

Formal Approaches to Information-Hiding (Tutorial*)

Romain Beauxis¹, Konstantinos Chatzikokolakis¹, Catuscia Palamidessi¹,
and Prakash Panangaden²

¹ INRIA and LIX, École Polytechnique, Palaiseau, France
{beauxis, kostas, catuscia}@lix.polytechnique.fr
² Mc Gill University, Montreal, Canada
prakash@cs.mcgill.ca

Abstract. In this survey paper we consider the class of protocols for information-hiding which use randomization to obfuscate the link between the observables and the information to be protected. We focus on the problem of formalizing the notion of information hiding, and verifying that a given protocol achieves the intended degree of protection. Without the pretense of being omni-comprehensive, we review the main approaches that have been explored in literature: possibilistic, probabilistic, information-theoretic, and statistical.

1 Introduction

During the last decade, internet activities have become an important part of many people's lives. As the number of these activities increases, there is a growing amount of personal information about the users that is stored in electronic form and that is usually transferred using public electronic means. This makes it feasible and often easy to collect, transfer and process a huge amount of information about a person. As a consequence, the need for mechanisms to protect such information is compelling.

A recent example of such privacy concerns are the so-called “biometric” passports. These passports, used by many countries and required by all visa waiver travelers to the United States, include a RFID chip containing information about the passport's owner. These chips can be read wirelessly without any contact with the passport and without the owner even knowing that his passport is being read. It is clear that such devices need protection mechanisms to ensure that the contained information will not be revealed to any non-authorized person.

In general, privacy can be defined as the ability of users to stop information about themselves from becoming known to people other than those they choose to give the information to. We can further categorize privacy properties based on the nature of the hidden information. *Data protection* usually refers to confidential data like the credit card number. *Anonymity*, on the other hand, concerns the identity of the user who performed a certain action. *Unlinkability* refers to the link between the information and the user, and *unobservability* regards the actions of a user.

* This work has been partially supported by the INRIA DREI Équipe Associée PRINTEMPS and by the INRIA ARC project ProNoBiS.

Information-hiding protocols aim at ensuring a privacy property during an electronic transaction. For example, the voting protocol Foo 92 ([1]) allows a user to cast a vote without revealing the link between the voter and the vote. The anonymity protocol Crowds ([2]) allows a user to send a message on a public network without revealing the identity of the sender.

Several anonymity protocols use randomized primitives to obtain the obfuscation of the information to be protected. This is the case, for instance, of the Dining Cryptographers [3], which use coin-flipping, Crowds [2] and Onion Routing [4], which select randomly another user of the network to forward the message to, and Freenet [5]. In this survey, we restrict our attention to the case in which the use of randomization is critical to achieve the intended security properties.

2 The Possibilistic Approaches

These are by far the approaches which have been explored the most in literature. Various formal definitions and frameworks for analyzing information-hiding have been developed. Some examples of these approaches are those based on epistemic logic ([6,7]), on “function views” ([8]), and on process-calculi ([9,10]). Here we focus on the last kind of approach.

Often possibilistic approaches rely on *nondeterminism*: a protocol provides protection if the set of possible observable outcomes is saturated with respect to the secrets. More precisely, if in one computation the instance of the secret to protect is s and the observable outcome is o , then for every other instance s' there must be a computation where, with secret s' , the observable is still o . Formally:

$$f^{-1}(f(P)) \sim P$$

where P is the protocol, and f is a relabeling function that maps all the secrets into a dummy, and \sim is a chosen equivalence relation [9].

A related approach is the one by [11,12], where the authors specify privacy in electronic voting (protection of the secrecy of the vote) as the property that if we swap the way in which two users, A and B , vote, then the resulting system is weakly bisimilar to the original one. Formally:

$$C[A[a/v]|B[b/v]] \approx C[A[a/v]|B[b/v]]$$

where a, b represent the votes of A and B respectively, and the context $C[\]$ represents the rest of the protocol.

This kind of approach is reasonable, as long as the protocols of interest do not involve the use of randomization. In case they do, then we have a problem, because the pure possibilistic approach is unable to cope with probabilities. So, the choice is either to move to a probabilistic approach, or to try to abstract from probabilities. The second choice is explored in [9]: In that paper, the authors replace probabilistic choice by nondeterministic choice, and then apply the usual possibilistic definition.

We now illustrate the above idea on the example of the dining cryptographers.




Fig. 1. Chaum's protocol for the dining cryptographers [3]

2.1 The Dining Cryptographers' Problem

This problem, described by Chaum in [3], involves a situation in which three cryptographers are dining together. At the end of the dinner, each of them is secretly informed by the master whether he should pay the bill or not. So, either the master will pay, or he will ask one of the cryptographers to pay. The cryptographers, or some external observer, would like to find out whether the payer is one of them or the master. However, if the payer is one of them, the cryptographers wish to maintain anonymity over the identity of the payer. Of course, we assume that the master himself will not reveal this information, and also we want the solution to be distributed, i.e. communication can be achieved only via message passing, and there is no central memory or central 'coordinator' which can be used to find out this information.

