



Compositionality Results for Quantitative Information Flow

Yusuke Kawamoto, Konstantinos Chatzikokolakis, Catuscia Palamidessi

► To cite this version:

Yusuke Kawamoto, Konstantinos Chatzikokolakis, Catuscia Palamidessi. Compositionality Results for Quantitative Information Flow. Proceedings of the 11th International Conference on Quantitative Evaluation of SysTems (QEST 2014), Sep 2013, Florence, Italy. hal-01006381v1

HAL Id: hal-01006381

<https://inria.hal.science/hal-01006381v1>

Submitted on 16 Jun 2014 (v1), last revised 16 Jun 2014 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Compositionality Results for Quantitative Information Flow^{*}

Yusuke Kawamoto^{1,2}, Konstantinos Chatzikokolakis^{2,3}, Catuscia Palamidessi^{1,2}

¹INRIA, France

²École Polytechnique, France

³CNRS, France

Abstract. In the min-entropy approach to quantitative information flow, the leakage is defined in terms of a minimization problem, which, in case of large systems, can be computationally rather heavy. The same happens for the recently proposed generalization called *g*-vulnerability. In this paper we study the case in which the channel associated to the system can be decomposed into simpler channels, which typically happens when the observables consist of several components. Our main contribution is the derivation of bounds on the *g*-leakage of the whole system in terms of the *g*-leakage of its components.

1 Introduction

The problem of preventing confidential information from being leaked is a fundamental concern in the modern society, where the pervasive use of automatized devices makes it hard to predict and control the *information flow*. While early research focussed on trying to achieve *non-interference* (i.e., no leakage), it is nowadays recognized that, in practical situations, some amount of leakage is unavoidable. Therefore an active area of research on information flow is dedicated to the development of theories to *quantify* the amount of leakage, and of methods to minimize it. See, for instance, [15,5,20,18,10,11,25,6].

Among these theories, min-entropy leakage [25,7] has become quite popular, partly due to its clear operational interpretation in terms of one-try attacks. This quite basic setting has been recently extended to the *g*-leakage framework [2]. The main novelty consists in the introduction of gain functions, that permit to quantify the vulnerability of a secret in terms of the gain of the adversary, thus allowing to model a wide variety of operational scenarios.

While *g*-leakage is appealing for its generality and flexible operational interpretation, its computation is not trivial. Like most of the quantitative approaches, its definition is based on the probabilistic correlation between the secrets and the observables. Such correlation is usually expressed in terms of an *information-theoretic channel*, where the secrets constitute the input and the observables the output. The channel is characterized by the *channel matrix*, namely

^{*} This work has been partially supported by the project ANR-12-IS02-001 PACE, by the INRIA Equipe Associée PRINCESS, by the INRIA Large Scale Initiative CAP-PRIS, and by EU grant agreement no. 295261 (MEALS). The work of Y. Kawamoto has been supported by a postdoc grant funded by the IDEX Digital Society project.

Kinds of systems	small systems	large systems	large unknown systems
Input distribution π	known	known	known
Component channels C_i	known	known	approx. statistically
Leakage of C_i with π_i	computable	computable	approx. statistically
Composed channel C	computable	unfeasible	unfeasible
Leakage of C with π	computable	unfeasible	unfeasible

Table 1. Computation of information leakage measures in various scenarios

the conditional probabilities of each output for any given input. The computation of the channel matrix from the system can be performed via model checking (see, e.g., [3]), if a system is completely specified and it is not too complicated. Once the matrix is known, the computation of the g -leakage involves solving an optimization problem. This can be quite costly when the matrix is large.

Worse yet, in many cases it is not possible to compute the channel matrix exactly, for instance because the system may be too complicated, or because the conditional probabilities are partially determined by unknown factors. Fortunately, there are statistical methods that allow to approximate the channel matrix and the leakage [9,12]. There is also a tool, `leakiEst` [14], which allows to estimate min-entropy leakage from a set of trial runs [13]. However, if the cardinality of secrets and observables is large, such estimation becomes computationally heavy, due to the huge amount of trial runs that need to be performed.

In this paper we determined bounds on g -leakages in compositional terms. More precisely, we consider the parallel composition of channels, defined on the cross-products of the inputs and of the outputs. Then, we derive lower and upper bounds on the g -leakage of the whole channel in terms of the g -leakages of the components. Since the size of the whole channel is the product of the sizes of the components, there is an evident benefit in terms of computational cost. Table 1 illustrates the situation for the various kinds of channel matrices (small, large, unknown): the first three rows characterize the situation, and the last three express the feasibility of computing the leakage of the components, the matrix of the whole system, and the leakage of the whole system, respectively. This computation is meant to be exact in the first two columns, and statistical in the last one. The number of components is assumed to be huge. Note that the size of the whole channel increases exponentially with the number of the components.

We evaluate our compositionality results on randomly generated channels and on `Crowds`, a protocol for anonymous communication, run on top of a mobile ad-hoc network (MANET). In such a network users are mobile, can communicate only with nearby nodes, and the network topology changes frequently. As a result, `Crowds` routes can become invalid forcing the user to re-execute the protocol to establish a new route. These protocol repetitions, modeled by the composition of the corresponding channels, lead to more information being leaked. Although the composed channel quickly becomes too big to compute the leakage directly, our compositionality results allow to obtain bounds on it.




Fig. 1. The two kinds of parallel compositions on channels, \times and \parallel .

The rest of the paper is organized as follows: Section 2 introduces basic notions of information theory, defines compositions of channels, and presents information leakage measures. Section 3 presents lower/upper bounds for g -leakages in compositional terms. Section 4 instantiates these results to min-entropy leakages. Section 5 introduces a transformation technique which improves the precision of our method. Section 6 evaluates our results by experiments.

All proofs can be found in the report version [17] of this paper.

2 Preliminaries

In this section we recall the notion of information-theoretic channels, define channel compositions, and recall some information leakage measures.

2.1 Channels

A *discrete channel* is a triple $(\mathcal{X}, \mathcal{Y}, C)$ consisting of a finite set \mathcal{X} of secret input values, a finite set \mathcal{Y} of observable output values, and an $|\mathcal{X}| \times |\mathcal{Y}|$ matrix C , called *channel matrix*, where each element $C[x, y]$ represents the conditional probability $p(y|x)$ of obtaining the output $y \in \mathcal{Y}$ given the input $x \in \mathcal{X}$. The input values have a probability distribution, called *input distribution* or *prior*. Given a prior π on \mathcal{X} , the joint distribution for X and Y is defined by $p(x, y) = \pi[x]C[x, y]$. The output distribution is given by $p(y) = \sum_{x \in \mathcal{X}} \pi[x]C[x, y]$.

