VULNERABILITIES, ATTACKS AND EQUILIBRIUM IN INFORMATION INFRASTRUCTURES

Fabrizio Baiardi

Dipartimento di Informatica, Università di Pisa[1]

Keywords: vulnerability, vulnerability window, open source, equilibrium

Abstract

The search for vulnerabilities of ICT components is fundamental to assess the security of several critical infrastructures. To optimise the investment in the search, we introduce a zero sum game between an attacker and a defender, each managing a fixed amount of resources to be allocated to the search of vulnerabilities in distinct infrastructure components. To prevent attacks, the resources allocated by the defender search for vulnerabilities to remove them. Instead, the attacker resources search for vulnerabilities to exploit them and attack the infrastructure. Attacks results in a loss for the defender that is proportional to the time in-between the discovery of a vulnerability by an attacker resource and the discovery of the same vulnerability by a defender one. A loss is avoided only if and when a defender resource discovers a vulnerability before than an attacker one. We define conditions for Nash equilibrium where a player cannot improve its utility by changing its move only. We show that the corresponding allocation requires a large defender investment with a low return. We also discuss how the adoption of open code components may reduce the defender investment.

Introduction

Most critical infrastructures include an ICT network that interconnects the resources of the infrastructure to a set of control rooms that manage the overall infrastructure to achieve some predefined goals. This ICT network will be referred to as the **information infrastructure** paired with the considered critical infrastructure. An information infrastructure consists of several interconnected components, some developed for the infrastructure and other commercial of the shelf, cots, components adopted to reduce the overall cost. A vulnerability [7, 8, 24, 25] is any defect of a component that enables an attack against the information infrastructure. If the attack is successful, the information infrastructure is controlled, to some extent, by the attacker rather than by the owner. When the attacker controls the information infrastructure, it can control and manage the whole infrastructure. Hence, the overall security of the infrastructure strongly depends upon that of the information infrastructure that should be considered in any risk assessment of the whole infrastructure [1-6].

By extending the approach proposed in [10], this paper considers cyber attacks against the information infrastructure and the search for vulnerabilities that enable such attacks. In particular, we assume that for each component of the information infrastructure there are two sets of people searching for component vulnerability. The goal of people in one set is to remove vulnerabilities, i.e. to patch the component. Instead, people in the other for vulnerabilities to attack the infrastructure. We assume that the patching [11] of the infrastructure invalidate any previous attack so that the loss due to an attack is proportional to

[1] Dipartimento di Informatica, Universita di Pisa
L.go B.Pontecorvo 3 56127 Pisa. f.baiardi@unipi.it

the vulnerability window [8, 24, 25] of the vulnerability V enabling the attack, i.e. to time in-between the discovery of V by the attacker and the discovery of V by the defender. The window is larger than zero iff those interested in attacking the infrastructure discover V before than those that patch the infrastructure. We introduce a mathematical model that defines the average impact as a function of the vulnerability window and of the number of people searching for a vulnerability. Then, we define a strategy to allocate resources, i.e. people, to the search. To define this strategy, we introduce a zero sum game between an attacker and a defender, each managing a set of resources to be allocated to the components of the information infrastructure. Each instance of the game consists of a pair of allocations, one for each player and that defines the resources the player allocates to each component. By exploiting the mathematical model previously defined, the allocations are mapped into the average vulnerability windows and the corresponding loss of the defender. This loss is the utility of the attacker and the inverse of the defender one. We define the condition for Nash equilibrium of the game in the case where the probability that a player finds a vulnerability in a component is geometrically distributed in the number of resources it has allocated to the component [12, 16]. In Nash equilibrium neither player can improve its utility by changing its allocation only, i.e. both allocations have to be changed to improve the utility of any player.

This paper is structured as follows. At first, it introduces the framework of the mathematical model to that relates the player utilities to the allocations, i.e. to the number of resources assigned to search for vulnerabilities in the components. We also discuss the optimisation of an allocation if the one of the opponent is known. Then, equilibrium conditions are discussed together with some preliminary applications to the adoption of open code components and some generalization of the player utility functions. Lastly we draw a first set of conclusions and outline some future developments.

