Identification of sources of failures and their propagation in critical infrastructures from 12 years of public failure reports

Hafiz Abdur Rahman*, Konstantin Beznosov and José R. Martí

Department of Electrical and Computer Engineering
University of British Columbia, Vancouver, Canada
E-mail: rahmanha@ece.ubc.ca
E-mail: beznosov@ece.ubc.ca
E-mail: jrms@ece.ubc.ca
*Corresponding author

Abstract: Understanding the origin of infrastructure failures and their propagation patterns in critical infrastructures can provide important information for secure and reliable infrastructure design. Among the critical infrastructures, the Communication and Information Technology Infrastructure (CITI) is crucial, as it provides the basic mechanism for sharing information among all infrastructures. Failures in CITI can disrupt the effective functionality of the other critical infrastructures. Conversely, failures in the other infrastructures can also propagate to CITI, and hence disrupt the operation of all systems. In this study, we used public domain failure reports to identify the origin of these failures and their propagation patterns. We analysed 347 infrastructure failure cases reported from 1994 to 2005 in the Association for Computing Machinery's (ACM) RISKS forum. We studied these reports to determine the causes of infrastructure failures and their impact on CITI and other critical infrastructures in a number of dimensions, such as the origin of failures, impacts of failures in spatial and temporal dimensions, their effect on public safety and how failures propagate from one infrastructure to another. The results obtained from the analysis of these real-life failure cases, which occurred over a considerable timespan, should be useful to researchers and practitioners. This paper also discusses the difficulties and limitations of using public domain data in academic research.

Keywords: critical infrastructures; infrastructure failure; telecommunications; infrastructure interdependencies; failure identification; failure propagation; communication and information technology infrastructure; CITI; information and communication technology; ICT; cyber interdependency.

Reference to this paper should be made as follows: Rahman, H.A., Beznosov, K. and Martí, J.R. (2009) 'Identification of sources of failures and their propagation in critical infrastructures from 12 years of public failure reports', *Int. J. Critical Infrastructures*, Vol. 5, No. 3, pp.220–244.

Biographical notes: Hafiz Abdur Rahman is working for his PhD in Electrical Engineering at the University of British Columbia (UBC), Canada. He received his MSECE in 2000 from Purdue University, USA. Prior to his PhD study, he worked for Purdue University and the Ministry of Planning of the government of Bangladesh to develop large-scale computerised systems. He was also a founding partner and Chief Technical Officer of a software

company named CyberSoft. His present research focuses are critical infrastructure simulation and modelling the interdependencies of critical infrastructures with telecommunications networks. He received his BSEE in 1988 from the Chittagong University of Engineering and Technology (CUET), Chittagong, Bangladesh.

Konstantin Beznosov is an Assistant Professor at the Department of Electrical and Computer Engineering, UBC, Canada, where he directs the Laboratory for Education and Research in Secure Systems Engineering. His research interests are distributed systems security, usable security, secure software engineering and access control. Prior to UBC, he was a Security Architect at Hitachi Computer Products (USA) and Concept Five. Besides many academic papers on security engineering in distributed systems, he is also a co-author of *Enterprise Security with EJB and CORBA* and *Mastering Web Services Security*, as well as XACML and several CORBA security specifications.

José R. Martí is a Professor at the Department of Electrical and Computer Engineering of the UBC, Canada, and a Fellow of the Institute of Electrical and Electronic Engineers (IEEE). He is a graduate of the Rensselaer Polytechnic Institute (RPI) in New York (Masters) and UBC (PhD). He has over 25 years of experience in industry and academics on the real-time simulation of large power system networks. Currently, he is the leader of UBC's Complex Interdependent Integrated Systems research group developing the I2Sim simulator for real-time decision support in multiple infrastructure systems.

1 Introduction

Understanding the origin of infrastructure failures and their propagation patterns in critical infrastructures is important for secure and reliable infrastructure design. Among the critical infrastructures, Communication and Information Technology Infrastructure (CITI) is crucial, as it provides the basic mechanism for sharing information among other critical infrastructures (Moteff and Parfomak, 2004) such as electricity, water supply, oil and gas networks, transportation, financial services, *etc.* Over the years, integration of these infrastructures with CITI has become pervasive, extensive, and complex. Failure in CITI, either due to an accident or a malicious action, can propagate to other infrastructures and degrade or disrupt their functionalities. Conversely, failures in other infrastructures can also propagate to CITI and hence disrupt the operation of many of these interconnected systems. Such disruptions may lead to substantial disturbances in the public life of modern nation states. Volatile world situations increase these threats even further. As a result, there are enormous concerns for secure and reliable operation of different critical infrastructures (Moteff and Parfomak, 2004; Bush, 2002). Smooth operation of these interconnected infrastructures requires an understanding of their interdependencies. By studying the origin of infrastructure related failures and their propagation patterns, we can develop a better understanding of their interdependencies, which can be useful to decision makers and infrastructure operators in policy making and system design (JIIRP, 2004).

Since 1992, US telephone companies have been required to submit major failures information to the US Federal Communications Commission (FCC). Using FCC outage reports, a study was done on the failure patterns of Public Switched Telephone Networks (PSTN) (Kuhn, 1997). To the best of our knowledge, no other critical infrastructure providers in North America are obliged to disclose their failure-related information. However, such data could help the research community develop a better understanding of failure patterns and their interdependencies among different critical infrastructures. Data from infrastructure service providers is especially helpful, because they may give detailed information about system states, control parameters, input and output specifications, operating assumptions, backup facilities, management procedures and practices, and other physical and environmental constraints.[1] Unfortunately, both public and private infrastructure operators are reluctant to share this information with the research community (Spafford, 2001). The FCC outage report mentioned above is accessible only to the FCC officials and the US Department of Homeland Security (DHS).[1]

Given this reality, one possible alternative is to use public domain infrastructure failure reports, such as newspaper or other mass media reports to develop an understanding of infrastructure interdependencies. There are two major difficulties in this approach:

1 These failure reports normally provide only brief amounts of information.

2 They do not have any regular structure.

Even though individual reports may not give much information about a specific failure, by studying a large number of cases we can trace common trends among similar classes of failures. To address the second obstacle, we have classified these reports based on their failure type, and extracted meaningful information through some critical attributes. We collected 347 cases of 12 years of failure data (1994 to 2005) from the Association for Computing Machinery's (ACM) RISKS forum,[2] which is the largest known public repository of this kind of report. The RISKS forum was started more than 20 years ago and is still very active. Posting to this forum is moderated, which ensures a certain level of quality and reliability. Since the cases reported to this forum represent only a fraction of all actual events, there may be a concern about the usefulness of the statistics we derive from our analysis. However, the trend of reporting to this forum is related to the public perception of risks, and research shows that despite partial information public perception of risks is fairly accurate (Fischhoff *et al.*, 1982; Rowe and Wright, 2001). To ensure the authenticity of our selected cases, we have paid particular attention to the verifiability of our selected report's sources. For instance, we gave more weight newspaper reports than to reports by private individuals. Our methodology is discussed in detail in the Approach and Methods section.

In this work, we have identified interdependencies among CITI and other infrastructures based on some key factors, such as origin of failures, their impact in spatial and temporal dimensions, their effect on public safety, and their propagation from CITI to other critical infrastructures and vice versa. More specifically, we would like to determine the main causes of infrastructure failures and the nature of their impact; the type of locality affected and their geographical location; how their degree of fatality changed over time; and how infrastructures are related to each other. In the absence of any formal model of interdependencies among CITI and other critical infrastructures, our findings should be useful to policy makers, practitioners, and researchers. In the Related Work section, we discuss previous efforts to classify and interpret infrastructure-related failures. In the Approach and Methods section, we give a brief overview of our methodology. In the Failure Database section, we describe our failure database. In the Results section, we summarise the results of our analysis. Finally, in the Conclusions section, we discuss the contributions of this research as well as future research directions.

2 Related work

There are three major approaches to classifying and interpreting infrastructure related failures. The first approach focuses on the failures and their impacts in relation to CITI (Kuhn, 1997; Neumann, 1994; Chakrabarti and Manimaran, 2002). The second approach focuses on understanding failures in any computer-based system, and is not limited to CITI (Neumann, 1994; Avizienis *et al.*, 2004). The third approach classifies failures and interdependencies among the critical infrastructures in a general system agnostic way (Rinaldi *et al.*, 2001).

