Understanding and Managing the Threat of Disruptive Events to the Critical National Infrastructure

Professor Kim Warren and Dr Robert Thurlby

Kim Warren: Teaching Fellow, London Business School kim@strategydynamics.com

Robert Thurlby: Utility Industry Consultant bob@thurlby.me.uk

Address for correspondence:

Kim Warren
Two Farthings
Aylesbury Rd
Princes Risborough
BUCKS HP27 0JS
+44 1844 274061

Please note that the model supporting this paper is available – its structure can be inspected in a browser at RRR Utility Model. To inspect its equations and parameters requires registration with the online application in which it was built (free). However, please note that these steps will both make clear to the reviewers who the authors are, and make clear to the authors who the reviewers are.

Understanding and Managing the Threat of Disruptive Events to the Critical National Infrastructure

Abstract

Concern about the vulnerability of utility networks, (electricity gas and water) and other infrastructures, including transport and telecommunications, to environmental, terrorist and other threats has increased in recent years, motivated both by a perceived increase in such threats and by recognition that the commercial pressures and regulation of companies operating these infrastructures could unintentionally have increased that risk. Powerful simulation tools already help utility operators assess the physical consequences of disruptive events on their networks, whilst others have helped increase their capability to respond efficiently when such events occur. However, better understanding is needed of the relationship between operational, commercial and regulatory pressures, the strategic choices these lead to on the part of infrastructure operators and the long-run consequences for the resilience of these systems and hence for service continuity. This paper describes a high-level model portraying these relationships, and early findings from testing alternative strategies, both over the long and short term.

Keywords

utilities, critical national infrastructure, risk, resilience, disruption, service continuity.

Introduction

Utilities, and other organisations who operate a country's critical national infrastructure (CNI), historically have a high level of competence concerning the assessment and mitigation of threats to the security and integrity of infrastructure they manage. They also have well established and tested procedures for recovering the infrastructure in the aftermath of a disruptive event.

However in recent years new threats have emerged and the probability of known threats occurring has increased. As a result, traditional approaches to risk assessment and recovery planning have started to be questioned as to their continued relevance and fitness for purpose. Additionally users of the CNI, typically the general public, have become far less tolerant of a loss of service and, encouraged by customer focused regulators, are demanding ever higher levels of performance from the CNI.

Prevention, mitigation and management of disruptive events has therefore become an even higher priority for utilities and the rest of the CNI and new methods and tools are being sought to help improve capability. Many powerful simulation tools already exist to understand how networks could be affected physically by major incidents. Others , help organisations develop the readiness to respond to such incidents, often by war-gaming approaches. The part of the problem that is less well understood is the relationship between long-term, strategic choices and the ability of infrastructure networks to withstand disruptive events. Those choices concern investment in the assets themselves, in the IT infrastructure, especially the network control systems, and in the people managing the system. Whilst it is clear enough that "spending less on assets, systems and people will degrade the system", it is not so obvious how much impact any particular choice will have over long periods of time, nor how choices on different issues will interact.

The issues that need to be better understood are therefore:

- how long-term choices on strategic issues make the network more resilient (less likely to be damaged by a disruptive event)
- how these and other choices can minimise the service loss when disruptive events do occur
- how strategic and operational choices can minimise the time for the network to recover, and thus the total cumulative loss of service

In the past, leading utilities, transport operators and the oil industry have successfully used system dynamics modelling to help address policy and strategy concerning the investment in infrastructure assets. This paper describes initial findings from a model intended to help organisations in the CNI explore and understand how such strategies affect the resilience of their networks to disruptive threats and the consequences for service continuity and financial performance.

The Changing Nature of Threats

Until recently, the nature of disruptive threats to any part of the CNI was well understood. Infrastructure was designed to be able to resist the events to a certain point and then behave in such a way that failure was orderly and somewhat controllable, so that recovery after the event would be efficient. Maintenance programmes were designed to minimise failures and infrastructure was built to be resilient, through redundancy, interconnection and higher operational specification.

As a result, the failure of infrastructure components under normal operating conditions was rare. Disruptive events such as storms, human intervention and error were equally rare, and emergency procedures could minimise the effect of these and accelerate recovery. This resilience, though, came at a high cost, and the wider external environment started to change:

- The privatisation of much of the CNI that has occurred in many countries had substantial benefits, but also had less desirable consequences. The resulting drive for efficiency and profitability led many organisations to adopt a policy of "sweating the assets" by reducing maintenance and delaying replacement programmes. Whilst this worked for a time, the infrastructure became increasingly aged and unreliable. This also made it more vulnerable to external disruption - what has become known as the asset time bomb[1]. In this context, "assets" also include the organisations' human resources, because many organisations reduced staffing levels as part of their efficiency improvement strategies.
- Privatisation also shortened the time-horizon for both financial and regulatory objectives. A financial cycle of 1 year and regulatory cycle of 4 years led to strategies focused on performance targets for similarly short periods. But with an asset life cycle of 40 years or more, these strategies reduced still further the weakened further impact on the long term health of the infrastructure.
- Climate Change caused weather patterns to become both more erratic and more extreme. Utility networks, built to withstand everything but the "100 year storm", were being subject to severe events far more frequently.
- A new threat of terrorism and in particular cyber-terrorism also emerged. Agencies aiming to cause damage for economic or political reasons realised that disrupting a country's infrastructure is a powerful option. They also realised that damaging the relevant control system is often the easiest and least risky method for doing this.
- Lastly, as the CNI has become more sophisticated and more complex, its various parts have become more dependent on each other. This increased the risk of an event spreading across different parts of the CNI.

The result has been that in the last 20 years both the threat of disruptive events damaging the CNI and the risk that they will occur has increased in both diversity and probability.

