Chapter 25

VISUALIZING CASCADING FAILURES IN CRITICAL CYBER INFRASTRUCTURES

Jason Kopylec, Anita D'Amico and John Goodall

Abstract This paper explores the relationship between physical and cyber infrastructures, focusing on how threats and disruptions in physical infrastructures can cascade into failures in the cyber infrastructure. It also examines the challenges involved in organizing and managing massive amounts of critical infrastructure data that are geographically and logically disparate. To address these challenges, we have designed Cascade, a system for visualizing the cascading effects of physical infrastructure failures into the cyber infrastructure. Cascade provides situational awareness and shows how threats to physical infrastructures such as power, transportation and communications can affect the networked enterprises comprising the cyber infrastructure. Our approach applies the concept of punctualization from Actor-Network Theory as an organizing principle for disparate infrastructure data. In particular, the approach exposes the critical relationships between physical and cyber infrastructures, and enables infrastructure data to be depicted visually to maximize comprehension during disaster planning and crisis response activities.

Keywords: Cyber infrastructure, infrastructure dependencies, cascading failures, actor-network theory, situational awareness

1. Introduction

Research efforts by the critical infrastructure protection (CIP) community that focus on the cyber infrastructure are primarily directed at vulnerabilities that expose cyber assets to software-based attacks by hackers, viruses, worms and denial-of-service attacks, and the effects of the digital threats on physical infrastructures. Less attention has been directed at the impact of physical infrastructures on the cyber infrastructure [17]. This paper explores how disruptions to physical infrastructures can cascade to the cyber infrastructure. Also, it examines how the cyber infrastructure can be better incorporated into

Kopylec, J., D'Amico, A. and Goodall, J., 2008, in IFIP International Federation for Information Processing, Volume 253, Critical Infrastructure Protection, eds. E. Goetz and S. Shenoi; (Boston: Springer), pp. 351–364.

the larger context of CIP, and presents the design of a software system that integrates information from physical and cyber infrastructures.

The intricate web of dependencies between cyber assets and physical infrastructures enables the cyber assets to function and communicate. From the power grid that provides electricity to roads that deliver workers to the data center, a complex orchestration of services exists to keep an enterprise network up and running [16]. In addition, categories of vulnerabilities that are tied to geographic locations (e.g., earthquake faults and flood plains) must be considered when assessing risk and planning for recovery. The individuals who maintain critical information technology (IT) systems must understand both the internal (cyber) and external (physical) infrastructures on which their assets rely.

We began our efforts by studying how IT disaster planners and crisis responders analyze the effects of other infrastructures on the cyber infrastructure. In addition to reviewing existing technologies and literature, we interviewed CIP experts and IT professionals, evaluating their work practices and the challenges they face. The interviewees were drawn from federal, state and local government agencies as well as from academia and commercial entities. They stressed the need to make time-sensitive decisions based on critical infrastructure data and diverse sensor data, both for proactive disaster planning as well as reactive crisis response. Although they came from a variety of backgrounds and had different responsibilities, these diverse individuals were linked by their shared concern for the planning, protection and recovery of critical cyber infrastructures. Collectively, we refer to this group as “IT crisis managers.”

IT crisis managers are required to protect large-scale enterprise networks from cyber threats (viruses, worms and targeted attacks) as well as from physical threats (hurricanes, floods and acts of terror). Nevertheless, we found that the IT crisis managers focused their efforts almost exclusively on cyber threats, largely ignoring the effects that disruptions to physical infrastructures could have on their systems. They had a poor understanding of the dependencies between infrastructures, which are complex and difficult to comprehend, especially in crisis situations. They did not adequately comprehend the cascading effects that disruptions in other critical infrastructures have on the cyber infrastructure. Moreover, they are deluged with massive volumes of disparate data that must be considered for effective crisis planning and response.

These challenges guided the requirements and use cases involved in our design of Cascade, a software system that visually presents the physical vulnerabilities of an enterprise network and how the vulnerabilities can propagate due to the network’s dependence on other critical infrastructures (e.g., electrical power). The design incorporates the locations of critical computing assets and man-made or natural threats specific to the geographic regions that could affect the enterprise network. The system presents information to IT crisis managers to support rapid vulnerability analysis and course-of-action evaluation when planning responses to potential threats, as well as command and control activities for individuals engaged in crisis management.

