We are IntechOpen,
the world's leading publisher of
Open Access books
Built by scientists, for scientists

4,800
Open access books available

122,000
International authors and editors

135M
Downloads

Our authors are among the

154
Countries delivered to

TOP 1%
most cited scientists

12.2%
Contributors from top 500 universities

BOOK CITATION INDEX
CLARIVATE ANALYTICS
INDEXED

WEB OF SCIENCE™

Selection of our books indexed in the Book Citation Index
in Web of Science™ Core Collection (BKCI)

Interested in publishing with us?
Contact book.department@intechopen.com

Numbers displayed above are based on latest data collected.
For more information visit www.intechopen.com

Security as a Game – Decisions from Incomplete Models

Stefan Rass, Peter Schartner and Raphael Wigoutschnigg
Alpen-Adria Universität Klagenfurt
Austria

1. Introduction

Securing a computer system is always a battle of wits: the adversary tries to locate holes to sneak in, whereas the protector tries to close them. Transmitting messages through a publicly accessible medium whilst having the content concealed from the adversary's eyes is traditionally accomplished using mathematical transformations. These are practically irreversible, unless some additional information – called the key – is available, making the secret accessible for the legitimate holder(s) of the key. Ever since the concept of perfect secrecy has rigorously been formalized by Shannon (1949), it has been known that unbreakable security is bought at the cost of keys that equal the message in terms of length. In addition, the key is required to be random and must be discarded immediately after usage. This pushed the concept of unconditional security out of reach for implementation in computer networks (though diplomatic and military applications existed), until 1984, where the idea of quantum cryptography was born by Bennett & Brassard (1984). The unique feature of this novel type communication is its usage of information carriers other than electrical pulses. By encoding bits in the polarization plane of single photons, the information becomes essentially not cloneable, as Wootters & Zurek (1982) have shown, and any attempt can be detected. This rendered the one-time pad practical in real-life electronic networks and unconditional security no longer needed to remain a dream.

Classical cryptography widely relies on unproven conjectures regarding the difficulty of solving computational problems. The field of public key cryptography draws its power from the infeasibility of reverting simple algebraic operations within large finite groups, but no proof has yet been discovered that rules out the existence of efficient algorithms to solve those problems. The sole indicator of security is thus the absence of any publication proving the assumptions wrong. But there is yet no other indication than pure hope for this to be true. Symmetric techniques, although conceptually different, come with no better arguments to support their security. Although these may lack much of the structure that public key systems enjoy and are thus harder to analyze, a rigorous proof of security or mathematical framework for proving security is also not available.

In this work, we attempt taking a step towards providing a rigorous and easy-to-use decision-theoretic framework for proving security. Results are formulated with applications to quantum networks, but we emphasize that the framework is in no way limited to these.

1.1 The problem of perfect end-to-end secrecy

Quantum cryptography claims to bring perfect secrecy to a given line, but speaking honestly, it is no more than this. Using a carrier that is sufficiently fragile to rule out copying it, naturally raises the question of how much distance can be bridged? In fact, nowadays available quantum cryptography allows for communication over a distance of up to 144 km, as demonstrated by Schmitt-Manderbach et al. (2007), but arbitrary distances can yet not be bridged. Although theoretical results due to Lo & Chau (1999) indicate that the noise problem can be overcome, making arbitrary ranges theoretically possible, building networks is inevitable for a global roll-out. Existing solutions mostly rely on trusted relay for that matter. However, why attack the quantum line, if attacking a relay node is fully sufficient?

Under the assumption of perfectly protected lines, recent results indicate that without pre-existing secrets that are exclusively known to the sender and the receiver, end-to-end-security is only achievable under hard constraints on the network topology. To be more precise, let G be a graph that models a network. Let $V(G)$, $E(G)$ be the sets of vertices and edges of G, and assume the sender s and receiver r to be parts of G, that is $\{s, r\} \subseteq V(G)$. The adversary can be modelled by a set $A \subseteq 2^{V(G)\setminus\{s, r\}}$ (the powerset of $V(G)\setminus\{s, r\}$), that is we assume that a selection of subsets of vertices can be compromised. If k such sets can become conquered simultaneously, then we face a k-active adversary. An infected vertex v is assumed fully under the adversary's control, so a message passing through v can be read, blocked or modified and v is free to create as many new messages as desired. There is no limitation on computational power or knowledge of the adversary.

If removing from G the vertices in any k sets in the adversary structure A cannot disconnect s and r in G, then we call the graph $A^{(k)}(s,r)$-subconnected. If, by doing so, the network cannot be disconnected at all, then the graph is said to be $A^{(k)}$-subconnected.

Referring to these notions, a network permits perfectly secure message delivery from s to r if and only if the graph G is $A^{(2)}(s,r)$-subconnected. The reader may consult Ashwin Kumar et al. (2002) for a proof. Different, yet no less stringent requirements are imposed by Wang & Desmedt (2008): among related results, the following necessary condition best highlights the difficulty of achieving unconditional security in a real-life network: if for $u \geq 1$, $3(k - u) + 1 \geq k + 1$ directed node-disjoint paths from s to r exist, then a necessary condition for perfectly secure message transmission from s to r against a k-active adversary is that there are u directed node disjoint paths (these u paths are also disjoint from the $3(k - u) + 1$ paths from s to r) from r to s.

The described adversary model applies to many situations, as for example machines running certain software may all suffer from the same security holes. Networks equipped with devices from different vendors may be considered vulnerable if one vendor's devices turn out to be insecure. A k-active adversary would correspond to k vendors cooperating, or equivalently arise, if k vendors obtained the same malicious module from a single fraudulent manufacturer, turning a heterogeneous set of products into a possible backdoor for an adversary.

