CRITICAL INFORMATION INFRASTRUCTURE PROTECTION (CIIP)

Peter WESTRIN

1. Introduction

One of the remarkable features of modern, computer-based society is that so many things must *work right*. Seemingly endless small details must function correctly and in co-operation in order to maintain processes, which we take for granted. A single "bug," the smallest aberration, so subtle as to be virtually impossible to foresee, can initiate a complex chain of events, the effects of which can manifest themselves at a national or global level.

An example of a small cause, which can lead to a large effect, is the case of the digital group selector in an electronic telephone exchange. A single erroneous binary digit in a particular shift register can instantaneously break off ongoing telephone conversations and recouple them randomly. Thousands of callers can suddenly be directing their conversations to complete strangers. The distance between a digital parity error and its social consequences may be incredibly short.

That such embarrassing situations seldom occur is because of the well-specified nature of the telephone system, its construction and rigorous built-in controls, and its careful testing before mass use. This is the case, generally, for most commercial computerized products, if we disregard computer games and related programs. If these latter systems do not function sufficiently well or otherwise lack reliability, they will soon fall to the competition. System programs inhabiting our PC's are examples of products, which are considered to function "well enough" in order to be acceptable to the average user. One reason for this is that the consequences of program failure are usually tolerable for the user.

Operating and monitoring systems for nuclear power plants, or transaction systems in the world of banking, are quite another matter and require high reliability. Even here, though, errors occur. Recently, the payment system of a major Swedish bank broke

INFORMATION & SECURITY. An International Journal, Vol. 7, 2001, 67-79.

down repeatedly during a two-week period, causing considerable trouble for millions of customers nation-wide. The bank in question reported that the reason for the disruption was the "human factor." No further details have hitherto been released.

Was this too the case of a minute detail causing an entire system to collapse? Any given system can, per se, function sufficiently well and perform reliably after being tested and "run-in"—which, of course, can take its good time. It is therefore understandable that large, complex systems, which cannot be tested fully by way of simulation, often have (seemingly endless) running-in problems, in which unexpected "features" arising out of millions of minute details can lead to high-level system consequences. This is something that we will have to learn to live with. For even as our knowledge and competence in regard to system reliability increases, new demands of functionality will likewise increase, and thereby even system complexity.

However, despite the fact that breakdowns in banking and payment systems can have nation-wide consequences, or that running-in disorders in a subway system can affect millions (as was the case in Stockholm last year), such disruptions are, in substance, *local occurrences*. That is, the disruptions are contained within a given, restricted system. There is a certain delimited, more or less well defined function or service, which is affected, and there are usually more or less acceptable reserve procedures or backup-functions. In short, there are ways to get around such problems, and one can hardly maintain that they constitute a serious threat to society, let alone threaten society's very existence.

And with this in mind, we seem to have identified something of a paradox as concerns our perceptions of modern, "high-tech" society: namely, its apparent *robustness*. Certainly there are disruptions—e.g. in traffic systems, electricity distribution and banking transactions. And accidents do happen—dams burst, airplanes crash, trains collide and ships sink. But on the whole, and in light of the sheer amount of activity at hand, our modern, technology-based society would seem to function exceedingly well.

Modern technology has been developed and exploited to the affect of creating both a safer and more comfortable society. Crisis management becomes more effective when technology creates increased redundancy and flexibility. Margins of safety, buffering us from catastrophes such as floods, famine, earthquakes and epidemics, have become wider in those areas of the world where modern technology has been most widely applied to societal development. The disruptions we do experience are most often local, the consequences of which are understood and relatively limited, and with known procedures of mitigation.

Once in a while however disruptions occur which we can designate as constituting major disturbance for an entire nation or region. The power failure in Auckland, New

Zealand, and the so-called ice storm in Canada in 1998 are examples of (relatively) catastrophic disruptions at the urban and regional levels, respectively. In the former event, an urban center's commercial activity was paralyzed by protracted power shortages caused by repeated power cable failures. The disruptions had relatively far-reaching economic and demographic consequences for an entire urban area.

