INFORMATION-LEAKAGE IN HYBRID RANDOMIZED PROTOCOLS

Stefan Rass and Peter Schartner
Institute of Applied Informatics, Universität Klagenfurt, Universitätsstrasse 65-67, Klagenfurt, Austria

Keywords: Information-theoretic security, Leakage-resilient cryptography, Protocol design, Game-theory, Multipath transmission.

Abstract: In light of the vast number of existing cryptographic protocols, performance tradeoffs become a major obstacle when selecting one for practical usage. For instance, protocols known to be secure but inefficient compete with others being efficient but offering less security. We tackle such tradeoffs by investigating sequences of random protocol instances (*randomized protocols*) and analyzing the rate at which information leaks from such a sequence. Remarkably, it can be demonstrated that the mutual information between an eavesdropped ciphertext and the plain text decays exponentially fast with the length of the protocol sequence. Using simple tools from game-theory, we devise a generic technique to assemble several protocols of different performance into a single protocol, unifying the advantages of its ingredients. We call this a *hybrid* randomized protocol. We illustrate our technique by using simplified multipath transmission as an example, while observing that our general construction is in no way restricted to this scenario.

1 INTRODUCTION

Imagine an engineer being in charge of designing the cryptographic protection for a given infrastructure. It is most advisable to choose from the vast amount of available cryptographic primitives rather than developing a novel proprietary solution. Existing protocols have the advantage of mostly being scientifically well-studied and justified. However, most of them are hardly comparable, so how should one choose the best among the many? For instance, there is AES with RSA-based key exchange, offering high efficiency but merely conjectured security (resting on computational intractability assumptions or empirical evidence). On the contrary, there is multipath transmission and quantum key distribution, known to offer perfect secrecy at comparably high cost. A technique to get the best of both worlds for a reasonable price is certainly demanding, yet seems missing up to now.

Our contribution in this paper is a general construction that provides an affirmative answer to the following problem:

> Given a protocol Π_1 which is secure but demands much bandwidth, and another protocol Π_2 that is less secure but enjoys good efficiency, can we somehow assemble those two into a protocol that is efficient *and* secure?

Our solution will be a generic framework protocol that takes several private message transmission protocols $\Pi_1,\ldots,\Pi_n$ as input and combines them into a single protocol Π. Provided that at least one of the "input protocols", say Π_1, is perfectly secure, the resulting scheme can be shown to be almost as secure as Π_1 (theorem 4.2) and almost as efficient (in terms of bandwidth demand) as the cheapest among $\Pi_1,\ldots,\Pi_n$ (proposition 4.3). As the construction employs a random ensemble of instances from a class of input protocols, we call this a *hybrid randomized protocol*.

Research in the field of information-theoretic security has brought up a considerable number of cryptographic primitives whose security does not hinge on computational intractability conjectures. Often, this comes at the price of increased efforts for key management, key establishment (e.g. via quantum key distribution) or bandwidth demand (multipath transmission is often based on Shamir's threshold secret sharing; it is known that every perfectly secure one-round message transmission has a minimum relative computational overhead of $n/(n-3t)$ for n bits to be transmitted over t node-disjoint channels (cf. (Fitzi et al., 2007)). Even if bandwidth is not the problem, then finding t node-disjoint channels is yet another significant obstacle in many networks). So one would

Rass S. and Schartner P..
INFORMATION-LEAKAGE IN HYBRID RANDOMIZED PROTOCOLS.
DOI: 10.5220/0003516401340143
In *Proceedings of the International Conference on Security and Cryptography* (SECRYPT-2011), pages 134-143
ISBN: 978-989-8425-71-3

be willing to trade perfect secrecy for efficiency, and our analysis shows how to do it safely.

The key observation made in this paper is that perfect security can be created from arbitrarily weak primitives in the sense that we do not need to be secure in all cases, but only in an arbitrarily small fraction of cases. In other words, once a – perhaps expensive – information-theoretically secure primitive is available, we can easily turn it into a *perfectly secure and cheap* system. The example application in this paper will be *private communication* based on *multipath transmission*.

Related Work: As we are after information-theoretic secrecy, this work is closely related to information-theory based cryptography, as well as conventional cryptography. While previous approaches to perfectly secure communication are mostly concrete ad-hoc solutions, our contribution is a *generic framework* to cook up a secure and efficient protocol from ingredients that do not enjoy both features at the same time. This is the major difference to the related literature, as generic constructions are hardly available, if not absent at all. Generic constructions (to which our results belong), are found rarely and widely scattered throughout the literature (see the work of (Fitzi et al., 2007; Srinathan et al., 2004) and references therein; both of which can be used with our generic construction). Among the few articles concerned with generic security constructions are (Gu et al., 2005), focusing on modeling mostly, (Noman, 2008), assuming a specific field of application, and (Poels et al., 2005), hinging on a particular computational model. Our work is not restricted in any of these ways, besides being about secure communication. In particular, our results can be used to make otherwise costly primitives efficient, e.g. (Alomair and Poovendran, 2009; Hanaoka et al., 2005; Kurosawa and Suzuki, 2007). As far as conventional cryptography is concerned, we will make use of a special block-cipher mode, known as *all-or-nothing transformation (AONT)*. This concept has been introduced in (Rivest, 1997) and further studied in (Stinson, 2001). Notably for us is also the field of leakage-resilient cryptography (Alwen et al., 2009), though we deal with Shannon-information leakage rather than physical attacks (through side-channels for instance). Our work pursues similar goals as leakage-resilient cryptography, yet uses different tools for that. In particular, we will utilize the framework of game-theoretic security analysis, as proposed in (Rass and Schartner, 2010). This approach will come handy and is, for convenience of the reader, sketched in section 3, stressing an example and repeating the basic results required here.

