



Carnegie Mellon
Software Engineering Institute


The Survivability of Network Systems: An Empirical Analysis

Soumyo D. Moitra
Suresh L. Konda

December 2000

TECHNICAL REPORT
CMU/SEI-2000-TR-021
ESC-TR-2000-021

DISTRIBUTION STATEMENT A
Approved for Public Release
Distribution Unlimited



Carnegie Mellon University does not discriminate and Carnegie Mellon University is required not to discriminate in admission, employment, or administration of its programs or activities on the basis of race, color, national origin, sex or handicap in violation of Title VI of the Civil Rights Act of 1964, Title IX of the Educational Amendments of 1972 and Section 504 of the Rehabilitation Act of 1973 or other federal, state, or local laws or executive orders.

In addition, Carnegie Mellon University does not discriminate in admission, employment or administration of its programs on the basis of religion, creed, ancestry, belief, age, veteran status, sexual orientation or in violation of federal, state, or local laws or executive orders. However, in the judgment of the Carnegie Mellon Human Relations Commission, the Department of Defense policy of "Don't ask, don't tell, don't pursue" excludes openly gay, lesbian and bisexual students from receiving ROTC scholarships or serving in the military. Nevertheless, all ROTC classes at Carnegie Mellon University are available to all students.

Inquiries concerning application of these statements should be directed to the Provost, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, telephone (412) 268-6684 or the Vice President for Enrollment, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, telephone (412) 268-2056.

Obtain general information about Carnegie Mellon University by calling (412) 268-2000.



Carnegie Mellon
Software Engineering Institute
Pittsburgh, PA 15213-3890

The Survivability of Network Systems: An Empirical Analysis

CMU/SEI-2000-TR-021
ESC-TR-2000-021

Soumyo D. Moitra
Suresh L. Konda

December 2000

Networked Survivable Systems

Unlimited distribution subject to the copyright.

20010420 021

This report was prepared for the

SEI Joint Program Office
HQ ESC/DIB
5 Eglin Street
Hanscom AFB, MA 01731-2116

The ideas and findings in this report should not be construed as an official DoD position. It is published in the interest of scientific and technical information exchange.

FOR THE COMMANDER



Norton L. Compton, Lt Col., USAF
SEI Joint Program Office

This work is sponsored by the U.S. Department of Defense. The Software Engineering Institute is a federally funded research and development center sponsored by the U.S. Department of Defense.

Copyright 2000 by Carnegie Mellon University.

NO WARRANTY

THIS CARNEGIE MELLON UNIVERSITY AND SOFTWARE ENGINEERING INSTITUTE MATERIAL IS FURNISHED ON AN "AS-IS" BASIS. CARNEGIE MELLON UNIVERSITY MAKES NO WARRANTIES OF ANY KIND, EITHER EXPRESSED OR IMPLIED, AS TO ANY MATTER INCLUDING, BUT NOT LIMITED TO, WARRANTY OF FITNESS FOR PURPOSE OR MERCHANTABILITY, EXCLUSIVITY, OR RESULTS OBTAINED FROM USE OF THE MATERIAL. CARNEGIE MELLON UNIVERSITY DOES NOT MAKE ANY WARRANTY OF ANY KIND WITH RESPECT TO FREEDOM FROM PATENT, TRADEMARK, OR COPYRIGHT INFRINGEMENT.

Use of any trademarks in this report is not intended in any way to infringe on the rights of the trademark holder.

Internal use. Permission to reproduce this document and to prepare derivative works from this document for internal use is granted, provided the copyright and "No Warranty" statements are included with all reproductions and derivative works.

External use. Requests for permission to reproduce this document or prepare derivative works of this document for external and commercial use should be addressed to the SEI Licensing Agent.

This work was created in the performance of Federal Government Contract Number F19628-00-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center. The Government of the United States has a royalty-free government-purpose license to use, duplicate, or disclose the work, in whole or in part and in any manner, and to have or permit others to do so, for government purposes pursuant to the copyright license under the clause at 52.227-7013.

For information about purchasing paper copies of SEI reports, please visit the publications portion of our Web site (<http://www.sei.cmu.edu/publications/pubweb.html>).

Table of Contents

Acknowledgements	vii
Abstract	ix
1 Introduction	1
2 Literature Review and Taxonomy	5
3 Model Development	7
4 Data Analysis	9
5 Data Description	11
6 Results	13
7 Simulation Results	21
8 Future Research	29
References	31
Appendix	33
Details of the Simulation Model and Survivability Measures	33
Additional Data Analysis	38
A. Summary Statistics	38
B. Frequency Distribution of Number of Incidents per Day	39
C. Average τ Disaggregated	40

List of Figures

Figure 1: The Components of the Simulation Model	3
Figure 2: The Marked Stochastic Point Process	7
Figure 3: Frequencies by Day	14
Figure 4: Frequencies by Week	14
Figure 5: Frequencies by Month	15
Figure 6: Frequencies by Quarter	15
Figure 7: Survivability/Cost Curves	27
Figure 8: Successive Stages of Compromise	36
Figure 9: Number of Incidents per Day	40

List of Tables

Table 1: Frequency Distribution of τ by Time Interval D (D in days)	13
Table 2: Frequency Distributions of τ 's by Year	15
Table 3: Average τ 's by Year	16
Table 4: Average τ 's by Type (j) and by Year	17
Table 5: Average τ (i,j)'s	18
Table 6a: Frequency of Incidents by Prior (i) and Subsequent (j) Types	18
Table 6b: Type-Switch Matrix (i)	19
Table 7: P(j)'s by Year	19
Table 8: Expected Survivability/Average Damage and P(j)	22
Table 9: Expected Survivability/Average Damage and π_1	22
Table 10: Expected Survivability/Average Damage and χ_1	23
Table 11: Expected Survivability/Average Damage and π_3	23
Table 12: Expected Survivability/Average Damage and χ_3	24
Table 13: Expected Survivability/Average Damage and AVS	24
Table 14: Average Damage and Arrival Rate a_0	25
Table 15: Average Damage and a'	25

Table 16: Relative Changes in Expected Survivability with Respect to Parameters	26
Table 17: Relative Changes in Average Damage with Respect to Parameters	26
Table 18: Frequency Distribution of Number of Incidents per Site	39
Table 19: Frequency Distribution of Number of Incidents per Day	40
Table 20: Average τ by Domain	41
Table 21: Average τ by Number of Sites	41
Table 22: Average τ by Year (Year 1 = 1988)	41

Acknowledgements

This work was done while the first author was a Visiting Scientist at CERT. We wish to acknowledge the support and encouragement of Tom Longstaff and also wish to thank the members of the support staff at CERT for helping us in many ways with this work.

Soumyo D. Moitra

Suresh L. Konda

December 2000

Abstract

This report presents an extended analysis of CERT Coordination Center[®] incidents data (from 1988 to 1995) and applies the results to simulate attacks and their impacts on network sites. The data were “sanitized” prior to the analysis to ensure complete anonymity. A model for the incidents process is discussed and extended. It consists of three parts: a stochastic process for the random occurrence of incidents at sites, a model for the state transition process for an attacked system given a level of defense, and a method of estimating the expected survivability of the system given possible degradations due to these attacks. This approach leads to the estimation of a survivability/cost function, which shows the tradeoffs involved between cost and system survivability. IS managers can use this to determine the most appropriate level of defense for the network systems of their organizations.

The stochastic process was simulated based on parameter values obtained from actual reported data. Extensive sensitivity analyses are reported that indicate how expected survivability would change with varying parameter analysis results values. The report concludes with a discussion of future work to be done and the appendix has details of the simulation model and further data.

[®] CERT and CERT Coordination Center are registered in the U.S. Patent and Trademark Office.

1 Introduction

The pervasiveness of network systems in today's world is certainly well recognized. Moreover, with the growth of the World Wide Web and other communications links between computers and information systems, many networks have become unbounded. That is, no one user has full knowledge of all the connections that comprise the network [Ellison 97, Fisher 99]. Concurrently, more and more elements of society are becoming increasingly dependent on information networks. The vulnerability of these networks is increasing because greater, open access necessarily subsumes greater access for potential attackers. [Baker 95, Gollman 99].

Since it is virtually impossible to control users in unbounded networks such as the Internet, malicious attacks will inevitably occur. Some of these attacks may cause damage to systems and loss to their owners. The magnitude of damages done and costs incurred as a result of such attacks have been estimated at varying levels, but it is clear that even by conservative estimates, they are considerable [CSI 98, Boni 99]. Therefore it has become imperative for systems managers and researchers to consider methods for improving the security of Information Systems (ISs).

It is probably futile to hope for an *absolute* security for any network system such that no possible attack will cause any damage. Almost certainly, as the sophistication of attackers increases (something we are witnessing) any open system can be compromised to some degree or other. The real issue is the level to which we deploy defense mechanisms against these random attacks. Stronger defenses will imply higher costs, and we have to consider tradeoffs between security and costs, where costs could include possible functional limitations to the system. That is, while we need to enhance security, we also need to decide by *how much* to enhance it—given the costs to the organization that owns the system. In other words, we have to determine how to enhance network security for ISs *efficiently*. We would like to achieve the most appropriate level of security based on the organizational needs, financial abilities, and potential threats [Cameron 98, Bernstein 96].

In view of this, a *cost/benefit analysis* of network systems security is clearly important. The costs will be those of deploying and maintaining various defense mechanisms to protect a system or site against attacks, including the costs of any constraints on the system imposed by the defense mechanism. For example, some desired characteristic such as remote accessibility or simple search abilities may conflict with increased security requirements. The benefits will be those of increased survivability of the system or site. Survivability means the ability of systems to recover from attacks, and in particular, the *degree* to which they recover

[Ellison 97]. However, there are many dimensions to survivability, as we shall discuss later. Cost/benefit analysis should lead to methods for improving the security and survivability of network systems in appropriate, cost-effective ways.

