Tools for Local Critical Infrastructure Protection: Computational Support for Identifying Safety and Security Interdependencies between Local Critical Infrastructures

Chris. W. Johnson and Kevin McLean

Department of Computing Science, University of Glasgow, Scotland, UK, Johnson@dcs.gla.ac.uk

Keywords: safety and security; civil resilience; infrastructures.

Abstract

Previous terrorist attacks, infrastructure failures and natural disasters have revealed the problems that States face in preparing for civil contingencies. One aspect of this is that the agencies which typically coordinate the protection of critical infrastructures have a national responsibility. However, the impact of particular failures is often focused at a local or regional level. For example, Hurricane Katrina was most acutely felt in the City of New Orleans (over 350,000 people affected), with concentrations in suburban Jefferson Parish (175,000) and St. Bernard Parish (53,000) and along the Mississippi Coast (54,000). The terrorist attacks of 2001 and the UK floods of 2007 also show how multiple localised contingencies can occur at the same time. National infrastructure protection agencies must, therefore, be prepared to provide simultaneous help to multiple local agencies. It is for this reason that national civil protection bodies provide national guidance but then devolve responsibility for the implementation of contingency plans to a local level. Unfortunately, many of the regional groups who are responsible for infrastructure protection have little or no idea about the detailed inter-relationships that exist between their own local infrastructures. For example, in the UK 'risk registers' enumerate local hazards without considering how, for example, an attack on a gas storage facility might damage power distribution infrastructures. Nor do they consider the knock-on effects that such damage might have upon water pumping and purification systems. This paper introduces a Geographic Information System that is intended to help identify dependencies between local critical infrastructures. Although we focus on supporting interaction between local and national contingency planning within the United Kingdom, similar problems affect many other nations. The goal is to support the 'joined up' thinking that is often recommended in the aftermath of previous failures.

1. Introduction

The US National Strategy for the Physical Protection of Critical Infrastructures and Key Assets focuses on the following sectors: agriculture and food; water; public health; emergency services; defence industrial base; telecommunications; energy; transportation; banking and finance; chemicals and hazardous materials; postal and shipping [2]. Similarly, the UK Centre for the Protection of National Critical Infrastructures identifies nine sectors which deliver essential services: energy, food, water, transport, telecommunications, government & public services, emergency services, health and finance. Although national bodies, such as the UK CPNI, establish strategies for improving the 'resilience' of these infrastructures, the burden of implementation often falls at a regional level. A wide range of stakeholders must work together to integrate the local tactical and operational response to a wide range of potential hazards. Unfortunately, previous incidents have illustrated significant gaps between national strategies and local implementation.

In the UK, there is a distinction between category 1 and category 2 responders. The former include the Police, Fire and Rescue Services, Emergency Medical Services, the Coastguard, Local Authorities, Primary Care Trusts, Acute Trusts, Foundation Trusts etc. Category 2 responders include Electricity Distributors and Transmitters, Gas Distributors, Water and Sewerage companies, Telephone service providers (fixed and mobile). They also include the transport sector, including Network Rail, Train Operating Companies (passenger and freight), Underground companies, the Highways Agency, Airport operators etc. There have been significant problems in communication and coordination between these different categories of responders. For example, the Knight report into the floods across England in 2007 argued that "the role of category 2 responders in all six phases of integrated emergency management (anticipation, assessment, prevention, preparation, response and recovery management) should be strengthened. The following points should be considered by the Cabinet Office in particular: How to ensure that category 2 responders are properly and consistently represented on Local and Regional Resilience Forums." (recommendation 28, [3]). The following pages argue that Geographical Information Systems (GIS) can be used to promote a more 'joined up' approach to resilience planning for local critical infrastructures.

2. The UK Civil Contingencies Act (2004)

In order to understand the importance of providing tools to promote local infrastructure protection it is necessary to summarize the implementation of the UK Civil Contingencies Act (2004). This legislation provides a local and regional framework for the implementation of civil protection. For example, the Act

requires that local authorities offer infrastructure providers, including businesses, with advice on contingency planning and business continuity management.

