



Probable Innocence Revisited

Konstantinos Chatzikokolakis, Catuscia Palamidessi

► **To cite this version:**

Konstantinos Chatzikokolakis, Catuscia Palamidessi. Probable Innocence Revisited. Theoretical Computer Science, Elsevier, 2006, 367 (1-2), pp.123-138. <10.1016/j.tcs.2006.08.033>. <inria-00201072>

HAL Id: inria-00201072

<https://hal.inria.fr/inria-00201072>

Submitted on 23 Dec 2007

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.

Probable Innocence Revisited[★]

Konstantinos Chatzikokolakis^a, Catuscia Palamidessi^a

^a *INRIA Futurs and LIX, École Polytechnique*

Abstract

In this paper we propose a formalization of probable innocence, a notion of probabilistic anonymity that is associated to “realistic” protocols such as Crowds. We analyze critically two different definitions of probable innocence from the literature. The first one, corresponding to the property that Reiter and Rubin have proved for Crowds, aims at limiting the probability of detection. The second one, by Halpern and O’Neill, aims at constraining the attacker’s confidence. Our proposal combines the spirit of both these definitions while generalizing them. In particular, our definition does not need symmetry assumptions, and it does not depend on the probabilities of the users to perform the action of interest. We show that, in case of a symmetric system, our definition corresponds exactly to the one of Reiter and Rubin. Furthermore, in the case of users with uniform probabilities, it amounts to a property similar to that of Halpern and O’Neill.

Another contribution of our paper is the study of probable innocence in the case of protocol composition, namely when multiple runs of the same protocol can be linked, as in the case of Crowds.

1 Introduction

Often we wish to ensure that the identity of the user performing a certain action is maintained secret. This property is called *anonymity*. Examples of situations in which we may wish to provide anonymity include: publishing on the web, retrieving information from the web, sending a message, etc. Many protocols have been designed for this purpose, for example, Crowds [15], Onion Routing [23], the Free Haven [7], Web MIX [1] and Freenet [4].

[★] This work has been partially supported by the Project Rossignol of the ACI Sécurité Informatique (Ministère de la recherche et nouvelles technologies) and by the INRIA/ARC project ProNoBiS.

Email addresses: kostas@lix.polytechnique.fr (Konstantinos Chatzikokolakis), catuscia@lix.polytechnique.fr (Catuscia Palamidessi).

Most of the protocols providing anonymity use random mechanisms. Consequently, it is natural to think of anonymity in probabilistic terms. Various notions of probabilistic anonymity have been proposed in literature, at different levels of strength. The notion of anonymity in [3], called conditional anonymity in [9,10], and investigated also in [2], describes the ideal situation in which the protocol does not leak any information concerning the identity of the user. This property is satisfied for instance by the Dining Cryptographers with fair coins [3]. Protocols used in practice, however, especially in presence of attackers or corrupted users, are only able to provide a weaker notion of anonymity.

In [15] Reiter and Rubin have proposed a hierarchy of notions of probabilistic anonymity in the context of Crowds. We recall that Crowds is a system for anonymous web surfing aimed at protecting the identity of the users when sending (originating) messages. This is achieved by forwarding the message to another user selected randomly, which in turn forwards the message, and so on, until the message reaches its destination. Part of the users may be corrupted (attackers), and one of the main purposes of the protocol is to protect the identity of the originator of the message from those attackers.

Quoting from [15], the hierarchy is described as follows. Here the *sender* stands for the user that forwards the message to the attacker.

Beyond suspicion From the attacker's point of view, the sender appears no more likely to be the originator of the message than any other potential sender in the system.

Probable innocence From the attacker's point of view, the sender appears no more likely to be the originator of the message than to not be the originator.

Possible innocence From the attacker's point of view, there is a nontrivial probability that the real sender is someone else.

In [15] the authors also considered a formal definition of probable innocence tailored to the characteristics of the Crowds system, and proved it to hold for Crowds under certain conditions. Later Halpern and O'Neill proposed in [10] a formal interpretation of the notions of the hierarchy above in more general terms. Their definitions are based on the confidence of the attacker. More precisely their definition of probable innocence holds if for the attacker, given the events that he has observed, the probability that an user i has performed the action of interest is no more than $1/2$.

However, the property of probable innocence that Reiter and Rubin express formally and prove for the system Crowds in [15] does not mention the user's probability of being the originator, but only the probability of the event observed by the attacker. More precisely, the property proved for Crowds is that the probability that the originator forwards the message to an attacker (given that an attacker receives eventually the message) is at most $1/2$. In other words, their definition expresses a

limit on the probability of detection.

The property proved for Crowds in [15] depends only on the way the protocol works, and on the number of the attackers. It is totally independent from the probability of each user to be the originator. This is of course a very desirable property, since we do not want the correctness of a protocol to depend on the users' intentions of originating a message. For stronger notions of anonymity, this abstraction from the users' probabilities¹ leads to the notion of probabilistic anonymity defined in [2], which is equivalent to the conditional anonymity defined in [9,10]. Note that this definition is different from the notion of *strong probabilistic anonymity* given in [9,10]: the latter depends, again, on the probabilities of the users to perform the action of interest.

Another intended feature of our notion of probable innocence is the abstraction from the specific characteristics of Crowds. In Crowds, there are certain symmetries that derive from the assumption that the probability that user i forwards the message to user j is the same for all i and j . The property of probable innocence proved for Crowds in [15] depends strongly on this assumption. We want a general notion that has the possibility to hold even in protocols which do not satisfy the Crowds' symmetries.

For completeness, we also consider the composition of protocols executions, with specific focus on the case that in which the originator is the same and the protocol to be executed is the same. This situation can arise, for instance, when an attacker can induce the originator to repeat the protocol (multiple paths attack). We extend the definition of probable innocence to the case of protocol composition under the same originator, and we study how this property depends on the number of compositions.

