

Of Threats and Costs: A Game-Theoretic Approach to Security Risk Management

Patrick Maillé, Peter Reichl, and Bruno Tuffin

1 Introduction

Telecommunication networks are becoming ubiquitous in our society, the most obvious example being the success of the Internet. One of the main reasons of this success is scalability, which means that a huge network can be managed properly at no – or no significant – additional cost compared to a small one. The key issue here is decentralization of decisions over all nodes of the network. On the other hand, it is often assumed that nodes cooperate by properly using the designed protocols, but playing with their parameters could improve one node’s position in the network, at the expense of the others. For this reason, non-cooperative game theory has recently come into play in the telecommunication community to analyse selfish behaviour and try to design mechanisms with appropriate incentives.

This chapter focuses on a specific aspect of telecommunication networks: their security. Network security mechanisms aim at protecting against “natural” failures and voluntary attacks. While the former risk is related to reliability issues and can be estimated through analytic or simulation methods, the latter implies that the actions of the attacker be foreseen and countered. The choice of a security mechanism therefore depends on the defender’s knowledge of the possible attacks. On the other hand, an attacker will take into account the target’s defense strategies when determining its own attack type. Each actor therefore considers the actions of the others when

Patrick Maillé

Institut Telecom; Telecom Bretagne, 2 rue de la Châtaigneraie CS 17607, 35576 Cesson-Sévigné Cedex, France

e-mail: patrick.maille@telecom-bretagne.eu

Peter Reichl

Telecommunications Research Center Vienna (ftw.), Donau-City-Str. 1, 1220 Wien, Austria

e-mail: reichl@ftw.at

Bruno Tuffin

INRIA Rennes – Bretagne Atlantique, Campus Universitaire de Beaulieu, 35042 Rennes Cedex, France

e-mail: bruno.tuffin@irisa.fr

striving to optimize her own objective. Such interactions between actors with conflicting interests are typically the object of non-cooperative game theory.

The goal of a security mechanism is to provide the highest possible level of protection; therefore, one might think that defenders should simply choose the most complete available protection. However, an improvement in the security level often has a counterpart cost in terms of bandwidth, computational power or money that decreases the overall performance or benefit from the service. Some real-time applications, for example, cannot use the most secure mechanisms because of delay constraints and, in general, choosing a defense strategy, i.e. an appropriate security mechanism for a given service, implies a trade-off between the costs of the mechanism and the incurred risks [12]. Likewise, the best strategy for an attacker might not be to develop all known attacks, because of the cost of running those attacks, and because this would increase the likelihood of being detected. It therefore appears that the “security game” played among attackers and defenders is not trivial in general since no dominant strategies can be exhibited.

There are indeed several kinds of situations in the network security domain where actors have conflicting interests and must deploy complex strategies in order to reach the most profitable outcome. Examples of such situations include worm propagation, creation of trust networks, intrusion detection scenario learning, and reputation mechanisms. Due to the nature and complexity of the interactions among actors, game theory is particularly well suited to analyze all those cases.

Again, the use of economical concepts – and particularly game theory – to study telecommunication networks has encountered a soaring interest for the last 15 years. Several kinds of applications of game theory have yielded important evolutions in different fields, including network routing, resource sharing and flow control, power control in wireless networks, pricing, as well as incentives to cooperate in ad hoc or peer-to-peer networks. Since the fundamental nature of telecommunication networks implies the fact that several agents share a set of resources and the actions of each one may affect the others, it seems natural that game theory is perfectly adapted to depict the network externalities and help predict the outcome of the interaction among self-interested users. Another relevant aspect mixing network security and game theory is the economic relationship between users and service providers. Indeed, providers have to define the right level of security to attract and sufficiently protect customers, but the introduction of security often implies a reduction of quality of service (QoS) and therefore potentially also of demand. This trade-off has to be analysed through a proper cost model. As a result, security can therefore be an important parameter in revenue management, especially in an environment where providers compete for customers.

This chapter focuses on the game-theoretic aspects of network security issues for a broad range of scenarios, which are placed right at the cornerstone between telecommunications, economics and applied mathematics. This interdisciplinarity is particularly critical during the modelling part of user preferences in network security situations: the challenge is then to convert certain technical factors such as security mechanisms, strategies and protocols and their consequences in terms of performance into the economical concept of utility. Likewise, studying the mod-

elled situations as non-cooperative games involves competences in several fields of applied mathematics and game theory, such as optimization theory, the theory of repeated games, Markov decision processes and eventually agent-based simulation.

The remainder of the chapter is organized as follows. In Section 2, we review the basic notions of game theory, by using some simple and illustrative examples from security. Specific security games and solutions are then described in Section 3. The implications and consequences on the economic interactions between users and providers are described in Section 4. We finally conclude and describe the main challenges to be addressed in the future in Section 5.

2 A Game-Theoretic Perspective on Network Security

This section presents the fundamental concepts of game theory that will be useful in our security context. The very general principles of game theory presented in Section 2.1 allow us to define a large number of game types, with specific forms or rules. In Section 3.1 we continue with discussing related work which focuses on the simplest (non-trivial) game models, where each user has just a finite number of choices and the game is played only once. Going one step further, we then introduce and discuss, respectively, three types of more complex games that have received specific attention in the context of security, namely repeated games, stochastic games and Bayesian games, in Sections 2.2, 2.3 and 2.4.

2.1 Fundamental concepts

Game theory is a mathematical framework which allows modelling conflict and cooperation between two or more separate parties, the *players*. Players are assumed to behave *rationally*, i.e. they are triggered by the selfish incentive of maximizing their individual benefit, which is usually expressed in terms of a *utility function*. During the game, which follows certain *rules*, players can choose and implement a *strategy* from a set of different behavioural options, the so-called *strategy space*, in order to maximize the *payoff* they are receiving as an outcome of the game.