As part of such an effort, we need to model the occurrence of attacks on systems and their impacts. Then we can simulate alternative scenarios to examine how different parameters affect system survivability. In this paper, we develop a model and simulate it to analyze network survivability based on available data.

A primary objective of this research is to develop and apply a reasonably realistic simulation model that can help systems managers and CIOs (Chief Information Officers) to understand survivability issues better and evaluate the tradeoffs involved in decisions about network systems design, including their defense mechanisms. One such model has already been developed and reported [Moitra 00]. Here we extend the model to make it more realistic, and we run the simulations based on parameter values estimated from analysis of data on actual incidents recorded at CERT[®] (at the Software Engineering Institute, Carnegie Mellon University). Furthermore, we wish to

- analyze CERT data at the site level rather than at incident level
- develop heuristics for moving towards “optimal” security strategies
- suggest how this simulation model could be embedded in a Decision Support Systems (DSS) to manage systems security and survivability

The incidents process can be viewed as a random process where a system is subjected to a series of random attacks over time (with incidents and attacks as defined above). Since we wish to assess the survivability of systems, we need to model the process of occurrence of incidents from the point of view of a system or site that experiences this process over time, as shown in Figure 1. This is equivalent to a stochastic point process where incidents occur at random points in time. Therefore we need to simulate a stochastic point process. The survivability also depends on how the system responds to an incident. This will depend on the system configuration, that is, its design and defense mechanisms as defined above. Thus we also need to model this response as a function of the incident type and configuration. The model will involve a transition matrix that will give the probabilities of the system ending up in any of its possible states after experiencing an incident. These probabilities will depend on the incident type and system configuration. A minor incident will probably cause relatively little damage. Also, the stronger the defense mechanism, the less the damage that will be inflicted by a given incident. Next, the degree to which the system has survived will have to be measured. This will be a function of the state in which it ends up and will be related to the amount of compromise that has occurred. For this purpose, we use a new survivability measure [Moitra 00] that takes into account the different dimensions of survivability, that is, the different functionalities and services that can be compromised.

[®] CERT and CERT Coordination Center are registered in the U.S. Patent and Trademark Office.

Modeling Approaches for Managing Survivability of Networks

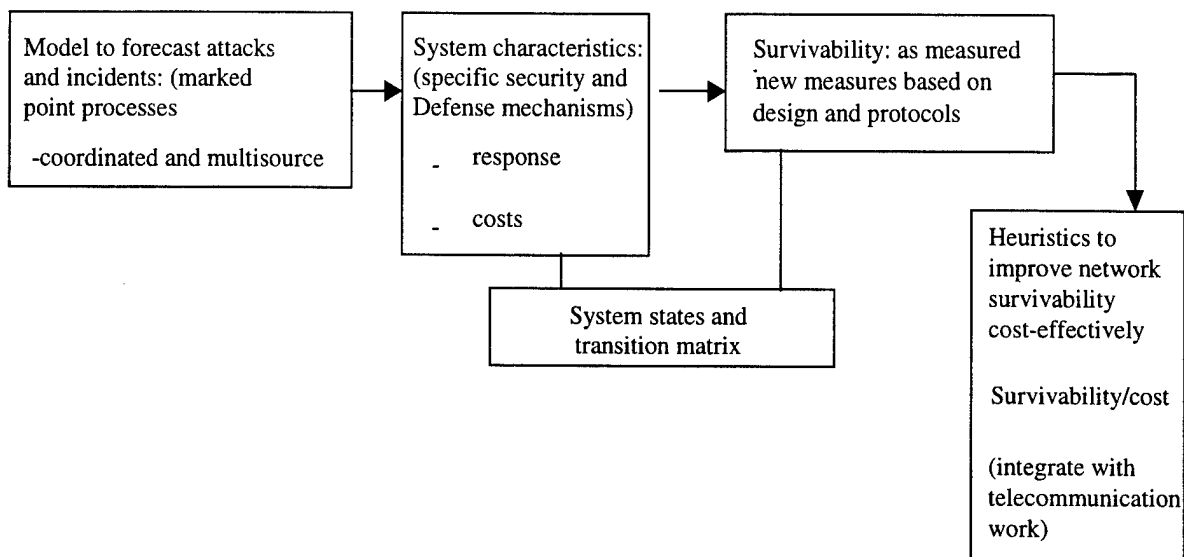


Figure 1: The Components of the Simulation Model

With this simulation model we can analyze the costs and benefits of alternative defense mechanisms under various scenarios. Such analyses can assist systems managers in making decisions regarding the system configuration that best suits their needs. The advantage of this systems simulation approach is that a large variety of scenarios may be explored. Alternative incident processes, different systems configurations, various state transition probabilities, and additional survivability measures may all be investigated with such a model. Thus, given the high degree of uncertainty regarding future attacks and their impacts, this method provides a practical approach to assess and manage survivability.

The rest of the paper is organized as follows. The next section reviews the literature in this area and the following two sections summarize the model that has been previously reported and the input data required to run it. After that we describe the data analysis and the main results. The results of the simulation are discussed next. The final section outlines some of the future work that should be done. Further details of the model and data analysis are given in the appendix.

2 Literature Review and Taxonomy

Survivability has been studied in telecommunications where the impacts of link or node failures have been considered [Moitra 97]. There is also some literature on the survivability of network information systems [Howard 95, Ellison 97, Linger 98, Fisher 99]. Howard has undertaken an extensive survey of the nature of attacks on computer systems and reports on the analysis of data on computer security violations. Ellison discusses the issue of survivability in the context of network systems and proposes a set of future research needs. Linger expands on the problems in analyzing survivability of network systems and explains the requirements for system survivability. Fisher suggests an alternative method for increasing survivability by applying emergent algorithms based on a distributed system.

A simulation model to track the impact of attacks on network systems has been proposed [Cohen 99] and it is demonstrated that shortening the response-to-attack time can have major benefits. A detailed model of the incidents process has been developed [Moitra 00] and has been simulated to provide a method to explore the costs and benefits of varying defense mechanisms. Moitra and Konda also provide a methodology to measure survivability. In addition there is a burgeoning literature on e-commerce and the security issues related to it [Baker 95, Bernstein 96, Boni 99, Cameron 98, Gollman 99].

One of the key issues is the taxonomy that is to be used in discussing network security incidents and survivability. Many alternative terms have been used, and here we follow the attempt to develop a common language for computer security incidents Howard and Longstaff define a number of terms, in particular, an attack and an incident [Howard 98]. An attack is defined as

“a series of intentional steps taken by an attacker to achieve an unauthorized result.”


An incident is defined as

“a group of related attacks that can be distinguished from other attacks because of the distinctiveness of the attackers, attacks, objectives, sites, and timing.”

It would be useful to add some additional terms. We define an *episode* as the combination [incident + response], that is the whole process of a set of attacks and the system's response to the incident. We consider a *system* to be the collection of all the relevant computers and network elements at a *site*. Finally we will refer to the system's *configuration* as the combination of its design and its defense mechanism.

3 Model Development

As illustrated in Figure 1, we shall develop models of the incidents process, the response of the system, and its survivability. In order to forecast incidents, we model the process as a marked, stochastic point process, where the incidents are the events that occur at random points in time, and the event type is the mark associated with an incident [Snyder 91]. The mark is used to identify random quantities associated with the point it accompanies. The mark, or event type in our case, has to take into account the severity of the incident and the possibility of single, or multiple and simultaneous attacks. This is because we are modeling a process that is taking place in an unbounded environment [Ellison 97]. Therefore the mark space will be two-dimensional, characterized by type (severity) and number of attackers. However, since no data on the distribution of the number of attackers per incident were available, only severity was used in the simulations. A marked point process is illustrated in Figure 2.



- τ ~ inter-incident time;
- t ~ times at which incidents occur;
- j ~ marks associated with each incident (incident type).

Figure 2: The Marked Stochastic Point Process

Next we need to characterize the system designs under consideration and the potential defense mechanisms that may be employed within the systems. That is, we need to define the set of designs/architectures of the system, and the defense mechanisms. The combination of a system design and defense mechanism will be called the configuration (or posture). The design could include distributed subsystems with different defenses for the subsystems. Each possible combination of design and defense would be a configuration. In this paper we consider only one design since specific data on system designs were not available. When information exists on different designs, any number of designs may be analyzed. We also assume five hypothetical levels of defense mechanisms, and cost increasing with the strength of the defense. In general, many complex designs and defense mechanisms can exist, and our model can accommodate such complexity whenever the data are available.

The response prediction model will predict the transition of the system to a new state after an attack/incident has occurred, and will be a function of the incident type and the configuration. Thus, given an incident-type j and initial system state r , the subsequent state s may be any one of the set $\{S\}$ of possible states that the system can be in, such as normal, under attack, compromised, recovered, or nonfunctional. The actual states may be different, of course, for different configurations. The transition matrix T will probabilistically map r to s given j . That is, each element of T is the probability of the system of that configuration going to another (possibly compromised) state when subjected to an incident of type j . In general, the incident type j will be a vector of severity level and number of attackers. But as mentioned above, since data on the number of attackers were not available, j is taken to be severity only in the simulations conducted here.

As discussed earlier, survivability is the key issue we wish to investigate with the simulation model. Therefore it is necessary to develop a measurable concept of survivability. There has been considerable work done on survivability in telecommunications [Moitra 97] that is essentially at the network topology level. Here we employ analogous measures of survivability based on concepts suitable to information systems and networks.