2.2 Composition of Channels

We now introduce the two kinds of composition which will be considered in the paper. We assume that the channels are *independent*, in the sense that, given the respective inputs, the outcome of one channel does not influence the outcome of the other. We start with defining *parallel composition with separate inputs* \times (*parallel composition* for short). Note that the term “parallel” here does not carry a temporal meaning: the actual execution of the corresponding systems could take place simultaneously or in any order.

Definition 1 (Parallel composition (with distinct inputs)). Given two discrete channels $(\mathcal{X}_1, \mathcal{Y}_1, C_1)$ and $(\mathcal{X}_2, \mathcal{Y}_2, C_2)$, their *parallel composition (with*

distinct inputs) is the discrete channel $(\mathcal{X}_1 \times \mathcal{X}_2, \mathcal{Y}_1 \times \mathcal{Y}_2, C_1 \times C_2)$ where $C_1 \times C_2$ is the $(|\mathcal{X}_1| \cdot |\mathcal{X}_2|) \times (|\mathcal{Y}_1| \cdot |\mathcal{Y}_2|)$ matrix such that $(C_1 \times C_2)[(x_1, x_2), (y_1, y_2)] = C_1[x_1, y_1] \cdot C_2[x_2, y_2]$ for each $x_1 \in \mathcal{X}_1$, $x_2 \in \mathcal{X}_2$, $y_1 \in \mathcal{Y}_1$ and $y_2 \in \mathcal{Y}_2$.

The condition $(C_1 \times C_2)[(x_1, x_2), (y_1, y_2)] = C_1[x_1, y_1] \cdot C_2[x_2, y_2]$ is what we mean by “the channels are independent”. Note that, although the output distributions Y_1 and Y_2 may be correlated, they are *conditionally independent*, in the sense that $p(y_1, y_2 | x_1, x_2) = p(y_1 | x_1)p(y_2 | x_2)$.

Next, we define the parallel composition with shared input \parallel .

Definition 2 (Parallel composition with shared input). Given two discrete channels $(\mathcal{X}, \mathcal{Y}_1, C_1)$ and $(\mathcal{X}, \mathcal{Y}_2, C_2)$, their *parallel composition with shared input* is the discrete channel $(\mathcal{X}, \mathcal{Y}_1 \times \mathcal{Y}_2, C_1 \parallel C_2)$ where $C_1 \parallel C_2$ is the $|\mathcal{X}| \times (|\mathcal{Y}_1| \cdot |\mathcal{Y}_2|)$ matrix such that $(C_1 \parallel C_2)[x, (y_1, y_2)] = C_1[x, y_1] \cdot C_2[x, y_2]$ for each $x \in \mathcal{X}$, $y_1 \in \mathcal{Y}_1$ and $y_2 \in \mathcal{Y}_2$.

Note that \parallel is a special case of \times . In fact, $(C_1 \parallel C_2)[x, (y_1, y_2)] = C_1[x, y_1] \cdot C_2[x, y_2] = (C_1 \times C_2)[(x, x), (y_1, y_2)]$.

Fig. 1 illustrates these definitions. These two kinds of compositions are used to represent different situations. For example, in the Crowds protocol (explained in Section 6.1) repeated executions of the protocol with different senders are described by the parallel composition (\times), while repeated executions with the same sender are described by the parallel composition with shared input (\parallel).

2.3 Quantitative Information Leakage Measures

The *information leakage* of a channel is measured as the difference between the *prior uncertainty* about the secret value of the channel’s input and the *posterior uncertainty* of the input after observing the channel’s output. The uncertainty is defined in terms of an attacker’s operational scenario. In this paper we will focus on *min-entropy leakage*, in which such measure, min-entropy, represents the difficulty for an attacker to guess the secret inputs in a single attempt.

Definition 3. Given a prior π on \mathcal{X} and a channel $(\mathcal{X}, \mathcal{Y}, C)$, the *prior vulnerability* and the *posterior vulnerability* are defined respectively as

$$V(\pi) = \max_{x \in \mathcal{X}} \pi[x] \quad \text{and} \quad V(\pi, C) = \sum_{y \in \mathcal{Y}} \max_{x \in \mathcal{X}} \pi[x] C[x, y].$$

Definition 4. Given a prior π on \mathcal{X} and a channel $(\mathcal{X}, \mathcal{Y}, C)$, the *min-entropy* $H_\infty(\pi)$ and *conditional min-entropy* $H_\infty(\pi, C)$ are defined by:

$$H_\infty(\pi) = -\log V(\pi) \quad \text{and} \quad H_\infty(\pi, C) = -\log V(\pi, C)$$

and the *min-entropy leakage* $I_\infty(\pi, C)$ and *min-capacity* $C_\infty(C)$ are defined by:

$$I_\infty(\pi, C) = H_\infty(\pi) - H_\infty(\pi, C) \quad \text{and} \quad C_\infty(C) = \sup_{\pi'} I_\infty(\pi', C).$$

Min-entropy leakage has been generalized by *g-leakage* [2], which allows a wide variety of operational scenarios. These are modeled using a set \mathcal{W} of possible *guesses*, and a *gain function* $g : \mathcal{W} \times \mathcal{X} \rightarrow [0, 1]$ such that $g(w, x)$ represents the gain of the attacker when the secret value is x and he makes a guess w on x .

Then *g-vulnerability* is defined as the maximum expected gain of the attacker:

Definition 5. Given a prior π on \mathcal{X} and a channel $(\mathcal{X}, \mathcal{Y}, C)$, the *prior g-vulnerability* and the *posterior g-vulnerability* are defined respectively by

$$V_g(\pi) = \max_{w \in \mathcal{W}} \sum_{x \in \mathcal{X}} \pi[x] g(w, x) \quad \text{and} \quad V_g(\pi, C) = \sum_{y \in \mathcal{Y}} \max_{w \in \mathcal{W}} \sum_{x \in \mathcal{X}} \pi[x] C[x, y] g(w, x).$$

We now extend Definition 4 to the *g*-setting:

Definition 6. Given a prior π on \mathcal{X} and a channel $(\mathcal{X}, \mathcal{Y}, C)$, the *g-entropy* $H_g(\pi)$, *conditional g-entropy* $H_g(\pi, C)$, *g-leakage* $I_g(\pi, C)$ and *g-capacity* $C_g(C)$ are defined by: $H_g(\pi) = -\log V_g(\pi)$, $H_g(\pi, C) = -\log V_g(\pi, C)$, $I_g(\pi, C) = H_g(\pi) - H_g(\pi, C)$, $C_g(C) = \sup_{\pi'} I_g(\pi', C)$.

The min-entropy notions are particular cases of the *g*-entropy ones, obtained by instantiating *g* to the identity function g_{id} defined as $g_{id}(w, x) = 1$ if $w = x$ and $g_{id}(w, x) = 0$ otherwise. Then we have $H_\infty = H_{g_{id}}$, $I_\infty = I_{g_{id}}$ and $C_\infty = C_{g_{id}}$.

3 Compositionality Results on *g*-Leakage

In this section we introduce joint gain functions for composed channels and present compositionality results for *g*-leakage.