The importance of a quantitative evaluation of the relation between attack impacts and the investments in the search for vulnerabilities has often been stressed [5, 17, 18, 22, 26]. [23] presents a survey of current approaches to evaluate attack impacts and introduces the notion of market price of vulnerability. In an infrastructure, this price mostly depends upon the service supported by the infrastructure. [17] applies game theory to information warfare. [20] defines an insurance inspired allocation of resources to minimize the impact of a terrorist attack. The competition between defenders and attackers in the search for vulnerabilities and an optimal disclosure policy is considered in [9, 24, 25]. [9] considers the search for vulnerabilities and a social planner that decide when a vulnerability is disclosed. However, it does not consider the case where a vulnerability is discovered by an attacker.

Modelling Resource Allocation As a Game

After discussing the main underlying assumptions and constraints, we present in some details the mathematical model that will be used to define the game to evaluate alternative allocation of resources searching for vulnerabilities in the components of the information infrastructure.

Underlying Assumptions

The proposed mathematical framework is similar to the one of the *zero delay* model [10] as it assumes that both the patching and the attacks on the infrastructure occur **as soon as a vulnerability has been found**. A further significant assumption of the proposed framework is that the impact of a successful attack, i.e. the loss of the infrastructure owner, increases with the size of the window of a vulnerability V, i.e. with the time in between the discovery of V by those interested in attacking the infrastructure and the discovery of V by those interested in patching the infrastructure. This assumption is satisfied any time the goal of the attackers is an economic benefit rather the disruption of the infrastructure or the loss of human lives. Furthermore, for the moment being we assume that in each component of the information infrastructure there is only one vulnerability. This is relaxed in the following.

Our model assumes that, for each component there are two sets of people searching for the vulnerability V of the component. Let us consider at first the search of one set. The size of the set is *N* and it is fixed. The probability that each of the people in the set finds V in one time unit is *p* and it does not change with time. The assumption that *p* does not change with time is realistic provided that people searching for the vulnerability have been trained before starting the search. Furthermore, two people in the set have the same probability of finding V because the same information is available to all the people in the set.

Since the probability that at last one people in the set finds V in a time unit is $1-(1-p)^N$, the one that V is found after exactly *t* unit of time is geometrically distributed and equal to $q^{t-1}(1-q)$ where $q=(1-p)^N$. Hence, on the average, V is found at time AT where

$$AT = \frac{1}{1-q} = \frac{1}{1-(1-p)^N}$$

We consider now that there are two sets, S_1 and S_2 including, respectively, N_1 and N_2 people with probabilities p_1 and p_2 of finding V. The two searches for V are independent because the two sets are disjoint and there is no exchange of information between them, hence the average difference between the times when V is discovered by the two sets is

$$\frac{1}{1-(1-p_1)^{N_1}} - \frac{1}{1-(1-p_2)^{N_2}}$$

Since no information flows between the two sets, no information from other people is available to speed up the search. The average difference between the two times can be larger or smaller than zero depending upon the size of the two sets. By exploiting the approximation

$$(1-p)^N \approx (1-pN)$$

the average difference may be rewritten as $\frac{1}{(N_1 p_1)} - \frac{1}{(N_2 p_2)}$

if this value is larger than zero then people in S_2 discover V before than those in S_1.

Allocating Resources to Infrastructure Components

Consider now an information infrastructure including C components that are instances of *n*

component types T_1, … , T_n. Let M_i, $1 \leq i \leq n$ denotes the number of components that are instances of T_i. According to the previous assumption, for each type T_i:

- there is just one vulnerability V_i,
- two sets of peoples, A_i and D_i, search for V_i. Let Na_i and Nd_i be the sizes of the two sets and pa_i and pd_i the probabilities that a person in the corresponding set finds V_i.