Organization of the paper: we describe our security and adversary model in section 2. In particular, we will be after information-theoretic security, not calling for any infeasibility conjectures. As the employed framework is non-standard and thus not well-known, we will use section 3 to describe how game-theory can be used for security analysis of multipath transmission. In particular, section 3.2 contains a brief introduction to the theory of matrix games, which is needed in the sequel. Our main result is the generic construction given in section 4, along with the security analysis culminating in theorem 4.2.

2 ADVERSARY AND SECURITY MODEL

We assume a computationally unbounded passive threshold adversary. That is, given a graph $G = (V,E)$ modeling the network at hand, at most $t \leq |V \setminus \{s,r\}|$ nodes (excluding the sender s and receiver r) are under control of an attacker. Without loss of generality, we may assume the channels $(u,v) \in E$ to be perfectly protected, for otherwise we could simply insert an intermediate node w as u—w—v, whose compromission models an attack on the u—v-link.

The honest parties transmit a secret message M from a sender $s \in V$ to a non-adjacent receiver $r \in V$, according to some protocol $\Pi^x(M)$, taking random coin-flips x. Let $\text{trans}(\Pi^x(M))$ denote the protocol transcript, i.e. the set of all data that travels over the network when the secret message M is sent using the parameters x. The adversary is assumed to know the protocol Π, the network infrastructure, but *not* the random coin-flips x (taking multipath transmission as our case-study, this assumption is justified since the parameter x is a local and thus secret input at the sender's side). Having a subset of nodes in G under control, the adversary's *view* is $C \subseteq \text{trans}(\Pi^x(M))$

As the adversary is passive, we are not concerned with reliability, and our main goals are privacy and low bandwidth demand for the transmission from s to r. In alignment with Shannon's model of secret communication (Shannon, 1949), the source s is a Markov-process, drawing secret messages from its stationary distribution.

Definition 2.1 (Security)**.** *Let $\Pi^x(M)$ be a protocol using random coin-flips (protocol parameters) x for transmitting a secret message M over a network. Let the adversary's view be $C \subseteq \text{trans}(\Pi^x(M))$.*

We call the protocol Π secure, *if for any given $\varepsilon > 0$, one can choose (protocol parameters) x so that the mutual information is $I(M;C) < \varepsilon$.*

The idea behind randomized protocols is increasing security by repeatedly executing a protocol while enforcing the adversary to break all executions rather than only a single one. This is an application of all-or-nothing-transformations (Rivest, 1997). We will describe the construction for the particular case of a linear transformation, emphasizing that alternative constructions might do equally well. It is important to notice that AONTs are usually a computationally secure primitive, yet in the form we use it below, are *not* resting on any complexity-theoretic intractability assumption.

For a general protocol Π^x with some fixed parameters and random coin-flips all represented by the input $x \in PS$, call PS the *parameter space* of Π.

Definition 2.2 (Randomized Protocol)**.** *Let a protocol Π with parameter space PS be given. A* randomized protocol *is a sequence $\{\Pi^{x_i} | x_i \sim F, i \in \mathbb{N}\}$, where the distribution F is supported on PS, and $x \sim F$ denotes the fact that the random parameter vectors $x_1, x_2, x_3, \ldots$ are drawn from the probability distribution F.*

In section 3.2, we demonstrate how game theory can be used to determine the optimal choice distribution F. In alignment with the terminology from game-theory, we shall also call a particular parameter $x \in PS$ a *strategy* henceforth.

Security of a randomized protocol is ensured by constructing it in a way such that the adversary must extract information from all instances in the sequence $\{\Pi(x_i) | x_i \sim F, i \in \mathbb{N}\}$ to gain insight on the secret information.

Bandwidth in our context captures the entire communication overhead, and as such is not directly related to efficiency in the sense of short key-usage or message-tags (Safavi-Naini et al., 2007).

3 SECURITY FROM GAME-THEORY

For our construction and security analysis, we will partially employ the framework introduced in (Rass and Schartner, 2010). We will thus briefly review the required concepts in this section.

3.1 Secure Multipath Transmission

It is known that multipath transmission (Ashwin Kumar et al., 2002; Fitzi et al., 2007; Wang and Desmedt, 2008) is one way of creating perfectly secure end-to-end communication in the presence of threshold adversaries. We will use such a scheme in a very simplified version to describe the ideas underlying the general construction in section 4. Similar considerations, though not of game-theoretic flavor, are found as well in (Srinathan et al., 2007; Araki, 2008).

Assume a multipath transmission from Alice to Bob: Alice puts her secret message through a secret sharing scheme with threshold k, and transmits each share over its own distinct channel to the receiver, where all these channels are assumed to be mutually node-disjoint. The scheme obviously works, as long as the adversary Eve has less than k paths intercepted[1]. Technically, Alice and the adversary Eve play a *game* against each other, where Alice's move is the selection of paths (wires, channels), and Eve's move is a selection of compromised nodes. In this simplified setting, the protocol Π^x is the multipath transmission scheme, and the parameter x is the particular choices of paths (explicitly including the random coin flips used within the sharing is not required for the example).

The problem is selecting these paths in the optimal manner (provided that this is possible). In terms of definition 2.2, if paths are selected on a random basis to repeatedly transmit messages, then our task is to find the optimal distribution F, according to which these paths are chosen.

If Alice has n_1 choices for the paths, and Eve has $n_2 = \binom{|V|-2}{k}$ choices of subsets to be compromised, it is easy to set up a $(n_1 \times n_2)$-tableau (matrix) A over the set $\{0,1\}$ with

$$a_{ij} = \begin{cases} 0, & \text{if Eve conquered} \geq k \text{ shares} \\ 1, & \text{otherwise,} \end{cases} \quad (1)$$

so that the entry a_{ij} indicates success ($a_{ij} = 1$) or failure ($a_{ij} = 0$) of Alice's attempt to securely deliver her secret over disjoint paths. The decision upon the value of a_{ij} is depending on the threshold of the secret-sharing scheme and the adversary, and the communication overhead and computational effort for reconstruction is depending on the chosen scheme.