Weaknesses of Current Response Capabilities

Organisations that manage and operate the CNI all recognise that risks and threats to their networks are unavoidable. Consequently their networks are constructed and operated to be able to resist disruption and recover from events either automatically or efficiently. This means that for minor events that are quite likely, such as wind damage to an overhead line or another organisation cutting through an underground cable, well-established and tested generic responses are in place. Where the risk is much larger, risking damage to a large or critical part of the network, such as loss of supply to a large financial or industrial centre, strategies and plans are in place for both mitigating the event and recovering from it when the limits of mitigation are exceeded. Mitigation includes reinforcing and protecting the infrastructure to make it more resistant to failure. Additionally, better prepared organisations' mitigation strategies have been integrated with the asset maintenance strategies to ensure that the impact of known risks and weaknesses on overall performance is minimised (Figure 1).

Attention has thus focused on high probability, high impact threats that are somewhat obvious. This strategy is, though, no longer adequate. Known threats have migrated towards the high impact, high probability quadrant and new threats have emerged, particularly in the low probability, high impact quadrant of Figure 1. What is more, the possible consequences from these increasingly unpredictable risks can be much more severe and widespread than the more obvious examples for which contingencies already exist. A serious flood, for example, may not only disrupt parts of a power network, but also its interconnections with other parts of the CNI, threatening still-more serious consequences. Loss of a power substation for a long period due to a flood, for example, could disrupt water supply services that would not be affected by a serious but shorter loss of power.

Figure 1: Risk Matrix

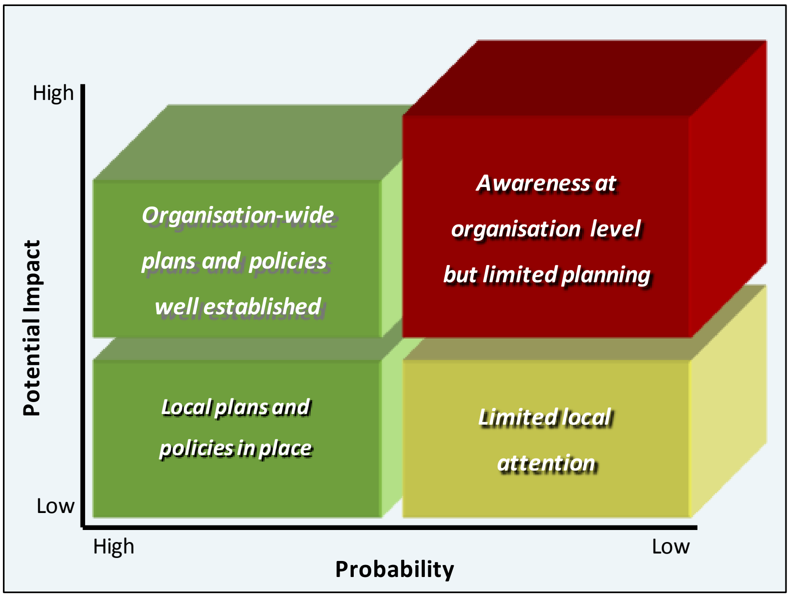
High
Potential Impact
Low
Organisation-wide plans and policies well established
Awareness at organisation level but limited planning
Local plans and policies in place
Limited local attention
High
Low
Probability

The existing approach to risk assessment is no longer adequate for the more dynamic environment in which the CNI exists, especially as it is now older and more vulnerable to age-related threats. That this situation existed in reality was confirmed in the UK by events such as the Severn Floods in 2007, where water levels came within centimetres of taking out a power substation serving the whole south-west of the country, which in turn would have cut off water and telecommunications services to many parts of the region. This demonstrated the increased vulnerability of the CNI to a weather event, the event having occurred in the summer period, and the weaknesses of organisations involved to effectively co-ordinate recovery. Risk, resilience and recovery strategies and policies, defined for a known set of risks to a strong and resilient infrastructure, have thus become unsuited and inappropriate for the increased and more diverse threat levels to infrastructures that are now less resilient to those threats.

The main weakness was a failure to explore in advance different risk scenarios, establish the probability of such scenarios occurring and put in place, also in advance, strategies that would mitigate the consequences and facilitate recovery. The challenge is that, since the source of threats is increasinly diverse, it is no longer sufficient to focus on making specific parts of the system resilient that are known to be at risk. Rather, strategy needs to raise the resilience of the whole network to whatever events might occur.

The requirement to understand how a current situation has come about, the direction in which it might develop in future, and how to change that future for the better is not unique to risk and resilience, but is fundamental to all policy and strategy situations.

Prior work on modelling risk and resilience in utilities

As might be expected, the issue of threat, risk and resilience of the utilities that make up the CNI has become a topic of considerable interest and concern since the terror attack of 9/11 and the evidence of more frequent and serious natural events. A multitude of groups, from widely diverse disciplines have applied a bewildering array of methods to study the issue. A comprehensive, though now somewhat dated survey (given the rapid developments in modelling) was carried out by the Idaho National Laboratory in 2006[2].

Agent-based and network modelling methods are widely deployed, often in considerable detail and including geospatial data about network assets, in an effort to capture accurately the response of particular infrastructure networks. Many such models are intended for developing organisational readiness to respond to events, and for training, often by means of war-gaming exercises. Models range from detailed attention to occurrences and responses of individual assets up to whole-system models for an entire infrastructure system (power, telecommunications and so on). Some go so far as to capture the interactions between different infrastructure networks.