2. Related Work

Several researchers have examined how digital attacks can disrupt the cyber infrastructure and how these disruptions cause failures in other critical infrastructures [4, 20]. Others have attempted to provide quantitative metrics for measuring risk associated with digital threats (see, e.g., [5, 10]). In our work, we examine the relationships between cyber and physical infrastructures from the opposite perspective. Instead of investigating how cyber threats affect other critical infrastructures, we focus on how disruptions to physical infrastructures cascade into and interact with the cyber infrastructure.

Infrastructure interdependence is fundamental to the propagation of threats between infrastructures. Therefore, understanding and documenting infrastructure dependencies is an essential step in coordinating disaster planning and emergency response activities [19]. There are two main approaches to understanding these dependencies and their role in infrastructure failure: surveying historical disasters, and modeling and simulating disasters.

Much of what is known about infrastructure failure comes from actual disasters. Identifying the causes and effects of previous failures and the infrastructures involved helps to better plan for the future. Zimmerman [18] has conducted extensive research in this area, surveying a large number of disasters in various infrastructures; her results strongly support the use of infrastructure dependency information in decision-making. Rinaldi and co-workers [15] have developed a foundation for learning from disasters and mapping the results into a framework of interdependent infrastructures.

The development of computer models and simulations for critical infrastructure dependencies is a new and rapidly evolving area of research that has yielded a number of techniques and tools with varying maturity levels. Robinson, *et al.* [16] describe the benefits of simulation-based infrastructure models. Pederson and colleagues [13] have completed an extensive survey of work in the area. Dudenhoeffer, *et al.* [1] have designed a simulation framework (CIMS) for multiple interacting infrastructures. CIMS introduces disaster scenarios on the modeled infrastructures and simulates the effects of infrastructure failures.

Unfortunately, the results of disaster studies and simulations rarely reach IT crisis managers and emergency responders who can benefit from them. Indeed, the individuals we interviewed were unaware of the work and had never used any infrastructure simulation technologies. This is unfortunate because much of this work is directly applicable to the cyber infrastructure, and the results of infrastructure simulations could help disaster planners better understand infrastructure dependencies and vulnerabilities. The Cascade system described in this paper helps translate simulation results into actionable information for IT crisis managers.

3. Linking Infrastructure Data

There are key challenges to linking cyber and physical infrastructures, mainly due to the deluge of data and the unique aspects of cyber data. To overcome these challenges, we propose the process of "induced depunctualization" as an organizing principle for linking cyber and physical infrastructures. We demonstrate how this principle can be used to organize and filter infrastructure data.

3.1 Physical Infrastructure Data Challenges

There is a concerted effort by federal, state and county government agencies to collect data about critical physical infrastructures. Geographic Information Systems (GISs) are often used to provide the robust storage, visualization and analysis solutions that are required. A GIS allows for the use of geographic location as a baseline for bringing data from different infrastructures together. Within these geodatabases, infrastructure information takes the form of map layers, where each layer depicts some aspect of an infrastructure. For example, when storing information about the telecommunications infrastructure, multiple map layers separately show the locations of telephone switching stations, fiber optic lines, telephone poles and cell phone towers. Surprisingly, very few layers are dedicated directly to capturing data about the cyber infrastructure (e.g., locations of government data centers). Without such location information, it is difficult to determine whether a flood, explosion or power outage will damage or impede access to important cyber assets.

The collection of physical infrastructure data can be thought of as a large stack of map layers, growing taller as new layers are added. When historical data is included, the number of layers grows even faster, making it difficult to discern the unfolding of a crisis. It is difficult, if not impossible, to view all of these layers at once; nor can one easily select those most likely to affect the cyber infrastructure. As the information density grows, users are overloaded with data and potentially important data is occluded. This makes it difficult to find the information most relevant to any single infrastructure. For example, an IT crisis manager may have to decide where to place a back-up facility, or determine which data centers are at risk during a hurricane. When presented with hundreds of infrastructure map layers, it is an arduous task to home in on the layers that provide relevant information.

Another problem is that there is no straightforward method for connecting map layers and, therefore, no way to relate different infrastructures. States like New York [12] and Montana [2] have created massive databases of infrastructure map layers and have begun efforts to provide search capabilities for map layers of interest. Still, these systems lack support for associating map layers from different infrastructures.