The 2004 Act was intended to replace the provisions contained in the Civil Defence Act 1948 and the Emergency Powers Act 1920. These described the ways in which local and public bodies should prepare for external attacks. They also provided for extra powers that could be granted to the government in order to ensure the provision of essential services. However, neither the 1920 nor the 1948 Acts were deemed sufficient for the UK government to coordinate the national response to a range of more recent events including the 2000 fuel protests, the floods of 2000 and outbreaks of Foot and Mouth disease. The 2004 Act creates a framework for the establishment of Local Resilience Forums. This enshrines the *local focus* that was emphasised in the opening sections of this paper. The responses to contingencies, including those that threaten critical infrastructures, are to be coordinated within the existing regional boundaries for local police forces. Responders in these areas are to prepare for contingencies by compiling a Community Risk Register. This provide information about any site that could be involved in a 'major emergency'. Plans must be developed that are proportionate to the risks associated with the entries in the local register.

The 2004 Civil Contingencies Act is structured around three parts. The second deals with the provision of emergency powers. The third describes supplementary legislation to support the implementation of the Act. However, the first part of this legislation explicitly addresses local arrangements for civil protection. The Category 1 responders, enumerated in the introduction, must conduct a risk assessment, develop plans and exercise for emergencies. These drills must include provision for infrastructure maintenance and business continuity. The public and other agencies are to be warned of any hazards. There is also a legal obligation for co-operation and information sharing between category 1 agencies and also with the category 2 responders that include infrastructure providers. The criticisms of the Knight report and the subsequent recommendation that the Cabinet Office must ensure category 2 responders are adequately represented on Local and Regional Resilience Forums are, therefore, particularly important [3].

3. Local Critical Infrastructure Clustering

The Civil Contingencies Act provides the framework that is intended to support the UK response to terrorist actions, infrastructure failures, natural disasters etc. Tools and techniques must be developed to help category 1 and 2 responders fulfil their statutory obligations under this legislation. Previous events have revealed the inadequacies of existing provision. For example, 'the 2007 floods exposed the fact that there is no systematic approach to reduce the vulnerability of the critical local infrastructure' [3]. Similarly, the Pitt Review identified the need for Government to: "establish a systematic, co-ordinated cross-sector campaign to reduce the disruption caused by natural events" [4].

Figure 1 provides a snap-shot of the Local Infrastructure Dependencies Geographical Information System. The LID-GIS is intended to encourage communication and planning between the different groups of responder involved in local infrastructure protection. It illustrates the road layout and principle geographic features of one of the main population centres in Scotland. The user can manipulate the image to show the location of all major items currently held in the Community Risk Register. It also shows the position of these sites in relation to physical features, such as rivers, and transportation links, including roads, railways and underground systems, that run close to many other sites of potential hazard.

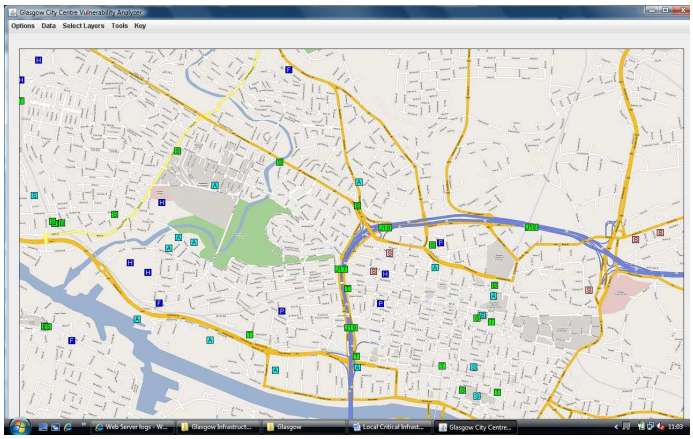

Figure 1: Overview of Critical Infrastructure GIS

The LID-GIS also provides information about the location of resources that can be called upon to mitigate particular risks. These include first responders, such as Fire and Rescue Service personnel, police etc. The system also records the location of other assets that might be required in the aftermath of an adverse event; these include hospitals as well as local authority depots containing heavy equipment. It is important to consider the deployment of these assets because they can be used to mitigate the impact of any contingency. Subsequent sections show how the LID-GIS can be used to explore the 'what if' scenarios that must be considered when these local assets can themselves become the target of terrorist attacks or natural disasters. These scenarios can be based on previous events within the same region or on contingencies in other countries. For example, the system illustrated in Figure 1 has been used to assess what might happen if the UK suffered flooding similar to that experienced by Houston in 2001. As with many NHS facilities, economies of scale had encouraged many healthcare providers from across Texas to group facilities within the same area of the city. This created vulnerabilities that were exposed following a period of extremely heavy rainfall. Highly localised flooding closed

