Average Impact of Attacks on Billing Infrastructures

F.Baiardi, C.Telmon
Dipartimento di Informatica
Università di Pisa

Abstract

A billing infrastructure is a networked system developed to bill a set of users for some service. Users may exploit vulnerabilities of the infrastructure to implement some attacks and reduce their bills. We consider the search of vulnerabilities in these infrastructures and assume that there are two sets of people, the attackers and the defenders, that compete in the search. The attackers are interested in the attacks enabled by a vulnerability, instead the defenders are interested in a patch that removes a vulnerability. The 0-delay model is a probabilistic model to evaluate the average impact of attacks. The main assumption underlying the model concerns the timing of the actions of the attackers and the defenders, respectively the attack or the patching, that are executed as soon as the vulnerability is discovered. Obviously, the action that is executed depends upon who finds the vulnerability. The main model parameters are the number of attackers and that of defenders. Starting from these parameters, the model evaluates the window of vulnerability, that is the time between the discovery of the vulnerability by an attacker and the discovery of the same vulnerability by a defender. Starting from the size of the window, the average impact of an attack is computed. The model also supports the computation of the probability that an attack has occurred before the system has been patched and the impact of such an attack. Initially, we consider an infrastructure with just one vulnerability and then the most general case of several, independent, vulnerabilities.
After discussing some generalizations of the 0-delay model, where the time to either patch or attack the infrastructure is larger than zero, we show that the model may also be applied to evaluate of the advantages of open source components vs. that of proprietary components with a security through obscurity approach.

Keywords : Infrastructure, vulnerability, attack, impact, mathematical model, open source, security through obscurity

1. Introduction

A billing infrastructure is any networked system deployed to bill a set of users for some service supplied either by the same or by a distinct system. Well-known examples are public utility infrastructures such as those for the distribution of electric power or of water, where a meter measures the amount of power or water distributed to the user. Through the infrastructure, the meter sends the running total to a database that is used to compute the user bill. The revenue of the infrastructure owner is the overall amount of the bills. The lifetime of a billing infrastructure is fairly long because a large number of its components are physically distributed on a wide area and their update is expensive.

We are interested in a mathematical model to optimise the investment of the infrastructure owner in the search and the elimination of vulnerabilities [1- 3, 5, 7, 12-15, 17, 19, 25] after the infrastructures has been deployed. For this reason, we are focused on vulnerabilities enabling attacks [3, 5, 7, 19] resulting in losses in the revenue and neglect other impacts, such as denials of service. We assume that two sets of people compete in the

search of vulnerabilities: attackers, a.k.a. white hats, and defenders, a.k.a. black hats. The goal of a defender is to patch the infrastructure [10] to prevent an attack. That of attackers, instead, is an attack resulting in lower bills. The number of defenders depends upon the investment in security of the owner after deploying the infrastructure. In a billing infrastructure, the loss of revenues depends upon the vulnerability window [8, 26, 27] of each vulnerability. For a vulnerability V, this window is the interval of time from when an attacker discovers V till when a defender finds V. The proposed model, 0-delay model, evaluates the loss in the revenue in terms of the window size and of the numbers of attackers and of defenders. For the sake of simplicity, at first, an infrastructure with just one vulnerability is considered, and then the case of several vulnerabilities is discussed. In such a case, the model may be paired with a game theory [21] approach to define an optimal allocation of attackers and defenders to the search of distinct vulnerabilities. The model also enables the owner to determine whether to deploy the infrastructure even if some vulnerabilities have not been removed because he/she is willing to accept the risk of the average impact of an attack. Lastly, the model supports a comparison of the advantages of open source components vs. proprietary ones with a "security through obscurity" approach [5, 6].

The importance of a quantitative evaluation of attack impacts has often been stressed [5, 17, 18, 23, 25]. A survey of current approaches is presented in [25] together with the notion of market price of vulnerability. This notion cannot be immediately applied to a billing infrastructure where the price of a vulnerability depends upon the service billed through the infrastructure rather than on the infrastructure itself. [16] applies game theory to information warfare while [21] discusses the optimal allocation of defenders to minimize the impact of a terrorist attack. The competition between benign users and attackers in the search for vulnerabilities has been considered in [9, 23] but these papers are focused on the disclosure policy rather than on attack impacts. Some assumptions of our model are similar to the one in [23] to compute the probability of finding a vulnerability. [9] considers the search for vulnerabilities and a social planner that decide when a vulnerability is disclosed. Coherently with the evaluation of disclosure policies, it assumes that a benign user always discovers a vulnerability before an attacker. Furthermore, most of the work on vulnerabilities and attacks considers general-purpose systems rather than billing infrastructure.