System dynamics appears not to have been used extensively for modelling the dynamics of infrastructure risk, resilience and response, possibly because of the perceived need to deal with geospatial aspects of the problem (the proximity of assets to each other and to residential and commercial neighbourhoods) and the mechanical behaviours of physical assets. One reported model does utilise Vensim to simulate the dynamics of individual infrastructures and links between them[3].

Although existing simulations for modelling risk and resilience are powerful and highly-developed, the present work is motivated by two perceived shortcomings. First, few models appear to address the long-term strategic choices leading to a network's resilience and speed of recovery. Secondly, most models are extremely large, detailed and complex, requiring considerable expertise and experience from their users. Some work has been done on the use of system control design to assess infrastructure resilience[4], but we believe system dynamics offers additional advantages on these two issues. Our aim, therefore, is to develop a model that demonstrates the impact of the few large, strategic choices made by infrastructure operating companies over long periods of time.

Apart from the obvious concern with disruption and danger to society, a related motivation for modelling the relationship between strategic choices and risk/resilience comes from those concerned with the commercial regulation of utilities. Regulators are now much more aware of the link between the regulatory regimes they impose on utilities, their impact on the investment and other decisions taken by those utilities, and the potential for adverse consequences. Guthrie (2006[5]) reviews the literature on this issue, and highlights the relevance of modern investment theory, which which highlights risk and the irreversibility and delays involved in investment decisions. The paper discusses the impact of different regulatory regimes and the length of the regulatory cycle on utilities' investment decisions.

Before embarking on an effort to model risk and resilience, it is important to clarify the definition of terms that arise in the discussion. The term "risk" is used in multiple fields, with widely differing meanings. In the context of infrastructure, not only is the term itself used somewhat inconsistently, but diverse meanings are also applied to related terms, such as "susceptibility", "vulnerability" and "resilience". In the context of the present topic, it is important to define these terms appropriately as they apply to the issue of infrastructure management.

The first step in this task is to distinguish between the likelihood of adverse *events* that have the potential to cause disruption, and the likelihood of adverse *consequences* arising from those events. The Oxford English Dictionary defines risk as "*Exposure to the possibility of loss, injury or other adverse or unwelcome circumstance*", implying that risk relates to consequences, rather than the events that may cause them. This is supported by terminology in the ISO31000:2009 Risk Management Standard, which defines risk as the "*effect of uncertainty on objectives*"[6]. More specifically as concerns infrastructure, Lowrance (1976[7]) - widely cited – defines risk as "*a measure of the probability and severity of adverse effects*". This definition, too, focuses on the adverse consequences arising from a disruptive event.

Definitions of other terms are not so clear or consistent. "Resilience" refers variously to the ability of the infrastructure system itself to withstand a disruptive event without damage, its ability to recover from such damage, or the ability of the system, or of the community or society of which it is a part, to recover from the consequences of the damage (Boin and McConnell, 2007[8]). "Vulnerability" may refer to a particular type of disruptive event, a specific weakness in the system, the general exposure to hazard, or simply to a collection of risks (Ezell, 2007[9]).

In an attempt to avoid ambiguity in this paper and the supporting model, "disruptive events" are occurrences outside of the infrastructure system itself that have the potential to cause assets to fail. The "severity" of the event indicates its scale – the strength and duration of a storm, the quantity and extent of flood-water, or the number and complexity of hits in a cyber-attack. "Risk" refers to the probability and scale of failures that result from a disruptive event of a certain severity.

"Resilience" is *not* used to refer to the ability of the system – either the infrastructure itself or the wider community – to bounce back from a failure, but much more narrowly to specify the resistance of an individual asset to damage by an event. Each asset has a probability of being damaged or destroyed by an event of a certain severity, but hardening or protecting it reduces that probability.

A Systemic Approach

As was identified by the Pitt Review[10], any improved approach to risk and resilience must be systemic in nature. The Review defined this requirement by several criteria. It should consider the system as a whole, not just individual components or subsets in isolation. The forces that cause the disruptive events and the events should themselves be part of the system. It must capture the feedback between parts of the system (i.e. the effects caused by actions on the system can themselves be the cause of further effects). It should recognise that actions can both improve the situation and make it worse (i.e. the feedback can be positive and negative).

- Actions taken in advance to create or eliminate an effect may have delays in achieving this.
- The effect of an action can change over time as the system changes.
- The process of development and adaption of the system is continuous.

Meeting these criteria requires a process similar to that shown in Figure 2. The approach requires, at its core, a modelling and simulation engine which will support the systemic requirements identified above. System dynamics modelling has already demonstrated its suitability for strategic planning in this context, and is already used by a number of European Utilities to improve their asset investment planning strategies. Probably the best example is that of RWE Energie[11]. (Ref.5)

Figure 2: A Systemic Approach to Risk and Resilience Policy Modelling[12].