admissions to 3 hospitals and forced the evacuation of more than 2,000 general and 500 ICU beds. The buildings were designed to be 2 feet above the 100-year flood plain and were protected by flood prevention systems that had been developed after a 1976 storm.

The LID-GIS can also be used to consider knock-on effects that propagate between local critical infrastructures and the assets that are intended to protect them. For example, the 2007 floods not only affected the transportation and power infrastructures in many areas of Hull. They also forced the evacuation of more than 200 prisoners and staff from Humberside Police Headquarters. The level of flooding was categorised as a 'once in 150 years occurrence'. However, the loss of critical communication assets and information systems during the response to the floods in Hull has persuaded the Police Authority to invest around £1 million in protecting the headquarters from future floods.

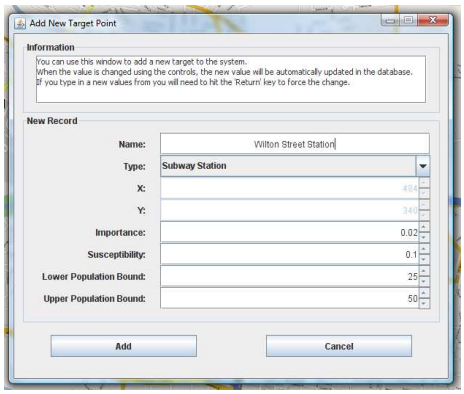

Figure 2: Adding an Infrastructure Component

Users of the LID-GIS can add infrastructure components either by loading them from an existing file or by using a mouse to select a location on the map. Figure 2 illustrates the information that must be provided for each infrastructure component. The user must specify what type of node is being added from a pull-down menu. This is important because the system collates different data depending on the nature of the infrastructure component that is being considered. For instance, if the user selects an electricity substation then information can be entered about the operating characteristics of the transformer. If a hospital is added to the system then data can be provided about the number of in-patients and out-patients that can be accommodated. This data, in turn, is very different from the parameters that are required for components of the transportation, gas or water infrastructures.

Figure 2 shows the user providing information about a subway station. The final two fields are specific to this particular type of infrastructure node. These parameters can be used to specify upper and lower bounds on the numbers of passengers that might be affected by any attack. Monte Carlo techniques can then be used to assess the impact of a potential failure. Future work intends to develop more sophisticated models where passenger distributions vary with time. From this it will be possible to model the ways in which the outcome of any infrastructure failure will vary with the time of day. A further limitation with the approach illustrated in Figure 2 is that the user defines critical infrastructures in terms of particular nodes. It is possible to extend this approach by associating values with the routes or edges between these individual locations on the map. This is necessary to model the flow of people on the trains between individual subway stations. Over 100,000 people were trapped on trains during the Italian blackout of 2003.

Figure 2 also shows how two qualitative values record the 'importance' and susceptibility of each component. Importance can be thought of as a measure of the utility or value of a potential target. Several research programmes have developed specialist techniques for performing these utility calculations, for instance Apostolakis and Lemon's measures for valued worth [5]. A simplified approach has been adopted in the prototype implementation and further work is urgently required to determine whether the additional complexity of these alternative approached can be demonstrated to outweigh more direct subjective, expert judgements. A further justification for our approach is that the subjective measure of utility differs between infrastructures. For example, it can be calibrated to reflect the connectivity of a segment of road or the importance of a bridge as the primary means of crossing from one part of a city to another. Alternatively, the importance value can be used to characterise the power distribution network in terms of different kV line capacity.