Sect. 2 introduces the 0-delay model and shows how it defines the average impact of an attack as a function of the numbers of attackers and defenders and of the vulnerability window. Sect. 3 considers an infrastructure with several vulnerabilities and shows that the impact is always a function of the numbers of attackers and defenders searching for distinct vulnerabilities. Lastly, we apply the model to the debate on "protection through obscurity" or on the adoption of open source components.

2. The 0-delay Model

After discussing the main assumptions underlying the 0-delay model, we present the model in some details. Lastly some generalizations of the models are discussed.

2.1 Underlying Assumptions

Besides the one implied by its name, the most important assumptions underlying the 0-

delay model concerns the existence of one vulnerability, denoted by V, and that the billing infrastructure is deployed even if it includes V. The former will be discussed in the next section. The latter, in general, it is satisfied because it may be not cost effective to deploy the infrastructure only after removing any vulnerability. Assuming the existence of V, two sets of people are searching for V, the attackers and the defenders. The attackers exploit V to define and implement an attack. The defenders, instead, search V to patch the infrastructure.

Time is modelled as a sequence of intervals with a size δt, in the following ***at time t*** means ***during the t-th interval.*** If a defender finds V, in the same interval, the patch is defined and applied to the infrastructure. We assume that the time to develop a patch is independent of the number of defenders and that δt is larger than the interval to start and complete the patching process. If a defender finds V at time t, any attack implemented after t fails. If, instead, an attacker finds V before any defender, then in the same interval the attack occurs and the loss begins. The loss ends only when, and if, the defender finds V and patches the infrastructure. Notice that δt cannot be reduced at pleasure because one interval suffices to define and execute an attack or to define and apply the patch. This also implies that δt depends upon the considered infrastructure. The probability of discovering V is the same for any interval, although it is different for an attacker and for a defender.

A further assumption concerns the absence of communication between the attackers and the defenders or within each set during the search. Hence, no information from other people is available to speed up the search. However, as soon as the attack has been discovered, it is immediately broadcasted to anyone that can implement it and that all the attacks are immediately executed. This is a worst-case for the defenders, as any delay in the execution of attacks reduces the loss. Furthermore, if not all the attacks are executed simultaneously and one of them is detected, this increases the probability that a defender finds V.

The model assumes that the impact of an attack is proportional to the size of the vulnerability window and that lifetime of the infrastructure is unbounded, i.e. the infrastructure is updated only to remove any vulnerability. The latter is realistic only for the long-term components of the infrastructure, such as the hardware of an ATM or a meter in the user house. Hence, the model should be applied to these components only. Notice that these two assumptions imply that the impact of a successful attack may become unbounded because it is proportional to the vulnerability window size but this size may be unbounded if a vulnerability that is not removed has been discovered by an attacker.

2.2 The 0-Delay Model

According to the 0-delay model, *I(na, nd),* the impact of an attack is a function of *na* and *nd* , the numbers of attackers and of defenders. *I(na, nd)* is positive if and only if the size of the vulnerability window is positive and it is proportional to this size as well as to the number of successful attacks. This is summed up in the relation:

$$I(na, nd) = \begin{cases} Pf_A \, Expl\,(td(nd) - ta(na)) & \text{if} \quad td(nd) - ta(na) > 0 \\ 0 & \text{if} \quad 0 \geq td(nd) - ta(na) \end{cases}$$

where:

- *ta(na)* is the time when one of the *na* attacker discovers *V* and *Att,* the attack enabled by V;
- *td(nd)* is the time when one of the *nd* defenders finds *V* and the patch is applied,
- *td(nd) – ta(na)* is the size of the vulnerability window,
- *Expl* is the number of successful attacks, each an instance of *Att*. *Expl* cannot be smaller than *na*, that is $Expl = \psi\, na$, $\psi \geq 1$,
- ψ is a decreasing function of the resources and the skills to execute *Att* and it reaches its maximum if *Att* can be fully automated by proper programming tools [29]
- Pf_A is the loss in the revenue for unit of time due to a single attack.

