Chapter 4

MODELING AND MANAGING RISK IN BILLING INFRASTRUCTURES

Fabrizio Baiardi, Claudio Telmon and Daniele Sgandurra

Abstract This paper discusses risk modeling and risk management in information and communications technology (ICT) systems for which the attack impact distribution is heavy tailed (e.g., power law distribution) and the average risk is unbounded. Systems with these properties include billing infrastructures used to charge customers for services they access. Attacks against billing infrastructures can be classified as peripheral attacks and backbone attacks. The goal of a peripheral attack is to tamper with user bills; a backbone attack seeks to seize control of the billing infrastructure. The probability distribution of the overall impact of an attack on a billing infrastructure also has a heavy-tailed curve. This implies that the probability of a massive impact cannot be ignored and that the average impact may be unbounded – thus, even the most expensive countermeasures would be cost effective. Consequently, the only strategy for managing risk is to increase the resilience of the infrastructure by employing redundant components.

Keywords: Risk modeling, risk management, billing infrastructures

1. Introduction

This paper describes the modeling and management of risk in an information and communications technology (ICT) infrastructure where the average impact of an attack is unbounded. A mathematical model is developed to express the impact of attacks and is then applied to a billing infrastructure. A billing infrastructure is an ICT infrastructure that is designed, constructed and managed to bill a large set of customers for services they access or consume. Such an infrastructure comprises a set of peripheral nodes and an intelligent backbone [2]. An example is a metering infrastructure in which the peripheral nodes measure the amount of a good (e.g., water or electricity) distributed to customers, and the backbone records, delivers and updates customer bills [9]. In general, the inner structure of a peripheral node depends on the service

C. Palmer and S. Shenoi (Eds.): Critical Infrastructure Protection III, IFIP AICT 311, pp. 51–64, 2009.

that is offered and billed. The intelligent backbone connects the peripheral nodes and includes additional computing nodes that manage and update the information shared by the peripheral and backbone nodes.

The impact of an attack against a peripheral node in a billing infrastructure is bounded. However, the impact of an attack against the intelligent backbone cannot be bounded because it depends on the infrastructure cost and/or the value of the business processes that use the infrastructure. Mathematical models for estimating the overall impact of attacks on peripheral nodes and on the backbone may be defined as the sum of two random processes described by a normal distribution and a power law distribution, respectively [5, 12–14, 17]. This paper discusses the models and their implications on risk management for the overall infrastructure. In particular, it considers the problem that arises when the average impact of an attack is unbounded, and demonstrates that, in such a case, it is difficult to predict the impact of attacks even when historical attack data is available. The paper also discusses how this result influences the selection of countermeasures [22, 23].

2. Billing Infrastructures

This section briefly describes a billing infrastructure, which corresponds to an abstract model of an ICT infrastructure [1, 2, 15]. A billing infrastructure is characterized by the types of attacks and their impact on the infrastructure rather than the specific ICT components used in the infrastructure. The section also presents some real-world infrastructures that match the abstract billing infrastructure [9].

2.1 Infrastructure Overview

A billing infrastructure charges customers for a service that they receive. The service is supplied by the same infrastructure or by a different infrastructure; the service provider is also the infrastructure owner. The infrastructure consists of a set of peripheral nodes, one for each customer (in general), along with an intelligent backbone. Peripheral nodes may be distributed across a large region (e.g., a country); each node stores, manages and updates information about the quantity of service received by a user. The backbone interconnects the peripheral nodes and other computing nodes. The computing nodes store information about the peripheral nodes in order to manage the overall service distribution and to bill users.

As example is a content distribution service, where the billing infrastructure charges each user for the content that has been accessed; the content may be distributed by another infrastructure as in the case of pay-per-view movies. Another example is a metering infrastructure, where each peripheral node is connected to a meter that measures the quantity of some good (e.g., water, gas or electricity) that is distributed to a customer. In this type of infrastructure, each peripheral node computes and transmits to the backbone the amount of good consumed by each user and the corresponding bill.

An infrastructure can terminate the distribution of a good when a condition related to the quantity of the good consumed and/or consumer status is met. For example, the infrastructure may prevent a customer who has not paid his bill from further resource usage. In an advanced metering infrastructure, a peripheral node can also program the behavior of other devices to optimize the overall amount of the resource that is consumed or to optimize a combination of parameters such as the overall amount of the consumed resource and the customer bill. This can happen, for example, if a peripheral node schedules multiple devices in a home to minimize the overall amount of electricity that is consumed.