All the notions developed in this paper are defined by using a model, for protocols and systems, based on a simplified version of Probabilistic Automata ([18]). Probabilistic Automata, and similar models like the Concurrent Markov Chains, are now a mature field of research with a solid theory and well established model checking tools like PRISM [13]. This opens the way to the automatic verification of our notion of probable innocence. We refer to [5] for various examples of verification, using PRISM, of the related notion of weak anonymity developed within the same framework of simplified Probabilistic Automata. Furthermore, we are currently developing a model checker for the probabilistic π -calculus [11,14]. This is a formalism whose semantics is again based on simplified Probabilistic Automata and it is a natural language for expressing protocols running on distributed systems like Crowds. We aim in particular at developing efficient model checking techniques for computing the conditional probability of events, which constitute the only kind of quantitative information needed for proving the formula expressing our notion of probable innocence.

¹ For simplicity sometime we will refer to the users' probability of performing the action of interest as "users' probabilities"

1.1 Contribution

The main goal of this paper is to establish a general notion of probable innocence which combines the spirits of the approaches discussed above, namely it expresses a limit both on the attacker's confidence and on the probability of detection. Furthermore, we aim at a notion that does not depend on symmetry assumptions and on the probabilities of the users to perform the action of interest.

We show that our definition, while being more general, corresponds exactly to the property that Reiter and Rubin have proved for Crowds, under the specific symmetry conditions which are satisfied by Crowds. We also show that in the particular case that the users have uniform probability of being the originator, we obtain a property similar to the definition of probable innocence given by Halpern and O'Neill.

A second contribution is the analysis of the robustness of probable innocence under multiple paths attacks, which induce a repetition of the protocol. We show a general negative result, namely that no protocol can ensure probable innocence under an arbitrary number of repetitions, unless the system is strongly anonymous. This generalizes the result, already known in literature, about the fact that Crowds cannot guarantee probable innocence under unbound multiple path attacks.

1.2 Plan of the paper

In next section we recall some notions which are used in the rest of the paper: the Probabilistic Automata, the framework for anonymity developed in [2], and the definition of (strong) probabilistic anonymity given in [2]. In Section 3 we illustrate the Crowds protocol, we recall the property proved for Crowds and the definition of probable innocence by Halpern and O'Neill, and we discuss them. In Section 4 we propose our notion of probable innocence and we compare it with those of Section 3. In Section 5 we consider the repetition of an anonymity protocol and we show that we cannot guarantee probable innocence for arbitrary repetition unless the protocol is strongly anonymous. In Section 6 we discuss some related work from the literature. Section 7 concludes.




Fig. 1. Examples of probabilistic automata

2 Preliminaries

2.1 Probabilistic Automata

In our approach we consider systems that can perform both probabilistic and nondeterministic choice. Intuitively, a probabilistic choice represents a set of alternative transitions, each of them associated to a certain probability of being selected. The sum of all probabilities on the alternatives of the choice must be 1, i.e. they form a *probability distribution*. Nondeterministic choice is also a set of alternatives, but we have no information on how likely one alternative is selected.

There have been many models proposed in literature that combine both nondeterministic and probabilistic choice. One of the most general is the formalism of *probabilistic automata* proposed in [18]. In this work we use this formalism to model anonymity protocols. We give here a brief description of it.

A probabilistic automaton consists in a set of states, and labeled transitions between them. For each node, the outgoing transitions are partitioned in groups called *steps*. Each step represents a probabilistic choice, while the choice between the steps is nondeterministic.

Figure 1 illustrates some examples of probabilistic automata. We represent a step by putting an arc across the member transitions. For instance, in (a), state s_1 has two steps, the first is a probabilistic choice between two transitions with labels a and b , each with probability $1/2$. When there is only a transition in a step, like the one from state s_3 to state s_6 , the probability is of course 1 and we omit it.

In this paper, we use only a simplified kind of automaton, in which from each node we have either a probabilistic choice or a nondeterministic choice (more precisely, either one step or a set of singleton steps), like in (b). In the particular case that the choices are all probabilistic, like in (c), the automaton is called *fully probabilistic*.

Given an automaton M , we denote by $etree(M)$ its unfolding, i.e. the tree of all possible executions of M (in Figure 1 the automata coincide with their unfolding because there is no loop). If M is fully probabilistic, then each execution (maximal branch) of $etree(M)$ has a probability obtained as the product of the probability of the edges along the branch. In the finite case, we can define a probability measure for each set of executions, called *event*, by summing up the probabilities of the elements². Given an event x , we will denote by $p(x)$ the probability of x . For instance, let the event c be the set of all computations in which c occurs. In (c) its probability is $p(c) = 1/3 \times 1/2 + 1/6 = 1/3$.

When nondeterminism is present, the probability can vary, depending on how we *resolve* the nondeterminism. In other words we need to consider a function ς that, each time there is a choice between different steps, selects one of them. By pruning the non-selected steps, we obtain a fully probabilistic execution tree $etree(M, \varsigma)$ on which we can define the probability as before. For historical reasons (i.e. since nondeterminism typically arises from the parallel operator), the function ς is called *scheduler*.

It should then be clear that the probability of an event is relative to the particular scheduler. We will denote by $p_\varsigma(x)$ the probability of the event x under the scheduler ς . For example, consider (a). We have two possible schedulers determined by the choice of the step in s_1 . Under one scheduler, the probability of c is $1/2$. Under the other, it is $2/3 \times 1/2 + 1/3 = 2/3$. In (b) we have three possible schedulers under which the probability of c is 0, $1/2$ and 1, respectively.

2.2 Anonymity systems

The concept of anonymity is relative to the set of anonymous users and to what is visible to the observer. Hence, following [17,16] we classify the actions of the automaton into the three sets A , B and C as follows:

- A is the set of the anonymous actions $A = \{a(i) \mid i \in I\}$ where I is the set of the identities of the anonymous users and a is an injective function from I to the set of actions, which we call *abstract action*. We also call the pair (I, a) *anonymous action generator*.
- B is the set of the observable actions. We will use b, b', \dots to denote the elements of this set.
- C is the set of the remaining actions (which are unobservable).