A possible solution to this problem, described in [3], is the following: Each cryptographer tosses a coin, which is visible to himself and to his neighbor to the right. Each cryptographer then observes the two coins that he can see, and announces *agree* or *disagree*. If a cryptographer is not paying, he will announce *agree* if the two sides are the same and *disagree* if they are not. However, if he is paying then he will say the opposite. It can be proved that if the number of *disagrees* is even, then the master is paying; otherwise, one of the cryptographers is paying. Furthermore, if one of the cryptographers is paying, then neither an external observer nor the other two cryptographers can identify, from their individual information, who exactly is paying.

In order to specify formally the protocol, we use a probabilistic version of the π -calculus, π_p , which is essentially the π -calculus enriched with a probabilistic choice operator \oplus_p . For a precise definition of the semantics of π_p we refer to [13].

The protocol can be described as the parallel composition of the master process *Master*, the cryptographers processes *Crypt_i*, of the coin processes *Coin_h*, and of a process *Collect*¹ whose purpose is to collect all the declarations of the cryptographers,

¹ The presence of the process *Collect* is due to technical reasons that have to do with the control of the power of the scheduler, and are out of the scope of this paper.

Table 1. The dining cryptographers protocol expressed in π_p

$$\begin{aligned}
Master &= \overline{m}_0\langle 0 \rangle . \overline{m}_1\langle 0 \rangle . \overline{m}_2\langle 0 \rangle \oplus_p \bigoplus_0^2 p_i \overline{m}_{0+i}\langle 1 \rangle . \overline{m}_{1+i}\langle 0 \rangle . \overline{m}_{2+i}\langle 0 \rangle \\
Crypt_i &= c_{i,i}(x_0) . c_{i,i+1}(x_1) . m_i(x) . \overline{p}a\overline{y}_i\langle x \rangle . \overline{out}_i\langle x_0 + x_1 + x \rangle \\
Coin_h &= \overline{c}_{h-1,h}\langle 0 \rangle . \overline{c}_{h,h}\langle 0 \rangle \oplus_{p_h} \overline{c}_{h-1,h}\langle 1 \rangle . \overline{c}_{h,h}\langle 1 \rangle \\
Collect &= out_0\langle y_0 \rangle . out_1\langle y_1 \rangle . out_2\langle y_2 \rangle . \overline{outall}\langle y_0, y_1, y_2 \rangle \\
DC &= (\nu \vec{c})(\nu \vec{m})(\nu \vec{out})(Master \mid \prod_i Crypt_i \mid \prod_h Coin_h \mid Collect)
\end{aligned}$$

Table 2. The nondeterministic version of the dining cryptographers protocol expressed in π

$$\begin{aligned}
Master &= \overline{m}_0\langle 0 \rangle . \overline{m}_1\langle 0 \rangle . \overline{m}_2\langle 0 \rangle + \sum_0^2 p_i \overline{m}_{0+i}\langle 1 \rangle . \overline{m}_{1+i}\langle 0 \rangle . \overline{m}_{2+i}\langle 0 \rangle \\
Crypt_i &= c_{i,i}(x_0) . c_{i,i+1}(x_1) . m_i(x) . \overline{p}a\overline{y}_i\langle x \rangle . \overline{out}_i\langle x_0 + x_1 + x \rangle \\
Coin_h &= \overline{c}_{h-1,h}\langle 0 \rangle . \overline{c}_{h,h}\langle 0 \rangle \oplus_{p_h} \overline{c}_{h-1,h}\langle 1 \rangle . \overline{c}_{h,h}\langle 1 \rangle \\
Collect &= out_0\langle y_0 \rangle . out_1\langle y_1 \rangle . out_2\langle y_2 \rangle . \overline{outall}\langle y_0, y_1, y_2 \rangle \\
DC &= (\nu \vec{c})(\nu \vec{m})(\nu \vec{out})(Master \mid \prod_i Crypt_i \mid \prod_h Coin_h \mid Collect)
\end{aligned}$$

and output them in the form of a tuple. See Table 1. In this protocol, the secret actions are $\overline{p}a\overline{y}_i\langle x \rangle$, and the observable actions are $\overline{outall}\langle y_0, y_1, y_2 \rangle$.

2.2 Nondeterministic Version of the Dining Cryptographers

In the approach of [9] the dining cryptographers are formalized as a purely nondeterministic system: the coins are approximated by nondeterministic coins, and the choice on who pays the bill is also nondeterministic. The specification of the solution can be given in π -calculus as illustrated in Table 2 (in the original work [9] the authors used CSP [14]).

Let f be the function $f(\overline{p}a\overline{y}_i) = \overline{p}a\overline{y}$ and $f(\alpha) = \alpha$ for all the other actions. It is possible to check that $f^{-1}(f(DC)) \sim_T DC$, where we recall that \sim_T stands for trace equivalence. Hence the nondeterministic notion of anonymity, as defined at the beginning of this section, is satisfied.