The ice storm in Canada, in which a whole region went without electric power during severe weather conditions, involved an even greater population than in Auckland, and required rescue operations on a wartime scale in order to keep the situation under control.

Both cases involve infrastructure failure. In the case of Auckland, the cause of the disruption involved inadequate infrastructure maintenance, whereas in Canada it concerned "forces of nature" for which the infrastructure—in this case the electricity distribution system—simply was not designed to weather.

At this point, and in the context of Information Technology and Critical Information Infrastructure, the question arises: Are we evolving towards an ever more robust society, or are we heading towards a situation where the risk of *a really major, society-threatening chain reaction of IT-related events is increasing*?

2. Societal Infrastructure

All of the disruptions hitherto referred to have involved societal infrastructure systems. Although the concept of *societal infrastructure* can be defined in a number of different ways, for the present application we find it most appropriately defined as: The totality of publicly utilized functions and services which constitute the conditions for the maintenance of social and productive relationships, as well as the framework for further societal development.

Certain forms of infrastructure, or infrastructure sectors, are of special importance for modern society. These so-called *critical infrastructures*, which are also critically interrelated and interdependent, include electricity production and distribution, transport, telecommunications and water supplies. Emergency services and government or administrative services can also be included. If any of these *infrastructures* ceases to function for a prolonged period, society will be hard pressed to maintain its functioning as a whole.

With current and future rapid developments in society's dependence on IT, this list of critical infrastructures will have to be extended to other sectors. And as this very fact attests to, one of these critical infrastructures distinguishes itself from the others: data-communication and its associated computers (in the wide sense of the word) and (world-wide) networks.

The *information infrastructure* is the term usually used to describe the totality of such interconnected computers and networks, and the essential information flowing through them. The distinguishing characteristic of the information infrastructure is that it is all embracing—it links other infrastructure systems together. Take away the information infrastructure and many other critical infrastructure systems will shut down relatively quickly.

Electricity supply is in many ways as all embracing as the information infrastructure. However, one can compensate for power failures by means of reserve generators placed at strategic locations. In many cases, this is not possible with the loss of critical information flows.

There are certain parts of the information infrastructure, which are especially critical. These are the data networks which monitor and control important societal function and services. These include electricity distribution, telecommunications, banking services, rail and air traffic control and emergency management systems, as well as stock exchange and securities management. Presently, many of these systems are relatively isolated and thus (relatively) secure from intrusion. However, with the accelerated pace of development within the IT-sector it will be all the more difficult for collective systems to isolate themselves from the outside world, and to maintain the boundaries between "inside" and "outside."

2.1 The Network Society

What we call the *Internet* is the top of an iceberg, which is currently in the process of changing society the world over. While creating vast new opportunities it is creating, and will continue to create, new risks and threats that will be difficult to anticipate.

The Internet is primarily employed as a means for transferring information between people. There are also dedicated networks for monitoring and controlling all types of technical systems and computerized processes. In such networks, data flows directly into control systems and affects their physical functions. Technically, there is no obstacle to using the Internet even for such purposes. And in this event, local, dedicated networks will become integrated into the whole of the Internet.

It is in no way unthinkable that, within the not so distant future, every person on Earth can, in principle, reach and influence every other person, as well as a good portion of society's collective technical infrastructure. If, added to this, the mutual interaction between such systems and networks continues to increase at the present, or even accelerated rate, then we are going to be faced with an extremely complex system of problems which will have bearing on the function and stability of the world system as a whole.

Where does one place responsibility for the maintenance of critical societal functions if these become mutually dependent and complex to the extent that there is no longer any way to understand how such a complex will behave, or how to exercise control over it.