Peoples in A_i, the attackers, exploit V_i to immediately attack the infrastructure while those in D_i the defenders, immediately patch the infrastructure. If an attacker discovers V_i before than the defenders, then the attack is successfully executed and it has an impact, i.e. a loss for the infrastructure owner. The impact is proportional to the size of the window for V_i and to M_i. According to the model previously discussed, the size of the window is

$$\frac{1}{(pd_i Nd_i)} - \frac{1}{(pa_i Na_i)}$$

On the average, if $pa_i\, Na_i \leq pd_i\, Nd_i$ then the defenders discover V_i before than the attackers so that $\text{Loss}(T_i)$, the loss of the owner due to V_i is zero. Instead, anytime

$$pd_i\, Nd_i < pa_i\, Na_i \qquad (1)$$

the attackers find V_i before than the defenders. In this case, the average impact is equal to

$$TypeLoss\ \ (T_i) = \alpha_i M_i \left(\frac{1}{pd_i Nd_i} - \frac{1}{pa_i Na_i} \right)$$

where α_I is a constant depending upon the infrastructure. The overall average loss AoLo is the sum of all the average losses for each type,

$$AoLo\ = \sum_{i=1}^{n} TypeLoss\ \ (T_i)$$

Obviously, only the types satisfying (1) contribute to the sum.

We define now the allocation of resources as a game with two players, the attacker and the defender, i.e. the owner. The attacker manages a pool with A resources, the defender manges one with D resources. A player allocates each resource to the search for the vulnerability of one type. An instance of the game is given by a pair of moves, i.e. of allocations, one of the attacker and one for the defender. These allocation define both Na_i and Nd_i, for each i where $1 \leq i \leq n$. Obviously, the overall number of resources a player may allocate is equal to the size of its pool. Each pair of allocations defines a value of AoLo that is equal to AU, the attacker utility. If DU is the utility of the defender, then $DU=-AoLo$. Notice that D, the size of the defender pool should be at least equal to n, the number of types. In fact, if the defender cannot allocate at least one resource to each type, the corresponding loss cannot be bounded.

Equilibrium and Optimal Allocation

This section defines how a player can optimise its allocation as soon as it knows that of the opponent. Taking into account this strategy, we consider Nash equilibrium of the game and the condition for such equilibrium.

Optimal Allocation

Consider an instance of the game previously described and suppose that both player moves are known. We say that the allocation of a player P is an optimal one if P cannot improve its utility given the N resources it manages and the opponent allocation.

We define now a procedure to compute an optimal allocation for a player P starting from the current one. After a first step, the procedure iteratively updates the allocation by shifting one resource managed by P from a type T_i to a type T_j $i \neq j$ so that a local condition is satisfied. The condition is local because it depends upon the resources allocated to T_j and T_i only. The number of iterations is bounded by the product of the number of resources and that of types.

At first, let us assume that P is the defender and consider an allocation where there is a type T_i such $pd_i\, Nd_i \geq pa_i\, Na_i$. In this case there is an **excess of defender resources** for T_i equal to

$$Nd_i - \frac{pa_i Na_i}{pd_i}$$

If in an allocation there is an excess of defender resources for T_i, then T_i does not contribute to the overall loss of the defender. When the excess for T_i is zero, there is equilibrium for T_i between the two allocations. While an excess of resources is always possible, the feasibility of equilibrium for T_i depends upon pa_i and pd_i. The following theorems hold.

Theorem 1

In an allocation where there are two types T_i and T_j such that

- *there is an excess of defenders resources for T_j larger than one, i.e.* $Nd_i - \frac{pa_i Na_i}{pd_i} > 1$
- $pd_j\, Nd_j < pa_j\, Na_j$,

the shift of one resource from T_i to T_j reduces the loss.

Theorem 2

An allocation where there is an excess of defender resources for any type is optimal for the defender.

Theorem 2 can be exploited to define an optimal allocation for the defender any time the defender pool includes at least $OD = \sum_{k=1}^{n} pa_k \frac{Na_k}{pd_k}$ resources because in this case the defender to choose an allocation where there is an excess of defender resources for any type. As proven in the following, OD is the lower amount of resources that guarantees a zero loss for the defender.