Survivability is the degree to which a system has been able to withstand an attack or attacks, and is still able to function at a certain level in its new state after the attack. This new state s will generally be a compromised state, and is the state in which the system ends up before any full-fledged recovery or repairs are done to restore it to its normal state. At the conceptual level, we propose that survivability be measured as

$$\text{SURV} = (\text{performance level at new state } s) / (\text{normal performance level})$$

The main issue is the measurement of performance levels. In telecommunications, it is generally taken as the traffic that is still carried relative to the offered traffic the network could carry under normal conditions. An analogous approach could be taken for computer systems, in that the different functionalities and services could be considered separately, and an assessment could be made as to what extent each functionality has survived in the new system state after an attack. For example, if a given functionality has survived intact, its value would be 1, and if the system were completely nonfunctional with respect to that service, then its value would be 0. Intermediate states would have values in between. Further details are given in the appendix. For additional measures of survivability see Moitra, Oki and Yamanaka [Moitra 97].

4 Data Analysis

We shall use the following notation. Additional notation will be introduced as needed.

- i] i, j = index for incident type, i, j in $\{J\}$. We consider actual, unauthorized incidents only. i denotes the prior incident and j the subsequent (or current) one.
- ii] $P(j)$ = probability that an incident is of type j .
- iii] $\tau(i, j)$ = inter-incident times between incidents i and j .
- iv] a = arrival rate of incidents = $1/\tau$.
- v] r, s = index for system state, r, s in $\{S\}$.
- vi] T = transition probability matrix with elements $\{p(r, s)\}$, where $p(r, s)$ is the probability of going from state r to state s .

The data required for the simulation is as follows:

- the distribution of the τ 's to determine the functional form of $f(t)$
- the parameters for $f(t)$: $\{a\}$. For example, a = arrival rate, a' = trend (if any), etc.
- total number of incident types = J , and a taxonomy of types $\{j\}$
- probabilities $P(j)$ of incidents of each type occurring for all j in J
- number of defense mechanisms and the costs for each. (Costs will be scaled from 1 to 100)
- number of possible states of the system = S . At this stage, we will consider only "long-term" end states, not the "transient" ones that a system may go through when attacked.
- transition probabilities in T : $p(r, s | j)$. This may be obtained from observed data, from expert judgement, or from a (tree) model that estimates the probabilities of end states based on transitions through all possible intermediate states.
- vector of $SURV(s)$ for all (end) states

With the data collected by CERT, we only have information on recorded incidents. We can use this data to estimate (and forecast) the incidents process that an individual network site would experience. However, detailed data on system responses and defense mechanisms are not available to date. Therefore we first focus on estimating the model for the stochastic point process for incidents. Then we simulate the incidents process based on these estimates.

We estimate the following from CERT data:

1. the functional form for inter-event times (f)

2. the parameters of “f”
3. whether $f = f(i,j)$ or $f(j)$ or $f(i)$
4. the stationarity of the form of the distribution function
5. correlations between inter-incident times and incident types, $\tau(i,j)$
6. the stationarity of τ
7. the dependence of τ on (victim) site type
8. correlations between consecutive incident types $\rho(i,j)$

However, there are a number of ambiguities in the data and taxonomy. For example, it is not always clear whether a reported incident is a truly unauthorized event. Also there may be ambiguities in incident identification, MO (method of operation) identification, etc. We should also note that the incidents in data are twice filtered; that is, they are based upon detection and then reporting. Also such recorded data is typically both right-censored and left-censored. That is, data collection starts at some point in time after the process has started, and ends at a point in time while the process is still going on. This creates some biases in the estimation process. To partially overcome this problem we have selected those sites that have experienced at least three incidents, thus providing us with at least two τ 's. Other issues will be discussed in their contexts and also in Section 5 on future research.

In order to simulate the incidents process realistically, we need to estimate the relevant parameters and correlations as delineated above. First we briefly describe the CERT data that is available and then we describe the data analysis.

5 Data Description

The CERT data analyzed here presents reported incidents between 1988 and 1995. For each incident, the variables recorded are: SD (start date), ED (end date), NS (number of sites involved), NM (number of messages), LV (level of the incident), MO (a vector of methods of operation used), CA (corrective action), NT (notes), RS (a vector of reporting sites), and OS (a vector of other sites involved). The data are described in detail in Howard [Howard 95].

For the purposes of this analysis only the following variables were needed:

SD

NS

LV (coded from [1-7])

MO

RS + OS (which were considered together as victim sites).

The attack characteristics are given by {SD, NS, LV, MO}, and the site characteristics are given by the domain identified in RS and OS. There is a likelihood of high correlation between LV and the MOs because the most frequently occurring MOs correspond to one or more “level” categories, as shown below:

Level	MO
1 = root break-in	001
2 = account break-in	002
3 = denial-of-service	017, 221, etc.
4 = corruption of information	several
5 = access attempt	several
6 = disclosure of information	several

Since there is a redundancy in the information contained in “level” and the MOs, we shall use the variable “level” only.

6 Results

We emphasize again that we are looking at the attack-incidents from the point of view of a victim site. We are essentially reconstructing and forecasting the victimization experience of individual sites over time. Also the data analysis reported here is at the exploratory level and further statistical testing will be required for confirmatory analysis in the future.

The first issue we investigate is the functional form of the distribution for the inter-incident times. To this end, we plot histograms of the frequency distributions of the τ 's for different time units. The smallest granularity of time is a day. So Figure 3 (next page) shows the frequencies for $\tau = 0$ days (number of incidents occurring the same day), 1 day (incidents occurring one day apart), 2 days, etc. up to 9 days. We see that the distribution is not very different from an exponential in shape. The mean is at 84 days. The values are given in Table 1 for up to 29 days.

Table 1: *Frequency Distribution of τ by Time Interval D (D in days)*

D	$f(\tau)$	D	$f(\tau)$	D	$f(\tau)$
0	1723	10	303	20	187
1	869	11	292	21	229
2	786	12	291	22	177
3	693	13	333	23	188
4	576	14	304	24	165
5	543	15	256	25	130
6	493	16	207	26	178
7	507	17	227	27	152
8	374	18	191	28	181
9	466	19	248	29	168

We also plot the frequency distribution for time intervals equal to weeks, months and quarters in Figures 4 to 6 respectively. All these figures support the hypothesis of an exponential distribution.

However, the tail of the distribution is rather long, as can be noticed from Table 1. That is, the frequencies decay very slowly. Therefore it might be worth investigating a mixed exponential, or even a triangular distribution for the τ 's in the future. In particular, we note in Figure 3 that there is a slight spike at $D = 0$, which means there is a higher probability of two incidents

occurring the same day than would be predicted by an exponential distribution. This has a substantive interpretation in that if a site is attacked, its vulnerability might be exposed, and this information may attract another attack on the same site almost immediately (within one day). This may be simulated by a “point mass” probability for $\tau = 0$, together with an exponential (or triangular) distribution for all values (including $\tau = 0$). For now we will use the exponential distribution since that is quite close to the actual distributions.

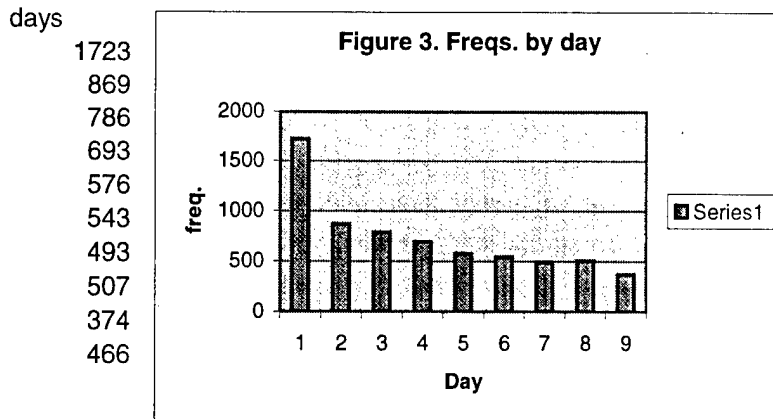


Figure 3: Frequencies by Day

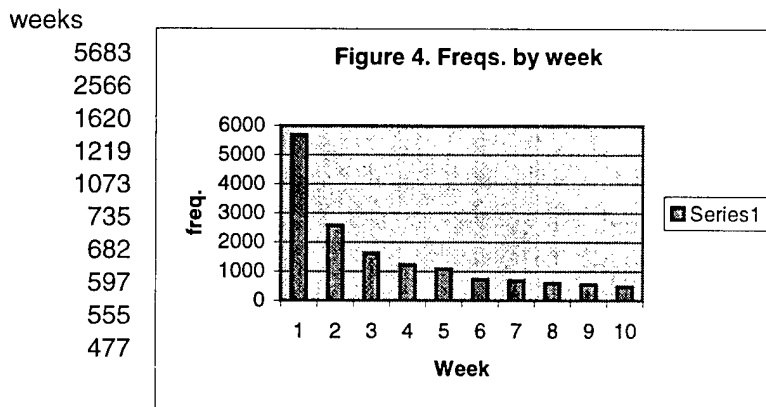


Figure 4: Frequencies by Week

month
 11437
 3070
 1859
 1202
 754
 612
 461
 430
 235
 210




Figure 5: Frequencies by Month

qtr
 16417
 2540
 1124
 577
 371
 250
 176
 90
 70




Figure 6: Frequencies by Quarter

Next we need to examine whether the distribution (that is the shape of the histograms) changes over time (from year to year). Table 2 gives the frequencies by year for each of the eight years for which we have data. Aside from 1988 and 1989, for which we have very little data, all the distributions look exponential and their forms are quite stable. Therefore we will use the exponential form throughout the simulation horizon.