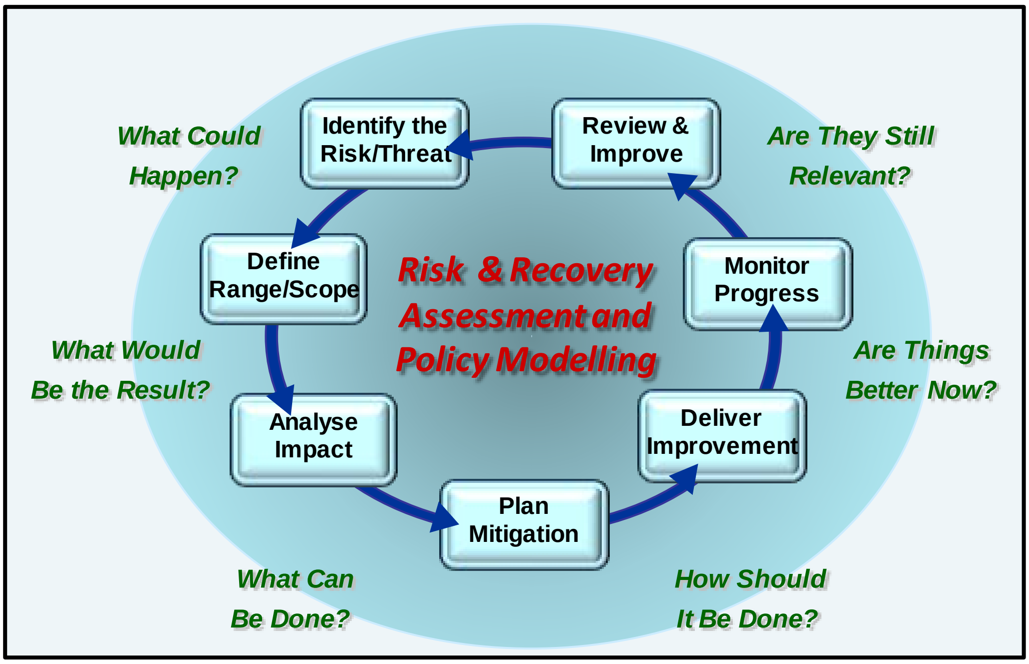

1. Modelling a strategic approach to risk and resilience

The core of the model reported here starts with a stock of 100,000 units, all of which are working before any disruption takes place (Figure 3). When an event occurs, a fraction of units fail, some being repairable, and others beyond repair. The event can lead to an immediate follow-on disruption, and the failure of some units can ‘infect’ others, by exposing them to more stress. Both repairs and the replacement of dead units are carried out by technicians, supervised by engineers. Each technician can carry out only a limited number of repairs each day, and replacements are only possible if spares are available. The normal level of spares is too low to cope with more than common, minor disruptions, so larger failures require more spares to be ordered. These arrive after a delivery lead-time, which gets longer as the number of units ordered rises. At times when no emergency has to be dealt with, technicians carry out routine maintenance and upgrading of equipment, which is reflected in the reliability of the assets in the system.

Figure 3: The asset-damage chain

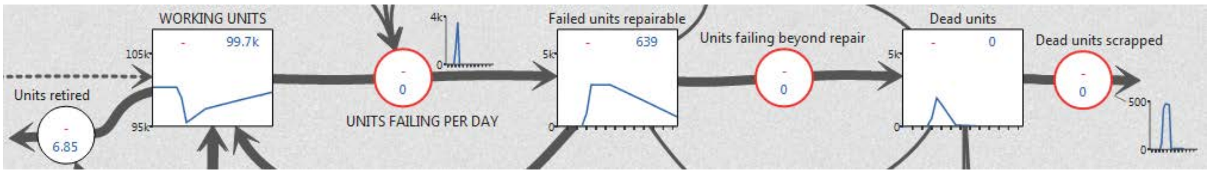

The model is initialised with parameters reflecting decisions – assumed to have been followed for a long period of history – concerning spending on the assets, systems and staffing. The infrastructure’s initial condition, then, reflects whether management has been spending and investing adequately or not, up to the model’s start date (Table 1).

From the start date, the model runs in days, and the time-scale can be set short, to inspect behaviour around a specific event, or longer, to investigate costs, performance and cash-flow consequences for different scenarios of disruptions, and for different management decisions, over many quarters or years. The following description concerns a 25-day period, with a disruption on day 5, followed by a period in which damaged units are repaired or replaced.

Table 1 summarises the impact of each decision, both on the initial state of the assets and the system, and on how that state changes after the model starts to run. Since a disruption occurs and is recovered over a very short period, the model automatically diverts activity and spending onto that task until the recovery is complete. The consequences of an event on day 5, therefore, can only reflect the initialisation decisions, because any changes to those decisions will not have had time to alter the overall state of the system. The consequences of an event occurring later, however, can reflect changed decisions to some degree, depending on how long after the simulation’s start the event occurs.

Table 1: The impact of decision-items on the initial state and continuing development of the infrastructure system

Decision	**Impact on initial state**	**Impact after the model starts**
Normal Capex	Determines the average age of assets in the network (more spend = younger units).	Determines if this age gradually increases or falls.
Capex on redundant units	Provides redundant backups for especially critical units.	Continues or raises the number of redundant units.
Capex on physical resilience	Protects or hardens a fraction of units against failure.	Continues or raises the fraction of protected or hardened units.
Spares levels (days of cover at normal usage)	Enables the immediate replacement of a limited number of dead units, after an event.	Higher spares levels can be set, but achieved only after a delivery delay.
Maintenance Opex	Determines the health of the assets, and hence their ability to resist damage.	Continues or raises the health of the assets.
ICT Opex (information and communications technology)	Makes units resilient through automatic shut-down or reconfiguration responses.	Continues or raises ICT-based resilience.
Technician numbers	With maintenance spending, starts assets in a healthy state + determines the time to fix breakdowns.	Continues or raises healthy asset condition + determines time to fix breakdowns.
Engineer numbers	Determines stress after a disruption, quality of decision-making, and time to fix breakdowns.	Continues or raises quality of decision-making.

The more units in the system that fail, the more customers are cut off, leading to a loss of revenue and to regulatory penalties. As units are repaired or replaced, the number of customers who are cut off gradually falls.

The company's revenue is assumed to come from a fixed charge rate per day to all consumers. Opex – both the decision-items above and other unspecified operating costs – is deducted from revenue to compute quarterly operating profit and Capex is deducted to calculate cash flow.