In his book on "normal accidents," Perrow[1] argues that in an interactively complex system two or more discrete failures can interact in unexpected ways, thereby affecting supposedly redundant sub-systems. A sufficiently complex system can in fact be *expected* to have many such unanticipated failure mode interactions, making it vulnerable to inevitable accidents.

2.2 The Threat to the Information Infrastructure

Modern society's infrastructure has always been, and still is, vulnerable to physical threat. Severe weather conditions, earthquakes, floods and sabotage are examples. Threats of these types can be categorized and analyzed, and given their own special defensive or mitigating strategies. They can be made intelligible, their consequences described, and they originate, in a seeming well-defined way, from the "outside."

The threats that we may face as concerns the information infrastructure are of another kind. They are not well-defined or specified beforehand, we cannot take in their potential consequences and in the developing, all embracing network society these threats may be seen as originating form the "inside."

As concerns sabotage, the information infrastructure can be employed as a means to bring about the disruption of critical infrastructure—including the information infrastructure itself. Information can be stolen or manipulated. Computers can be infected with malicious programs, which can disrupt not only software and immediate associated hardware, but also adjoining or bordering technical systems—as well as trust and confidence in society as a whole.

The network society bears within itself the seeds of a crisis of confidence, as the individual member of that society finds it more and more difficult to gain an overall understanding of the social and technical environment, or to identify responsibility for its maintenance.

2.3 Critical Nodes and Links

An important question arises: will the IT-based network society become increasingly unstable on the basis of its increased complexity alone and, if so, how will these instabilities express themselves?

The network society is characterized by a system of integrated networks consisting of nodes and links. How can we identify those nodes, which are "critical" for the network society itself. One way is to designate a node as critical if either:

- It alone can exert such influence on other nodes that a serious disruption of societal infrastructure can occur; or
- It forms an integral part of an ensemble of nodes, which can be attacked or otherwise influenced in a similar manner, such that the aggregate malfunction can lead to serious disruptions.

An examination of all likely nodes in order to estimate criticality would however be exceedingly time-consuming and only give results with an early expiration-date, both because of the rapid rate of development within the IT-sector and the fact that such an examination would involve a myriad of details in system construction and implementation. On the other hand, it may be the case that no single node can ever be disqualified as being non-critical!

An example of the fact that many similar nodes can be critical at the same time comes from the collapse of AT&T´s long-distance telephone switching system in January 1990. Because of a "bug" in an updated portion of a systems program, put into operation on 80 of AT&T's switching systems nation-wide in the US, a chain reaction of shutdowns occurred. The culprit was a specific piece of status information exchanged between stations.

In this case, no single node was "more critical" than any other node. The defect was in the system as a whole. A penetrating account of the course of events and its underlying causes is given by Bruce Sterling.[2]

2.4 Complexity and Vulnerability

Two seemingly conflicting forces are at work in the network society. On the one hand, new means of communication make accessible to the average citizen an almost unimaginable array of new sources of information and services—as well as the prospect of becoming an active party in countless new collectivities and processes. At the same time, the increasing supply of information and the escalating technical complexity of the network society make it all the more difficult to identify potential malfunctions, find their sources and treat them in an adequate manner. The *consequences* of such potential malfunctions—above all their indirect or wider sociological effects—are becoming increasingly difficult to foresee.

The concept of the network is thus central to all discussions of society's intrinsic vulnerabilities. The combination of an exponentially increasing number of human-computer and computer-computer transactions, and the coupling of communication

networks on a global scale open up new possibilities for faulty instructions or malicious code—in whatever form—to spread globally.

In the case of the above mentioned digital telephone switches, in which a single binary digit could create such bizarre effects, it is relatively easy for systems engineers to determine the consequences beforehand. For sufficiently complex systems however it is virtually impossible to anticipate all the potential consequences of errors occurring at the micro-level. Many such errors may be controllable. Some will emerge at the highest system level and give rise to local disruptions. Will some slip out of the local system and propagate unrestrained on a global level?