The model assumes that $Pf_A\, Expl$ is a constant.

If $\mathbf{Av}(R)$ denotes the average value of the random variable *R*, then

$$\mathbf{Av}\,(Pf_A \cdot Expl \cdot (td(nd) - ta(na)) = Pf_A \cdot Expl \cdot \mathbf{Av}\,(td(nd) - ta(na)).$$

We are interested in the positive values of the vulnerability windows because only in this case *Att* has been successful. All the cases where $td(nd) < ta(na)$ are mapped into a zero size because in all these the loss is zero. In the following, we drop the dependency from *na* and *nd* and replace *td(nd)–ta(na)* by *td–ta* or by *vw* .

$\mathbf{Av}(vw)$ the average size of the window depends upon $P(vw = i > 0 \mid na, nd)$, the probability that $vw = i$ conditioned to the existence of na attackers and of nd defenders. In turns, this probability is a function of *Pd(nd)* (*Pa(na)*), the probability that, in each interval, at least one of the *nd* defenders (*na* attackers) finds V. Since both *Pd(nd)* and *Pa(na)* are time independent, the probability that the defenders (attackers) find V exactly at time t i.e. that $td=t$ ($ta=t$) is

$$(1- Pd(nd))^{t-1} \cdot Pd(nd) \qquad \left((1- Pa(na))^{t-1} \cdot Pa(na) \right)$$

Taking into account that each attacker and each defender works in isolation,

$$Pd(nd) = 1-(1-Pd(1))^{nd}$$
$$Pa(na) = 1-(1-Pa(1))^{na}$$

where *Pd(1)* and *Pa(1)* are, respectively, the probabilities that a defender and an attacker finds V in a time interval. We assume that an attacker and a defender have the same probability of finding the vulnerability in one interval so that

$$Pd(1) = Pa(1) \qquad (1)$$

This assumption neglects the larger amount of the information on the infrastructure that the defender can access and that should, at least in principle, simplify the search. The 0-delay model represents this asymmetry by multiplying the number of defenders by an equivalence ratio φ so that we may assume that (1) holds. For the sake of simplicity, we assume that the

number of defenders has already been multiplied by φ and drop the dependency of the values from *Pd(1)* and *Pa(1)*.

The previous consideration shows that the following relation holds:

$$P(vw=td\text{-}ta= i >0 \mid na, nd) = \lim_{N\to\infty} \sum_{ta=1}^{N-i} (1 - \mathrm{Pa(na)})^{\mathrm{ta\text{-}1}} \cdot \mathrm{Pa(na)} \cdot (1 - \mathrm{Pd(nd)})^{\mathrm{ta\text{-}1+i}} \cdot \mathrm{Pd(nd)}$$

In other words, the probability that *vw* = i, i>0, is the limit as N goes to infinity of the sum of all the cases where:

1. an attacker finds V at *ta*,
2. *td=ta+i*,
3. both *ta* and *td* belong to the range 1..N.

Notice that *ta* cannot be larger than N-i because *td* belongs to 1..N. We can consider the limit of the sum as $N\to\infty$ because the life of the infrastructure is unbounded. This approximation is acceptable anytime the infrastructure will be operational for a time much larger than δt.

It may be proved that:

$$\begin{aligned} &P(vw= td\text{-}ta = i>0 \mid na, nd) \\ &= \lim_{N\to\infty} Pa(na) \cdot Pd(nd) \cdot (1\text{-}Pd(nd))^{i} \frac{1-((1-Pa(na))\cdot(1-Pd(nd)))^{N-i}}{1-(1-Pa(na))\cdot(1-Pd(nd))} \\ &= Pa(na) \cdot Pd(nd) \frac{(1-\mathrm{Pd(nd)})^{i}}{1-(1-\mathrm{Pa(na)})\cdot(1-\mathrm{Pd(nd)})} \end{aligned}$$