A small number of key parameters determines how the model responds to a disruptive event. A "failed fraction infection rate" sets how many more units are damaged because of units initially knocked out. A "risk of follow-on disruption" causes an initial disruption to be followed immediately by another in a certain fraction of cases. A "vulnerability multiplier", based on the already-failed fraction of units, increases the number of additional units damaged by such a follow-on event.

Loss of service reflects the three main classes of impact that spending choices have – the number of units damaged or destroyed (resilience), the daily loss of service that results from that damage, and the time taken to recover and restore the system to full working order. Each decision item in the model has a distinctive impact on those three consequences (Table 2).

Table 2: Contribution of increased expenditure on the main objectives of the policy.

Decision	Resilience	Reduced loss of service	Shorter recovery time
Normal Capex	✓	✓	-
Capex on redundant units	✓	✓✓✓	-
Capex on physical resilience	✓✓✓	✓✓	-
Spares levels	-	-	✓✓
Maintenance Opex	✓	✓✓	-
ICT Opex	✓✓	✓✓✓	-
Technician numbers	✓✓	✓✓✓	✓✓✓
Engineer numbers	-	✓	✓✓

The principal feedback structure of the operational parts of the model is shown in Figure 4. When a disruption occurs, two powerful reinforcing effects occur. First, units that are directly damaged put additional stress on others, adding to the number that fail. In severe cases, this can cascade to the point where whole regions cease to receive service as happened in the East Coast of the USA in 2003, and Northern Italy in the same year. In less severe cases, the cascade is stopped by the reducing number of further units that can be damaged (this balancing mechanism is omitted, for clarity).

Secondly, the damaged units make the network more vulnerable to any subsequent disruptions that occur before the consequences of the original event are rectified. The more severe those initial consequences – either due to the severity of the event, or the poor state of the network when it occurred – the longer it takes to rectify and so the longer the network is at further risk.

The ability to rectify all the damage is constrained by a simple operational constraint, in the number of staff available to do the work, but can also be hit by a strong balancing feedback – replacement-unit delivery delays. The more units are destroyed, and so need total replacement, the longer the delivery delay for those replacements becomes. Since both utilities and their suppliers have strong financial incentives to minimise the value of costly un-used units, these delays can escalate very sharply. This extends the time needed to totally rectify the problem.

Figure 4: Principal operational feedback in the risk/resilience model

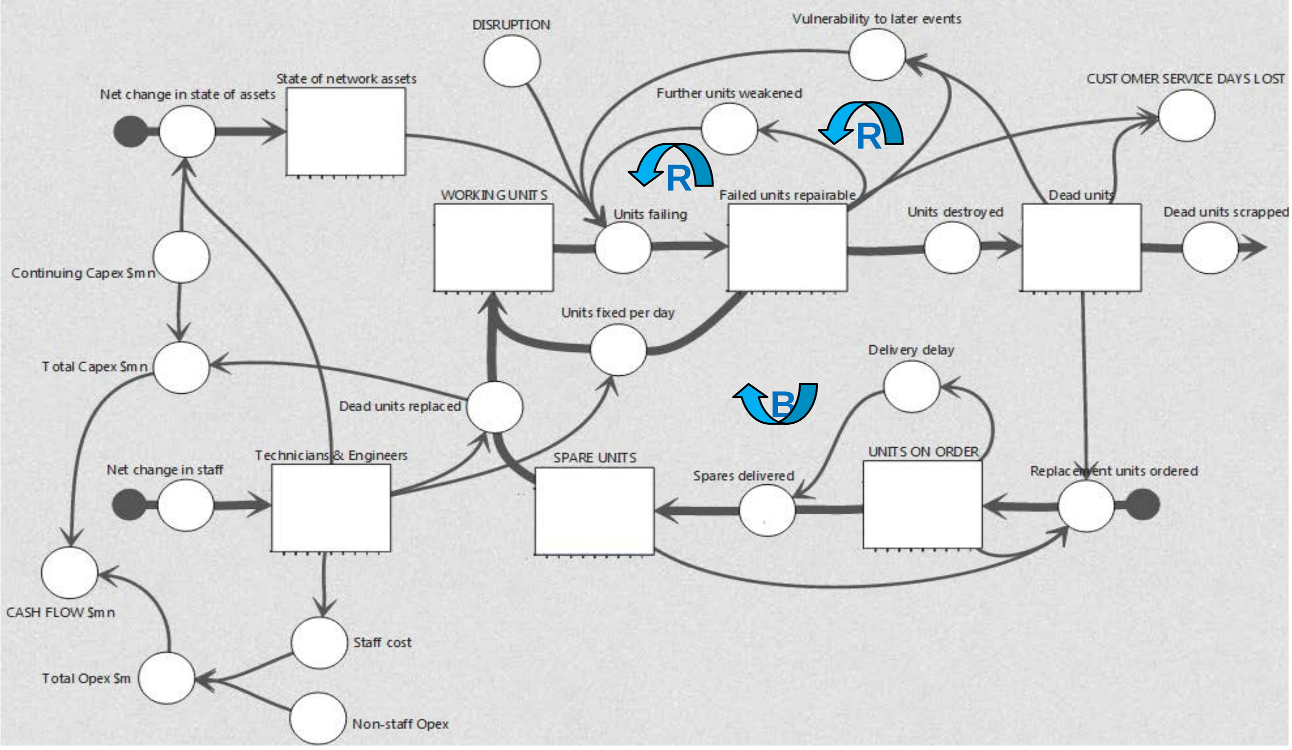

This operational feedback structure summarises the reaction challenge faced by management when a disruption occurs, and is embedded in the formulation of the model. The strategic challenge, however, is at a higher level and concerns the balancing of financial and service level objectives. This is shown in the implicit structure of Figure 5.