1.2 Decision theory and system security

Many results either guarantee or rule out perfectly secret communication, but this might not be satisfactory. If perfectly secure communication is not possible, then how much is achievable with the given resources? A variety of security metrics has been proposed, but a measure of security is yet missing. This work summarizes a decision-theoretic approach to quantifying risk in terms that can be specified to best suit the application at hand.

Protecting business assets is the core goal that security engineers are in pursuit of, so measuring the quality of a protection mechanism in terms of values of the protected asset is certainly a more convincing argument than hoping that relaying nodes in quantum networks are trustworthy, or no efficient solution algorithm for some computational problem has yet been discovered.

The problem of measuring security has been tackled by a vast number of authors. Assessing security is commonly achieved by security metrics or scores, whereas the latter is considered for sole comparative purposes and does not have an interpretation on its own. The common vulnerability scoring scheme (see Houmb & Franqueira (2009) is one example for a scoring technique. Other taxonomies like proposal of Innerhofer-Oberperfler & Breu (2009) are as well subjective and may help decision-makers, but are not designed to support a further mathematical treatment.

Decision support systems like CAULDRON by Mas (2008) cook up reports generated by vulnerability scanners to boil down a vast amount of information to a manageable lot of recommendations. The models we describe can naturally benefit from these systems, and are thus considered an add-on for a standard topological vulnerability analysis (cf. Jajodia et al. (2005) for details on the latter). In particular, our results will generalize the assertions about the (im)possibility of perfectly secure communication as cited above. An approach that is closely related to our model has been given by Ying et al. (2006) and Mavronicolas et al. (2005). These approaches consider less general scenarios than we do, and suffer from the need for accurate adversary models. We demonstrate how this requirement can elegantly be dropped, while simultaneously simplifying a subsequent analysis.

2. Modelling

The modelling approach proposed in the following requires identification of security primitives of a given network. Its core ingredient is an enumeration of possibilities for transmission and parameter selection, and its output will be a game-theoretic model. For convenience of the reader, we review some necessary basics of game theory and multipath transmission, to illustrate the required input for a powerful model

2.1 Game-theoretic foundations

It is useful to collect some tools from game-theory that will help establishing the results presented here. A *(non-cooperative n-person) game* $\Gamma = (N, S, H)$ is a triple composed from a set $N = \{1, 2, \ldots, n\}$ of players being able to choose actions from their corresponding strategies within the set of sets $S = \{PS_1, PS_2, \ldots, PS_n\}$, such that the *i*-th player, when taking the action $s_i \in PS_i$ from his set PS_i of possible pure strategies, receives the payoff $u_i(s_i, s_{-i})$, where $u_i \in H$ and s_{-i} denotes the vector of pure strategies chosen by *i*'s opponents. The set H thus comprises the set of payoff functions for each player. A probability distribution over the set PS_i is called a *(mixed) strategy*. We will exclusively refer to mixed strategies in the following, and denote the set of distributions over PS_i as S_i (note that the set of pure strategies is included in the set of distributions by considering a pure strategy as a Dirac-mass located at the pure strategy s_i). A *(Nash-)equilibrium* is a strategy-profile $s^* = (s_1^*, \ldots, s_n^*)$ such that

$$u_i(s_i, s^*_{-i}) \leq u_i(s^*_i, s^*_{-i}) \quad \forall i \in N.$$

In other words, no player can benefit by solely deviating from the equilibrium strategy. The possibility of a gain when several players cooperate is not ruled out however. This is not topic of this work.

If all strategy sets are finite (assumed in the following), then the utility for a mixed strategy is the expected (average) utility over an infinite number of repetitions of the game. In other words, if (x, y) are the strategies (discrete probability distributions) of player 1 and 2, respectively, then the expected utilities are given by the bilinear forms

$$u_1(x, y) = x^T A y \text{ and } u_2(x,y) = x^T B y,$$

where $A \in \mathbb{R}^{n \times m}$, $B \in \mathbb{R}^{n \times m}$ (for $n = |PS_1|$, $m = |PS_2|$) are the *game-matrices*. The full two-player game is denoted as the triple $\Gamma = (\{1, 2\}, \{S_1, S_2\}, \{A, B\})$. If $A = -B$, then the game is called zero-sum, and its *value* $v(\Gamma)$ is given as the value of the function at the saddle-point, which is

$$v(\Gamma) = \max_x \min_y x^T A y.$$

It can be determined upon linear optimization, as described by Schlee (2004).

2.2 Multipath transmission

We have already summarized two results characterizing the possibility of perfectly secure message transmission in the introduction. A popular for circumventing the person-in-the-middle attack is relying on several paths, over which messages are propagated independently. We shall not burden ourselves with the intricate details of error-correcting codes and how these relate to the concepts of secret sharing, and refer the reader to McElice & Sarwate (1981) for details. Recent protocols embodying the ideas of correctable shares for multipath transmission are found in the work of Fitzi et al. (2007) and Wang & Desmedt (2008), and we confine ourselves to remarking that perfectly secure message transmission is possible under a few assumptions:

1. An encoding is available that allows to divide a message into pieces such that any subset (of pieces) of limited cardinality does not provide any information about the secret itself. This is achieved by standard secret-sharing, as we will summarize later.
2. The network topology ensures the existence of several node-disjoint paths that connect any two nodes in the network. Results from graph theory (see Chartrand & Zhang (2005)) characterize suitable networks. Procedures for building such topologies from scratch have been developed in Rass (2005), and algorithms for determining optimal extensions of existing networks have been devised by Rass & Schartner (2009b).