From our point of view, the internal behavior of peripheral nodes and the backbone are not fundamental because the important properties of the two subsystems are related to the attack impact. In particular, we are interested in billing infrastructures where the impact of an attack against a peripheral node is bounded whereas the impact of an attack on the backbone is unbounded (e.g., if the attacker seeks to control the overall infrastructure). In practice, the impact may be bounded by the cost of the overall business process that uses the infrastructure. However, because this cost depends on the infrastructure that is considered, no bound may exist in the general case. Furthermore, the overall impact may also depend on other infrastructures that are connected to the infrastructure under consideration. This problem is discussed below in the context of developing a mathematical model for attack impacts.

Note that the two types of attacks considered in this paper are distinguished by the goal of the attacker instead of the subsystems that are involved. Thus, an attack that targets a peripheral node as the first step of an attack against the backbone is considered to be an attack against the backbone.

2.2 Threat Model

The threat model considers two types of attacks against the infrastructure: (i) peripheral attacks that attempt to reduce user bills by attacking peripheral nodes; and (ii) backbone attacks that seek to control the infrastructure.

A peripheral attack that attempts to reduce a customer's bill is typically executed by an unethical costumer. We assume that general statistics about the customer population are available, which implies that the percentage of customers who may behave in an unethical manner is also known.

A backbone attack may seek to reduce the bills of a large number of customers, access confidential information about a set of customers, or control the use of a resource or service. Such an attack may be executed by a competitor, organized crime group or terrorist entity.

For both types of attacks, we distinguish between an attack that requires skill and knowledge that cannot be encapsulated in a tool that automates the attack, and an attack that can be fully automated so that its execution does not require any knowledge about the implementation of the infrastructure and nodes, only the availability of an attack tool. The two cases correspond to distinct pools of attackers because if an attack cannot be automated, then only

a customer/attacker with the requisite skill and knowledge who is willing to act in an unethical manner can execute the attack. If a tool that implements the attack is available (e.g., downloadable from the Internet), then any unethical customer can launch the attack. For both types of attacks, the impact of a peripheral attack is bounded by the customer bill plus the cost to replace the peripheral node.

In the case of backbone attacks, we also distinguish between automated attacks and non-automated attacks. However, the impact is not related to customer bills because the goal of the attack may be to control the entire infrastructure or distinct systems connected to the infrastructure and managed by the owner. As described below, the notion of an average impact is questionable when modeling backbone attacks because the average impact of distinct sets of backbone attacks may converge to distinct values. Another difference between the two classes of attacks is related to the discovery of a vulnerability after the infrastructure has been deployed. If a newly-discovered vulnerability only enables peripheral attacks, then it increases the probability of one of these attacks but not the largest impact, which is always bounded by the customer bill. On the other hand, a vulnerability that enables a backbone attack may increase the probability of a successful attack and, thus, increase the overall impact or the overall value at risk.

3. Modeling Attacks and Attack Impact

This section discusses the modeling of peripheral attacks and backbone attacks, and the impact of these attacks on the billing infrastructure. The attack impact is modeled by considering a time interval and attempting to predict all the impacts of interest in this interval and the information needed for prediction.

3.1 Peripheral Attacks

If an attack can only be executed manually (i.e., it cannot be automated), then the average number of attacks is proportional to the number of unethical customers who have the knowledge and the skills required to execute the attack. On the other hand, the number of automated attacks is proportional to the number of unethical customers. This also covers the case where unethical customers contract external agents to execute attacks on their behalf.

In the following, Npa denotes the expected number of potential attackers (this is always very large, but is much larger in the case of automated attacks). The impact D_i due to customer C_i is described by a random process with a probability distribution $Imp_i(D)$ specifying the probability that C_i executes an attack that yields an impact of D. $Imp_i(D)$ is larger than zero if $D \in 0..M_i$, where M_i is C_i's largest bill. The shape of Imp_i cannot be easily deduced as it depends on the amount of resources and skill C_i can summon to execute the attack. Hence, the variance of Imp_i is unknown and its rigorous approximation depends on several factors, including: (i) the ability of C_i to implement the

attack that, in turn, influences the probability of a successful attack; and (ii) the gap between the time distribution of the attacks of C_i and the resource usage of C_i.