To optimise the defender allocation, at first we define Benefit(T_i), the **benefit of a defender resource for T_i** as the difference in the overall loss achieved by assign a further resource T_i:

$$Benefit(T_i) = \frac{\alpha_i M_i}{pd_i}\left(\frac{1}{Nd_i + 1} - \frac{1}{Nd_i}\right)$$

An increase of the defender resources allocated to T_i reduces the overall loss of the defender, provided that there is not an excess of resources for T_i. The change in the defender utility due to the further resource is equal to -Benefit(T_i), a monotone decreasing function of Nd_i. A further definition is that of Dloss(T_i), the **loss of a defender resource for T_i**, the increase in the loss if a defender resource allocated to T_i is shifted to a distinct type. We have that

$$DLoss\ (T_i) = \frac{\alpha_i M_i}{pd_i}\left(\frac{1}{Nd_i - 1} - \frac{1}{Nd_i}\right)$$

If there is not an excess of resources for T_i, Dloss(T_i) is always positive because the reduction of the defender resources for T_i increases the loss. Since Dloss(T_i) is proportional to Nd_i, the following theorem holds.

Theorem 3

Both the benefit and the loss of a defender resource allocated to T_j decrease with Nd_j, the number of defender resources currently allocated to T_j .

If the defender shifts one resource from T_i to T_j, the difference in the overall loss is equal to

$$Shift(i, j) = - Benefit(T_j) + DLoss(T_i)$$

We are interested in negative values of *Shift(i, j)* that reduce the overall loss. Because of Theorem 1, a negative value is achieved in the following case:

a. there is not an excess of resources for both types before or after the shift

b. $$\frac{\alpha_i M_i}{pd_i}\frac{1}{(Nd_i-1)Nd_i} > \frac{\alpha_j M_j}{pd_j}\frac{1}{(Nd_j+1)Nd_j}$$

If no pair of types satisfies condition b), the allocation is optimal for the defender. Instead, if there is a pair of types that satisfies both conditions, then a resource shift improves the defender allocation. A particular case is the one where the shift from T_i eliminates an excess of defender resources and $Dloss(T_i)$ is the contribution of T_i to the overall loss after the shift:

$$Dloss(T_i) = \alpha_i M_i \left(\frac{1}{pd_i(Nd_i-1)} - \frac{1}{pa_i Na_i} \right)$$

Since, the shift of one resource removes the excess, we have that

$$pd_i\,(Nd_i\text{-}1) \le pa_i\,Na_i \le pd_i\,Nd_i$$

This implies that

$$\frac{\alpha_i}{pd_i}\frac{M_i}{(Nd_i-1)Nd_i} > \frac{\alpha_j}{pd_j}\frac{M_j}{(Nd_j+1)Nd_j}$$

is an upper bound on $Dloss(T_i)$ and the previous condition still guarantees that the shift decreases the overall loss. To define, in the general case, the resource shifts that improve the defender allocation, consider that the following cases are possible:

1) $Dloss(T_i) > 0$, $Benefit(T_j) < 0$,
2) $Dloss(T_i) = 0$, $Benefit(T_j) < 0$,
3) $Dloss(T_i) > 0$, $Benefit(T_j) = 0$,
4) $Dloss(T_i) = 0$, $Benefit(T_j) = 0$.

Case 1) has been previously discussed. In case 2), the shift always increase the loss because it moves a resource from a type where there is not an excess to one where already there is an excess. In case 3), a resource is shifted from a type where there is an excess before and after the shift to one where there is not an excess. Since in case 4) $Dloss(T_i) = 0$ and $Benefit(T_j) = 0$, there is an excess of resources for both types before and after the shift. Hence, the shift does no influence the overall loss. Only case 1) can improve the defender allocation because a resource shift can reduce the loss provided that there is a pair of types satisfying condition b).

However, we have neglected a shift involving several resources and types simultaneously. To be able to neglect these shift, together with those that do not improve the utility but are useful an elementary step of a larger shift, we transform each defender allocation *DA* into Min(DA), the minimization of DA. Min(Da) is an allocation that minimizes the excess of resources so that, for each type T_k there is an excess of resources for T_k in DA there is also an excess in Min(DA) but the amount of resources allocated to T_k is the smallest one resulting in an excess. The saved resources are assigned, one at the time, to a type T_w such that the value of

$$\frac{\alpha_w M_w}{pd_w Nd_w (Nd_w+1)}$$

is the largest one among all the types such that there is not an excess of resources for the type. Only if there is an excess of resources for any type, the assignment can increase the excess for some type. The minimization operator saves some resources from types where there is an excess to transfer them to types that contribute to the overall loss and to the attacker utility. After applying the minimization operator, we can consider only the shift of one resource.

