



HAL
open science

Probabilistic Information Flow

Mário S. Alvim, Miguel E. Andrés, Catuscia Palamidessi

► **To cite this version:**

Mário S. Alvim, Miguel E. Andrés, Catuscia Palamidessi. Probabilistic Information Flow. 25th Annual IEEE Symposium on Logic in Computer Science (LICS 2010), Jul 2010, Edinburgh, United Kingdom. pp.314-321, 10.1109/LICS.2010.53 . hal-00548200

HAL Id: hal-00548200

<https://hal.science/hal-00548200>

Submitted on 19 Dec 2010

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.

Probabilistic Information Flow

(Invited Lecture)

Mário S. Alvim
LIX, Ecole Polytechnique
France
msalvim@lix.polytechnique.fr

Miguel E. Andrés
Inst. for Comp. and Inf. Sciences
The Netherlands
mandres@cs.ru.nl

Catuscia Palamidessi
INRIA & LIX, Ecole Polytechnique
France
catuscia@lix.polytechnique.fr

Abstract—In recent years, there has been a growing interest in considering the probabilistic aspects of Information Flow. In this abstract we review some of the main approaches that have been considered to quantify the notion of information leakage, and we focus on some recent developments.

I. INTRODUCTION

One of the concerns in the use of computer systems is to avoid the leakage of secret information through public observables. Ideally we would like systems to be completely secure, but in practice this goal is often impossible to achieve. Therefore it is important to express the amount of leakage in quantitative terms, so to be able to assess whether a system is better than another, although they may both be insecure.

Several works in literature use an Information Theoretic approach to model the problem and define the leakage in a quantitative way, see for example [1], [2], [3], [4], [5], [6], [7]. The idea is that the system is seen as a *channel*. The input represents the secret, the output represents the observable, and the correlation between the input and output (*mutual information*) represents the information leakage. The worst case leakage corresponds then to the *capacity* of the channel, which is by definition the maximum mutual information that can be obtained by varying the input distribution.

In the works mentioned above, the notion of mutual information is based on *Shannon entropy*, which (because of its mathematical properties) is the most established measure of uncertainty. From the security point of view, this measure corresponds to a particular model of attack and a particular way of estimating the security threat (vulnerability of the secret). Other notions have been considered, and argued to be more appropriate for security in certain scenarios. These include: *Rényi min-entropy* [8], [9], *Bayes risk* [10], [11], *guessing entropy* [12], and *marginal guesswork* [13]. In Section II we will discuss their meaning and show how they relate (or do not relate) to each other and to Shannon entropy.

This work has been partially supported by the project ANR-09-BLAN-0169-01 PANDA and by the INRIA DRI Equipe Associée PRINTEMPS.

Whatever definition of uncertainty (i.e. vulnerability) we want to adopt, the notion of leakage is inherent to the system as can be expressed in a uniform way as the difference between the initial uncertainty, i.e. the degree of ignorance about the secret *before* we run the system, and the remaining uncertainty, i.e. the degree of ignorance about the secret *after* we run the system and observe its outcome. Following the principle advocated by Smith [9], and by many others:

$$\text{information leakage} = \frac{\text{initial uncertainty}}{\text{remaining uncertainty}} \quad (1)$$

In (1), the initial uncertainty depends solely on the input distribution, aka *a priori distribution*. Intuitively, the more uniform this is, the less we know about the secret (in the probabilistic sense). After we run the system, if there is a probabilistic correlation between input and output, then the observation of the output should increase our knowledge of the secret. This is determined by the fact that the distribution on the input changes: in fact we can update the probability of each input with the corresponding conditional probability of the same input, given the output. The new distribution is called *a posteriori distribution*. In case input and output are independent, then the *a priori* and the *a posteriori* distributions coincide and the knowledge should remain the same. In the following, we will use the attributes “*a priori*” and “*a posteriori*” to refer to before and after the observation of the output, respectively.

The above intuitions should be reflected by any reasonable notion of uncertainty: it should be higher on more uniform distributions, and it should decrease or remain equal with the observation of related events.

If the uncertainty is expressed in terms of Shannon entropy, then the initial uncertainty is the entropy of the input, the remaining uncertainty is the conditional entropy of the input given the output, and (1) matches exactly the definition of mutual information. This justifies the notion of leakage adopted in the works mentioned before ([1], [2], [3], [4], [5], [6], [7]).

The analogy between information flow in a system and a (simple) channel works well when

- (i) there is no nondeterminism, i.e. either the system is deterministic, or purely probabilistic, and
- (ii) there is a precise temporal relation between secrets and observables in the computations; namely, the value of the secret is chosen at the beginning of the computation, and, the computation of the system produces an observable outcome, with a probability that depends solely on the chosen input and on the system. Furthermore, each new run of the system is independent from the previous ones.

Restriction (i) implies that for each secret there is exactly one conditional probability distribution on the observables, where the condition is the secret value. Restriction (ii) ensures that this conditional distribution depends uniquely on the system (not on the input distribution). These conditional probabilities constitute the so-called *matrix* of the channel. Note that in a (basic) information-theoretic channel the matrix must be invariant with respect to the input distribution, which is exactly what condition (ii) guarantees.