In summary, higher expenditure is desired in order to build and maintain the state of the network at a high level, and any shortfall in service level exerts pressure to raise expenditure. High expenditure also increases the level of cumulative investment in items – hardware, software and people – whose existence helps sustain service levels for any given state of the core system. However, high expenditure reduces cash flow, relative to a desired rate, creating pressure to reduce that spending.

At present, this policy feedback is implemented manually, though a development of the model will attempt to implement effective policy-feedback rules to arrive at high-performance outcomes. This is complicated by the number of distinct expenditure categories available, their differential impact and their interdependence.

Figure 5: Implied policy feedback structure of the model

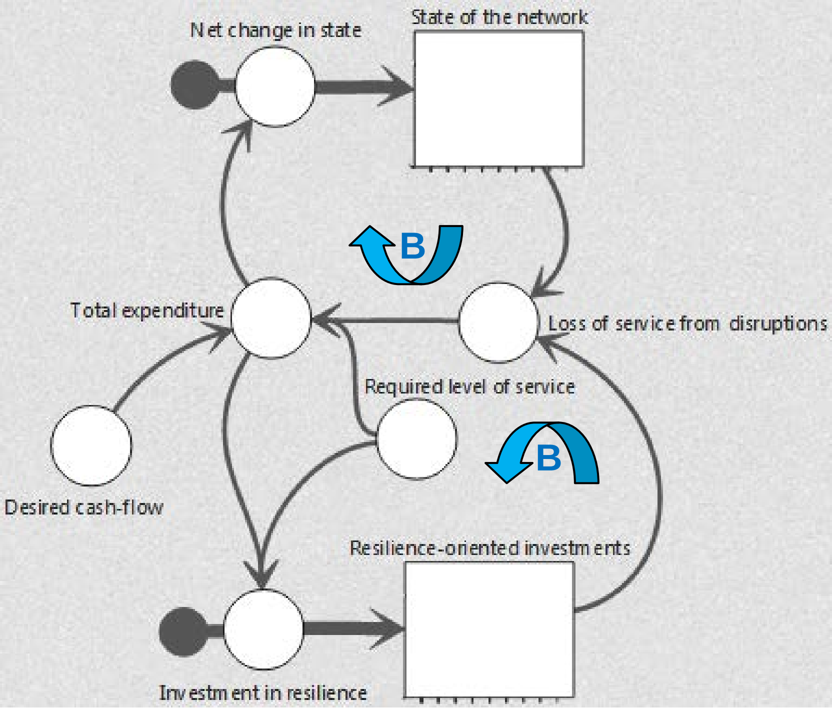

The model is used to test policy options – both long-term choices made from the start that determine the network's initial state, and continuing choices that may slowly improve that state and/or reduce the recovery time after a disruption. The policy options can be tested against three alternative scenarios – a single event of some desired severity, a fixed sequence of disruptions of varying frequency and severity over a longer period, or fully randomised scenarios, which differ on every experiment. The last of these is designed for game-playing, so the results that follow report results only for a single disruptive event and a single fixed sequence of events.

The environment in which asset intensive businesses have operated over the last 20 years or so have generally encouraged them to minimise Capex and Opex, both in order to remain profitable, and to do so against a background of falling unit prices for delivering their service. This raises two questions which the model is suited to answering. First, for differing levels of such spending constraint over many years, just how much loss of service might result from the vulnerability of the network to disruptive events? Secondly, what mix of spending and resource-decisions is best for the future to reduce that consequential trouble most rapidly and affordably?

This situation therefore offers strongly conflicting performance indicators – minimising the loss of service to customers, and sustaining a strong financial cash-flow. The financial harm experienced by such companies as a direct result of service cuts is relatively small, consisting only of the revenue from each customer cut off, for the time they are cut off, plus immediate costs of carrying out repairs. Consequently, most regulatory regimes impose substantial performance penalties on companies (essentially fines), that escalate as the rate of service loss escalates. This has the helpful effect of incorporating the inconvenience experienced by customers into the financial results of the business, and thus producing a single objective function for the model – the cumulative cash-flow of the business over the simulated time-scale. (Strictly, this cash-flow should be discounted to arrive at its net present value, but the relatively short time-scale under investigation makes this unnecessary).

The model's results are assessed against a background of two historic strategies:

1. sustained investment in both Capex and Opex to keep the network in a good state
2. under-investment in Capex and Opex to sustain medium-term cash flows

For each historic strategy, the system is tested against two scenarios of 200 days:

A. A benign external environment, during which only a single, mid-scale disruption occurs (which gives time for problems to be rectified and removes any chance that the system is hit by more disruptions when already weakened)

B. A more challenging environment in which disruptions of varying frequency and severity occur, leading to repeated damage and the danger of overlapping service failures.

<u>The model is constructed to allow different scenarios to be predefined so the scenario required can be selected at run time and the model automatically sets its parameters for the required scenario.</u>

1. <u>Sustained historic investment</u>

Table 3 shows the outcome for service loss and cumulative cash flow over the 200-day period, following a long prior history of sustained Capex and Opex investment. It can be seen that both cash flow and service performance are high when the system is subjected to a single disruption.

Table 3: Cash flow and service loss performance following a long-run strategy of sustained Opex and Capex investment.