Table 2: Frequency Distributions of τ 's by Year

D	Year							
	88	89	90	91	92	93	94	95
0	16	136	552	968	1828	2904	5525	4488
1	4	67	129	224	321	592	777	426
2	8	34	67	106	161	228	403	117
3	1	6	52	70	100	89	243	16
4	5	11	40	33	82	58	142	0
5	4	3	52	31	67	20	73	0
6	1	4	24	19	67	22	39	0
7	0	2	24	20	29	12	3	0
8	1	2	24	19	12	12	0	0
9	1	13	91	59	21	5	0	0

Having decided upon the distributional form, we now consider the mean inter-incident time, or, equivalently, the arrival rate (τ) of incidents. The first question that arises is whether they are constant over time, or whether they have trends. Table 3 has the average τ 's by year, and it is clear that there is a decreasing trend in the τ 's, which implies an increasing trend in the arrival rates. This should be reflected in the simulation.

Table 3: Average τ 's by Year

Yr	Freq.	Average τ
88	41	243.6
89	278	185.9
90	1055	249.0
91	1549	144.8
92	2688	116.0
93	3942	74.6
94	7205	70.1
95	5047	34.5

The estimated regression equation is

$$\text{Ave. } \tau = 278 - 30 * \text{Year},$$

with the slope coefficient significant at the .001 level. It is important to realize that there could be a bias in the trend value (slope coefficient) due to the data being right-censored. Those τ 's included in the sample that is cut off in Year 8 (1995) will tend to have a disproportionate number of small values in them, since larger values would not be in such a sample. Nevertheless, the data do suggest a consistent trend. The actual trend value should be computed more carefully for future work. For our purposes, we are simply interested in detecting the existence of a trend. In the appendix, we discuss the issues further.

A related question is whether this trend varies by incident type. For example, is the trend for root break-ins different from that of account break-ins? Table 4 gives the average inter-incident times by year and incident type, and we see no significant difference in trend by type. Apart from 1988 and 1989, when we have very little data, all types have a generally decreasing trend, and. The three major types (1, 2, and 5) have a similar trend.

Table 4: Average τ 's by Type (j) and by Year

(j)	Year							
	88	89	90	91	92	93	94	95
1	262.1	224.0	364.1	149.0	129.0	87.6	80.2	41.3
2	199.2	134.1	174.5	158.1	103.4	61.6	53.9	26.5
3	-	-	68.2	4.5	106.9	67.1	62.1	32.0
4	-	22.5	225.9	79.6	106.4	55.5	50.3	33.2
5	-	174.3	145.7	109.4	92.2	58.8	51.1	32.7
6	40.0	-	8.3	240.8	111.1	59.5	52.6	32.7
7	-	217.5	112.6	189.1	136.5	95.1	62.0	28.7

The inter-incident times may in general depend on both the prior incident type and the subsequent type. That is, the time between pairs of incidents may depend on the incident types. To explore this, we show the average $\tau(i,j)$'s in Table 5 where $\tau(i,j)$ is a time interval between an incident of type i and type j . The result is somewhat difficult to interpret because the numbers in many cells are small, and the types are dominated by 1 (root break-in) and 2 (account break-in). The only other type that is relatively frequent is 5 (access attempt). However, it appears that the τ 's do depend on incident type; if confirmed this has important implications because it means that the interval between two incidents is correlated with incident type, particularly with the *next* incident type. In general, it indicates that the past history of attacks at a site may be a predictor of future experience, and thus is important to model. The grand mean is 84 days, which is the order of magnitude of a quarter. The row marginals give the average time *from* a type, and the column marginals give the average *to* a type. Thus the average time from a Type 1 incident (root break-in) to any other type is relatively long at 97 days, which is the same as the average time to a Type 1 incident from any other type. The inter-incident times between Type 1 incidents are also quite long. The average times between Type 1 and Type 2 incidents is about the same (90 and 94 days). However the times between Type 2 incidents is relatively short at 40.5 days. Thus for these two types, both the prior and subsequent types affect the inter-incident times. The average time from a Type 5 incident to Type 1 is 94 days, 69 days to a Type 2 incident, and 46 days to another Type 5 incident. The diagonal elements represent times between the same types of attacks, and they would be expected to be shorter than average. This is true for Types 2 and 5 but not for Type 1. The frequencies of other types of incidents are much smaller than these three, and no firm conclusions can be drawn about them. Type 7 is actually a false alarm, and it is not surprising that its marginals are large, as well as the average time to a Type 1 incident.

Table 5: Average τ (i,j)'s

(i)	j=1	2	3	4	5	6	7	All
1	103.5	90.3	78.3	72.8	85.2	99.6	94.9	97.0
2	93.9	40.5	54.1	67.2	79.5	83.5	58.6	71.0
3	50.8	46.9	35.3	62.5	46.5	57.8	133.0	51.0
4	61.1	57.3	35.2	37.6	44.5	56.8	64.0	53.0
5	93.5	69.0	27.3	66.7	45.6	65.3	64.9	71.0
6	69.1	54.6	44.7	44.5	78.7	37.9	52.0	60.0
7	114.1	84.9	37.3	15.3	57.2	11.9	66.1	88.0
M	97.0	67.0	60.0	63.0	69.0	81.0	79.0	84.1

We also looked at the average τ as a function of the domain type, but for the four major domains (edu, com, net, and gov) there was no significant difference. Finally we looked at the average τ for cases where there was only one site involved, two sites involved, and more than two sites involved. When more than two sites are involved, the time tends to be slightly longer.

Now we turn to the incident types. These are the “marks” of the point process; that is, whenever an incident occurs, there is a type associated with it. Here we take “level” as representing the type of the incident. Alternatively, the MOs might also have been taken as indicating type; and finally, the number of attackers should also be considered as another dimension for incident type. However, the set of MOs is very large and the most common MOs are collinear with “level.” Also, there were no direct data on the number of attackers in the CERT data analyzed here. Therefore we took “level” to be type. The type of incident has very important consequences for the attacked system of course, so it has to be included in the simulation model. The frequencies of and interaction between the types are best analyzed with the contingency table and the type-switch matrix respectively. The contingency table has elements $N(i,j)$ where $N(i,j)$ = the number of pairs of incidents where type j occurred after type i . The type-switch matrix gives the probability of type j occurring next, given type i has occurred. These are shown in Tables 6a and 6b. The last row of Table 6a gives the totals, and the last row of Table 6b gives the marginal probabilities; that is, the fraction of times each type occurred.

Table 6a: Frequency of Incidents by Prior (i) and Subsequent (j) Types

(i)	j=1	2	3	4	5	6	7	All
1	6407	2184	153	127	1401	592	307	11171
2	2212	1797	51	58	588	274	109	5089
3	134	53	7	11	31	23	5	264
4	109	55	11	29	53	13	11	281
5	1350	604	36	37	1024	116	115	3282
6	531	290	26	19	91	184	23	1164
7	278	95	9	9	112	14	37	554
T	11021	5078	293	290	3300	1216	607	21805

Table 6b: Type-Switch Matrix (i)

(i)	j=1	2	3	4	5	6	7
1	0.57	0.20	0.01	0.01	0.13	0.05	0.03
2	0.43	0.35	0.01	0.01	0.12	0.05	0.02
3	0.51	0.20	0.03	0.04	0.12	0.09	0.02
4	0.39	0.20	0.04	0.10	0.19	0.05	0.04
5	0.41	0.18	0.01	0.01	0.31	0.04	0.04
6	0.46	0.25	0.02	0.02	0.08	0.16	0.02
7	0.50	0.17	0.02	0.02	0.20	0.03	0.07
P(j)	0.51	0.23	0.01	0.01	0.15	0.06	0.03

There are a number of points to be made regarding these tables. As far as the frequencies are concerned, we have already noted that Types 1 (root break-in), 2 (account break-in), and 5 (access attempt) dominate the sample. Examining the type-switch matrix, we can see that the marginals are indeed quite different between different types. We also note that the columns generally have the same values except for the diagonal elements, which are larger. Thus the probability of an incident of type j appears to be independent of the previous type (i). Since the types are ordered by severity with 1 being the most severe and 7 the least, the lower triangle below the diagonal represents escalation: from less serious to more serious. But in this sample it is somewhat confounded by the fact that the two most serious types are also the most frequent. However, there still seems to be some evidence of a general escalation in seriousness of incidents but this needs to be verified more rigorously.

Finally, we check whether the marginal probabilities vary with time. Their values by year are given in Table 7, and we can see that they are reasonably stable over time. In particular, after 1989, the values are quite stable.

Table 7: $P(j)$'s by Year

(j)	88	89	90	Year 91	92	93	94	95
1	0.81	0.50	0.46	0.34	0.51	0.48	0.62	0.44
2	0.12	0.33	0.18	0.36	0.26	.027	0.13	0.31
3	0.00	0.00	0.01	0.00	0.00	0.01	0.02	0.02
4	0.00	0.01	0.01	0.01	0.01	0.01	0.01	0.03
5	0.00	0.11	0.31	0.23	0.14	0.15	0.14	0.11
6	0.05	0.00	0.00	0.02	0.06	0.05	0.06	0.06
7	0.00	0.05	0.03	0.03	0.02	0.03	0.02	0.02

This completes the exploratory data analysis needed for estimating the inputs to the simulation model as far as the incidents process is concerned. As noted above, this is the information that can be extracted from the CERT data now available. The rest of the required simula-

tion inputs will be estimated from secondary or tertiary data and expert judgement. In the next section, we proceed with the simulation.

7 Simulation Results

The simulation model has been developed and reported previously [Moitra 00]. Here we run the model with inputs estimated from the CERT data. One extension to the previous simulation model is that trends are included in the incidents process. Also, other changes have been instituted such as three incident-types (root break-in, account break-in, access attempt), and the actual correlations between the inter-incident times and incident types, $\tau(i,j)$. The unit of time is taken as one quarter, since the average τ is 84 days.

Our interest is to observe how well a system survives when subjected to a series of attacks. This will obviously depend on both the severity levels of the attacks as well as the level of defense that is built into the system. The stronger the defense system, the more likely it is to withstand an attack; that is, to stay in its normal state, and the less likely it is to end up in a compromised state. In other words, the transition probabilities of the system are a function of the defense mechanism, and this functional relationship drives the expected survivability of the system in any attack scenario. Therefore simulation was carried out for different probabilities of the attack types, and different relationships between the cost of the defense mechanism and the probabilities of the system ending in the various possible states (from normal to nonfunctional).