Kuhn (1997) analyses PSTN failure data based on the following six failure categories: human errors, acts of nature, hardware failures, software failures, overloads, and vandalism. Using this scheme, Kuhn (1997) analyses two years of PSTN failure data (1992–1994) from the US FCC and shows the impact of different types of failures on PSTN operation. Howard (1997) proposes a taxonomy based on attack type in computer networks and uses that taxonomy to perform a frequency analysis of more than 4000 security-related incidents reported to the Computer Emergency Readiness Team Coordination Center (CERT/CC). From the results of this analysis, he proposes a set of recommendations for government, vendors, the CERT/CC, and individual users to improve security practices. Howard and Longstaff (1998) further extend this taxonomy by incorporating additional terms, such as additional objects and attributes like site name, attack date, reporting time, *etc.* Chakrabarti and Manimaran (2002) propose a taxonomy to classify internet infrastructure security failures, based on a survey of different intrusion detection and prevention techniques. They classify internet infrastructure failures into four categories: DNS hacking, routing-table poisoning, packet mistreatment, and denial of service attacks.

Neumann initiated the ACM's RISKS forum[2] in 1985 to compile computer-related system mishaps that affect public life. In 1994, Neumann (1994) published the book *Computer-Related Risks*, in which he selectively compiled a large collection of RISKS forum reports based on problem sources. These problems included problems in requirement definition, system design, hardware implementation, software implementation, system use and operation, environmental problems, *etc.* Through his analysis, Neumann draws attention to the safety and security issues associated with each type of failure. Avizienis *et al.* (2004) propose another generalised taxonomy based on the reliability and security aspects of computer systems. Their approach is to compile a few key definitions as a set of generalised concepts and to extend those concepts with extended sets of definitions. The main objective of this taxonomy is to be able to use these concepts in a wide variety of cases.

Rinaldi *et al.* (2001) address classification of failures and interdependencies among the critical infrastructures in a system agnostic way. Their taxonomy is based on six functional dimensions to determine cross-infrastructure interdependency issues. These are: type of interdependency (*e.g.*, physical, cyber, logical); infrastructure environment (*e.g.*, business, economic, healthcare); coupling and response behaviour (*e.g.*, adaptive, loose, tight); infrastructure characteristics (*e.g.*, temporal, spatial, organisational); type of failure (*e.g.*, common cause, cascading, escalating); and state of operations. (*e.g.*, normal, stressed, repaired). However, their failure source classification is very restrictive (*e.g.*, common cause, cascading, escalating) and gives a very limited number of options for analysing the RISKS forum failure reports.

In our research, we used Kuhn's (1997) approach for CITI and related critical infrastructure failure classification. However, we added the following two additional categories to Kuhn's original six: malicious logic fault and authorisation violation fault. Even though these two faults are software related, due to their intentional and malicious nature, we placed them in separate categories. In recent years, these two faults have become of increasing concern for critical infrastructure management. Their remedial methodologies are also different from those used for traditional software failures.

3 Approach and methods

We have followed a four-step methodology in collecting and analysing failure reports. We started by systematically collecting failure cases from the RISKS forum. We then categorised these reports based on their failure type, extracted useful information from them, and then performed an analysis of the extracted information. The following sections discuss these steps in detail.

3.1 Data collection

Postings in the RISKS forum cover a wide range of computer related risk topics, including system failure reports, conference announcements, book reviews, *etc.* Collecting useful infrastructure failure reports from RISKS forum data (we scanned more then 10 000 records) was the most difficult time consuming but important step. During our selection process, we selected only those reports where the failure originated from CITI and affected other critical infrastructures (including CITI), or the failure originated from some other critical infrastructure and affected CITI. Failure is defined as the inability of a system or component to perform its required functions within the specified performance requirements,[3] and may be the result of one or many faults. A fault is defined as a defect in a hardware device or component, or an incorrect step, process, or data definition in a computer program.[3] In our study, failure is attributed to critical infrastructures. The following infrastructures are considered critical, based on a US Congress document on critical infrastructures identification (Moteff and Parfomak, 2004):

- IT infrastructure
- telecommunication infrastructure
- water supply

- electrical power system
- oil and gas
- road transportation
- railway transportation
- air transportation
- banking and financial services
- public safety services
- healthcare system
- administration and public services.

We collected a unique report for each incident. However, in cases of widespread failure, we collected unique reports from different affected sites. One example of such widespread failure is an internet-wide worm attack. The apparently simple task of selecting appropriate sets of reports became quite complicated due to the subtleties associated with each of the reports; the following three examples reveal this intricacy. The first is an example of a report that we chose to select:

> "On November 19, 1994, Iowa City's US West telephone system shut down at about 3:30 p.m., local time, and service was gradually restored between 7:30 and 9:30 p.m, affecting about 60,000 people. Analysis showed that a new switching system had been installed in July 1994. In removing the old system, an electrical grounding cable had been inadvertently removed."[4]

The above example clearly shows that a fault in the electrical power system due to human error caused a failure in the telecommunications infrastructure. The report has a clear reference to a newspaper source. In contrast to the above report, the following is an example of a report that we omitted:

> "I suppose I shouldn't be surprised, but the power went out for 17,000 here in our small town (38,000) last week. The local newspaper first reported that the power company didn't know why it went out, but that it 'may be related to someone digging in their back yard'. A week later they fixed the blame. A phone call (by the power company), supposedly to one substation, (completely automated judging by the tone of the article) went instead to a different substation (for unexplained reasons) and shut that substation down. It was down for 1.5 hours."[5]

In the above report, the failure in the electrical power system is not clearly related to CITI. There is no clear reference to where it occurred and there is an undefined term 17 000 in the report. Similarly, we avoid survey reports, as they are not attributed to any particular failure case. The following is an example of such a survey report:

> "The Federal Trade Commission says that complaints of Internet-related identity theft more than tripled last year, to 2,352 last year from the year before. Jay Foley of the Identity Theft Resource Center says, 'Online fraud is becoming as big an issue for eBay and AOL as security is for Microsoft.' Typically, eBay covers buyers or sellers for up to $200 (or $500 for some listings) if an item is not delivered or is in bad condition, though there is a $25 processing fee."[6]

3.2 Fault classification

Our next step was to categorise collected reports based on the nature of the failure. As explained, we used eight fault classes, most of which were derived from the taxonomies proposed by Kuhn (1997). The major advantage of Kuhn's approach is that the failure sources are orthogonal and as such, can be dealt with independently. A similar approach was used by Chillarege *et al.* (1992) for software defect classification. For instance, the risk of hardware faults can be minimised by using a redundant physical channel, redundant backup power supply, *etc.* (Balkovich and Anderson, 2004), which are independent of other types of fault consideration. Similarly, for malicious logic faults, different kinds of protection techniques can be used. These include secure routing protocols, secure domain name systems, firewalls and anti-virus tools (Chakrabarti and Manimaran, 2002). Sometimes failure management can be infrastructure dependent, which requires more specialised tools and techniques. For instance, air transportation services require specialised hardware and software tools to ensure their systems' reliability (Kirwan, 2001). Table 1 shows the different fault classes used in our study.

Table 1 Fault classes related to critical infrastructures

Fault name	*Description*
Hardware fault	All fault classes that affect hardware.
Software fault	Fault caused by an error in the software system.
Human error	Non-deliberate faults introduced by a mistake.
Natural fault	Physical faults that are caused by natural phenomena without human participation.
Overload	Service demand exceeds the designed system capacity.
Vandalism	Sabotage or other intentional damage.
Malicious logic fault	These include trojan horses, logic or timing bombs, viruses, worms, zombies or DoS attack.
Authorisation violation	Attempt by an unauthorised person to access or damage network resources, but does not exclude the possibility of authorised users who are exceeding their rights. This also includes unauthorised sharing of digital contents, like audio, video or software.

Faults that trigger infrastructure failure can be mapped to different generic fault types. These generic faults belong to one of the eight fault classes mentioned above. As we analysed failure cases and identified the root cause of each failure, we tried to identify type of fault source for each of these root causes. These generic fault sources are similar to Kuhn's decomposition of fault classes into finer detail, such as failure of cable component or power supplies, software version mismatch, *etc.* Table 2 lists the generic faults and the fault classes they belong to.

Table 2 Generic faults related to each fault class

Fault class	*Generic fault*
Hardware fault	Physical link failure.
	Hardware design or implementation flaw.
	Failure due to external operating environment exceeds predefined limit.
	Device failure due to lack of backup power supply.
	Device failure due to lack of proper maintenance.
	Device failure, origin unknown.
Software fault	System failure due to software glitch.
	Software design or implementation flaw.
	System failure due to software configuration or update error.
	System failure due to weak encryption algorithm.
Human error	Usability factors not considered in system design.
	Inadequate safety measures.
	Careless mistake.
	Data entry error.
	Lack of proper user training or documentation.
Natural fault	Natural calamity.
	Resource unusable due to natural cause.
Overload	User request failed due to inadequate system capacity.
Vandalism	Intentional breakage of physical links or devices.
Malicious logic fault	System failure due to malicious logic.
	Misguiding using malicious logic.
Authorisation violation	Unauthorised access by the outsider.
	Access right violation by authorised user.
	Unauthorised use of technology for malicious intension.
	Identity theft through authorisation violation.
	Unauthorised capture or sharing of digital contents.