If a system is deterministic, then under the same input each run produces always the same output, with probability 1. Therefore the matrix contains only 0's and 1's. The problem of inferring the secret is still interesting though, because the same output may correspond to different inputs. If the system is probabilistic, i.e. it uses some randomized mechanisms, then the matrix usually contains probabilities different from 0 and 1.

Unfortunately, for real-life systems usually conditions (i) and (ii) are too restrictive:

- Specifications typically need to use nondeterminism in order to abstract from implementation details. This is particularly compelling in the case of concurrent and distributed systems: The order in which the various components get executed, and their interactions, depend on scheduling policies that may differ from implementation to implementation. Furthermore, even if the scheduling policy is fixed, there are run time circumstances that may influence the relative speed of the processes. Nondeterminism is an unavoidable aspect of concurrency.
- Secrets and observables often alternate and interact during an execution. In particular, the choice of a new secret may depend on previous observables. Furthermore, new execution of the systems may depend on previous ones. This may be due to the way the system works, or to the presence of an active adversary that may use the knowledge derived from previous observations to try to tamper with the mechanisms of the system, with the purpose of increasing the leakage. Examples of such systems, that we call here *interactive* systems (where interaction refers to the interplay between secrets and observables), can be found in the area of game theory, auction protocols, web servers, GUI applications, etc.

In this paper, we consider the challenges of extending the information-theoretic approach when conditions (i) and (ii) are lifted, and we illustrate two approaches that we have recently proposed in [14] and [15].

The rest of the paper is organized as follows: in next section we discuss and compare various notions of uncertainty proposed in literature. In Section III we illustrate our proposal for modeling interactive systems and defining the notion of leakage. In Section IV we recall briefly the *possibilistic* approaches proposed in literature to characterize the absence of leakage in nondeterministic systems, we discuss the problems that arise with respect to implementation refinement, and those caused by the presence of omniscient schedulers. We then illustrate our proposal to cope with nondeterminism.

II. UNCERTAINTY AND LEAKAGE

In this section we recall various definitions of uncertainty proposed in literature, and we discuss the relation with security attacks and the way of measuring their success. In general we consider the kind of threats that in the model of [16] are called *brute-force guessing attacks*, which can be summarized as follows: The goal of the adversary is to determine the value of a random variable. He can make a series of queries to an oracle. Each query must have a yes/no answer. In general the adversary is *adaptive*, i.e. he can choose the next query depending on the answer to the previous ones. We assume that the adversary knows the probability distribution.

In the following, A, B denote two discrete random variables with carriers $\mathcal{A} = \{a_1, \dots, a_n\}$, $\mathcal{B} = \{b_1, \dots, b_m\}$, and probability distributions $p_A(\cdot)$, $p_B(\cdot)$, respectively. We will use $A \wedge B$ to represent the random variable with carrier $\mathcal{A} \times \mathcal{B}$ and joint probability distribution $p_{A \wedge B}(a, b) = p_A(a) \cdot p(b | A = a)$, while $A \cdot B$ will denote the random variable with carrier $\mathcal{A} \times \mathcal{B}$ and probability distribution defined as product, i.e. $p_{A \cdot B}(a, b) = p_A(a) \cdot p_B(b)$. Clearly, if A and B are independent, we have $A \wedge B = A \cdot B$. We shall omit the subscripts on the probabilities when they are clear from the context. In reference to a channel, in general A will denote the input (secret), and B the output (observable).

A. Shannon entropy

The Shannon entropy of A [17] is defined as

$$H(A) = - \sum_{a \in \mathcal{A}} p(a) \log p(a)$$

The minimum value $H(A) = 0$ is obtained when $p(\cdot)$ is concentrated on a single value (i.e. when $p(\cdot)$ is a delta of Dirac). The maximum value $H(A) = \log |\mathcal{A}|$ is obtained when $p(\cdot)$ is the uniform distribution. Usually the base of the logarithm is set to be 2 and, correspondingly, the entropy is measured in *bits*.

The *conditional entropy* of A given B is

$$H(A | B) = \sum_{b \in \mathcal{B}} p(b) H(A | B = b) \quad (2)$$

where

$$H(A | B = b) = - \sum_{a \in \mathcal{A}} p(a|b) \log p(a|b)$$

We can prove that $0 \leq H(A | B) \leq H(A)$. The minimum value, 0, is obtained when A is completely determined by B . The maximum value $H(A)$ is obtained when B reveals no information about A , i.e. when A and B are independent.

The *mutual information* between A and B is defined as

$$I(A; B) = H(A) - H(A | B) \quad (3)$$

and it measures the amount of information about A that we gain by observing B . It can be shown that $I(A; B) = I(B; A)$ and $0 \leq I(A; B) \leq H(A)$.

Meaning in security: To explain what $H(A)$ represents from the security point of view, consider a partition $\{\mathcal{A}_i\}_{i \in I}$ of \mathcal{A} . The adversary is allowed to ask questions of the form “is $A \in \mathcal{A}_i$?” according to some strategy. Let $n(a)$ be the number of questions that are needed to determine the value of a , when $A = a$. Then $H(A)$ represents the lower bound to the expected value of $n(\cdot)$, with respect to all possible partitions and strategies of the adversary [13], [16].