Error-correction facilities that are inherently available within some secret-sharing schemes can be exploited to further increase security and reliability, however, are not a must for our upcoming considerations.

Let us review a simple form of secret sharing here that will become a theoretical asset for later conclusions. Given a secret $s \in [0, n-1]$, choose t random numbers $r_i \in [0, n-1]$, and set $r_{t+1} = s \oplus r_1 \oplus \ldots \oplus r_t$, where $\oplus$ denotes the bitwise exclusive or operation. It is evident that unless all values $r_1, \ldots, r_t, r_{t+1}$ are known, the secret remains one-time pad encrypted with the exclusive-or of the unknown components, and thus perfectly concealed. The values $r_1, \ldots, r_t, r_{t+1}$ are the *shares* that arose from s. The core idea of multipath transmission is to send each share over its own channel that does not intersect any other channel in the network. Unless an adversary has $(t+1)$ nodes conquered, no information about s can leak out. Practical multipath transmission protocols utilize a more sophisticated form of secret-

sharing, where shares are created as points on a chosen (random) polynomial. Unless a sufficient number of such points are known, the polynomial, and therefore the secret it embeds in its constant term, remains protected from the adversary's eyes. The advantage over the previously sketched scheme is its resilience against loss of shares up to the extent of the threshold. This comes at the price of higher computational effort, as calculations have to be performed in large finite fields.

The methodology that is presented in the following naturally captures a much wider range of situations; however, we stick with a multipath scenario for illustrative purposes.

2.3 Setting up the model

Given a network at hand, mapping it into a model that permits decision-theoretic treatment proceeds in several steps. Each step is expanded below, starting with a definition capturing some terminology. Notice that in the sequel, we explicitly consider secret and reliable transmission, which is assumed available in various ways over the given network. Degrees of freedom that are available to the sender of a message comprise the following: transmission paths, encoding schemes (including encryption) and protocol parameters. We will assume a multipath transmission scenario (for otherwise perfect secrecy over multi-hop connections is ruled out under weak conditions as shown previously), and take the encoding to be fixed (as prescribed by the hardware devices). However, we can determine the path through the network.

Definition 1: A *pure strategy* is a set of node-disjoint paths that connect a sender Alice to a receiver Bob. The set of all pure strategies is denoted as *PS*. A *mixed strategy* is a probability distribution over *PS*. We denote the set of all such (discrete) distributions over *PS* as *S*, and refer to $x \in S$ simply as a *strategy*.

Speaking in game-theoretic terms, we refer to a pair of honest instances Alice and Bob as player 1, and call player 2 the adversary. Consequently, the sets of pure strategies are PS_1 and PS_2, with corresponding strategy sets S_1 and S_2. The methodology comprises five steps:

Step 1. **Identification of pure transmission strategies**: The expert shall enumerate all degrees of freedom that a sender enjoys when initiating a transmission. This in particular includes all sets of node-disjoint paths that can be used for multipath transmission. All meaningful choices are collected in the set PS_1 of pure strategies.

Step 2. **Identification of pure attacks strategies**: The expert shall enumerate all nodes that are vulnerable to an attack. This could be an assumption of the number of nodes that can simultaneously be compromised or a more complex adversary structure. In particular, this analysis should account for software security flaws that could be exploited. The finite set of options that are open to the adversary makes up the set PS_2 of pure strategies.

Step 3. **Setting up the utility taxonomy**: The expert shall specify a scoring scheme that applies to the outcome of a transmission. Examples include the binary set $I = \{0, 1\}$ with 0 meaning failure of a transmission, and 1 indicating a successful secret delivery. Finer discrete or even continuous scales can be based on a message priority ranking, or on the amount of Shannon-entropy that a message is tied to. We will frequently use the set $I = \{0, 1\}$ in the following.

Step 4. **Setting up the model matrix**: For every combination $(s_i, s_j) \in PS_1 \times PS_2$, identify the outcome of an attack and assign to the variable u_{ij} the score according to the ranking *I*. Notice that $u_{ij} = u_1(s_i^1, s_j^2)$ is the utility for the honest parties. For instance,

with $I = \{0, 1\}$, if a successful transmission using the pure strategy $s_i \in PS_1$ would resist an attack via strategy $s_j \in PS_2$, then we set $u_{ij} = 1$. Otherwise we would write $u_{ij} = 0$ to indicate the adversary's success. The *model matrix* is denoted as A and has the entries u_{ij}.

Step 5. Analysis and Conclusions: The model matrix is the sole ingredient for any further analysis of the model. Conclusions are obtained from the results to follow.

Example 1: To illustrate the modelling process, consider a network topology as shown in Fig. 1 with two instances Alice and Bob who wish to communicate.

Modeling step 1: Assume that Alice and Bob have picked three pairs of shortest node-disjoint paths, disregarding other possibly longer paths. So player 1's set of pure strategies is denoted as $PS_1 = \{s_1^1, s_2^1, s_3^1\}$, and given by

- s_1^1: Use paths (s, 1), (1, 2), (2, t) and (s, 3), (3, 6), (6, t),
- s_2^1: Use paths (s, 1), (1, 2), (2, t) and (s, 4), (4, 5), (5, t)
- s_3^1: Use paths (s, 3), (3, 6), (6, t) and (s, 4), (4, 5), (5, t)

Modeling step 2: Eve strategies for attacking are given by $PS_1 = \{s_1^2, s_2^2, s_3^2\}$, where

- s_1^2: Compromise nodes 1 and 3,
- s_2^2: Compromise nodes 1 and 4,
- s_3^2: Compromise nodes 3 and 4,

modeling a situation in which two out of three vulnerable nodes {1, 3, 4} can be compromised simultaneously.