3.2 Cyber Infrastructure Data Challenges

The cyber infrastructure has certain characteristics that affect its total representation within a GIS: it is geographically dispersed, it incorporates components beyond the IT crisis manager's control, and it is often dynamically reconfigured. Large enterprise networks have mission-critical servers in geographically dispersed locations. These servers may support one organizational mission, yet they are housed in separate locations and may be vulnerable to quite different physical threats (e.g., hurricanes on the Gulf Coast and earthquakes on the Pacific Coast). Displaying such widely dispersed assets within a single GIS display would require a scale that affords little space for details. The other side of this issue is that a single facility may incorporate systems with very different missions. Separate database servers containing medical records and transportation records may be co-located and, therefore, share a common physical vulnerability even though they have no logical relationship. Furthermore, large enterprise networks rely on other entities (e.g., Internet service providers and backbone providers) that are outside the enterprise owners' control; moreover, the locations and status of the assets may be unknown. Finally, large enterprise networks are dynamic. Networks are reconfigured with new hardware, software is updated or replaced, and file content is changed at a frequency that far exceeds any configuration document or disaster plan. Thus, the current state of the system is often partially unknown. Consequently, it is important to allow for frequent display refreshes and to provide the IT crisis manager with information about the age and reliability of network-related data.

Whereas physical infrastructure data is collected and managed as GIS map layers, cyber data is gathered by sensors such as network monitors and intrusion detection systems. This data is collected at different rates from the various sensors and is often stored in multiple formats. Some of these systems can generate huge amounts of data. All this data must be linked to the physical infrastructure to fully understand threats to the cyber infrastructure. But this is very difficult because cyber data is typically not stored in the GIS format of physical data.

3.3 Infrastructure as an Actor-Network

This section describes a methodology for organizing the massive, complex data discussed in the previous section in a way that highlights only the relevant interaction effects between infrastructures. This principle forms the basis for our design and allows IT crisis managers to rapidly home in on the data they require while filtering out irrelevant details. Malone and Crowston's coordination theory [11] supports these requirements, helping address the important and pervasive need to study the dependence between interacting systems. We apply concepts from Actor-Network Theory (ANT) [6, 7] to address these challenges. ANT provides a perspective on how to view and analyze complex systems and interactions with disparate, yet coordinated, parts. In particular, ANT combines processes seamlessly with the objects and interactions that constitute them. Law has used ANT to study disasters [9] and system failures [8].

A key concept in ANT is punctualization [7], where different, interacting parts of a complex system are abstracted and named by their collective emergent behavior [3]. In a punctualized system, the individual parts are hidden. The concept of punctualization can be applied very effectively to the problem of infrastructure protection. For example, an IT crisis manager may view the electrical infrastructure as a single entity whose mission is to provide reliable power. In actuality, it comprises thousands of power lines, generators and transformers, all working together to supply electricity. As long as these components work seamlessly to provide the needed power, they remain concealed.

This process of hiding component parts and only acknowledging the larger whole contributes to the challenge of studying infrastructure interactions and dependencies. Due to punctualization, interactions within and between infrastructures are hidden, so identifying vulnerabilities and threats to these invisible systems is extremely difficult. Perrow [14] defines the complexities of such physical systems, outlining the visible and hidden interactions among them, motivating the question of how to make the hidden interactions visible. Our work also attempts to understand why we cannot see some interactions and what we can do to make them visible.

Returning to our example, when there is a power outage at a critical data center, the IT crisis manager no longer sees the electrical infrastructure as a single entity. Downed power lines, back-up generators, utility companies and repairmen that go unseen during normal operation all become visible, exposing the infrastructure's parts, couplings and dependencies. The hidden elements are rediscovered when an actor-network suffers from disruption or failure.

Although not explicitly described in ANT, but essential to the study of infrastructure dependencies, is that not all the parts are revealed when a failure is introduced into a punctualized system. For example, if a critical data center loses power, only those systems that rely on that power become important. The status of back-up generators and possible failure of critical computer systems become the focus of attention. Data center operations may also rely on other elements (e.g., staff and telecommunications), but they remain hidden during the power infrastructure failure. In fact, a failure causes a partial depunctualization of the system, where the parts that become visible are those that are directly relevant to and affected by the failure; the rest of the punctualized system remains hidden.