	1 event	**Repeated events**		
Long-run policy:		**1a**	**1b**	**1c**
Maintenance spend ($mpa)	10	10	10	10
Normal Capex spend ($mpa)	500	500	300	500
Capex spend (redundancy) ($mpa)	30	30	30	30
Capex spend (resilience) ($mpa)	80	80	80	80
ICT spend (resilience) ($mpa)	3	3	3	3
Number of Skilled Technicians	120	120	120	180
Skilled Engineers	30	30	30	30
Spares cover days	20	20	20	20
Outcomes				
Cumulative Cash Flow ($m)	360	236	100	270
Cumulative Customer Days Lost (000)	151	3,520	9,770	2,030

When subjected to a continuing sequence of disruptions over the 200-day period, service losses escalate sharply, because of both the service loss of early events and the additional problems caused when events hit an already-weakened system. To illustrate the impact of specific policy-components, strategy 1b shows that the problem is seriously worsened by a long-run cut in Capex – the substantial savings in capital expenditure are far outweighed by the financial penalties from much higher service losses. Strategy 1c, in contrast, shows that a significant increase in technician numbers (which is not especially costly compared with Capex changes) nearly halves the extent of service losses and actually improves cash flow, relative to the base case strategy 1a.

2. Constrained historic investment

Table 4 shows the service losses and cash flow for the business after a period of sustained under-investment in both Capex and Opex.

Table 4: Cash flow and service loss performance following a long-run strategy of sustained under-investment

	STRATEGY Run			
	1 event	Repeated events		
Long-run policy:		2a	2b	2c
Maintenance spend ($mpa)	10	10	10	10
Normal Capex spend ($mpa)	350	350	600	350
Capex spend (redundancy) ($mpa)	10	10	10	10
Capex spend (resilience) ($mpa)	50	50	50	50
ICT spend (resilience) ($mpa)	5	5	5	5
Number of Skilled Technicians	90	90	90	150
Skilled Engineers	30	30	30	30
Spares cover days	10	10	10	10
Outcomes				
Cumulative Cash Flow ($m)	457	-1950	-1500	-26
Cumulative Customer Days Lost (000)	394	133,000	104,000	13,100

Unsurprisingly, if just a single disruption occurs, the system generates a high rate of service loss. However, the additional costs arising from this service loss are rather small, so cash flow is substantially improved – the business "got away with" the strategy of lower spending over many years. Furthermore, although the service loss is more than doubled, compared with the sustained investment case, those losses are still relatively modest. The network is serving about 2 million customers, so on average each experiences a loss of service of about 5 hours.

The weakness of this low investment strategy is dramatically exposed, however, if the 200-day period features the same series of frequent and serious distuptions. Strategy 2a shows a very large and continuing loss of service for customers. The financial penalties, plus the cost of recovery, lead to serious negative cash flows for the business.

The last two columns look at two long-run choices that could have been chosen to mitigate these serious outcomes. Strategy 2b is, again, a long-run under-investment strategy, but with higher levels of routine Capex than strategy 1a. The assets are therefore less aged, and fewer of them fail, so the service loss is reduced by about 25% and the cash flow over the period is actually better, in spite of the large increase in Capex.

Strategy 2c, on the other hand, repeats the low-Capex policy and all other under-investments, except that it deploys larger numbers of technicians. Although the network is still quite aged, therefore, the high level of maintenance and repair means it is less vulnerable to failure. Service losses are barely 10% of those in strategy 1a, and seriously negative cash flows are avoided.

<u>Short-term response to historic under-investment</u>

The final question to consider is what to do if the dangers of long-run under-investment are recognised and the business wants to react quickly to minimise its exposure to unanticipate escalation in the number and severity of disruptions. This could arise, for example, either because new management is appointed and quickly appreciates the risk they have taken on, or because regulators come to appreciate the risk and relax companies' expenditure constraints and allow them to spend more to mitigate the danger.

Table 5 shows the impact over the 200-day period of immediate changes, either increasing numbers of technicians, a focused increase in Capex, or both. the higher Capex is not spent on general upgrades to the equipment, but targeted at making vulnerable equipment more resilient and building in redundancy. Note that not all additional technicians can be hired immediately, but increase over an average hiring lead-tiem of 90 days.

Table 5: Service loss and cash flow over a 200-day period from reacting to a historic underinvestment

	STRATEGY 2 + Response			
Short-term policy response:	**2 base**	**↑ people**	**↑ capex**	**↑ both**
Maintenance spend ($mpa)	10	10	10	10
Normal Capex spend ($mpa)	350	350	350	350
Capex spend (redundancy) ($mpa)	10	10	30	30
Capex spend (resilience) ($mpa)	50	50	100	100
ICT spend (resilience) ($mpa)	5	5	5	5
Number of Skilled Technicians	90	150	90	150
Skilled Engineers	30	30	30	30
Spares cover days	10	10	10	10
Outcomes				
Cumulative Cash Flow ($m)	-1950	-631	-937	-357
Cumulative Customer Days Lost (0000)	133000	43000	78800	30000

The immediate increase in people is more effective alone, mostly because it reduces the recovery time after each disruption, and at relatively little cost. The focused Capex response is more costly, and less effective during this 200-day period. However, thereafter, the network would continue to be more resilient to subsequent events, so the longer-term benefit of focuses Capex increases could be considerable.

The combination of higher technician numbers and focused capex is highly effective, even during this limited 200-day period, cutting service lossesto about a quarter of what they would otherwise have been, and improving cash flow substantially, in spite of much higher spending. Furthermore, the service and financial losses are largely concentrated in the early part of the period, before the investment in redundancy and resilience has had a chance to make the system more robust – by the end of the period, service losses are substantially reduced.

Conclusions

Useful conclusions arise from this initial work, both in relation to CNI Capex and Opex policies aimed at improving resilience and recovery from disruptions and regarding the potential value of system dynamics to understand these issues.