Hence, formally a game is described by the number n of players, their strategy spaces and their payoff functions, S_i and u_i , respectively, for each player i ($1 \leq i \leq n$):

$$G = \{n; S_1, S_2, \dots, S_n; u_1, u_2, \dots, u_n\}. \quad (1)$$

Based on that description, game-theoretic analysis attempts to understand the probable behaviour of the players, regarding their strategy choice, and thus to determine the presumable outcome of the game. In some cases, this work relatively straight forwardly, for instance, if each of the players can identify a “dominant strategy”, i.e. a strategy with which this player is better off independently of the behaviour of his opponents and which directly leads to an equilibrium situation. A much broader

equilibrium concept, the so-called *Nash equilibrium*, is achieved if an operation point is reached where each player is giving her best response facing her opponents' strategies, i.e. for none of the players there is a unilateral incentive to change her strategy, given that the strategies chosen by all opponents are fixed. Formally, if $s = (s_1, \dots, s_n)$ is the profile of strategies with $s_i \in S_i$ and if s_{-i} stands for the profile of strategies excluding player i , a Nash equilibrium is a profile s (with $s = (s_i; s_{-i}) \forall i$) such that $\forall 1 \leq i \leq n$,

$$u_i(s) \geq u_i(t; s_{-i}) \quad \forall t \in S_i.$$

In other words, the strategy of each player i is a best reply to the strategies of the others.

Note that individual elements of the strategy space S_i are called *pure* strategies, whereas a *mixed* strategy can be described as a linear combination of two or more pure strategies, with weights summing up to 1 which may be interpreted as the probability distribution $\pi_i = (\pi_{i,t})_{t \in S_i}$ for player i choosing randomly among the pure strategies involved. The goal is then to determine for each player i the probability distribution maximizing the *expected* utility $\sum_{j=1}^n \sum_{s_j \in S_j} u_i(s_1, \dots, s_n) \prod_{k=1}^n \pi_{k,s_k}$.

A Nash equilibrium in mixed strategies is then a set of probability distributions $(\pi_i)_{1 \leq i \leq n}$ such that $\forall i$ and any other probability vector $\tilde{\pi}_i = (\pi_{i,t})_{t \in S_i}$,

$$\sum_{j=1}^n \sum_{s_j \in S_j} u_i(s_1, \dots, s_n) \prod_{k=1}^n \pi_{k,s_k} \geq \sum_{t \in S_i} \sum_{j \neq i} \sum_{s_j \in S_j} u_i(t; s_{-i}) \tilde{\pi}_{i,t} \prod_{k \neq i} \pi_{k,s_k}.$$

Whereas the existence of a pure strategy Nash equilibrium cannot be guaranteed, it can be demonstrated that any static game with a finite number of players and finite strategy spaces has at least one Nash equilibrium in mixed strategies, i.e. a profile of distributions such that the choice of each player maximizes its expected payoff or utility.

Within the taxonomy of games, the distinction between static and dynamic games is worth mentioning: whereas in static games, all players simultaneously choose a strategy without further knowledge about their opponents' decisions, dynamic games are characterized by a sequence of moves, where each player gives an answer based on the entire history of the game. Repeated games are a specific class of dynamic games and basically represent a sequence of static games. More details on those games are given in Section 2.2.

Based on these introductory remarks, we will discuss more specific game-theoretic issues at various places in the rest of the chapter; for a detailed comprehensive introduction we kindly refer to standard textbooks like [10, 21]. Similarly, for applications in telecommunication networks, involving a lot of participants with different interests, the reader is advised to consult [3] and references therein for an overview of the type of problems that can be modelled in this way, including routing, resource allocation and queueing management. Game theory has indeed appeared as

a promising tool to study interactions in that context. A last related issue is *network pricing*, which can fruitfully be studied using game theory tools (see [8]).

In the rest of this section, we will introduce the different kinds of security-related games in more detail and discuss specific examples.

2.2 Repeated Games

Repeated games are a simple way to include the time aspect into game theory. Based on a classical one-shot game, we assume that the same game is played repeatedly for a number T (possibly infinite or random) of periods.

Traditionally, it is assumed that players value more the present than the future periods, which is modeled by considering the discounted sum $\sum_{t=1}^T \delta^{t-1} u_i(s_t)$ as the overall payoff function of each player i , where $\delta \in [0, 1]$ is called the discount factor, s_t is the (possibly mixed) strategy profile played at period t and u_i the corresponding utility for player i .

In repeated games, a strategy for a user describes the action choice she should make depending on the whole history of past actions. Repeated games thus allow us to model some kind of reputation effects, where the past actions of a player can be sanctioned or rewarded by her opponents. Such games have interesting properties, such as the fact that the set of equilibria can become quite large (this result is known as the “Folk Theorem” [21]).

2.3 Stochastic Games

Like repeated games, stochastic games also have a form of memory, but in a more complex fashion. The game is still played over (discretized) time, but memory is represented by a *state*, which in the context of network security can, for example, describe the current feature of the data (not compromised, compromised, stolen), the type of application used, the ongoing attacks and activated countermeasures.

The state space traditionally is assumed to be finite, where each state corresponds to

- some payoff value for each player,
- an action set for each player, and
- some *transition probability* value for each state that depends on the actions taken by the players at the current period. The value corresponding to a state is the probability that the system be in that state for the next period.

As for repeated games, one has to define an overall payoff function for each player, that is classically chosen as a discounted sum of the per-period payoffs.

Stochastic games are hard to study analytically, since the number of states increases exponentially with the number of players and of strategy choices. Therefore, the equilibria of such games are often computed numerically. On the other hand, those games allow us to model some quite rich scenarios. They are therefore

well suited for some types of attacks like intrusion, which usually implement a sequence of attack steps before reaching their goal.

2.4 Bayesian Games

Bayesian games are characterized by *incomplete* information about the opponents (e.g. their payoff functions). In a security context, this is used to model the difference between malicious attackers and non-malicious ordinary users who are accessing the system regularly. To model this, we assume that in the game the system is facing a subject, where the subject can be of one out of several types (for instance, malicious or not in the most simple case, but there can be more types, such as malicious users with different interests), and users with different available security services. This type (and the corresponding payoff function) is assumed to be private knowledge to the subject, whereas the system only can have a certain belief on that, e.g., a probability distribution between malicious/non-malicious users. In the course of the game, players can update their initial beliefs because of the actions observed according to Bayes' rule.

The so-called signalling games are a particularly interesting example of Bayesian games. Here, there is an informed player (agent) knowing the type of the opponent (principal). The principal is unaware of the agent's type and has to start the game with an initial belief. During the game, however, the principal is able to update her initial belief based on signals originating from actions of the agent, until the principal eventually manages to deduce the type of her opponent.

