

Enhancing Infrastructure Resilience Under Conditions of Incomplete Knowledge of Interdependencies

L. Andrew Bollinger¹ and Gerard P.J. Dijkema²

¹Delft University of Technology, The Netherlands

²University of Groningen, The Netherlands

ABSTRACT

Today's infrastructures — such as road, rail, gas, electricity and ICT — are highly interdependent, and may best be viewed as multi-infrastructure systems. A key challenge in seeking to enhance the resilience of multi-infrastructure systems in practice relates to the fact that many interdependencies may be unknown to the operators of these infrastructures.

How can we foster infrastructure resilience lacking complete knowledge of interdependencies? In addressing this question, we conceptualize the situation of a hypothetical infrastructure operator faced with incomplete knowledge of the interdependencies to which his infrastructure is exposed. Using a computer model which explicitly represents failure propagations and cascades within a multi-infrastructure system, we seek to identify robust investment strategies on the part of the operator to enhance infrastructure resilience.

Our results show that a strategy of constructing redundant interdependencies may be the most robust option for a financially constrained infrastructure operator. These results are specific to the infrastructure configuration tested. However, the developed model may be tailored to the conditions of real-world infrastructure operators faced with a similar dilemma, ultimately helping to foster resilient infrastructures in an uncertain world.

Keywords: Interdependencies, Resilience, Networks, Electricity Infrastructures, Incomplete Knowledge, Modeling and Simulation

INTRODUCTION

Hurricane Sandy, which struck the Northeast coast of the U.S. on 29 October, 2012, was the second costliest Atlantic hurricane in history and the largest Atlantic hurricane ever¹. The infrastructure impacts of Hurricane Sandy were immense — 8.7 million customers lost power; 25% of customers lost mobile, landline, Internet and cable television service; New York City's subway services were completely shut down; and multiple oil and gas refineries/pipelines were disabled². These impacts were not only a direct consequence of meteorological conditions. They were also a result of the (inter)dependencies between infrastructures, which allowed failures to cascade from one infrastructure to another — e.g. from the electricity infrastructure to the gas and oil infrastructures, and from the gas and oil infrastructures to the road infrastructure — as well as beyond the geographical scope of the hurricane itself.

1 NOAA. NOAA's National Weather Service Newport/Morehead City, NC event summaries / case studies. Web site, October 2012. URL <http://www.erh.noaa.gov/mhx/EventReviews/20121029/20121029.php>. Accessed May 2014.

2 Comes, T. and van de Walle, B. Measuring disaster resilience: The impact of Hurricane Sandy on critical infrastructure systems. In *Proceedings of the 11th International ISCRAM Conference*, 2014.

Due to the many (inter)dependencies between different types of infrastructures and the potential for disruptions to cascade between infrastructures, it is of limited use to assess the resilience of different infrastructures independently. Often, it is more productive to view infrastructures as *multi-infrastructure systems* – interacting networks of infrastructure components featuring various (types of) dependencies. However, a key challenge in seeking to enhance the resilience of multi-infrastructure systems in practice relates to the fact that many interdependencies may be unknown to the actors responsible for operating and safeguarding these infrastructures³. Part of this challenge may be attributed to the organizationally fragmented nature of multi-infrastructure systems, which limits the control and knowledge of individual actors⁴.

Actions by infrastructure operators to identify and catalogue infrastructure interdependencies can be an important step to enhancing resilience. However, complete knowledge of infrastructure interdependencies is an elusive goal, as interdependencies are inherently “dynamic and situational”⁵. Infrastructures are constantly changing and co-evolving⁶, meaning that new (hard and soft) interdependencies may continuously arise. Moreover, precisely what constitutes an interdependency may depend on the situation. These realities limit our capacity to develop accurate models of multi-infrastructure systems, thus restricting our ability to effectively foster resilient infrastructures.