Some additional notation will be needed to discuss the simulation results, and we list them below.

a_0 = initial arrival rate of incidents (per quarter),

a' = trend in arrival rate,

$\pi_1, \chi_1, \pi_3, \chi_3$ = parameters determining the system transition probabilities,

SURV(s) = vector of survivability ratings of the states, initially = {1, .8, .6, .4, 0}

surv = expected survivability

dmg = average damage

AVS = average of the survivability ratings of state 2, 3, and 4.

A large number of simulations can be carried out with our model to investigate a wide variety of issues related to managing survivability, since we can observe the impact of any model parameter on the system survivability. Another quantity of interest is the "average damage caused per unit time." This is computed by taking the total damage ($= \sum (1 - \text{survivability})$) over all episodes) and dividing by the total time elapsed during the simulation. In this paper, we present some of the possible sensitivity analyses to illustrate what can be done. In the absence

of data for the transition probabilities, the model described in the appendix was used to generate them, and values for survivability ratings were assumed. The results are presented in Tables 8 to 17 where the survivabilities are given as fractions. First we investigate the impact of varying the relative probabilities of the serious and mild incidents. The results are given in Table 8.

Table 8: Expected Survivability/Average Damage and P(j)

$a_0 = 1.0, a' = .001, \pi_1 = .15, \chi_1 = .008, \pi_3 = .25, \chi_3 = .075$

cost	P(1)=.55		P(1)=.65	
	surv	dmg	surv	dmg
5	0.704	0.473	0.6908	0.479
10	0.7364	0.421	0.7192	0.434
25	0.7686	0.369	0.7512	0.385
50	0.7868	0.34	0.7708	0.354
75	0.8056	0.31	0.79	0.324
100	0.8224	0.283	0.8094	0.294

While the survivability increases with the cost of the defense mechanism as expected from the relation of the transition probabilities to cost, the survivability does not appear to decrease significantly with increases in the probability of occurrence of serious incidents. This is somewhat surprising, and this particular result is most likely related to the method of generation of the $p(r,s)$'s in T. This method does not vary the $p(r,s)$'s very much with j . With some other set of $\{p(r,s)\}$, we may well find greater sensitivity of survivability to the $P(j)$'s since there is a nonlinear relationship between them.

Table 9 shows how survivability changes with the parameter π_1 , which determines $p(1,1|m)$, the probability of remaining normal under attack given a defense mechanism m .

Table 9: Expected Survivability/Average Damage and π_1

$P(1)=.55, a_0 = 1.0, a' = .001, \chi_1 = .008, \pi_3 = .25, \chi_3 = .075$

cost	$\pi_1 = .1$		$\pi_1 = .15$		$\pi_1 = .2$	
	surv	dmg	surv	dmg	surv	dmg
5	0.6848	0.504	0.704	0.473	0.719	0.449
10	0.722	0.444	0.7364	0.421	0.748	0.402
25	0.7628	0.379	0.7686	0.369	0.7702	0.367
50	0.7864	0.341	0.7868	0.34	0.7868	0.34
75	0.8054	0.31	0.8056	0.31	0.8056	0.31
100	0.8224	0.283	0.8224	0.283	0.8224	0.283

At low cost levels survivability increases slightly and damage decreases slightly as survivability ratings rise. However, both become stationary at high cost levels, indicating that survivability has most likely saturated.

Table 10 gives the survivabilities as χ_1 varies. χ_1 determines $p(1,slm)$ for $s>1$; that is, the probabilities of going to compromised states, including the nonfunctional state.

Table 10: Expected Survivability/Average Damage and χ_1

$P(1)=.5, a_0 = 1.0, a' = .001, \pi_1 = .15, \pi_3 = .25, \chi_3 = .075$

cost	$\chi_1 = .006$		$\chi_1 = .008$		$\chi_1 = .010$	
	surv	dmg	surv	dmg	surv	dmg
5	0.703	0.475	0.704	0.473	0.7054	0.471
10	0.7352	0.423	0.7364	0.421	0.7382	0.418
25	0.764	0.377	0.7686	0.369	0.7722	0.364
50	0.7794	0.352	0.7868	0.34	0.7968	0.324
75	0.7932	0.33	0.8056	0.31	0.8186	0.289
100	0.8056	0.31	0.8224	0.283	0.8402	0.255

Here the variation of survivability with χ_1 is slight at low cost values but is significantly higher at high cost values. Overall it increases with cost and χ_1 as we would expect.

Tables 11 and 12 show the effect of varying π_3 and χ_3 respectively.

Table 11: Expected Survivability/Average Damage and π_3

$P(1)=.55, a_0 = 1.0, a' = .001, \pi_1 = .15, \chi_1 = .008, \chi_3 = .075$

cost	$\pi_3 = .20$		$\pi_3 = .25$		$\pi_3 = .3$	
	surv	dmg	surv	dmg	surv	dmg
5	0.7094	0.464	0.704	0.473	0.6992	0.481
10	0.7438	0.409	0.7364	0.421	0.7332	0.426
25	0.774	0.361	0.7686	0.369	0.7628	0.379
50	0.7944	0.328	0.7868	0.34	0.7818	0.348
75	0.8114	0.301	0.8056	0.31	0.7994	0.32
100	0.8292	0.272	0.8224	0.283	0.8172	0.291

π_3 determines the levels of the transition probability $p(1,1)$ as cost changes. Thus the higher the value of π_3 , the higher the chances that the system will stay in the normal state; thus the survivability will be higher. This is what we observe from Table 11, where survivability de-

creases with π_3 , and we also notice that the impact is relatively greater at lower costs than at higher costs.

Table 12: *Expected Survivability/Average Damage and χ_3*

$P(1)=.55, a_0 = 1.0, a' = .001, \pi_1 = .15, \chi_1 = .008, \pi_3 = .25.$

	$\chi_3 = .070$		$\chi_3 = .075$		$\chi_3 = .080$	
	surv	dmg	surv	dmg	surv	dmg
5	0.7094	0.464	0.704	0.473	0.6992	0.481
10	0.7438	0.409	0.7364	0.421	0.7332	0.426
25	0.774	0.361	0.7686	0.369	0.7628	0.379
50	0.7944	0.328	0.7868	0.34	0.7818	0.348
75	0.8114	0.301	0.8056	0.31	0.7994	0.32
100	0.8292	0.272	0.8224	0.283	0.8172	0.291

Table 12 shows the variations in expected survivability and average damage with χ_3 .

χ_3 determines the levels of the transition probabilities $p(1,s)$ for $s > 1$, that is, the compromise probabilities. Thus a (slightly) higher value of χ_3 leads to lower values of survivability as should be expected. The relative impact is not insignificant, since the change in χ_3 is very small, and the impact is constant over the values of cost.

Table 13: *Expected Survivability/Average Damage and AVS*

$P(1)=.55, a_0 = 1.0, a' = .001, \pi_1 = .15, \chi_1 = .008, \pi_3 = .25.$

cost	AVS = .60		AVS = .37		AVS = .30	
	surv	dmg	surv	dmg	surv	dmg
5	0.704	0.473	0.5489	0.722	0.5045	0.793
10	0.7364	0.421	0.601	0.638	0.5602	0.704
25	0.7686	0.369	0.6495	0.56	0.6131	0.619
50	0.7868	0.34	0.6772	0.516	0.644	0.569
75	0.8056	0.31	0.705	0.471	0.6742	0.521
100	0.8224	0.283	0.7308	0.43	0.7024	0.476

Next we investigate the effect of an increase in the arrival rate. If the rate of arrivals of incidents is increased, this simply amounts to accelerating the time scale, and expected survivability remains the same. This is because expected survivability is measured per incident, and increasing the number of incidents makes no difference to this measure. However, the *average damage done* changes because this does depend on time, and this can be seen in Table 14.

Table 14: Average Damage and Arrival Rate a_0

$P(1)=.55, a' = .001, \pi_1 = .15, \chi_1 = .008, \pi_3 = .25, \chi_3 = .075$

cost	$a_0=0.75$	$a_0=1.00$	$a_0=1.25$
5	0.548	0.473	0.402
10	0.488	0.421	0.358
25	0.428	0.369	0.314
50	0.394	0.34	0.29
75	0.359	0.31	0.264
100	0.328	0.283	0.241

A similar situation arises when considering the impact of a possible correlation between the arrival rates and the incident type. So far we have not assumed any such correlation, but Table 5 suggests that it does exist. In fact, τ appears to depend on both the prior and the subsequent incident types. Therefore, we included that effect in our simulation and the arrival rates were adjusted according to the data in Table 5. Again, the expected survivability does not change, because the (less serious) incidents simply happen faster and the system responds in the same way. However, the average damage done increases, and would have increased even more if the more serious incidents had occurred faster (instead of the less serious ones). In this case, the damage done saturates with higher defense levels, and does not increase. The results are shown in Table 15.

Table 15: Average Damage and a'

$P(1)=.55, a_0 = 1.0, \pi_1 = .15, \chi_1 = .008, \pi_3 = .25, \chi_3 = .075$

cost	a'		
	0.00075	0.001	0.00125
5	0.449	0.473	0.495
10	0.4	0.421	0.441
25	0.351	0.369	0.386
50	0.323	0.34	0.356
75	0.295	0.31	0.324
100	0.269	0.283	0.296

The tables above show the absolute changes in expected survivability when some parameter is varied. To fully understand the impact of a parameter on expected survivability, we need to examine the relative changes. These are reflected in the elasticities, which give the percent changes in expected survivability when the parameters are varied by 100 percent. Thus these relative changes give a more accurate measure of the impacts of the parameters and allow the impacts to be compared. The relative changes or elasticities are given in Table 16.