Example. Consider the network depicted in figure 1(a), and assume an adversary with threshold 2 and Alice (node s) executing a $(2,2)$-secret sharing scheme to transmit her message to the receiver Bob (node r). Her protocol parameter vector contains a selection of two disjoint paths. Eve succeeds in extracting the message if and only if she has compromised both paths simultaneously. For the network shown in

[1] Various criteria regarding the interplay between network connectivity and the thresholds of the secret-sharing and the adversary are found in (Franklin and Yung, 1995; Franklin and Wright, 2000). We shall not go into details here.

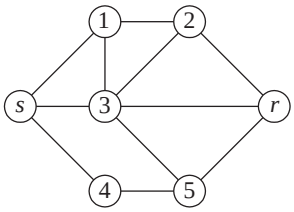

(a) Transmission network

A	s_1^E	s_2^E	s_3^E
s_1^A	0	1	1
s_2^A	1	0	1
s_3^A	1	1	0

(b) Game payoffs

Figure 1: Transmission game example.

figure 1(a), we end up identifying 9 admissible combinations of two paths, and Eve could have conquered one pair from $\binom{5}{2} = 10$ possibilities (the complete list is displayed in table 1). Making the decision for all $9 \times 10 = 90$ scenarios, we end up with a 9×10-matrix (we used a prototype implementation written in Java to run examples).

If Alice chooses a strategy that performs worse than another regardless what Eve does, then this strategy would be a dumb choice and can (should) be ignored. For player i, a strategy s_1 is said to *dominate* another strategy s_0, if the outcome under s_1 is at least as good as the outcome under s_0, for all (strategic) choices of the opponent, with strict inequality in at least one case. Formally, if s_i, s_j are attack strategies with corresponding (payoffs) columns $c_i = (a_{1i}, \ldots, a_{ni})$ and $c_j = (a_{1j}, \ldots, a_{nj})$, then s_i dominates s_j, if and only if $c_i \neq c_j$ and $c_i \leq c_j$ componentwise (recall that player 2's objective is to *minimize* the outcome, as he is the adversary and therefore the "column-player"). Similarly, Eve will not attack in a way that gives her a smaller chance of succeeding than another in every possible scenario. To exemplify this, consider the following two strategies for Alice, referring to our multipath transmission example:

- s_0^A: use paths s—3—5—r and s—1—2—r
- s_1^A: use paths s—3—r and s—1—2—r

giving different outcomes depending on how Eve behaves; table 1 shows the complete list. Strategy s_1^A obviously outperforms strategy s_0^A, as giving as least as good results in every case. Dominated strategies are thus no more profitable than some other strategy, and can thus safely be ignored either. Strategies with all identical payoffs are obviously redundant and can

Table 1: Dominated strategies.

strategy	compromised nodes	payoff	
		s_0^A	s_1^A
s_1^E	1,3	0	0
s_2^E	1,4	1	1
s_3^E	3,4	1	1
s_4^E	1,2	1	1
s_5^E	2,3	0	0
s_6^E	2,4	1	1
s_7^E	1,5	0	1
s_8^E	3,5	1	1
s_9^E	4,5	1	1
s_{10}^E	2,5	0	1

be ignored, thus the inequality condition $c_i \neq c_j$ above is of no harm.

An analogous argument can be applied to the columns. Carrying out this elimination of dominated strategies repeatedly, we end up with a labeled 3×3-matrix, showing the outcomes for each combination of relevant (i.e. undominated) strategies. This matrix is depicted in figure 1(b). The corresponding strategies for Alice in addition to s_1 from above are:

- s_2^A: use paths s—4—5—r and s—1—2—r
- s_3^A: use paths s—4—5—r and s—3—r

We will revisit the concept of dominance in section 4.4. Although the number of strategies can become quite large for bigger networks, the number of undominated strategies often remains small, as various experiments using the Java-based implementation indicated. This drastically reduces the computational efforts tied to a game-theoretic treatment of cryptographic protocols.

3.2 Elements of Matrix-Games

Let us introduce some basics of game-theory, stressing the example for illustrative purposes. A *finite non-cooperative n-person game* is a triple $\Gamma = (N, PS, H)$, where $N = \{1, 2, \ldots, n\}$ is the set of players, $PS = \{PS_1, \ldots, PS_n\}$ is the family of *finite* strategy spaces from which the players can choose their actions, and $H := \{u_i(x,y) | i \in N\}$ is the set of payoffs for each player. These functions have the form $u_i : PS_i \times PS_{-i} \to \mathbb{R}$, where PS_{-i} is the cartesian product $\prod_{j \neq i} PS_j$, i.e. the joint strategy choices of i's opponents. A game is called *zero-sum*, if $\sum_{i \in N} u_i = 0$. In our further treatment, we will exclusively refer to two-person games, with player 1 being the honest pair (s,r), and player 2 being the adversary. Furthermore, our games will be zero-sum, adopting the worst-case assumption that the adversary's sole purpose is de-

feating the intended privacy of the transmission from s to t.

Let us illustrate these concepts in terms of the previous example, i.e. one-round multipath transmission over the network $G = (V,E)$: we have the set of two players $N = \{s, \text{adversary}\}$ with strategy spaces PS_1 being Π's parameter space (ensembles of non-intersecting paths) and $PS_2 = \{U \subseteq V \setminus \{s,r\} \,||U| \leq t\}$ (recall that the adversary has threshold t). The payoff function for the honest party (player 1) is exactly the indicator variable a_{ij} given in equation (1). Hence, for this game (as well as for finite games in general), the payoffs can be specified using matrices (hence the name *matrix games*).