The vulnerability field in Figure 2 is intended to capture the susceptibility of an infrastructure component to a potential hazard. The initial focus of our work was on improving urban resilience to terrorist attacks. Hence, susceptibility was formally derived from an analysis of the physical protection that could be provided to a location, for example by access control measures, surveillance cameras etc. The same approach can be extended to represent the vulnerability of infrastructure components to a far wider range of potential threats. For example, many subway systems are susceptible to high levels of rainfall. In other meteorological conditions, these vulnerability assessments might record the vulnerability of electricity distribution networks to ice damage and so on. In each case, the user would reload the same infrastructure elements but they would then edit the vulnerability assessments to those that are appropriate for the scenario being considered. As we shall see, the tool can then be used to identify different knock-on effects between infrastructures that would hold in each different threat scenario.

Community Risk Registers have helped to identify sites of concern within a local area. However, they provide few insights into the interdependencies that exist between regional infrastructures. For instance, they cannot easily be used to consider the ways in which the loss of an electricity substation

might affect water treatment and pumping installations. The development of integrated contingency plans is, therefore, often an ad hoc process. Practice varies considerably between different areas within the UK. Some knock-on consequences are considered in detail for some items in the register while others are hardly mentioned at all. Figure 3 shows how the LID-GIS prototype encourages a more systematic approach to local resilience planning. The system uses the geographical proximity of different infrastructures to identify high-value clusters within the local area. A ranked list is then produced to illustrate areas where the co-location of critical infrastructures combines to create a potential risk which is significantly greater than might otherwise be apparent from their individual entries in the Community Risk Register. As mentioned above, the nature of these clusters will vary as the user enters different vulnerability assessments for particular hazard scenarios. Hence, we would expect that different critical groupings might be identified for terrorist attacks compared with extreme weather scenarios. Even a cursory use of this system in this case study has yielded significant insights, such as the collocation of two major transportation hubs close to a petrol station and two other key infrastructure nodes in a counter terrorism scenario. It has also revealed the vulnerability of local healthcare resources as more and more facilities are centralised in a small number of major hospitals. These units are so dependent on a relatively small number of transportation nodes that access can become a significant problem in case of extreme weather events. Similar problems can be anticipated if relatives and media rush to gain access following an adverse event involving one of these healthcare facilities.

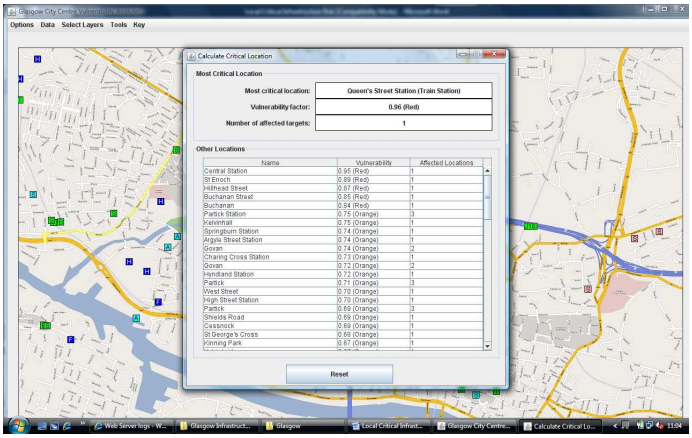

Figure 3: Calculating Co-Located Infrastructures

This proximity analysis can be tailored and calibrated in a number of ways. For instance, users can search across infrastructures to look for significant elements of gas transportation infrastructure within 1-2 kilometres of any site where more than 500 people may be gathered. Similarly, users can search for shopping centres and places of worship with more than 200 people that are more than 10 kilometres from the nearest emergency personnel. This is useful in helping to plan for the future location of Fire and Rescue Service resources. The value systems used to identify high-risk targets can be mapped directly from the individual entries in the Community Risk Registers or using tailored approaches determined by domain specialists. The sensitivity or area of the proximity analysis can also be specified by the user to bring in more and more infrastructures when identifying potential clusters.