Applied to CIP, this partial depunctualization is useful because even though the infrastructure interactions may be too complex to fully understand, the most relevant interactions are exposed. So although all the interactions between complex infrastructures may be difficult to define, it is possible to discern the interactions that are of most interest by studying and simulating failures in these systems. By purposefully inducing or simulating failure into punctualized systems, the relevant facets and connections between infrastructures can be uncovered while keeping the non-relevant portions hidden.

We refer to the process that purposefully deconstructs an entity into its separate, dependent parts as "induced depunctualization." This process can be accomplished through either of the two methods discussed previously: surveying historical disasters or computer simulation. To illustrate the use of induced depunctualization to reveal the cascading effects of other infrastructures on the cyber infrastructure, consider the example of a hurricane hitting a critical data

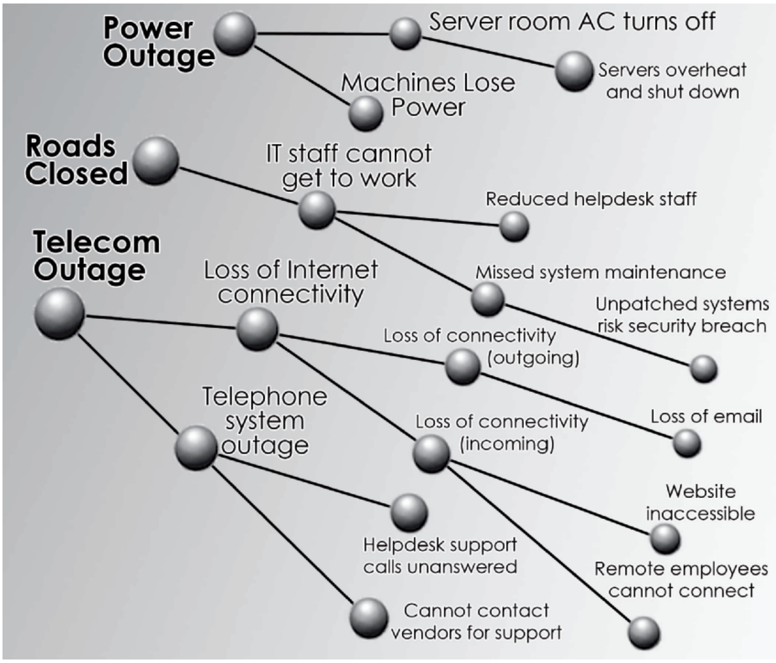
Power Outage
Server room AC turns off
Servers overheat and shut down
Machines Lose Power
Roads Closed
IT staff cannot get to work
Reduced helpdesk staff
Missed system maintenance
Unpatched systems risk security breach
Telecom Outage
Loss of Internet connectivity
Loss of connectivity (outgoing)
Loss of email
Loss of connectivity (incoming)
Website inaccessible
Remote employees cannot connect
Telephone system outage
Helpdesk support calls unanswered
Cannot contact vendors for support

Figure 1. Cascading effects on a data center resulting from a hurricane.

center. Figure 1 shows the cascading failures that can result from disruptions to the electrical power, transportation and telecommunications infrastructures.

The disaster scenario shows how three separate physical infrastructure failures, namely electrical power, transportation and telecommunications, can affect an enterprise computer network. Failures propagate across infrastructures, exposing otherwise hidden portions of the infrastructures. For example, data centers often connect critical servers to back-up power supplies, but do not provide back-up power to the air conditioning units that cool the servers. When a power failure occurs, the air conditioning goes down, causing the servers to overheat and shut down, which reduces the effectiveness of the back-up power. Depunctualization reveals this hidden dependency. Induced depunctualization provides a method for determining the relevant component dependencies and cascading disruptions of a physical infrastructure failure.

3.4 Organizing Infrastructure Data

As discussed earlier, there are massive collections of infrastructure data. As more sensors are added to critical computing networks and other infrastructures, the deluge of incoming data will increase. Missing from these collections is a filtering mechanism or organizing principle that can guide an IT crisis manager to the right information in a timely manner. Induced depunctualization

Table 1. Infrastructure disruptions with associated data sources.

Infrastructure Disruption	Associated Data Source
Power outage occurs	Outage location map Backup generator status sensor UPS status sensor
Server room AC shuts down	Server room temperature sensor
Servers overheat and shut down	Server status sensor
Machines lose power	Network status sensor Router status sensor
Roads are blocked	Snow accumulation map Traffic map
IT staff cannot get to work	IT staff house locations map IT staff route to work map Traffic map
Help desk staff is reduced	Trouble ticket status Help desk on-hold wait time
System maintenance is missed	System maintenance schedule
Unpatched systems are breached	Intrusion detection sensor

analysis is useful because it shows the potential disruptions that could cascade from an infrastructure failure.