3.3 Feature extraction

Once we categorised a failure report to a particular failure class, we extracted key features from each of these reports using a set of key attributes, which sometimes judgmental. For example, Degree of Impact is a feature that is intended to capture the severity of a failure. Reading the failure case, we tried to understand how many people or systems were affected and how that number affected the overall functionality of an organisation. The degree of impact assigned to a rating 'High' indicates a massive effect on the functionality of CITI and other critical infrastructures. Similarly, ratings of

'Medium' and 'Low' indicate moderate and low impacts, respectively. Clemen *et al.* (2000) show that in the absence of detailed information, a judgmental approach can be followed to predict risk. The following three examples illustrate the assignment of Degree of Impact by subjective judgment.

> Degree of Impact – High (Report ID # 5) – "On November 19, 1994, Iowa City's US West telephone system shut down at about 3:30 p.m., local time, and service was gradually restored between 7:30 and 9:30 p.m, affecting about 60,000 people. Analysis showed that a new switching system had been installed in July 1994. In removing the old system, an electrical grounding cable had been inadvertently removed."
>
> Degree of Impact – Medium (Report ID # 3) – "MCI's inbound internet gateways were saturated during July 1994, resulting in days of delay in delivering e-mail to MCI customers. A fix was considered to be months in the offing."
>
> Degree of Impact – Low (Report ID # 8) – "A software glitch on March 10, 1995, caused Prodigy's e-mail system to send 473 e-mail messages to incorrect recipients and to lose 4,901 other messages. The system had to be shut down for five hours."

In the first report, the telephone service of 60 000 people was affected, which was an important consideration in designating it as a high-impact case. In the second report, even though e-mail service was delayed; other means of communication were available, therefore it was assumed gateway saturation was a moderate inconvenience for the MCI customers and the report was designated as a medium-impact case. In the third report, it seems that after 5 h the e-mail service was fixed and misdirected e-mails were redispatched to the recipients. Since people did not check their e-mail very often and the number of recipients was only 473, we concluded that this failure modestly affected the users, and assigned it a low degree of impact.

Another key feature of a failure report is the Report Accuracy, where we assigned an accuracy rating on a scale of 10 based on the source type. For each of these reports, the information source was given. If the information was released from an official source and had other supporting references for validation, we assigned it 9 or 10 points. If it was from an official source but no further details were given, we assigned it 7 or 8 points. All newspaper reports were given 5 or 6 points. Reports from individuals, which were difficult to verify, were normally given less than 5 points. Higher ratings were given to reports of a particular class if the reports fulfilled most of our additional criteria. For instance, if a newspaper report had most of the required information, such as severity, duration, financial impact, description of fault origin, *etc.*, then it was given 6 points; otherwise, it was given 5 points. In the future, we would like to use this accuracy rating to conduct a reliability analysis of the collected cases. Table 3 lists the extracted key features of a failure report and their meanings.

Many of these features are intended to capture the extent of the failures, their impact in spatial and temporal dimensions, their effect on public safety, and the propagation of the failures from CITI to other critical infrastructures and vise versa. Table 4 groups these features into different categories according to the intended use.

Table 3 Extracted features and their meaning

Feature name	*Meaning*
Title	Title of the report. Most often this is the same as the original report.
Date	Date of the failure report.
Country	Country where fault incident originated. For global fault, it is world.
Impact scale	Size of area affected. It could be an organisation, a city, a region (a big part of the country), a country, a continent, or the whole world.
Degree of impact	Failure impact. Could be high, medium or low.
Simulation	Indicates if the fault conditions can be simulated within a lab environment using I2Sim (Rahman *et al.*, 2008) critical infrastructure simulator.
Fault intent	Fault could be intentional, due to deliberate and malicious attempts by any individual or groups, or unintentional due to human error or system flaw.
Duration	Time from the start of the fault to its full recovery.
Financial impact	Amount of financial loss in million USD.
Public safety	Any public safety concern associated with a particular fault incident, such as failure of 911 service, medical emergency service, fire rescue service, or police service.
Affected sites	Number of sites or locations affected by a particular fault incident.
Description	Description of the failure (report text).
Report source	Reference of the report collected from RISKS forum and referred to as RISKS (i, j), where i is the volume number and j is the issue number within the volume.
Report accuracy	Based on the source type, an accuracy rating on a scale of 10.
Fault class	Fault type is one of the eight types mentioned in Table 1.
Generic fault	A qualitative assessment of the origin of the fault.
Source infrastructure	One of the critical infrastructures discussed in the Data Collection section.
Affected infrastructures	One of the critical infrastructures discussed in the Data Collection section.
Affected industry sectors	Description of the industry sectors affected by the failure.
Comment	Comments on specific interesting aspects of these faults.

Table 4 Features that capture different failure dimensions

Analysis dimension	*Feature names*
Extent of failure	Fault class, degree of impact, fault intent, fault type
Impact (spatial)	Country, impact scale, affected sites
Impact (temporal)	Date, duration
Public safety	Public safety
Failure propagation	Source infrastructure, affected infrastructures

3.4 Data analysis

Many of the public domain failure reports we collected had some missing attributes. For example, the duration of a failure and the number of sites affected by failures were not clearly specified for almost half of the cases. The financial impact of failures is mentioned in fewer than 10% of the cases. As a result, we could not use concepts like ‘Customer Minutes’ (product of average number of customers affected and average outage duration) as used by Kuhn (1997) to measure the severity of failures in PSTN networks. Use of such concepts is possible for FCC reports, as each FCC report has to include date, time, failure duration and the number of affected customers (Kuhn, 1997). Unlike FCC, however, our failure reports did not have such uniformity and universal impact dimensions. To compensate for this, we used a frequency-based approach to quantify results from the extracted features of the failure database (next section). This way, we tried to determine the most likely cause of infrastructure failure, the types of localities affected, and the implications for public safety. As mentioned before, due to the absence of clearly specified values for many key attributes, we had to use our own judgment to estimate some of the values of these key attributes. In doing so, we were limited by the description of the data, and there were no mechanisms for obtaining further detail.

4 Failure database

The collected cases and their extracted features were compiled in a MS Excel database. A sample record from this database is shown below (Figure 1). Each record represents a single row in the MS Excel spreadsheet. The analysis performed on these records was done in another sheet within the same MS Excel file. Each record in this database has a report ID. A report ID is a sequential number assigned based on the incidence date. Other fields have their own set of valid values. Table 5 summarises acceptable values for each of these attributes.

Figure 1 A sample database record

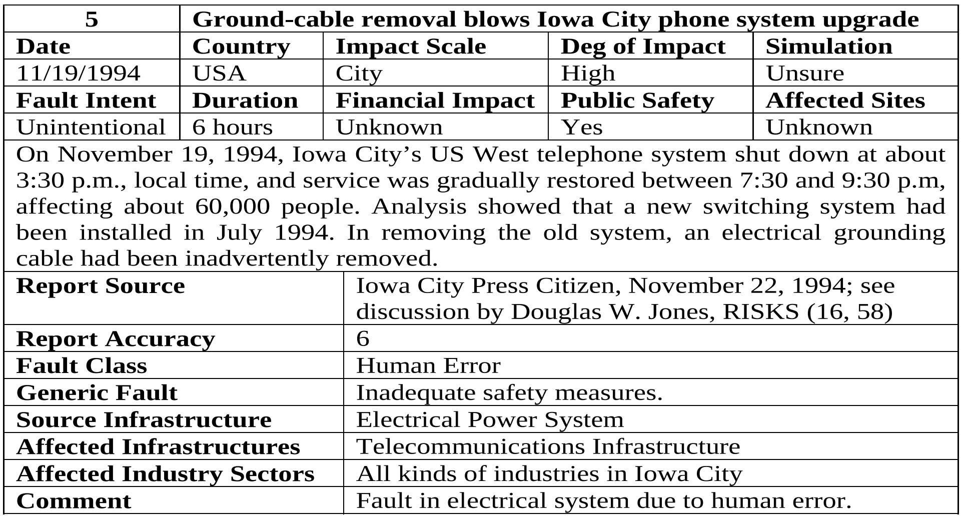

5	**Ground-cable removal blows Iowa City phone system upgrade**			
Date	**Country**	**Impact Scale**	**Deg of Impact**	**Simulation**
11/19/1994	USA	City	High	Unsure
Fault Intent	**Duration**	**Financial Impact**	**Public Safety**	**Affected Sites**
Unintentional	6 hours	Unknown	Yes	Unknown
On November 19, 1994, Iowa City’s US West telephone system shut down at about 3:30 p.m., local time, and service was gradually restored between 7:30 and 9:30 p.m, affecting about 60,000 people. Analysis showed that a new switching system had been installed in July 1994. In removing the old system, an electrical grounding cable had been inadvertently removed.				
Report Source	Iowa City Press Citizen, November 22, 1994; see discussion by Douglas W. Jones, RISKS (16, 58)			
Report Accuracy	6			
Fault Class	Human Error			
Generic Fault	Inadequate safety measures.			
Source Infrastructure	Electrical Power System			
Affected Infrastructures	Telecommunications Infrastructure			
Affected Industry Sectors	All kinds of industries in Iowa City			
Comment	Fault in electrical system due to human error.			