3.1 Joint Gain Functions for Composed Channels

To formalize the *g*-leakages of composed channels, we need to know in advance a *joint gain function* g that is defined as a function from $(\mathcal{W}_1 \times \mathcal{W}_2) \times (\mathcal{X}_1 \times \mathcal{X}_2)$ to $[0, 1]$. When a joint secret input is $(x_1, x_2) \in \mathcal{X}_1 \times \mathcal{X}_2$ and the attacker's joint guess is $(w_1, w_2) \in \mathcal{W}_1 \times \mathcal{W}_2$, the attacker's joint gain from the guesses is represented by $g((w_1, w_2), (x_1, x_2))$.

For the sake of generality, we do not assume any relation between g and the two gain functions g_1 and g_2 , except for the following: a joint guess is worthless iff at least one of the single guesses is worthless. Formally: $g((w_1, w_2), (x_1, x_2)) = 0$ iff $g_1(w_1, x_1)g_2(w_2, x_2) = 0$.¹

We say that g_1 and g_2 are *independent* if $g((w_1, w_2), (x_1, x_2)) = g_1(w_1, x_1)g_2(w_2, x_2)$ for all x_1, x_2, w_1 and w_2 .

¹ This property holds, for example, when g, g_1, g_2 are the identity gain functions.

3.2 Jointly Supported Input Distributions

Given a joint prior π on $\mathcal{X}_1 \times \mathcal{X}_2$, the *marginal distribution* π_1 on \mathcal{X}_1 is defined as $\pi_1[x_1] = \sum_{x_2 \in \mathcal{X}_2} \pi[x_1, x_2]$ for all $x_1 \in \mathcal{X}_1$. The *marginal distribution* π_2 on \mathcal{X}_2 is defined analogously. Note that $\pi_1[x_1] \cdot \pi_2[x_2] = 0$ implies $\pi[x_1, x_2] = 0$. The converse does not hold in general, but occasionally we will assume it:

Definition 7. A prior π on $\mathcal{X}_1 \times \mathcal{X}_2$ is *jointly supported* if, for all $x_1 \in \mathcal{X}_1$ and $x_2 \in \mathcal{X}_2$, $\pi_1[x_1] \cdot \pi_2[x_2] \neq 0$ implies $\pi[x_1, x_2] \neq 0$.

Essentially, this condition rules out all the distributions in which there exist two events that happen with a non-zero probability, but that never happen together, i.e., events that are incompatible with each other.

If π_1 and π_2 are independent, i.e., $\pi[x_1, x_2] = \pi_1[x_1] \cdot \pi_2[x_2]$ for all $x_1 \in \mathcal{X}_1$ and $x_2 \in \mathcal{X}_2$, then we denote π by $\pi_1 \times \pi_2$. Note that $\pi_1 \times \pi_2$ is jointly supported.

3.3 The g -Leakage of Parallel Composition

In this section we present a lower and an upper bound for the g -leakage of $C_1 \times C_2$ in terms of the g -leakages of C_1 and C_2 . We first introduce some notation.

Definition 8. Let π be a prior on $\mathcal{X}_1 \times \mathcal{X}_2$, and $g : (\mathcal{W}_1 \times \mathcal{W}_2) \times (\mathcal{X}_1 \times \mathcal{X}_2) \rightarrow [0, 1]$ be a joint gain function. For $w_1 \in \mathcal{W}_1$ and $w_2 \in \mathcal{W}_2$, their *support with respect to* g is defined as: $\mathcal{S}_{w_1, w_2} = \{(x_1, x_2) \in \mathcal{X}_1 \times \mathcal{X}_2 \mid \pi[x_1, x_2] \cdot g((w_1, w_2), (x_1, x_2)) \neq 0\}$.

The lower and the upper bounds are based on the following two measures.

Definition 9. Let g be a joint gain function from $(\mathcal{W}_1 \times \mathcal{W}_2) \times (\mathcal{X}_1 \times \mathcal{X}_2)$ to $[0, 1]$. Let g_1, g_2 be two gain functions from $\mathcal{W}_1 \times \mathcal{X}_1$ to $[0, 1]$ and from $\mathcal{W}_2 \times \mathcal{X}_2$ to $[0, 1]$ respectively. Given a prior π on $\mathcal{X}_1 \times \mathcal{X}_2$, we define M_π^{\min} and M_π^{\max} :

$$M_\pi^{\min} = \min_{w_1 \in \mathcal{W}_1, w_2 \in \mathcal{W}_2} \min_{(x_1, x_2) \in \mathcal{S}_{w_1, w_2}} \frac{\pi_1[x_1] g_1(w_1, x_1) \cdot \pi_2[x_2] g_2(w_2, x_2)}{\pi[x_1, x_2] \cdot g((w_1, w_2), (x_1, x_2))}$$

$$M_\pi^{\max} = \max_{w_1 \in \mathcal{W}_1, w_2 \in \mathcal{W}_2} \sum_{(x_1, x_2) \in \mathcal{S}_{w_1, w_2}} \frac{\pi_1[x_1] g_1(w_1, x_1) \cdot \pi_2[x_2] g_2(w_2, x_2)}{\pi[x_1, x_2] \cdot g((w_1, w_2), (x_1, x_2))}.$$

When π_1 and π_2 are independent and g_1 and g_2 are independent, $M_\pi^{\min} = M_\pi^{\max} = 1$. In addition, for any prior π , M_π^{\min} is strictly positive.

We now show compositionality results for generalized information measures.

Posterior g -Entropy of Parallel Composition

Lemma 1. For any prior π on $\mathcal{X}_1 \times \mathcal{X}_2$ with marginals π_1 and π_2 , and two channels $(\mathcal{X}_1, \mathcal{Y}_1, C_1)$, $(\mathcal{X}_2, \mathcal{Y}_2, C_2)$,

- $H_g(\pi, C_1 \times C_2) \geq H_{g_1}(\pi_1, C_1) + H_{g_2}(\pi_2, C_2) + \log M_\pi^{\min}$
- if π is jointly supported, then $H_g(\pi, C_1 \times C_2) \leq H_{g_1}(\pi_1, C_1) + H_{g_2}(\pi_2, C_2) + \log M_\pi^{\max}$.

The equalities hold if the priors and the gain functions are independent:

Corollary 1. If $g((w_1, w_2), (x_1, x_2)) = g_1(w_1, x_1) g_2(w_2, x_2)$ for all x_1, x_2, w_1 and w_2 , then, for any π_1 and π_2 , $H_g(\pi_1 \times \pi_2, C_1 \times C_2) = H_{g_1}(\pi_1, C_1) + H_{g_2}(\pi_2, C_2)$.

