

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/233202111>

Controlling Cascading Failure: Understanding The Vulnerabilities of Interconnected Infrastructures

Article in *Journal of Urban Technology* · April 2002

DOI: 10.1080/106307302317379855

CITATIONS

233

READS

1,734

1 author:



Richard G. Little

Rensselaer Polytechnic Institute

43 PUBLICATIONS 647 CITATIONS

SEE PROFILE

Some of the authors of this publication are also working on these related projects:



the impact of disruptions in interdependent civil infrastructure during extreme weather events on the delivery of critical commercial services [View project](#)

Controlling Cascading Failure: Understanding the Vulnerabilities of Interconnected Infrastructures*

***Editor's Note:** Although this paper does not explicitly address the infrastructure failures that followed the World Trade Center attacks on September 11, those attacks provide an additional and obvious context for the author's arguments.

Richard G. Little

CIVIL infrastructures are vital public artifacts that support a nation's economy and quality of life. They represent a massive capital investment, and, at the same time, constitute an economic engine of enormous power. Modern economies rely on the ability to move goods, people, and information safely and reliably. Consequently, it is of the utmost importance to government, business, and the public at-large that the flow of services provided by a nation's infrastructure continues unimpeded in the face of a broad range of natural and man-made hazards.

This linkage between systems and services is critical to any discussion of infrastructure. Although it may be the hardware (i.e., the highways, pipes, transmission lines, communication satellites, and network servers) that is the initial focus of discussions of infrastructure, it is actually the services that these systems provide that are of real value to the public. Therefore, high among the concerns in protecting these systems from harm is ensuring the continuity (or at least the rapid restoration) of service.

Causes and Consequences of Infrastructure Failure

The built environment must be designed to resist a formidable array of natural and man-made hazards over its lifetime. In the natural realm, earthquakes, extreme winds, floods, snow and ice, volcanic

Journal of Urban Technology, Volume 9, Number 1, pages 109-123.

Copyright © 2002 by The Society of Urban Technology.

All rights of reproduction in any form reserved.

ISSN: 1063-0732 paper/ISSN: 1466-1853 online



Carfax Publishing
Taylor & Francis Group

activity, landslides, tsunamis, and wildfires all pose some degree of risk to infrastructure systems. To this list of natural hazards, we can add terrorist acts, design faults, excessively prolonged service lives, aging materials, and inadequate maintenance. Although analysis of past events, improved prediction and forecasting methods, and engineering approaches to design and construction have improved the ability of infrastructure systems to withstand natural hazards, crippling failures continue to occur.

Mileti

The consequences of infrastructure failure can range from the benign to the catastrophic. For example, whereas a power outage or water main break may cause only minor annoyance, a street closure due to the formation of a sinkhole may cause major disruption. If the same sinkhole were to cause simultaneous failures in the water and natural gas systems, and resultant fires could not be fought effectively due to inadequate water supply or pressure, possible loss of life and property damage could far exceed expectations from the initial cause. Obvious examples of how a single hazard event can have consequences far beyond the initial damage are the fires that followed the earthquakes in San Francisco, U.S. in 1906 and in Kobe, Japan in 1995. Although hazard mitigation has moved beyond purely life-safety issues, the protection of lifeline infrastructures has generally focused on first-order effects—designing systems to resist the loads imparted by extreme natural events, and more recently, malevolent acts such as sabotage and terrorism. However, as these systems become increasingly complex and interdependent, hazard mitigation must also be concerned with secondary and tertiary effects.

Interdependent Infrastructures

Mitigating damage to infrastructure and ensuring continuity of service is complicated by the interdependent nature of these systems. For example, although the interdependence of many systems is straightforward (e.g., the role played by electric power in providing other services is obvious), the interdependencies of other systems are no less real if not as visible.

Interdependent effects occur when an infrastructure disruption spreads beyond itself to cause appreciable impact on other infrastructures, which in turn cause more effects on still other infrastructures. When an infrastructure system suffers an outage, it is often possible to estimate the impact of that outage on service delivery. These are the “directly dependent effects” of the outage. However, that outage may also diminish the ability of other infrastructures, through no malfunction of their own, to deliver the level of services that they

normally provide. These indirect effects make up a first-order interdependent effect.