Modeling step 3: The utility taxonomy is chosen as $I = \{0, 1\}$, where 0 indicates failure of a secret transmission, and 1 means success. This scale considers loss of any secret content equally harmful.


Fig. 1. Example Network Topology

Modeling step 4: Writing down every possible combination of pure strategies in a matrix, with entry 1 if the attack fails, we end up with the following table, directly representing the utility function $u_1: PS_1 \times PS_2 \to \{0, 1\}$, specified by a tableau (game-matrix, model-matrix):

u_1	s_1^2	s_2^2	s_3^2
s_1^1	0	1	1
s_2^1	1	0	1
s_3^1	1	1	0

The final step is the formal analysis of the model. We defer this until the formal results have been presented.

3. Decisions from incomplete models

An accurate game-theoretic model would call for the specification of the adversary's payoffs in order to optimally count his intrusion attempts. Unfortunately, we have no method of reliably eliciting the intentions and benefits that an attacker gains. Furthermore, we may be unable to observe our unknown opponent's payoffs at all, which rules out any chance of learning the adversary's payoff structure from experience. The game-theoretic model is thus *incomplete* in two respects:

1. We have no way of reliably determining the utility for the second player.
2. We have no mechanism of detecting our own success, nor can we observe the adversary's success. This may not apply for scenarios in which the adversary is active, so that an intervention can be detected, but a passively eavesdropping intruder will most likely remain undetected.

The remedy is switching to a zero-sum game, assuming the adversary's intentions and payoffs to be the precise opposite of our own ones. Though intuitively evident, the validity of this approach is formally founded (see Rass & Schartner (2009a) for a proof):

Lemma 1: Let $\Gamma = (\{1, 2\}, \{S_1, S_2\}, \{A, B\})$ be a bi-matrix game. Set $n = |PS_1|$, $m = |PS_2|$ and let A, B be the payoff matrices $A \in \mathbb{R}^{n \times m}$ and $B \in \mathbb{R}^{n \times m}$ for players 1 and 2, respectively. Let $\Gamma_0 = (\{1, 2\}, \{S_1, S_2\}, \{A, -A\})$ be the zero-sum game from player 1's perspective, i.e. with the payoff of player 2 being the negative payoff of player 1 (disregarding the unknown matrix B), and let v denote the value of the game Γ_0. Then

$$v \le (x^*)^{\mathrm{T}} A y^*$$

for all existing Nash-Equilibria (x^*, y^*) of the game Γ.

This is the formal permission to use $(-A)$ as a substitute for the adversary's payoff, for getting a lower bound on the achievable utility. In other words, unless the adversary's purpose is truly opposite to our own one, we can only be off better than expected. Also, the bound cannot be improved, as examples by Rass (2009) demonstrate. In the following, we denote a random variable X with discrete distribution x by writing $X \sim x$.

Definition 2 (Loss): Let $i \in \{1, 2\}$ denote a player in a two-person game with pure strategy set PS_i, and S_i denoting the associated (mixed) strategy space. Assume that the utility function $u: PS_1 \times PS_2 \to I \subset \mathbb{R}^+$ to be a mapping into a compact set I. The *loss* is a random variable L measuring the difference between the actual and the possible outcome under the chosen pure strategies. It is defined as $L := (\max I) - u(X, Y)$, where $X \sim x \in S_1$, $Y \sim y \in S_2$.

Based on this, we can define *risk* as the expected loss. This is in alignment with the definition of risk as the product of probability and damage, as used by the German BSI (2008), as well as Hammer (1999).

Definition 3 (Risk): With the notation from Definition 2, player i's *risk* (for $i \in \{1, 2\}$) when choosing the strategies $x \in S_1$ and $y \in S_2$, is defined as the expected loss under this strategy choice, namely $r_i(x, y) = E(L) = (\max I) - E_{x,y}(u_i(X, Y))$, where for the random variables X and Y have the discrete distributions x and y, respectively, and the risk is dependent on the choices of player i's opponent.

It is straightforward to reformulate Definition 2 and Definition 3 for more than two entities. However, this general formulation is not required in the sequel, and thus omitted. The core concept upon which we can analyze security in a decision-theoretic sense is introduced through

Definition 4 (Vulnerability): Let $A \in I^{n \times m}$ be the model matrix set up by Alice and Bob, where these two have $n = |PS_1|$ pure strategies for communicating, facing an adversary with $m = |PS_2|$ pure strategies to choose from. Assuming that the set $I \subset \mathbb{R}^+$ to be compact, the *vulnerability* $\rho(A)$ is defined as

$$\rho(A) := (\max I) - v(\Gamma_0),$$

where $v(\Gamma) = \max_x \min_y x^T A y$ is the value of the associated zero-sum game (see Lemma 1) $\Gamma_0 = (\{1, 2\}, \{S_1, S_2\}, \{A, -A\})$ and S_1, S_2 are the probability spaces over PS_1, PS_2, respectively.

The vulnerability is directly derived from the game-matrix A, which we shall refer to as the model matrix in the following. Summarizing the construction, this matrix displays the utility for the honest parties, for each possible (pure) transmission and (pure) attack strategy.

Example 2: The value of the game with the matrix given previously is found as $v = 1/3$, so that the vulnerability comes to $\max I - v = 1 - 1/3 = 2/3$.

4. Results

An immediate consequence of the definition of vulnerability and Lemma 1 is the following result. We refrain from stating the proofs for the cited assertions and refer the interested reader to the work of Rass (2009).

Theorem 1 (Rass (2009)): If the message valuation scale is binary in the sense that every message scores 1 when being delivered successfully, and zero when getting deciphered by Eve, then $\rho(A)$ is the maximum probability of a concealed message becoming disclosed.