B. Rényi min-entropy

In [8], Rényi introduced a one-parameter family of entropy measures, intended as a generalization of Shannon entropy. The Rényi entropy of order α ($\alpha > 0$, $\alpha \neq 1$) of a random variable A is defined as

$$H_\alpha(A) = \frac{1}{1-\alpha} \log \sum_{a \in \mathcal{A}} p(a)^\alpha$$

Rényi’s motivations were of axiomatic nature: Shannon entropy satisfies four axioms, namely symmetry, continuity, value 1 on the Bernoulli uniform distribution, and the chain rule¹:

$$H(A \wedge B) = H(A | B) + H(B) \quad (4)$$

(The entropy of the joint probability, $H(A \wedge B)$, is more commonly denoted by $H(A, B)$. We will use the latter notation in the following.)

Shannon entropy is also the *only* function that satisfies those axioms. However, if we replace (4) with a weaker property representing the additivity of entropy for independent distributions:

$$H(A \cdot B) = H(A) + H(B) \quad (5)$$

then there are more functions satisfying the axioms, among which all those of the Rényi’s family.

¹The original axiom, called the grouping axiom, does not mention the conditional entropy. However it corresponds to the chain rule if the conditional entropy is defined as in (2).

Shannon entropy is obtained by taking the limit of H_α as α approaches 1. In fact we can easily prove, using l’Hôpital’s rule, that

$$H_1(A) \stackrel{\text{def}}{=} \lim_{\alpha \rightarrow 1} H_\alpha(A) = - \sum_{a \in \mathcal{A}} p(a) \log p(a)$$

We are particularly interested in the limit of H_α as α approaches ∞ . This is called *min-entropy*. It can be proven that

$$H_\infty(A) \stackrel{\text{def}}{=} \lim_{\alpha \rightarrow \infty} H_\alpha(A) = - \log \max_{a \in \mathcal{A}} p(a)$$

Rényi considered also the α -generalization of the Kullback-Liebler divergence, which is defined as (assuming that p and q are distributions on the same set \mathcal{X}):

$$D_{KL}(p \| q) = \sum_{x \in \mathcal{X}} p(x) \log \frac{p(x)}{q(x)}$$

Rényi’s α -generalization is:

$$D_\alpha(p \| q) = \frac{1}{1-\alpha} \log \sum_{x \in \mathcal{X}} p(x)^\alpha q(x)^{\alpha-1}$$

The standard case, i.e. the Kullback-Liebler divergence, is again obtained by taking the limit of D_α as $\alpha \rightarrow 1$.

The interest of the above for our purposes lies on the fact that Shannon mutual information can equivalently be defined in terms of the Kullback-Liebler divergence (see for instance [10]):

$$I(A; B) = D_{KL}(A \wedge B \| A \cdot B)$$

Therefore, it seems natural to define the α -generalization of the mutual information as:

$$I_\alpha(A; B) = D_\alpha(A \wedge B \| A \cdot B) \quad (6)$$

Other α -generalizations of the mutual information, based on the same idea, are explored in [18].

As $\alpha \rightarrow \infty$, the above definition gives the following min-version of the mutual information:

$$I_\infty(A; B) \stackrel{\text{def}}{=} \lim_{\alpha \rightarrow \infty} I_\alpha(A; B) = \log \max_{a,b} \frac{p(a,b)}{p(a)p(b)} \quad (7)$$

Another natural way to generalize $I(A; B)$ would be to replace H by H_α in Definition (3). However, Rényi did not define the α -generalization of the conditional entropy, and there is no agreement on what it should be.

Various researchers, including Cachin [19], have considered the following definition, based on (2):

$$H_\alpha^{\text{Cachin}}(A | B) = \sum_{b \in \mathcal{B}} p(b) H_\alpha(A | B = b) \quad (8)$$

which, as $\alpha \rightarrow \infty$, becomes

$$H_\infty^{\text{Cachin}}(A | B) = - \sum_{b \in \mathcal{B}} p(b) \log \max_{a \in \mathcal{A}} p(a | b) \quad (9)$$

An alternative proposal for $H_\infty(\cdot | \cdot)$ came from Smith [9]:

$$H_\infty^{\text{Smith}}(A | B) = -\log \sum_{b \in \mathcal{B}} \max_{a \in \mathcal{A}} p(a, b) \quad (10)$$

Using (9), (10), and the analogue of (3) we can define I_∞^{Cachin} and I_∞^{Smith} .

Meaning in security: Rényi min-entropy can be related to a model of adversary who is allowed to ask exactly one question, which must be of the form “is $A = a$?” (one-try attacks). More precisely, $H_\infty(A)$ represents the (logarithm of the inverse of the) probability of success for this kind of attacks and with the best strategy, which consists, of course, in choosing the a with the maximum probability.