A *Nash-equilibrium* for a (n-person) game Γ is a joint strategy profile $(s_1^*, \ldots, s_n^*) \in PS_1 \times \cdots \times PS_n$ such that

$$u_i(s_i^*, s_{-i}^*) \geq u_i(s_i, s_{-i}^*) \qquad \forall s_i \in PS_i \, \forall i \in N, \quad (2)$$

i.e. no player can gain by unilaterally deviating from the equilibrium (as before, s_{-i} denotes the vector with the strategies of i's opponents). It is easy construct examples where such profiles do not exist, but this can be fixed by switching to *randomized strategies:* assume an infinite number of repetitions of the game, giving a sequence of payoffs for player 1. A *mixed strategy* for the i-th player is a probability distribution over PS_i such that the *expected* payoff (i.e. the long-run average) is maximized. Games as defined above are trivially reformulated in terms of mixed strategies, by replacing the payoff functions in (2) by their expectations w.r.t. the mixed strategy, and replacing $PS_1, \ldots, PS_n$ by the ensemble of probability distributions supported on $PS_1, \ldots, PS_n$. We denote these sets of mixed strategies by $S(PS_1), S(PS_2), \ldots, S(PS_n)$. For a two-player zero-sum matrix-game, as our example gave rise to, the equilibrium condition (2) boils down to

$$u(s_1, s_2^*) \leq u(s_1^*, s_2^*) \leq u(s_1^*, s_2), \quad (3)$$

where $u = u_1 = -u_2$ and s_1, s_2 are discrete probability distributions over PS_1, PS_2. Nash's famous result (Nash, 1951) ensures the existence of such mixed strategies in all finite games. The quantity

$$v(\Gamma) = \max_{s_1 \in S(PS_1)} \min_{s_2 \in S(PS_2)} u(s_1, s_2),$$

i.e. the best that player 1 can do, when player 2 tries to maximize his own payoff (zero-sum regime), is called the *value* of the game Γ. For matrix games, it is easy to calculate this quantity by means of linear optimization (Gibbons, 1992). Obviously, using zero-sum games might not accurately reflect the real behavior of the adversary, yet it is a valid worst-case scenario, since if the adversary has other intentions, he will deviate from the 0-sum-equilibrium and thus causing less damage that expected.

Hence, if Alice follows a zero-sum equilibrium strategy when selecting her protocol parameters, then the protocol will perform optimally in the long run. If the payoff is an indicator variable, this means that Alice can use the zero-sum equilibrium to find the optimal rule for choosing the parameters for her protocol, such that the *probability of failure* is minimized. In other words, F (in definition 2.2) is precisely Alice's part in the Nash-profile. The quantity $\rho(A) = 1 - v(A)$ can be interpreted as *vulnerability*, as it measures the average rate of failure when executing the protocol Π_x repeatedly. For our example, we find $v(A) = 2/3$ and thus $\rho(A) = 1 - v(A) = 1/3$, provided that Alice selects her paths with uniform probability (this is the distribution F in definition 2.2).

Consequently, the average success rate in the zero-sum regime lower-bounds the true success-rate, or, by interpreting this success-rate as a probability, we have

$$\Pr[\text{Eve attacks successfully}] \leq 1 - v(A) = \rho(A), \quad (4)$$

whenever Alice draws her protocol parameters from her equilibrium distribution F (cf. definition 2.2).

4 CONSTRUCTION OF RANDOMIZED PROTOCOLS

Let us sketch the idea underlying our construction, letting the details follow in the next section. Consider Alice transmitting a sequence of plaintexts $M_1, \ldots, M_n$ using n repetitions of a protocol Π with a finite parameter space. Suppose that the protocol permits game-theoretic modeling so that in each repetition $i = 1, 2, \ldots, n$, we can choose the parameters x_i according to an optimal (Nash-equilibrium) rule (i.e. maximizing the average success rate). Using an all-or-nothing transformation on the plain texts $M_1, \ldots, M_n$, we can enforce the adversary to succeed in all n repetitions in order to extract a single message. With an indicator variable for the k-th round, defined as $L_k = 1$ if and only if Π^{x_k} *fails* to protect M_k, we have the following result (cf. (Rass and Schartner, 2010)):

Theorem 4.1. *Fix any $\varepsilon \geq 0$. Let a protocol Π with parameter space PS_1 be given whose game-theoretic payoff matrix is $A \in \{0,1\}^{|PS_1| \times |PS_2|}$ for player 1 (honest team "Alice and Bob"), where PS_2 comprises all attack strategies on Π. Let $v(A)$ be the saddle-point value of A, and put $\rho(A) := 1 - v(A)$. Suppose that Alice runs n repetitions of the protocol Π, with parameters $x_1, \ldots, x_n$ drawn from the zero-sum (Nash-)equilibrium rule for her. For the k-th round, define*

the indicator variable $L_k := 1 - a_{ij}$ with i, j being the strategies of player 1 and 2 in that round, so that $L_k = 1$ if and only if Π^{x_k} fails. Then,

$$\Pr\left[\min_{1\le k\le n} L_k \ge \rho(A) + \varepsilon\right] \le \exp\left(\frac{-n\varepsilon^2}{2}\right).$$

Proof sketch. This result is established by first observing that the sum of n losses is upper-bounded by $n\rho(A)$, using (4), and dividing by n for getting the average loss. The minimum of losses obviously lower-bounds the average, and the sum of losses can be upper-bounded by a generalized version of Hoeffding's inequality, found in (Smith, 2005). The latter inequality is basically identical to Hoeffding's result, except for its weaker stochastic independence assumption, which is trivially satisfied in our setting. □

Our construction involves three steps:

1. Take a protocol Π allowing for a game-theoretic modeling such as outlined in section 3. Make sure that for the corresponding payoff matrix A, we have $\rho(A) = 1 - v(A) < 1$.
2. Apply an all-or-nothing transformation to a sequence of plaintexts to be transmitted, and transmit each resulting block using another instance of Π with randomly chosen parameters.
3. Use theorem 4.1 to assure an exponentially small chance to reveal any transmitted block, thus giving an overwhelming chance of keeping everything secret (thanks to the all-or-nothing transformation).

In the next section, we put this to work, achieving both, bandwidth efficiency and security at the same time. The trick is to repeatedly and randomly switch between protocols, one being secure and the other being efficient.