In two interesting articles, Pastor-Satorras and Vespignani [3] and Duncan Watts[4] address two aspects of error propagation in networks. Pastor-Satorras and Vespignani analyze cases of network virus infections on the Internet and examine their average lifetimes and persistence. On this basis, they then describe a dynamic model for virus propagation in "scale-free" networks and discover that they cannot find any epidemic threshold or associated critical behavior involved in such propagation. If this model in fact captures the essence of the dynamics of such a propagation process, then there is no "virus"—no matter how poorly constructed—which cannot propagate on the Internet.

Watts brings up another aspect of error propagation, namely, how small, local disruptions (chocks) can—in singular cases—trigger widespread cascades in a network consisting of interacting agents. A possible explanation for such processes is described in a model in which each agent's decisions are dependent upon its nearest neighbor's actions—in accordance with a simple threshold rule. Watts investigates the conditions for such a cascade and why it is difficult to anticipate. The model covers a wide range of cascading phenomena, including cultural fads, innovations and social movements, as well as error propagation in infrastructure networks.

3. How Do We Assure the Information Infrastructure?

Since the putative new societal risks and vulnerabilities are directly or indirectly related to the development and utilization of new technologies, it would therefore seem natural to follow a chain of analysis beginning with technical specifications and casually running "up" through systems, actors, threats, vulnerabilities, consequences and, finally, counter measures/ mitigation.

However, in view of the rapid technological developments constantly taking place, and the particular nature of their implementations, one can raise certain objections to such a synthetic scheme. If, for instance, one carefully examines a relatively localized subsystem from the point of view of risks and threats, thereby identifying certain of

its vulnerabilities, in what way can these insights be generalized and established in order to utilize them "beyond" the subsystem itself, on a higher system level?

One might hope that certain "typical" system components or operations might be found in many subsystems, but in order to identify these one would need to have access to a good number of such systems for comparative studies. This however would be extremely time-consuming, and the rapid development of new systems and networks would quickly render such comparisons obsolete.

Furthermore, it is highly unlikely that detailed access to more than a few such systems will be available to research directed towards this end. Systems for such services as finance and security exchange, or data communication in general, will most probably remain inaccessible for analysis.

What would be required is a filtering mechanism by which the technological background noise could be eliminated to the benefit of those more enduring, central factors—which need not at all be "technical" in nature. If such a selection process is impossible to devise—perhaps because no single bit of information can, in advance, be characterized as irrelevant—then we will need to gain insights into the problem complex by working with its different levels of causal action *in parallel*, and attempt to put each of these into mutual context.

It may very well be that critical vulnerabilities, and even the worst consequences of infrastructure disruptions, will not be traceable in any useful way to single technical subsystems—perhaps as a consequence of an already overwhelming system complexity. Perhaps the analysis of vulnerability should be based instead on *functional units*, whose interactions with each other and with the environment as a whole can best be described by way of their societal manifestations as a whole, with less emphasis placed on the technical.

To the extent that this is the case, one of the most important problems for CIIP research is to identify relevant *functional units* and to describe their mutual relations. This perspective also implies that it will be difficult to differentiate between "insiders" and "outsiders"—in some sense we will all be insiders.

3.1 Unforeseeable Consequences of Disruptions in the Information Infrastructure

When we talk about the consequences of disruptions to the infrastructure, we usually think about the more established, direct effects, quantifiable in the form of injuries to people, damage to the (built up and natural) environment, and—of course—in terms of dollars and cents. Other, more indirect and/or non-quantifiable manifestations can, in fact, create the really dangerous consequences for society. One of the conditions of

a secure society is a measure of basic trust among the citizenry for the mechanisms, which govern it—i.e. that one has confidence in its inherent stability.

At some point, there will be a limit to a population's tolerance towards IT-related disruptions—especially when these seem to have inexplicable or unintelligible causes. Tolerance will turn into doubt, suspicion and anger directed towards a network society seen as having become uncontrollable.