We can exploit *P(vw= i>0 | na, nd)* to compute *P(vw=0 | na, nd)*, the probability that *ta>td* because that the defenders have discovered the vulnerability before the attackers:

$$\begin{aligned} P(vw=0 \mid na, nd) &= 1\text{-}\ P(vw>00 \mid na, nd) \\ &= lim\ n \to \infty (1\text{-} \sum_{i=1}^{n} \mathrm{P(vw = i)}) = \frac{Pd(nd)}{1-(1-Pa(na))\cdot(1-Pd(nd))} \end{aligned}$$

Taking into account that a loss occurs if and only if *vw>0*, we have that

$$\begin{aligned} \mathbf{Av}(I(na, nd)) &= Pf_A \cdot Expl \cdot \sum_{i=1,\infty} iP(vw = i \mid na, nd) \\ &= Pf_A \cdot Expl \cdot Pa(na) \cdot \frac{1-Pd(nd)}{Pd(nd)\cdot(1-(1-Pa(na))\cdot(1-Pd(nd)))} \end{aligned}$$

By deriving $\mathbf{Av}(I(na, nd))$ with respect to *Pd(nd),* we see that it reaches a maximum or a minumum if

$$Pd(nd) = 1 \pm 1/\sqrt{(1-Pa(na))}\ .$$

Since *Pd(nd)* belongs to (0,1), in this range $\mathbf{Av}(I(na, nd))$ is a decreasing function of *Pd(nd)*. By replacing $Pd(nd)=1-(1-p)^{nd}$ and $Pd(nd)=1-(1-p)^{na}$, we have that

$$\mathbf{Av}(I(na, nd)) = \mathrm{Pf}_A \cdot \mathrm{Expl} \cdot \frac{(1-(1-p)^{na}) \cdot (1-p)^{nd}}{(1-(1-p)^{nd}) \cdot (1-(1-p)^{na+nd})}$$

Since p is fairly small, because δt is small, we exploit the approximation $(1\text{-p})^{n} \approx 1 - p \cdot n$,

$$\mathbf{Av}(I(na,nd)) \approx Pconst_A \cdot \frac{1 - p \cdot nd}{p \cdot nd \cdot (1 + \frac{nd}{na})}$$

To increase the accuracy of the approximation, δt may be reduced because this reduces *p* too. However, δt cannot be arbitrary small because both an attack and the patching require δt.

By exploiting the same approximation in the formula for *P(vw=0|na, nd)*, we have that

$$P(vw=0|na, nd) \approx \frac{1}{1 + \frac{na}{nd}}$$

This shows that the probability that no loss occurs

- depends upon the **ratio between the number of attackers and of defenders** rather than upon the numbers of attackers and defenders
- is **independent of the probability that an attacker or a defender finds V**.

By deriving $\mathbf{Av}(I(na, nd))$ with respect to *nd* and *na*, we can verify that, as expected, lower number of defenders and a larger number of attackers always result into a larger vulnerability window and a larger impact.

2.3. Loss as a Function of the Time of the Discovery

We apply the 0-delay model to compute the average loss as a function of *td*, the time when a defender discovers V. This average loss is equal to

$$\mathbf{Av}(Svw(k \mid t, na, nd)) \cdot Pfa \cdot Expl$$

where *Svw(k |t, na, nd)* is the probability that *vw=k* provided that *td=t.* The loss defines an upper bound on the investment in the checks to discover attacks that may have occurred before patching [22]. These checks are the first step to recover the loss of the attacks but, since the checks may be rather expensive, an estimate of the loss enable the defender to choose whether it is more convenient to simply accept any loss occurred before t.

To compute $\mathbf{Av}(Svw(k \mid t, na, nd))$, we consider *P(vw=k>0 | td=t, na, nd)*, the probability that *vw* = k provided that there are *na* attackers, *nd* defenders and *td=t*. Since *td=t* and *vw=k* jointly imply that *ta=t-k*, i.e. attackers discovers V at t-k, we have that

$$P(vw=k| td=t, na, nd) = P(ta=t-k \mid td=t, na, nd)$$

Since the probability that an attackers finds V is independent of the one that a defender finds V we have that:

$$P(ta=t-k \mid td=t, na, nd) = P(ta=t-k \mid na, nd) \cdot P(td=t \mid na, nd)$$