The g -Leakage of Parallel Composition

Theorem 1. *Let π be a jointly supported prior on $\mathcal{X}_1 \times \mathcal{X}_2$ with marginals π_1 and π_2 . Let $(\mathcal{X}_1, \mathcal{Y}_1, C_1)$, $(\mathcal{X}_2, \mathcal{Y}_2, C_2)$ be two channels. Then:*

$$I_{g_1}(\pi_1, C_1) + I_{g_2}(\pi_2, C_2) - \log \frac{M_{\pi}^{\max}}{M_{\pi}^{\min}} \leq I_g(\pi, C_1 \times C_2) \leq I_{g_1}(\pi_1, C_1) + I_{g_2}(\pi_2, C_2) + \log \frac{M_{\pi}^{\max}}{M_{\pi}^{\min}}$$

Again, the equality holds if the priors and the gain functions are independent:

Corollary 2. *If g_1 and g_2 are independent, then $I_g(\pi_1 \times \pi_2, C_1 \times C_2) = I_{g_1}(\pi_1, C_1) + I_{g_2}(\pi_2, C_2)$.*

These results can be naturally extended to the composition of n channels; this extension can be found in the report version of this paper [17].

3.4 The g -Leakage of Parallel Composition with Shared Input

In this section we present compositionality results for g -leakage when two channels share the same input value.

The parallel composition with shared input corresponds to the parallel composition with two identical inputs values: $(C_1 \parallel C_2)[x, (y_1, y_2)] = C_1[x, y_1]C_2[x, y_2] = (C_1 \times C_2)[(x, x), (y_1, y_2)]$. To give the same input value x to both C_1 and C_2 , the prior π^\dagger on $\mathcal{X} \times \mathcal{X}$ is defined from a prior π on \mathcal{X} by:

$$\pi^\dagger[x, x'] = \begin{cases} \pi[x] & \text{if } x = x' \\ 0 & \text{otherwise} \end{cases}$$

Then $H_g(\pi, C_1 \parallel C_2) = H_g(\pi^\dagger, C_1 \times C_2)$. In addition, $\pi_1^\dagger[x] = \pi_2^\dagger[x] = \pi[x]$.

As we see in the definition, the attacker's gain is determined solely from a secret input x and his guess w on x (and independently of channels that receive x as input). Let g be a gain function from $\mathcal{W} \times \mathcal{X}$ to $[0, 1]$. Since C_1 and C_2 receive input from the same domain \mathcal{X} , we use the same gain function g to calculate both the g -leakages of C_1 and C_2 . Since an identical input value x is given to C_1 and C_2 in the composed channel $C_1 \parallel C_2$ and the attacker makes a single guess w on the secret x , we define the joint gain function $g^\dagger: \mathcal{W} \times \mathcal{W} \times \mathcal{X} \times \mathcal{X} \rightarrow [0, 1]$ from g by: $g^\dagger((w, w'), (x, x')) = g(w, x)$ if $w = w'$ and $x = x'$ and $g^\dagger((w, w'), (x, x')) = 0$ otherwise. If $\pi^\dagger[x, x'] \cdot g^\dagger((w, w'), (x, x')) \neq 0$, then $w = w'$ and $x = x'$. Let $(\mathcal{W} \times \mathcal{X})^+ = \{(w, x) \in \mathcal{W} \times \mathcal{X} \mid \pi[x]g(w, x) \neq 0\}$. By $\pi_1^\dagger[x] = \pi_2^\dagger[x] = \pi[x]$, $M^{\min}(\pi^\dagger) = \min_{(w, x) \in (\mathcal{W} \times \mathcal{X})^+} \pi[x]g(w, x)$ and $M^{\max}(\pi^\dagger) = \max_{w \in \mathcal{W}} \sum_{x \in \mathcal{X}} \pi[x]g(w, x)$. Then $H_g(\pi) = -\log M^{\max}(\pi^\dagger)$.

To describe compositionality results, we introduce the following notation.

Definition 10. For any prior π on \mathcal{X} and any gain function g , we define $H_g^{\min}(\pi)$ by: $H_g^{\min}(\pi) = -\log \min \{\pi[x]g(w, x) : x \in \mathcal{X}, w \in \mathcal{W}, \pi[x]g(w, x) \neq 0\}$.

Then, for any prior π , $H_g^{\min}(\pi) = -\log M^{\min}(\pi^\dagger)$ and $H_g^{\min}(\pi) \geq H_g(\pi)$.

Since π^\dagger is *not* jointly supported, we can instantiate compositionality results in the previous sections only on a lower bound for the posterior g -entropy and upper bounds for g -leakage and g -capacity.

The posterior g -entropy $H_g(\pi, C_1 \parallel C_2)$ of a channel composed in parallel with shared inputs is lower-bounded by the summation of $\log H_g^{\min}(\pi)$ and the posterior g -entropies of its two components:

Theorem 2. *For any prior π on \mathcal{X} and channels $(\mathcal{X}, \mathcal{Y}_1, C_1)$ and $(\mathcal{X}, \mathcal{Y}_2, C_2)$,*

$$H_g(\pi, C_1 \parallel C_2) \geq H_g(\pi, C_1) + H_g(\pi, C_2) - H_g^{\min}(\pi).$$

An upper bound of the g -leakage $I_g(\pi, C_1 \parallel C_2)$ of a channel composed in parallel with shared inputs is described using the g -leakages of its two components:

Theorem 3. *For any prior π on \mathcal{X} and channels $(\mathcal{X}, \mathcal{Y}_1, C_1)$ and $(\mathcal{X}, \mathcal{Y}_2, C_2)$,*

$$I_g(\pi, C_1 \parallel C_2) \leq I_g(\pi, C_1) + I_g(\pi, C_2) + H_g^{\min}(\pi) - H_g(\pi).$$

We emphasize this result holds for any prior. Note that in the right-hand side of the above inequality, $H_g^{\min}(\pi) - H_g(\pi)$ is necessary as the following illustrates.

Example 1. Let us consider the channel $(\mathcal{X}, \mathcal{Y}, C)$ where $\mathcal{X} = \mathcal{Y} = \{0, 1\}$ and C is the 2×2 matrix defined by $C[0, 0] = C[1, 1] = 0.9$ and $C[0, 1] = C[1, 0] = 0.1$. Let g be the identity gain function g_{id} and π be the prior on \mathcal{X} such that $\pi[0] = 0.1$ and $\pi[1] = 0.9$. Then $H_g(\pi) = H_g(\pi, C) = -\log 0.9$, $H_g(\pi, C \parallel C) = -\log 0.972$. Therefore $I_g(\pi, C \parallel C) = \log 1.08 > 0 = I_g(\pi, C) + I_g(\pi, C)$.

Note that the inequality of Theorem 3 does not give a useful upper bound when the prior π is far from the uniform distribution. In this example, by $H_g^{\min}(\pi) - H_g(\pi) = \log 9$, the left-hand side is $\log 1.08 \approx 0.111$ while the right-hand side is $\log 9 \approx 3.170$.