The impact of the outage may not stop at these first-order effects. They may go on to adversely affect still other critical infrastructure components, including even the infrastructure that was the original source of the problem, further aggravating the situation. These effects become second-order effects, which can propagate still further, causing yet another round of effects. How far these effects propagate, and how serious they become, depends on how tightly coupled the infrastructure components are, how potent the effects are, and whether or not countermeasures such as redundant capacity are in place. Either the outage effects will die out as they move further away from the base outage, limiting overall damage, or they will gather force in successively stronger waves of cascading effects until part, or all, of the infrastructure network breaks down. In the latter case, losing a key component creates a much broader failure that is out of proportion to the original failure. Given the linkages among infrastructures, a cascading failure could well cross infrastructure boundaries, as demonstrated by the 1998 Galaxy IV satellite failure.

When the PanAmSat Galaxy IV communication satellite rotated out of its orbital position in May 1998, over 80 per cent of the digital pagers in the United States went off-line. Cable and broadcast transmissions were affected, as were credit card authorizations and ATM transactions. This event could have had serious human effects as many hospitals and health care providers in the United States faced a crisis in emergency communications when they could not page doctors and other care givers. This was particularly critical in a health care system that, in the quest for increased efficiency and productivity like much of the economy, relies on just-in-time service delivery.

The Galaxy IV failure was not unique in either cause or consequence. Solar flares play havoc with satellite systems as do spikes in the Van Allen radiation belts. Since 1971, over 4,500 incidents of satellite malfunction have been traced to the natural radiation environment. Other satellite failures have been ascribed to mechanical or other equipment breakdown.

The interdependency problem is further compounded by the extensive linkage of physical infrastructure with information technology systems. Communication and information technologies (ICT) are already affecting infrastructure system design, construction, maintenance, operations, and control, and more change appears inevitable. Potential applications include coupled sensing, monitoring, and management systems, distributed and remote wireless control devices, Internet-based data systems, and multimedia information systems.

Although the coupling of physical infrastructure with information technology promises improved reliability and efficiency at reduced cost, there is surprisingly little known about the behavior of these coupled systems, and thus, their potential for cataclysmic failure is high. Experience has shown that software is fragile by nature, and the software element of control and data acquisition systems is usually the least robust part of an integrated system.


Although recognized as a serious concern, the issue of infrastructure interdependency has received little or no attention. The potential for failures in one infrastructure system to cause disruptions in others that could ultimately cascade to still other systems with unanticipated consequences is very real. In truth, beyond a certain rudimentary level, the linkages between infrastructures, their interdependencies, and possible failure mechanisms are not well understood.

Understanding Interdependency

As a first approach, the multi-ordered implications of infrastructure failure can be generalized using a probabilistic model similar to that developed by Baisuck and Wallace to analyze marine accidents. As depicted in Figure 1, the first stage, or CAUSE, could be a natural hazard such as an earthquake or a technological hazard such as equipment or material failure. This is followed by the INCIDENT, in the examples above, the actual failure of the infrastructure with loss of water pressure and venting of natural gas. Stage 3, the EVENT, would be the resultant fires leading to Stage 4 PHENOMENON with property damage and loss of life.

Baisuck and Wallace

FIGURE 1
A Model for Depicting the Linked Relationships Between Hazards and Their Ultimate Outcomes
 Each stage in the process link is connected to the preceding and



following stages by a probabilistic function based on the frequency of occurrence for any two linked stages. Thus, gas line ruptures in certain soil types (INCIDENT) can be linked to earthquakes of a certain magnitude (CAUSE) by obtaining the frequency with which gas line ruptures occurred as a result of an earthquake. If sufficient data exist, similar probabilistic analyses can be carried through the entire chain of events. Although this type of model can be useful for predicting outcomes when there is much historical data or when frequency relationships can be developed by other means, it is of lesser value when attempting to understand the extreme events that occur at the tails of probability functions.

Closely Coupled Complex Systems

*These occur where the systems involved are sufficiently complex to allow unexpected interactions of failures to occur such that safety systems are defeated, and sufficiently tightly coupled to allow a cascade of increasingly serious failures ending in disaster.

Perrow "The Vulnerability of Complexity"

In his book, *Normal Accidents*, Charles Perrow described numerous failures of tightly coupled, complex systems.* In the search for speed, volume, efficiency, and the ability to operate in hostile environments, he maintains, we have neglected the kind of system designs that provide reliability and security. A particularly troubling characteristic of these tightly coupled, complex systems is that they predictably fail but in unpredictable ways. Similar chains of events do not always produce the same phenomena, but system-level or "normal" accidents of major consequence continuously recur.