However, the variance of Imp_i can be approximated if a representative sample of peripheral attacks is available. In this case, we can compute the largest customer bill M, which represents the upper bound of the impact. Obviously, M is finite and any error in the approximation of Imp_i is bounded because the distance between impacts is always bounded by M.

The overall impact of a collection of peripheral attacks Ima is a stochastic process that is the union of the processes $D_1, \ldots, D_n$ corresponding to the impacts of the individual attacks by customers. Since $Ima = \Sigma_{i=1..n} D_i$, whenever the number of customers is very large, the shape of Ima can be approximated by assuming that the impacts of distinct customers are independent and applying the central limit theorem. Under these assumptions, Ima is normally distributed with a mean and variance equal to the sums of the means and variances of D_i, respectively. Since each sum can be restricted to Npa unethical customers (the only individuals who can execute attacks), the mean of Ima is bounded by $Npa \times M$ where M is the largest customer bill. By profiling unethical customers, we can replace M by Mun, the largest bill in the group of unethical customers. Similar considerations apply to the estimation of the upper bound for the variance of Ima. Thus, the estimate of the overall impact of peripheral attacks improves if it is possible to profile unethical customers.

Obviously, Npa strongly increases when attacks can be automated. Nevertheless, Ima can always be approximated by a normal distribution when the number of attackers or (from another point of view) the number of unethical customers is so large that the error due to the application of the central limit theorem is acceptable. Note that the independence property of customer attacks is fundamental. We assume that this property holds even if some customers belong to social networks and exchange information about vulnerabilities and attacks. Thus, the relevance of social networks is ignored when computing the overall impact. However, even if social networks are considered, and the number of successful attacks increases and the parameters of the normal distribution change, the overall impact and the approximation error are still bounded.

3.2 Backbone Attacks

The approach adopted for modeling peripheral attacks cannot be used for backbone attacks because it is not possible to approximate the largest impact or the average impact of attacks. In fact, any successful attack against the backbone has an unbounded impact if the appropriate backbone components are controlled as a result of the attack. While some attackers would be interested in achieving as large an impact as possible, other attackers may be interested in a bounded impact in order to achieve their goals (i.e., they execute attacks to control infrastructure components that may cause larger impacts than the attacks that interest them).

In order to model the impact probability distribution, we assume that the backbone is often optimized to minimize its overall cost [6] and that this may result in the adoption of a preferential attachment strategy to define the interconnections among backbone components at distinct implementation levels (ranging from the physical interconnections to the services offered by software components) [1, 16]. In this case, the impact of an attack depends on the component that is the target of the attack, and the impact probability distribution is a power law of the form:

$$\frac{C}{x^{1+a}}$$

where C is a normalizing constant.

In general, the probability distribution of an impact x assumes arbitrary values if x is in the range $0..xmax$, and subsequently follows a power law. If the sum of the values that x assumes before the power law behavior is γ, i.e.,

$$\Sigma_{x \in 0..xmax} p(x) = \gamma,$$

then, for $x \geq xmax + 1$, the distribution has the form:

$$\frac{a \cdot (1-\gamma)}{(xmax+1)} \cdot \left(\frac{xmax+1}{x} \right)^{1+a}$$

This also covers the more interesting case where $x = 0$ is the only impact with a non-null probability in the range $0..xmax$ because every successful backbone attack has an impact larger than $xmax$. It is important to note that the impact probability distribution has power law behavior whenever the parameter to be optimized is the overall cost (or return on investment) even if faults or external attacks are considered. As an example, the high optimized tolerance methodology introduces components into a system to minimize the impact of faults [4, 18]. However, because this methodology optimizes the return on investment, the impact distribution also has power law behavior. Furthermore, any error in the approximation of the fault distribution strongly reduces the effectiveness of the optimization and may give rise to unbounded impacts.

We assume that the overall impact due to a single attacker also has a power law distribution. This implies that the attacker targeting the backbone is interested in executing just one attack, but the most powerful attack he can implement. Obviously, the actual impact would depend on the attacker's motive and knowledge, but this only influences the parameters in the power law equation. Therefore, in the worst case, the probability distribution of the impact x, $Psa_i(x)$, is a heavy-tailed power law [12, 13, 21] given by:

$$Psa_i(x) = \frac{C}{x^{1+a}}$$

where $0 < a < 1$. This is the worst case because the probability of a large impact decreases very slowly.