As a consequence of approximating the coins by nondeterministic coins, we cannot differentiate between a fair coin and a biased one. However, it is evident that the fairness of the coins is essential to ensure the anonymity property in the system, as illustrated by the following example.

Example 1. Assume that, whenever a cryptographer pays, an external observer obtains *almost always* one of the three outcomes represented in Figure 2, where a stands for *agree* and d for *disagree*. More precisely, assume that these three outcomes appear with a frequency of 33% each, while the missing configuration, d, a, a , appears with a frequency of only 1%. What can the observer deduce? By examining all possible cases, it is easy to see that the coins must be biased, and more precisely, $Coin_0$ and $Coin_1$




Fig. 2. Illustration of Example 1: The results that are observed with high frequency

must produce almost always *head*, and *Coin*₂ must produce almost always *tail* (or vice-versa). From this estimation, it is immediate to conclude that, in the first case, the payer is *almost for sure* *Crypt*₁, in the second case *Crypt*₂, and in the third case *Crypt*₀.

In the situation illustrated in the above example, clearly, the system does not provide anonymity. However the nondeterministic definition of anonymity is still satisfied (and it is satisfied in general, as long as “almost always” is not “always”, i.e. the fourth configuration *d, a, a* also appears, from time to time). The problem is that the nondeterministic definition can only express whether or not it is possible to have a particular outcome, but cannot express whether one outcome is more likely than the other.

3 The Probabilistic Approaches

The probabilistic approaches have been investigated in particular in the field of anonymity, and almost exclusively in the strongest form, namely to express the property that the observables reveal no (quantitative) information about the secrets (*strong anonymity*).

There are essentially three probabilistic notions considered in literature: one based on the equality of the *a posteriori* probabilities, one based on the equality between the *a posteriori* probability and the *a priori* probability, and one based on the equality of the *likelihoods*. All of them involve the notion of conditional probability $p(a|b)$, which represents the probability of the event *a*, given the event *b*. We recall the equality known as Bayes theorem:

$$p(a|b) = \frac{p(b|a)p(a)}{p(b)}$$

These probabilistic notions also require that the secrets are mutually exclusive events, and that $\sum_s p(s) = 1$. The same for the observables.

Equality of the *a posteriori* probabilities. The idea is to say that a system is strongly secure if, for any observable *o*, the *a posteriori* probability of a secret *s* (namely $p(s|o)$) is the same as the *a posteriori* probability of any other secret *s'*. Formally:

$$p(s|o) = p(s'|o) \quad \text{for all observables } o, \text{ and secrets } s \text{ and } s' \quad (1)$$

This is similar to the definition of *strong anonymity* by Halpern and O'Neill [7], although their setting is different, being based on a probabilistic version of epistemic logic.

Equality of the *a posteriori* and *a priori* probabilities. The idea is to say that a system is strongly secure if, for any observable o , the *a posteriori* probability of a secret s is the same as its *a priori* one. In other words, the observation does not add anything to the expectation that the secret is s . Formally:

$$p(s|o) = p(s) \quad \text{for all observables } o, \text{ and secrets } s \quad (2)$$

This is the definition of *anonymity* adopted by Chaum in [3]. He also proved that the DC satisfies it if the coins are fair. Halpern and O'Neill also consider a similar property in their epistemological setting, and they call it *conditional anonymity* [7].

Equality of the likelihoods. The idea is to say that a system is strongly secure if, for any observable o , the likelihood of a secret s given o (namely $p(o|s)$) is the same as the same as the likelihood of any other secret s' . Formally:

$$p(o|s) = p(o|s') \quad \text{for all observables } o, \text{ and secrets } s \text{ and } s' \quad (3)$$

This was proposed as definition of *strong anonymity* by Bhargava and Palamidessi [15].

3.1 Comparison

It is easy to see that definitions (2) and (3) are equivalent. In fact:

(2) \Rightarrow (3)

$$\begin{aligned} p(o|s) &= \frac{p(s|o)p(o)}{p(s)} && \text{by Bayes theorem} \\ &= p(o) && \text{by (2)} \\ &= \frac{p(s'|o)p(o)}{p(s')} && \text{by (2)} \\ &= p(o|s') && \text{by Bayes theorem} \end{aligned}$$

(2) \Leftarrow (3)) We prove that $p(o|s) = p(o)$ for all observables o , and secrets s . From this it is immediate to derive (2) by applying Bayes theorem.

$$\begin{aligned} p(o) &= \sum_s p(o \text{ and } s) && \text{by the disjointness of the secrets} \\ &= \sum_s p(o|s) p(s) && \text{by definition of conditional probability} \\ &= p(o|s) \sum_s p(s) && \text{by (3)} \\ &= p(o|s) && \text{since } \sum_s p(s) = 1 \end{aligned}$$

Definition (3) has the advantage that it makes clear that depends only on the protocol, not in the distribution on the secrets, and, more important, it does extend in a natural way to the case in which the choice of the secret is done nondeterministically. See [15] for more details.