² In the infinite case things are more complicated: we cannot define a probability measure for all sets of execution, and we need to consider as event space the σ -field generated by the *cones* of $etree(M)$. However, in this paper, we consider only the finite case.

Note that the actions in A normally are not visible to the observer, or at least, not for the part that depends on the identity i . However, for the purpose of defining and verifying anonymity we model the elements of A as visible outcomes of the system.

Definition 1 *An anonymity system is a tuple (M, I, a, B, Z, p) , where M is a probabilistic automaton, (I, a) is an anonymous action generator, B is a set of observable actions, Z is the set of all possible schedulers for M , and for every $\varsigma \in Z$, p_ς is the probability measure on the event space generated by $etree(M, \varsigma)$.*

For simplicity, we assume the users to be the only possible source of nondeterminism in the system. If they are probabilistic, then the system is fully probabilistic, hence Z is a singleton and we omit it.

We introduce the following notation to represent the events of interest:

- $a(i)$: all the executions in $etree(M, \varsigma)$ containing the action $a(i)$;
- a : all the executions in $etree(M, \varsigma)$ containing an action $a(i)$ for an arbitrary i ;
- o : all the executions in $etree(M, \varsigma)$ containing the sequence of observable actions o (where o is of the form $b_1 b_2 \dots b_n$ for some $b_1, b_2, \dots, b_n \in B$). We denote by O (*observables*) the set of all o 's of interest.

We use the symbols \cup , \cap and \neg to represent the union, the intersection, and the complement of events, respectively.

We wish to keep the notion of observables as general as possible, but we still need to make some assumptions on them. First, we want the observables to be execution-disjoint events, in the sense that no execution can contain both o_1 and o_2 if $o_1 \neq o_2$. Second, they must cover all possible outcomes. Third, an observable o must indicate unambiguously whether a has taken place or not, i.e. it either implies a , or it implies $\neg a$. In set-theoretic terms it means that either o is a subset of a or of the complement of a . Formally³:

Assumption 1 (on the observables)

- (1) $\forall \varsigma \in Z. \forall o_1, o_2 \in O. o_1 \neq o_2 \Rightarrow p_\varsigma(o_1 \cup o_2) = p_\varsigma(o_1) + p_\varsigma(o_2)$
- (2) $\forall \varsigma \in Z. p_\varsigma(O) = 1$
- (3) $\forall \varsigma \in Z. \forall o \in O. (p_\varsigma(o \cap a) = p_\varsigma(o)) \vee p_\varsigma(o \cap \neg a) = p_\varsigma(o)$

Analogously, we need to make some assumption on the anonymous actions. We consider first the conditions tailored for the nondeterministic users: each scheduler

³ Note that the intuitive explanations here are stronger than the corresponding formal assumptions because, in the infinite case, there could be non-trivial sets of measure 0. However in the case of anonymity we usually deal with finite scenarios. In any case, these formal assumptions are enough for the ensuring the properties of the anonymity notions that we need in this paper.

determines completely whether an action of the form $a(i)$ takes place or not, and in the positive case, there is only one such i . Formally:

Assumption 2 (on the anonymous actions, for nondeterministic users)

$$\forall \varsigma \in Z. p_{\varsigma}(a) = 0 \vee (\exists i \in I. (p_{\varsigma}(a(i)) = 1 \wedge \forall j \in I. j \neq i \Rightarrow p_{\varsigma}(a(j)) = 0))$$

We now consider the case in which the users are fully probabilistic. The assumption on the anonymous actions in this case is much weaker: we only require that there be at most one user that performs a , i.e. $a(i)$ and $a(j)$ must be disjoint for $i \neq j$. Formally:

Assumption 3 (on the anonymous actions, for probabilistic users)

$$\forall i, j \in I. i \neq j \Rightarrow p(a(i) \cup a(j)) = p(a(i)) + p(a(j))$$

2.3 Strong probabilistic anonymity

In this section we recall the notion of strong anonymity proposed in [2].

Let us first assume that the users are nondeterministic. Intuitively, a system is strongly anonymous if, given two schedulers ς and ϑ that both choose a (say $a(i)$ and $a(j)$, respectively), it is not possible to detect from the probabilistic measure of the observables whether the scheduler has been ς or ϑ (i.e. whether the selected user was i or j).

Note that ς chooses a if and only if $p_{\varsigma}(a) = 1$ or, equivalently, if and only if $p_{\varsigma}(a(i)) = 1$ for some i .

Definition 2 A system (M, I, a, B, Z, p) with nondeterministic users is anonymous if

$$\forall \varsigma, \vartheta \in Z. \forall o \in O. p_{\varsigma}(a) = p_{\vartheta}(a) = 1 \Rightarrow p_{\varsigma}(o) = p_{\vartheta}(o)$$

The probabilistic counterpart of Definition 2 can be formalized using the concept of *conditional probability*. Recall that, given two events x and y with $p(y) > 0$, the conditional probability of x given y , denoted by $p(x | y)$, is equal to $p(x \cap y)/p(y)$.

Definition 3 A system (M, I, a, B, p) with probabilistic users is anonymous if

$$\forall i, j \in I. \forall o \in O. (p(a(i)) > 0 \wedge p(a(j)) > 0) \Rightarrow p(o | a(i)) = p(o | a(j))$$

The notions of anonymity illustrated so far focus on the probability of the observables. More precisely, it requires the probability of the observables to be indepen-

dent from the selected user. In [2] it was shown that Definition 3 is equivalent to the notion adopted implicitly in [3], and called *conditional anonymity* in [9]. As illustrated in the introduction, the idea of this notion is that a system is anonymous if the observations do not change the probability of the $a(i)$'s. In other words, we may know the probability of $a(i)$ by some means external to the system, but the system should not increase our knowledge about it.