As for $H_\infty(A | B)$ and $I_\infty(A; B)$, the most interesting versions, in terms of security, seem to be those of Smith: In fact, $H_\infty^{\text{Smith}}(A | B)$ represents the inverse of the (expected value of the) probability that the same kind of adversary succeeds in guessing the value of A *a posteriori*, i.e. after observing the result of B . The complement of this probability is also known as *probability of error* or *Bayes risk*. Since in general B and A are correlated, observing B increases the probability of success. In fact we can prove formally that $H_\infty^{\text{Smith}}(A | B) \leq H_\infty^{\text{Smith}}(A)$, with equality if and only if A and B are independent. $I_\infty^{\text{Smith}}(A; B)$ corresponds to the *ratio* between the probabilities of success a priori and a posteriori, which is a natural notion of leakage. ($I_\infty^{\text{Smith}}(A; B)$ is in the format of (1), but the difference becomes ratio due to the presence of the logarithms.) Note that $I_\infty^{\text{Smith}}(A; B) \geq 0$, which seems desirable for a good notion of leakage.

The definition of I_∞ in (7) has also an interpretation in security: it represents the maximum gain in the probability of success, i.e. the maximum ratio between the a posteriori and the a priori probability. Note that also $I_\infty(A; B)$ is always non-negative and it is 0 if and only if A and B are independent. Furthermore $I_\infty(A; B)$ coincides with I_∞^{Smith} if B is uniformly distributed. More in general, $D_{KL}(p \| q)$ and its α -extension $D_\alpha(p \| q)$ should represent the “inefficiency” of an adversary who bases its strategy on the distribution q , when in fact the real distribution is p . Hence $I_\alpha(A; B)$ defined as $D_\alpha(A \wedge B \| A \cdot B)$ should represent the gain of the adversary in revising his strategy according to the knowledge of the correlation between A and B .

Concerning H_α^{Cachin} and I_α^{Cachin} , they have some nice properties. For instance they enjoy weak versions of the chain rule (4). More precisely, the “=” in (4) becomes “ \geq ” for $\alpha < 1$, and “ \leq ” for $\alpha > 1$. However, there is no general relation between $H_\infty^{\text{Cachin}}(A | B)$ and $H_\infty(A)$, and therefore I_∞^{Cachin} is not guaranteed to be non-negative.

C. Guessing entropy

The notion of guessing entropy was introduced by Massey in [12]. Let us assume, for simplicity, that the elements of \mathcal{A} are ordered by decreasing probabilities, i.e. if $1 \leq i < j \leq n$

then $p(a_i) \geq p(a_j)$. Then the guessing entropy is defined as follows:

$$H_G(A) = \sum_{1 \leq i \leq |\mathcal{A}|} i p(a_i)$$

Massey did not define the notion of conditional guessing entropy. In some works, like [19], [16], it is defined analogously to (2):

$$H_G(A | B) = \sum_{b \in \mathcal{B}} p(b) H_G(A | B = b)$$

Meaning in security: Guessing entropy represents an adversary who is allowed to ask repeatedly questions of the form “is $A = a$?”. More precisely, $H_G(A)$ represents the expected number of questions that the adversary needs to ask to determine the value of A , assuming that he follows the best strategy, which consists, of course, in choosing the a 's in order of decreasing probability.

$H_G(A | B)$ represents the expected number of questions *a posteriori*, i.e. after observing the value of B and reordering the queries according to the updated probabilities (i.e. the queries will be chosen in order of decreasing a posteriori probabilities).

Also in this case, $H_G(A | B)$ is not necessarily smaller than or equal to $H_G(A)$, so the corresponding notion of mutual information is not guaranteed to be non-negative².

D. Marginal guesswork

The marginal guesswork is a variant of guessing entropy that was proposed by Pliam [13]. It is parametric to a number $\varepsilon > 0$, and is defined as follows. Again, we assume that the elements of \mathcal{A} are ordered by decreasing probabilities.

$$H_\varepsilon(A) = \min\{j \mid \sum_{1 \leq i \leq j} p(a_i) > \varepsilon\}$$

Pliam did not define the conditional version of marginal guesswork, but in [16] it is defined following (2):

$$H_\varepsilon(A | B) = \sum_{b \in \mathcal{B}} p(b) H_\varepsilon(A | B = b)$$

Meaning in security: Consider again an adversary who is allowed to ask repeatedly questions of the form “is $A = a$?”. $H_\varepsilon(A)$ represents the minimum number of questions that the adversary needs to ask to determine the value of A with probability at least ε .

$H_\varepsilon(A | B)$ represents the same notion, but using the a posteriori probabilities. Again, it is not necessarily the case that $H_\varepsilon(A | B) \leq H_\varepsilon(A)$.

²This problem is inherent to the probabilistic case, and therefore it does not occur in [16], since that work considers only deterministic systems.

E. Comparison and Discussion

The various notions of entropy discussed in this section have been carefully compared with Shannon entropy, to conclude that in general there is no tight relation. Fano's inequality gives a lower bound to the Bayes risk in terms of (conditional) Shannon entropy, and Rényi [20], Hellman-Raviv [21], and Santhi-Vardi [22] give upper bounds as well, but all these are rather weak. Smith has shown in [9] that the orderings induced on channels by the Bayes risk and by Shannon entropy are in general unrelated.