These compositionality results are naturally extended to n channels composed in parallel; the extension can be found in the report version of this paper [17]. On the other hand, the result may not hold when the composition of channels is done in a dependent way (i.e., it is not a parallel composition). The following is a counterexample:

Example 2. Let $\mathcal{X} = \mathcal{Y}_1 = \mathcal{Y}_2 = \{0, 1\}$, π be the uniform distribution on \mathcal{X} and g be the identity gain function. We consider the channel that, given an input $x \in \mathcal{X}$, outputs a bit y_1 uniformly drawn from \mathcal{Y}_1 and the exclusive OR y_2 of x and y_1 . Then the g -leakage of the channel is 1 while both of the g -leakages from \mathcal{X} to \mathcal{Y}_1 and from \mathcal{X} to \mathcal{Y}_2 are 0 and $H_g^{\min}(\pi) - H_g(\pi) = 0$. Hence the property expressed by Theorem 3 in general does not hold if we replace \parallel with some other kind of composition.

4 Compositionality Results on Min-Entropy Leakage

In this section we present compositionality results for min-entropy leakage, which yield compositionality theorems for min-capacity.

4.1 Leakage of Parallel Composition

In this section we derive bounds for min-entropy, which, we recall, is a particular case of g -leakage obtained when g is the identity gain function.

We start by remarking that, when the gain functions are identity gain functions, M_π^{\min} and M_π^{\max} reduce to $M_{\infty,\pi}^{\min}$ and $M_{\infty,\pi}^{\max}$ defined as:

$$M_{\infty,\pi}^{\min} = \min_{(x_1,x_2) \in (\mathcal{X}_1 \times \mathcal{X}_2)^+} \frac{\pi_1[x_1] \cdot \pi_2[x_2]}{\pi[x_1,x_2]}, \quad M_{\infty,\pi}^{\max} = \max_{(x_1,x_2) \in (\mathcal{X}_1 \times \mathcal{X}_2)^+} \frac{\pi_1[x_1] \cdot \pi_2[x_2]}{\pi[x_1,x_2]}$$

The next results are consequences of the results of Section 3:

Corollary 3. *For any prior π on $\mathcal{X}_1 \times \mathcal{X}_2$ and channels $(\mathcal{X}_1, \mathcal{Y}_1, C_1)$, $(\mathcal{X}_2, \mathcal{Y}_2, C_2)$,*

- $H_\infty(\pi, C_1 \times C_2) \geq H_\infty(\pi_1, C_1) + H_\infty(\pi_2, C_2) + \log M_{\infty,\pi}^{\min}$.
- *If π is jointly supported, $H_\infty(\pi, C_1 \times C_2) \leq H_\infty(\pi_1, C_1) + H_\infty(\pi_2, C_2) + \log M_{\infty,\pi}^{\max}$.*
- *If $\pi = \pi_1 \times \pi_2$, then $H_\infty(\pi_1 \times \pi_2, C_1 \times C_2) = H_\infty(\pi_1, C_1) + H_\infty(\pi_2, C_2)$.*

Corollary 4. *For a jointly supported prior π on $\mathcal{X}_1 \times \mathcal{X}_2$, channels $(\mathcal{X}_1, \mathcal{Y}_1, C_1)$, $(\mathcal{X}_2, \mathcal{Y}_2, C_2)$ and $F = \log \frac{M_{\infty,\pi}^{\max}}{M_{\infty,\pi}^{\min}}$,*

- $I_\infty(\pi_1, C_1) + I_\infty(\pi_2, C_2) - F \leq I_\infty(\pi, C_1 \times C_2) \leq I_\infty(\pi_1, C_1) + I_\infty(\pi_2, C_2) + F$
- *If $\pi = \pi_1 \times \pi_2$, then $I_\infty(\pi_1 \times \pi_2, C_1 \times C_2) = I_\infty(\pi_1, C_1) + I_\infty(\pi_2, C_2)$.*

The min-entropy leakage coincides with the min-capacity when the prior π is uniform. Thus we re-obtain the following result from the literature [4]: $C_\infty(C_1 \times C_2) = C_\infty(C_1) + C_\infty(C_2)$.

4.2 Leakage of Parallel Composition with Shared Input

As corollaries of Theorems 2 and 3 we obtain the compositionality results for the posterior min-entropy and the min-entropy leakage by taking g as the identity gain function g_{id} . For any prior π on \mathcal{X} , let $H^{\min}(\pi) = -\log \min\{\pi[x] \mid x \in \mathcal{X}, \pi[x] \neq 0\}$. Then $H^{\min}(\pi) \geq \log |\mathcal{X}| \geq H_\infty(\pi)$.

Corollary 5. *For any prior π on \mathcal{X} and channels $(\mathcal{X}, \mathcal{Y}_1, C_1)$ and $(\mathcal{X}, \mathcal{Y}_2, C_2)$,*

- $H_\infty(\pi, C_1) + H_\infty(\pi, C_2) - H^{\min}(\pi) \leq H_\infty(\pi, C_1 \parallel C_2) \leq \min\{H_\infty(\pi, C_1), H_\infty(\pi, C_2)\}$
- $\max\{I_\infty(\pi, C_1), I_\infty(\pi, C_2)\} \leq I_\infty(\pi, C_1 \parallel C_2) \leq I_\infty(\pi, C_1) + I_\infty(\pi, C_2) + H^{\min}(\pi) - H_\infty(\pi)$.

The min-entropy leakage coincides with the min-capacity when the prior π is uniform. If π is uniform we have $H^{\min}(\pi) = H_\infty(\pi)$. Thus we re-obtain the following result from the literature [16]: $C_\infty(C_1 \parallel C_2) \leq C_\infty(C_1) + C_\infty(C_2)$.

The following is an example of the above inequality.

Example 3. Consider the channel $(\mathcal{X}, \mathcal{Y}, C)$ shown in Example 1. Let π be the uniform prior on \mathcal{X} . Then $H_\infty(\pi) = 1$, $H_\infty(\pi, C) = H_\infty(\pi, C \parallel C) \approx 0.152$. Hence $C_\infty(C \parallel C) = H_\infty(\pi) - H_\infty(\pi, C \parallel C) \approx 0.848$ while $C_\infty(C) + C_\infty(C) \approx 1.696$.

5 Improving Leakage Bounds by Input Approximation

The compositionality results for g -leakage shown in a previous section may not give good bounds when the prior is far from the uniform distribution, as illustrated in Example 1. In particular, probabilities that are closer to 0 in priors make our leakage bounds much worse. Since such small probabilities do not affect true g -leakage values much, they can be removed from the priors while this may cause little error on g -leakage values. In the following we present a way of improving bad g -leakage bounds by removing small probabilities. We call it *input approximation* technique. We will only consider the case of min-entropy leakage, i.e., when g is the identity gain function.

The idea of removing small entropies is reminiscent of the notion of *smooth entropy* [8], although the motivation and technicalities are different.