Proposition 4 ([2]) *A system (M, I, a, B, p) with probabilistic users is anonymous iff*

$$\forall i \in I. \forall o \in O. p(o \cap a) > 0 \Rightarrow p(a(i) | o) = p(a(i) | a)$$

Note 1 *To be precise, the probabilistic counterpart of Definition 2 should be stronger than that given in Definition 3, in fact it should be independent from the probabilities of the users to perform the action of interest, like Definition 2 is. We could achieve this by assuming the system to be parametric with respect to the probability distribution of the users, and then require the formula to hold for every possible distribution. Proposition 4 should be modified accordingly.*

Note 2 *The large number of anonymity definitions often leads to confusion. In the rest of the paper we will refer to Definition 3 as (strong) probabilistic anonymity. By conditional anonymity we will refer to the condition in Proposition 4 which corresponds to the definition of Halpern and O'Neill ([9]). Finally by strong anonymity we will refer to the corresponding definition in [9] which can be expressed as:*

$$\forall i, j \in I. \forall o \in O : p(a(i) | o) = p(a(j) | o) \tag{1}$$

3 Probable Innocence

Strong and conditional anonymity are notions which are usually difficult to achieve in practice. For instance, in the case of protocols like Crowds, the originator needs to take some initiative, thus revealing himself to the attacker with greater probability than the rest of the users. As a result, more relaxed levels of anonymity, such as probable innocence, are provided by real protocols.

Probable innocence is verbally defined by Reiter and Rubin ([15]) as “the sender (the user who forwards the message to the attacker) appears no more likely to be the originator than not to be the originator”. Two different approaches to formalize this notion exist. The first focuses on the probability of the observables and constraints the probability of detecting a user. The second focuses on the probability of the users and constraints the attacker’s confidence that the detected user is the originator.

In this section we first present the Crowds protocol. Then we discuss the two exist-

ing definitions in literature, corresponding to the approaches above, and we argue that each of them has some shortcoming: the first does not seem satisfactory when the system is not symmetric. The second depends on the users (their probability to perform the action) while, intuitively, anonymity should be a property of the protocol only. In Section 4 we will present a new definition which combines the spirit of the existing ones, and that at the same time overcomes the above shortcomings.

3.1 *The Crowds protocol*

This protocol, presented in [15], allows Internet users to perform web transactions without revealing their identity. The idea is to randomly route the request through a crowd of users. Thus when the web server receives the request he does not know who is the originator since the user who sent the request to the server is simply forwarding it. The more interesting case, however, is when an attacker is a member of the crowd and participates in the protocol. In this case the originator is exposed with higher probability than any other user and strong anonymity cannot be achieved. However, it can be proved that Crowds provides probable innocence under certain conditions.

More specifically a crowd is a group of m users who participate in the protocol. Some of the users may be corrupted which means they can collaborate in order to reveal the identity of the originator. Let c be the number of such users and p_f a parameter of the protocol, explained below. When a user, called the *initiator* or *originator*, wants to request a web page he must create a *path* between him and the server. This is achieved by the following process:

- The initiator selects randomly a member of the crowd (possibly himself) and forwards the request to him. We will refer to this latter user as the *forwarder*.
- A forwarder, upon receiving a request, flips a biased coin. With probability $1 - p_f$ he delivers the request directly to the server. With probability p_f he selects randomly, with uniform probability, a new forwarder (possibly himself) and forwards the request to him. The new forwarder repeats the same procedure.

The response from the server follows the same route in the opposite direction to return to the initiator. It must be mentioned that all communication in the path is encrypted using a *path key*, mainly to defend against local eavesdroppers (see [15] for more details). In this paper we are interested in attacks performed by corrupted members of the crowd to reveal the initiator's identity. Each member is considered to have only access to the traffic routed through him, so he cannot intercept messages addressed to other members.

3.2 Definition of probable innocence

3.2.1 First approach (limit on the probability of detection):

Reiter and Rubin ([15]) give a definition which considers the probability of the originator being observed by a corrupted member, that is being directly before him in the path. Let I denote the event “the originator is observed by a corrupted member” and H_{1+} the event “at least one corrupted member appears in the path”. Then probable innocence can be defined as

$$p(I | H_{1+}) \leq 1/2 \quad (2)$$

In [15] it is proved that this property is satisfied by Crowds if $n \geq \frac{p_f}{p_f-1/2}(c+1)$.

For simplicity, we suppose that a corrupted user will not forward a request to other crowd members, so at most one user can be observed. This approach is also followed in [15,21,24] and the reason is that by forwarding the request the corrupted users cannot gain any new information since forwarders are chosen randomly.

We now express the above definition in the framework of this paper (Section 2.2). Since $I \Rightarrow H_{1+}$ we have $p(I | H_{1+}) = p(I)/p(H_{1+})$. If A_i denotes that “user i is the originator” and D_i is the event “the user i was observed by a corrupted member (appears in the path right before the corrupter member)” then $p(I) = \sum_i p(D_i \wedge A_i) = \sum_i p(D_i | A_i)p(A_i)$. Since $p(D_i | A_i)$ is the same for all i then the definition (2) can be written $\forall i : p(D_i | A_i)/P(H_{1+}) \leq 1/2$.

Let A be the set of all crowd members and $O = \{o_i | i \in A\}$ the set of observables. Essentially $a(i)$ denotes A_i and o_i denotes D_i . Note that D_i is an observable since it can be observed by a corrupted user (remember that corrupted users share their information). Also let $h = \bigvee_{i \in A} o_i$, meaning that some user was observed. The definition (2) can now be written:

$$\forall i \in A : p(o_i | a(i)) \leq \frac{1}{2}p(h) \quad (3)$$

This is indeed an intuitive definition for Crowds. However there are many questions raised by this approach. For example, we are only interested in the probability of one specific event, what about other events that might reveal the identity of the initiator? For example the event $\neg o_i$ will have probability greater than $p(h)/2$, is this important? Moreover, consider the case where the probability of o_i under a different initiator j is negligible. Then, if we observe o_i , isn't it more probable that user i sent the message, even if $p(o_i | a(i))$ is less than $p(h)/2$?

If we consider arbitrary protocols, then there are cases where the condition (3) does not express the expected properties of probable innocence. We give two examples of such systems in Figure 2 and we explain them below.