Concerning definition (1), it probably looks at a first site the most natural, but it actually turns out to be too strong, that one is strictly stronger than (2) and (3). In fact it is equivalent to (2) and (3), *plus* the condition that the probability distribution of the secrets is uniform, namely

$$p(s) = p(s') \quad \text{for all secrets } s \text{ and } s' \quad (4)$$

(1) \Rightarrow (4))

$$\begin{aligned} p(s) &= \sum_o p(s \text{ and } o) && \text{by the disjointness of the secrets} \\ &= \sum_o p(s|o) p(o) && \text{by definition of conditional probability} \\ &= \sum_o p(s'|o) p(o) && \text{by (1)} \\ &= \sum_o p(s' \text{ and } o) \\ &= p(s') \end{aligned}$$

(1) \Rightarrow (3))

$$\begin{aligned} p(o|s) &= \frac{p(s|o) p(o)}{p(s)} && \text{by Bayes theorem} \\ &= \frac{p(s'|o) p(o)}{p(s)} && \text{by (1)} \\ &= \frac{p(s'|o) p(o)}{p(s')} && \text{by (4)} \\ &= p(o|s') && \text{by Bayes theorem} \end{aligned}$$

(1) \Leftarrow (3),(4)

$$\begin{aligned} p(s|o) &= \frac{p(o|s) p(s)}{p(o)} && \text{by Bayes theorem} \\ &= \frac{p(o|s') p(s)}{p(o)} && \text{by (3)} \\ &= \frac{p(o|s') p(s')}{p(o)} && \text{by (4)} \\ &= p(s'|o) && \text{by Bayes theorem} \end{aligned}$$

It is interesting to notice that (4) can be split in two orthogonal properties: one which depends only in the protocol ((3)), and one which depends only in the distribution on the secrets ((4)).

In our opinion condition (4) is not a suitable condition for defining the notion of protection provided by a protocol, because it only depends on the distribution on the

secret data, which can be influenced by the users, but not by the protocol. We believe that a good notion of protection should abstract from such distribution. In this sense we consider (1) too strong.

There are also weaker notions of protection, still based on the comparison between conditional probabilities, which have been investigated in literature. In particular, Rubin and Reiter proposed the concepts of *possible innocence* and of *probable innocence* [2]. See also [16] for a generalization of the latter.

The need for formalizing weaker forms of protection comes from the fact that the strong properties discussed above are almost never achieved in practice. Hence the need to express in a quantitative way the *degree* of protection. Researchers have been exploring for suitable notions within the well-established fields of Information Theory and of Statistics.

4 Information Theory

Recently it has been observed that at an abstract level information-hiding protocols can be viewed as *channels* in the information-theoretic sense. A channel consists of a set of input values \mathcal{S} , a set of output values \mathcal{O} (the observables) and a transition matrix which gives the conditional probability $p(o|s)$ of producing o as the output when s is the input. In the case of protocols for information hiding, \mathcal{S} contains the secret information that we want to protect and \mathcal{O} the facts that the attacker can observe.

Let us revise some of the basic concepts of Information Theory: Let X be a random variable. The *entropy* $H(X)$ of X is defined as

$$H(X) = - \sum_{x \in \mathcal{X}} p(x) \log p(x)$$

The entropy measures the uncertainty of a random variable. It takes its maximum value $\log |\mathcal{X}|$ when X 's distribution is uniform and its minimum value 0 when X is constant. We usually take the logarithm with a base 2 and measure entropy in *bits*. Roughly speaking, m bits of entropy means that we have 2^m values to choose from, assuming a uniform distribution.

The *relative entropy* or *Kullback–Leibler distance* between two probability distributions p, q on the same set \mathcal{X} is defined as

$$D(p \parallel q) = \sum_{x \in \mathcal{X}} p(x) \log \frac{p(x)}{q(x)}$$

It is possible to prove that $D(p \parallel q)$ is always non-negative, and it is 0 if and only if $p = q$.

Now let X, Y be random variables. The *conditional entropy* $H(X|Y)$ is

$$H(X|Y) = - \sum_{y \in \mathcal{Y}} p(y) \sum_{x \in \mathcal{X}} p(x|y) \log p(x|y)$$

Conditional entropy measures the amount of uncertainty of X when Y is known. It can be shown that $0 \leq H(X|Y) \leq H(X)$. It takes its maximum value $H(X)$ when Y

reveals no information about X , and its minimum value 0 when Y completely determines the value of X .

Comparing $H(X)$ and $H(X|Y)$ gives us the concept of *mutual information* $I(X; Y)$, which is defined as

$$I(X; Y) = H(X) - H(X|Y)$$

Mutual information measures the amount of information that one random variable contains about another random variable. In other words, it measures the amount of uncertainty about X that we lose when observing Y . It can be shown that it is symmetric ($I(X; Y) = I(Y; X)$) and that $0 \leq I(X; Y) \leq H(X)$.

The maximum mutual information between X and Y over all possible distributions $p(x)$ is known as the channel's *capacity*:

$$C = \max_{p(x)} I(X; Y)$$

The capacity of a channel gives the maximum rate at which information can be transmitted using this channel.