This type of games requires an extension to the concept of an equilibrium: Thus, a Bayesian Nash equilibrium is defined as a strategy profile together with a probability distribution characterizing the belief about the types of the opponents, which maximizes the expected payoff assuming that strategies and beliefs of the other players are fixed. Note that, similar to normal-form games, the existence of one or more Bayesian Nash equilibria can be proved when the numbers of pure strategies are finite, since Bayesian strategies can be interpreted as mixed strategies.

3 A Closer Look into Specific Security Games

In this section, we present some security games/contexts that have been introduced in the literature. As we shall see, several types of interactions can be modelled as a non-cooperative game, depending on the aggregation level, services, timescale and types of attacks considered. We are going to present specific games illustrating the most significant interactions when dealing with security. The most basic kind of game between an attacker and a defender is analysed in Section 3.1, where each player has the choice between two strategies: doing nothing or launching a costly attack or detection procedure. The Nash equilibrium (in mixed strategies) is determined. A more complicated game is then described in Section 3.2, where more actions are possible for each player and the information available to players may be incomplete. The interactions between the attacker and the defender can then be

studied as a repeated Bayesian game, whose study allows us to identify the most relevant parameters in the attack and defense strategies. In the same family of games, one can try to incentivize the defenders in a network to participate to a common defense strategy, for the best of the whole network. Such games are described in Section 3.3. While those three sections are for direct interactions among defenders and attackers, no information about the network topology is used either in player payoffs or in their strategies. Section 3.4 considers games on networks, where attacks and defenses have to be placed appropriately on the different links. Likewise, attacks can be performed by worms, for which the trade-off is between a discrete and a fast propagation to maximize the dissemination. Such worm propagation games are presented in Section 3.5. Section 3.6 highlights the problem (not fully modelled yet) of interdomain incentives for confidentiality when an intermediate node is supposed to forward the traffic of his neighbours.

Note that all the games presented in the previous sections have applications in security modelling. Repeated game models have been used in other works, for instance, in [1] to represent the interactions of nodes in a wireless mobile ad hoc network (MANET). Here the attacks considered are only passive attacks, which consist of some nodes refusing to act as a relay for the communication flows of the others. The repeated feature of the game allows us to build mechanisms that sanction non-contributing nodes, in order to create incentives for collaboration. Likewise, in [28] the fact that the game has a form of memory is used to detect and isolate malicious nodes among a whole set of wireless ad hoc nodes. Similarly, the stochastic game approach is used in [17] to model attacks directed against specific applications, against communication capacities, or against databases. Numerical studies lead to some mixed Nash equilibria, i.e. at some state(s) players should choose their action according to a given probability distribution. References [25, 29] apply stochastic game models as well, to study other specific attack scenarios, and numerically compute mixed Nash strategies. The Bayesian approach is also considered in security games, for example, in [23]. There, the authors consider two types of attackers, namely “normal users” whose behaviour is only driven by selfishness and “malicious nodes” that intend to maximize the damage done. The defender has an a priori belief about the probability of its opponent being of each type and updates that probability as it observes the attacker moves. That belief is also used to determine the probability of using the detection mechanisms.

3.1 Play Detection/Attack or Not: Mixed Nash Equilibria

In [2], Alpcan and Başar¹ introduce a model where the attacker has two choices, i.e. launching an attack or doing nothing, while the defender’s choices are to trigger or not its (costly) detection scheme. The authors observe that for reasonable values

¹ In [2], the authors actually first define a cooperative game where nodes in a sensor network should collaborate to improve intrusion detection, but we do not describe this model since this chapter focuses on noncooperative games.

of the payoffs for each outcome, the game has no Nash equilibrium in pure strategies. This is easy to see: if the attacker always attacks, then the victim always defends, so the attacker would be better off not attacking; on the contrary if the attacker never attacks, then defending is only costly for the victim that should thus never defend, which precisely makes attacks profitable to the attacker. A convenient way to visualize this is to represent player utilities depending on their actions in a matrix where player 1 (defender) actions correspond to rows, player 2 (attacker) to columns and the terms in the matrix are written in the form (u_1, u_2) with u_i the utility (payoff) of player i . An example of this so-called *normal form representation* is given in Fig. 1. In that example, a, b, c, α, γ are all positive numbers. We assume here that triggering the detection mechanism (resp. launching the attack) is costly to the defender (resp. the attacker), whereas doing nothing has no cost. In general, the cost for the defender of missing an attack is much larger than the cost of running the detection scheme, i.e. $c \gg b$. We now see that as soon as $c > 2b$ there exists a unique Nash equilibrium in mixed strategies: we denote by π_{def} the probability of the defender triggering the detection scheme and by π_{att} the probability of the attacker launching the attack. To have a Nash equilibrium with $0 < \pi_{def} < 1$, i.e. positive probability for both possible choices, the utilities are

	Attack	No attack
Trigger detection	$(a, -\alpha)$	$(-b, 0)$
No detection	$(-c, \gamma)$	$(0, 0)$

Fig. 1 A two-player attacker–defender game in normal form: the defender chooses a line and the attacker a column

$$u_{def}(\pi_{def}, \pi_{att}) = a\pi_{def}\pi_{att} - b\pi_{def}(1 - \pi_{att}) - c\pi_{att}(1 - \pi_{def}),$$

$$u_{att}(\pi_{def}, \pi_{att}) = -\alpha\pi_{def}\pi_{att} + \gamma\pi_{att}(1 - \pi_{def}).$$

Computing the conditions $\partial u_{def}/\partial \pi_{def} = 0$ and $\partial u_{att}/\partial \pi_{att} = 0$ gives, respectively, $\pi_{att} = \frac{b}{a+b+c}$ and $\pi_{def} = \frac{\gamma}{\alpha+\gamma}$. Note that another interesting view/interpretation to get the relations obtained from $\partial u_{def}/\partial \pi_{def} = 0$ and $\partial u_{att}/\partial \pi_{att} = 0$ is that the defender should be indifferent between triggering the detection scheme or not, in terms of expected payoff (otherwise he would simply choose the best strategy). This also gives

$$\pi_{att}a - (1 - \pi_{att})b = -\pi_{att}c + (1 - \pi_{att}) \times 0, \quad (2)$$

where the left-hand (resp. right-hand) side of (2) is the defender's expected payoff if he triggers the detection scheme (resp. does nothing). Similarly we also get from the opposite side

$$-\pi_{def}\alpha + (1 - \pi_{def})\gamma = \pi_{def} \times 0 + (1 - \pi_{def}) \times 0.$$

The (existing and unique) Nash equilibrium of the game, therefore, corresponds to

- the defender choosing to trigger the detection mechanism or do nothing with respective probabilities $\frac{\gamma}{\alpha+\gamma}$ and $\frac{\alpha}{\alpha+\gamma}$ and
- the attacker choosing to launch the attack or do nothing with respective probabilities $\frac{b}{a+b+c}$ and $\frac{a+c}{a+b+c}$.