Fig. 2. Examples of arbitrary (non symmetric) protocols. The value at position i, j represents $p(o_j | a(i))$ for user i and observable o_j .

Example 5 On the left-hand side of Figure 2, m users are participating in a Crowds-like protocol. The only difference, with respect to the standard Crowds, is that user 1 is behind a firewall, which means that he can send messages to any other user but he cannot receive messages from any of them. In the corresponding table we give the conditional probabilities $p(o_j | a(i))$, where we recall that o_j means that j is the user who sends the message to the corrupted member, and $a(i)$ means that i is the initiator. When user 1 is the initiator the probability of observing him is $\frac{c}{m-p_f}$ (there is a c/m chance that user 1 sends the message to a corrupted user and there is also a chance that he forwards it to himself and sends it to a corrupted user in the next round). All other users can be observed with the same probability l . When any other user is the initiator, however, the probability of observing user 1 is 0, since he will never receive the message. In fact, the protocol will behave exactly like a Crowd of $m - 1$ users as it is shown in the table.

Note that Reiter and Rubin's definition (3) requires the diagonal of this table to be less than $p(h)/2$. In this example the definition holds provided that $m - 1 \geq \frac{p_f}{p_f - 1/2}(c + 1)$. In fact, for all users $i \neq 1$, $p(o_i | a(i))$ is the same as in the original Crowds (which satisfies the definition) and for user 1 it is even smaller. However, if a corrupted member observes user 1 he can be sure that he is the initiator since no other initiator leads to the observation of user 1. The problem here is that Reiter and Rubin's definition constraints only the probability of detection of user 1 and says nothing about the attacker's confidence in case of detection. We believe that totally revealing the identity of the initiator with non-negligible probability is undesirable and should be considered as a violation of an anonymity notion such as probable innocence.

Example 6 *On the right-hand side we have an opposite counter-example. Three users want to communicate with a web server, but they can only access it through a proxy. We suppose that all users are honest but they do not trust the proxy so they do not want to reveal their identity to him. So they use the following protocol: the initiator first forwards the message to one of the users 1, 2 and 3 with probabilities $2/3, 1/6$ and $1/6$ respectively, regardless of which is the initiator. The user who receives the message forwards it to the proxy. The probabilities of observing each user are shown in the corresponding table. Regardless of which is the initiator, user 1 will be observed with probability $2/3$ and the others with probability $1/6$ each.*

In this example Reiter and Rubin’s definition does not hold since $p(o_1 | a(1)) > 1/2$. However all users produce the same observables with the same probabilities hence we cannot distinguish between them. Indeed the system is strongly anonymous (Definition 3 holds)! Thus, in the general case, we cannot adopt (3) as the definition of probable innocence since we want such a notion to be implied by strong anonymity.

However, it should be noted that in the case of Crowds the definition of Reiter and Rubin is correct, because of a special symmetry property of the protocol. This is discussed in detail in Section 4.1.

Finally, note that the above definition does not mention the probability of the users to be the originator. It only considers such events as conditions in the conditional probability of the event o_i given that i is the originator. The value of such conditional probability does not imply anything for the user, he might have a very small or very big probability of initiating the message. This is a major difference with respect to the next approach.

3.2.2 Second approach (limit on the attacker’s confidence):

Halpern and O’Neill propose in [9] a general framework for defining anonymity properties. We give a very abstract idea of this framework, detailed information is available in [9]. In this framework a system consists of a group of agents, each having a local state at each point of the execution. The local state contains all information that the user may have and does not need to be explicitly defined. At each point (r, m) user i can only have access to his local state $r_i(m)$. So he does not know the actual point (r, m) but at least he knows that it must be a point (r', m') such that $r'_i(m') = r_i(m)$. Let $K_i(r, m)$ be the set of all these points. If a formula ϕ is true in all points of $K_i(r, m)$ then we say that i knows ϕ . In the probabilistic setting it is possible to create a measure on $K_i(r, m)$ and draw conclusions of the form “formula ϕ is true with probability p ”.

To define probable innocence Halpern and O’Neill first define a formula $\theta(i, a)$ meaning “user i performed the event a ”. We then say that a system has probable innocence if for all points (r, m) , the probability of $\theta(i, a)$ in this point for all users j (that is, the probability that arises by measuring $K_j(r, m)$) is less than one half.

This definition can be expressed in the framework of Section 2.2. The probability of a formula ϕ for user j at the point (r, m) depends only on the set $K_j(r, m)$ which itself depends only on $r_j(m)$. The latter is the local state of the user, that is the only things that he can observe. In our framework this corresponds to the observables of the probabilistic automaton. Thus, we can reformulate the definition of Halpern and O’Neill as:

$$\forall i \in I, \forall o \in O : p(a(i) | o) \leq 1/2 \quad (4)$$

This definition is similar to the one of Reiter and Rubin but not the same. The difference is that it considers the probability that, given a certain observation, the user has performed the action of interest, not the opposite. If this probability is less than one half then intuitively i appear less likely to have performed o than not to.

The problem with this definition is that the probabilities of the users are not part of the system and we can make no assumptions about them. Consider for example the case where we know that user i visits very often a specific web site, so even if we have 100 users, the probability that he performed a request to this site is 0.99. Then we cannot expect this probability to become less than one half under all observations. A similar remark about strong anonymity led Halpern and O’Neill to define conditional anonymity. If a user i has higher probability of performing the action than user j then we cannot expect this to change because of the system. Instead we can request that the system does not provide any new information about the originator of the action.

4 A new definition of probable innocence

In this section we give a new definition of probable innocence that combines the spirit of the two existing ones. The spirit of Reiter and Rubin’s definition is to constraint the probability of detection of a user, which is captured in our Definition 8. The spirit of Halpern and O’Neill’s definition is to constrain the attacker’s confidence, which is captured in our Definition 7. The new definition combines both spirits in the sense that Definitions 7 and 8 are equivalent. Moreover it overcomes the shortcomings discussed in previous section, namely, it does not depend on the symmetry of the system and it does not depend on the users’ probabilities. We also show that our definition is a generalization of the existing ones since it can be reduced to them under the assumption of symmetry for the first, and of uniform users’ probability for the second.