By replacing the probabilities in the right hand size, we have that

$$P(ta=t-k \mid td=t, na, nd) = (1- Pd(nd))^{t-1} \cdot Pd(nd) \cdot (1-Pa(na))^{t-k-1} Pa(na).$$

We apply the 0-delay model to compute the average size of the window:

$$\mathbf{Av}(Svw(k \mid t, na, nd)) = \sum_{k=1}^{t-1} k(1-Pd(nd))^{t-1} \cdot Pd(nd)(1-Pa(na))^{t-k-1} Pa(na)$$

$$= (1-Pa(na)) \cdot (1-Pd(nd))^{t-1} \cdot Pa(na) \cdot Pd(nd) \cdot \sum_{k=1}^{t-1} k \cdot \frac{1}{(1-Pa(na))^{k}}$$

To simplify this expression, we notice that $\mathbf{Av}(Svw(k \mid t, na, nd))$ is interesting only when V has been discovered after the infrastructure has been working for a fairly long time. If, instead, it has been patched shortly after being deployed, the loss cannot be very large because vw is lower than the time from the deployment. Hence, we are interested in large values of t and the following approximation holds

$$\sum_{k=1}^{t-1} \frac{k}{(1-\mathrm{Pa(na)})^{k}} \approx \sum_{k=0}^{t-1} \frac{k}{(1-\mathrm{Pa(na)})^{k}} = \frac{\frac{1}{1-\mathrm{Pa(na)}}}{(1-\frac{1}{1-\mathrm{Pa(na)}})^{2}} = \frac{1-Pa(na)}{Pa(na)^{2}}$$

By applying this approximation, we have that

$$\mathbf{Av}(Svw(k \mid t, na, nd)) \approx ((1-Pd(nd)) \cdot (1-\mathrm{Pa(na)}))^{t-1} \cdot \mathrm{Pa(na)} \cdot Pd(nd) \cdot \frac{1-\mathrm{Pa(na)}}{\mathrm{Pa(na)}^{2}}$$

By simplifing the right hand size and by exploiting again the equalities $Pd(nd) = 1-(1-p)^{nd}$, $Pa(na) = 1-(1-p)^{na}$ and the approximation $(1-p)^{n} \approx (1-n\,p)$, we have that

$$\mathbf{Av}(Svw(k \mid t, na, nd)) \approx (nd \cdot (t-1) + na \cdot t) \cdot \frac{nd}{na} \approx (nd+na) \cdot \frac{nd}{na} \cdot t$$

This show that the average loss due to attacks occurring before t may be approximated by

$$Pf_{A} \cdot (nd+na) \cdot nd \cdot t$$

2.4. Generalization of the Model

This section generalizes the 0-delay model by relieving some of its assumptions.

At first, we consider the time between the discovery of the vulnerability and the patching of the infrastructure. In most cases, the time to produce and validate the patch or to update some components will be larger than zero. The delay increases with the number of the infrastructure components to be corrected. Consider, as an example the vulnerabilities in the WEP authentication scheme. Hence, the delay DP between the discovery of the vulnerability and the complete patching of the infrastructure may be fairly larger than zero. We assume that DP is not fixed but that it does not depend upon other parameters of the model. Let M_{DP} be an upper bound on DP.

To take this delay into account, we increase the size of the vulnerability window. As a matter of fact, if the defenders discover the vulnerability at *td* and the infrastructure is patched at $td+M_{DP}$ then $vw= td-ta+M_{DP}$. Obviously, the average of the new size can be computed by adding M_{DP} to value previously computed. In the same way, we can handle a delay DA between the discovery of V and the execution of the attacks exploiting V. If M_{DA} is the upper bound on the time to discover an attack, in the most general case, we have that

$$vw=td -ta+M_{DP} - M_{DA}= td -ta -(M_{DA} -M_{DP})$$

To compute the average loss, we take into account that $td -ta -(M_{DA} -M_{DP})(td -ta)$ has the same probability of $(td -ta)$ in the 0-delay model.

The previous discussion shows that the framework of the 0-delay model can handle constant delays both in the patching and in the attack, provided that all the attacks are executed simultaneously. Hence, *constant delay* may be a more appropriate name for the model.