Table 5 Legal values for failure database

Field name	*Legal values*
Report ID #	A sequential number assigned based on the report's date.
Title	Text string
Date	MM/DD/YYYY
Country	Country name/world
Impact scale	Organisation/City/Region/Country/Continent/World
Degree of impact	High/Medium/Low
Simulation	Yes/No/Unsure
Fault intent	Intentional/Unintentional/Unknown
Fault class	One of the eight fault class as (Table 1)
Generic fault	Origin of the fault that belongs to one fault class (Table 2)
Duration	# hour
Financial impact	# million USD
Public safety	Yes/No/Unknown
Affected sites	#/Unknown
Description	Text string
Report source	Text string
Report accuracy	#
Fault origin	Text string
Source infrastructure	One of the infrastructures from Data Collection section
Affected infrastructures	One of the infrastructures from Data Collection section
Affected industry sectors	Text string
Comment	Text string

5 Results

The failure data we collected from the RISKS forum came from two different sources. The first type of data are those events that received much attention and were conveyed to the readers through global news distribution networks, such as the Associated Press, Reuters, *etc.* The second type of data are those events that did not receive similar attention but were made public through regional newspapers, radio or television stations, or different organisations' websites. We found that in 20% of cases, the reports we collected were broadcast through large news networks. The other 80% came from national or local news sources. These sources included major national newspapers like the Washington Post, the New York Times, USA Today, the Guardian, Toronto Star, the Vancouver Sun, New Zealand Herald, *etc.* These reports were forwarded to the RISKS forum by forum users. Figure 2 shows different report sources based on their contributing ratio. In this figure, 'Others' (66%) is a category, which includes all sources that individually contributed 2% or less. It appears that most of the reports in our study (about 60%) came from this type of small, local-level sources. Because our study draws upon a wide range of sources, it can be considered to be broadly representative of actual failure scenarios.

Figure 2 Distribution of report sources (see online version for colours)

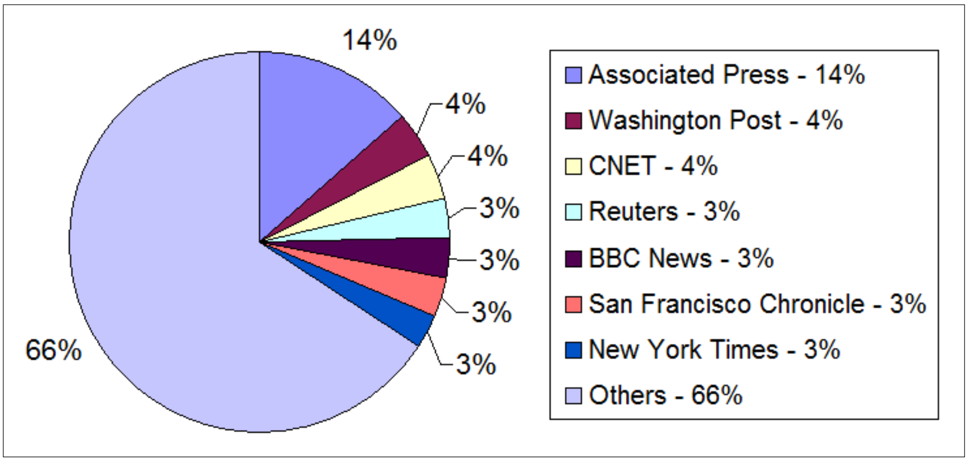

Figure 3 Reported failures over time (see online version for colours)

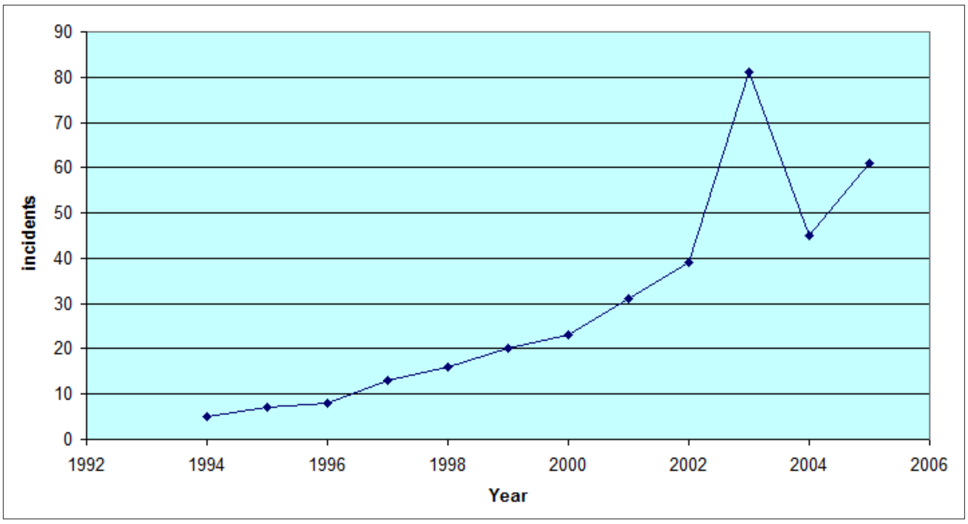

We analysed 347 cases that occurred in a 12 year period (1994 to 2005). Figure 3 shows that the frequency of reporting infrastructure failure to the RISKS forum changed during this period. The trend is nearly linear, except for the year 2003. The linear increase of CITI and other critical infrastructure failure reports may imply that these infrastructures are becoming increasingly dependent on CITI services. However, the sudden rise of failure cases during 2003 was due to the significant escalation of malicious attacks against IT infrastructure (Figure 10). These attacks included different kinds of worm attacks (Slammer, MSBlaster, Nachi) and DoS attacks. We also observed a significant number of Authorisation Violation cases (22 reports) in 2003. These cases included major identity thefts (Report ID# 172, 183), unauthorised digital content sharing (Report ID# 189, 208), change of stock market index using malicious techniques (Report ID# 199), the presence of a fake US government organisation on the internet (Report ID# 167), online auction frauds (Report ID# 207), and similar other cases. The trend was worldwide, and was visible in the remote parts of the world (Report ID# 211, 217). One explanation for the large number of failures during this period is that corporate cyber security mechanisms were not mature enough to compete with the power and availability of automated hacking tools. The increase of malicious attacks during 2003 can be cross-checked by performing a trend analysis on other bounded repositories similar to the RISKS forum (unlike the internet, which is unbounded).

5.1 Failure by category

Figure 4 shows percentages of failures by fault class. It is interesting to note that software-related failures constituted more than 65% of all reported failures (if we include malicious logic and authorisation violation within this group).

Each fault class can be further categorised by their generic fault type. Figure 5 shows software and hardware faults related to all infrastructures classified according to their generic type. The most common cause of software failures is software glitch (45%), followed by software design or implementation flaw (29%). Software glitch is a generic term we use to indicate software failure due to unknown reasons; most often they are related to design or implementation flaws. The high percentage of failures due to software design and implementation defects suggests that better software engineering practices are essential to increase the CITI infrastructure safety and reliability. Similar results were obtained for hardware failures, where the most common cause is device failure due to unknown origin (45%).