3.2 Where Rests the Responsibility for Assuring the Information Infrastructure?

Who is responsible for the Internet? This is not primarily a question of who is responsible for maintaining the Internet's *technical* functions, but rather for the enormous amount of information flowing in this worldwide network.

Since the very idea of the Internet is based on free, anonymous flows of information, every sender or poster of information is responsible for what he or she sends, and every receiver of information is responsible for interpreting and making use of this information. In this sense, everyone, and no one, is responsible.

How does this tally with other information systems and networks? The more local, bounded and (relatively) simple a system is, the easier it is to define what is *correct* and what is *incorrect* input and output. As long as there is a specification, such that any state of the system can be tested against it, and as long as it is meaningful to define an outer interface to the system, then some consequential form of responsibility for the system in question can be positioned within its system boundaries.

When systems—including infrastructure systems—begin to blend into one another due to increasing IT-utilization and increasing functional demands, then it is useless to attempt to maintain the fiction of separate systems, each with own internally demarcated mode of responsibility. The distinction between *inside* and *outside* the system, and even the concept of *systems boundaries* as such, becomes blurred.

No firewall, security system, control system or certificate in the world will help when it is no longer possible to determine what is correct or incorrect, before a disruption propagates up through the system structure and manifests itself on the social or political-ideological plane.

This argument concerns primarily so-called *soft* information. As concerns purely technical functions, we may hope that—even in the future—it will be possible to demand responsibility from an electricity supplier when the lights go out, or from the banks when your e-payments fail to go through and you end up with bad credit ratings.

The possibilities of national or local governments regulating the network society, in order to better assure future information infrastructure, would also seem to be minimal. No central authority can control a network—a state of affairs that is, so to speak, built into the very concept of network society.

4. The Vulnerability of the Information Infrastructure to Intentional Disruption

Who, can we imagine, would attempt to damage society by way of attacking the information infrastructure? The outline of possible actors includes hostile states, terrorist groups and fanatical religious movements, criminal organizations and extremist political parties as well as discontented insiders and irresponsible hackers and crackers.

An aggressor, or group of such, who would attack society through its information infrastructure has, in principle, adequate opportunities to cause major damage. However, they will be confronted by a number of difficult practical problems. Our attacker must work secretly and exploit the complexity, speed and opacity of the computerized systems at hand. He (or she) must attempt to calculate the consequences of the contemplated attack, which can itself be a very complex matter and will require a number of correct assumptions concerning countermeasures and operator intervention during the process.

One important factor, which may increase an attacker's chances of success, is that the mental preparedness of non-specialists—as concerns managing computer-related disorders—decreases in relation to increases in computer reliability, a condition that may provide a false sense of security. In addition to this, those still occurring, but all the more exceptional, computer errors often resemble one another structurally, thus increasing the risk of stereotype reaction from users, and thus rendering the discovery of, and measures against, IT-related attacks all the more difficult.

With current developments in IT, it follows that information sent from person to person is seldom sent directly, but flows through a number of anonymous, intervening links and processes. Information injected through evil intent, or even by mistake, can spread through systems in which human operator-control is becoming all the more rare, and the possibility of tracing the source of the "error" all the more difficult.

In the context of conventional threats, accessing the vulnerability of an IT-based system to "external" attacks amounts to evaluating the necessary physical violence required to penetrate a node's (physical) defense, and the effect of the information reduction resulting from its disruption. In the case of an info-logical threat, we need to know how an aggressor can penetrate the node's info-logical shell (or its "protection in depth"), the effect of reduced information—*and* the effect of (further

disrupting) false information emanating from the attacked node and how this may effect a wider system context. This last point makes the problem considerably more complex, and demands much more foresight as concerns analysis and preparedness planning.

5. CIIP-Research in the Future

The question of generalizing and establishing over time the results of studies involving information infrastructure protection is itself a fundamental issue. Does the area of CIIP have a classifiable structure and content which is sufficiently stable in time, such that it will provide a foundation for durable protection and preparedness planning?