In the following, we use a power law rather than a heavy-tailed power law. The corresponding results hold for the larger class of probability distribution

functions (i.e., subexponential functions [7]), which includes any function that decreases slower than an exponential function. A process X has a subexponential distribution if:

$$limit_{d\rightarrow\infty}\frac{Prob\,(X_1+...X_n>d)}{Prob\,(max(X_1,...,X_n)>D)}=1$$

where any X_i is distributed as X and all the X_i are independent.

This condition implies that the sum is large because of the large contribution provided by only one term of the sum. In a billing infrastructure, the condition implies that an attacker is interested in causing one large impact rather that several average impacts. This is often the case because low impact attacks are of limited interest to several classes of attackers. For example, a terrorist entity or a competitor would be interested in executing one large impact attack that results in considerable publicity (and loss of credibility for the owner of the targeted infrastructure) rather than several low impact attacks that also increase the probability of being detected and apprehended.

The class of subexponential functions strictly includes the class of heavy-tailed functions. The class of heavy-tailed functions includes any distribution of a process X where for any h:

$$limit_{D\rightarrow\infty}\frac{Prob\,(X>D+h)}{Prob\,(X>D)}=1.$$

In other words, as D increases, for any h, the probability that X is larger than D is that same as the probability that X is larger than $D+h$. This implies that an attack that produces an impact D can also produce an impact $D+h$. Note that this class faithfully models the case under consideration because, as D increases, the backbone components that must be attacked to produce an impact D make it possible to achieve an even larger impact.

Alternatively, a process X is deemed to have a heavy-tailed distribution if a value V exists such that, if $X\geq V$, then the ratio:

$$\frac{Prob\,(X>nD)}{Prob\,(X>D)}$$

is independent of D for any $n>0$. In our case, this again expresses the fact that if the impact of an attack is larger than D, then it may be unbounded. Another reason to describe the impacts of backbone attacks using a power law is that the impacts may be proportional to the overall value of the infrastructure. Additional reasons for using a power law are discussed in [10, 21].

An important consequence of a power law distribution of impacts is that, depending on the exponent, it may be impossible to build a representative sample of backbone attacks. This implies that the central moment estimators (e.g., mean and variance) of finite-sized samples drawn from the impact distribution may not converge to a value when data is accumulated. This is because no moment is defined for the distribution and a key property of a billing infrastructure is that the impact of just one attack may be unbounded. The overall impact of

attacks on the infrastructure is the sum of the impacts of all the attacks; the corresponding random process is the union of all the processes corresponding to the individual attacks. Note that the probability distribution of a process created by the union of several processes, each described by a power law, is also a power law whose exponent is equal to the minimum of the exponents of the individual power laws [24]. Informally, the differences between the summands may be so large that the behavior of the sum largely depends on the maximum term and the probability distribution of this term is a power law.

Another implication is that the attack impacts are distributed according to a power law even if only a few attackers are interested in very large impacts because the overall impact mostly depends on these attackers. In other words, a general model of backbone attacks against a billing infrastructure assumes that there are at least two sets of attackers who are interested in finite impacts and unbounded impacts, respectively. The behavior of attackers in the first class is described by a normal distribution that can be handled in a manner similar to that for peripheral attacks. However, a new problem posed by backbone attacks is that attackers are interested in impacts that are distributed according to a power law, which determines the overall impact for the infrastructure owner. It is possible to introduce an upper bound also on the impact of backbone attacks by summing a negative exponential term to the power law to quickly cut off the probability of an impact that is larger than a threshold. However, this solution increases the complexity of the model without increasing its accuracy, especially for a large threshold. Note that in many instances it is almost impossible to determine a proper threshold value.

As an example, consider the case where a billing infrastructure is connected to other infrastructures outside the control of the owner or, even worse, where the existence of such a connection may not be known but cannot be excluded *a priori*. In this case, the impacts on other infrastructures must be considered, but they cannot be estimated easily. Therefore, in the next section, we assume that the probability distribution $Iia(D)$ of the random process that describes the overall impact of backbone attacks follows a power law. Based on historical data about infrastructure attacks, power law behavior may occur only for values larger than a positive threshold, while a distinct distribution models lower impacts. One of the key issues related to adopting a power law is discussed in the following section – it concerns the interpolation of the characteristic parameters of the probability distribution of the overall impact $Iia(D)$.

3.3 Overall Attack Impact

As described above, the overall impact of attacks against a billing infrastructure is a random process $OvImp$ that is the sum of two processes:

- ***Imppa:*** Impact of peripheral attacks, which has a normal distribution Nld.
- ***Impba:*** Impact of backbone attacks, which has a power law distribution Pld.

