalto H. Real-time receding horizon optimisation of gas pipeline networks. PhD thesis, Department of Automation and Systems Technology, Helsinki University of Technology, Espoo, Finland; May 2005.

Albert R, Barabási A-L. Statistical mechanics of complex networks. Reviews of Modern Physics 2002;74:47–97.

Albert R, Jeong H, Barabási A-L. Error and attack tolerance of complex networks. Nature 2000;406(6794):378–82.

Amin M. Toward self-healing infrastructure systems. IEEE Computer Magazine 2000;33(8):44–53.

Bang-Jensen J, Gutin G. Digraphs: theory, algorithms and applications. 1st ed. Springer-Verlag; 2006.

Barton DC, Stamber KL. An agent-based microsimulation of critical infrastructure systems. Technical report, Sandia National Laboratories, Infrastructure Surety Department, Albuquerque, NM, USA; 2000.

Broder A, Kumar R, Maghoul F, Raghavan P, Rajagopalan S, Stata R, et al. Graph structure in the web. Computer Networks 2000;33:309–20.

Brömmelhörster J, Fabry S, Wirtz N, editors. Internationale Aktivitäten zum Schutz Kritischer Infrastrukturen. Bonn, Germany: Bundesamt für Sicherheit in der Informationstechnik; 2004.

Callaway DS, Newman MEJ, Strogatz SH, Watts DJ. Network robustness and fragility: percolation on random graphs. Physical Review Letters 2000;85(25):5468–71.

Casselman W. Networks. Notices of the American Mathematical Society 2004;51(4):392–3.

Chen Q, Chang H, Govindan R, Jamin S, Shenker SJ, Willinger W. The origin of power laws in Internet topologies revisited. In: INFOCOM 2002: proceedings of the 21st annual joint conference of the IEEE computer and communications societies, vol. 2. New York, NY, USA: IEEE; 2002.

Cohen R, Erez K, Ben-Avraham D, Havlin S. Resilience of the Internet to random breakdowns. Physical Review Letters 2000; 85(21):4626–8.

Cohen R, Erez K, Ben-Avraham D, Havlin S. Breakdown of the Internet under intentional attack. Physical Review Letters 2001;86(16):3682–5.

Critical infrastructures: protecting America's infrastructures. In: Marsh RT, editor. Washington, DC, USA: United States Government Printing Office; 1997 [report of the President's Commission on critical infrastructure protection].

Dorogovtsev S, Mendes J. Evolution of networks, from biological nets to the Internet and WWW. Oxford University Press; 2003.

Dorogovtsev SN, Mendes JFF. Effect of the accelerating growth of communications networks on their structure. Physical Review E 2001;63:025101.

Edwards Jr CH, Penny DE. Differential equations and boundary value problems, computing and modeling. 1st ed. Prentice Hall International; 1995.

E.ON Netz GmbH. Bericht über den Stand der Untersuchungen zu Hergang und Ursachen der Störung des kontinentaleuropäischen Stromnetzes am Samstag, 4. November 2006 nach 22:10 Uhr. E.ON Netz GmbH, Bayreuth, Germany; November 2006.

Faloutsos M, Faloutsos P, Faloutsos C. On power-law relationships of the Internet topology. In: Chapin L, Sterbenz JPG, Parulkar G, Turner JS, editors. Proceedings of the conference on applications, technologies, architectures, and protocols for computer communication. Cambridge, MA, USA: Association for Computing Machinery; 1999.

Freeman RL. Fundamentals of telecommunications. In: Wiley series in telecommunications and signal processing. 1st ed. John Wiley & Sons, Inc.; 1999.

Gerhart PM, Gross RJ, Hochstein JI. Fundamentals of fluid mechanics. 2nd ed. Addison–Wesley Publishing Company; 1992.

Goh K-I, Kahng B, Kim D. Spectra and eigenvectors of scale-free networks. Physical Review E 2001;64:051903.

Hæsken OM, Olsen TG, Fridheim H. Beskyttelse av samfunnet (bas) – sluttrapport. Technical report. FFI/RAPPORT-97/01459, Norwegian Defence Research Establishment, Norwegian Defnece Reserach Establishment, Postboks 25, 2007 Kjeller, Norway; 1997.

Hilt D. Technical analysis of the August 14, 2003, blackout. Technical report, North American Electric Reliability Council, Princeton, NJ, USA; July 2004.

James J, Mabry F. Building trustworthy systems: guided state estimation as a feasible approach for interpretation, decision and action based on sensor data. In: Proceedings of the 37th Annual Hawaii International Conference on System Sciences (HICSS'04). Big Island, HI, USA: IEEE Computer Society Press; 2004.

Karmarkar N. A new polynomial time algorithm for linear programming. Combinatorica 1984;4(4):373–95.

Kralik J, Stiegler P, Vostry Z, Zavorka J. A universal dynamic simulation model of gas pipeline networks. IEEE Transactions on Systems, Man and Cybernetics 1984;14(4):597–606.

Newman MEJ. The structure and function of complex networks. SIAM Review 2003;45(2):167–256.

North MJ. Agent-based modeling of complex infrastructures. In: Proceedings of the simulation of social agents: architectures and institutions workshop, Chicago, IL, USA; 2000.

Pasqualini D, Witkowski M. System dynamics approach for critical infrastructure and decision support. In: Proceedings of the fall meeting 2005 of the American Geophysical Union. San Francisco, CA, USA: American Geophysical Union; 2005.

Pastor-Satorras R, Vespignani A. Evolution and structure of the Internet: a statistical physics approach. Cambridge University Press; 2004.

Rinaldi SM. Modeling and simulating critical infrastructures and their interdependencies. In: Proceedings of the 37th annual Hawaii international conference on system sciences (HICSS'04). Big Island, HI, USA: IEEE Computer Society Press; 2004.

Schmitz W. Analysis and assessment for critical infrastructure protection (ACIP). Final report IST-2001-37257 deliverable D7.5. In: ACIP consortium, Ottobrunn, Germany; June 2003.

Schrijver A. Combinatorial optimization. Berlin, Germany: Springer-Verlag; 2003 [three volumes].

Skvortsov AV, Sarychev DS. Modeling of pipeline sections. Russian Physics Journal 2002;45(2):155–62.

Sullivan K, Knight JC, Du X, Geist S. Information survivability control systems. In: Proceedings of the 21st international conference on software engineering. Los Angeles, CA, USA: IEEE Computer Society Press; 1999.

Svendsen NK, Wolthusen SD. Multigraph dependency models for heterogeneous critical infrastructures. In: Proceedings of the First Annual IFIP WG 11.10 International Conference on Critical Infrastructure Protection. Hanover, NH, USA: Springer-Verlag; 2007.

Thomas WH, North MJ, Macal CM, Peerenboom JP. From physics to finances: complex adaptive systems representation of infrastructure interdependencies. Technical report, Naval Surface Warfare Center Technical Digest, Dahlgren, VA, USA; 2003.

Watts D, Strogatz S. Collective dynamics of small-world networks. Nature 1998;393:440–2.

Yook S-H, Jeong H, Barabási A-L. Modeling the Internet's large-scale topology. Proceedings of the National Academy of Sciences 2002;99(21):13382–6.

Zhu G, Henson MA, Megan L. Dynamic modeling and linear model predictive control of gas pipeline networks. Journal of Process Control 2001;11:129–48.

Zumdal SS. Chemical principles. 3rd ed. Hughton Mifflin Company; 1998.