Table 16: Relative Changes in Expected Survivability with Respect to Parameters

Parameter	P(1)	π_1	π_3	χ_1	χ_3	SURV(s)
Relative change in survival	-0.123	0.002	0.104	0.036	-0.128	.317

Table 16 confirms what we noticed before with respect to the insensitivity of survivability to P(1). Survivability appears to be most sensitive to π_3 and χ_3 . That is, the initial level of the transition probabilities is most important, rather than how they change with m.

Table 17: Relative Changes in Average Damage with Respect to Parameters

Parameter	a_0	a'	SURV(s)
Relative change in average damage (η)	0.556	0.179	-0.883

The relative changes or elasticities of average damage with respect to the parameters a_0 , a' and SURV(s) are given in Table 17. Here we see that the signs are in the expected directions, and average damage is quite sensitive to a_0 and SURV(s).

In the above simulations, we have assumed a Poisson process for the incidents, which is a flexible model and commonly used in point processes. However, any other distribution may be used in the model, and the distribution of the τ 's at time intervals of one day suggests that perhaps a mixed distribution should be tried. This distribution would be as follows.

After a τ has been generated, with some probability, say β , the next τ will be set to 0. Otherwise (with probability $1-\beta$) another random τ will be generated for the time to the next incident. The parameter β reflects the probability that the first attack spurs another attack soon afterwards (in one day, since that is the granularity of our data).

Also, the mixed exponential may be investigated. For example, a mixture of two exponentials may be reasonable. This could arise from there being two types of attackers (amateur and experienced), each with their own rate of attacking. Then $\{f(t) = f(t; a_1, a_2, \alpha)\}$, where a_1 may be the aggregate attack rate for amateurs, a_2 the rate for experienced attackers, and α the proportion of amateurs to experienced attackers.

The above results are just a small subset of all the possible analyses that can be done with this simulation model but they demonstrate the potential of this model and this approach. Any incidents process can be generated, and any system-response may be inserted in the model through the transition matrix T. Thus we can investigate the survivability and the damage done for any scenario for any set of defense mechanisms. Given the costs of these mechanisms, we can derive a survivability/cost function as shown below, and achieve a cost-effective level of security.

In Figure 7, we have plotted the expected survivability against cost for $P(1)=.55$ and $P(1)=.65$ with other parameter values as in Table 1. The plot shows the relationship between cost and survivability. As cost increases, survivability increases rapidly at first, and then more slowly. Such a plot can provide a systems manager with the ability to make an informed decision about the level of defense that is most appropriate for his or her organization since it shows the tradeoff involved between cost and expected survivability.

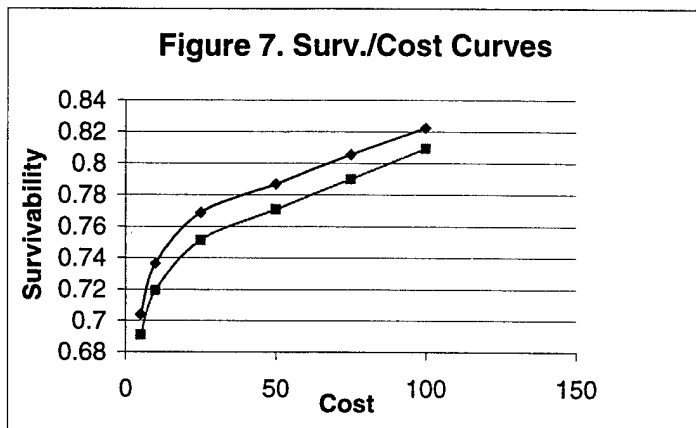


Figure 7: Survivability/Cost Curves

When survivability is not critical, the organization may choose a lower point on the tradeoff curve, but when survivability is critical, the organization may well choose a point higher up on the curve. In the case when the “indifference curve” can be estimated, we can actually choose an optimal or “best” point on the curve. However, even if we are not aiming for optimality, we can still use the curve to find the most appropriate point in the tradeoff between cost and survivability.

8 Future Research

There are many areas in which further work would be useful. They can be classified into the following groups: further data analysis, modeling, and the value of security to business.

In the area of data analysis, it would be worthwhile to analyze more recent data to identify current trends. More statistical testing should be done to follow up on the exploratory analysis done here. In particular, standard deviations and confidence intervals need to be estimated. The wide variation in the inter-incident times should also be analyzed. For example, there might be a difference associated with long intervals versus short intervals. Generally, further pattern analysis should be done to find out which kinds of incidents different sites/systems (i.e., organizations/environments) are likely to experience. For example, do certain sites "attract" certain attacks (such as abuse of trust relationships)? We can investigate such questions by mapping patterns of incidents to sites, and also by some other methods such as cluster analysis. The various possible interactions among the variables in the data may be explored through techniques like data mining. Ideally we would like a model of *patterns of incidents* that could be forecast, so that managers could use the information to decide on appropriate defenses. The simulation model should also be run to conduct more sensitivity analysis and to understand the interactions among the parameters.

Further modeling would provide a better understanding of the whole process. For example, one might consider simulating a pattern where the "end" state of a system influences the time to the next attack. This is plausible if success in penetrating a system instigates another attack sooner than otherwise. However, further modeling will require more detailed data. The data that is needed to better model and understand survivability has been outlined in another report [Moitra 00]. In particular, we should try alternative models for the incidents process, get a better understanding of the transitions of systems when under attack, and develop a viable measure for survivability. We need good and reliable metrics for security and survivability of network systems.

This model should eventually be embedded into a DSS that IS managers could use to manage the security and survivability of their ISs. For this, we need to understand how managers in organizations view their information systems and what is valuable or critical to them. It would be extremely useful to conduct a survey of managers in various organizations and to model their decision-making processes. A number of methodologies in decision analysis and operations research exist which could be utilized for such applications.

Bibliography

- [Baker 95] Baker, R.H. *Network Security: How to Plan for it and Achieve it*. New York, NY: McGraw-Hill, 1995.
- [Basawa 80] Basawa, I.V. & Prakasa Rao, B.L.S. *Statistical Inference for Stochastic Processes*. New York, NY: Academic Press, 1980.
- [Bernstein 96] Bernstein, T. (Ed.) *Internet Security for Business*. New York, NY: John Wiley and Sons, 1996.
- [Boni 99] Boni, W.C. & Kovacich, G.L. *I-Way Robbery: Crime on the Internet*. London, England: Butterworth-Heinemann, 1999
- [Cameron 98] Cameron, D. *E-commerce Security Strategies: Protecting the Enterprise*. Charleston, NC: Computer Technology Research Corporation, 1998.
- [Cohen 99] Cohen, F. *Simulating Cyber Attacks, Defenses, and Consequences*. Livermore, CA: Fred Cohen & Associates, 1999.
- [CSI 98] Computer Security Institute. *Computer Security Issues and Trends*. 4, 1 (Winter 1998).
- [Daley 88] Daley, D.J. & Vere-Jones, D. *An Introduction to the Theory of Point Processes*. New York, NY: Springer-Verlag, 1988
- [Ellison 97] Ellison, R.J.; Fisher, D.A.; Linger, R.C.; Lipson, H.F.; Longstaff, T.; & Mead, N.R. *Survivable Network Systems: An Emerging Discipline*. (CMU/SEI-97-TR-013 ADA 341963) Pittsburgh, PA: Software Engineering Institute, Carnegie Mellon University, 1987. Available WWW<URL <http://www.sei.cmu.edu/publications/documents/97.reports/97tr013/97tr013abstract.html>>
- [Fisher 99] Fisher, D.A. "Emergent Algorithms—A New Method for Enhancing Survivability in Unbounded Systems," *IEEE Proceedings of the Hawaii International Conference on Systems Sciences*. Wailea, HI, Jan. 5-7, 1999. New York, NY: IEEE Computer Society Press, 1990.

- [Gollman 99]** Gollman, D. *Computer Security*. New York, NY: John Wiley and Sons, 1999.
- [Howard 95]** Howard, J. *An Analysis of Security Incidents on the Internet (1989-1995)*. Ph.D. Dissertation, Carnegie Mellon University, Pittsburgh, PA, 1995.
- [Howard 98]** Howard, J. & Longstaff, T. *A Common Language for Computer Security Incidents*. (SAND98-8667) Livermore, CA: Sandia National Laboratories, 1998.
- [Law 82]** Law, A.M. & Kelton, W.D. *Simulation Modeling and Analysis*. New York, NY: McGraw-Hill, 1982.
- [Lilien 92]** Lilien, G.L.; Kotler, P.; & Moorthy, K.S. *Marketing Models*. Englewood Cliffs, NJ: Prentice Hall, 1992.
- [Linger 98]** Linger, R.C.; Mead, N.R.; & Lipson, H.F. *Requirements Definition for Survivable Network Systems*. Proceedings of the International Conference on Requirements Engineering, Colorado Springs, CO: April 6-10, 1998. New York, IEEE Computer Society Press. Also published Pittsburgh, PA: Software Engineering Institute, Carnegie Mellon University, 2000. Available WWW<URL <http://www.sei.cmu.edu/programs/nss/icre.html>>
- [Moitra 97]** Moitra, S.D.; Oki, E.; & Yamanaka, N. *Some New Survivability Measure for Network Analysis and Design*. IEICE Transactions on Communications, *E80-B*, 4, April 1997.
- [Moitra 00]** Moitra, S.D. & Konda, S. *A Simulation Model for Managing Survivability of Networked Information Systems*. (CMU/SEI-00-TR-0020) Pittsburgh, PA: Software Engineering Institute, Carnegie Mellon University, 2000.
- [Snyder 91]** Snyder, D.L. & Miller, M.I. *Random Point Processes in Time and Space*. New York, NY: Springer-Verlag, 1991.