Capturing utility in terms of entropy permits quantification of the expected leak of information. In other words, the decision-theoretic approach yields a measure of secrecy-capacity of a network via a corollary to Theorem 1:

Theorem 2 (Rass (2009)): For a (random) message M with Shannon-entropy $H(M)$, the amount h by which the adversary's uncertainty (Shannon-entropy) is decreased upon a (secret) communication between Alice and Bob satisfies

$$h \leq \rho(A) \cdot H(M)$$

where $\rho(A)$ is the vulnerability, and the model matrix is set up with the binary scale $I = \{0, 1\}$ (i.e. a 1 in the matrix corresponds to one successful secret delivery, and 0 means failure).

Example 3: Knowing that the vulnerability of our example model is 0.667, this is the maximum probability of having a secret communication disclosed when communicating over the network. Theorem 2 states that no more than two thirds of any secret information will leak out in the long run average.

Whereas Theorem 1 and Theorem 2 capture long-term average secrecy of a channel, the decision on whether or not the next transmission should be started calls for a measure of risk for a single concrete transmission. If L is the random variable that measures the loss (i.e. the difference between the maximum utility and the actual utility) of the next transmission, then the next result upper-bounds the probability for loosing more than indicated by the vulnerability (in accordance with Theorem 1).

Theorem 3 (Rass (2009)): Let the secret communication between Alice and Bob be modelled by a bi-matrix game Γ and let Γ_0 be the associated zero-sum game as in Lemma 1. Let A be

the model-matrix. If Alice and Bob play an equilibrium strategy for Γ_0, then for $\varepsilon \geq 0$, the random loss $L \in I$ satisfies

$$\Pr(L > \rho(A) + \varepsilon) \leq \exp\left(-\frac{2\varepsilon^2}{C}\right)$$

where $\rho(A)$ is the vulnerability and the constant C is the solution of the optimization problem

$$C = \min_{t \in [0, \min(I) + \max(I)]} \sum_{i=1}^{n} \left(\max_{j} |a_{ij} - t|\right)^2$$

Determining the constant C is easy and requires only polynomial effort. The optimization problem is convex so the solution is unique. Asking for the loss that a sequence of messages can cause requires taking possible interdependencies of the messages into account. This rules out applying Theorem 3 repeatedly to upper-bound the probability of the joint loss. Instead, one can prove the following assertion to hold for several, possibly interdependent transmissions.

Theorem 4 (Rass (2009)): Let A denote the model matrix of the honest participants Alice and Bob. Assume $I \subset \mathbb{R}^+$ to be compact. If n (possibly interdependent) messages are transmitted over the network, then for any $\varepsilon \geq 0$,

$$\Pr(\min_{1 \leq i \leq n} L_i \geq \rho(A) + \varepsilon) \leq \exp\left(\frac{-n\varepsilon^2}{2(\max(I) - \min(I))^2}\right)$$

if L_i denotes the loss if the adversary mounts an attack on the i-th transmission.

5. Applications

The framework sketched here is general and applies to a wide range of scenarios. Despite its initial purpose being the security assessment of quantum networks, the results apply to any finite two-person game. Future research includes applications to classical networks, as well as considering more general communication scenarios like broadcasting.

5.1 Perfectly secure transmission

A conclusion that can be drawn from Theorem 4 is the possibility of perfectly secure communication over an arbitrarily insecure channel. Assume $I = \{0, 1\}$, so that utility 1 (cf. step 3 of the modelling process) corresponds to secure and secret delivery, and 0 corresponds to successful adversarial extraction of the secret. Consider the event that for n messages, $\min_i L_i \geq 1$. Since the i-th loss $L_i \in I = \{0, 1\}$, this implies $L_i = 1$ for all $i = 1, 2, \ldots, n$. In other words, the upper bound given by Theorem 4 refers to the even that the adversary extracted all messages from the sequence. Letting n become large, this probability will decay exponentially fast, which means that with overwhelming probability, at least one message will remain concealed and secure. If an (n,n)-secret sharing as described in previous paragraphs is employed, then the probability of extracting any secret content is negligible. Notice that none of the results presented relies on a hypothesis about the adversary's

intentions or a mechanism of detecting the success of a transmission. Hence, we can draw strong conclusions from a game-theoretic model that we cannot even fully specify. The formal statement of the above intuition is

Theorem 5 (Rass (2009)): Let the pair (Alice, Bob) set up their model matrix with binary entries $u_{ij} \in \{0, 1\}$, where $u_{ij} = 1$ if and only if a message can securely be delivered by choosing the i-th pure strategy, and the adversary uses his j-th pure strategy for attacking. Then $\rho(A) \in [0, 1]$, and

- for any $\varepsilon > 0$, if $\rho(A) < 1$, then a protocol exists so that Alice and Bob can communicate with an eavesdropping probability of at most ε.
- if $\rho(A) = 1$, then the probability of the message becoming extracted by the adversary is equal to 1.

Rephrasing the implication of Theorem 5 reveals that the possibility of secure communication is completely characterized by the vulnerability. This is a significantly stronger result than the ones presented by Wang & Desmedt (2008), as its assertion is valid for any given network topology. In particular, it opens the possibility of optimizing a given topology, as we will show later. The vulnerability is compatible with the security concept as given by

Definition 5 (Wang & Desmedt (2008)): Let a message transmission protocol be given, and call *adv* the adversary's transcript of an eavesdropped execution. Assume the transmission protocol to take a message m and a random sequence r of coin-flips as input, and denote the adversary's information as $adv(m, r)$. Furthermore, let m_A denote Alice's input, and let m_B denote Bob's final output of the protocol.