Using the cascading effects from an induced depunctualization of the hurricane scenario in Figure 1, each step in the scenario can be paired with infrastructure GIS map layers or network sensor data. Table 1 shows the failures from the hurricane scenario with the associated data sources (map layers or cyber sensors). For example, an electrical outage map from the utility company would show if a data center is in danger of losing power; this can be coupled with the status of back-up power supply and generator sensors to provide better situational awareness. On their own, the individual physical and cyber components do not describe the power outage threat, but in combination they can help define the threat to IT systems.

Table 1 shows that at each possible disruption point, there are map layers or cyber sensors that provide insight about how a network could be, or is being, affected. In addition, the large number of data sources can be organized by pairing them only with the relevant failure entries. Combining the physical and cyber infrastructure data enables IT crisis managers to fully understand the threats to their cyber assets. However, the data can be difficult to comprehend without visual aids. The next section demonstrates how the data can be displayed using the organizing principle of induced punctualization in a manner that assists IT crisis managers in planning for and responding to threats to their cyber assets.

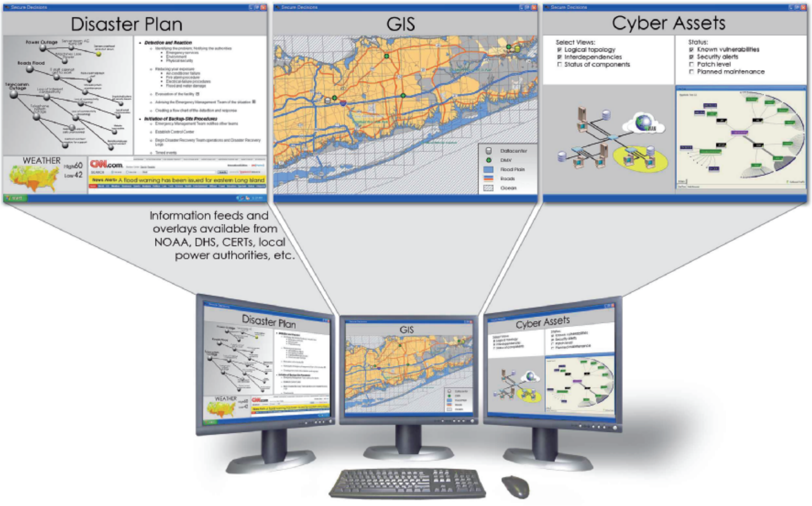

Figure 2. Coordinated views of cascading effects.

4. Visualizing Interdependent Infrastructures

Presenting information about cyber and physical infrastructures to IT crisis managers in an intuitive manner is of paramount importance. This section describes our design for providing this capability. Figure 2 shows the organization of the Cascade user interface. The design provides multiple coordinated views, which present potential infrastructure disruptions and their cascading effects, and support GIS infrastructure map layers and network topology.

Combining physical and cyber infrastructure data within these views enables IT crisis managers to easily determine if a threat or disruption is occurring or may occur. Specifically, the design incorporates: (i) cascading infrastructure failures that show cause-effect relationships of what can go wrong; (ii) disaster plan documents that suggest what to do when failures occur; (iii) infrastructure GIS data that describes the status of physical threats to the network; and (iv) network topology that connects infrastructure data to affected network function.

4.1 Cascading Effects and Disaster Plans

The first view, shown in Figure 3, provides information about what can fail and what to do about it. Presenting the cascading effects of vulnerabilities on network operations illuminates the possible failures. Specific scenarios – such as hurricane, fire or pandemic – can be chosen and displayed. These scenarios can either be hand-crafted or generated from underlying infrastructure dependency simulations.