One of the goals of the new definition is to abstract from the probabilities of the users to perform the action of interest. These probabilities, although they affect the probability measure p of the anonymity system, are not part of the protocol and can vary in different executions. To model this fact, let u be a probability measure on the set I of anonymous users. Then, we suppose that the anonymity system is

equipped with a probability measure p_u , which depends on u , satisfying the following conditions:

$$p_u(a(i)) = u(i) \tag{5}$$

$$p_u(o | a(i)) = p_{u'}(o | a(i)) \tag{6}$$

for all users i , observables o and user distributions u, u' such that $u(i) > 0, u'(i) > 0$. Condition (5) requires that the selection of user is made using the distribution u . Condition (6) requires that, having selected a user, the distribution u does not affect the probability of any observable o . In other words u is used to select a user and only for that. This is typical in anonymity protocols where a user is selected in the beginning (this models the user's decision to send a message) and then some observables are produced that depend on the selected user. We will denote by $p(o | a(i))$ the probability $p_u(o | a(i))$ under some u such that $u(i) > 0$.

In general we would like our anonymity definitions to range over all possible values of u since we cannot assume anything about the probabilities of the users to perform the action of interest. Thus, Halpern and O'Neill's definition (4) should be written: $\forall u \forall i \forall o : p_u(a(i) | o) \leq 1/2$ which makes even more clear the fact that it cannot hold for all u , for example if we take $u(i)$ to be very close to 1. On the other hand, Reiter and Rubin's definition contains only probabilities of the form $p(o | a(i))$. Crowds satisfies condition (6) so these probabilities are independent from u .

In [9], where they define conditional anonymity, Halpern and O'Neill make the following remark about strong anonymity. Since the probabilities of the users to perform the action of interest are generally unknown we cannot expect that all users appear with the same probability. All that we can ensure is that the system does not reveal any information, that is that the probability of every user before and after making an observation should be the same. In other words, the fraction between the probabilities of any couple of users should not be one, but should at least remain the same before and after the observation.

We apply the same idea to probable innocence. We start by rewriting relation (4) as

$$\forall i \in A, \forall o \in O : 1 \geq \frac{p_u(a(i) | o)}{p_u(\bigvee_{j \neq i} a(j) | o)} \tag{7}$$

As we already explained, if $u(i)$ is very high then we cannot expect this fraction to be less than 1. Instead, we could require that it does not surpass the corresponding fraction of the probabilities before the execution of the protocol. So we generalize condition (7) in the following definition.

Definition 7 *A system (M, I, a, B, p_u) has probable innocence if for all user dis-*

tributions u , users $i \in I$ and observables $o \in O$, the following holds:

$$(n - 1) \frac{p_u(a(i))}{p_u(\bigvee_{j \neq i} a(j))} \geq \frac{p_u(a(i) | o)}{p_u(\bigvee_{j \neq i} a(j) | o)}$$

where $n = |I|$ is the number of anonymous users.

In probable innocence we consider the probability of a user to perform the action of interest compared to the probability of all the other users together. Definition 7 requires that the fraction of these probabilities after the execution of the protocol should be no bigger than $n - 1$ times the same fraction before the execution. The $n - 1$ factor comes from the fact that in probable innocence *some* information about the sender's identity is leaked. For example, if users are uniformly distributed, each of them has probability $1/n$ before the protocol and the sender could appear with probability $1/2$ afterwards. In this case, the fraction between the sender and all other users is $\frac{1}{n-1}$ before the protocol and becomes 1 after. Definition 7 states that this fraction can be increased, thus leaking some information, but no more than $n - 1$ times.

Definition 7 generalizes relation (4) and can be applied in cases where the distribution of users is not uniform. However it still involves the probabilities of the users to perform the action of interest, which are not a part of the system. What we would like is a definition similar to Def. 3 which involves only probabilities of events that are part of the system. To achieve this we rewrite Definition 7 using the following transformations. For all users we assume that $u(i) > 0$. Users with zero probability to perform the action could be removed from Definition 7 before proceeding.

$$\begin{aligned} (n - 1) \frac{p_u(a(i))}{\sum_{j \neq i} p_u(a(j))} &\geq \frac{p_u(a(i) | o)}{\sum_{j \neq i} p_u(a(j) | o)} \Leftrightarrow \\ (n - 1) \frac{p_u(a(i))}{\sum_{j \neq i} p_u(a(j))} &\geq \frac{\frac{p_u(o | a(i)) p_u(a(i))}{p_u(o)}}{\sum_{j \neq i} \frac{p_u(o | a(j)) p_u(a(j))}{p_u(o)}} \Leftrightarrow \\ (n - 1) \sum_{j \neq i} p_u(o | a(j)) p_u(a(j)) &\geq p_u(o | a(i)) \sum_{j \neq i} p_u(a(j)) \end{aligned}$$

We obtain a lower bound of the left clause by replacing all $p_u(o | a(j))$ with their minimum. So we require that

$$(n - 1) \min_{j \neq i} \{p_u(o | a(j))\} \sum_{j \neq i} p_u(a(j)) \geq p_u(o | a(i)) \sum_{j \neq i} p_u(a(j)) \Leftrightarrow \quad (8)$$

$$(n - 1) \min_{j \neq i} p_u(o | a(j)) \geq p_u(o | a(i)) \quad (9)$$

Condition (9) can be interpreted as follows: for each observable, the probability that user i performs the action should be balanced by the corresponding probabilities of

the other users. It would be more natural to have the sum of all $p_u(o | a(j))$ at the left side, in fact the left side of (9) is a lower bound of this sum. However, since the probabilities of the users are unknown, we have to consider the “worst” case where the user with the minimum $p_u(o | a(j))$ has the greatest probability of appearing.