The ability to map local critical infrastructures onto a common GIS can help local resilience forums identify critical clusters in the Community Risk Register. However, with computational support it is possible to provide a range of additional facilities. Users do not need to identify the many detailed causal mechanisms that, for example, cause particular Internet routers to fail during a blackout. However, this information can be integrated into the modelling if it is available. In contrast, the intention is that local planners provide subjective probability distributions that can help drive 'what if' scenarios. In other words, they have to estimate how likely it would be that the loss of a particular electricity transformer might affect the local water treatment plant or how likely it would be that the closure of a road might prevent Fire and Rescue service personnel from reaching a hospital. Once these inter-dependencies have been established it is possible to perform more complex calculations, such as identifying the likelihood that the closure of multiple roads would prevent access to a healthcare department or the likelihood that the loss of major elements of the power distribution network might propagate across both transportation systems and communications infrastructures. The central contribution of our work is that increased local resilience cannot simply be calculated in terms of individual systems but must consider the interconnections with other local infrastructures.

4. Further Work

This work is in its infancy and that much remains to be done before the LID-GIS and similar tools can provide adequate support for Local Resilience Forums. We are particularly concerned to identify what might happen in *multiple contingent scenarios*. The electricity distribution industry is only one of many that use 'N-1' criteria. In other words, the system as a whole must be designed and operated in such a way that it is resilient to the loss of any one component. If an element does fails, then steps must be taken to restore N-1 capacity as soon as possible, either by restoring the failed component or by making alternate arrangements. However, things are considerably more complex for local resilience forum who must consider what might happen under N-2 or N-3 where the failed components occur across very different infrastructures. There is little reason to believe that the impact of any single contingency will reduce the likelihood of other potential scenarios. For example, the infrastructure dependencies in our tools might be used to consider the interactions between a terrorist attack on a major

infrastructure component during a period of extreme weather. Alternatively, the tool might be used to analyse the potential knock-on effects of a failure in electricity distribution at the same time as a major structural collapse involving critical components of the rail system. These N-2 scenarios are extremely unlikely to occur. However, many previous studies have repeated the sense of 'surprise' that has characterised the immediate response to infrastructure failures [1, 7]. By preparing for this more extreme class of adverse events, we may be better prepared for the more likely class of infrastructure failures.

Much of our previous work has focused on the development of causal models that can help us understand the ways in which human error, systems failure and managerial decision making combine to create the context in which accidents and incidents are more likely to occur [8, 9, 10]. This work has illustrated the difficulty of identifying the precise ways in which failures will propagate between and within systems of systems. It often takes several years to investigate aviation accidents. This work is very relevant for the analysis of infrastructure failure. In aviation, rail, or maritime accident investigation we only have to identify the causal mechanisms that led to one particular, known outcome. In contrast, the prediction of local and national infrastructure failures is much more complex. We must anticipate the causal mechanisms that could stem from multiple and concurrent adverse events, the precise hazards or threats that generate these failures cannot easily be enumerated.

It is also difficult to predict the different failure modes that are associated with infrastructure components. The LID-GIS system models the probability that particular nodes will fail given that other nodes have failed in associated infrastructures. However, some local systems may continue to provide limited levels of service under contingency. For example, traffic may continue to flow even though part of the road is blocked. Similarly, redundant topologies enable some residual current to flow through a power distribution network even though part of it may be damaged by an adverse event. Our work can, therefore, only be seen as a first step towards the development of more integrated tools for local critical infrastructure protection [7].

Further problems arise because the topology and composition of local infrastructures are continually evolving. The rapid deployment of fibre optic and mobile communications systems, the gradual introduction of Internet based systems in SCADA applications, the development of local and renewable power generation systems are all changing the interdependencies between infrastructures. Similarly, changes in transportation and population patterns affect the numbers of people that might be involved in an incident. Changes in demographics can trigger the redeployment of fire and rescue services, new power supplies, water and transportation infrastructures etc. It can be difficult to ensure that the changes in the underlying models keep pace with the changes in the associated infrastructures; this was a key lesson from the difficulty that reliability organisations had in combating the 2003 US power failure [7]. Hence considerable further work is needed to identify appropriate means for integrating causal information that captures the mechanisms for knock-on failures between infrastructures at a level of abstraction that is both useful for future predictions and can be maintained over time.