- Let $\delta < \frac{1}{2}$. A message transmission protocol is δ-*reliable* if, with probability at least $1 - \delta$, Bob terminates with $m_A = m_B$. The probability is taken over the possible choices of m_A and the coin-flips of all nodes. If $\delta = 0$, then the protocol is said to be *reliable*.
- A message transmission protocol is ε-*private* if, for every two messages m_0, m_1 and every r, $\sum_c |\Pr(adv(m_0, r) = c) - \Pr(adv(m_1, r) = c)| \le 2\varepsilon$ that is, if the distributions of the adversary's views for transmissions of m_0, m_1 differ at most by 2ε in the 1-norm. The probabilities are taken over the coin-flips of the honest parties, and the summation is over all possible values of the adversaries view. A 0-private protocol is called *perfectly private*.
- A message transmission protocol is (ε, δ)-*secure*, if it is δ-reliable and ε-private. A $(0,0)$-secure protocol is called *perfectly secure*.

This definition is perfectly compatible with our understanding of vulnerability, as indicated by the following

Theorem 6 (Rass (2009)): The vulnerability is a measure of privacy and reliability in the sense of Definition 5 because if Alice and Bob set up their model matrix with entries

$$u_{ij} = \begin{cases} 1, & \text{if the message is delivered successfully;} \\ 0, & \text{otherwise} \end{cases}$$

for every strategy combination of the honest parties and the adversary, then the protocol is ρ-reliable, where $\rho = \rho(A)$ is the vulnerability. If Alice and Bob set up their model matrix with entries u_{ij} such that

$$u_{ij} = \begin{cases} 1, & \text{if the adversary learns nothing about the secret content;} \\ 0, & \text{otherwise} \end{cases}$$

for every strategy combination of the honest parties and the adversary, then the protocol is (2ρ)-private, where $\rho = \rho(A)$ is the vulnerability.

5.2 Denial-of-service resilience

Denial of service scenarios are of particular interest in the field of quantum cryptography, because communication is aborted upon detecting the presence of an adversary. Since there is no evasive mechanism that could ensure the function of a link under eavesdropping, a denial of service is most easily mounted in a quantum network. Any small physical damage may raise the error rate sufficiently to logically cut the link. Other examples include classical (electrical) networks, or even distributed attacks on computer networks via bot-nets. All of these can easily be modelled and analyzed with our approach.

Modelling scenarios with random influences (such as intrusion detection systems that have only a high probability of preventing an attack, but offer no provable security assurance) is straightforward by switching from a deterministic utility function to a random one. Basically, this amounts to replacing a random outcome by its expectation. Examples are networks that employ intrusion detection and prevention mechanisms. These mechanisms evade an intruder with a certain probability, but not with certainty, so the possible outcomes $u_{ij} = 1$ (successful transmission), or $u_{ij} = 0$ (transmission failure) occur with probabilities p and $1 - p$, respectively. In that case, one would set up a matrix over the set $I = [0, 1]$ instead of $I = \{0, 1\}$, and replace each random utility U_{ij} with its expected value $\mathrm{E}(U_{ij}) = p \cdot u_{ij}$.


Fig. 2. Network in jeopardy of a DoS-Attack

Example 4: Take the network topology as shown in Fig. 2, and assume the adversary with threshold two, i.e. no more than two links can be attacked simultaneously. Moreover, assume that only the dashed links are assumed vulnerable, which corresponds to the assumption that the neuralgic point is the last mile connecting the peer's machine to the (quantum) backbone. The links inside the network are assumed protected from unauthorized access. This time, we are interested in whether or not Alice (node 1) and Bob (node 4) are in danger of suffering a denial-of-service attack, that is, the adversary's purpose is cutting the channel between the two by exploiting the eavesdropping detection facility of quantum cryptography. Setting up the game with a binary matrix with entries

$$u_{ij} = \begin{cases} 1, & \text{if removing the edges in } s_j^2 \text{ leaves the chosen path } s_i^1 \text{ intact;} \\ 0, & \text{otherwise (Eve has blocked the path)} \end{cases}$$

we end up with a 9×15 matrix. Solving the game yields the value $v = 1/3$, so $\rho = 1 - v = 2/3$. The assertion of Theorem 1 is not limited to pure communication, and we may directly

conclude that the probability of a successful denial-of-service is at most $\rho = \,^2/_3$. Let Alice and Bob retransmit their message in case of failure. Then the probability of mounting a denial-of-service is for, say 100 messages, by Theorem 5 no more than $\exp(-100(^1/_3)^2/2) \approx 0.00386$ (choose the maximal ε, which is $\varepsilon = \,^1/_3$).

5.3 Constructing networks with optimal security

Decision-makers that ought to assess several security enhancements for an existing network may be interested in an objective measure of security. The vulnerability as given in Definition 4 provides a natural scoring functional that can be optimized under given budget constraints.

Sticking with an eavesdropping scenario for simplicity, consider a network whose topology does not permit perfectly secure message transmission. This could be the case if a company owns fibre-optic lines and wants to enter the market as a backbone provider for quantum networks. Such secure delivery services are most interesting in R&D-scenarios, where several spatially separated departments work on highly valued research projects with the need to exchange sensitive data regularly. Different, yet equally important, examples are secure backups, which should be located outside the company's premises (due to fire protection requirements).

In this section, we consider the first of the following two questions, where the second is straightforward to tackle.

- Given a set of environmental and monetary constraints, what is the best security we can achieve under the given conditions?
- Given a minimum security level, what is the cheapest extension to the network that achieves the desired security?