Interestingly, in such games the mixed strategy choice of each player is made such that its opponent has no preference among its possible actions, so that it can also choose a mixed strategy.

Alpcan and Başar then extend that kind of model by considering several types of attacks and the corresponding defense type has to be chosen by the defender to detect the intrusion. Another interesting extension proposed in [2] consists in considering that before choosing its defense strategy, the defender has an imperfect knowledge of the type of attack chosen by the attacker: the set of possible attacks is partitioned into sets and the defender knows to which set the attack (if any) belongs. Some payoffs for each player correspond to each situation in terms of attack presence, attack type, defense trigger and defense type.

Jormakka and Mölsä [14] present some very similar but concrete situations of network security (to be more precise: information warfare) games, with specific numerical values, that also lead to simple strategy sets. The specific examples introduced allow us to exhibit some particular outcomes and phenomena. The so-called *evildoer game* has the same form as the basic model of [2] (see Fig. 1), i.e. it has two players – an attacker and a victim – with two possible choices each and no pure Nash equilibrium. The conclusions for that game also hold for another interpretation of the game, called *vandal game* in [14]: here Jormakka and Mölsä do not consider defense strategy, but only the fact that the victim will simply not use the service (say, a network), and thus not suffer from the attack. Then the same reasoning as in Fig. 1 is valid: since the attacker’s objective is to maximize victim’s harm, then it should not always attack (but only with some probability) because then the victim would simply avoid the service.

The same kind of attacker–defender game is studied in [5]. The number of strategies for each player is larger than 2: several attack and countermeasure types are considered. Moreover, the actual payoffs corresponding to some given strategic choices are not deterministic, since attacks are supposed to succeed with a probability that depends on the activated countermeasures. Nevertheless, players are assumed risk-neutral, i.e. only sensitive to payoff expectations, so introducing success probabilities does not change the game type. The modelling effort made in this chapter to quantify the payoffs for each player is worth mentioning:

- The attacker is assumed to be sensitive to a *return on attack* criterion that involves some financial equivalents of the value of a successful attack, the costs of building and launching it and its success probability.

- The defender acts so as to maximize some *return on investment* that is calculated based on the monetary cost of the countermeasures, the value of the good to protect, the potential impact of an attack and the attack success probability.

All the games mentioned above have no pure Nash equilibrium, thus only mixed strategies lead to equilibria: players would then randomize their action choice, according to a specific probability distribution as we did for the example of Fig. 1.

The game presented in [22, 23] has the same features than [2, 14], but introduces an interesting refinement: the defender might not know what kind of attacker he is facing. More precisely, the “attacker” can either be a regular network user that simply may not want to offer some service (the *passive attack* discussed above) or a badly intentioned actor that possibly launches *active attacks*. The defender has an a priori knowledge of the probability of the attacker being of one type or the other and may update those probability values based on some observations of the attacker actions or messages, according to Bayes’ rules. Again, the resulting Bayesian game applied to intrusion detection does not exhibit any pure Nash equilibrium.

3.2 Incentive-Based Attacker Modelling

In their fundamental paper [16], Liu et al. introduce a systematic method to model attacker intent, objectives and strategies (AIOS), based on combining the incentives of an attacker as well as his cost into a single utility function. Moreover, they propose a game-theoretic formulation of AIOS in order to capture also the relationship to the objectives and strategies of the defender and allow for inferring AIOS automatically.

To this end, Liu et al. [16] start from the basic assumptions that security attacks are usually intentional (i.e., planned), that both attacker and defender only possess incomplete information about their respective opponent and that the success of an attack is always relative to the protection level of the attacked system (and vice versa). The attacker intents can vary widely, but may be subsumed under the notion of an incentive which is assumed to be quantifiable. Typical examples are, to be quantified with the same units, the amount of profit earned, the amount of terror or damages caused, directly or due to no-show of users because of the threat. Together with certain constraints like attack cost or risk of detection, the resulting utility function describes the objective of the attacker and is supposed to be maximized by the attacker.

Modelling attacker strategies is considered to be more sophisticated, as they have to account for a sequence of potentially very different actions which determine a series of battles between attacker and system. This may lead to extraordinarily complex strategy spaces, and also comparing different attack strategies is far from being trivial, as the efficiency in terms of system security degradation strongly depends on countermeasures performed by the system.

The formalization of these AIOS models starts from perceiving the attacker and likewise the environment (comprising the non-malicious users) as peers of the

system under attack. The system is separated into a production-oriented service part and a security-related protection part and is assumed to actively take defense actions. Then, attacks are described as games between rational attackers and defenders whose Nash equilibria allow us to infer attacker strategies, whereas deriving the intention and objectives of the attacker is based on detecting strategic patterns which are matched against insights gained during a learning phase. Together with related accuracy and sensitivity analyses, this is supposed to significantly advance the risk assessment of security attacks.

Eventually, this general approach leads to a more fine-grained taxonomy of AIOS models along two orthogonal dimensions, i.e. the correlation among attack actions and the accuracy of intrusion detection. Whereas a low correlation of attack actions suggests the application of Bayesian repeated games, high correlation leads to (potentially multi-stage) dynamic game models.

The paper concludes with an instructive case study modelling attacker strategies for a distributed denial-of-service (DDOS) attack on a system which is countered by the popular *pushback* mechanism, i.e. by identifying and rate limiting those packet flows that cause the DDOS attack. To this end, user traffic is classified as either good (non-malicious), bad (malicious) or poor (non-malicious, but with the same properties as malicious traffic). Assuming a unique attacker together with multiple legitimate users, in the corresponding repeated Bayesian game (see Section 2.4) the system is uncertain about the type of each user and may only resort to a respective probability distribution. The action space of the attacker consists of several DDOS attacks, the action space of the legitimate user includes a variety of network applications and services and the action space of the system is determined by the potential defense postures of each router (specified by a large set of characteristic parameters like congestion checking time, target drop rate, rate-limit time, maximum session number). As far as the utility functions are concerned, the attacker's utility depends on the impact of the attack on both the system and the legitimate users, whereas the utility of the non-malicious users boils down to the relative availability of the system. Finally, the utility function of the system is determined by the trade-off between the absolute impact of the DDOS attack and its relative impact on the system availability.