Note that the power law behavior may start at any positive integer value and that the mean of the normal distribution may be strongly shifted towards large positive values when the percentage of unethical customers and the maximum bill of the group of unethical customers are both very large. If we assume that Ima and Iia are mutually independent, then the probability distribution of $OvImp$, $Iia(D)$, is the convolution of Nld and Pld. Unfortunately, the mean and other moments of $Iia(D)$ cannot be computed because these statistics do not exist for Pld.

First, we consider the interpolation of the parameters of $Iia(D)$ and the component distributions, Nld and Pld, using actual attack data. The complexity of the interpolation strongly depends on the amount of information that is available. It has been shown [22] that this problem is extremely complex when only a sequence of outputs of the overall process ($OvImp$, in our case) is available because it is almost impossible to determine which component process (impacts of peripheral and backbone attacks, in our case) generates each output. This occurs when the two process domains overlap in a manner that prevents the pairing of some outputs with the corresponding processes.

Since the outputs are impacts, this means that we can only observe a sequence of impacts, i.e., a decrease in revenue for the owner. Also, the overlap of Pld and Nd may prevent us from recognizing their relative contributions to each observed impact and, thus, from approximating the parameters of each distribution. Moreover, the time frame available for impact data collection may be too short to cover a number of backbone attacks completely, which would make it impossible to deduce the parameters of the corresponding processes [22]. This is an important, but pessimistic, result because it means that the properties of $Iia(D)$ cannot be deduced even when a large sequence of attack impacts is used. The impossibility of forecasting future attacks and their impacts arises not only because of the lack of data about previous attacks but also because the distributions of interest cannot be approximated from the available data.

While the previous considerations hold for the abstract case of stochastic processes and a sequence of impacts, attacks on a billing infrastructure (as with most physical systems) often leave evidence in certain infrastructure components. Furthermore, some infrastructure components may be designed to facilitate the discovery of evidence (e.g., log files that record infrastructure activities and intrusion/anomaly detection systems that analyze the interactions between infrastructure components). This evidence may be used to pair an impact with the corresponding attack and to discover the relative contributions of attacks. From a probability point of view, proper attack classification makes it possible to analyze (separately) the probability distribution of each process and attempt to approximate Nld and Pld instead of using the distribution of the union process $Iia(D)$. In other words, a forensic analysis of attacks can help pair each impact with a successful attack against the infrastructure in order to deduce the properties of each distribution. This implies that the infrastructure

should be designed to facilitate the forensic analysis of successful attacks as well as attempted attacks.

4. Risk Management Strategies

This section examines the implications of the attack impact probability distribution on risk management for a billing infrastructure and on the return on security investments.

The primary problem related to managing risk in a billing infrastructure is the evaluation of the cost effectiveness of countermeasures implemented against peripheral and backbone attacks. Since the impact of a single peripheral attack and that of the entire class of peripheral attacks can be bounded, it is possible to determine the conditions that guarantee the cost effectiveness of countermeasures in terms of the impact probability distribution for peripheral attacks, the bounds on attack impacts and the cost of the countermeasures. The cost effectiveness of countermeasures for a single peripheral node depends on the average loss for the node, and the overall impact places a bound on the largest return on investment in security for all the peripheral nodes. Knowledge of the normal probability distribution of the overall impact of peripheral attacks can be used to fine-tune the choice of countermeasures by taking into account the distribution variance and the exponential decrease of the probability of very large impacts. The error in approximating the actual probability distribution as a normal distribution should also be taken into account.

Backbone attacks are more complex because of the power law distribution of their impact. A heavy-tailed distribution makes it almost impossible to evaluate the cost effectiveness of a countermeasure because very little information is available about the expected impacts of the attacks that are foiled by the countermeasure. In particular, a power law distribution implies that even if the probability of an impact is very low, its relative weight that cannot be ignored. Since the relative impact of some attacks cannot be easily bounded, the overall impact strongly depends on these attacks. This results in an unmanageable situation from the point of view of cost effectiveness because the impact justifies extremely expensive countermeasures while the probability of the attacks does not justify such an expense and no information about the average impact is available.