In the following we recall some of the notions of protection, based on information-theoretic notions, which have been proposed in literature.

In [17,18] the authors propose a notion of anonymity based on the entropy of the users. The idea is to represent the lack of information that an attacker has about the secrets. Note that this is not in line with our point of view: in our opinion the interesting thing is to model the capability of the protocol to conceal the secret information despite of the observables that are made available to the attacker.

Zhu and Bettati propose in [19] a definition of anonymity based on mutual information.

In [20,21] the authors study the ability to have covert communication as a result of non-perfect anonymity. In [21] the authors suggest that the channel's capacity can be used as an asymptotic measure of the worst-case loss of anonymity.

In [22] we explore the implications of adopting the (converse of the) notion of capacity as measure of the degree of protection, and we introduce a more general concept that we call *conditional capacity*.

Note that the capacity is an abstraction of mutual information obtained by maximizing over the possible input distributions. As a consequence, we get a measure that depends only on the protocol and not on the input distribution, which is an advantage with respect to the mutual-information approach because in general we don't know the input distribution, and it also may change over time. Of course, in case we know the input distribution, then the mutual-information approach is more precise because it gives the exact loss of anonymity for the specific situation.

In [23] the authors use the Kullback-Leibler distance to perform a metric analysis of anonymity.

In [24] the authors define as information leakage the difference between the a priori accuracy of the guess of the attacker, and the a posteriori one, after the attacker has made his observation. The accuracy of the guess is defined as the Kullback-Leibler

distance between the *belief* (which is a weight attributed by the attacker to each input hypothesis) and the true distribution on the hypotheses.

In the field of information flow and non-interference there is a line of research which is closely related. There have been various works [25,26,27,28,29,30] in which the *high information* and the *low information* are seen as the input and output respectively of a channel. From an abstract point of view, the setting is very similar; technically it does not matter what kind of information we are trying to conceal, what is relevant for the analysis is only the probabilistic relation between the input and the output information.

5 Hypothesis Testing

In information-hiding systems the attacker finds himself in the following scenario: he cannot directly detect the information of interest, namely the actual value of the random variable $S \in \mathcal{S}$, but he can discover the value of another random variable $O \in \mathcal{O}$ which depends on S according to a known conditional distribution. This kind of situation is quite common also in other disciplines, like medicine, biology, and experimental physics, to mention a few. The attempt to infer S from O is called *hypothesis testing* (the ‘‘hypothesis’’ to be validated is the actual value of S), and it has been widely investigated in statistics.

In this section we discuss possible methods by which an adversary can try to infer the secrets from the observables, and consider the corresponding probability of error, that is, the probability that the adversary draws the wrong conclusion. We regard the probability of error as a representative of the degree of protection provided by the protocol, and we study its properties with respect to the associated matrix.

We start by recalling the notion of *decision function*, which represent the guess the adversary makes about the secrets, for each observable: a decision function is simply any function $f : \mathcal{O} \rightarrow \mathcal{S}$.

The *probability of error* associated to a decision function f is the probability of guessing the wrong hypothesis by using f , averaged on all possible observables. In general the probability of error depends on the input distribution and on the channel’s matrix. We will use the notation $\mathcal{P}(f, M, \vec{p})$ to represent the probability of error associated to the decision function f , the channel’s matrix M , and the input distribution \vec{p} . The following characterization of $\mathcal{P}(f, M, \vec{p})$ is well-known in literature, see for instance [31].

$$\mathcal{P}(f, M, \vec{p}) = 1 - \sum_{\mathcal{O}} p(o|f(o))p_{f(o)} \quad (5)$$

Given a channel $(\mathcal{S}, \mathcal{O}, M)$, the best decision function that the adversary can use, namely the one that minimizes the probability of error, is the one associated to the so-called MAP rule, which prescribes to choose the hypothesis s which has *Maximum A Posteriori Probability* (for a given $o \in \mathcal{O}$), namely the s for which $p(s|o)$ is maximum. The fact that the MAP rule represent the ‘best bet’ of the adversary is rather intuitive, and well known in literature. We refer to [31] for a formal proof.

The MAP rule is used in the so-called *Bayesian approach* to hypothesis testing, and the corresponding probability of error is also known as *Bayes risk*. We will denote it by $\mathcal{P}_{MAP}(M, \vec{p})$. The following characterization is an immediate consequence of (5) and of the Bayes theorem $p(s|o) = p(o|s)p_s/p(o)$.

$$\mathcal{P}_{MAP}(M, \vec{p}) = 1 - \sum_{\mathcal{O}} \max_s (p(o|s)p_s)$$

In [32] we have proposed to express the degree of protection Pt provided by a protocol T in terms of the probability of error of the corresponding matrix $M(T)$:

$$Pt_{MAP}(T, \vec{p}) = \mathcal{P}_{MAP}(M(T), \vec{p})$$

The problem with the MAP rule is that it assumes that the input distribution is known to the adversary. This is often not the case, so it is natural to try to approximate it with some other rule. One such rule is the so-called ML rule, which prescribes to choose the s which has *Maximum Likelihood* (for a given $o \in \mathcal{O}$), namely the s for which $p(o|s)$ is maximum. The name comes from the fact that $p(o|s)$ is called the *likelihood* of s given o . We will denote the corresponding probability of error by $\mathcal{P}_{ML}(M, \vec{p})$. The following characterization is an immediate consequence of (5).

$$\mathcal{P}_{ML}(M, \vec{p}) = 1 - \sum_{\mathcal{O}} \max_s (p(o|s))p_s$$

It has been shown (see for instance [22]) that under certain conditions on the matrix, the ML rule approximates indeed the MAP rule, in the sense that by repeating the protocol the adversary can make the probability of error arbitrarily close to 0, with either rule.