Massey has shown that the exponential of the Shannon entropy is a lower bound for the guessing entropy, and that, in case of a geometric distribution, the bound is tight. However Massey has also shown that in the general case the Shannon entropy can be arbitrarily close to 0 while the guessing entropy is constant [12].

As for the marginal guesswork. Pliam has shown that it is essentially unrelated with Shannon entropy [13].

We conclude this section with an observation about the principle (1). As we have seen above this principle prescribes that, given a model of attack (and a measure of success), one should find the corresponding notion of entropy and conditional entropy, which will then be considered as the initial and the residual uncertainty, respectively. The tendency in literature is to define the conditional entropy following the formula (2).

It is important to realize that the notion of probability enters the definition of entropy in two ways: one corresponds to the use of them made by the adversary to decide its strategy. The other is for averaging purposes. While the distribution used in the first way depends on the knowledge of the adversary, and changes from a priori to a posteriori with the revelation of the observable, the distribution used in the second way should always be the real one, i.e. the a posteriori one. Some of the definitions of entropy given previously do not satisfy this rule, and in fact we obtain the counterintuitive consequence that the a posteriori uncertainty may be higher than the a priori one.

An alternative of the principle (1) would be to define the leakage as information flow directly by using a suitable variant of the Kullback-Liebler divergence, like in (6). As discussed at the end of Section II-B, this divergence represents the "inefficiency" of an adversary who bases its strategy on the distribution q , when in fact the real distribution is p . So, the new principle would be:

$$\begin{aligned} \text{Leakage} &= \text{effectiveness of the attack using the} \\ &\quad \text{a posteriori distribution} \\ &\quad - \\ &\quad \text{effectiveness of the attack using the} \\ &\quad \text{a priori distribution} \end{aligned}$$

where the computations are done taking into account that the real distribution is the a posteriori one. In the case of

Shannon entropy this new principle coincides with (1), but in general they are different.

For instance, in the case of the guessing entropy, we should define $I_G(A; B)$ as the expected value (averaged using the a posteriori probabilities) of the divergence between the number of queries when the a 's are ordered using the a priori distribution on A , and the number of queries when the a 's are ordered using the a posteriori distribution on A .

$$I_G(A; B) = \sum_{b \in \mathcal{B}} p(b) \sum_{1 \leq i \leq |\mathcal{A}|} i (p(a_i | b) - p(a_{k_i} | b)) \quad (11)$$

where k_i is a permutation which reorders the elements of \mathcal{A} so that their a posteriori probabilities are decreasing, namely if $1 \leq i \leq j \leq n$, then $p(a_{k_i}) \geq p(a_{k_j})$.

Note that, if we applied (3) with the definitions of guessing entropy given above, we would obtain, instead:

$$I_G(A; B) = \sum_{b \in \mathcal{B}} p(b) \sum_{1 \leq i \leq |\mathcal{A}|} i (p(a_i) - p(a_{k_i} | b))$$

It is possible to prove that, if defined as in (11), then $I_G(A; B)$ is always non-negative.

III. INTERACTIVE INFORMATION FLOW

In this section we consider the applicability of the information-theoretic approach to interactive systems, i.e. those systems in which there can be an alternation of secrets and observables during the computation, and they influence each other. The conditional probabilities $p(b | a)$ can be computed as the ratio between the probability that a computation has trace (a, b) , given that it has secret trace a [23]. This is natural and correct, as it follows the definition of conditional probability in terms of joint and marginal probability. However, as shown by the example below, it does not help to define an information-theoretic channel because by definition a channel should be invariant with respect to the input distribution, and such construction is not.

Example: Consider the protocol represented in Figure 1. This protocol is used in a website with one seller and two possible buyers (one of them is *poor* and the other *rich*). The sale starts with the seller offering a product, which can be either *cheap* (with probability r) or *expensive* (with probability $1 - r$). In case the offered product is cheap, the poor buyer acquires the product with probability s whereas the rich buyer does it with probability $1 - s$. Similarly (with probabilities t and $1 - t$) for the case on which the offered product is expensive. We assume that the offered product is observable, since it is visible to everyone in the website, while the identity of the buyer is secret. In the following, we use the notation \bar{r} to represent $1 - r$.

If we want to build the channel matrix, one could think of using the standard formula for conditional probability $p(b | a) = \frac{p(a,b)}{p(a)}$ to fill each entry of the matrix. That is what is




Figure 1. Interactive System

	cheap	expensive
poor	$\frac{rs}{rs+\bar{r}t}$	$\frac{\bar{r}t}{rs+\bar{r}t}$
rich	$\frac{r\bar{s}}{r\bar{s}+\bar{r}t}$	$\frac{\bar{r}t}{r\bar{s}+\bar{r}t}$

Table I
CHANNEL MATRIX

indeed proposed in [23]. Proceeding this way we obtain the matrix on Table I.

However, the entries of the matrix are not invariant with respect to the input distribution. If we fix the parameters $r = \bar{r} = 0.5$ and make two different assignments to the values of s, \bar{s}, t, \bar{t} we induce two different input distributions, with the associated matrices, as shown in Table II.