Bak and Paczuski

Bak and Paczuski

Bak developed the concept of self-organized criticality to explain how large dynamic systems can self-organize into a highly interactive critical state where even minor perturbations can lead to events, or "avalanches" of all sizes. His work is particularly valuable to the study of interdependent infrastructures and extreme events because the tails of the relevant frequency distributions behave in accordance with power laws that relate the number of events of different sizes by a constant proportion or, in other words, "...large catastrophic events occur as a consequence of the same dynamics that produce small, ordinary events" (6690). On this basis, the catastrophic system failures that Perrow calls normal accidents cannot be dismissed as statistical anomalies—unique intersections of random events—but rather as the expected behavior of closely coupled, complex systems. Taken together, the work of Perrow and Bak supports a discomfiting premise that although it may not be possible to predict the precise nature of the next Chernobyl or Bhopal, a cascading failure of a similar magnitude is destined to occur if we continue to rely on the types of critical-state systems underlying these disasters.

Complex Adaptive Systems

Understanding how complex, interconnected infrastructure systems behave when subjected to the external stresses of natural and technological hazards presents enormous challenges. Managing such systems under these circumstances is even more difficult. This is a world at the edge of stability, where the environment is constantly changing, and systems are continuously adapting to the situation and each other. To provide a framework for understanding and acting on these types of events, Axelrod and Cohen developed a theory of Complex Adaptive Systems. Their premise is that complex systems exist at the edge of chaos, which is disordered and unmanageable. Although their behavior is hard to predict because of the many interacting agents, these systems can be understood, improved, and exploited.

The work of Axelrod and Cohen provides a useful structure for understanding how systems might be designed to lessen the frequency and impact of cascading failures, and Three Mile Island and Chernobyl provide useful case studies. In both cases, it was the intersection of concurrent failures in technology and human performance that was the key factor because neither failure alone would have produced the ultimate disastrous outcome. Perrow believes that such failures are the inevitable consequence of closely-coupled complex systems and, as previously noted, this premise is supported by Bak's self-ordered criticality. There are aspects of Complex Adaptive Systems that can aid in understanding these and similar disasters.

In Complex Adaptive Systems there are many participants, often many kinds of participants, who interact in complicated ways that continuously reshape the future. The three key processes are Variation, Interaction, and Selection. *Variation* in an interactive system, as in a biological community, reduces the vulnerability to single-point failures. The reduced efficiency brought about by independent elements (or evolutionary paths) is balanced by increased robustness of the system. By studying how interactive communities adapt, thrive, or perish, we can learn much about what types of systems are inherently safer in practice. Similarly, *interactions* between members of the same group or social framework, while enhancing communication and simplifying information transfer, can have disastrous consequences when the jointly held information is wrong. At both Three Mile Island and Chernobyl, commonly held views of the situation were uniformly wrong and ultimately contributed to the system breakdowns. Fortunately, in the case of Three Mile Island, an outside agent who had not been influenced by observing the

Perrow *Normal Accidents...*
Chiles

Chiles emerging events, was able to intervene before the system failed totally. Finally, *selection* deals with choosing successful strategies and rejecting those that lead to failure. The key here is learned behavior that will enable participants to survive in a complex, evolving environment. In the absence of actual conditions in which to learn adaptive behavior (such as warfare for the military) there is a need to train the participants by other means, e.g., gaming or simulation. None of the workers at Three Mile Island had been trained to expect anything resembling the types of problems that they actually had to confront. They had no successful patterns or strategies to call upon and were unable to adapt to the rapidly changing conditions.

Other Infrastructure Failures

Disastrous infrastructure failures with similar but subtler links between technology and human performance abound in the literature. The collapse of the Mianus River, Schoharie Creek, and Hatchie River Bridges and the Hyatt Regency Skywalk are illustrative in this regard. The Mianus River Bridge in the State of Connecticut carried Interstate 95. In 1983, a rusted hanger pin and hanger failed and caused a two-lane section of the roadway to fall into the river below, resulting in the loss of three lives. Excessive rust had developed due to paved-over road drains and went unobserved because of poor inspection practices. The Schoharie Creek Bridge, which carried the New York State Thruway, failed in 1987 after a pier was undercut by scour and fell into the creek. The bridge girders slipped off their supports and caused a section of the roadway to fall into the creek, killing ten people. Despite a report almost ten years earlier calling for replacement of missing riprap around the failed pier, the work was deleted from a maintenance contract. In 1989, an 85-foot section of the bridge carrying U.S. Route 51 over the Hatchie River in Tennessee fell into the river after two columns supporting three bridge spans collapsed. Eight people were killed in an accident whose primary causes were a lack of redundancy in design and poor inspection and maintenance practices that failed to detect a developing problem.