From a mathematical point of view, this situation strongly resembles the St. Petersburg paradox regarding a lottery with an expected unbounded payoff. In our model, if the impact probability distribution is a power law and if a proper condition on the $1 + a$ exponent holds (i.e., $a \in 0..1$), then the average impact of infrastructure attacks is infinite. This implies that we cannot claim that a set of countermeasures is optimal because the overall cost of any set of countermeasures is less than the impact it is intended to prevent.

A problem also arises when attempting to approximate the probability distribution parameters based on the available attack data. We have shown that even if a forensic data collection system has been implemented, a large amount of evidence about attempted and successful attacks and their impacts may

be required to approximate the distribution. Moreover, small data errors can produce large differences in the parameters of the impact distribution. Consequently, a risk management strategy based on cost effectiveness of countermeasures cannot be adopted in the majority of scenarios.

The only feasible risk management strategy is to minimize the probability of successful attacks while recognizing that some attacks will be successful and minimize their impact. According to this strategy, in the worst case, a successful attack should cause a graceful degradation of infrastructure performance and functionality, which is measured in terms of the ability of the billing infrastructure to meter service usage and charge customers. In other words, risk management should increase infrastructure resilience in order to minimize the probability of successful attacks and their impact [1, 11, 15].

A fundamental issue is the absence of singularity points of catastrophic failure at any level – from hardware components to the personnel responsible for infrastructure management – because any of these points is an ideal target to maximize the attack impact. In general, an approach that attempts to increase infrastructure resilience cannot be cost effective (based on the simple view of cost effectiveness described above) because it involves the addition of redundant components in the infrastructure. Furthermore, such an approach avoids large optimizations that result in scale-free networks.

Instead of introducing a few components with a large number of connections, a redundancy-based approach would distribute the same number of connections among a larger number of interconnected components, with an increase in the cost of connections and components. In terms of probability, redundancy implies the independence of the random variables used to model the components of interest. Therefore, whenever two random variables used to model infrastructure components are not independent, some dependencies exist among the components so that a successful attack against one component may simplify attacks against another. The adoption of redundancy at the software level may be even more costly than at the hardware level because (as far as reliability is concerned) two instances of the same software module will always be affected by the same faults or vulnerabilities. Hence, the adoption of redundant active software components implies the presence of distinct providers for each copy of a component to guarantee independence of both vulnerabilities and faults. Note also that to prevent the introduction of a single point of catastrophic failure, the redundant components may have to be executed in parallel by different computing resources and they have to be properly synchronized, which increases the execution time. This also contributes to increased overall cost and reduced cost effectiveness.

Consider, for example, a standard implementation of triple-modular redundancy with three components and a voter, where the voter is a point of catastrophic failure [8]. If the threat model assumes erroneous but not malicious behavior of a component and possible voter failure, then a spare replacement for the voter increases the overall redundancy. On the other hand, if the threat model covers both erroneous and malicious behavior, then a distributed im-

plementation of the voter is required where all the consumer components (i.e., components that receive the output of the components that act as producers) need to exchange the received values to compute their correct input [19]. However, a solution that is correct independently of the behavior of each producer can be defined only if at least five consumers exist, so that at least five instances of each module are required to discover malicious behavior in just one producer. This simple example shows that failure independence implies that redundancy is effective only if the failure of each instance is independent of the failure of other instances.

Consider also the case where two copies of a database reside on physical systems maintained at different locations. The databases may be independent with respect to physical threats such as earthquakes or floods, but they can be attacked by the same malware or infected by the same virus and are, therefore, are not independent in general. The incorporation of components to discover attacks and their impact can further reduce the cost effectiveness because they are not required for normal infrastructure operations. Note also that a rigorous approach to risk assessment, security and integrity of an infrastructure may distinguish between the strategies to manage the risk due to unethical customers, a customer that attacks the entire infrastructure, and business continuity. While there are good management reasons for the approach underlying these, or similar, classifications, it is important to recognize that a modular approach to risk management should not ignore the fact that several threats may result in similar impacts, and that it is complex (if not impossible) to assess the probability that one of these threats implements a successful attack.

5. Conclusions

Attacks that target billing infrastructures have heavy-tailed impact probability distributions, typically power law distributions. This implies that the mean value of the impact of attacks cannot be computed and that the choice of countermeasures cannot be made on the basis of cost effectiveness. As a consequence, the only risk management strategy appropriate for a billing infrastructure is one that introduces redundant components to increase the resilience of the infrastructure and decrease the probability of successful attacks.

References

[1] R. Albert, H. Jeong and A. Barabasi, Error and attack tolerance of complex networks, *Nature*, vol. 406, pp. 378–382, 2002.