(a) $r = \frac{1}{2}, s = \frac{2}{5}, t = \frac{3}{5}$

	cheap	expensive	Marginal $p_A(\cdot)$
poor	$\frac{2}{5}$	$\frac{3}{5}$	$\frac{1}{4}$
rich	$\frac{8}{15}$	$\frac{7}{15}$	$\frac{3}{4}$

(b) $r = \frac{1}{2}, s = \frac{1}{10}, t = \frac{3}{10}$

	cheap	expensive	Marginal $p_A(\cdot)$
poor	$\frac{1}{4}$	$\frac{3}{4}$	$\frac{1}{2}$
rich	$\frac{9}{16}$	$\frac{7}{16}$	$\frac{1}{2}$

Table II
TWO DIFFERENT CHANNEL MATRICES DEPENDING ON THE INPUT DISTRIBUTION

As shown by this example, when secrets occur *after* observables the conditional probabilities depend on the distribution on secrets and, thus, so it does the matrix making it unsound to analyze such systems using information-theoretical approaches.

In [14], we investigate an extension of the theory of simple channels so to make the information-theoretic approach applicable also the case of interactive systems. It turns out that a richer notion of channels, known in Information Theory as *channels with memory and feedback*, serves our purposes. Indeed the dependence of inputs on past outputs corresponds exactly to feedback, and the dependence of the output on all previous inputs and outputs corresponds to memory.

A. Applications

Interactive systems can be found in a variety of disparate areas such as game theory, auction protocols, and zero-knowledge proofs. We now present two examples of interactive systems.

- In the area of auction protocols, consider the cocaine auction protocol [24]. The auction is organized as a succession of rounds of bidding. Round i starts with the seller announcing the bid price b_i for that round. Buyers have t seconds to make an offer (i.e. to say *yes*, meaning “I am willing to buy at the current bid price b_i ”). As soon as one buyer says *yes*, he becomes the winner w_i of that round and a new round begins. If nobody says anything for t seconds, round i is concluded by timeout and the auction is won by the winner w_{i-1} of the previous round.

The identities of the buyers in each round constitute the input of the channel, whereas the bid prices constitute the output of the channel. Note that inputs and outputs alternate so the system is interactive. It is also easy to see that inputs depend on past outputs (feedback): the identity of the winner of each round depends on the previous bid prices. Furthermore, outputs depend on the previous inputs (memory): (in some scenarios) the bid price of round i may depend on the identity of previous winners. For more details on the modeling of this protocol using channels with memory and feedback see [14].

- In the area of game theory, consider the classic prisoner’s dilemma (the present formulation is due to Albert W. Tucker [25], but it was originally devised by Merrill Flood and Melvin Dresher in 1950). Two suspects are arrested by the police. The police have insufficient evidence for a conviction, and, having separated both prisoners, visit each of them to offer the same deal. If one testifies (defects from the other) for the prosecution against the other and the other remains silent (cooperates with the other), the betrayer goes free and the silent accomplice receives the full 10-year sentence. If both remain silent, both prisoners are sentenced to only six months in jail for a minor charge. If each betrays the other, each receives a five-year sentence. Each prisoner must choose to betray the other or to remain silent. Each one is assured that the other would not know about the betrayal before the end of the investigation.

In the iterated prisoner’s dilemma, the game is played repeatedly. Thus each player has an opportunity to punish the other player for previous non-cooperative play. In this case the strategy (cooperate or defect) of each player is the input of the channel and the sentence is the output. Once again, it is easy to see that the system is interactive: inputs and outputs alternate.

Furthermore, inputs depend on previous outputs (the strategy depend on the past sentences) and outputs depend on previous inputs (the sentence of the suspects depend on their declarations - cooperate or defect).

IV. NONDETERMINISM AND INFORMATION FLOW

As seen in previous Section II, the *noise* of the channel, namely the similarity between the rows of the channel matrix, helps preventing the inference of the secret from the observables. In practice noise is created by using randomization, see for instance the DCNet [26] and the Crowds [27] protocols.

In the literature about the foundations of Computer Security, however, the quantitative aspects are often abstracted away, and probabilistic behavior is replaced by nondeterministic behavior. Correspondingly, there have been various approaches in which information-hiding properties are expressed in terms of equivalences based on nondeterminism, especially in a concurrent setting. For instance, [28] defines *anonymity* as follows³: A protocol S is anonymous if, for every pair of culprits a and b , $S^{[a/x]}$ and $S^{[b/x]}$ produce the same observable traces. A similar definition is given in [29] for *secrecy*, with the difference that $S^{[a/x]}$ and $S^{[b/x]}$ are required to be bisimilar. In [30], an electoral system S preserves the *confidentiality of the vote* if for any voters v and w , the observable behavior of S is the same if we swap the votes of v and w . Namely, $S^{[a/v | b/w]} \sim S^{[b/v | a/w]}$, where \sim represents bisimilarity.