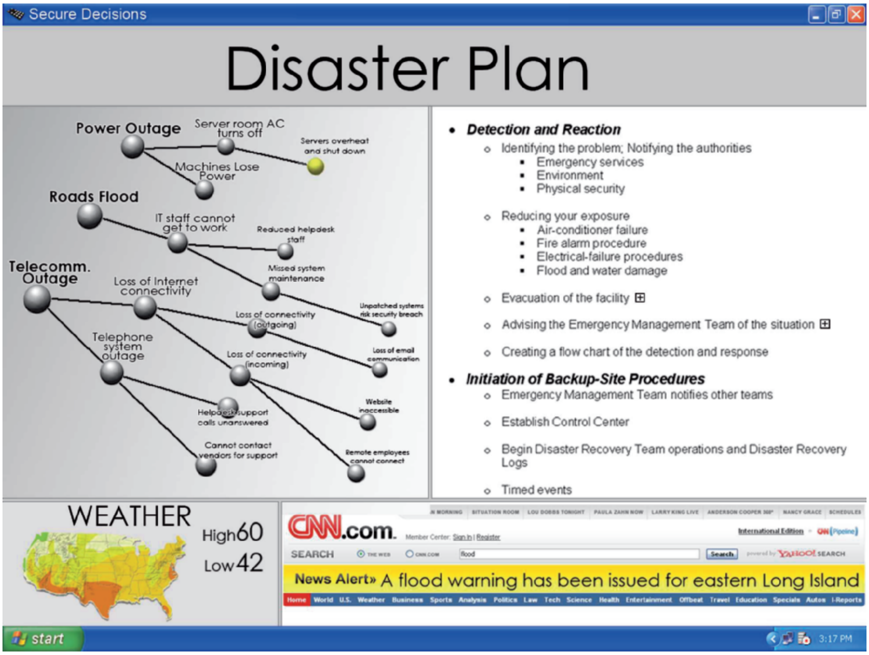

Disaster Plan

- ***Detection and Reaction***
  - Identifying the problem; Notifying the authorities
    - Emergency services
    - Environment
    - Physical security
  - Reducing your exposure
    - Air-conditioner failure
    - Fire alarm procedure
    - Electrical-failure procedures
    - Flood and water damage
  - Evacuation of the facility ⊞
  - Advising the Emergency Management Team of the situation ⊞
  - Creating a flow chart of the detection and response
- ***Initiation of Backup-Site Procedures***
  - Emergency Management Team notifies other teams
  - Establish Control Center
  - Begin Disaster Recovery Team operations and Disaster Recovery Logs
  - Timed events

News Alert» A flood warning has been issued for eastern Long Island

Figure 3. Cascade view of failures and disaster plans.

Including disaster planning documents directly into the interface puts them at the fingertips of IT crisis managers and affords coordination with the other views. For example, a user can link to staff contact lists, news feeds or weather reports. A user who clicks on the node for the failure "Help desk staff is reduced" is directed to the portion of the disaster plan that outlines how to deal with the problem.

4.2 Viewing GIS Infrastructure Data

As discussed earlier, much of the critical infrastructure data is stored in the form of GIS map layers. The second view, shown in Figure 4, incorporates map layers in the presentation. The advantage of map displays lies in the ability to overlay very different kinds of information in the same space, using physical location as the underlying connection. GIS displays and analysis tools have a central role in collecting and using critical infrastructure information.

Cascade leverages GIS technology to present a familiar view of infrastructure data. By coordinating the disaster plan view with the GIS view, failure-to-data associations can be used to organize the map layers that should be viewed. This provides the fundamental mechanism for organizing large catalogs of map layers and implicitly shows the dependencies between infrastructures.