In any case where the attacker allocation is known, the procedure to compute the optimal defender allocation may be outlined as follows:

a. if $\sum_{h=1}^{n} \frac{pa_h Na_h}{pd_h} \leq D$ then we allocate to each type T_k the smallest amount of resources larger than $\frac{pa_k Na_k}{pd_k}$

b. otherwise
   - apply the minimization operator to shift the defender resources from types with an excess of resources to those without an excess;
   - shift one at the time a defender resource between types satisfying both conditions of Theorem 1. If no pair of types satisfies both conditions then an optimal allocation for the defender has been computed.

The overall complexity is $O(D^3 \log D)$, because each defender resource may be sequentially assigned to, at most, all the types. Furthermore, if step b) of the procedure is applied, the complexity of each assignment is $O(D \log D)$ because types have to be ordered. Lastly, the number of types is bounded by D.

To prove that the procedure is correct, we prove that distinct steps of the procedure cannot be shifted back and forward between the same types. To this purpose, consider that a resource may be shifted from T_i to T_j and then to T_i only if condition b) is satisfied between T_i and T_j before the first shift and between T_j and T_i before the second one. This occurs only if

$$\frac{\alpha_i M_i}{pd_i(Nd_i - 1)Nd_i} > \frac{\alpha_j M_j}{pd_j(Nd_j + 1)Nd_j} > \frac{\alpha_i M_i}{pd_i(Nd_i - 1)Nd_i}$$

that is impossible.

We do not discuss in details the optimisation of the attacker allocation because the cases to be considered are similar to the previous ones as the only difference between players is the sign of the utility function, as *AD=-DU*. As an example, we can define the notion of excess of attacker resources as for the defender ones and introduce the shift of attacker resources.

Equilibrium

A pair of allocations defines a game equilibrium if no shift of resources can improve a player utility. This imply that the allocations are optimal the players because they cannot increase their utility by shifting some resource.

To define conditions for equilibrium, consider an optimal allocation for the defender and assume the attacker shift all its resources to one type only. In general, the utility of each player changes. It is trivial to see that a case where the utility function does not change is the one where the defender has allocated to each type an amount of resources such that there is an excess of defender resources independently of the attacker allocation, that is where $\frac{D_i}{A} \geq \frac{pa_i}{pd_i}$ holds for any $i \in 1..n$. This is possible only if $D \geq A\sum_{i=1}^{n} \frac{pa_i}{pd_i}$.

If this condition is satisfied, the defender can allocate to each type a number of resources that guarantees that, on the average, it will find any vulnerability before than the attacker. In the following, this allocation is denoted as the **overflow allocation**. The overflow allocation is not cost effective because it assumes the worst case where the attacker focuses all its resources on just one type. Hence, if the attacker distributes its resources across several types,

some of the resources the defender allocates to a type are ineffective, i.e. they could be removed from the defender pool without increasing the overall loss.

Hence, any pair of allocations where the defender allocation is an overflow one always defines a game equilibrium. Any pair of allocations with an excess of defender resources for all the types but where the defender pool is too small for an overflow allocation is not a game equilibrium because the attacker may improve its utility by assigning all its resources to just one type. This always increases the attacker utility. In this case there is not equilibrium because one allocation is optimal for a player but the opponent on is not optimal.

It can be shown that there may be an equilibrium that not correspond to an overflow allocation. This is the one where there any pair of types T_i and T_j satisfies the condition:

$$\frac{pa_i(Na_i-1)Na_i}{pd_i(Nd_i-1)nD_i}=\frac{pa_j(Na_j+1)Na_j}{pd_j(Nd_j+1)Nd_j}$$

However, a player can exploit this condition to achieve equilibrium only if it knows the opponent allocation. Instead, an overflow allocation is independent of the opponent allocation and requires that the player knows the size of the opponent pool only.