4.1 A Generic Protocol

In this section, we shall devise a generic construction that allows for assembly of several protocols into a single one such that retains security and improves the performance. Let Π_1, Π_2 denote two protocols, one of which is secure but requires much bandwidth, the other one being less secure but substantially more efficient. Assume that both protocols can be modeled as described in section 3, giving rise to two binary matrices A_1, A_2 and quantities $\rho(A_1) = 1 - v(A_1), \rho(A_2) = 1 - v(A_2)$ corresponding to Π_1, Π_2, respectively. Candidate protocols could, for instance, be multipath transmissions using different thresholds or even different forms of secret sharing (such as polynomial or Chinese remaindering based ones, etc.). Due to the so-constructed protocol arising from two or more other protocols, we call them *hybrid* (our construction will be illustrated with two protocols, leaving the generalization to three or more as an obvious extension). It is important to notice that the sender needs to follow the Nash-profiles induced by the zero-sum games A_1, A_2 when running either protocol.

Let a sequence of stochastically independent messages $M_1, \ldots, M_n \in \mathbb{F}$ be given, where $\mathbb{F}$ is some (fixed) finite field. Without loss of generality, suppose that all M_i are equally long and drawn from the same source (Alice) with entropy $H(M_1) = H(M_2) = \ldots = H(M_n)$.

Initialization: Select a switching probability p, satisfying $0 < p < 1$. Let $Q \in \mathbb{F}^{n\times n}$ be an invertible matrix over $\mathbb{F}$ having no zero entry. The matrix $\mathbb{F}$ does not need to be secret, and can most trivially be constructed by doing a Gauss-Jordan elimination "backwards" (or by constructing Hadamard-matrices if $n-1$ is prime or a power of 2). Applying Q to the vector of messages $(M_1, \ldots, M_n)$, we end up with a vector of "ciphertexts" $C_1, \ldots, C_n$.

Transmission: For $i = 1, 2, \ldots, n$, transmit C_i by randomly running either Π_1 with probability p, or by executing Π_2 with probability $1 - p$. In any case, she selects current protocol parameters according to her Nash-equilibrium strategy obtained from the zero-sum games corresponding to the game-matrices A_1, A_2 for Π_1, Π_2 respectively. The receiver, upon getting all "ciphertexts" $C_1, \ldots, C_n$ can apply Q^{-1} and the combining algorithm to recover all plain text messages.

Figure 2: Randomized protocol Π composed from protocols Π_1, Π_2.

The randomized protocol Π runs as shown in figure 2. The distribution F introduced in definition 2.2 is for both protocols obtained from their respective Nash-equilibria.

Notice that assuming stochastic independence of plaintexts might appear unusual, but is justifiable if the sequence of "messages" constitute a bunch of encryption keys for subsequent encryption of payload.

4.2 Security Analysis

Theorem 4.2 constitutes our main result, regarding the efficiency and security of a randomized protocol.

Theorem 4.2. *Let Π_1, Π_2 be two protocols, with game-matrices A_1, A_2 over $\{0,1\}$ of the form as sketched in section 3, giving the quantities $\rho(A_1), \rho(A_2)$. If $\min\{\rho(A_1), \rho(A_2)\} < 1$, then the protocol Π as shown in figure 2 is secure in the following*

sense: call C the adversary's view on all n repetitions, and let $M = (M_1, \ldots, M_n)$ denote the vector of n plain text messages. Then, the information leakage $I(M;C)$ vanishes in the long run, i.e. $I(M;C) \to 0$ as $n \to \infty$.

Proof. The average rates of failure for both protocols are $\rho(A_1) = 1 - v(A_1)$ and $\rho(A_2) = 1 - v(A_2)$, assuming optimal (Nash-)strategies whenever running Π_1 or Π_2. Without loss of generality, assume $\rho(A_1) < 1$, and (pessimistically), assume $\rho(A_2) = 1$. This would correspond to a case where Π_2 is doomed to failure (in terms of secrecy) with probability 1. By switching between Π_1 and Π_2, we either get the (equilibrium) outcome of Π_1 with probability p or observe a failure (zero payoff) if Π_2 is chosen in the i-th round. The resulting game-matrix A for Π (assuming $\rho(\Pi_2) = 1$) comes to $A = pA_1$. The saddle-point mapping $v(A)$ is a homogeneous functional in A, so consequently, $\rho(A) = 1 - pv(A) < 1$ because $0 < pv(A_1) = v(pA)$. This average rate of failure cannot increase if Π_2 performs better than we assumed, so the actual $\rho(A)$ for the real protocol Π_2 as well satisfies $\rho(A) < 1$ under a random switching regime. To ease notation, we write ρ for the rate of failure of the randomized protocol Π, and note that $\rho < 1$ by our hypothesis.

By construction, the overall information flowing over the network (disregarding any invertible transformation performed by Π_1 or Π_2) is $(C_1, \ldots, C_n)^T = Q \cdot (M_1, \ldots, M_n)^T$. Notice two core facts about the so-constructed C_i's:

1. C_i and C_j are stochastically *dependent* for each i, j
2. If C_{-i} denotes the vector of all C_j's *excluding* the i-th entry, then
$$H(M_j|C_{-i}) = H(M_j)$$
for all $j = 1,2,\ldots,n$, as every M_j is dependent on M_i because Q has no zero entries. Hence, if the adversary's view is C_{-i}, he does not learn anything, thanks to the all-or-nothing transformation via Q.

Suppose $C \subseteq \{C_1, \ldots, C_n\}$ to be the adversary's view on the entire transmission (including all repetitions). For an arbitrary $i \leq n$, consider the information C provides about M_i. We have (by the second of the above observations),

$$H(M_i|C) = \begin{cases} 0, & \text{if } C \text{ suffices to disclose } M_i; \\ H(M_i), & \text{otherwise.} \end{cases}$$

Viewing these two cases within a probability tree permits calculating the entropy as the sum of the root-choice ("attack successful or not?") and the weighted sum of successive entropies (Shannon, 1948):

$$H(M_i|C) = h(p_{\text{attack}}) + (1 - p_{\text{attack}})H(M_i), \qquad (5)$$

where p_{attack} is the likelihood of Eve disclosing M_i upon her knowledge C, and $h(x) := -[x\log(x) + (1-x)\log(1-x)]$ is the binary entropy function. Our task is now upper-bounding p_{attack}. This is done using theorem 4.1.