NTSB 1984
NTSB 1988
NTSB 1990
Levy and Salvadori

In 1981 a failure occurred that was described at that time as “the worst structural disaster in the United States.” The Skywalk at the Hyatt Regency Hotel in Kansas City, Missouri collapsed, killing 114 people and injuring more than 200. Through an unfortunate and bizarre sequence of events, a design that did not meet the applicable building code was produced by the structural engineer and was subsequently modified and *made weaker* by the contractor. The

contractor's shop drawings were later approved by the structural engineer, and the effects of the change were never noticed (although it was never clear whether they were actually reviewed). The walkway was opened for use despite several instances during construction of the hotel when deficiencies were noted but were not acted upon. Although not on the scale of a Three Mile Island or Chernobyl, what arguably places these four examples in the same context is the recurring intersection of technical faults and human performance failure. The critical role played by the human component of technological systems needs to be far better understood in the context of managing interdependent infrastructures in times of stress or crises.

Petroski *To Engineer is Human...*

Learning from Failure

Some form of structural failure analysis has probably existed since the time of Hammurabi, if not before. Contract disputes over shoddy work or construction failures required that someone conduct an investigation and determine, as best they were able, the cause of failure and who was at fault. Forensic engineering is now a healthy, mature discipline, and much knowledge has been gained, and advances made, from the study of engineering failures. Engineering approaches to hazard-resistant design for structures and lifeline systems have improved continuously from the observation of past failures, assessment of their causes, and improvements in techniques and materials. However, despite the value of forensic engineering to the advancement of engineering practice, the system is far from ideal. Much work of value exists only in court records, sealed by litigation settlements. Nothing analogous to the Air Safety Reporting System (ASRS)* exists for engineering practice although the Near-Miss Project at the Wharton School of the University of Pennsylvania is an attempt to develop a similar reporting framework for other industries. There are also conceptual concerns with commonly used forensic techniques. In its study of errors in the health care industry, *To Err Is Human*, the Institute of Medicine noted that:

Petroski *To Engineer is Human...*
Petroski *Design Paradigms...*

Mileti
NRC 1994

Phimister et al.

*The ASRS is a voluntary program administered by NASA, wherein air safety-related incidents and near accidents can be reported without fear of self-incrimination. The program is credited with facilitating beneficial change throughout the airline industry. (Perrow *Normal Accidents*)

The complex coincidences that cause systems to fail could rarely have been foreseen by the people involved. As a result, they are reviewed only in hindsight; however, knowing the outcome of an event influences how we assess past events. *Hindsight bias* means that things that were not seen or understood at the time of the accident seem obvious in retrospect. Hindsight bias also misleads a reviewer into simplifying the

Institute of Medicine

causes of an accident, highlighting a single element as the cause and overlooking its multiple contributing factors. Given that the information about an accident is spread over many participants, none of whom may have had complete information, hindsight bias makes it easy to arrive at a simple solution or to blame an individual, but difficult to determine what really went wrong (53).

In light of this, care needs to be taken so that “lessons learned” programs (or other forms of adaptive learning for understanding the failure mechanisms of interdependent infrastructures) are designed to capture the influence of all contributing factors, not merely the obvious or easy.

Assessing and Managing Infrastructure Risk

NRC 1996

Risk gives meaning to things, forces, or circumstances that pose danger to people or what they value. Descriptions of risk are typically stated in terms of the likelihood of harm or loss from a hazard and usually include an identification of what is “at risk” and may be harmed or lost; the hazard that may occasion this loss; and a judgment about the likelihood that harm will occur. In the context of physical infrastructure, *risk* connotes the likelihood and level of failure of a critical physical or operational system that would prevent an infrastructure element from fulfilling its primary mission, i.e., providing services. To assess these risks, systemic quantitative risk assessment and management is necessary.