Let us consider now the assumption on the simultaneous execution of attacks. As already mentioned, this is a worst case for the defenders because any delay in the execution of attacks reduces the loss. By removing this assumption, the overall number of attacks does not change but the attacks may occur at distinct times. As an example, at each interval, someone could implement *Att* and then inform *i* other people so that the number of attacks at *t* is *i* times that at *t-1*. If V has been discovered at *ta* and *Natt(t)* is the number of attacks executed at *t*, $t>ta$ we have that

$$NAtt(t)=\frac{i^{t-ta+1}-1}{i-1}$$

In the most general case, if *fa(t)* is the number of attacks executed at *t*, $t>ta$

$$NAtt(t)=\sum_{tv=o}^{t-ta} fa(ta+tv)$$

δaa, the interval to execute all the attacks, is computed by solving the equation

$$NAtt(\delta aa+ta) = Expl.$$

To compute the loss, we notice that, if $vw>0$, two cases have to be considered:

a) $td>ta+\delta aa$, if the defender discovers V after all the attacks have been executed,

b) $ta + \delta aa > td$, if the defender discovers V before all the attacks have been executed.

In case a), the overall loss is the sum of two losses. The first one occurs in the interval $(ta+ \delta aa, td)$ and it is equal to

$$Pf_A \cdot Expl \cdot (td-ta- \delta aa)$$

The other loss occurs in the interval $(ta, ta+stca)$ and it is equal to

$$Pf_A \cdot \sum_{t=0}^{\delta aa} fa(t) \cdot (\delta aa - t)$$

because it is proportional to $(\delta aa-t)$.

In case b), the overall loss is

$$Pf_A \cdot \sum_{t=0}^{td-ta} fa(t)(td - ta - t)$$

This shows that, as in the 0-delay model, we can pair each window size with a loss. Since the size and the loss have the same probability, we can compute the average loss.

In a further case, the number of attacks reaches *Expl* asymptotically. As an example, the number of attacks in an interval sharply increases after discovering V and then approaches zero in a few intervals after this maximum. This case may be modelled by a Weibull distribution so that the number of attacks executed at $ta+\delta t$, $\delta t > 0$, is Expl*W(δt) where

$$W(\delta t) = 1 - e^{-\left(\frac{\delta t}{\alpha}\right)^{\gamma}}$$

α and γ determine the shape of W(t) and the standard deviation that goes to zero as γ increases. In this case, the loss may be approximated as

$$\mathrm{Pf_A} \cdot \mathrm{Expl} \cdot \mathrm{vw} \cdot (1 - e^{-\left(\frac{vw}{\alpha}\right)^{\gamma}})$$

Again, the average value can be computed starting from the probability distribution of vw.

3. Further Applications of the 0-Delay Model

After discussing the case of an infrastructure with several vulnerabilities, we shows how the 0-delay model can contribute to the debate about "security through obscurity" as well as to the one on the security advantages of open source components.

3.1. Infrastructure with Several Vulnerabilities

Let us consider an infrastructure with several, mutually independent, vulnerabilities. The independence is a worst case for the defender, because the discovery of a vulnerability does not increase the probability of discovering another one. First of all, we assume that attackers and defenders may be assigned to a vulnerability. This could be considered as a contradiction

because no a priori information on the vulnerabilities is available. To solve the contradiction, we assume that attacker and defenders are allocated to components of the infrastructure. Hence, two defenders or two attackers are assigned to distinct vulnerabilities if they consider distinct components. In this way, each vulnerability is paired with just one component even if it is due to the interactions among several components. The component a vulnerability V_i is paired with determines two important parameters: the profit of an attack exploiting V_i and the probability p_i of finding V_i. If these parameters are known, the 0-delay model can compute the average loss due to the vulnerability as well as the number of defenders to be assigned to the vulnerability so that the corresponding average impact is lower than some predefined threshold.

However, the most interesting problem is the relation among the loss due to a vulnerability and the overall allocation of attackers and defenders to distinct vulnerabilities. Two cases have to be considered. In the first one the number of attackers allocated to a vulnerability is known when choosing the number of defenders searching for the same vulnerability, and the other way around. In the other, more interesting, case the allocations of attackers and of defenders are chosen simultaneously. In these cases, the allocation of a resource, i.e. an attacker or a defender, to the search for vulnerabilities can be modelled as a strategy game with two players, the attacker and the defender. The attacker manages *na* resources, the attackers, while the defender manages a pool with *nd* resources, the defenders. The move of each player defines the resources allocated to each of the n vulnerabilities.