How can we foster infrastructure resilience lacking precise knowledge of interdependencies? In addressing this question, we conceptualize the situation of a hypothetical operator of an electricity network faced with incomplete knowledge of the interdependencies between his infrastructure and other infrastructures. Using a computer model which explicitly represents failure propagations and cascades within a multi-infrastructure system, we explore the effectiveness of different investment strategies on the part of the infrastructure operator to enhance the resilience of his infrastructure.

BACKGROUND – MODELING INTERDEPENDENT INFRASTRUCTURES

Due both to the socio-technical complexity of multi-infrastructure systems and the impracticality of experimenting with massive infrastructure failures in the real world, modeling and simulation can play an important role in helping to untangle the potential consequences of disruptions in multi-infrastructure systems. Ouyang⁷ differentiates between four different approaches within this body of work – empirical approaches, agent-based approaches, system dynamics approaches and network-based approaches. *Network-based approaches* normally involve the representation of infrastructure components as nodes and the links or interdependencies between them as edges. An example of this type of analysis is introduced by Lam et al.⁸, who use a systematic set of experiments to identify critical nodes in a generic multi-infrastructure network based on the giant component size.

The model introduced in this paper follows a network-based approach. However, rather than assuming complete knowledge of the multi-infrastructure network, we test different strategies for dealing with the problem of incomplete network knowledge. This problem has also been addressed by Tai, et al⁹, who suggests that, rather than seeking to determine the effects of possible failures in a network, we should solve the inverse problem – to identify those networks that may result in “the most extreme disruptions”. These networks may then be used to identify scenarios in the real world that could potentially lead to the realization of such disruptions⁹.

We adopt a different approach. We begin by conceptualizing the situation of a boundedly rational operator of a hypothetical electricity infrastructure. We assume that this operator lacks complete knowledge of the interdependencies between his infrastructure and a connected infrastructure, over which he has no control.

3 Robert, B. and Morabito, L. An approach to identifying geographic interdependencies among critical infrastructures. *International Journal of Critical Infrastructures*, 6(1):17–30, 2010.

4 De Bruijne, M. and van Eten, M. Systems that should have failed: Critical infrastructure protection in an institutionally fragmented environment. *Journal of Contingencies and Crisis Management*, 15(1):18–29, 2007.

5 U.S. DOE. Overview of response to Hurricane Sandy nor’easter and recommendations for improvement. Technical report, February 2013.

6 Nikolic, I. *Co-evolutionary Method for Modeling Large-scale Socio-technical Systems Evolution*. PhD thesis, TU Delft, 2009.

7 Ouyang, M. Review on modeling and simulation of interdependent critical infrastructure systems. *Reliability Engineering & System Safety*, 121:4360, 2014.

8 Lam, C.Y., Lin, J., Sim, M.S. and Tai, K. C. Y. Identifying vulnerabilities in critical infrastructures by network analysis. *International Journal of Critical Infrastructures*, 9:190–210, 2013.

9 Tai, K., Kizhakkedath, A., Lin, J., Tiong, R. and Sim, M. Identifying extreme risks in critical infrastructure interdependencies. In *International Symposium for Next Generation Infrastructure*, 2013.

Additionally, we assume that the operator faces significant uncertainty concerning the severity and frequency of future disruptive events to which his infrastructure may be exposed. Given the massive uncertainties surrounding looming threats such as climate change and (cyber-)terrorism, this is a very real challenge facing today's infrastructure operators.

Starting from this basis, we ask the following question: *How should an infrastructure operator ideally invest to ensure a resilient network in the face of these uncertainties?* In addressing this question, we draw from Lempert, et al¹⁰, who suggest that, in dealing with such situations, decision makers should not seek optimal strategies but rather robust strategies – “strategies that that perform ‘well enough’ by meeting or exceeding selected criteria across a broad range of plausible futures”. With this in mind, we place our hypothetical infrastructure operator in a simulated environment consisting of a (known) electricity network and an (unknown) interdependent second network, both of which may be subjected to a range of disturbances. We assign the operator four possible strategies for enhancing the resilience of his electricity network, and examine how these strategies perform under different conditions.