[2] F. Baiardi, C. Telmon and D. Sgandurra, Hierarchical, model-based risk management of critical infrastructures, *Reliability Engineering and System Safety*, vol. 94(9), pp. 1403–1415, 2009.

[3] P. Bernstein, *Against the Gods: The Remarkable Story of Risk*, Wiley, New York, 1996.

[4] J. Carlson and J. Doyle, HOT: A mechanism for power laws in designed systems, *Physical Review E*, vol. 60(2), pp. 1412–1427, 1999.

[5] A. Clauset, C. Shalizi and M. Newman, Power-law distributions in empirical data, arXiv:0706.1062v2, arXiv, Cornell University, Ithaca, New York (arxiv.org/PS_cache/arxiv/pdf/0706/0706.1062v2.pdf), 2007.

[6] R. D'Souza, C. Borgs, J. Chayes, N. Berger and R. Kleinberg, Emergence of tempered preferential attachment from optimization, *Proceedings of the National Academy of Sciences*, vol. 104(15), pp. 6112–6117, 2007.

[7] C. Goldie and C. Kluppelberg, Subexponential distributions, in *A Practical Guide to Heavy Tails: Statistical Techniques and Applications*, R. Adler, R. Feldman and M. Taqqu (Eds.), Birkhauser, Boston, Massachusetts, pp. 435–459, 1998.

[8] L. Lamport, R. Shostak and M. Pease, The Byzantine generals problem, *ACM Transactions on Programming Languages and Systems*, vol. 4(3), pp. 382–401, 1982.

[9] L. LeMay, R. Nelli, G. Gross and C. Gunter, An integrated architecture for demand response communication and control, *Proceedings of the Forty-First Annual Hawaii International Conference on System Sciences*, p. 174, 2008.

[10] T. Maillart and D. Sornette, Heavy-tailed distribution of cyber-risks, arXiv:0803.2256v2, arXiv, Cornell University, Ithaca, New York (arxiv.org/PS_cache/arxiv/pdf/0803/0803.2256v2.pdf), 2008.

[11] D. Maluf, Y. Gawdiak and G. Bell, On space exploration and human error: A paper on reliability and safety, *Proceedings of the Thirty-Eighth Annual Hawaii International Conference on System Sciences*, p. 79, 2005.

[12] B. Mandelbrot, *Fractals and Scaling in Finance: Discontinuity, Concentration, Risk*, Springer, New York, 1997.

[13] B. Mandelbrot, New methods of statistical economics revisited: Short versus long tails and Gaussian versus power law distributions, *Complexity*, vol. 14(3), pp. 55–65, 2009.

[14] M. Mitzenmacher, A brief history of generative models for power law and log-normal distributions, *Internet Mathematics*, vol. 1(2), pp. 226–251, 2003.

[15] National Infrastructure Protection Center, Risk Management: An Essential Guide to Protecting Critical Assets, Washington, DC, 2002.

[16] M. Newman, The structure and function of complex networks, *SIAM Review*, vol. 45(2), pp. 167–256, 2003.

[17] M. Newman, Power laws, Pareto distributions and Zipf's law, *Contemporary Physics*, vol. 46, pp. 323–351, 2005.

[18] M. Newman, M. Girvan and J. Doyne Farmer, Optimal design, robustness and risk aversion, *Physical Review Letters*, vol. 89(2), pp. 028301.1–028301.4, 2002.

[19] M. Pease, R. Shostak and L. Lamport, Reaching agreement in the presence of faults, *Journal of the ACM*, vol. 27(2), pp. 228–234, 1980.

[20] S. Resnick, *Heavy-Tail Phenomena: Probabilistic and Statistical Modeling*, Springer, New York, 2007.

[21] D. Sornette, *Critical Phenomena in Natural Sciences: Chaos, Fractals, Self-Organization and Disorder: Concepts and Tools*, Springer, Berlin-Heidelberg, Germany, 2006.

[22] N. Taleb, Black swans and the domains of statistics, *The American Statistician*, vol. 61(3), pp. 1–3, 2007.

[23] N. Taleb, *The Black Swan: The Impact of the Highly Improbable*, Random House, New York, 2007.

[24] C. Wilke, S. Altmeyer and T. Martinetz, Large-scale evolution and extinction in a hierarchically structured environment, *Proceedings of the Sixth International Conference on Artificial Life*, pp. 266–272, 1998.