The complete definition of the game requires that of the utility of each player. For both players, the utility always depends upon the resources allocated to each vulnerability, but alternative definitions are possible. As an example, the utility of the attacker may be equal to the average loss of the infrastructure, i.e. to the sum of the average impacts due to the vulnerabilities, while the utility of the defender may the inverse of that the attacker. This defines is a zero sum game where the loss of a player is the utility of the other one. Alternative definitions of utility may involve the probability that no loss occurs.

In all these cases, the main results of game theory, starting from the Nash equilibrium, can be exploited to define an optimal strategy for each player [22]. It is worth noticing that the worst case for the defender arises anytime a few of its resources and a large number of that of the attacker are assigned to the same vulnerability. This is a dangerous case because the 0-delay model shows that the loss sharply increases as the number of defenders goes to zero.

3.2. "Security through Obscurity" and Open Source

The 0-delay model supports the introduction of some mathematical considerations into the discussion on the "security through obscurity" philosophy. According to this philosophy, that favours proprietary solutions with respect to open source ones, security increases when no information on the infrastructure is available because this obstacles the search for vulnerabilities of the attackers. Furthermore, an attacker has to study a "live" system, which is much more dangerous. The 0-delay models the asymmetry between the attackers and the defenders by introducing the constant φ that multiplies the number of defenders so that an attacker and a defender have the same probability of finding a vulnerability. In a "closed"

solution, and if the resources of the attackers are constant, φ increases the resources of the defender to take into account the larger amount of information they can access. As a consequence, in an infrastructure exploiting a proprietary solution, if the technical skills of the attacker and defender resources are comparable, φ will be larger than one and inversely related to public information on the infrastructure or on the considered component.

On the other hand, if an open source, or at least an off-the-shelf, component is adopted, the number of defenders may become much larger because the search for the vulnerabilities may involve also other instances of the same component and other resources, besides those managed by the defender. As a counterpart, the number of people searching for a vulnerability may increase as well, because other people may be interested in attacking an instance of the same component in distinct system. In the case of a widely adopted open source component, the defender is fairly sure that, independently of the adopted strategy to allocate his/her resources, all the vulnerabilities in the component will be covered because other people are searching for them. Hence, it is highly unlikely that very few defenders are searching for a vulnerability and the dangerous case considered at the end of Sect. 3.1 will not arise. Notice that an open source component by itself can guarantee a larger numbers of attackers and of defenders only if it has been adopted in distinct systems, i.e. being open source is a necessary but not sufficient condition for larger number of attackers and defenders.

When adopting an off-the-shelf component, the number of resources searching for vulnerabilities can be actually so large that it is almost independent of those managed by, respectively, the attacker and the defender. This can be a noticeable advantage with respect to a proprietary solution anytime the number of defenders cannot be very large. Obviously, this advantage is even more critical for a small enterprise where the defenders may also have limited skills in this very specific field. On the other side, if the expected number of attackers is low and they are low skilled, moving to an off-the-shelf solution can be a disadvantage.

By considering both φ and the numbers of attackers and defenders, the 0-delay model makes it possible to compare in a quantitative way the advantages of a proprietary solution, i.e. a smaller number of attackers and defenders, against those of a widely adopted open source component, i.e. a larger numbers of both attackers and defenders. Even when the inputs of the model are a rough approximation of the real ones, some general guidelines about advantages and disadvantages of alternative solutions can be deduced from the mathematical framework underlying the models we have discussed.

5. References

1. A. Acquisti, *Privacy and security of personal information Economic incentives and technological solutions* Workshop on Economics of Information Security, University of California, Berkley, May 2002
2. R.Adkins, *An Insurance Style Model for Determining the Appropriate Investment Level against Maximum Loss arising from an Information Security Breach*, Workshop on Economics of Information Security, University of Minnesota, May 2004

3. C.J. Alberts, A.J. Dorofee. *An introduction to the OCTAVE method.* http://www.cert.org/octave/methodintro.html

4. R. J. Anderson, *Why Information Security is Hard-An Economic Perspective*, 17th Applied Computer Security Applications Conference , Dec. 2001

5. R. J. Anderson, *Security Engineering: A Guide to Building Dependable Distributed Systems*. John Wiley & Sons, Inc., first edition, 2001.