MODEL DESIGN

The starting point for the model is a network representation of a generic electricity infrastructure (infrastructure A). The setup of infrastructure A is based on the *IEEE 118 bus power system test case*¹¹, which represents an archetypal electricity transmission infrastructure. Infrastructure A includes three types of nodes – representing electrical substations, power generation facilities and power consuming facilities, or loads – and one type of link – representing power lines or electrical transformers. Infrastructure A is augmented with links to a second infrastructure network of a different unspecified type (infrastructure B). Infrastructure B could be conceptualized to represent a road infrastructure, natural gas infrastructure, IT infrastructure, etc. Infrastructure B is composed of a set of randomly connected nodes and edges, and features only a single type of node and a single type of edge. Infrastructures A and B feature a number of common links in the form of interdependencies. For the sake of simplicity, we assume that all of these interdependencies are bi-directional – that is, each represents a dependency of infrastructure A on infrastructure B and vice versa.

Each run of the simulation consists of 100 timesteps, each representing a timespan of one year. Every timestep during the course of a simulation, we introduce a set of failures, which correspond to the disabling of a set of links in the multi-infrastructure network. It is assumed that all of the failures within a given year occur simultaneously – that is, each year a single set of concurrent failures occur. These failures may affect links of infrastructure A or infrastructure B, or the interdependencies between them. The number of failures occurring each timestep is randomly determined according to a power law size-frequency distribution. Approximate power law size-frequency distributions are typical of many types of natural disasters¹², as well as electricity blackouts¹³. Precisely which links are affected by a failure is randomly determined, and all links are returned to working order at the start of the next timestep.

The failure of a single link may affect not only that link, but may also result in a cascade of failures through the multi-infrastructure network. Failure cascades through the electricity network (infrastructure A) are determined by an iterative power flow algorithm. Following a set of initial failures, the algorithm calculates anticipated power flows through the network. If the calculated power flow across any line exceeds its capacity, the line is assumed to fail due to overload. This results in an altered distribution of power flows, which may cause additional lines to fail. This process is repeated until no more power lines are overloaded. Failure cascades in infrastructure B are determined in a simpler manner – if a link fails, there is a 10% chance that each of its neighboring links will fail that timestep.

Via interdependency links, failure cascades may cross from infrastructure A to infrastructure B, and vice versa. Each time a failure cascades from one infrastructure to the other, it induces (a) failure(s) in the latter. This may cause the failure to cascade (via another route) back to the first infrastructure, and so on. In this manner, we represent not only the propagation of failures within a particular infrastructure, but also the potential for first, second, third and higher order interdependency effects.

¹⁰ Lempert, R., Popper, S. and Bankes, S. Shaping the next one hundred years: New methods for quantitative, long-term policy analysis. Technical report, RAND Pardee Center, 2003.

¹¹ Christie, R. Power systems test case archive. Web site, May 1993. URL http://www.ee.washington.edu/research/pstca/pf118/pg_tca118bus.htm. Accessed May 2014.

¹² Barton, C. and Nishenko, S. Natural disasters - forecasting economic and life losses. Web site, 2013. URL <http://pubs.usgs.gov/fs/natural-disasters/>. Accessed November 2013.

¹³ Dobson, I. Carreras, B.A., Lynch, V.E. and Newman, D.E.. Complex systems analysis of series of blackouts: Cascading failure, critical points, and self-organization. *Chaos*, 17:1–13, 2007.

The model is implemented in the agent-based modeling platform *Netlogo*¹⁴, and uses the MATLAB-based power system simulation package *Matpower*¹⁵. Runtime communication between these pieces of software is enabled via a custom-developed Netlogo extension, the *MatpowerConnect* extension.