Appendix

Details of the Simulation Model and Survivability Measures

We present the notation again.

- i] i, j = index for incident type, i, j in $\{J\}$. We consider actual, unauthorized incidents only. i denotes the prior incident and j the subsequent (or current) one.
- ii] $P(j)$ = probability that an incident is of type j .
- iii] $\tau(i, j)$ = inter-incident times between incidents i and j .
- iv] a = arrival rate of incidents = $1/\tau$.
- v] r, s = index for system state, r, s in $\{S\}$.
- vi] d = index for system design, d in design space $\{D\}$.
- vii] m = index for defense mechanism, m in the set $\{M\}$.
- viii] configuration = design \times mechanism in configuration space $\{D \times M\}$.
- ix] T = transition probability matrix with elements $\{p(r, s)\}$, where $\{p(r, s)\}$ possibly being functions of i, j, d, m .
- x] l = (victim) sites, l in $\{L\}$.
- xi] $h(l)$ = index for incidents at individual site l : $h(l) = 1, 2, 3, \dots$
- xii] $H(l)$ = total number of incidents at site l .
- xiii] $t(h(l), l)$ = time of h -th incident at l
$$= \sum_{k=1}^{k=h} \tau(k), \text{ where } \tau(k) = t(k) - t(k-1).$$
- xiv] n = number of simultaneous attacking sites in an incident.
- xv] $g(n | v)$ = probability density function for n with parameter v .

In order to forecast incidents, we model the process as a marked, stochastic point process, where the incidents are the events that occur at random points in time, and the event type is the mark associated with an incident [Snyder 91]. The mark is used to identify random quantities associated with the point it accompanies. As shown in Figure 2, each occurrence time t_k of the k^{th} incident in a temporal point-process has a mark j_k associated with it, where j_k will have values in a specified space. The mark, or event type in our case, has to take into account the severity of the incident and the possibility of single, or multiple and simultaneous attacks. This is because we are modeling a process that is taking place in an unbounded environment

[Ellison 97]. Therefore the mark space will be two-dimensional, characterized by type (severity) and number-of-attackers. That is, it will be in the $\{J \times N\}$ space. Although this 2-D marked point process model was developed, no data on the distribution of the number of attackers per incident were available, so only a 1-D mark space with severity was used in the simulations.

A stochastic point process can generally be represented as $\{x(t): t \in T\}$, that is, as a family of random variables indexed by a parameter t that takes values in a parameter set T called the index set of the process. In our case, t represents time, and since T is a subset of \mathbb{R} , it is a continuous-parameter process. For the purposes of this analysis, we limit our attention to the probability density function of the "inter-incident times" (τ 's) which we denote by $f(t)$. That is,

$$f(t) = \Pr\{t \leq \tau \leq t + dt\}.$$

When the process is Poisson, the density function is given by

$$f(t) = a * e^{-at}$$

where a is the rate of occurrence of incidents, and the distribution function is given by

$$F(t) = 1 - e^{-at}.$$

We should note here that the incidents recorded in the data are twice filtered: that is they are conditional upon detection, and then, reporting. Also, the data are doubly censored data, that is, both right and left censored. This means that the process had already started before data collection began, and the process had not finished when data collection was stopped. Censoring may introduce biases in parameter estimates, and it is important to take note of this. In the simulation we included a trend and a correlation between τ and j . This was done by having

$a = (a_0 - a' * t) * \rho(i,j)$, where $\rho(i,j)$ represents the correlation factor which is a function of both i and j .

Next we define the design/architecture space $\{D\}$ of the system, and the defense mechanism state space $\{M\}$. The combination of a system design and defense mechanism will be called the configuration (or posture) space, $\{D \times M\}$.

The response prediction model predicts the transition of the system to a new state after an attack/incident has occurred, and will be a function of the incident type and the configuration, or $p(r,s) = p(r,s | j,d,m)$. The transition matrix T will probabilistically map r to s given j, d, m . That is, each element of T is the probability of the system of design d and defense mechanism m going to another (possibly compromised) state when subjected to an incident of type j . In general, the incident type j could be a vector of any number of characteristics of the incident.

We assume the following structure for T. Without any loss of generality, we assume that the states are ordered by degree of compromise, that is, from $s = 1 = \text{normal}$ (totally functioning) to $s = S = \text{(totally) nonfunctional}$. Given an incident, the system can never go to a "better" state; therefore the lower triangle below the diagonals will have structural zeros as shown below.

$$\begin{pmatrix} p_{11} & p_{12} & p_{13} & p_{14} & p_{15} \\ 0 & p_{22} & p_{23} & p_{24} & p_{25} \\ 0 & 0 & p_{33} & p_{34} & p_{35} \\ 0 & 0 & 0 & p_{44} & p_{45} \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix}$$

We also impose the following constraints on the elements of T, $\{p(r,s)\}$, in terms of their dependence on s, j, and m.

$p(r,s) \downarrow s, \forall s > r$, holding j, m constant, that is, same severity level and same defense;

this implies graceful degradation: the probability of going to a much worse state is lower than going to a slightly worse state.

$p(r,s) \uparrow r, \forall s > r$, holding j, m constant, that is, same severity level and same defense;

vulnerability increases with level of degradation.

Assuming that the j's are ordered from most severe to least severe,

$p(1,1) \uparrow j$, holding m constant, that is, same defense level;

probability of staying normal is higher if the incident is less severe.

$p(1,s) \downarrow j, \forall s > 1$, holding m constant, that is, same defense level;

probability of degradation is lower if the incident is less severe.

$p(r,s) \downarrow m, \forall s > r$, holding j constant, that is, same severity level;

probability of degradation is lower if the defense is stronger.

$p(r,r) \uparrow m, \forall r$, holding j constant, that is, same severity level;

probability of staying in the same state and not degrading is higher if the defense is stronger.

$p(r,s) \uparrow n, \forall s > r$, holding all else constant;

probability of degradation increases with the number of attackers.

Finally, $\sum_s p(r,s) = 1, \forall r$. implies that the system must end up in some state or other.

Currently, no reliable data are available on the times to transition to different states, or the time to fully recover. The CERT data indicate that the mean time between incidents *at a site* is greater than one month. Since it may be reasonably expected that recovery times will be shorter than that on the average [Cohen 98], in these simulations we have assumed that the system would always fully recover before the next incident occurred. So the initial state r was always set equal to 1. However, the model includes the possibility of the system still being in a compromised state when the next incident occurs. We can simulate these conditions given data on system transition times.

The $p(r,s)$'s could be estimated from observations of the responses of actual systems to attacks. A simple representation of the successive stages of compromise that a system might undergo is given in Figure 8. This is based on the model of the three distinct phases on intrusion: penetration, recovery, and exploitation [Ellison 97].

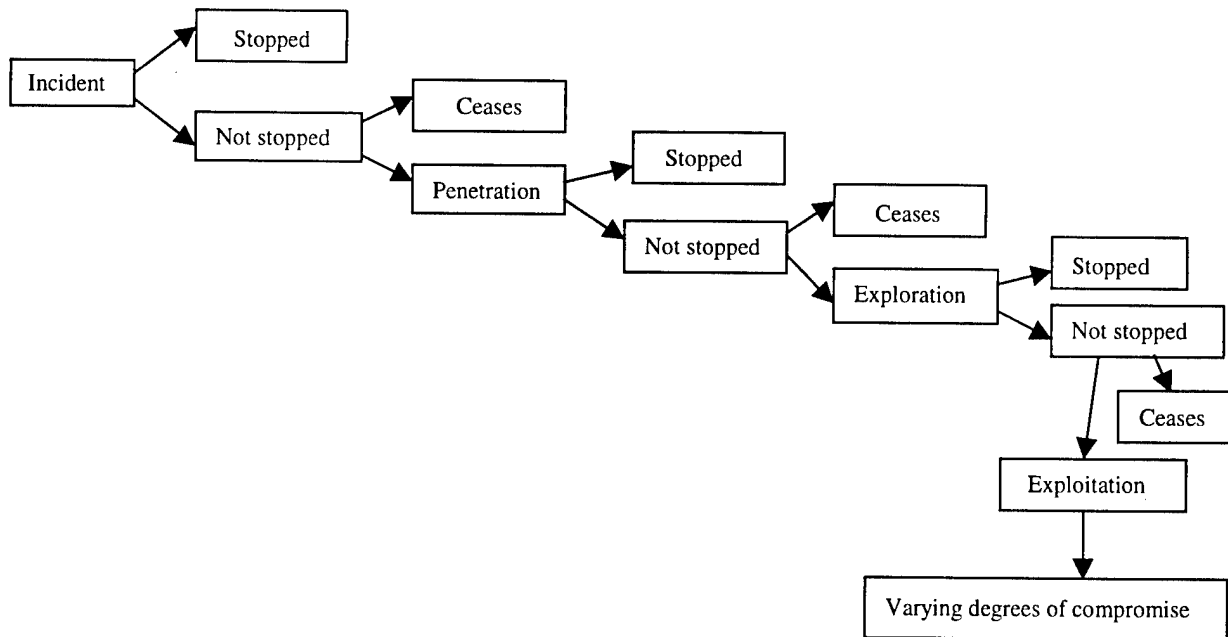


Figure 8: Successive Stages of Compromise

If we had data on the times these transitions occurred, we could estimate the $p(r,s)$'s. In the absence of data, we developed a model to generate the $p(1,s)$'s, such that

$$p(1,s) = p(s, j, \text{cost}(m); \pi_0, \chi_0, \pi_1, \chi_1, \pi_2, \chi_2).$$

There are two cases, $s = 1$ and $s > 1$.

$$p(1,1) = \pi_2 * (1 - e^{-\pi_1 (\text{cost}(m) - \pi_0)}) \quad \text{for } s=1, \text{ and}$$

$$p(1,s) = \chi_2 * (e^{-\chi_1 (\text{cost}(m) - \chi_0)}) \quad \text{for } s > 1.$$

These are simple but commonly used functional forms that are concave and convex respectively, and so reflect decreasing returns with cost. π_1 and χ_1 are the critical shape coefficients that determine the relationship of the transition probabilities with the cost of the defense mechanisms $\text{cost}(m)$. This in turn determines how the survivability varies with cost.