As this game is way too complex for an analytic treatment, Liu et al. [16] present extensive simulations based on ns-2 where a total of 11 defense strategies are investigated. Legitimate user traffic is based on real-world Internet traces, whereas the attacker's action space is determined by the number of "zombies" (i.e. hosts controlled by the attacker) as well as varying attack traffic patterns and total volumes. For the resulting 64 different possible attack strategies, the corresponding average payoffs for attacker, legitimate users and system are calculated and analysed. Whereas some resulting insights are widely consistent with the existing mainstream opinion, e.g. on the impact of total zombie number or drop rate preferences, also some surprising consequences may be drawn: For instance, neither rate nor pattern of the attacking traffic is of significant relevance for the attacker's payoff function, but only the number of zombies and the properties of the traffic aggregate matters. Similarly, the simulation results allow a clear identification of the relevant defense

parameters of the system. Finally, a total of 42 different Nash equilibria have been calculated and allow further inferences for the attacker's strategies, for instance, with respect to traffic patterns or the optimal ratio between bad and poor traffic, and even leads us to bounds for the attacking capacity of the attacker (i.e. the worst-case damage caused) and the assurance capacity for the defending system (i.e. the resilience against DDOS) which are of central relevance for any risk assessment purposes.

3.3 *Passive Attacks in Collaborative Networks: Enforcing Cooperation on Defenders*

As previously mentioned, a *passive attack* is the action of a network participant refusing to provide some service. In peer-to-peer file sharing networks, a passive attack would consist of offering no files to the community. Likewise, in wireless ad hoc networks, a node refusing to transfer packets is considered as making a passive attack.

It is true that those passive, free-riding attacks are not motivated by a desire to harm a machine, a network or a system, but rather simply by user selfishness. However, it is also reasonable to consider those noncooperative behaviours as attacks, since the system does not work anymore if too many participants do not contribute to it. More directly, a node which does not participate in the collective security by refusing to provide useful information can be considered as a passive attacker.

In that context, the objective is to incentivize players to contribute to the service, either through sanctions or through rewards. Some appropriate mechanism thus has to be defined, such that rational and selfish players are better off participating to the service provision. This implies, for example, building reputation scores based on past behaviour and using those scores to possibly exclude misbehaving nodes from the system [19]. The goal is to prevent passive DoS attacks that consist of simply not participating to security (if we place it into our context instead of specifically MANETs in [19]). If non-participating nodes are isolated from the network, a reputation mechanism can, at a low cost, enforce participation. Formally, assume we have N nodes (players) and that the utility of node i ($1 \leq i \leq N$) depends on both his payoff y_i and the relative share $\sigma_i = y_i / (\sum_{j=1}^N y_j)$ by

$$\alpha_i u(y_i) + \beta_i r(\sigma_i)$$

with $u()$ differentiable, strictly increasing and concave and $r()$ is differentiable, concave and maximized at $1/N$. Weights $\alpha_i, \beta_i \geq 0$ characterize node i . If k nodes cooperate, this induces a (network) benefit $B(k)$ (increasing and concave) and a cost $C(k)$ (such that $kC(k)$ is increasing) for implementing the procedures; thus, a payoff $y_k = B(k) - C(k)$. Reputation is included in functions $B()$ and $C()$. Conditions for a Nash equilibrium to occur can be derived. Under proper conditions on $B()$ and $C()$, it can be ensured that at least half of the nodes will cooperate.

Those attacks are also addressed in [1], where the proposed mechanism is evaluated using a repeated game model.

3.4 Routing Problems or “Cat-and-Mouse” Games

We now describe some intrusion detection games played on a physical network, where strategies involve some routing decisions. Since the paradigm and modelling are quite different, we present them in a separate section devoted to “security routing games”. In this section, we consider games that are played over the links (or node interfaces) of a network. The strategy sets are either a single link in the network (chosen to carry out an attack, or to put an attack detection device) or a whole routing strategy (choice of flow or attack spreading among different available paths). In those games, the attacker can try to intercept normal traffic (he is then the cat), or to reach a destination while avoiding detection (he is then the mouse). Likewise, according to the considered service, the network manager chooses to place specific detection mechanisms to protect important links and/or provide a higher security in general.

Those interactions can be modelled as two-player zero-sum games, i.e. games where the gain of one a player is necessarily the loss of the other one: if player 1 gets U_1 then $U_2 = -U_1$. The (possibly mixed) Nash equilibria for that game are such that the corresponding utilities (U_1^N, U_2^N) verify

$$U_1^N = -U_2^N = \max_{s_1} \min_{s_2} U_1(s_1, s_2) = \min_{s_2} \max_{s_1} U_1(s_1, s_2),$$

where $s_i, i = 1, 2$, is a mixed strategy for player i , i.e. a distribution probability over the strategy set of player i .

Kodialam and Lakshman [15] consider such a game between an attacker trying not to be detected and an active defender. The attacker is located at some point a of the network and his target location is denoted by t . The goal of the attacker is to select a path to send his malicious packet so as to minimize the detection probability. To do so, he might choose some highly loaded links in order to become less detectable. (The background traffic on each link e is denoted by f_e .) On the other hand, the defender’s objective is to select which links to scan so as to maximize that detection probability (subject to a constraint B in the total number of scanned bits per time unit).

Then the authors prove that the Nash equilibrium value of the detection probability is $B/M(f)$, where $M(f)$ is the maximum possible flow from a to t on a network assuming each link e has capacity f_e . Also, the player strategies at Nash equilibrium are derived:

- If m_i denotes the flow on the i th path from a to t for the maximum flow mentioned above, the attacker chooses to use that path with probability $m_i/M(f)$.

- The defender selects a *minimum cut* of that maximum flow, which is therefore made of links e where the maximum flow is f_e . Then the defender chooses to scan each of those links e with probability $Bf_e/M(f)$.