Taking into account that the notion of overflow allocation also applies to the attacker, and that a move of a player **forces equilibrium** if it defines a game **equilibrium independently of that of the other player**, we can introduce the following theorem.

Theorem 4

A player is sure of having forced a game equilibrium if and only if

1. *it knows the size of the pool of its opponent*
2. *it can play an overflow allocation.*

Examples

First of all we notice that in the fairly common case where, for any type, the attacker and the defender resources have the same probability of finding the vulnerability, an overflow allocation is possible if $D \geq nA$. This implies that the defender is sure of having avoid a loss only if the size of its pool is *n* times that of the attacker.

Let us consider now a simple case with two types T_1 and T_2 and where:

- M_1= 10, M_2=15,
- D=30, A=20
- $\alpha_1=\alpha_2$
- $pd_1=pd_2=1/10^5$
- $pa_1=pa_2=5/10^6$

Consider a pair of allocations where $Na_1=Na_2=10$, $Nd_1=Nd_2=15$. In this case, the overall average loss for the defender is zero, because there is an excess for both types. However, a loss is possible because a game equilibrium cannot be forced. As an example, there is a loss for the defender if $Na_1=0$ and $Nd_1=25$, because $\frac{1}{pa_2Na_2}=10^{-4}$ while $\frac{1}{pd_2Nd_2}=5\cdot10^{-4}$.

If, instead D=A=20, $pd_1=pa_1$ and $pd_2=pa_2$, the optimal allocation for the defender is the one where it matches the resources of the attacker, i.e. where $Nd_1=Na_1$ and $Nd_2=Na_2$.

Consider now the case where

- M_1= 10, M_2=15, M_3=20,
- D = 45,
- A = 30,
- $\alpha_1=\alpha_2=\alpha_3$
- $pd_1 = pa_1 = 1/10^5$,
- $pd_2 = pa_2 = 5/10^6$,
- $pd_3 = pa_3=1/10^4$.

Also in this case, the defender cannot force equilibrium because $D < A\sum_{i=1}^{3}\frac{pa_i}{pd_i}$. Consider a defender allocation where Nd_1=20, Nd_2=20 and Nd_3=5 while in the attacker one $Na_1=Na_2=Na_3=10$. In this instance of the game, there is an excess of resources for both T_1 and T_2 that do not contribute to the overall loss because there are at least six resources allocated to each type. Hence, when computing the optimal defender allocation, the minimization operation will shift resources from either T_1 or T_2 to T_3 till $Nd_3>10$.

Generalization and Preliminary Applications

This section discusses some generalizations of the model as far as concerns the number of vulnerabilities of a type and alternative utility functions. Then, a first application of the results is illustrated.

Generalization

At first, we consider the presence of several vulnerabilities in a component, a case that can be handled in several ways. As an example, if we assume that all the vulnerabilities enable attacks with the same severity, i.e. attacks with similar impacts, then the existence of several vulnerabilities can be handled by properly defining the probability value of finding just one vulnerability. If, instead, the vulnerabilities have distinct severities we can model the search by considering the goal of the attackers, Some attackers will be satisfied even if they find a vulnerability that enable low impact attacks only. Instead, other attackers will neglect such vulnerabilities and focus their search for those that enable high impact attacks. In this case, several searches occur simultaneously. The proposed model can be exploited to model each search in isolation and the information it returns on the optimal allocation of resources for a single search can be used to allocate resources among the various searches. In this way, we define a two level hierarchical model. The highest level allocates resources to distinct pools, one for each severity class. The second level allocates the resources of a pool to the types of the information infrastructure components.