We have also explored, in [32], the possibility of defining the degree of protection provided by a term T under the ML rule as $\mathcal{P}_{ML}(M(T), \vec{p})$, but it did not seem reasonable to give a definition that depends on the input distribution, since the main reason to apply a non-Bayesian approach is that we do not know the input distribution. Instead, we have defined the degree of protection associated to a process term as the *average* probability of error with respect to all possible distributions \vec{p} :

$$Pt_{ML}(T) = (m - 1)! \int_{\vec{p}} \mathcal{P}_{ML}(M(T), \vec{p}) d\vec{p}$$

In previous definition, $(m - 1)!$ represents a normalization function: $\frac{1}{(m-1)!}$ is the hyper-volume of the domain of all possible distributions \vec{p} on \mathcal{S} , namely the $(m - 1)$ -dimensional space of points \vec{p} such that $0 \leq p_s \leq 1$ and $0 \leq \sum_{s \in \mathcal{S}} p_s = 1$ (where m is the cardinality of \mathcal{S}).

Fortunately, it turns out that this definition is equivalent to a much simpler one: the average value of the probability of error, under the Maximum Likelihood rule, can be obtained simply by computing \mathcal{P}_{ML} on the uniform distribution $\vec{p}_u = (\frac{1}{m}, \frac{1}{m}, \dots, \frac{1}{m})$:

$$Pt_{ML}(T) = \min_{\zeta \in \mathcal{A}} \mathcal{P}_{ML}(M_{\zeta}(T), \vec{p}_u)$$

We believe that the probability of error is a central notion for information-hiding, and we expect that it will be thoroughly explored in the next future.

In [33] we have characterized the Bayes risk in terms of the solution of certain systems of equations derived from the matrix of the channel. This has led to an algorithm to compute the maximum value of the Bayes risk. Furthermore, it has allowed us to improve functional bounds on the Bayes risk.

In [32] we have studied how the operators of π_p affect the probability of error. In particular, we have characterized constructs that have the property of not decreasing the degree of protection, and that can therefore be considered safe in the modular construction of protocols. As a case study, we apply these techniques to the Dining Cryptographers, and we are able to derive a generalization of Chaum's strong anonymity result for the Dining Cryptographers. More precisely, we have shown that the Dining Cryptographers on an arbitrary graph (where the nodes are the cryptographers and the arcs are the coins) is strongly anonymous if and only there is a spanning tree formed entirely of fair coins.

6 Computing the Matrix Associated to a Protocol

In this section we show how to compute the matrix associated to a protocol specified in π_p , using our (very preliminary) π_p model checker VAMP (<http://vamp.gforge.inria.fr/>).

We consider the protocol for the DC represented in Table 1. Assume we want to compute $p(o|s_i)$, where s_i represents the fact that the cryptographer i is the payer, and o is of the form $outall\langle y_o, y_1, y_2 \rangle$. We redefine the *Master* to be

$$\bar{m}_i\langle 1 \rangle . \bar{m}_{i+1}\langle 0 \rangle . \bar{m}_{i+2}\langle 0 \rangle$$

Then, we run the resulting process *DC* in VAMP, with query o . VAMP gives as result the (unconditional) probability of executing o in the new specification, which corresponds to the conditional probability $p(o|s_i)$ in the original specification.

We have computed various channel matrices, for different values of the probability p that a coin gives heads (we assume that each coin is biased in the same way). The results are shown in Fig. 3.

Finally, from the matrix, we can compute the capacity. This can be done, in general, by using the Arimoto-Blahut approximation algorithm, or, under certain symmetry conditions, we can apply a formula (see [22] for more details).

In this case we could apply the formula because the conditions are satisfied. The resulting graph is displayed in Fig. 4. As expected, when $p = 0.5$ the protocol is strongly anonymous and the relative loss of anonymity is 0. When p approaches 0 or 1, the attacker can deduce the identity of the payer with increasingly high probability, so the capacity increases. In the extreme case where the coins are totally biased the attacker can be sure about the payer, and the capacity takes its maximum value of $\log 3$.