Figure 4 Faults that lead to infrastructure failure (see online version for colours)

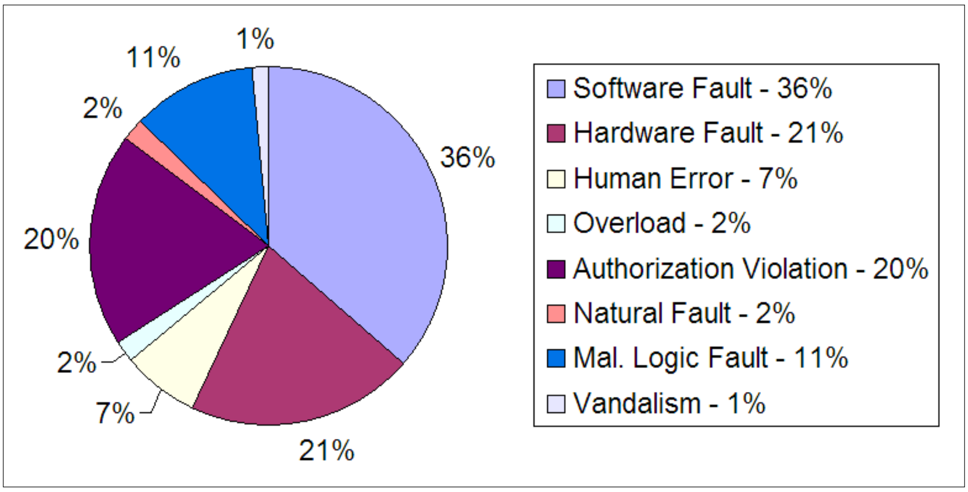

Figure 5 Software and hardware faults are further categorised by generic type (see online version for colours)

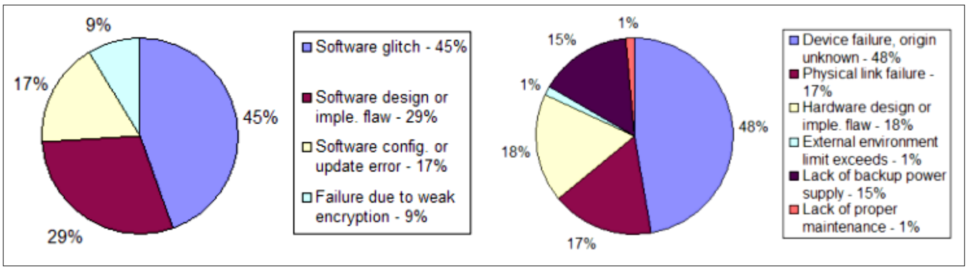

5.2 *Impact of failure*

The failure reports revealed that the root causes of most of the CITI and related infrastructure failures were unintentional or accidental (Figure 6). Critical infrastructures are the lifelines of modern societies. As such, even though their root causes may be unintentional, the impact of failures is significant (Figure 7). Accidental causes include hardware or software faults, configuration problems, human error, *etc.* In contrast, malicious logic faults, authorisation violation attempts, and vandalism account for fewer than 33% of the cases. This situation illustrates the subtle fact that system reliability deserves more attention than it is getting now in relation to system security. An example of this discrepancy is the air transportation industry. Of the 27 air transportation failure cases we reported, 26 were due to various non-malicious hardware and software faults in the air traffic control system. One such case (Report ID # 83) says:

> "On 17 Jun 2000, thousands of would-be passengers were stranded when the main air-traffic control computer collapsed. The National Air Traffic Services computer was fixed later in the day, but the resulting congestion caused many people to spend the night at airports around the UK, and many flights were cancelled the next day as well. Heathrow and Gatwick were hardest hit, although other UK airports experienced severe delays. This was the second time in a week that the computer system had failed."

Figure 6 Failure type distribution (see online version for colours)

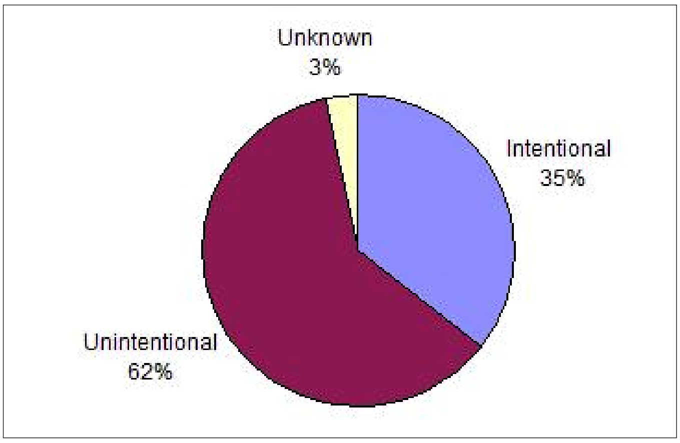

Figure 7 Infrastructure failure impact distribution (see online version for colours)

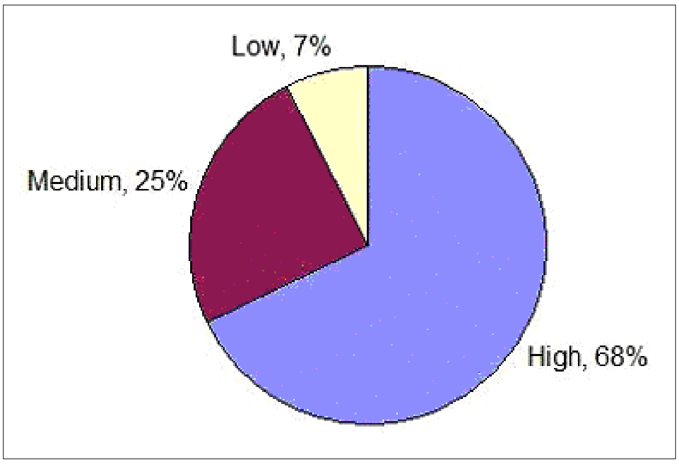

5.3 Public safety concerns

Public safety is a major concern for critical infrastructure failure. Although our study shows that a majority of failure cases did not have public safety implications (Figure 8), we observed that in nearly 20% of the cases, the failure affected public safety to some degree. Many times, these failures were due to improper design or set up of public safety-related devices, lack of backup power supply, *etc.* Since infrastructure failures are on the rise, there are increasing concerns for public safety. An example can be given from the following 911 systems report (Report ID # 230):

> "Houston has deployed a new 911 emergency response system which has had a number of failures since it went 'live' a week ago. Pictures of the new facility look somewhat like Mission Control - large consoles with multiple displays in front of each operator. It sure looks nice, but the system does not appear to work reliably. The latest incident occurred during the day when technicians were working on the link between the computers and units within the cars. To quote: When the system started slowing, technicians reverted to the backup, which crashed within minutes. From 9:50 a.m. to 10:30 a.m., dispatchers resorted to dispatching by radio instead of by computer. Without the computer's locator system, they frequently had to ask emergency workers to volunteer for individual assignments rather than assigning them to calls. Another notable quote is But city officials say the only way to test the system was by going 'live'."

Figure 8 Public safety impact distribution (see online version for colours)

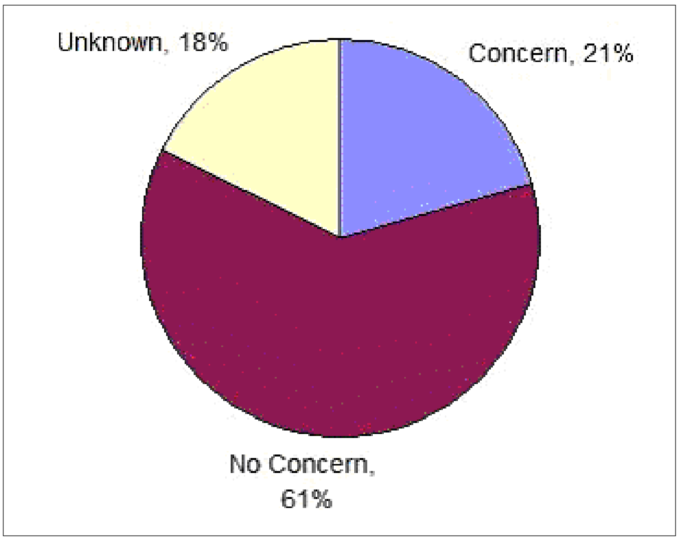

5.4 Change in degree of impact over time

Figure 9 shows that the frequency of high impact infrastructure failure is on the rise. Figure 10 shows that many of these failures are due to malicious intent. Examples of origin of failure include DoS attack against the internet infrastructure (Report ID # 156), worm or virus attack (Report # ID 163, 166) and identity theft (Report # ID 172). From 2001 on, failures due to intentional causes is on the rise. This change of trend is more apparent in recent years. For instance, during 2005, we had 30 intentional failure reports, most of which were related to identity theft, system hacking, phishing, and spamming. These kinds of cyber attacks have become the major form of threat against critical infrastructures. Other contributing factors include the emergence of automated and high-speed worms (*e.g.*, Code Red), increasing deployment of off-the-self software systems for critical infrastructure management (*e.g.*, MS SQL Server), and inadequate expert manpower to manage more complex interconnected infrastructure systems. This increasing dependency on IT infrastructure is making other critical infrastructures ever more vulnerable. There is no sign that this trend will change in the near future. The following example shows how a healthcare system can be affected due to its dependency on computerised prescription systems that depend on electrical power systems (Report ID # 186):

> "Thousands of patients could have received the wrong prescription drugs after a power outage at Kaiser Permanente's computer center in Southern California knocked the pharmacy's labeling system out of sync – printing the wrong labels on filled prescriptions. There were no reports yet of patients suffering from adverse reactions. About 4,700 patients from Fresno to the Oregon border were affected, including those ordering prescriptions by telephone. After the error was discovered on 14 Mar 2003, hospital officials attempted to contact the affected patients, although by 17 Mar, 152 remained uncontacted – including those for whom they had only PO-box addresses."