Alternative Utility Functions

As a first alternative utility function, assume that the defender is interested in minimizing the probability of a successful attack, i.e. the probability that the attacker resources find a vulnerability before the defenders one. To handle this case, at first we define $DFirst_i$ as the probability that a defender resource finds V_i before an attacker one. It may be proved [10] that

$$Dfirst_i = \frac{pd_i Nd_i}{pd_i Nd_i + pa_i Na_i}$$

We can now define the defender utility as $DU = \sum_{i=1}^{n} Dfirst_i$

Also in this case, the attacker utility is the inverse of the defender one, i.e. $AU = -DU$. $Dfirst_i$ is important because it allows us to define utility functions that take catastrophic attacks into account. In fact, the only defender strategy that can prevent such attacks is the one that finds any vulnerability of a component before than the attacker. The corresponding allocations can

be evaluated by utility functions defined in terms of $Dfirst_i$. Even in the case of these utilities, we can define an optimisation strategy based upon resource shifts and excess of resources. Consider, as an example, the shift of a resource from T_i to T_j such that there is no excess before and after the shift. In this case,

$$DLoss(T_i) \cong \frac{pd_i}{pd_i Nd_i + pa_i Na_i} \qquad Benefit(T_j) \cong -\frac{pd_j}{pd_j Nd_j + pa_j Na_j}$$

The update to the overall utility depends upon the sign of -Benefit(T_j) + Dloss(T_i).

A further definition of the defender utility may consider the probability that the defender finds all the vulnerabilities before than the attacker. As a consequence, the attacker utility is the probability that it finds at least one vulnerability before than the defender. The resulting game is not a zero sum one because we have that

- $DU = \prod_{i=1}^{n} DFirst_i$
- $AU=1-DU$.

In this case, a shift from T_i to T_j increases the defender utility if $\frac{Nd_i - 1}{Nd_j + 1} > \frac{Nd_i}{Nd_j}$

Some Preliminary Applications

A first important application of the previous results concerns the adoption of open code components in an information infrastructure. A component is **open code** if its code is available for peer review and analysis. This notion may be applied also to hardware-firmware components because even in this case the component behaviour depends upon some code in an internal memory. Obviously, open source components are open code ones but the inverse is not true. When evaluating the benefit of adopting an open code component vs. that of a closed code one, the main difference to be considered is the large number of resources that search for a vulnerability but that are not managed either by the attacker or by the defender. In the most general case, for each type T_i, there are two sets of people, Sa_i and Sp_i searching for the vulnerability. People in Sa_i are interested in selling the information on vulnerabilities to those interested in attacking the infrastructure, instead the results of people in Sp_i is public and may be exploited by both the attacker and the defender. Let Ca_i and Cp_i be the cardinalities of, respectively, Sa_i and Sp_i.

A quantitative evaluation of the benefits enabled by the adoption of open code components is possible only if we define the attitude of the owner with respect to the vulnerabilities found by people in Sp_i. If the corresponding patches are applied immediately, as those defined by people managed by the owner, then if T_i is open code, there are at least Cp_i defender resources and at least Ca_i attacker ones. This is rather important because, in most cases, $Cp_i >> Ca_i$ and the defender can allocate the resources it manages to other types. In general terms, the adoption of open code components may strongly reduce the overall cost of an overflow allocation so that this allocation becomes economically feasible because most of the defender resources allocated to types do not belong to the defender pool.

If, instead, the owner neglects the results of people in Sp_i because, as an example, it does not monitor the mailing lists or forum where these results are published, then the resources in Sp_i are to be considered as attacker resources only. Hence, the number of attacker resources allocated to T_i is larger than $Ca_i + Cp_i$ while the only defender resources allocated to T_i are those managed by the infrastructure owner. In this case, the investment of the owner in the resources searching for a vulnerability has to be rather large. To evaluate the overall return of

the investment, we have to take into account the reasons for not investing in discovering and evaluating the results of people in Sp_i.

Conclusion and Future Developments

We have considered the allocation of resources to the search for vulnerabilities and presented some formal results on optimal allocations and game equilibrium. These results have been applied to evaluate the adoption of open code components in information infrastructures. Other applications are possible for specific infrastructures or infrastructure components.

Among the possible developments of this work we have more accurate evaluation of the allocations chosen by the attacker and the defender that should take into account the distribution of the impact values rather than the average values only. Also other models for the search should be considered, besides the one leading to the geometric distribution of the success rate. As an example, the probability that a resource finds a vulnerability may increase with time because of the experience it has acquired. A further issue to be considered is the adoption of distinct utility functions for distinct components.