Introduce an indicator variable L_k for each repetition ($1 \leq k \leq n$) that measure the loss in the k-th round when transmitting C_k:

$$L_k = \begin{cases} 1, & \text{if Eve catches } C_k; \\ 0, & \text{otherwise.} \end{cases}$$

By the all-or-nothing transformation, Eve is successful in extracting M_i from her information C, if and only if $L_1 = L_2 = \ldots = L_n = 1$. If we choose $\varepsilon := 1 - \rho > 0$, then theorem 4.1 tells that

$$p_{\text{attack}} = \Pr\left[\min_{k=1,\ldots,n} L_k \geq \underbrace{\rho + \varepsilon}_{=1}\right] \leq \exp\left(-\frac{1}{2}n\varepsilon^2\right),$$

and hence the probability of catching all messages during an attack is exponentially small.

With this, we can lower-bound (5) as

$$H(M_i|C) \geq \left[1 - \exp\left(-\frac{1}{2}n\varepsilon^2\right)\right] H(M_i),$$

where we have used the fact that $h(p_{\text{attack}}) \geq 0$. Recalling that the mutual information is $I(M_i;C) = H(M_i) - H(M_i|C)$, we end up concluding that within n repetitions of a protocol, the mutual information between the secret M_i and the full transcript C (capturing all repetitions) satisfies

$$I(M_i;C) \leq 2^{-O(n)} \cdot H(M_i), \qquad (6)$$

Now, consider the (joint) mutual information $I(M;C) = I(M_1, \ldots, M_n;C)$ in the following: observe that $\Pr[C|M_i] = \Pr[C]$, since C involves more than just M_i, and can be considered as a "one-time pad encrypted" version of M_i for each i. Conversely, knowing the full transcript C indeed provides information on M_i, so we have the following view on the stochastic dependencies among $M_1, \ldots, M_n$ and C (Bayesian network):

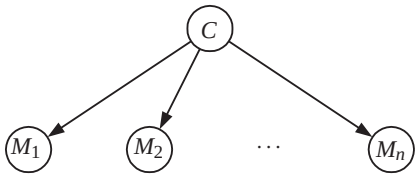

It is easy to verify that the variable particular M_j's are d-separated via C (cf. (Koski and Noble, 2009)), and hence conditionally independent given C, i.e.

$$\Pr[M_i, M_j|C] = \Pr[M_i|C]\Pr[M_j|C]. \qquad (7)$$

This implies that $I(M_1,\ldots,M_n;C) \leq \sum_{i=1}^{n} I(M_i;C)$ (we skip the tedious details here; the derivation is found in the appendix).

Applying (6) to upper-bound the terms on the right-hand side, we get

$$I(M;C) \leq n2^{-O(n)} \sum_{i=1}^{n} H(M_i) = n2^{-O(n)} H(M),$$

because the variables $M_1, M_2, \ldots, M_n$ are stochastically independent. Hence, the information leakage vanishes for long sequences of messages, because $n2^{-O(n)} \to 0$ as $n \to \infty$. □

The number of iterations, and therefore the size of the mixing matrix Q will eventually increase, and one is free to choose p to find a reasonable balance between security and computational cost. Still, the scheme becomes more and more efficient for long sequences of messages (the effort for applying Q and inverting it is polynomial anyway).

We wish to remark some important facts about theorem 4.2:

Remark 1. *The converse of theorem 4.2 is also true: if* $\min\{\rho(A_1),\rho(A_2)\} = 1$*, then* $v(A_1) = v(A_2) = 0$*. This means that an adversary actually playing a Nash-strategy in the zero-sum game will disclose any secret content with probability* 1*. Consequently, the joint randomized protocol* Π *can no longer be secure.*

Remark 2. *The dependencies introduced by the matrix Q make the scheme vulnerable against a denial-of-service, if only one of the* C_i*'s becomes damaged. So, the scheme is secure only against passive adversaries. Gaining resilience against an active adversary is a yet open issue, but detecting such an incident is very simple: any cryptographic checksum attached to any plain text message* M_i *(a single one suffices) will most likely become invalidated when any of the* C_j*'s is modified. This can be detected on the receiver's side, thus making the scheme kind of non-malleable at negligible additional cost.*

4.3 Performance

The bandwidth demand analysis of our generic construction is obvious, and hence its proof is omitted:

Proposition 4.3. *The bandwidth demand of the hybrid randomized protocol* Π*'s can be lowered arbitrarily close towards* $\min\{B(\Pi_1),B(\Pi_2)\}$*, where* $B(\Pi_1),B(\Pi_2)$ *denote the bandwidth demands of* Π_1,Π_2*, respectively.*

Comparing our protocol to others in terms of round- and communication-complexity is difficult, as both features are inherited from the underlying "input-protocols" and are thus subject to the particular implementation at hand. Hence, if the protocol of (Fitzi et al., 2007) is used, one can achieve the theoretical one-round communication overhead in each iteration, performing arbitrarily close to the theoretical minimum overhead. Similarly, one can instantiate our construction with the two-round communication scheme by (Srinathan et al., 2004), which is as well close to the theoretical optimum in terms of overhead for two-round schemes.

4.4 Attack Resilience from Dominance

The condition in theorem 4.2 that at least one of Π_1,Π_2 needs to exhibit a strictly positive chance of success (i.e. $v(A_i) > 0$ for at least one i) deserves a closer look.