6. R. Anderson, *Security in Open versus Closed Systems-The Dance of Boltzmann, Coase and Moore,* Conf. on Open Source Software Economics, Toulouse (France), June 2002

7. P.S. Anton, R.H. Anderson, R. Mesic, M. Scheiern, *Finding and fixing vulnerabilities in information systems: the vulnerability assessment and mitigation methodology*, 1MR-1601, Rand Corporation, 2003.

8. W.A. Arbaugh, W.L. Fithen, J. McHugh, *Windows of Vulnerability: A Case Study Analysis*, IEEE Computer, Dec. 2000, p. 52-59

9. A.Arora, R. Telang, H. Xu, Optimal Policy for Software Vulnerability Disclosure, Workshop on Economics of Information Security, University of Minnesota, May 2004

10. S. Beattie, S. Arnold, C. Cowan, et al. Timing the application of security patches for optimal uptime. *16th USENIX Systems Administration Conf. (LISA 2002)*, Nov. 2002.

11. D.A.Burke, Towards a game theory model of information warfare, Master Thesis, Air Force Institute of Technology, Dec. 1999.

12. B. Carini, *Dynamics and Equilibria of Information Security Investments*, Workshop on Economics of Information Security, University of California, Berkley, May 2002.

13. R.Deraison, *The Nessus Attack Scripting Language Reference Guide* www.nessus.org.

14. B.S.Frey, S.Luechinger, A.Stulzer, *Calculating Tragedy: Assessing the Cost of Terrorism,* Inst. for Empirical Research in Economics, University of Zurich, Sept. 2004.

15. L.A. Gordon, M.P. Loeb, *The economics of information security investment*, ACM Trans. on Information and System Security, Vol. 5. No.4, pp. 438-457, November 2002

16. S.N. Hamilton, W.L. Miller, A. Ott, O.S. Saydjari, O.S. *The role of game theory in information warfare*. 4th Information Survivability Workshop, Vancouver, B.C., Canada, March 2002.

17. K.S. Hoo, *How Much Is Enough? A Risk Management Approach to Computer Security* Ph.D. Thesis, Standford University.

18. K. Kannan, R.Telang, *An Economic Analysis of Market for Software Vulnerabilities*, Workshop on Economics of Information Security, University of Minnesota, May 2004

19. I.V.Krsul*, Software Vulnerability Analysis*, Ph.D. Thesis , Purdue University

20. J. A. Major, *Advanced Techniques for Modelling Terrorism Risk*, Journal of Risk Finance, Fall 2002

21. LC. Mercer, *Fraud detection via regression analysis.* Computers &Security, vol. 9, no.4, June 1990.

22. G.Owen, Game Theory, Academic Press, 1995, Third Edition

23. E.Rescorla, *Is Finding Security Holes a Good Idea?* Workshop on Economics of Information Security, University of Minnesota, May 2004

24. S. E. Schechter, *Quantitatively differentiating system security*, Workshop on Economics of Information Security, University of California, Berkley, May 2002.

25. S. E. Schechter, *Computer Security Strength & Risk: A Quantitative Approach*, Ph.D. thesis, Harvard University, May 2004

26. B. Schneier, *Full disclosure and the window of vulnerability*, Crypto-Gram available as http://www.counterpane.com/crypto-gram-0009.html#1, September 15, 2000.

27. B. Schneier, *Closing the Window of Exposure: Reflections on the Future of Security*, Securityfocus.com, 2000, http://www.securityfocus.com/templates/forum_message.html?forum=2&head=3384&id=3384

28. G.Stoneburner, A. Goguen, A.Feringa. *Risk management guide for information technology systems*, NIST, Special Publication 800-30, Oct.2001.

29. G.Schudel, B. Wood, *Adversary work factor as a metric for information assurance*, Workshop on New security paradigms, p.23-30, Sept. 2000, Ballycotton, County Cork, Ireland.