Risk assessment is commonly distinguished from, but is part of, the overall process of risk management. In *risk assessment* for infrastructure systems, the analyst attempts to answer three questions:

- *What can go wrong due to the interdependency and interconnectedness among critical infrastructures?*
- *What is the likelihood that the interdependency and interconnectedness among critical infrastructures will cause major unacceptable consequences?*
- *What might these consequences be?*

Risk management builds on the risk-assessment process by seeking answers to a second set of questions:

- *What can be done to better understand the interdependency and interconnectedness among critical infrastructures and to manage the adverse consequences from a threat?*
- *What organizational, institutional, and research and development options (among others) are available to add more surety and security to interdependent and interconnected critical infrastructures?*
- *What are the impacts of current decisions made on the interdependency and interconnectedness among critical infrastructures on future options?*

Any actions taken to develop and implement comprehensive hazard mitigation strategies for infrastructure must be based on a balanced assessment of all risks confronting the systems and the possible consequences of their failure, either singly or in combination with other, interconnected systems. These strategies must be informed by the best available information and carried out by people knowledgeable about the systems, their possible failure modes, the implications of concurrent system failures, and possible interventions that would allow systems to degrade gracefully and avoid catastrophic, multi-system failure.

Conclusion

Although recent events have focused on malevolent acts and how to prevent them, infrastructure faces other equally serious threats. In addition to natural hazards, the literature demonstrates that excessively prolonged service lives, aging materials, and inadequate maintenance all negatively affect infrastructure. Despite this formidable array of threats confronting our infrastructures, many problems will occur simply due to the complexity of these systems. Potential failure nodes are repeatedly created at the intersections of tightly coupled, highly sophisticated transportation, electric power, and telecommunications systems and are compounded by their reliance on information systems and software. As a first step in protecting these systems, the “vulnerability of complexity” must be resolved.

Beyond generic complexity issues, there are specific emerging threats that are not well understood. For example, commercial satellites are playing an increasingly important role in earth observation, communication, and geospatial positioning—activities that are

central to the control of many key civilian and military systems. This orbital infrastructure is vulnerable to natural events such as solar flares and radiation spikes as well as to man-made threats such as electromagnetic pulses. Its ground-based elements are vulnerable to physical threats and terrestrial natural hazards. Although the hazard community knows how to identify and assess these vulnerabilities, it must also understand that vulnerability assessments represent only part of a total systems solution.

Infrastructure protection is not seen as a purely developmental problem but one in which basic research is necessary and, to date, insufficient. Research needs range from a better understanding of networks and interconnections, to the impacts of deregulation, privatization, and globalization, to better software and system designs. Some of this needed work is underway, but there is still much to be done. A valuable first step would be a comprehensive review and assessment of ongoing research with the goal of identifying gaps in the knowledge base and establishing research priorities.

Opportunities for Collaboration

The issues outlined in this paper suggest a need for collaboration between the social and physical sciences and engineering. Some approaches may be straightforward such as those that call for reinstating “shock absorbers and circuit breakers” in both a physical and operational sense to increase the resilience and reliability of infrastructure systems. Others will be more esoteric and call for the application of sophisticated analytical, modeling, and forecasting tools to improve understanding of the systems and the modes and consequences of failure. There are many potential topics for research and they include such areas as:

Theoretical Foundations. Research into the complex and adaptive behaviors of infrastructures and the overall behavior and functioning of economies from an interdependent perspective is key if we are to understand how infrastructures will behave in the face of failure from a variety of causes—from physical or cyber attack, to a major earthquake, to failure of the network or its components.

Modeling and Simulation. Modeling and simulation of interconnected complex infrastructures is rudimentary today. More advanced models, using actual regional or national infrastructure data, network layouts, and operating conditions are needed to uncover critical nodes, behaviors, and vulnerabilities.

Mitigation, Response, and Recovery. In the event of a major infrastructure failure, isolating the affected portions of the system and preventing cascading failure will be important. Any mitigation actions will require accurate accounting of linkages among the infrastructures and the behaviors arising from such interdependencies. Appropriate and safe steps must also be identified for bringing the systems back on line.

Policy Research. Policies affecting one infrastructure may have unintended consequences in others, due to the linkages involved. Little is known of how this happens and how to reduce the likelihood of its occurring. Likewise, in some cases appropriate policy decisions can probably forestall the need to make costly infrastructure expenditures.

The Human/Technological Interface. Human error has played a major role in some of the most significant technological disasters of the past century. A better understanding of how systems can be designed to take human factors into account, as well as decision tools that enable people to structure rational choices for technological interaction, is needed.