The model and results are also extended to the case where the attacker can choose among several points to origin the packet from and to the case of several potential targets.

A model with roles somehow inversed is also of interest. In [6], Bohacek et al. consider a user willing to send some flow from one point to another, through a network with vulnerable links: if the attacker decides to attack a link ℓ (e.g. for eavesdropping) used by a user packet, then there is a probability p_ℓ that the packet gets intercepted. The strategies are thus as follows:

- The attacker has to spread his scanning effort among the links.
- The defender has to choose routes for his flow. He actually uses stochastic routing, i.e. determines a distribution over the next-hop possibilities for each node (avoiding cycle possibilities).

Two different types games are studied:

1. *Online games*: The attacker can scan one physical interface at each node and therefore chooses a probability distribution over the interfaces, for each node.

For each link, the transfer delay τ_ℓ is augmented by T_ℓ if the packet is intercepted. The objective of the attacker is to maximize the total expected transfer time. Then the authors express the Nash equilibrium as a saddle point and show that the corresponding equilibrium strategies can be computed in a distributed way.

2. *Off-line games*: Now the attacker only chooses *one link* to perform his attack. A strategy for the attacker is therefore a probability distribution over all links. The attacker's objective is to maximize the probability of intercepting user packets, which shall be minimized by the defender (zero-sum game).

To include path lengths into players objectives, the authors add a penalty related to the path length. More precisely, they define the variable χ_ε as being 0 if the packet does not get intercepted, and $(1 + \varepsilon)^{t-1}$ if it is during the t th hop. The attacker's (resp. defender's) objective is to maximize (resp. minimize) the expected value of χ_ε .

The authors show how the saddle point can be computed, using the solution of a flow maximization problem in a network where the capacity constraint on link ℓ is p_ℓ . This solution is interestingly similar to the one obtained in [15] for reversed roles.

Note that the ε parameter tunes the system according to the user preferences for short paths with respect to security. In particular, if ε is small then the defender will fully exploit the path diversity in the network by spreading his flow along all paths, whereas he concentrates on shortest paths when ε increases.

3.5 Worm Propagation Games

Another important domain where it is believed that game theory could be applied is the case of worm propagation [11]. Network worms are autonomous intrusion agents that have created tremendous financial losses to users due to their propagation through the Internet. The first major worm was the Morris worm in 1988 which crippled a substantial proportion of the Internet [26]. As another example, the Slammer SQL worm infected over 90% of the vulnerable hosts within just 10 min [20].

Security managers create patches, but in general those need to be developed manually and require some time: to first identify the problem, check that the patch does not have side effects, and then distribute it. For this reason, worm containment procedures are being developed. We focus here on scanning worms for which an infected node scans the address space at a given rate and infects nodes which it manages to locate. Indeed, to propagate, a worm tries out many IP addresses to be sent to and infect the corresponding host. Since those IP addresses are somehow randomly chosen, many of the ones tried do not respond.

The approach for representing and analysing worm propagation is characterized by fluid models, which can adequately represent a large population of vulnerable hosts. For a given population size N , assuming that once infected, a node remains infected forever, the evolution of the number of infected nodes at time t , I_t follows in its simplest form (this equation being potentially different depending on the kind of worm) the (epidemiologic) differential equation

$$\frac{dI_t}{dt} = \beta I_t (N - I_t).$$

In this equation, β is a parameter representing the rate of infection of vulnerable nodes by a given infected node. The equation depends on not only the number of infected nodes (who send the worm), but also the remaining nodes to be infected (which will become less likely to be reached). Some variations of this equation will, for instance, describe whether or not worms are sent to uniformly chosen IP addresses or “closely chosen” ones. This kind of equation is typical of a worm’s propagation when the effects of human counteractions and network congestion are ignorable. We then experience a slow-start phase due to few nodes sending the worm and a slow-finish phase because at the end the remaining nodes are very few. In the slow-start phase (the one of interest for detectors, since we are interested in finding times such that I_t/N reaches say 5%), $\frac{dI_t}{dt} = \beta N I_t$, whose solution is $I_t = I_0 e^{-\beta N t}$. Countermeasures affect the rate at which nodes are infected. This is represented, for instance, in [11] by a reduction factor θ_t which affects the scanning rate of a host which has been infected for t time units. The equation now becomes in the slow-start phase

$$\frac{dI_t}{dt} = \beta N \left[I_0 \theta_0 + \int_0^t I_s \theta_{t-s} ds \right],$$

whose solution is $I_t = I_0 + \int_0^t I_{t-s} \beta N \theta_s ds$. The epidemics will spread or die out exponentially fast depending on whether the integral $\int_0^t I_{t-s} \beta N \theta_s ds$ is larger or smaller than 1.

The interplay between worm strategies and detection/containment techniques can then be described as a game, the worm trying to infect the network as much as possible, while the network tries to slow it down. An important characteristic is that we are in presence of a Stackelberg game with the worm as leader, playing first its strategy, to which the detection and containment technique responds (the follower). This situation makes the worm powerful in the sense that, taking into account the optimal strategies of defenders, he can decide the strategy that optimizes his own interest, i.e. the infection rate. The typical goal of such an analysis is to prevent global spread before patches are developed and distributed (i.e. a given fixed amount of days).

In [11], the strategy of players is to choose the best quarantining strategy for the detector, while the worm chooses a scanning rate. Quarantining means that after some time τ , an infected host's connection attempts are blocked. We then have $\theta_t = P[\tau \geq t]$. There are also throttling mechanisms reducing the rate at which a node makes new connections when considered suspicious. For Williamson's throttle, connection requests are processed at rate c connections per second. If the rate for generating non-wormy connection attempts is w , the slow-down factor is $\theta_t = c/(\beta + w)$.

Payoff is the speed of spread (the growth exponent of the epidemic), which has to be maximized for the worms and minimized for the detector. The number of unsuccessful scans can therefore be used to detect worms, as was suggested by Ganesh et al. [11] who study the game played between the worm designer setting the scanning rate and the worm detector setting the detection threshold for considering a host as infected. Detection is performed through a CUSUM (cumulative sum) test, minimizing the time between infection and detection for a given false-positive rate. It declares a node infected at R if the log likelihood ratio of being infected to being uninfected over a length k interval in the past exceeds a threshold c : $R = \inf\{n : \max_{1 \leq k \leq n} \sum_{i=k}^n \ln(f_1(t_i)/f_0(t_i)) \geq c\}$, where the t_i inter-failure times and f_0 (resp. f_1) is the density, assumed here exponential, under normal (resp. infected) conditions. A detector can be designed to restrict the growth rate to no more than a value ν , while simultaneously ensuring that the false alarm probability over a specified time window T does not exceed a specified threshold. Interestingly in [11], the optimal detector is such that the worm growth exponent is insensitive to the scanning rate. As a consequence, the leader, the worm, does not have actually a significant influence.