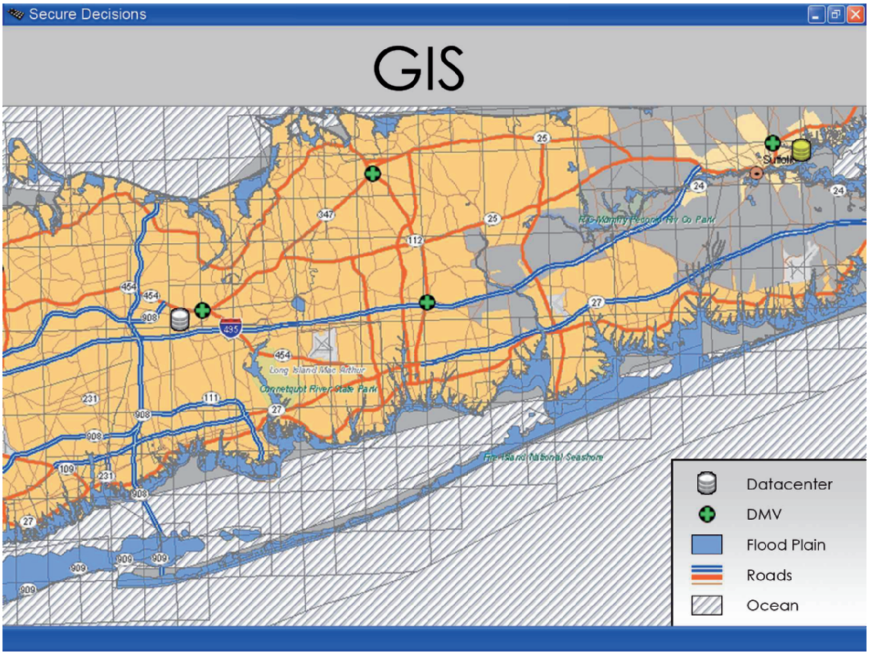

Figure 4. Associated infrastructure map layers.

4.3 Coordinating Physical and Cyber Views

The final view, presented in Figure 5, closes the loop between the physical locations of critical cyber assets and where they function in the network topology by showing the logical layout of the network. This view depicts how workstations, servers and network hardware are organized and connected into logical subnets, showing how connections can be made between machines and to Internet gateways. Additional information may be visually layered on this logical network base view, such as the status of software patches, power availability, temperature and connectivity.

Cascade combines the network topology view into a coordinated application with a disaster planning and infrastructure GIS, allowing interactive exploration of how infrastructure effects cascade to physical and cyber assets, and the network impact of failures. For example, an IT crisis manager in the midst of a hurricane might click on the failure "Server room AC shuts down." This brings up a GIS status map of all the data center's air conditioning systems. Spotting one that has failed in a particular building, he or she clicks on it. The corresponding critical servers in the network topology window light up, showing which servers are at risk of overheating. This intuitive and seamless integration of asset status, infrastructure data and network information provides the IT crisis manager with a comprehensive picture of the impact of failures on the network.

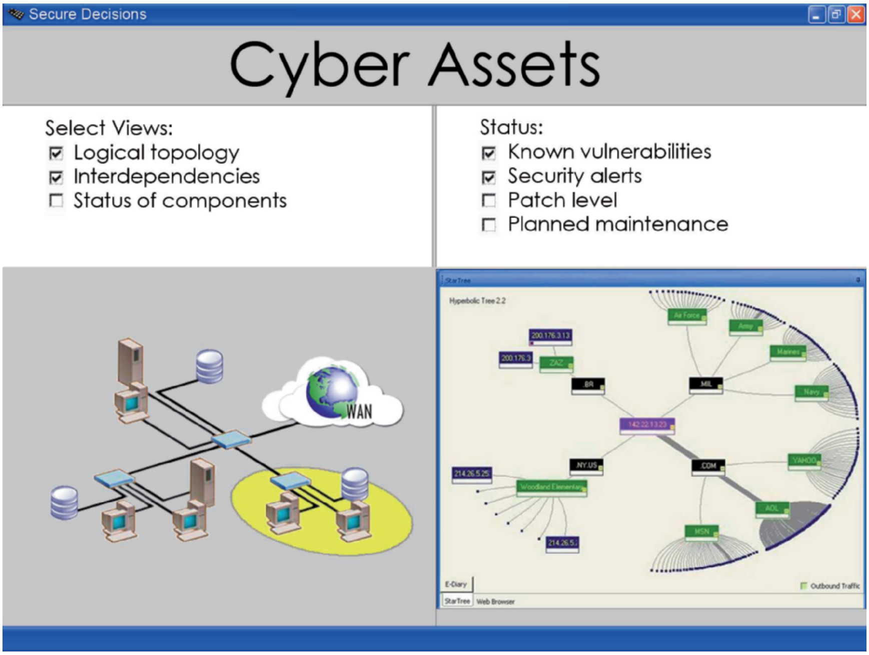

Secure Decisions

Cyber Assets

Select Views:

☑ Logical topology

☑ Interdependencies

☐ Status of components

Status:

☑ Known vulnerabilities

☑ Security alerts

☐ Patch level

☐ Planned maintenance

Figure 5. Network topology and critical cyber assets.