REPRESENTATION OF INVESTMENT STRATEGIES

A hypothetical operator of infrastructure A seeks to ensure the infrastructure's resilience in the face of (unknown) future failures. However, the operator's task is complicated by several factors. First, he does not know exactly how severe future failure scenarios may be, so he does not know precisely how robust his network must be to withstand these failures. Second, the operator does not know where or how many interdependencies exist between his infrastructure (infrastructure A) and infrastructure B. Third, the operator has no control over and limited knowledge of infrastructure B – that is, he cannot influence the design of that network nor can he predict possible failures within it.

The operator of infrastructure A is given a set of four possible strategies to enhance the resilience of his electricity network. The strategies are as follows:

- **Focus on the critical links:** Each time a link of the electricity infrastructure fails due to capacity overload, the operator increases the link's capacity by the magnitude of the overload.
- **Focus on the interdependencies:** Each time an unknown interdependency is revealed (due to the failure of an interdependency link), the operator constructs a redundant interdependency link, essentially obviating the possibility for future failure of that link.
- **Pre-emptively increase capacities:** The operator increases the capacity of all power lines by 50% at the start of the simulation.
- **Combination:** Combination of strategies 1 and 2.

In addition to these four strategies, we include a null strategy (strategy 0), corresponding to no action on the part of the infrastructure operator. Each investment made by the infrastructure operator is assumed to entail a certain cost. The cost of one additional redundant interdependency link or power line is assumed to be 1 monetary unit. The cost of one unit of additional power line capacity is assumed to be 1 monetary unit, divided by the mean capacity of power lines in the system.

EXPERIMENTS

We have carried out a set of 60 experiments with the developed model. The purpose of these experiments is to identify robust strategies for enhancing infrastructure resilience – strategies that perform sufficiently well across a range of possible futures. The performance of the different strategies is evaluated using two metrics – resilience and cost. *Resilience* is defined as the mean performance of infrastructure A across all timesteps of all simulations employing that particular strategy, with performance quantified in terms of the fraction of demand served – the mean fraction of power received by customers vs. that which was demanded. *Cost* is defined as the sum of all expenditures of the infrastructure operator in making upgrades to his network, again averaged across all simulations employing the respective strategy. In the course of experimentation, we vary the values of three key parameters: (1) the strategy employed by the infrastructure operator, (2) the number of interdependencies between the electricity infrastructure and the second infrastructure, and (3) the severity of failure scenarios.

14 Wilensky, U.. Netlogo. Web site, 2012. URL <http://ccl.northwestern.edu/netlogo/>. Accessed February 2013.

15 Zimmerman, R., Murillo-Sanchez, C., and Gan, D. Matpower: A MATLAB power system simulation package. Web site, 2011. URL <http://www.pserc.cornell.edu/matpower/>. Accessed February 2013.

RESULTS AND ANALYSIS

Figure 1 illustrates key results from the developed model. These plots show the spread of observed resilience and cost values observed across each of the tested strategies. In terms of resilience, the top performing strategies are strategies 4 and 2, corresponding to mean resilience values of 0.87 and 0.85, respectively. Both of these strategies involve focusing on constructing redundant interdependency links between infrastructure A and B, suggesting that failure cascades between infrastructures A and B play an important role in affecting the performance of infrastructure A. It is important to note, however, that – under certain circumstances – a strategy of focusing exclusively on the construction of redundant interdependency links (strategy 2) can still result in relatively low resilience values – as low as 0.66. A strategy combining the construction of redundant interdependencies with the construction of additional capacity for the critical links of infrastructure A (strategy 4) has the effect of increasing the minimum observed resilience value from 0.66 to 0.71.