	<i>daa</i>	<i>ada</i>	<i>aad</i>	<i>ddd</i>	<i>aaa</i>	<i>dda</i>	<i>dad</i>	<i>add</i>
c_1	0.25	0.25	0.25	0.25	0	0	0	0
c_2	0.25	0.25	0.25	0.25	0	0	0	0
c_3	0.25	0.25	0.25	0.25	0	0	0	0
m	0	0	0	0	0.25	0.25	0.25	0.25

	<i>daa</i>	<i>ada</i>	<i>aad</i>	<i>ddd</i>	<i>aaa</i>	<i>dda</i>	<i>dad</i>	<i>add</i>
c_1	0.28	0.24	0.24	0.24	0	0	0	0
c_2	0.24	0.28	0.24	0.24	0	0	0	0
c_3	0.24	0.24	0.28	0.24	0	0	0	0
m	0	0	0	0	0.28	0.24	0.24	0.24

	<i>daa</i>	<i>ada</i>	<i>aad</i>	<i>ddd</i>	<i>aaa</i>	<i>dda</i>	<i>dad</i>	<i>add</i>
c_1	0.37	0.21	0.21	0.21	0	0	0	0
c_2	0.21	0.37	0.21	0.21	0	0	0	0
c_3	0.21	0.21	0.37	0.21	0	0	0	0
m	0	0	0	0	0.37	0.21	0.21	0.21

	<i>daa</i>	<i>ada</i>	<i>aad</i>	<i>ddd</i>	<i>aaa</i>	<i>dda</i>	<i>dad</i>	<i>add</i>
c_1	0.52	0.16	0.16	0.16	0	0	0	0
c_2	0.16	0.52	0.16	0.16	0	0	0	0
c_3	0.16	0.16	0.52	0.16	0	0	0	0
m	0	0	0	0	0.52	0.16	0.16	0.16

	<i>daa</i>	<i>ada</i>	<i>aad</i>	<i>ddd</i>	<i>aaa</i>	<i>dda</i>	<i>dad</i>	<i>add</i>
c_1	0.73	0.09	0.09	0.09	0	0	0	0
c_2	0.09	0.73	0.09	0.09	0	0	0	0
c_3	0.09	0.09	0.73	0.09	0	0	0	0
m	0	0	0	0	0.73	0.09	0.09	0.09

Fig. 3. The channel matrices for probability of heads $p = 0.5$, $p = 0.6$, $p = 0.7$, $p = 0.8$, and $p = 0.9$




Fig. 4. The degree of anonymity in the Dining Cryptographers as a function of the coins' probability to yield heads

References

1. Fujioka, A., Okamoto, T., Ohta, K.: A practical secret voting scheme for large scale elections. In: ASIACRYPT 1992: Proceedings of the Workshop on the Theory and Application of Cryptographic Techniques, pp. 244–251. Springer, London, UK (1993)
2. Reiter, M.K., Rubin, A.D.: Crowds: anonymity for Web transactions. *ACM Transactions on Information and System Security* 1, 66–92 (1998)
3. Chaum, D.: The dining cryptographers problem: Unconditional sender and recipient untraceability. *Journal of Cryptology* 1, 65–75 (1988)
4. Syverson, P., Goldschlag, D., Reed, M.: Anonymous connections and onion routing. In: IEEE Symposium on Security and Privacy, Oakland, California, pp. 44–54 (1997)
5. Clarke, I., Sandberg, O., Wiley, B., Hong, T.W.: Freenet: A distributed anonymous information storage and retrieval system. In: Federrath, H. (ed.) *Designing Privacy Enhancing Technologies*. LNCS, vol. 2009, pp. 44–66. Springer, Heidelberg (2001)
6. Syverson, P.F., Stubblebine, S.G.: Group principals and the formalization of anonymity. *World Congress on Formal Methods (1)*, 814–833 (1999)
7. Halpern, J.Y., O'Neill, K.R.: Anonymity and information hiding in multiagent systems. *Journal of Computer Security* 13, 483–512 (2005)
8. Hughes, D., Shmatikov, V.: Information hiding, anonymity and privacy: a modular approach. *Journal of Computer Security* 12, 3–36 (2004)
9. Schneider, S., Sidiropoulos, A.: CSP and anonymity. In: Martella, G., Kurth, H., Montolivo, E., Bertino, E. (eds.) *ESORICS 1996*. LNCS, vol. 1146, pp. 198–218. Springer, Heidelberg (1996)
10. Ryan, P.Y., Schneider, S.: *Modelling and Analysis of Security Protocols*. Addison-Wesley, Reading (2001)
11. Delaune, S., Kremer, S., Ryan, M.D.: Verifying properties of electronic voting protocols. In: *Proceedings of the IAVoSS Workshop On Trustworthy Elections (WOTE 2006)*, Cambridge, UK, 45–52 (2006)