Figure 9 Change in degree of impact over time (see online version for colours)

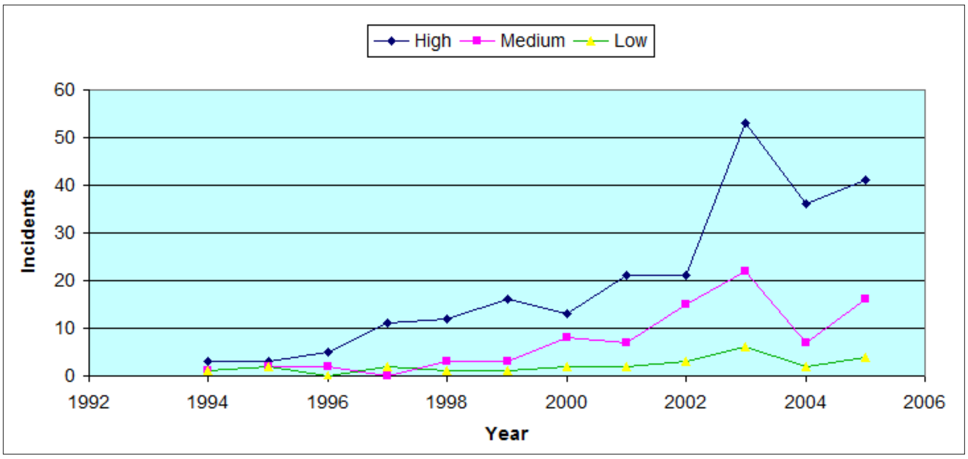

Figure 10 Change in intentional and unintentional failure over time (see online version for colours)

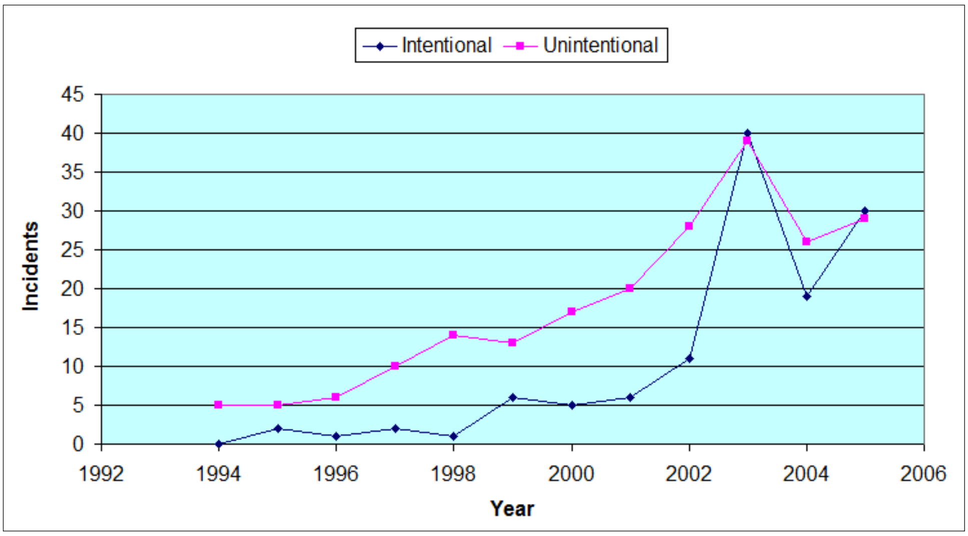

5.5 Localities affected by CITI failures

Figure 11 shows that almost half of the CITI and connected infrastructure failures studied propagated beyond organisational boundaries (47%). Crossing the national boundary was relatively rare, unless an attack was targeted internationally. Figure 12 shows that North America (USA/Canada) was the most vulnerable region for CITI infrastructure failure (63%) in our study. One possible explanation is that this region has a much higher proportion of computer use than any other part of the world. Figure 12 (left) includes worldwide failure cases (*e.g.*, worm attack), whereas the figure on the right excludes those cases. In both figures, most of the reported failures (above 60%) took place in North America (USA/Canada). Inclusion or exclusion of worldwide failure does not significantly change these patterns.

Figure 11 Localities affected by infrastructure failures (see online version for colours)

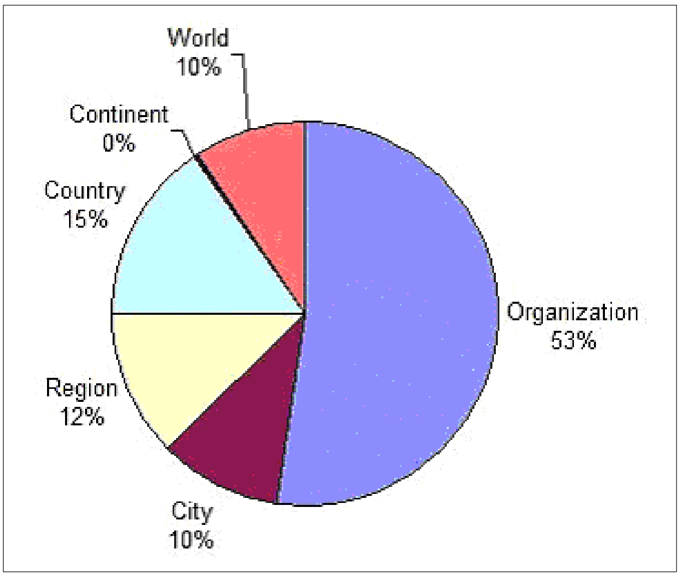

Figure 12 Failure location USA and Canada (see online version for colours)

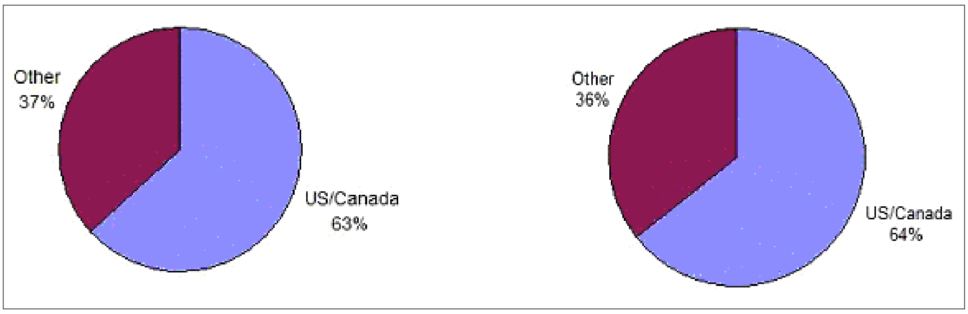

5.6 *Interdependencies among CITI and other infrastructures*

Figure 13 shows that in most of the cases studied, CITI failures originated from within the CITI infrastructure. The role of other infrastructures was relatively minor. Figure 14 shows that most of the CITI failures affected banking and financial services, administration and public services, and the CITI infrastructure itself (IT and telecommunications).

Figure 13 Source of failures affecting CITI (see online version for colours)

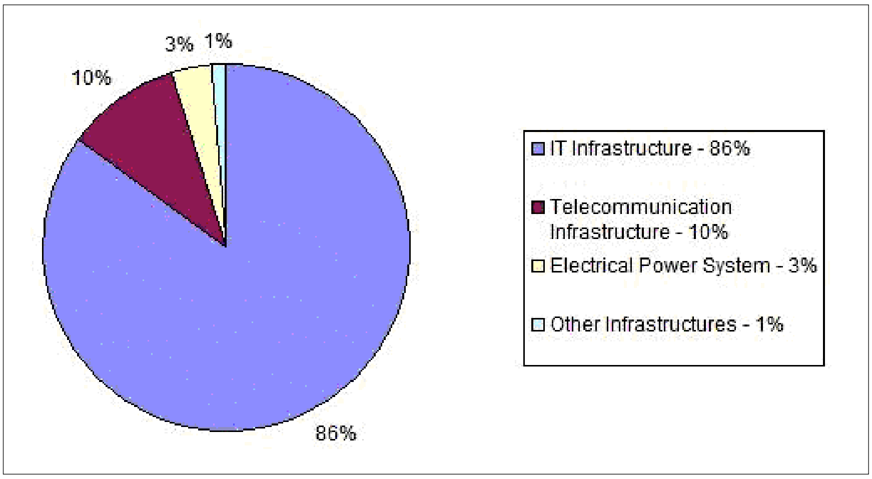

Figure 14 Infrastructures affected due to CITI failures (see online version for colours)