Finally, condition (9) is equivalent to the following definition that we propose as a general definition of probable innocence.

Definition 8 *A system (M, I, a, B, p_u) has probable innocence if for all observables $o \in O$ and for all users $i, j \in I$:*⁴

$$(n - 1)p(o | a(j)) \geq p(o | a(i))$$

The meaning of this definition is that in order for $p_u(a(i))/p_u(\bigvee_{j \neq i} a(j))$ to increase at most by $n - 1$ times (Def. 7), the corresponding fraction between the probabilities of the observables must be at most $n - 1$. Note that in probabilistic anonymity (Def. 3) $p(o | a(i))$ and $p(o | a(j))$ are required to be equal. In probable innocence we allow $p(o | a(i))$ to be bigger, thus losing some anonymity, but no more than $n - 1$ times.

Definition 8 has the advantage of including only the probabilities of the observables and not those of the users, similarly to the Definition 3 of probabilistic anonymity. It is clear that Definition 8 implies Definition 7 since we strengthened the first to obtain the second. Since Definition 7 considers all possible distributions of the users, the inverse implication also holds.

Proposition 9 *Definitions 7 and 8 are equivalent.*

Proof Def. 8 \Rightarrow Def. 7 is trivial, since we strengthen the second to obtain the first. For the inverse suppose that Def. 7 holds but Def. 8 does not, so there exist users k, l and observable o such that $(n - 1)p_u(o | a(k)) < p_u(o | a(l))$. Thus there exist an $\epsilon > 0$ s.t.

$$(n - 1)(p_u(o | a(k)) + \epsilon) \leq p_u(o | a(l)) \tag{10}$$

Def. 7 should hold for all user distributions u so we select one which assigns a very small probability δ to all users except k, l . That is $u(i) = \frac{\delta}{n-2} \forall i \neq k, l$. From Def. 7 (for $i = k$) we have:

⁴ Remember that $p_u(o | a(i))$ is independent from u so we can take any distribution such that $u(i) > 0$, for example a uniform one.

$$\begin{aligned}
& (n-1)(p_u(a(k))p_u(o|a(k)) + \sum_{j \neq k,l} \delta p_u(o|a(j))) \geq p_u(o|a(l))(\delta + p_u(a(k)) \stackrel{p_u(o|a(j)) \leq 1}{\Rightarrow}) \quad (11) \\
& (n-1)(p_u(a(k))p_u(o|a(k)) + \delta) \geq p_u(o|a(l))(\delta + p_u(a(k)) \stackrel{(10)}{\Rightarrow}) \\
& \quad p_u(a(k))p_u(o|a(k)) + \delta \geq (p_u(o|a(k)) + \epsilon)(\delta + p_u(a(k)) \Rightarrow \\
& \quad \delta(1 - p_u(o|a(k)) - \epsilon) \geq \epsilon p_u(a(k)) \stackrel{(10)}{\Rightarrow} \\
& \quad \delta \geq \frac{\epsilon p_u(a(k))}{1 - \frac{p_u(o|a(l))}{n-1}} \quad (12)
\end{aligned}$$

If $n > 2$ then the right side of inequality 12 is strictly positive so it is sufficient to take a smaller δ and end up with a contradiction. If $n = 2$ then there are no other users except k, l and we can proceed similarly. \square

Example 10 Recall now the two examples of Figure 2. If we apply Definition 8 to the first one we see that it doesn't hold since $(n-1)p(o_1|a(2)) = 0 \not\geq \frac{c}{n-p_f} = p(o_1|a(1))$. This agree with our intuition of probable innocence being violated when user 1 is observed. In the second example the definition holds since $\forall i, j : p(o_i|a(i)) = p(o_j|a(j))$. Thus, we see that in these two examples our definition reflects correctly the notion of probable innocence.

4.1 Relation to other definitions

4.1.1 Definition by Reiter and Rubin

Reiter and Rubin's definition can be expressed by the condition (3). It considers the probabilities of the observables (not the users) and it requires that for any user which originates the message, a special observable, representing the detection of the user by a corrupted member, has probability less than $p(h)/2$. As we saw at the examples of Figure 2 what is important is not the actual probability of an observable when a specific user is the originator, but its relation with the corresponding probabilities when the other users are the originators.

However in Crowds there are some important symmetries. First of all the number of the observables is the same as the number of users. For each user i there is an observable o_i meaning that the user i is observed. When i is the initiator, o_i has clearly a higher probability than the other observables. However, since forwarders are randomly selected, the probability of o_j is the same for all $j \neq i$. The same holds for the observables. o_i is more likely to have been performed by i . However all other users $j \neq i$ have the same probability of producing it. These symmetries can be expressed as:

$$\forall i \in I, \forall k, l \neq i : p(o_k | a(i)) = p(o_l | a(i)) \quad (13)$$

$$p(o_i | a(k)) = p(o_i | a(l)) \quad (14)$$

Because of these symmetries, we cannot have a situation similar to the ones of Figure 2. On the left-hand side, for example, the probability $p(o_1 | a(2)) = 0$ should be the same as $p(o_3 | a(2))$. To keep the value 0 (which is the reason why probable innocence is not satisfied) we should have 0 everywhere in the row (except $p(o_2 | a(2))$) which is impossible since the sum of the row should be $p(h)$ and $p(o_2 | a(2)) \leq p(h)/2$.

So the reason why probable innocence is satisfied in Crowds is not the fact that observing the initiator has low probability (what definition (2) ensures) by itself, but the fact that definition (2), because of the symmetry, forces the probability of observing any of the other users to be high enough.

Note that the number of anonymous users n is not the same as the number of users m in Crowds, in fact $n = m - c$ where c is the number of corrupted users.