A Closing Caution

In *Betrayal of Trust*, Laurie Garrett paints a grim picture of how, in the twentieth century, the public health infrastructure in the United States deteriorated from a formidable first-line defense against infectious disease to a struggling, under-funded, and under-appreciated appendage. Today's concerns with bio-terrorism have the public and policy makers alike wondering if the United States is capable of dealing with deliberately induced outbreaks of infectious disease. However, terrorism may not be the real threat. The global economy and the worldwide air transportation network have created a closely coupled system that makes it possible, and even likely, that people infected with highly contagious diseases unwittingly will spread the infection far beyond national borders. In the absence of a global public health infrastructure, the potential consequences are grim. As Garrett points out:

High-tech solutions, devices to "sniff out" nasty microbes in the air or detect them in the water supply are a technological solution to a public health threat. Were a biological attack to occur, or a naturally arising epidemic, the public would have

only one viable direction in which to place its trust: with its local, national, and global public health infrastructure. If such an interlaced system did not exist at a time of grave need, it would constitute an egregious betrayal of trust (585).

Hopefully, no bio-disasters will come to pass. But those concerned with physical infrastructure should take careful note of the warning implied. Our basic systems are at risk from threats we may not yet foresee. We need to anticipate these threats to our physical infrastructures, design systems that are inherently safer and more robust, and be prepared to restore them when they fail. In this regard, we should take counsel from this historical anecdote:

In 1346 a particular set of circumstances occurred, in a particular sequence, resulting in what may have been the first truly global epidemic. Perhaps only the Americas and Antarctica were spared humanity's globalized Black Death. With epidemics, timing is everything (545).

Garrett

Bibliography

R. Axelrod, and M.D. Cohen, *Harnessing Complexity: Organizational Implications of a Scientific Frontier* (New York, NY: Basic Books, 2000).

A. Baisuck and W.A. Wallace, "A Framework for Analyzing Marine Accidents," *Marine Technology Society Journal* 13:5 (1979) 8-14.

P. Bak and M. Paczuski, "Complexity, Contingency, and Criticality." *Proceedings of the National Academy of Sciences*, 92:6689-6696 (Washington, D.C.: National Academy of Sciences, 1995).

J.R. Chiles, *Inviting Disaster: Lessons from the Edge of Technology* (New York, NY: HarperCollins Publishers, 2001).

L. Garrett, *Betrayal of Trust: The Collapse of Global Public Health* (New York, NY: Hyperion Books, 2000).

Institute of Medicine, *To Err Is Human: Building A Safer Health System* (Washington, D.C.: National Academy Press, 2000).

M. Levy and M. Salvadori, *Why Buildings Fall Down* (New York, NY: W.W. Norton & Company, 1992).

D.S. Mileti, *Disaster by Design: A Reassessment of Natural Hazards in the United States* (Washington, D.C.: Joseph Henry Press, 1999).

National Research Council (NRC), *Practical Lessons From the Loma Prieta Earthquake* (Washington, D.C.: National Academy Press, 1994).

National Research Council (NRC), *Understanding Risk: Informing Decisions in a Democratic Society* (Washington, D.C.: National Academy Press, 1996).

National Transportation Safety Board (NTSB), *Collapse of New York Thruway (I-90) Bridge, Schoharie Creek, near Amsterdam, New York*, HAR-88/02 (Washington, D.C.: National Transportation Safety Board, 1988).

National Transportation Safety Board (NTSB), *Collapse of the Northbound U.S. Route 51 Bridge Spans over the Hatchie River near Covington, Tennessee*, HAR-90/01 (Washington, D.C.: National Transportation Safety Board, 1990).

National Transportation Safety Board (NTSB), *Collapse of a Suspended Span of Route 95 Highway Bridge over the Mianus River, Greenwich, Connecticut*, HAR-84/03 (Washington, D.C.: National Transportation Safety Board, 1984).

C. Perrow, *Normal Accidents: Living with High-Risk Technologies* (Princeton, NJ: Princeton University Press, 1999).

C. Perrow, "The Vulnerability of Complexity," paper presented at the planning meeting on the "Role of the National Academies in Reducing the Vulnerabilities of Critical Infrastructures," National Academy of Sciences, Washington, D.C. (April 28-29, 1999).

H. Petroski, *Design Paradigms: Case Histories of Error and Judgment in Engineering* (Cambridge, U.K.: Cambridge University Press, 1994).

H. Petroski, *To Engineer Is Human: The Role of Failure in Successful Design* (New York, NY: Vintage Books, 1992).

J.R. Phimister, U. Oktem, P.R. Kleindorfer, and H. Kunreuther, "Near-Miss System Analysis: Phase I" working paper of the Near-Miss Project, Wharton School, Center for Risk Management and Decision Processes, *online* <<http://grace.wharton.upenn.edu/risk/wp/wplist00.html>> (Accessed January 24, 2002).