These proposals are based on the implicit assumption that *all the nondeterministic executions present in the specification of S will always be possible under every implementation of S* . Or at least, that the adversary will believe so. In concurrency, however, as argued in [31], nondeterminism has a rather different meaning: if a specification S contains some nondeterministic alternatives, typically it is because we want to abstract from specific implementations, such as the scheduling policy. A specification is considered correct, with respect to some property, if every alternative satisfies the property. Correspondingly, an implementation is considered correct if all executions are among those possible in the specification, i.e. if the implementation is a refinement of the specification. There is no expectation that the implementation will actually make possible all the alternatives indicated by the specification.

We argue that the use of nondeterminism in concurrency corresponds to a *demonic* view: the scheduler, i.e. the entity that will decide which alternative to select, may try to choose the worst alternative. Hence we need to make sure that “all alternatives are good”, i.e. satisfy the intended property. In the above mentioned approaches to the formalization of security properties, on the contrary, the interpretation

³The actual definition of [28] is more complicated, but the spirit is the same.

of nondeterminism is *angelic*: the scheduler is expected to actually help the protocol to confuse the adversary and thus protect the secret information.

There is another issue, orthogonal to the angelic/demonic dichotomy, but relevant for the achievement of security properties: the scheduler *should not be able to make its choices dependent on the secret*, or else nearly every protocol would be insecure, i.e. the scheduler would always be able to leak the secret to an external observer (for instance by producing different interleavings of the observables, depending on the secret). This remark has been made several times already, and several approaches have been proposed to cope with the problem of the “almighty” scheduler (aka omniscient, clairvoyant, etc.), see for example [32], [33], [31], [34].

The risk of a naive use of nondeterminism to specify a security property, is not only that it may rely on an implicit assumption that the scheduler behaves angelically, but also that it is clairvoyant, i.e. that it peeks at the secrets (that it is not supposed to be able to see) to achieve its angelic strategy.

Consider the following system, in a CCS-like syntax:

$$\begin{aligned} S &\stackrel{\text{def}}{=} (c, \text{out})(A \parallel \text{Corr} \parallel H_1 \parallel H_2), \\ A &\stackrel{\text{def}}{=} \bar{c}\langle \text{sec} \rangle, \quad \text{Corr} \stackrel{\text{def}}{=} c(s).\overline{\text{out}}\langle s \rangle \\ H_1 &\stackrel{\text{def}}{=} c(s).\overline{\text{out}}\langle a \rangle, \quad H_2 \stackrel{\text{def}}{=} c(s).\overline{\text{out}}\langle b \rangle \end{aligned}$$

where \parallel is the parallel operator, $\bar{c}\langle \text{sec} \rangle$ is a process that sends *sec* on channel c , $c(s).P$ is a process that receives s on channel c and then continues as P , and (c, out) is the restriction operator, enforcing synchronization on c and out . In this example, *sec* represents a secret information.

It is easy to see that we have $S^{[a/\text{sec}]} \sim S^{[b/\text{sec}]}$. Note that, in order to simulate the third branch in $S^{[a/\text{sec}]}$, the process $S^{[b/\text{sec}]}$ needs to select its first branch. Viceversa, in order to simulate the third branch in $S^{[b/\text{sec}]}$, the process $S^{[a/\text{sec}]}$ needs to select its second branch. This means that, in order to achieve bisimulation, the scheduler needs to know the secret, and change its choice accordingly.

This example shows a system that intuitively is not secure, because the third component, *Corr*, reveals whatever secret it receives. However, according to the equivalence-based notions of security discussed above, *it is secure*. But it is secure thanks to a scheduler that angelically helps the system to protect the secret, and it does so by making its choices dependent on the secret! In our opinion these assumptions on the scheduler are excessively strong.

In a recent work [15] we address the above issue by defining a framework in which it is possible to combine both angelic and demonic nondeterminism in a setting in which also probabilistic behavior may be present, and in a context in which the scheduler is restricted (i.e. not clairvoyant). We propose safe versions of typical equivalence relations (traces and bisimulation), and we show how to use them to characterize information-hiding properties.