5.1 Bounds for Known Channels

We first consider the case in which the channel components are known. Let π be a prior on \mathcal{X} . Let \mathcal{X}' be a non-empty proper subset of \mathcal{X} such that $\max_{x' \in \mathcal{X}'} \pi[x'] \leq \min_{x \in \mathcal{X} \setminus \mathcal{X}'} \pi[x]$. Then $\max_{x \in \mathcal{X} \setminus \mathcal{X}'} \pi[x] = \max_{x \in \mathcal{X}} \pi[x]$. Let $\epsilon = \sum_{x' \in \mathcal{X}'} \pi[x']$. We define a function $\pi|_{\mathcal{X} \setminus \mathcal{X}'}$ from \mathcal{X} to $[0, 1]$ by:

$$\pi|_{\mathcal{X} \setminus \mathcal{X}'}[x] = \begin{cases} 0 & \text{if } x \in \mathcal{X}' \\ \pi[x] & \text{otherwise} \end{cases}$$

Then $\pi|_{\mathcal{X} \setminus \mathcal{X}'}$ is not a probability distribution, as it does not sum up to 1; i.e., $\sum_{x \in \mathcal{X}} \pi|_{\mathcal{X} \setminus \mathcal{X}'}[x] < 1$. However, the results in previous sections do not require π to be a probability distribution, and neither do the definitions of entropy and leakage. Errors caused by the above input approximation are bounded as follows:

Theorem 4. *For any prior π on \mathcal{X} and channel $(\mathcal{X}, \mathcal{Y}, C)$,*

$$I_\infty(\pi|_{\mathcal{X} \setminus \mathcal{X}'}, C) \leq I_\infty(\pi, C) \leq I_\infty(\pi|_{\mathcal{X} \setminus \mathcal{X}'}, C) + \log(1 + \frac{\epsilon}{V(\pi|_{\mathcal{X} \setminus \mathcal{X}'}, C)}).$$

So, the idea is to remove very small probabilities in priors and then apply our compositional approach to derive bounds illustrated in a previous section. This will allow to obtain better bounds, as small probabilities affect dramatically the precision of our approach, while removing them produces only relatively small errors as shown in Theorem 4.

More precisely, the technique works as follows. Consider a channel C composed of C_1 and C_2 in parallel and a joint prior π on $\mathcal{X}_1 \times \mathcal{X}_2$. We take $\mathcal{X}_1 \times \mathcal{X}_2$ as \mathcal{X} in the input approximation procedure and Theorem 4. Recall that the prior must be jointly supported in order to apply our compositional approach, therefore we take a $\mathcal{X}' \subseteq \mathcal{X}_1 \times \mathcal{X}_2$ so that $\pi|_{\mathcal{X} \setminus \mathcal{X}'}$ is jointly supported. Then we apply Corollary 4 to obtain a lower and an upper bound for $I_\infty(\pi|_{\mathcal{X} \setminus \mathcal{X}'}, C)$. Finally we apply Theorem 4 to obtain bounds for the original $I_\infty(\pi, C)$.

Example 4. Consider the channel $(\mathcal{X}, \mathcal{Y}, C)$ for $\mathcal{X} = \{x_0, x_1, x_2\}$, $\mathcal{Y} = \{y_0, y_1, y_2\}$ and C is given in Fig. 2. We assume the prior π such that $\pi(x_0) = 0.01$, $\pi(x_1) = 0.49$ and $\pi(x_2) = 0.50$, is shared among channels. Then the min-entropy leakage of the channel C^{10} composed of ten C 's in parallel is 0.1319, while our upper bound is 0.7444 when $\epsilon = 0.01$. On the other hand, the upper bound obtained using min-capacity [4] is 4.114, which is much larger than ours.

	y_0	y_1	y_2
x_0	0.50	0.23	0.27
x_1	0.20	0.40	0.40
x_2	0.21	0.43	0.36

Fig. 2. Channel matrix

5.2 Bounds for Channels Composed of Unknown Channels

In some situations an analyst may not know the channel matrices C_1 , C_2 and therefore cannot calculate $I_\infty(\pi|_{\mathcal{X} \setminus \mathcal{X}'}, C_i)$ or $V(\pi|_{\mathcal{X} \setminus \mathcal{X}'}, C_i)$ (necessary to apply Corollary 4), while he may know the information leakages $I_\infty(\pi_1, C_1)$ and $I_\infty(\pi_2, C_2)$. Our input approximation technique allows us to obtain bounds also in this case, although less precise than in the case of known channels. Hereafter we let $\pi' = \pi|_{\mathcal{X} \setminus \mathcal{X}'}$. From Theorem 4:

$$\begin{aligned} \textbf{Theorem 5.} \quad & I_\infty(\pi_1, C_1) + I_\infty(\pi_2, C_2) - \log \frac{M_{\infty, \pi'}^{\max}}{M_{\infty, \pi'}^{\min}} - \log \frac{V(\pi_1, C_1)}{V(\pi_1, C_1) - \epsilon} - \log \frac{V(\pi_2, C_2)}{V(\pi_2, C_2) - \epsilon} \\ & \leq I_\infty(\pi, C_1 \times C_2) \leq I_\infty(\pi_1, C_1) + I_\infty(\pi_2, C_2) + \log \frac{M_{\infty, \pi'}^{\max}}{M_{\infty, \pi'}^{\min}} + \log \frac{\max(V(\pi_1, C_1), V(\pi_2, C_2))}{\max(V(\pi_1, C_1), V(\pi_2, C_2)) - \epsilon}. \end{aligned}$$

$$\begin{aligned} \textbf{Theorem 6.} \quad & I_\infty(\pi, C_1 \| C_2) \leq I_\infty(\pi_1, C_1) + I_\infty(\pi_2, C_2) + \log \frac{\max(V(\pi_1, C_1), V(\pi_2, C_2))}{\max(V(\pi_1, C_1), V(\pi_2, C_2)) - \epsilon} \\ & + H^{\min}(\pi') - H_\infty(\pi'). \end{aligned}$$

When $\epsilon = 0$ these theorems coincide with Corollaries 4 and 5.

Note that $V(\pi_1, C_1)$ and $V(\pi_2, C_2)$ are calculated from $V(\pi_1)$, $V(\pi_2)$, $I_\infty(\pi_1, C_1)$ and $I_\infty(\pi_2, C_2)$. So it is sufficient for an analyst to know only π , $I_\infty(\pi_1, C_1)$ and $I_\infty(\pi_2, C_2)$ to calculate the above leakage bounds.