$\pi_2 = \pi_2(j)$ which is modeled as a linear function $= \pi_3 * j$, and

$\chi_2 = \chi_2(j,s) = \chi_3 * ((6-s) - (.4*j))$, again linear in s and j .

The scale coefficients π_3 and χ_3 as well as the constants were calibrated to give reasonable values of the transition probabilities subject to all the restrictions given above. The location coefficients π_0 and χ_0 were set to 0, and $\pi_1, \chi_1, \pi_3, \chi_3$ were varied during the simulation runs.

We measure survivability as

$$\text{SURV} = (\text{performance level at new state } s) / (\text{normal performance level})$$

The main issue is the measurement of performance levels. If a given functionality has survived intact, the value of its performance level could be set to 1. If the system is completely nonfunctional with respect to that service, then its value could be 0. Intermediate states would have values in between. Let $\varphi(s,k)$ be degree to which the compromised function/service k has survived in state s , and let $w(k)$ be the importance level of function/service. Then one possible measure of survivability might be in the form of a weighted sum:

$$\text{SURV}(s) = \sum_k w(k) * \varphi(s, k)$$

This assumes that a complete set of states $\{S\}$ of the system has been defined, and that a systems analyst or IS manager can assess $\varphi(s,k)$ for each s and k . In view of the data requirements, it may be necessary to aggregate the state space $\{S\}$, and the different functionalities and services $\{K\}$. The states in $\{S\}$ may be {normal, under attack, compromised, re-

covered, nonfunctional}, for example, or {normal, minor compromise, significant compromise, very serious compromise, nonfunctional}. Then $\varphi(s,k)$ could be the average level to which function or service k survives in each of those states s . This is a flexible approach, and can be applied in many situations. For example, there might be a particular function that an organization values very highly (such as protecting the confidentiality of a database in a financial services company). Then the weight on this would be very high and also the survivability of this function could be rated low even for a slight compromise. Then any defense mechanism that protected this function would give a high expected survivability, and thus a high benefit, while a defense that did not protect this function would give very low value for expected survivability, and thus very low benefits.

This is a standard multicriteria approach to assessing survivability. While this approach has been used widely, there can be difficulties and biases associated with such a measure. These can be mostly overcome through careful analysis. The weights $w(k)$ are such that

$$0 \leq w(k) \leq 1, \text{ and } \sum_k [w(k)] = 1;$$

The $\varphi(s,k)$'s may also be normalized measures $0 \leq \varphi(s,k) \leq 1$. Then $SURV(s)$ will be between 0 and 1, where 0 means total failure and 1 means completely normal.

Additional Data Analysis

This paper has presented analyses of incidents data collected by CERT. The body of the paper has the analyses relevant to estimating the parameters of the simulation model. Some additional data analysis was done and it is reported here since it provides a more complete understanding of the CERT data.

A. Summary Statistics

Starting dates = [1988 – 1995];

Number of sites per incident = [1 – 1699];

Modes of Operation [22 types reported more than a hundred times];

Top 5 MOs = [root break-in (1188), login attempt (1131), account break-in (865),

password file (598), and password cracking (450)];

Number of unique sites in the data set = 6684;

Number of sites with at least 3 incidents = 1818;

Maximum value of $\tau = 2056$ days;

Total number of inter-event times in sample of sites with at least 3 incidents = 21805.

Table 18 gives the frequency distribution of the number of incidents that individual sites experienced. Thus 3891 sites had experienced only one incident, 975 sites had experienced two incidents, and so on.

Table 18: Frequency Distribution of Number of Incidents per Site

<u>H</u>	<u>f(H)</u>
1	3891
2	975
3	481
4	280
5	180
6	126
7	94
8	72
9	56
10	43

The maximum number of incidents at any one site is 1675. The numbers of incidents experienced by the next top nine sites are [153, 155, 160, 167, 177, 206, 222, 321, 458].

B. Frequency Distribution of Number of Incidents per Day

In the paper the focus was on the inter-incident times. However, the count data is also important for understanding stochastic point processes. Count data refers to the number of occurrences within each unit of time. Here the unit of time is one day, so we would be interested in how many incidents occur each day. These counts will have a frequency distribution, and this is given in Table 19. We can see that there are 1298 days in which nothing occurred. There were 8 days when only one incident occurred, 342 days when two incidents occurred. These are for all sites in the sample. We see that there is high volatility in the data. However, if we smooth the data out, and consider the number of days when one or two incidents occurred, the number of days when three or four incidents occurred, and so on, the distribution looks much smoother. This is shown in Figure 9 and appears approximately Poisson. The great disparity between the frequency of one incident per day and two incidents per day warrants further study.

Table 19: Frequency Distribution of Number of Incidents per Day

Number of Incidents per day	Frequency
0	1298
1	8
2	342
3	86
4	197
5	95
6	111
7	78
8	68
9	49

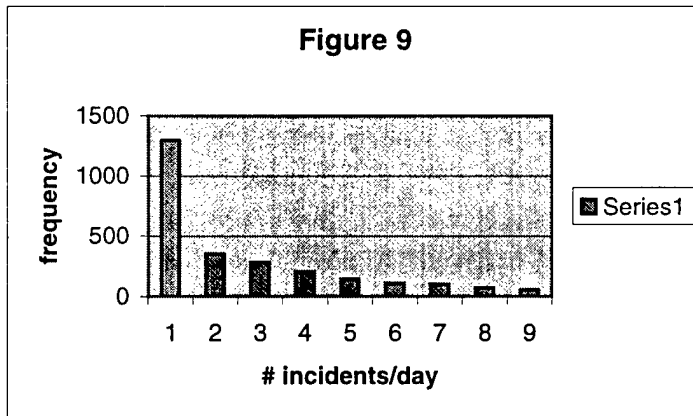


Figure 9: Number of Incidents per Day

C. Average τ Disaggregated

Next we present Average values of τ disaggregated by domain and the number of sites involved. The major domains were {com, net, gov, edu} and the frequencies of occurrences of the τ 's are also given.

Table 20: Average τ by Domain

Domain	Frequency	Average τ
edu	10158	64
com	2931	101
net	916	76
gov	881	79

Table 21: Average τ by Number of Sites

No. of sites	Frequency	Average τ
1	53	63.3
2	4552	63.7
3 or more	17205	89.5

The average time between incidents is longest for "com" and shortest for "edu." It is also relatively long for incident-pairs where the first incident involved three or more sites.

Finally, we present another view of the trends in average τ . Table 22 shows the trend in the τ 's when they are disaggregated by the time of the next incident. This removes the bias in the τ 's when they are disaggregated by the time of the previous incident. In that case, we had noted that there might be a bias in the trend because shorter τ 's would be disproportionately frequent in the later years. Now we see that there probably was a bias, and now the τ 's appear stationary; that is, constant over time. However, this procedure also has its own bias, and we tentatively conclude that there is a small trend towards shorter but it is not as strong as it appears to be from Table 3.

Table 22: Average τ by Year (Year 1 = 1988)

Year	Freq.	Ave. τ
1	13	11.1
2	147	74.1
3	750	73.7
4	1337	99.7
5	2431	90.4
6	3647	97.2
7	7165	77.4
8	6314	80.0
9	1	117.0

REPORT DOCUMENTATION PAGE	<i>Form Approved</i> <i>OMB No. 0704-0188</i>
----------------------------------	--

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (LEAVE BLANK)	2. REPORT DATE December 2000	3. REPORT TYPE AND DATES COVERED Final
4. TITLE AND SUBTITLE The Survivability of Network Systems: An Empirical Analysis	5. FUNDING NUMBERS C — F19628-95-C-0003	
6. AUTHOR(S) Soumyo D. Moitra, Suresh L. Konda	8. PERFORMING ORGANIZATION REPORT NUMBER CMU/SEI-2000-TR-021	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Software Engineering Institute Carnegie Mellon University Pittsburgh, PA 15213	9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) HQ ESC/XPB 5 Eglin Street Hanscom AFB, MA 01731-2116	
10. SPONSORING/MONITORING AGENCY REPORT NUMBER		

11. SUPPLEMENTARY NOTES	
12.A DISTRIBUTION/AVAILABILITY STATEMENT Unclassified/Unlimited, DTIC, NTIS	12.B DISTRIBUTION CODE

13. abstract (maximum 200 words)

This report presents an extended analysis of CERT Coordination Center® incidents data (from 1988 to 1995) and applies the results to simulate attacks and their impacts on network sites. The data were "sanitized" prior to the analysis to ensure complete anonymity. A model for the incidents process is discussed and extended. It consists of three parts: a stochastic process for the random occurrence of incidents at sites, a model for the state transition process for an attacked system given a level of defense, and a method of estimating the expected survivability of the system given possible degradations due to these attacks. This approach leads to the estimation of a survivability/cost function, which shows the tradeoffs involved between cost and system survivability. IS managers can use this to determine the most appropriate level of defense for the network systems of their organizations.

The stochastic process was simulated based on parameter values obtained from actual reported data. Extensive sensitivity analyses are reported that indicate how expected survivability would change with varying parameter analysis results values. The report concludes with a discussion of future work to be done and the appendix has details of the simulation model and further data.

® CERT and CERT Coordination Center are registered in the U.S. Patent and Trademark Office.

14. SUBJECT TERMS survivability, survivability measures, network systems, stochastic point process, transition probabilities, defense mechanisms, incident types	15. NUMBER OF PAGES 53
--	----------------------------------

16. PRICE CODE			
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. Z39-18
298-102