Since the vulnerability as defined up to now refers to only two players, one needs a more general tool: for a graph $G = (V, E)$, and a set of instances $U \subseteq V$, we will consider the maximum vulnerability over each pair of communicating nodes in the set U in the graph G. This quantity is

$$R(U,G) := \max_{s,t \in U, s \neq t} \rho(A(s,t))$$

where $\rho(A(s,t))$ is the vulnerability that is derived from the model matrix A, which now depends on the specific pair (s, t). For simplicity, we restrict ourselves to extensions on the link level of a network. That is, given a graph $G = (V, E)$, and a set E' of links that can technically be realized, we seek a minimum cost extension $\widetilde{E} \subseteq E' \subseteq V \times V$ such that the extended topology $G' = (V, E \cup \widetilde{E})$ has minimum vulnerability. Costs of various types (staff, maintenance, etc) can be captured through a vector-valued function $c: E' \to \mathbb{R}^d$, where $d \geq 1$. The components of c refer to different types of costs that cannot be merged into a single functional. Having specified some constraint $M \in \mathbb{R}^d$ we ought to solve the following nonlinear optimization problem:

$$\begin{aligned} &\min_{\widetilde{E} \subseteq E'} R(U, G(V, E \cup \widetilde{E})) \\ &s.t. \qquad c(\widetilde{E}) \leq M \end{aligned} \tag{1}$$

By reducing this problem to the 0-1-Integer programming problem, as done by Rass (2009), we obtain the following

Theorem 7: The optimization problem (1) is at least NP-hard.

Unfortunately, this result provides some evidence that we can hardly do better than solving the optimization procedure as follows:

1. Enumerate all feasible extensions to the network. That is, find all extensions that obey the cost constraints.
2. Determine the vulnerability of the extended network for each of these cases, and select the extension with the least vulnerability.

This method can be applied to illustrate the optimization through the following

Fig. 3. Initial Network

Example 5 (Rass & Schartner (2009b)): Given a graph G as shown in Fig. 3, suppose that the set E' of candidate extensions comprises the dashed links, i.e. $E' = \{(1, 6), (1, 11), (2, 5), (5, 6), (5, 7), (6, 11), (9, 11)\}$. The weights of the links are the costs for building them. It is easy to verify that the network vulnerability excluding the dashed links is 1, which is due to node 5's inability to communicate with any other node, once an adversary with threshold at least 1 compromises node 4. In our example, assume that each pair within the set of communicating nodes $C = \{1, 2, 5, 11\}$ uses a $(2, 2)$-secret sharing scheme, and that the adversary has threshold 2, i.e. Eve can compromise any two nodes in the network, except for $\{1, 2, 5, 11\}$. Enumerating the paths that a fixed pair can use by $i = 1, 2, \ldots,$ and Eve's possible attacks by $j = 1, 2, \ldots,$ the game-matrix has an entry $u_{ij} = 0$ if and only if Eve's attack is a cut between the i-th pair (alternatively, Eve mounted a person-in-the-middle attack), and 1 otherwise. The cost functional is assumed additive and scalar-valued. Given a budget limit of $M = 13$, we ask for a selection of dashed links that gives us the minimum achievable network jeopardy. Solving the optimization problem, we can find 14 different solutions, each of which satisfies the budget limit and provides a network jeopardy of $1/3$ (in a real-life situation, the set of admissible solutions may further be restricted to limitations on cable length, or other additional constraints). One such solution (the cheapest among the

candidates, with cost 9) is $\widetilde{E} = \{(1,11),(5,6),(5,7),(6,11)\}$, and the resulting, security enhanced network, is depicted in Fig. 4.


Fig. 4. Optimized Network

6. Conclusion

We devised a simple to use framework for analyzing complex scenarios with tools from game theory. Motivated by recent results regarding the impossibility of perfect end-to-end secrecy, we point out decision theory as a valuable tool for obtaining strong general results in the field of system security. Unlike many other approaches, our method is not limited to a specific scenario, and can equally well be applied to tackle confidentiality, authenticity, integrity and availability aspects of a system.

The framework has yet been formulated for the communication of two instances, and generalizations to broadcast scenarios are a direction for future research. The vulnerability measure that we obtained may also be used with time-series forecasting techniques to build an automatic alarming system that keeps track of ongoing evolution and predicts the future security of the given system.

A third avenue of future research is the analysis of the optimization problem. Some steps of turning the problem into a standard mixed integer optimization problem have been accomplished (see Rass (2009) for details), and commercial software packages can be used for tackling the arising problems. The results allow casting the problem of designing a secure network into a combinatorial optimization problem, using a widely automated procedure. Protocol parameters, protocols themselves, transmission paths, and most other parameters can be enumerated automatically. This way, we can automatically create the strategy sets for the honest instances. Network vulnerability scanners help identifying the attack strategies, so these can be set up in an automated manner as well. It is easy to set up the game-matrix, even if random influences are considered. Finally, the analysis, optimization and prediction of future values can also be handed over to software solutions, making the methodology flexible, efficient and a valuable add-on for security analysis in a broad range.

7. References

Ashwin Kumar, M.; Goundan, P. R.; Srinathan, K. & Pandu Rangan, C. (2002), On perfectly secure communication over arbitrary networks, *in 'PODC '02: Proceedings of the twenty-first annual symposium on Principles of distributed computing'*, ACM, New York, NY, USA, pp. 193–202.

Bennett, C. & Brassard, G. (1984), Public key distribution and coin tossing, in '*IEEE International Conference on Computers, Systems, and Signal Processing.*', IEEE Press, Los Alamitos.