It is easy to see that these bounds are not as good as those in Section 5.1. Also they are more sensitive to the choice of ϵ . If we take a very small ϵ , the input approximation does not improve substantially, as neither $\frac{M_{\infty, \pi'}^{\max}}{M_{\infty, \pi'}^{\min}}$ nor $H^{\min}(\pi') - H_\infty(\pi')$ decreases much. If we take a very large ϵ , then the error caused by the input approximation is also very large, while $\frac{M_{\infty, \pi'}^{\max}}{M_{\infty, \pi'}^{\min}}$ and $H^{\min}(\pi') - H_\infty(\pi')$ are close to 0. We will later present experiments on the input approximation and illustrate that we should take ϵ as a value less than $\max\{V(\pi_1, C_1), V(\pi_2, C_2)\}$.

The input approximation techniques illustrated in Sections 5.1 and 5.2 can be extended to n -ary channel parallel composition. We refer to [17] for the details.

6 Experimental Evaluation

In this section we evaluate our bounds in two use-cases: first, on the Crowds protocol for anonymous communication, running on a mobile ad-hoc network (MANET), and second, on randomly generated channels.

6.1 Crowds Protocol on a MANET

Crowds [22] is a protocol for anonymous communication, in which participants achieve anonymity by forwarding messages through other users. A group of n users, called the Crowd, participate in the protocol, and one of them, called the *initiator* decides to send a message to some arbitrary recipient in the network, called the *server*. The protocol works as follows: first the initiator selects randomly (with uniform distribution) a member of the crowd, called the *forwarder*, and forwards the message to him. A forwarder, upon receiving a message, throws a (biased) probabilistic coin: with probability p_f (a parameter of the system) he randomly selects a new forwarder and advances the message to him, and with probability $1 - p_f$ he delivers the message directly to the server. Replies from the server follow the inverse path to arrive to the initiator and future requests use the already established route, to avoid repeating the protocol.

The goal of the protocol is to provide sender anonymity w.r.t. an attacker who does not control the whole network, but controls only some of the nodes and can only see traffic passing through them. Still, if the attacker controls some members of the crowd, strong anonymity is not satisfied. A forwarding request from user i is evidence that i is the initiator of the message. However, some anonymity is still provided since user i can always claim that he was in fact only forwarding a message from user j . If the number of corrupted users is relatively small, it is more likely that i is innocent (i.e. the initiator is user $j \neq i$) than guilty, offering a notion of anonymity called *probable innocence* [22].

In this section we consider an instance of Crowds running on a mobile ad-hoc network, in which users are mobile and can communicate only to neighbouring nodes hence the network topology changes frequently. Due to the network changes, routes become invalid and the initiator needs to rerun the protocol to establish a new route, which causes further information leakage. Our goal is to measure how quickly the leakage increases as a function of the number of re-executions. Concerning the attacker model, we assume that the attacker (i) knows the network topology (this could be achieved using known protocols for MANETs, e.g. [21]), (ii) controls some members of the crowd and (iii) controls the server. For a given network topology, the system is modeled by a channel with inputs init_i , meaning that user i is the initiator. The observable events are $\text{forw}_{j,k}$, meaning that user j forwarded the message to the corrupted node k (possibly the destination server). A matrix element $C[\text{init}_i, \text{forw}_{j,k}]$ gives the probability that $\text{forw}_{j,k}$ happens when i is the initiator. Finally, for channels C_1, C_2 modeling the protocol under different network topologies, the repetition of the protocol is modeled as $C_1 \parallel C_2$.

As anonymity metric, we use g -leakage with the 2-tries gain function $g_{\mathcal{W}_2}$, modeling an attacker who can guess the initiator twice. Formally, \mathcal{W}_2 is the set of all subsets of \mathcal{X} with $\#\mathcal{X} = 2$, and $g_{\mathcal{W}_2}(w, x)$ is 1 if $x \in w$ and 0 otherwise.

We evaluate our compositionality results on a Crowds instance with 25 users, of which one is corrupted, and with $p_f = 0.7$. The network topology is generated by randomly adding a connection between any two users with probability 0.4. For a given topology, the matrix is computed by the PRISM model checker [19],

using a model similar to one of [24]. Although executions in Crowds can be infinite, a finite state model can be employed, keeping track of only the current forwarder instead of the full route. Then each element of the channel matrix can be computed by PRISM as the probability of reaching the corresponding state.

The g -leakage of a single execution can be directly computed from the channel; however, for multiple executions, the channel quickly becomes too big to be of practical use (already at 5 repetitions). On the other hand, g -leakage can be bounded using the results in Section 4. The obtained bounds for up to 9 protocol repetitions are shown in Fig. 3. Three variations are given in which the topology changes every 2 executions, every 3 executions or always stays the same. All bounds are computed using a uniform prior and some randomly generated channels. The experiments show that the compositionality technique allows us to obtain meaningful bounds when the system is too big to compute exact values.




Fig. 3. Numbers of observations and bounds

Note that the assumption of uniformly chosen forwarders is standard for the Crowds protocol, however it would be interesting to study how our results would change if we considered non-uniform distributions. For instance, we could have a non-uniform distribution if the possible forwarders were equipped with a notion of trust, like in [23]. We leave this for future work.

6.2 Evaluation on Randomly Generated Channels

In this section we evaluate our bounds on min-entropy leakage using randomly generated channels. In particular, we evaluate the improvement on the bounds due to the input approximation technique, and the efficiency of our approach, which we have implemented as a library in `leakiEst` version 1.3 [1].

We first compare the exact leakage values with their upper bounds calculated using the input approximation technique in the case of shared input. Fig. 4 shows the average upper bounds obtained from Theorem 4, that can be applied when we know the channel matrix. Fig. 5 shows those obtained from Theorem 6 that we can apply when we *do not know* it. For both experiments we use randomly generated 10×10 channel matrices C and a prior π that contains some input with very small probabilities. We set $\epsilon = 0.1$ in the first case and $\epsilon = V(\pi, C)/3$ in the second one. We calculated the min-entropy leakage $I_\infty(\pi, C \parallel C \parallel C \parallel C \parallel C \parallel C)$ (composition of six C 's), and its lower and upper bounds, using the n -ary generalizations of Theorems 4 and 6 (see [17] for the precise formulations.)

These cases give similar upper bounds as shown in Figs. 4 and 5. The x-axis represents *noise levels* of randomly generated matrices, which we define as




Fig. 4. Min-entropy leakages and their bounds for *known* channels




Fig. 5. Min-entropy leakages and their bounds for *unknown* channels




Fig. 6. Min-entropy leakages and their bounds when the analyst does *not* know the channels and chooses ϵ badly




Fig. 7. Upper bounds of min-entropy leakages as a function of the numbers of components

the maximum values (over rows of C) of the summations of the differences of probabilities from the uniform distributions. For instance, when the noise level is 0.10, the average upper bound is 1.699 in the first case (Fig. 4) while it is 1.701 in the second (Fig. 5).