This paper rightfully stresses the importance of game theory for worm containment, pursuing in that direction, where the impact should be important. Several directions can be exploited to study this kind of games, extending the set of available strategies to the attacker and/or the defender. For example, one could imagine that a properly chosen proportion of the IP addresses used by the worm are chosen from the host's recently contacted ones, in order to be detected later while still replicating at the same speed. Some other complicated strategies, involving scanning rates that

change over time, could also be considered. It is important to note that scanning worms are not the only kind of worms; there exist many other types. For instance, routing worms which uses BGP routing tables to only scan the Internet routable address space, which allows them to propagate three times faster than a traditional worm and which can produce selective attacks. In a similar way to what was done in [11] for scanning worms, it is clearly of interest to investigate the games that can, or more exactly need to, be introduced between worm mechanisms and detection procedures for each specific type of worms and design more efficient reaction and defense strategies.

Note that another potential level of game has been introduced in [27]. Instead of looking at the game between a given worm and security tools, we can also look at a larger timescale the race between worm writers and security managers. Indeed, when a worm is circumvented, a new one generally appears needing a new fight. Such a game is therefore a game for survival, in order to *stay in the game*. In [27], a parallel is made with biological nature, and evolutionary game theory is described as the appropriate tool.

3.6 Security/Confidentiality Issue in Interdomain and Ad Hoc Networks

Another issue brought in by Chandramouli [7] is the *inter-domain* and ad hoc network case. A user/domain is expecting its traffic to arrive at destination with an appropriate level of security/confidentiality. But this traffic often needs to be forwarded by other providers/nodes that could behave maliciously. How to create incentives for a proper behaviour in this case? This kind of problem has been extensively studied in the literature to yield economic incentives to indeed forward traffic (see, for instance, [4, 9, 13]), but not much was related to security/confidentiality incentives. It is suggested in [7], but not solved yet, to play a repeated game, such that if a node defects in providing the expected security/confidentiality, its own traffic will also be unsecurely forwarded as a sanction, for at least a fixed amount of time. If the sanction is long enough, this should prevent the nodes from misbehaving.

4 Economics of Security

Besides the direct game-theoretic modelling of interactions between malicious attacks and protection strategies, it has to be emphasized that security brings new economic issues to network service providers, because of its growing importance for companies or users. This has to be analysed mathematically and, again, can be treated by game theory not only to represent the business interactions between a provider and its customers first, but also to represent competition among providers for customers having to choose between different offers. The growth of a network such as the Internet has had a positive externality from a business point of view, but has also a negative externality when talking about security. As pointed out in

[18], “businesses have a strong incentive to seek profit from users (consumers) while cooperating – and competing – in the provision of privacy and security.” This common sense statement has to be verified though. Security can be provided at the network layer (with protocols such as Secure Sockets layer (SSL)) or at the application layer, but can be limited by the government public policy [18]. Note that security and economics bring the problem of secure payment that will not be dealt with here [24].

4.1 Model Based on Risk Percentage

We could, for instance, assume that a provider has different *initial* security levels (or classes) $\ell \in \{1, \dots, L\}$, to which an intrusion risk r_ℓ is associated, with $r_{\ell_1} < r_{\ell_2}$ for $\ell_1 < \ell_2$ and a price p_ℓ with $p_{\ell_1} > p_{\ell_2}$ for $\ell_1 < \ell_2$. Security levels may correspond to various options concerning the availability of hardware or software security. Demand splits among the different classes, but an important characteristic of security attacks is that the larger number of customers on a class, the more likely new attacks will happen (according to Metcalfe’s or a power law), decreasing therefore the *actual* security level. Assuming non-atomic users, demand can then be characterized by a so-called Wardrop equilibrium, i.e. a combination of price and actual security risk which is the same for all classes having positive demand (otherwise users would have an interest in switching) and some classes have a null demand because too expensive for the proposed level.

A typical situation for this kind of models is virus scan softwares, where different softwares can have different efficiencies but are also sold at different prices. If many users are known to use a typical software, then attacks will basically concentrate on this population in order to reach more people.

Two situations can be considered: first all the levels are managed by a single provider (a monopoly) which then tries to maximize its revenue by paying with prices and second the case where each security level is handled by an independent provider, and providers compete for customers at a higher level by playing on prices (an oligopoly).

Typical game-theoretic analysis of security management offers can be built this way.

4.2 Coalitions

In the case of competing security service providers, the question of cooperation is probably more relevant than in many other fields. Indeed, due to the interactions among users, low security provided by a competitor induces a risk for its own customers, and therefore a lower security level. Coalition formation can thus become efficient for providers, in terms of reputation and revenue. It is therefore very interesting to model and investigate the incentive for forming such coalitions, and whether or not a full cooperation is the best solution for all providers. Such studies

would then involve tools from collaborative game theory to study the sustainability of coalitions and the effect of revenue repartition on the sustainable coalitions.

5 Conclusions

Whereas it has become clear that modelling and analysis of telecommunication networks security through non-cooperative game theory is of paramount importance, this approach is nevertheless still in its infancy and has indeed attracted interest only recently. As one of the key issues, we have identified the understanding of the interactions between malicious users (attackers) and end users or the network manager who expect a secure connection. Different such types of interactions have been introduced and discussed in this chapter, dealing, for instance, with intrusion detection, denial-of-service attacks or worm propagation, to mention just a few examples. In any case, the ultimate goal is to understand the equilibrium situation and therefore to try to design schemes or strategies to drive this equilibrium towards the most secure situation. This can be done by the introduction of proper incentives, for instance. We have also highlighted that security additionally brings economical issues for the providers in their relationships with users to propose the most profitable contracts, as well as in the competition between providers; those relationships can again be analysed within the framework of non-cooperative game theory.