At the present time, it would appear that the answer to this question is "no." The problem complex that CIIP deals with represents one of the most dynamic social phenomena in history. Only when this area of research has gained a more stable scientific and methodological base will we be able to change this assessment.

Thus in the short and middle term, developments may dictate that we best direct our efforts towards mitigating—i.e. diminishing the *consequences*—of disruptions to the information infrastructure, rather than attempting to totally prevent their occurrence.

The United States was the first state to take particular notice of the IT-threat to critical infrastructure. The report from the PCCIP [5] (the President's Commission on Critical Infrastructure Protection) puts forward a complex threat assessment in order to discuss what must be done to assure critical infrastructure.

In Europe, both at the strategic/policy level and as concerns research, there are a number of activities in progress with strong association to the area of infrastructure protection. The European Dependability Forum [6] is a European Commission initiative promoting information exchange and discussions on the dependability of Information and Communications Technologies (ICT). The aim of the forum is to provide a platform for exchange of information over a wide range of technical and policy-related domains associated with the dependability of ICT-systems. One of the major concerns is the potential consequences of massive disruptions cascading through the different systems.

The Center for Security Studies and Conflict Research at the ETH in Zurich is developing the comprehensive Risk Analysis and Management Network CRN [7]—an electronic platform for promoting risk-profiling dialogue. Current project partners are the Swedish Agency for Civil Emergency Planning and the Swiss Federal Office for Civil Protection. The project is supported by the Swiss government and additional partners have been invited to participate.

At the Swedish Defence Research Agency (FOI), a long-term research program concerning "Critical Infrastructure Protection" is currently in progress. The program is sponsored by the Swedish Agency for Civil Emergency Planning and is focused on the evolution of the information infrastructure and IT-related threats and vulnerabilities.

Within a few years, and in co-operation with other research groups and other national programs, we hope to be able to establish a coherent plan of research for the study of the evolution of the IT-network society in general, and the development of threats and vulnerabilities to the information infrastructure in particular.

ACKNOWLEDGEMENT: This work is kindly sponsored by the Swedish Agency for Civil Emergency Planning.

DISCLAIMER. The opinions expressed in this article are those of the author and do not necessarily reflect the official standpoint of FOI.

Notes:

1 Charles Perrow, *Normal Accidents: Living with High-Risk Technologies* (New York: Basic Books, 1984).

2 Bruce Sterling, *Hacker Crackdown, Law and Disorder on the Electronic Frontier* (Bantam Books, 1993). Also available @ http://www.lysator.liu.se/etexts/hacker/.

3 Romualdo Pastor-Satorras and Allessandro Vespignani, "Epidemic Spreading in Scale-free Networks," *Physical Review Letter* 86, 14 (2001): 3200-3203.

[4] Duncan J. Watts, "A simple model of fads and cascading failures," (submitted to Physical Review Letter). Available @ http://www.santafe.edu/sfi/publications/Abstracts/00-12-062abs.html, December 2000.

[5] Preliminary Research and Development Roadmap for Protecting and Assuring Critical National Infrastructures, Report available @ http://www.ciao.gov/CIAO_Document_Library/Preliminary_RandD_exsum.htm.

[6] European Dependability Initiative: http://deppy.jrc.it/default.

[7] Comprehensive Risk Analysis and Management Network (CRN): http://www.isn.ethz.ch/crn/index.cfm.

PETER WESTRIN has a PhD in theoretical physics and long professional experience in the development and utilisation of complex simulation systems -- both in academic and industrial contexts. His earlier work in the area of electronic warfare -- which for security reasons has not be published -- has provided a special background for his present work concerning infrastructure vulnerability and the network society. He is currently director of a long-term research program for the study of critical infrastructure protection. Present affiliation: Swedish Defence Research Agency, Division of Defence Analysis, SE-172 90 Stockholm, Sweden. *E-mail*: peter.westrin@foi.se.