It is not surprising to discover that apparently stable networks can be wrecked by an increase in the frequency and severity of disruptive events, but in the absence of such events, it is not clear how great that vulnerability might be. Constrained investment would appear to offer little risk whether that constraint is modest or severe. It is only when a high-disruption is tested that the considerable difference in the scale of risk becomes apparent.

The model also shows that relatively modest investments in redundancy, resilience and ICT can offer substantial protection, if sustained over long periods. Furthermore, simply having enough technicians in place keeps the network in good shape (and so reduces the immediate impact of disruptions) but also helps consdierably when disruptions happen.

If starting from a situation of sustained under-investment, the risks of which may be invisible because the organisation has been lucky to have experienced no serious threats, focused investments into a poor network have disporportionate benefits. And once again, high levels of technical support provide good protection at relatively low cost.

This early demonstration also shows that system dynamics can add value at the policy level, and with simple models, relative to the large and detailed models that have been more widely used to date. Furthermore, the simplicity and compact structure of the model is highly transparent, enabling policy makers to explore unknowns and gain insight much more quickly that large models might allow.

Further Developments

The model, which is still in its early life, has been designed and built in response to the need for organisations operating the CNI to radically improve their risk and resilience policy development capabilities. As has been discussed, this need has arisen because the threat of disruption to the CNI has increased and will continue to increase as the intensity of existing threats becomes greater and the diversity of sources for those threats widens. Initial demonstrations of the model and its capabilities to planners and analysts working on the CNI have received both a positive response that the model is timely and useful, and suggestions to improve its functionality. Some of these, such as raising the probability of secondary events following more quickly after an initial event, have already been incorporated in the model.

Nevertheless the model must still regarded as only a proof of concept and has to overcome two major challenges to prove that it can be deployed as a routine management support tool. The first of these challenges is that system dynamics has not been widely used in previous risk and resilience modelling, either for individual utilities or the CNI as a whole. Therefore there will be some cautious about its use and suspicion of its capabilities. The second challenge is to show that the model outcomes model translate realistically into the specifics of actual cases. Naturally if an organisation is going to change a policy that will have an effect on the lives of millions of people as well as its financial performance, then it has to be sure that the basis for that decision is correct.

Two immediate developments will seek to address these challenges. Currently the model is populated with data typical for an average organisation (a power distribution business). This data was sufficient to meet the purpose of testing the logic and behaviour of the model. Replacing this with actual data from disruptive events experienced by real organisations is necessary to remove any doubts about system dynamics in general, and the accuracy of the model and its outcomes in particular. The second challenge will be to build management training based on the model that gives planners and analysts the opportunity to use it in a controlled environment. This will educate them about the use of system dynamics for modelling risk and resilience, and the capability of such models.

In the longer term there are a number of ways in which the model can be developed. These include:

- Extend the functionality to support more complex risk and resilience scenarios.
- Develop different versions of the model to support other parts of the CNI.

- Integrate the model with asset strategy planning models to create a more end to end policy modeller for infrastructure operators.
- Explore the requirement for a version of the model that, as well as focusing on a single organisation's assets, will consider all the assets of all the organisations that operate in a defined physical location.

We also recognise that, as awareness and use of the model increases and risk and resilience continues to rise up the agenda of both government and CNI operating organisations, other requirements will emerge. It is therefore part of our longer term plans to use the model to actively raise the profile of system dynamics for risk and resilience policy planning across the CNI.

References

[1] Thurlby R. 1999. The Asset Time Bomb Threat or Reality? *Proceedings of the EEI/AGA Utilities for the New Millennium Conference*. Atlanta, USA.

[2] Pederson P, Dudenhoeffer D, Hartley S and Permann M. 2006. *Critical Infrastructure Interdependency Modelling: A survey of US and International Research*. Idaho National Laboratory.

[3] Michelsen R, Brown T. 2006. *The Critical Infrastructure Protection Decision Support System*. Sandia National Laboratories.

[4] Vugrin ED, Warren DE, Ehlen MA and Camphouse RC. 2010. A Framework for Assessing the Resilience of Infrastructure and Economic Systems. In *Sustainable and Resilient Critical Infrastructure Systems: Simulation, Modelling, and Intelligent Engineering.* Gopalakrishnan K and Peeta S (eds). Springer-Verlag.

[5] Guthrie G. 2006. Regulating Infrastructure: The Impact on Risk and Investment, *Journal of Economic Literature*. **44**(4). 925-972.

[6] International Organisation for Standardization. 2009. *ISO 31000:2009 Risk Management – Principles and Guidelines*. http://www.iso.org/iso/catalogue_detail?csnumber=43170 Retrieved 8-Mar-2012.

[7] Lowrance W. 1976. *Of Acceptable Risk: Science and the Determination of Safety.* William Kaufmann Inc. Los Altos: CA.

[8] Boin A, McConnell A. 2007. Preparing for Critical Infrastructure Breakdowns: The Limits of Crisis Management and the Need for Resilience. *Journal of Contingencies and Crisis Management*. **15**(1). 50–59.

[9] Ezell BC. 2007. Infrastructure vulnerability assessment model (I-VAM). *Risk Assessment.* **27**(3). 571-583.

[10] HM Government (UK). 2008. Learning Lessons from the 2007 Floods. Cabinet Office. http://webarchive.nationalarchives.gov.uk/20100807034701/http://archive.cabinetoffice.gov.uk/pittreview/thepittreview/final_report.html. Retrieved 15-Mar-2012

[11] Gaul A, Thurlby R. 2005. Strategic Investment Planning. *Proceedings of the 18th International Conference on Electricity Distribution*. Turin, Italy.

[12] Thurlby R. 2009. Understanding and Reducing Vulnerabilities in Utilities. *Proceedings of the 10th IDERR Conference*, Karlstadt, Sweden.