References

1. Acquisti, A. (2002) Privacy and security of personal information. Economic incentives and technological solutions, *Workshop on Economics of Information Security*, University of California, Berkley, CA, USA

2. Akins, J. (2004) 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, Twin Cities, USA.

3. Alberts, J. Dorofee, A.J. (2000) An introduction to the OCTAVE Method. http://www.cert.org/octave/methodintro.html

4. Anderson, R. J. (2001) Why Information Security is Hard-An Economic Perspective, *17th Applied Computer Security Applications Conference*, New Orleans, Louisiana, USA.

5. Anderson, R. J. (2001) *Security Engineering: A Guide to Building Dependable Distributed Systems*. John Wiley & Sons, Inc, ISBN: 0-471-38922-6

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

7. Anton, S. Anderson, R.H. Mesic, R. Scheiern, M. (2003) Finding and fixing vulnerabilities in information systems, MR-1601, Rand Corporation.

8. Arbaugh, A. Fithen, W.L. McHugh, J. (2002) Windows of Vulnerability: A Case Study Analysis, *IEEE Computer*, vol.12, p. 52-59

9. Arora, R. Telang, R. Xu, H. (2004) Optimal Policy for Software Vulnerability Disclosure, *Workshop on Economics of Information Security*, University of Minnesota, Twin Cities, USA.

10. Baiardi, F. Telmon, C. (2005) Theoretical Model for the Average Impact of Attacks Against Billing Infrastructures. *Mathematical Methods and Model for Advance Computer Network Security*, S.Petersburg. LNCS, Vol. 3685, ISBN: 3-540-29113-X

11. Beattie, S. Arnold, S. et al. (2002) Timing the Application of Security Patches for Optimal Uptime. *16th USENIX Sys. Administration Conf.* (LISA 2002), Berkley, CA, USA

12. Burke, D.A. (1999) Towards a game theory model of information warfare, Master Thesis, Air Force Institute of Technology. USA

13. Carini, B (2002) Dynamics and Equilibria of Information Security Investments, Workshop on Economics of Information Security, University of California, Berkley, CA, USA

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

15. Gordon, L.A. Loeb, M.P. (2002) The Economics of Information Security Investment, *ACM Trans. on Information and System Security*, Vol. 5. No.4, pp. 438-457.

16. S.N. Hamilton, W.L. Miller, A. Ott, O.S. Saydjari, (2002) The Role of Game Theory in Information Warfare. *4th Information Survivability Workshop*, Vancouver, B.C., Canada.

17. Hoo, K.S. (2000) How Much Is Enough? A Risk Management Approach to Computer Security, Ph.D. Thesis, Standford University, Standford CA, USA

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

19. Krsul, I.V. (1998) Software Vulnerability Analysis, Ph.D. Thesis , Purdue University Purdue West Lafayette, IN,USA

20. Major, J. A. (2002), Advanced Techniques for Modeling Terrorism Risk, *Journal of Risk and Finance*, Vol. 4, No. 1, pp. 15-24

21. G.Owen, *Game Theory*, Academic Press, 1995, Third Edition, ISBN 0-12-531151-6.

22. Rescorla, E.(2004) Is Finding Security Holes a Good Idea?, *Workshop on Economics of Information Security*, University of Minnesota, Twin Cities, USA

23. Schechter, S. E. (2004) Computer Security Strength & Risk: A Quantitative Approach, Ph.D. thesis, Harvard University, Boston USA.

24. Schneier, B (2000) Full disclosure and the window of vulnerability, Crypto-Gram http://www.counterpane.com/crypto-gram-0009.html.

25. Schneier, B. (2000) Closing the Window of Exposure: Reflections on the Future of Security, Securityfocus.com,http://www.securityfocus.com.

26. Schudel, G. Wood, B. (2000) Adversary work factor as a metric for information assurance, *Workshop on New Security Paradigms*, Ballycotton, County Cork, Ireland.

Fabrizio Baiardi is a full professor with the Dipartimento di Informatica, Università di Pisa where he coordinates several activities in the field of computer security. He has been involved in the evaluation of several computer infrastructures with high security requirements.