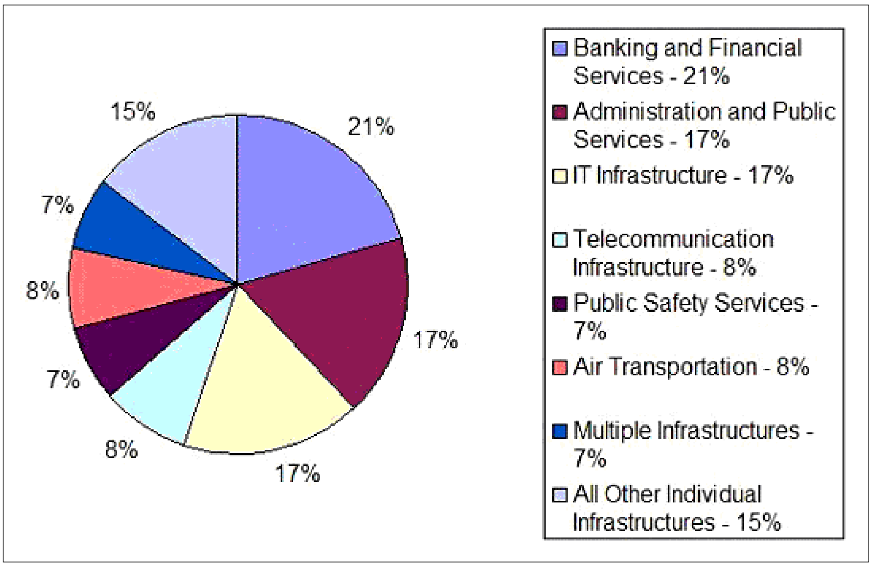

Results of a more detailed analysis of the failures of the first four infrastructures shown in Figure 14 are presented in Figures 15 to 18. Figure 15 shows that software systems were the most vulnerable points for banking and financial services and that a large percentage of these failures had malicious origins (45%). Better software engineering practices and incorporation of adequate security measures can improve the reliability of banking and financial services.

Figure 15 Generic faults that led to banking and financial services failures (see online version for colours)

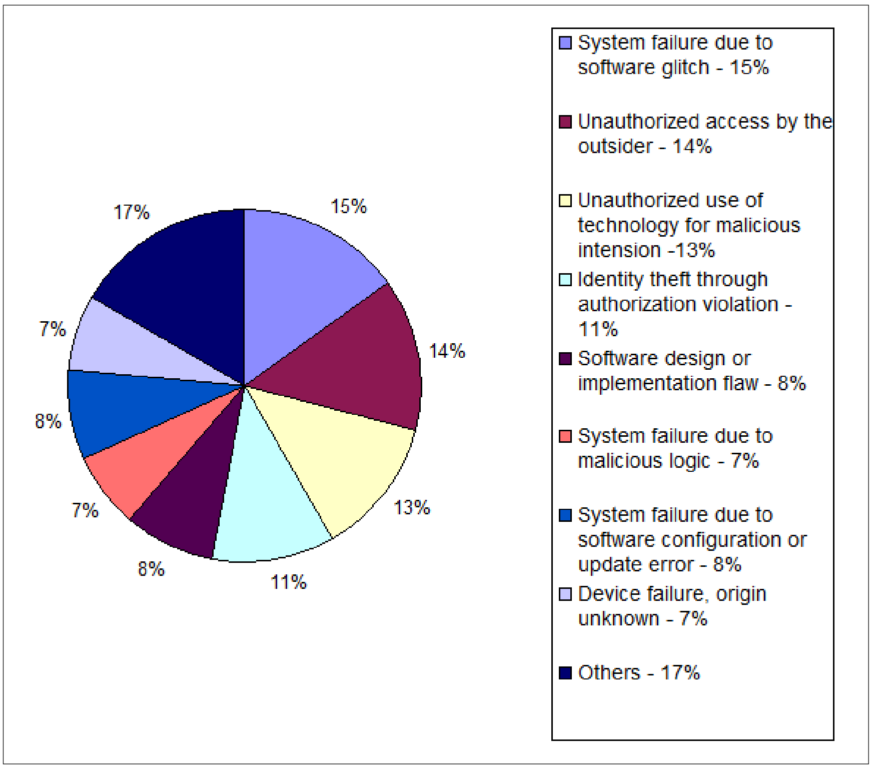

Administration and public services infrastructure includes large government organisations, universities, and other educational institutions. Figure 16 shows that these organisations were also susceptible to software-related failures in the reports we studied, and that large percentage of them had malicious origins (37%).

Figures 17 and 18 show most of the CITI failures originated within CITI. Detailed analysis reveals that IT infrastructure in our study was largely vulnerable to software-related failures, whereas telecommunications infrastructure was vulnerable to different hardware-related failures. Therefore, improved software and hardware related techniques in CITI infrastructure design, implementation, and management will ensure a greater stability in its operation, which will eventually improve the reliability of other connected infrastructures.

Figure 16 Generic faults that led to administration and public services failures (see online version for colours)

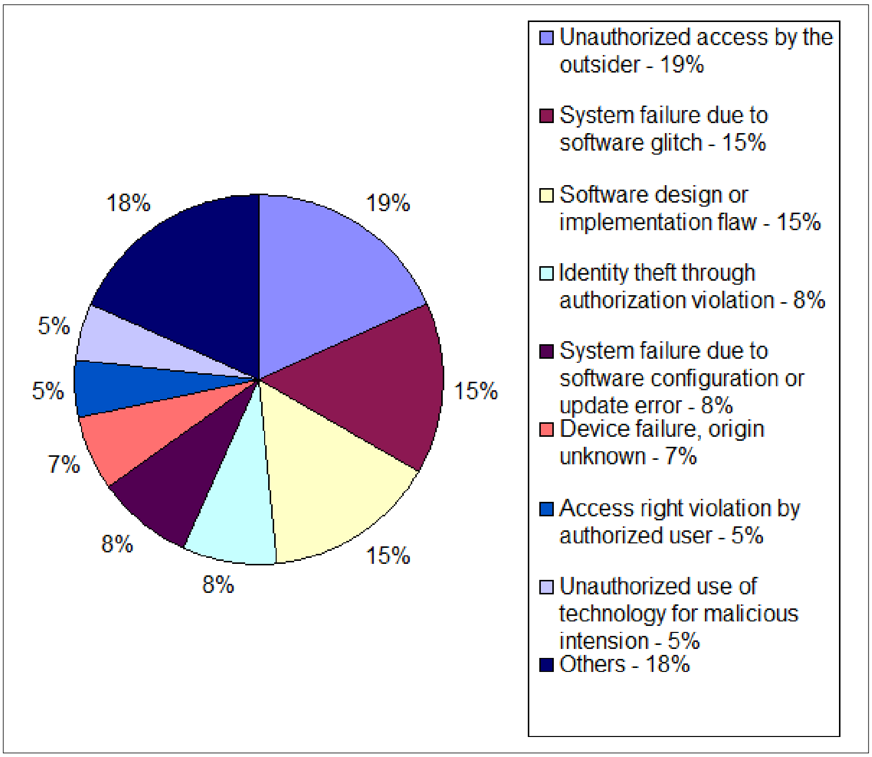

Figure 17 Generic faults that led to IT infrastructure failure (see online version for colours)

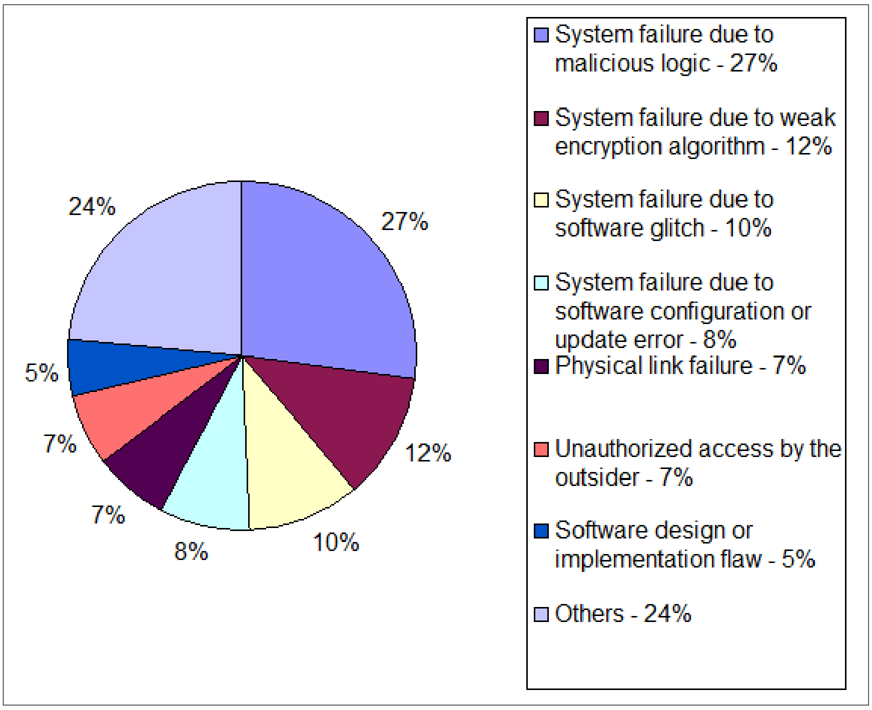

Figure 18 Generic faults that led to telecommunications infrastructure failure (see online version for colours)