BSI (2008), *IT-Grundschutz-Kataloge – 10. Ergänzungslieferung*, Bundesamt für Sicherheit in der Informationstechnik. http://www.bsi.bund.de/gshb/, English version (from 2005) available at http://www.bsi.de/gshb/intl/index.htm.

Chartrand, G. & Zhang, P. (2005), *Introduction to Graph Theory*, Higher education, McGraw-Hill, Boston.

Fitzi, M., Franklin, M. K., Garay, J. & Vardhan, S. H. (2007), Towards optimal and efficient perfectly secure message transmission, in S. Vadhan, ed., '*Theory of Cryptography, 4th Theory of Cryptography Conference, TCC 2007*', Lecture Notes in Computer Science LNCS 4392, Springer, pp. 311–322.

Hammer, V. (1999), *Die 2. Dimension der IT-Sicherheit: Verletzlichkeitsreduzierte Technikgestaltung am Beispiel von Public Key Infrastrukturen*, DuD-Fachbeiträge, Vieweg.

Houmb, S. H. & Franqueira, V. N. L. (2009), Estimating ToE risk level using CVSS, in '*Proceedings of the International Conference on Availability, Reliability and Security*', IEEE Computer Society Press, pp. 718–725.

Innerhofer-Oberperfler, F. & Breu, R. (2009), An empirically derived loss taxonomy based on publicly known security incidents, in '*Proceedings of the International Conference on Availability, Reliability and Security*', IEEE Computer Society Press, pp. 66–73.

Jajodia, S., Noel, S. & O'Berry, B. (2005), *Massive Computing*, Springer US, chapter Topological Analysis of Network Attack Vulnerability, pp. 247–266.

Lo, H.-K. & Chau, H. F. (1999), 'Unconditional security of quantum key distribution over arbitrarily long distances', Science 283, 2050–2056. arXiv:quant-ph/9803006.

Mas (2008), 'Combinatorial analysis utilizing logical dependencies residing on networks (CAULDRON)'. http://ait.gmu.edu/~csis/.

Mavronicolas, M., Papadopoulou, V., Philippou, A. & Spirakis, P. (2005), *Internet and Network Economics*, LNCS, Springer, chapter A Graph-Theoretic Network Security Game, pp. 969–978.

McElice, R. & Sarwate, D. (1981), 'On sharing secrets and Reed-Solomon codes', *Communications of the ACM* 24(9), 583–584.

Rass, S. (2005), *How to send messages over quantum networks in an unconditionally secure manner*, Technical Report TR-syssec-05-05, Universität Klagenfurt, Forschungsgruppe Systemsicherheit.

Rass, S. (2009), *On Information-Theoretic Security: Contemporary Problems and Solutions*, PhD thesis, Klagenfurt University, Institute of Applied Informatics (to appear).

Rass, S. & Schartner, P. (2009a), Game-theoretic security analysis of quantum networks, in '*Proceedings of the Third International Conference on Quantum, Nano and Micro Technologies*', IEEE Computer Society, pp. 20–25.

Rass, S. & Schartner, P. (2009b), Security in quantum networks as an optimization problem, in '*Proceedings of the International Conference on Availability, Reliability and Security*', pp. 493–498.

Schlee, W. (2004), *Einführung in die Spieltheorie*, Vieweg.

Schmitt-Manderbach, T., Weier, H., Fürst, M., Ursin, R., Tiefenbacher, F., Scheidl, T., Perdigues, J., Sodnik, Z., Kurtsiefer, C., Rarity, J. G., Zeilinger, A. & Weinfurter, H. (2007), 'Experimental demonstration of free-space decoy-state quantum key distribution over 144 km', *Physical Review Letters* 98(1), 010504. http://link.aps.org/abstract/PRL/v98/e010504

Shannon, C. (1949), 'Communication theory of secrecy systems', *Bell System Technical Journal* 28, 656–715.

Wang, Y. & Desmedt, Y. (2008), 'Perfectly secure message transmission revisited', *IEEE Transactions on Information Theory* 54(6), 2582–2595.

Wootters, W. K. & Zurek, W. H. (1982), 'A single quantum cannot be cloned', *Nature* 299(802), 802–803.

Ying, Z., Hanping, H. & Wenxuan, G. (2006), 'Network security transmission based on bimatrix game theory', *Wuhan University Journal of Natural Sciences* 11(3), 617–620.

Decision Support Systems

Edited by Chiang S. Jao

Decision support systems (DSS) have evolved over the past four decades from theoretical concepts into real world computerized applications. DSS architecture contains three key components: knowledge base, computerized model, and user interface. DSS simulate cognitive decision-making functions of humans based on artificial intelligence methodologies (including expert systems, data mining, machine learning, connectionism, logistical reasoning, etc.) in order to perform decision support functions. The applications of DSS cover many domains, ranging from aviation monitoring, transportation safety, clinical diagnosis, weather forecast, business management to internet search strategy. By combining knowledge bases with inference rules, DSS are able to provide suggestions to end users to improve decisions and outcomes. This book is written as a textbook so that it can be used in formal courses examining decision support systems. It may be used by both undergraduate and graduate students from diverse computer-related fields. It will also be of value to established professionals as a text for self-study or for reference.

How to reference

In order to correctly reference this scholarly work, feel free to copy and paste the following:

Stefan Rass, Peter Schartner and Raphael Wigoutschnigg (2010). Security as a Game – Decisions from Incomplete Models, Decision Support Systems, Chiang S. Jao (Ed.), ISBN: 978-953-7619-64-0, InTech, Available from: http://www.intechopen.com/books/decision-support-systems/security-as-a-game-decisions-from-incomplete-models

INTECH
open science | open minds