5. Conclusions

IT crisis managers, who must keep mission-critical enterprise networks operating during all types of disasters, need accurate, timely information about how vulnerabilities and failures in other critical infrastructures can cascade to their networks. To address this issue, we have engaged Actor-Network Theory, which provides powerful constructs for organizing diverse critical infrastructure data and deconstructing how the cyber infrastructure can be affected by failures in other critical infrastructures. The resulting Cascade system accommodates massive amounts of infrastructure sensor and GIS data, and provides sophisticated visualization facilities for understanding how failures in physical infrastructures can cascade to cyber assets. Cascade's coordinated geographic and network topological views provide situational awareness about the physical and logical aspects of large-scale enterprise networks. Furthermore, the intuitive, interactive visualization of disaster plans illuminates the cascading effects of infrastructure failures, which is essential to maintaining the stability and survivability of critical cyber assets. The implementation and operational use of tools like Cascade coupled with maturing infrastructure simulation systems and risk management tools will contribute to enhancing the reliability and trust of all critical infrastructures.

Acknowledgements

This research was funded by the Office of Naval Research under Contract N00014-06-M-0146.

References

[1] D. Dudenhoeffer, M. Permann and M. Manic, CIMS: A framework for infrastructure interdependency modeling and analysis, *Proceedings of the Winter Simulation Conference*, pp. 478–485, 2006.

[2] E. Eidswick, Montana spatial data infrastructure: Enhancing an all-hazards approach to emergency preparedness, *Proceedings of the ESRI Homeland Security GIS Summit*, 2006.

[3] J. Goldstein, Emergence as a construct: History and issues, *Emergence*, vol. 1(1), pp. 49–72, 1999.

[4] V. Kumar, J. Srivastava and A. Lazarevic (Eds.), *Managing Cyber Threats: Issues, Approaches and Challenges*, Springer, New York, 2005.

[5] G. Lamm and Y. Haimes, Assessing and managing risks to information assurance: A methodological approach, *Systems Engineering*, vol. 5(4), pp. 286–314, 2002.

[6] B. Latour, *Reassembling the Social: An Introduction to Actor-Network Theory*, Oxford University Press, Oxford, United Kingdom, 2005.

[7] J. Law, Notes on the theory of actor-network: Ordering, strategy and heterogeneity, *Systems Practice and Action Research*, vol. 5(4), pp. 379–393, 1992.

[8] J. Law, Ladbroke Grove, or how to think about failing systems, Technical Report, Lancaster University, Lancaster, United Kingdom, 2000.

[9] J. Law, Disasters, a/symmetries and interferences, Technical Report, Lancaster University, Lancaster, United Kingdom, 2003.

[10] T. Longstaff, C. Chittister, R. Pethia and Y. Haimes, Are we forgetting the risks of information technology? *IEEE Computer*, vol. 33(12), pp. 43–51, 2000.

[11] T. Malone and K. Crowston, The interdisciplinary study of coordination, *ACM Computing Surveys*, vol. 26(1), pp. 87–119, 1991.

[12] New York State, Geographic Information Systems Clearinghouse (www.nysgis.state.ny.us).

[13] P. Pederson, D. Dudenhoeffer, S. Hartley and M. Perman, Critical Infrastructure Interdependency Modeling: A Survey of U.S. and International Research, Technical Report INL/EXT-06-11464, Idaho National Laboratory, Idaho Falls, Idaho, 2006.

[14] C. Perrow, *Normal Accidents: Living with High-Risk Technologies*, Princeton University Press, Princeton, New Jersey, 1999.

[15] S. Rinaldi, J. Peerenboom and T. Kelly, Identifying, understanding and analyzing critical infrastructure interdependencies, *IEEE Control Systems*, vol. 21(6), pp. 11–25, 2001.

[16] C. Robinson, J. Woodward and S. Varnado, Critical infrastructure: Interlinked and vulnerable, *Issues in Science and Technology*, vol. 15(1), pp. 61–67, 1998.

[17] U.S. Department of Homeland Security, National Infrastructure Protection Plan, Washington, DC (www.dhs.gov/nipp), 2006.

[18] R. Zimmerman, Decision-making and the vulnerability of interdependent critical infrastructure, *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics*, vol. 5, pp. 4059–4063, 2004.

[19] R. Zimmerman, Critical infrastructure and interdependency, in *The McGraw-Hill Homeland Security Handbook*, D. Kamien (Ed.), McGraw-Hill, New York, pp. 523–545, 2006.

[20] R. Zimmerman and C. Restrepo, The next step: Quantifying infrastructure interdependencies to improve security, *International Journal of Critical Infrastructures*, vol. 2(2-3), pp. 215–230, 2006.