12. Delaune, S., Kremer, S., Ryan, M.: Coercion-resistance and receipt-freeness in electronic voting. In: Computer Security Foundations Workshop, pp. 28–42. IEEE Computer Society, Los Alamitos (2006)
13. Chatzikokolakis, K., Palamidessi, C.: A framework for analyzing probabilistic protocols and its application to the partial secrets exchange. *Theoretical Computer Science* (2005); In: De Nicola, R., Sangiorgi, D. (eds.) TGC 2005. LNCS, vol. 3705, pp. 146–162. Springer, Heidelberg (2005), <http://www.lix.polytechnique.fr/~catuscia/papers/PartialSecrets/TCSreport.pdf>
14. Hoare, C.A.R.: *Communicating Sequential Processes*. Prentice-Hall, Englewood Cliffs (1985)
15. Bhargava, M., Palamidessi, C.: Probabilistic anonymity. In: Abadi, M., de Alfaro, L. (eds.) CONCUR 2005. LNCS, vol. 3653, pp. 171–185. Springer, Heidelberg (2005), <http://www.lix.polytechnique.fr/~catuscia/papers/Anonymity/concur.pdf>
16. Chatzikokolakis, K., Palamidessi, C.: Probable innocence revisited. *Theoretical Computer Science* 367, 123–138 (2006), <http://www.lix.polytechnique.fr/~catuscia/papers/Anonymity/tcsPI.pdf>
17. Serjantov, A., Danezis, G.: Towards an information theoretic metric for anonymity. In: Dingleline, R., Syverson, P.F. (eds.) PET 2002. LNCS, vol. 2482, pp. 41–53. Springer, Heidelberg (2003)
18. Díaz, C., Seys, S., Claessens, J., Preneel, B.: Towards measuring anonymity. In: Dingleline, R., Syverson, P.F. (eds.) PET 2002. LNCS, vol. 2482, pp. 54–68. Springer, Heidelberg (2003)
19. Zhu, Y., Bettati, R.: Anonymity vs. information leakage in anonymity systems. In: Proc. of ICDCS, pp. 514–524. IEEE Computer Society, Los Alamitos (2005)
20. Moskowitz, I.S., Newman, R.E., Crepeau, D.P., Miller, A.R.: Covert channels and anonymizing networks. In: Jajodia, S., Samarati, P., Syverson, P.F. (eds.) WPES, ACM, pp. 79–88. ACM, New York (2003)
21. Moskowitz, I.S., Newman, R.E., Syverson, P.F.: Quasi-anonymous channels. In: IASTED CNIS, pp. 126–131 (2003)
22. Chatzikokolakis, K., Palamidessi, C., Panangaden, P.: Anonymity protocols as noisy channels. In: *Information and Computation* (to appear, 2007), <http://www.lix.polytechnique.fr/~catuscia/papers/Anonymity/Channels/full.pdf>
23. Deng, Y., Pang, J., Wu, P.: Measuring anonymity with relative entropy. In: Dimitrakos, T., Martinelli, F., Ryan, P.Y.A., Schneider, S. (eds.) FAST 2006. LNCS, vol. 4691, pp. 65–79. Springer, Heidelberg (2007)
24. Clarkson, M.R., Myers, A.C., Schneider, F.B.: Belief in information flow. *Journal of Computer Security*. Available as Cornell Computer Science Department Technical Report TR 2007-207 (to appear, 2008)
25. McLean, J.: Security models and information flow. In: *IEEE Symposium on Security and Privacy*, pp. 180–189 (1990)
26. Gray III, J.W.: Toward a mathematical foundation for information flow security. In: *Proceedings of the 1991 IEEE Computer Society Symposium on Research in Security and Privacy (SSP 1991)*, pp. 21–35. IEEE, Washington - Brussels - Tokyo (1991)
27. Clark, D., Hunt, S., Malacaria, P.: Quantitative analysis of the leakage of confidential data. In: *Proc. of QAPL 2001*. *Electr. Notes Theor. Comput. Sci*, vol. 59 (3), pp. 238–251. Elsevier Science, Amsterdam (2001)
28. Clark, D., Hunt, S., Malacaria, P.: Quantified interference for a while language. In: *Proc. of QAPL 2004*. *Electr. Notes Theor. Comput. Sci*, vol. 112, pp. 149–166. Elsevier Science, Amsterdam (2005)

29. Lowe, G.: Quantifying information flow. In: Proc. of CSFW 2002, pp. 18–31. IEEE Computer Society Press, Los Alamitos (2002)
30. Boreale, M.: Quantifying information leakage in process calculi. In: Bugliesi, M., Preneel, B., Sassone, V., Wegener, I. (eds.) ICALP 2006. LNCS, vol. 4052, pp. 119–131. Springer, Heidelberg (2006)
31. Cover, T.M., Thomas, J.A.: Elements of Information Theory. John Wiley & Sons, Chichester (1991)
32. Compositional Methods for Information-Hiding. In: FOSSACS 2008, LNCS, vol. 4962, pp. 443–457, Springer, Heidelberg (to appear, 2008)
33. Chatzikokolakis, K., Palamidessi, C., Panangaden, P.: Probability of error in information-hiding protocols. In: Proceedings of the 20th IEEE Computer Security Foundations Symposium (CSF20), pp. 341–354. IEEE Computer Society, Los Alamitos (2007), <http://www.lix.polytechnique.fr/~catuscia/papers/ProbabilityError/full.pdf>