Note that most of the games presented here and in the literature are of a rather basic form, mainly due to the novelty of the issue. We have pointed out that therefore a lot of work remains to be done to represent practical scenarios as closely as possible. On the other hand, the constant evolution of networking technologies requires to adapt the presented issues and raises new challenges to be tackled. Thus, summarizing what has been said so far, we consider this game-theoretic perspective on various security issues of significant interest for the research community as well as of key practical importance for future industrial applications, and sincerely hope that the presented survey will manage to further stimulate research in this seminal field.

Acknowledgments The authors acknowledge the support of European initiative COST IS0605, Econ@tel. Part of this work has been supported by the Austrian government and the city of Vienna in the framework of the COMET competence centre program and by the French research agency through the FLUOR project.

References

1. Agah A, Das SK (2007) Preventing DoS attacks in wireless sensor networks: A repeated game theory approach. *Int J Netw Secur* 5(2):145–153
2. Alpcan T, Başar T (2003) A game theoretic approach to decision and analysis in network intrusion detection. In: *Proceedings of the 42nd Conference on Decision and Control*, Maui, HI
3. Altman E, Boulogne T, El-Azouzi R, Jiménez T, Wynter L (2006) A survey on networking games in telecommunications. *Comput Oper Res.* 33(2)

4. Andereg L, Eidenbenz S (2003) Ad hoc-VCG: A truthful and cost-efficient routing protocol for mobile ad hoc networks with selfish agents. In: Proceedings of the 9th Annual International Conference on Mobile Computing and Networking (MobiCom 2003), San Diego, CA, USA, pp 245–259
5. Bistarelli S, Dall’Aglia M, Peretti P (2006) Strategic games on defense trees. In: Proceedings of the 4th International Workshop on Formal Aspects in Security and Trust (FAST’06), LNCS 4691, Hamilton, Ontario, Canada, pp 1–15
6. Bohacek N, Hespanha JP, Lee J, Lim C, Obraczka K (2007) Game theoretic stochastic routing for fault tolerance and security in computer networks. *IEEE Trans Parallel Distrib Syst* 18(9):1227–1240
7. Chandramouli R (2007) Economics of security: Research challenges. In: Proceedings of the 16th International Conference on Computer Communications and Networks (ICCCN’2007), Hawaii, USA
8. Courcoubetis C, Weber R (2003) Pricing communication networks—economics, technology and modelling. Wiley, Chichester
9. Feigenbaum J, Papadimitriou C, Sami R, Shenker S (2002) A BGP-based mechanism for lowest-cost routing. In: Proceedings of the 21st ACM Symposium on Principles of Distributed Computing, Monterey, California, USA, pp 173–182
10. Fudenberg D, Tirole J (1991) Game theory. MIT, Cambridge, MA
11. Ganesh A, Gunawardena D, Jey P, Massoulié L, Scott J (2006) Efficient quarantining of scanning worms: Optimal detection and co-ordination. In: Proceedings of IEEE INFOCOM 2006, Barcelona, Spain
12. Gordon LA, Loeb MP (2002) The economics of information security investment. *ACM Trans Inf Syst Secur* 5(4):438–457
13. Hershberger J, Suri S (2001) Vickrey prices and shortest paths: What is an edge worth? In: Proceedings of the 42nd IEEE Symposium on Foundations of Computer Science, Las Vegas, Nevada, USA, pp 252–259
14. Jormakka J, Mölsä J (2005) Modelling information warfare as a game. *J Inf Warf* 4(2):12–25
15. Kodialam M, Lakshman TV (2003) Detecting network intrusions via sampling: A game theoretic approach. In: Proceedings of IEEE INFOCOM, San Francisco, CA, USA
16. Liu P, Zang W, Yu M (2005) Incentive-based modeling and inference of attacker intent, objectives, and strategies. *ACM Trans Inf Syst Secur* 8(1):78–118. doi: <http://doi.acm.org/10.1145/1053283.1053288>
17. Lye KW, Wing JM (2005) Game strategies in network security. *Int J Netw Secur* 4(1–2):71–86
18. McKnight L, Solomon R, Reagle J, Carver D, Johnson C, Gerovac B, Gingold D (1997) Information security for internet commerce. In: McKnight LW, Bailey JP (eds) Internet economics. MIT, Cambridge, MA, pp 435–452
19. Michiardi P, Molva R (2002) Game theoretic analysis of security in mobile ad hoc networks. Tech. Rep. RR-02–070, Institut Eurécom
20. Moore D, Paxson V, Savage S, Shannon C, Staniford S, Weaver N (2003) Inside the slammer worm. *IEEE Secur Priv* 1(4):33–39
21. Osborne MJ, Rubinstein A (1994) A course in game theory. MIT, Cambridge, MA
22. Pacha A, Park JM (2004) A game theoretic approach to modeling intrusion detection in mobile ad hoc networks. In: Proceedings of IEEE Workshop on Information Assurance and Security, West Point, NY, USA, pp 30–34
23. Pacha A, Park JM (2006) A game theoretic formulation for intrusion detection in mobile ad hoc networks. *Int J Netw Secur* 2(2):131–137
24. Racz P, Stiller B (2006) A service model and architecture in support of ip service accounting. In: Management of integrated end-to-end communications and services, Proceedings of the 10th IEEE/IFIP Network Operations and Management Symposium, NOMS 2006, Vancouver, Canada, April 3–7, 2006. IEEE, pp 1–12
25. Sallhammar K, Helvik BE, Knapskog SJ (2006) A game-theoretic approach to stochastic security and dependability evaluation. In: Proceedings of the 2nd IEEE Intl Symposium on Dependable, Autonomic and Secure Computing (DASC), Indianapolis, IN, USA

26. Seeley D (1989) A tour of the worm. In: Proceedings of the Winter USENIX Conference, San Diego, California, USA
27. Somayaji A (2004) How to win an evolutionary arms race. *IEEE Secur Priv*, 2(6):70–72
28. Theodorakopoulos G, Baras JS (2008) Game theoretic modeling of malicious users in collaborative networks. *IEEE J Select Areas Commun* 26(7):1317–1327
29. Wang H, Liang Y, Liu X (2008) Stochastic game theoretic method of quantification for network situational awareness. In: Proceedings of the International Conference on Internet Computing in Science and Engineering (ICICSE), Harbin, Leilongjiang, China, pp 312–316