The previous paragraphs suggest ways in which the scope of our work might be extended to cover more complex failure mechanisms and multiple contingent scenarios. There are also areas of the existing simulators that might be developed to improve our analysis of relatively simple adverse events. In particular, existing implementations do not capture the temporal properties that can determine the effectiveness of any response to major infrastructure failures. For instance, the provision of Uninterruptible Power Supplies (UPSs) helps to delay the knock-on effects of some problems. Battery power can sustain mobile telecommunications base stations for several hours. Hospitals and other key assets have independent generating capacity, although a key lesson of Hurricane Katrina is that these cannot be relied upon in all potential scenarios. At present, the tools described in this paper do not account for the temporal aspects of infrastructure failure. Further work could consider the impact of UPS' and similar systems through the introduction of more complex stochastic approaches based on Markov chains.

6. Conclusions

Previous terrorist attacks, infrastructure failures and natural disasters have revealed the problems that States face in preparing for national civil contingencies. One aspect of this is that the agencies which typically coordinate the protection of critical infrastructures have a national responsibility while the impact of particular failures is often focused at a local or regional level. National infrastructure protection agencies must, therefore, be prepared to provide simultaneous help to multiple local agencies. It is for this reason that national civil protection bodies provide national guidance but then devolve responsibility for the implementation of contingency plans to a local level. Unfortunately, many of the regional groups who are responsible for infrastructure protection have little or no idea about the detailed inter-relationships that exist between their own local infrastructures. For example, UK Community Risk Registers often simply enumerate local hazards without considering how, for example, an explosion involving a gas storage facility might damage power distribution infrastructure that removes supply from water pumping and purification systems. This paper has introduced a Geographic Information System that is intended to help plan for the knock-on effects that propagate between local infrastructures. The aim is to support the 'joined up' thinking that has been advocated in the aftermath of previous failures.

Acknowledgement

The original idea for the systems described in this paper came from Mike Corcoran, DSTL. All errors of omission and commission are, however, entirely those of the co-authors.

References

[1] Department of Homeland Security, The Federal Response to Hurricane Katrina, Lessons Learned, Washington DC, February 2006. Available on http://www.whitehouse.gov/reports/katrina-lessons-learned.pdf, last accessed March 2008.

[2] National Infrastructure Advisory Council (NIAC) , U.S. Department of Homeland Security, US National Strategy for the Physical Protection of Critical Infrastructures and Key Assets, Washington DC, February 2003. Available on: http://www.whitehouse.gov/pcipb/physical.html last accessed March 2008.

[3] Sir K. Knight, Facing the Challenge: The Chief Fire and Rescue Adviser's review of the operational response by the Fire and Rescue Service to the widespread flooding in England during 2007, Department for Communities and Local Government, London, UK, March 2008.

[4] The Pitt Review learning Lessons from the 2007 Floods (Interim report), Cabinet Office, London, UK, December 2007.

[5] G.E. Apostolakis and D.M. Lemon, A Screening Methodology for the Identification and Ranking of Infrastructure Vulnerabilities due to Terrorism, Risk Analysis, 25:361-376, 2005.

[6] J. Zhuang and V.M. Bier, Balancing Terrorism and Natural Disasters—Defensive Strategy with Endogenous Attacker Effort, Operations Research 55(5): 976-991, 2007.

[7] C.W. Johnson, Public Policy and the Failure of National Infrastructures, International Journal of Emergency Management, (1)4:18-32, 2007.

[8] C.W. Johnson and C.M. Holloway, A Longitudinal Analysis of the Causal Factors in Major Maritime Accidents in the USA and Canada (1996-2006). In F. Redmill and T. Anderson (eds.), The Safety of Systems: Proceedings of the 15th Safety-Critical Systems Symposium, Springer, London UK, 85-104, 2007.

[9] C.W. Johnson and C. Shea, The Contribution of Degraded Modes to Accidents in the US, UK and Australian Rail Industries. In A.G. Boyer and N.J. Gauthier (eds.), Proceedings of the 25th International Systems Safety Conference, Baltimore, USA, International Systems Safety Society, Unionville, VA, USA, 626-636, 2007.

[10] C.W. Johnson and C. Shea, The Contribution of Degraded Modes of Operation as a Cause of Incidents and Accidents in Air Traffic Management. In A.G. Boyer and N.J. Gauthier (eds.), Proceedings of the 25th International Systems Safety Conference, Baltimore, USA, International Systems Safety Society, Unionville, VA, USA, 616-626, 2007.