These upper bounds depend on how we choose the parameter ϵ for the input approximation technique. In particular upper bounds strongly depend on ϵ in the case of unknown channels. In Fig. 5 we chose $\epsilon = V(\pi, C)/3$ which gives a relatively good upper bound. On the other hand, if we choose an ϵ too large we may obtain useless bounds. Indeed, if we set for instance $\epsilon = 0.2$, then we obtain upper bounds above the maximum possible leakage, which is the min-entropy, and is always $\log 10 \approx 3.322$ (as shown in Fig. 6) since the input is shared.

Fig. 7 shows average upper bounds of min-entropy leakages of randomly generated 100×100 channels, with randomly generated priors, noise level 0.1, and $\epsilon = 0.005$. As we can see from the figure, the gap between the lower and upper bounds increases with the number of components.

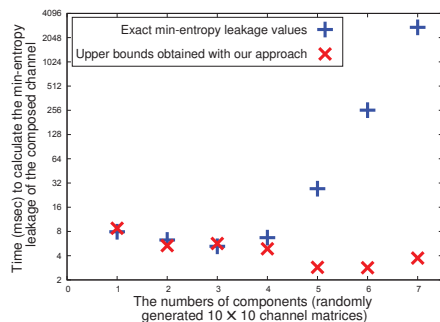


Fig. 8. Average time to calculate min-entropy leakages and their upper bounds




Fig. 9. Average time to calculate upper bounds as a function of the number of components

Finally we evaluate the efficiency of our method. We consider here the min-entropy leakage. Fig. 8 shows the execution time on a laptop (1.8 GHz Intel Core i5) for `leakiEst` to compute the exact min-entropy leakages of the channels composed of randomly generated 10×10 component channels, in comparison with the time to compute their upper bounds. To compute the exact leakages, we used `leakiEst` with an option that calculates the leakages from exact matrices. As we can see, the execution time for the exact values increases rapidly. In fact, the size of composed channel increases exponentially with the number of components, so the complexity of this computation is at least exponential.

For a large number of components, the time to calculate upper bounds increases linearly as shown in Fig. 9. As for the computation of the exact values with `leakiEst`, we expected an exponential blow-up, but we could not check it since we run out of memory because of the size of the matrices.

7 Conclusion and Future Work

We have investigated compositional methods to derive bounds on g -leakage. To improve the precision of the bounds, we have proposed a technique based on the idea of approximating priors by removing small probabilities up to a parameter ϵ . From our experimental results we have found that the dependency of the precision on ϵ is not straightforward. We leave for future work the problem of determining optimal values for ϵ . We also want to explore a possible relation between our technique and the notion of smooth entropies from the information theory literature [8]. This could allow us to develop a more principled approach to the input approximation technique.

References

1. `leakiEst`, <http://www.cs.bham.ac.uk/research/projects/infotools/leakiest/>
2. Alvim, M.S., Chatzikokolakis, K., Palamidessi, C., Smith, G.: Measuring information leakage using generalized gain functions. In: Proc. of CSF. pp. 265–279. IEEE (2012)

3. Andrés, M., Palamidessi, C., van Rossum, P., Smith, G.: Computing the leakage of information-hiding systems. In: Proc. of TACAS. LNCS, vol. 6015, pp. 373–389. Springer (2010)
4. Barthe, G., Köpf, B.: Information-theoretic bounds for differentially private mechanisms. In: Proc. of CSF. pp. 191–204. IEEE (2011)
5. Boreale, M.: Quantifying information leakage in process calculi. In: Proc. of ICALP. LNCS, vol. 4052, pp. 119–131. Springer (2006)
6. Boreale, M., Pampaloni, F., Paolini, M.: Asymptotic information leakage under one-try attacks. In: Proc. of FOSSACS. LNCS, vol. 6604, pp. 396–410. Springer (2011)
7. Braun, C., Chatzikokolakis, K., Palamidessi, C.: Quantitative notions of leakage for one-try attacks. In: Proc. of MFPS. ENTCS, vol. 249, pp. 75–91. Elsevier (2009)
8. Cachin, C.: Smooth entropy and rényi entropy. In: EUROCRYPT. pp. 193–208 (1997)
9. Chatzikokolakis, K., Chothia, T., Guha, A.: Statistical Measurement of Information Leakage. In: Proc. of TACAS. pp. 390–404. LNCS, Springer (2010)
10. Chatzikokolakis, K., Palamidessi, C., Panangaden, P.: Anonymity protocols as noisy channels. *Inf. and Comp.* 206(2–4), 378–401 (2008)
11. Chatzikokolakis, K., Palamidessi, C., Panangaden, P.: On the Bayes risk in information-hiding protocols. *J. of Comp. Security* 16(5), 531–571 (2008)
12. Chothia, T., Kawamoto, Y., Novakovic, C., Parker, D.: Probabilistic point-to-point information leakage. In: Proc. of CSF. pp. 193–205. IEEE (June 2013)
13. Chothia, T., Kawamoto, Y.: Statistical estimation of min-entropy leakage (April 2014), <http://www.cs.bham.ac.uk/research/projects/infotools/>, manuscript
14. Chothia, T., Kawamoto, Y., Novakovic, C.: A tool for estimating information leakage. In: Proc. of CAV 2013 (2013)
15. Clark, D., Hunt, S., Malacaria, P.: Quantitative analysis of the leakage of confidential data. In: Proc. of QAPL. ENTCS, vol. 59 (3), pp. 238–251. Elsevier (2001)
16. Espinoza, B., Smith, G.: Min-entropy as a resource. *Information and Computation* (2013)
17. Kawamoto, Y., Chatzikokolakis, K., Palamidessi, C.: Compositionality Results for Quantitative Information Flow. Tech. rep., INRIA (2014), <http://hal.inria.fr/hal-00999723>
18. Köpf, B., Basin, D.A.: An information-theoretic model for adaptive side-channel attacks. In: Proc. of CCS. pp. 286–296. ACM (2007)
19. Kwiatkowska, M.Z., Norman, G., Parker, D.: PRISM 2.0: A tool for probabilistic model checking. In: Proc. of QEST. pp. 322–323. IEEE (2004)
20. Malacaria, P.: Assessing security threats of looping constructs. In: Proc. of POPL. pp. 225–235. ACM (2007)
21. Nassu, B., Nanya, T., Duarte, E.: Topology discovery in dynamic and decentralized networks with mobile agents and swarm intelligence. In: Proc. of ISDA. pp. 685–690. IEEE (2007)
22. Reiter, M.K., Rubin, A.D.: Crowds: anonymity for Web transactions. *ACM Trans. on Information and System Security* 1(1), 66–92 (1998)
23. Sassone, V., Hamadou, S., Yang, M.: Trust in anonymity networks. In: CONCUR. pp. 48–70 (2010)
24. Shmatikov, V.: Probabilistic analysis of anonymity. In: Proc. of CSFW. pp. 119–128. IEEE (2002)
25. Smith, G.: On the foundations of quantitative information flow. In: Proc. of FOSSACS. LNCS, vol. 5504, pp. 288–302. Springer (2009)