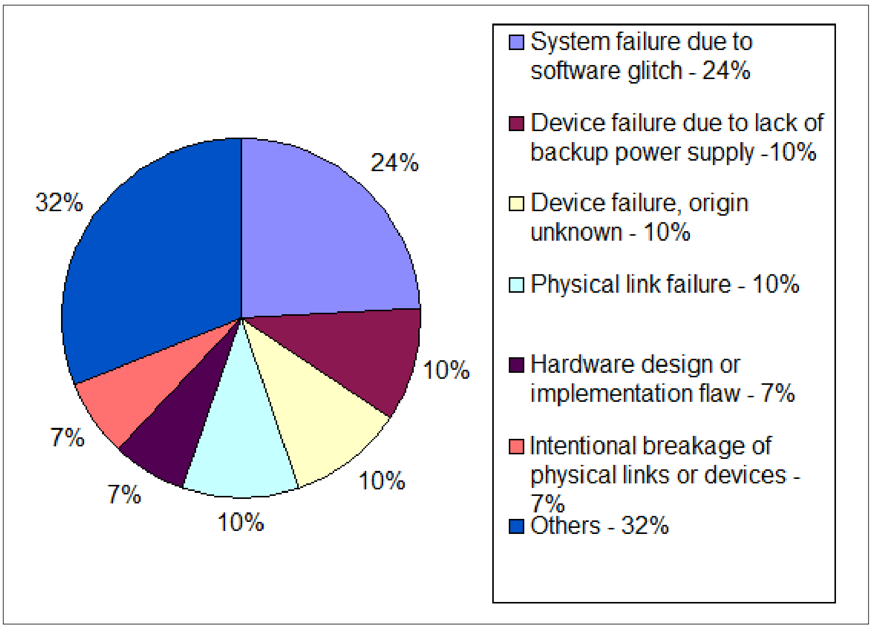

6 Conclusions

Our society relies upon continued services from interdependent critical infrastructures to function. CITI failures are particularly pervasive in their penetration of all infrastructures, and can have a very large impact on the workings of society. Understanding and classifying patterns of CITI failures is an important step towards quantifying interdependency analyses in CITI-dependent infrastructure systems and identifying preparedness and mitigation strategies to ameliorate the impact of system-wide failures. In this study, we have used public domain data over 12 years to understand CITI interdependencies. To our knowledge, this is the first attempt of this type of analysis in this area (using either public or private data sources). In our analysis, we identified infrastructure failure patterns, propagation, impacts on public life, and historical trends. We have developed a CITI failure database for our Infrastructures Interdependencies Coordination (I2C) research group at UBC in setting up realistic test-case scenarios involving CITI failures during large-scale system failure situations (Martí *et al.*, 2008; Hollman *et al.*, 2007).

Acknowledgements

The work presented in this paper is part of the effort currently in progress at the University of British Columbia, Vancouver, Canada, to better understand infrastructure interdependencies and their effects during large-scale disasters. This work is sponsored by the Joint Infrastructures Interdependencies Research Program (JIIRP, 2004) of the Government of Canada. The UBC Infrastructures Interdependencies Coordination group (I2C) (I2SIM, 2005) includes fourteen researchers from a number of disciplines in engineering, computer science, commerce, and psychology. CITI interdependencies are an important concern of this group. Suggestions received from I2C and LERSSE[7] meetings and discussion on some aspects of this work with Professor Lee Iverson in the initial phase were helpful. Interest and help received from Professor Virgilio A. Centeno of Virginia Tech was one of the key reasons to make this work published. We would like to thank the anonymous reviewers for providing helpful feedback. ICICS copy editor Craig Wilson was instrumental to make this paper more readable.

References

Avizienis, A., Laprie, J-C., Randell, B. and Landwehr, C. (2004) 'Basic concepts and taxonomy of dependable and secure computing', *IEEE Transactions on Dependable and Secure Computing*, January, Vol. 1, No. 1, pp.11–33.

Balkovich, E.E. and Anderson, R.H. (2004) 'Critical infrastructures will remain vulnerable: neighbourhoods must fend for themselves', *Int. J. Critical Infrastructures*, Vol. 1, No. 1 pp.8–19.

Bush, G.W. (2002) 'Executive order on critical infrastructure protection', *Proceedings of the 12th Annual Conference on Computers, Freedom and Privacy*, New York, NY: ACM Press, pp.1–10.

Chakrabarti, A. and Manimaran, G. (2002) 'Internet infrastructure security: a taxonomy', *IEEE Network*, November–December, Vol. 16, No. 6, pp.13–21.

Chillarege, R., Bhandari, I.S., Chaar, J.K., Halliday, M.J., Moebus, D.S., Ray, B.K. and Wong, M-Y. (1992) 'Orthogonal defect classification – a concept for in-process measurements', *IEEE Transactions on Software Engineering*, November, Vol. 18, No. 11.

Clemen, R.T., Fischer, G.W. and Winkler, R.L. (2000) 'Assessing dependence: some experimental results', *Management Science*, August, Vol. 46, No. 8, pp.1100–1115.

Fischhoff, B., Slovic, P. and Lichtenstein, S. (1982) 'Lay foibles and expert fables in judgments about risk', *The American Statistician*, *Part 2: Proceedings of the Sixth Symposium on Statistics and the Environment*, August, Vol. 36, No. 3, pp.240–255.

Hollman, J.A., Martí, J.R., Jatskevich, J. and Srivastava, K.D. (2007) 'Dynamic islanding of critical infrastructures: a suitable strategy to survive and mitigate critical events', *Int. J. Emergency Management*, Inderscience, Vol. 4, No. 1, pp.45–58.

Howard, J.D. (1997) 'An analysis of security incidents on the internet 1989–1995', PhD thesis, Carnegie Mellon University.

Howard, J.D. and Longstaff, T.A. (1998) 'A common language for computer security incidents', Sandia National Laboratories Techinical Report SAND98-8997.

Infrastructure Interdependencies Simulation (I2SIM) (2005) University of British Columbia, http://www.ece.ubc.ca/~jiirp/.

Joint Infrastructure Interdependencies Research Program (JIIRP) (2004) Government of Canada, http://www.nserc.ca/programs/jiirp_e.htm.

Kirwan, B. (2001) 'The role of the controller in the accelerating industry of air traffic management', *Safety Science*, Vol. 37, Nos. 2–3, pp.151–185.

Kuhn, D.R. (1997) 'Sources of failure in the public switched telephone network', *IEEE Computer*, April, Vol. 30, No. 4, pp.31–36.

Martí, J.R., Hollman, J.A., Ventura, C. and Jatskevich, J. (2008) 'Dynamic recovery of critical infrastructures: real-time temporal coordination', *Int. J. Critical Infrastructures*, Inderscience, Vol. 4, No. 1, pp.17–31.

Moteff, J.D. and Parfomak, P. (2004) *Critical Infrastructure and Key Assets: Definition and Identification*, Congressional Research Service, Library of Congress.

Neumann, P.G. (1994) *Computer-Related Risks*, Addison-Wesley Professional, 18 October, ISBN: 020155805X.

Rahman, H.A., Armstrong, M., Mao, D. and Martí, J.R. (2008) 'I2Sim: a matrix-partition based framework for critical infrastructure interdependencies simulation', *8th IEEE Electrical Power & Energy Conference*, Vancouver, Canada, EPEC, 6–7 October, http://www.ece.ubc.ca/~rahmanha/I2Sim_paper_ECEC2008.pdf.

Rinaldi, S.M., Peerenboom, J.P. and Kelly, T.K. (2001) 'Identifying, understanding, and analyzing critical infrastructure interdependencies', *IEEE Control Systems Magazine*, December, Vol. 21, No. 6, pp.11–25.

Rowe, G. and Wright, G. (2001) 'Differences in expert and lay judgments of risk: myth or reality?', *Risk Analysis*, April, Vol. 21, No. 2, pp.341–356.

Spafford, E.H. (2001) 'Congressional testimony', 10 October, http://usacm.acm.org/usacm/crypto/spaf.pdf.

Notes

1 FCC Network Outage Reporting System – User Manual, http://www.fcc.gov/oet/outage/nors_manual.pdf.

2 The RISKS Forum, http://catless.ncl.ac.uk/Risks.

3 IEEE Standard Glossary of Software Engineering Terminology, IEEE Standard 610.12-1990, 1990.

4 'Extended phone failure in Iowa City', *RISKS*, Vol. 16, No. 58.

5 'Make a call, turn off the power', *RISKS*, Vol. 17, No. 4.

6 'Internet fraud update', *RISKS*, Vol. 22, No. 98.

7 LERSSE, http://lersse.ece.ubc.ca.