Loosely speaking, if we underestimated the adversary's capabilities, then there might be a way of entirely breaking the protocol Π. In the game-theoretic model, this manifests itself as a winning strategy for player 2, making $\rho(A_i) = 1$. We will derive a general result (theorem 4.4) that provides hints on how to fix the protocol (more precisely its game-theoretic payoff matrix) in order to avoid winning strategies for the adversary.

Dominating (i.e. winning) strategies arise if we underestimated the adversary's capabilities, i.e. threshold. Indeed, if the adversary does not have a threshold or the threshold has mistakenly been assumed too low, then one column in A_i appears for the adversary that outperforms all others in the sense that $v(A_i) = 0$ in this case. This new column corresponds to a strategy that entirely breaks the protocol and makes the whole transmission insecure. If we can avoid this situation in *at least one* scenario, then security can be fully restored. Our next result makes this intuition rigorous.

Theorem 4.4. *Take a real number* $M > 0$*. Let* $A \in [0,M]^{n\times m}$ *be a game-matrix with value* $v(A) > 0$*. Extend A by one column c for the adversary, such that c does not dominate all columns in A, and write* $(A|c)$ *for the so-extended matrix. Then* $v((A|c)) > 0$*.*

Proof. Put $B := (A|c)$ with the column $c = (c_1,\ldots,c_n)$, and write e_i for the i-th unit vector. Because c does not dominate every column in B, some $j > 0$ with corresponding column $c = (a_{1j},\ldots,a_{mj})$ exists, having $a_{ij} < c_j$ at some index $i \in \{1,\ldots,n\}$. Let (x,y) be any strategy profile for the game induced by the matrix B. If $y = (0,\ldots,0,1)$, then $v(B) = x^T By = x^T c = c_j > a_{ij} \geq 0$ if $x = e_i$ is played. Now, suppose that $y \neq (0,0,\ldots,1)$ and write $y = (y_1,\ldots,y_m,y_{m+1}) = (y_{1:m},y_{m+1})$. Furthermore, let

(x^*, y^*) be a Nash-equilibrium of the game induced by the matrix A. Assume that player 1 plays x^*, then $(x^*)^T By$ equals

$$(x^*)^T (A|c) \begin{pmatrix} y_{1:m} \\ y_{m+1} \end{pmatrix} = (x^*)^T A y_{1:m} + \overbrace{y_{m+1}(x^*)^T c}^{\geq 0} \geq (x^*)^T A y_{1:m}.$$

From $y_{m+1} < 1$ we deduce $y_{1:m} \neq 0$, and some $\lambda := (\sum_{i=1}^{m} y_i)^{-1} > 0$ exists, rendering $\lambda y_{1:m}$ a probability distribution. From the saddle-point property (3) of the equilibrium (x^*, y^*) in the game with matrix A, we get

$$v(A) = (x^*)^T A y^* \leq (x^*)^T A(\lambda y_{1:m})$$

and subsequently $(x^*)^T A y_{1:m} \geq \frac{1}{\lambda} v(A) > 0$. So, for any adversarial strategy y, player 1 can find some x giving positive revenue, so $v(B) > 0$. □

Rephrasing the result, this means that if there is a winning strategy for the adversary, then only *one* scenario needs to be fixed in order to restore security of the overall protocol.

Stressing the example from section 3 a last time, a winning strategy s_4^E for the adversary would correspond to a new column in A, dominating all others. We denote this matrix by A_1. If there is at least one entry (scenario) such that the adversary looses, then security can be restored. In the example, the matrix A_1 has $v(A_1) = 0$, while the matrix A_2 has $v(A_2) = 1/2$.

A_1	s_1^E	s_2^E	s_3^E	s_4^E
s_1^A	0	1	1	0
s_2^A	1	0	1	0
s_3^A	1	1	0	0

A_2	s_1^E	s_2^E	s_3^E	s_4^E
s_1^A	0	1	1	1
s_2^A	1	0	1	0
s_3^A	1	1	0	0

5 REMARKS AND CONCLUSIONS

It is known that the linear transformation of $M_1, \ldots, M_n$ via the matrix Q used in the proof of theorem 4.2 is insecure if the M_i's are in some well-known redundancy scheme (the argument is the same as for the insecurity of Vernam's one-time pad when a key is used twice). Hence, the framework protocol is best suited for key-exchange and a later cryptographic application (such *as* the one-time pad for instance). It should, however, be noticed that the security proof does not apply to lengthy messages being split into several blocks, due to the assumption of stochastic independence among these blocks.

Therefore, a natural application of our construction is within infrastructures with back-bone networks, claiming to provide perfect security. If we exclude attacks from an insider of an infrastructure, then the outgoing gateway might collect traffic and transmit a whole chunk of messages to another network interface gateway, doing the decryption (according to our framework protocol) and further dispatching of messages within the other secure environment.

An open issue with our construction of the randomized protocol is making it resilient against active adversaries, as remark 2 merely permits detecting a modification, but in no way prevents it. Finally, it is interesting to notice a further conclusion drawn from theorem 4.4 in light of remark 1: if the game-theoretic model comes up with the average success-rate 1, then this proves that there is an adversary, able to break the entire scheme. Hence, if the threshold assumptions is mistaken, theorem 4.4 tells that pretty much the only way of remaining secure against a non-threshold adversary is having an auxiliary security mechanism, which with some positive probability does work against this adversary. In the latter case, however, theorem 4.2 ensures that we can fully recover from such strong attacks.

REFERENCES

Alomair, B. and Poovendran, R. (2009). Information theoretically secure encryption with almost free authentication. *Journal of Universal Computer Science*, 15(15):2937–2956.

Alwen, J., Dodis, Y., and Wichs, D. (2009). Survey: Leakage resilience and the bounded retrieval model. In Kurosawa, K., editor, *Proceedings of the 4th International Conference on Information-Theoretic Security (ICITS)*, pages 1–18. Springer.

Araki, T. (2008). Almost secure 1-round message transmission scheme with polynomial-time message decryption. In Safavi-Naini, R., editor, *Proceedings of the 3rd International Conference on Information-Theoretic Security (ICITS)*, pages 2–13. Springer.