As illustrated in Figure 1, the best performing strategy in terms of cost (excluding strategy 0) is strategy 2, which corresponds to a mean investment of approximately 46 monetary units. On average, strategy 1 is significantly more expensive, featuring a mean investment of approximately 114 units. However, it is interesting to note that, under certain conditions, strategy 1 can result in lower costs than strategy 2 – as low as 7 monetary units. A closer look at the results reveals that these conditions occur only in situations with a very low severity of failure scenarios. Thus, if the infrastructure operator is relatively confident in the limited severity of future failure scenarios, this may be an attractive strategy.

Which strategy is most robust depends on the priorities of the infrastructure operator and which values for the different metrics may be considered “good enough”. If the operator prioritizes resilience and prefers minimum resilience values above 0.7, the most robust option would be to follow a strategy combining the construction of redundant interdependencies with the construction of additional capacity for the critical links of infrastructure A (strategy 4). If the operator seeks to minimize investment costs while still maintaining relatively high resilience values, the most robust option would be to follow strategy 2, and focus exclusively on constructing redundant interdependencies.

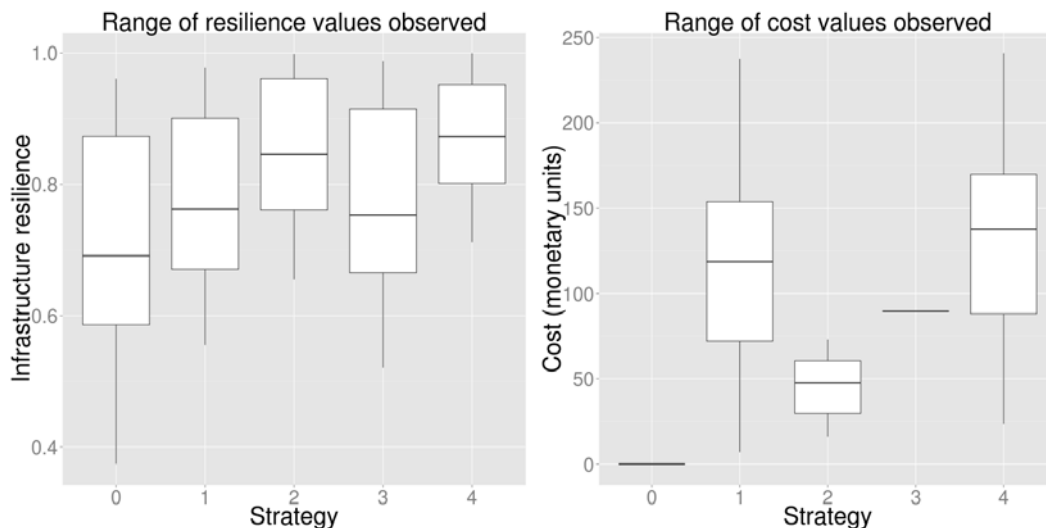


Figure 1: Results of the developed model: box plots showing the range of observed resilience values (left pane) and cost values (right pane). The ends of the box indicate the 25th and 75th percentiles.

CONCLUSIONS

Uncertainty concerning both the locations/types of infrastructure interdependencies and the severity/frequency of future failure scenarios are very real problems faced by today's infrastructure operators. In this paper, we have introduced a model exploring different long-term strategies for an infrastructure operator to invest in light of these uncertainties. The results from the developed model show that a strategy focusing on the construction of redundant interdependencies may be the most robust option for a financially constrained infrastructure operator. It must be kept in mind that these results are specific to the infrastructure configuration tested and rely on a set of assumptions that may not hold in the real world. However, the approach used and the developed model may be useful insofar as they can be tailored to the specific conditions of real-world infrastructure operators faced with a similar dilemma. For instance, the generic network structures used in the current model can be replaced by real-world networks of different types (e.g. ICT, transport, water), and a broader array of investment strategies may be assessed. In this manner, the developed model can be a stepping stone towards fostering resilient infrastructures in an uncertain world.

ACKNOWLEDGEMENTS

This research is supported by the 'Knowledge for Climate' research programme in the Netherlands.