REFERENCES

- [1] Y. Zhu and R. Bettati, "Anonymity vs. information leakage in anonymity systems," in *Proc. of ICDCS*. IEEE, 2005, pp. 514–524.
- [2] D. Clark, S. Hunt, and P. Malacaria, "Quantitative information flow, relations and polymorphic types," *J. of Logic and Computation*, vol. 18, no. 2, pp. 181–199, 2005.
- [3] P. Malacaria, "Assessing security threats of looping constructs," in *Proc. of POPL*. ACM, 2007, pp. 225–235.
- [4] P. Malacaria and H. Chen, "Lagrange multipliers and maximum information leakage in different observational models," in *Proc. of PLAS*. ACM, 2008, pp. 135–146.
- [5] I. S. Moskowitz, R. E. Newman, and P. F. Syverson, "Quasi-anonymous channels," in *Proc. of CNIS*. IASTED, 2003, pp. 126–131.
- [6] I. S. Moskowitz, R. E. Newman, D. P. Crepeau, and A. R. Miller, "Covert channels and anonymizing networks," in *Proc. of PES*. ACM, 2003, pp. 79–88.
- [7] K. Chatzikokolakis, C. Palamidessi, and P. Panangaden, "Anonymity protocols as noisy channels," *Inf. and Comp.*, vol. 206, no. 2–4, pp. 378–401, 2008.
- [8] A. Rényi, "On Measures of Entropy and Information," in *Proc. of the 4th Berkeley Symposium on Mathematics, Statistics, and Probability*, 1961, pp. 547–561.
- [9] G. Smith, "On the foundations of quantitative information flow," in *Proc. of FOSSACS*, ser. LNCS, vol. 5504. Springer, 2009, pp. 288–302.
- [10] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. John Wiley & Sons, Inc., 1991.
- [11] K. Chatzikokolakis, C. Palamidessi, and P. Panangaden, "On the Bayes risk in information-hiding protocols," *Journal of Computer Security*, vol. 16, no. 5, pp. 531–571, 2008.
- [12] Massey, "Guessing and entropy," in *Proc. of ISIT*. IEEE, 1994, p. 204.
- [13] Pliam, "On the incomparability of entropy and marginal guesswork in brute-force attacks," in *Proc. of INDOCRYPT*, ser. LNCS, no. 1977. Springer-Verlag, 2000, pp. 67–79.
- [14] M. S. Alvim, M. E. Andrés, and C. Palamidessi, "Information Flow in Interactive Systems," Tech. Rep., 2010, submitted for publication. [Online]. Available: <http://hal.archives-ouvertes.fr/inria-00479672/en/>
- [15] M. S. Alvim, M. E. Andrés, C. Palamidessi, and P. van Rossum, "Safe Equivalences for Security Properties," in *Proc. of IFIP TCS*, vol. To appear, 2010, submitted for publication. [Online]. Available: <http://hal.archives-ouvertes.fr/inria-00479674/en/>
- [16] B. Köpf and D. A. Basin, "An information-theoretic model for adaptive side-channel attacks," in *Proc. of CCS*. ACM, 2007, pp. 286–296.
- [17] C. E. Shannon, "A mathematical theory of communication," *Bell System Technical Journal*, vol. 27, pp. 379–423, 625–56, 1948.
- [18] I. Csiszár, "Generalized cutoff rates and Rényi's information measures," *Transactions on Information Theory*, vol. 41, no. 1, pp. 26–34, 1995.
- [19] C. Cachin, "Entropy measures and unconditional security in cryptography," Ph.D. dissertation, 1997.
- [20] A. Rényi, "On Measures of Entropy and Information," in *Proc. of the 4th Berkeley Symposium on Mathematics, Statistics, and Probability*, 1960, pp. 547–561.
- [21] M. Hellman and J. Raviv, "Probability of error, equivocation, and the Chernoff bound," *IEEE Trans. on Information Theory*, vol. IT-16, pp. 368–372, 2007.
- [22] N. Santhi and A. Vardy, "On an improvement over Rényi's equivocation bound," 2006, presented at the 44-th Annual Allerton Conf. on Communication, Control, and Computing, September 2006. Available at <http://arxiv.org/abs/cs/0608087>.
- [23] J. Desharnais, R. Jagadeesan, V. Gupta, and P. Panangaden, "The metric analogue of weak bisimulation for probabilistic processes," in *Proc. of LICS*. IEEE, 2002, pp. 413–422.
- [24] F. Stajano and R. J. Anderson, "The cocaine auction protocol: On the power of anonymous broadcast," in *Information Hiding*, 1999, pp. 434–447.
- [25] W. Poundstone, *Prisoners Dilemma*. Doubleday NY, 1992.
- [26] D. Chaum, "The dining cryptographers problem: Unconditional sender and recipient untraceability," *Journal of Cryptology*, vol. 1, pp. 65–75, 1988.
- [27] M. K. Reiter and A. D. Rubin, "Crowds: anonymity for Web transactions," *ACM Transactions on Information and System Security*, vol. 1, no. 1, pp. 66–92, 1998.
- [28] S. Schneider and A. Sidiropoulos, "CSP and anonymity," in *Proc. of ESORICS*, ser. LNCS, vol. 1146. Springer, 1996, pp. 198–218.
- [29] M. Abadi and A. D. Gordon, "A calculus for cryptographic protocols: The spi calculus," *Inf. and Comp.*, vol. 148, no. 1, pp. 1–70, 1999.
- [30] S. Delaune, S. Kremer, and M. Ryan, "Verifying privacy-type properties of electronic voting protocols," *Journal of Computer Security*, vol. 17, no. 4, pp. 435–487, 2009.
- [31] K. Chatzikokolakis, G. Norman, and D. Parker, "Bisimulation for demonic schedulers," in *Proc. of FOSSACS*, ser. LNCS, vol. 5504. Springer, 2009, pp. 318–332.
- [32] R. Canetti, L. Cheung, D. Kaynar, M. Liskov, N. Lynch, O. Pereira, and R. Segala, "Task-structured probabilistic i/o automata," in *Proc. of WODES*, 2006.
- [33] K. Chatzikokolakis and C. Palamidessi, "Making random choices invisible to the scheduler," *Inf. and Comp.*, 2010.
- [34] M. E. Andrés, C. Palamidessi, P. van Rossum, and A. Sokolova, "Information hiding in probabilistic concurrent systems," Tech. Rep., 2010.