Ashwin Kumar, M., Goundan, P. R., Srinathan, K., and 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*, pages 193–202, New York, NY, USA. ACM.

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

Franklin, M. and Wright, R. (2000). Secure communication in minimal connectivity models. *Journal of Cryptology*, 13(1):9–30.

Franklin, M. and Yung, M. (1995). Secure hypergraphs: privacy from partial broadcast. In *Proceedings of the twenty-seventh annual ACM symposium on Theory of computing*, STOC '95, pages 36–44, New York, NY, USA. ACM.

Gibbons, R. (1992). *A Primer in Game Theory*. Pearson Education Ltd.

Gu, Y., Fu, Y., Li, Y., and Dong, X. (2005). A generic model for symbolic analyzing security protocols. In *The Fifth International Conference on Computer and Information Technology (CIT)*, pages 680–684.

Hanaoka, G., Shikata, J., Hanaoka, Y., and Imai, H. (2005). Unconditionally secure anonymous encryption and group authentication. *The Computer Journal*, 49(3):310–321.

Koski, T. and Noble, J. M. (2009). *Bayesian Networks*. Wiley Series in Probability and Statistics. Wiley.

Kurosawa, K. and Suzuki, K. (2007). Almost secure (1-round, *n*-channel) message transmission scheme. In Desmedt, Y., editor, *Proceedings of the Second International Conference on Information-Theoretic Security (ICTIS)*, pages 99–112. Springer.

Nash, J. F. (1951). Non-cooperative games. *Annals of Mathematics*, 54:286–295.

Noman, A. (2008). A generic framework for defining security environments of wireless sensor networks. In *International Conference on Electrical and Computer Engineering (ICECE)*, pages 924–929.

Poels, K., Tuyls, P., and Schoenmakers, B. (2005). Generic security proof of quantum key exchange using squeezed states. In *Proceedings of the International Symposium on Information Theory (ISIT)*, pages 1612–1616.

Rass, S. and Schartner, P. (2010). A unified framework for the analysis of availability, reliability and security, with applications to quantum networks. *IEEE Transactions on Systems, Man, and Cybernetics – Part C: Applications and Reviews*, 40(5):107–119.

Rivest, R. L. (1997). All-or-nothing encryption and the package transform. In *Proceedings of the 4th International Workshop on Fast Software Encryption*, FSE '97, pages 210–218, London, UK. Springer-Verlag.

Safavi-Naini, R., Wang, S., and Desmedt, Y. (2007). Unconditionally secure ring authentication. In *Proceedings of the 2nd ACM symposium on Information, computer and communications security*, ASIACCS '07, pages 173–181, New York, NY, USA. ACM.

Shannon, C. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27:379–423 and 623–656.

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

Smith, W. D. (2005). Tail bound for sums of bounded random variables. http://www.math.temple.edu/∼wds/homepage/works.html.

Srinathan, K., Narayanan, A., and Pandu Rangan, C. (2004). Optimal perfectly secure message transmission. In *CRYPTO*, pages 545–561.

Srinathan, K., Pandu Rangan, C., and Kumaresan, R. (2007). On exponential lower bound for protocols for reliable communication in networks. In Desmedt, Y., editor, *Proceedings of the Second International Conference on Information-Theoretic Security (ICTIS)*, pages 89–98. Springer.

Stinson, D. R. (2001). Something about all or nothing (transforms). *Designs, Codes and Cryptography*, 22:133–138.

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

APPENDIX

Details to the Proof of Theorem 4.2. Recall the conditional independence of M_i and M_j given C, which manifests itself though formula 7, which was

$$\Pr[M_i, M_j|C] = \Pr[M_i|C]\Pr[M_j|C]\,.$$

For the conditional entropy $H(M_i, M_j|C)$ we find (denoting the respective densities as π, and letting all sums run over all values of m_i, m_j, c)

$$\begin{aligned} H(M_i, M_j|C) &= -\sum \pi(m_i, m_j, c)\log\pi(m_i, m_j|c) \\ &\overset{(7)}{=} \sum \pi(m_i, m_j, c)\log[\pi(m_i|c)\pi(m_j|c)] \\ &= -\sum \pi(m_i, m_j, c)\log\pi(m_i|c) \\ &\quad -\sum \pi(m_i, m_j, c)\log\pi(m_j|c). \end{aligned} \tag{8}$$

Equation (7) implies for the joint distribution

$$\begin{aligned} \Pr[M_i, M_j, C] &= \Pr[M_i, M_j|C]\Pr[C] = \\ &\Pr[M_i|C]\Pr[M_j|C]\Pr[C]\,. \end{aligned}$$

Plugging this into (8) gives

$$\begin{aligned} H(M_i, M_j|C) &= -\sum \underbrace{\pi(m_j|c)}_{\leq 1}\underbrace{\pi(m_i|c)\pi(c)}_{=\pi(m_i,c)}\log\pi(m_i|c) \\ &\quad -\sum \underbrace{\pi(m_i|c)}_{\leq 1}\underbrace{\pi(m_j|c)\pi(c)}_{=\pi(m_j,c)}\log\pi(m_j|c) \\ &\geq H(M_i|C) + H(M_j|C) \end{aligned} \tag{9}$$

From this, we obtain for the conditional information for any fixed choice of i, j

$$\begin{aligned} I((M_i, M_j); C) &= H(M_i, M_j) - H(M_i, M_j|C) \\ &\overset{(9)}{\leq} H(M_i) + H(M_j) - (H(M_i|C) + H(M_j|C)) \\ &= I(M_i; C) + I(M_j; C) \end{aligned}$$

and on iterating, $I(M_1, \ldots, M_n; C) \leq \sum_{i=1}^{n} I(M_i; C)$, recalling that the adversary does not conquer the full C by our construction and theorem 4.1 applied to the randomized protocol in figure 2.