Proposition 11 *Under the symmetry requirements (13) and (14), Definition 8 is equivalent to the one of Reiter and Rubin.*

Proof Due to the symmetry it is easy to see that there are only two possible values for $p(o_i | a(j))$. Namely when i is the sender, the probability to observe i is the same for all i . Similarly the probability of observing a different user $j \neq i$ is the same for all j . So

$$p(o_i | a(j)) = \begin{cases} \phi & \text{if } i = j \\ \chi & \text{if } i \neq j \end{cases}$$

Note that $\phi + (n - 1)\chi = p(h)$. So Def. 8 for o_i becomes

$$\begin{aligned} p(o_i | a(i)) \leq (n - 1)p(o_i | a(j)) &\Rightarrow \\ \phi \leq (n - 1)\chi &\Rightarrow \\ \phi \leq p(h) - \phi &\Rightarrow \\ p(o_i | a(i)) \leq \frac{1}{2}p(h) \end{aligned}$$

which corresponds to Reiter and Rubin's definition. □.

4.1.2 Definition of Halpern and O'Neill

One of the motivations behind the new definition of probable innocence is that it should make no assumptions about the probabilities of the users. If we assume a uniform distribution of users then it can be shown that our definition becomes the same as the one of Halpern and O'Neill.




Fig. 3. Relation between the various anonymity definitions

Proposition 12 *The definition of Halpern and O’Neill can be obtained by Definition 7 if we consider a uniform distribution of users, that is a distribution u such that $\forall i, j \in I : u(i) = u(j) = 1/n$.*

Proof Trivial. Since all users have the same probability then $\forall i \in I : p(a(i)) = 1/n$ and the left side of definition 7 is equal to 1. \square

Note that the equivalence of Def. 7 and Def. 8 is based on the fact that the former ranges over all possible distributions u . Thus Def. 8 is strictly stronger than the one of Halpern and O’Neill.

4.1.3 Probabilistic anonymity

It is easy to see that strong anonymity (equation (1)) implies Halpern and O’Neill’s definition of probable innocence. Definition 8 preserves the same implication in the case of probabilistic anonymity.

Proposition 13 *Probabilistic anonymity implies probable innocence (Definition 8).*

Proof Trivial. If Definition 3 holds then $p(o | a(j)) = p(o | a(i)) \forall o, i, j$. \square

The relation between the various definitions of anonymity is summarized in Figure 3. The classification in columns is based on the type of probabilities that are considered. The first column considers the probability of different users, the second the probability of the same user before and after an observation and the third the probability of the observables. Concerning the lines, the first corresponds to the

strong case and the second to probable innocence. It is clear from the table that the new definition is to probable innocence as conditional anonymity is to strong anonymity.

5 Protocol Composition

In protocol analysis, it is often easier to split complex protocols in parts, analyze each part separately and then combine the results. In this section we will consider the case where a protocol is “repeated” multiple times but with only one user-selection phase in the beginning. This situation arises when an attacker can force a user to repeat the protocol many times. We will examine the anonymity guarantees of the resulting protocol with respect to the existing one, obtaining a general result for a class of attacks that appear in protocols such as Crowds.

First, we define the “sequential composition” of two anonymity systems.

Definition 14 *Let $A_1 = (M_1, I, a_1, B_1, p_1)$, $A_2 = (M_2, I, a_2, B_2, p_2)$ be two anonymity systems with the same set of anonymous users I . The sequential composition of A_1 and A_2 , denoted as $A_1; A_2$ is an anonymity system (M, I, a, B, p) such that:*

$$\text{exec}(M) \subseteq \text{exec}(M_1) \times \text{exec}(M_2) \quad (15)$$

$$a_1^{-1}(\xi_1) = a_2^{-1}(\xi_2) \quad \forall \xi_1 \xi_2 \in \text{exec}(M) \quad (16)$$

$$p(o_1 o_2 | a(i)) = p_1(o_1 | a(i)) \cdot p_2(o_2 | a(i)) \quad \forall o_1 o_2 \in O_1 \times O_2 \quad (17)$$

where $\text{exec}(M)$ is the set of all executions in $\text{etree}(M)$, a_i^{-1} is the inverse function of a_i and O_i is the set of observables of A_i .

Intuitively, $A_1; A_2$ emulates A_1 in the beginning. When A_1 terminates then it emulates A_2 but without re-selecting a user, keeping the same user that was selected in A_1 . So the executions of $A_1; A_2$ are of the form $\xi_1 \xi_2$, where ξ_i is an execution of A_i , with the constraint that ξ_1, ξ_2 should correspond to the same user. Since the user is selected once, the probability of the event $o_1 o_2$ given a user i is the product of the corresponding probabilities of each system. We are not interested in the exact structure of the automaton M , however it should be relatively simple to construct it from M_1 and M_2 .

Repetition is a special case of sequential composition when the two systems are the same.

Definition 15 *Let A be an anonymity system. We define the m -repetition of A as $A^m = A; \dots; A$, m times.*

Let A be an anonymity system and O its set of observables. We will examine the

anonymity guarantees of A^m with respect to the ones of A . From Definition 3 and equation (17) it is easy to conclude that A^m is strongly anonymous if and only if A is strongly anonymous too, which is expected since the probability of each single event is the same under any user. However, the case of probable innocence is more interesting since an event might have greater probability under user i than under user j .

Consider a system with three users, and one event o with probabilities $p(o | a(1)) = 1/2$ and $p(o | a(2)) = p(o | a(3)) = 1/4$. This system satisfies Definition 8 thus it provides probable innocence. If we repeat the protocol two times then the probabilities for the event oo will be $p(oo | a(1)) = 1/4$ and $p(oo | a(2)) = p(oo | a(3)) = 1/16$, but now Definition 8 is violated. In the original protocol the probability of o under user 1 was two times bigger than the corresponding probability of the other users, but after the repetition it became 4 times bigger and Definition 8 does not allow it.

In the general case, the system A^m satisfies (by definition) probable innocence if

$$(n - 1)p(o_1 \dots o_m | a(i)) \geq p(o_1 \dots o_m | a(j)) \quad \forall o_1, \dots, o_m \in O, \forall i, j \in I \quad (18)$$

The following lemma states that it is sufficient to check only the events of the form $o \dots o$ (the same event repeated m times), and expresses the probable innocence of A^m using probabilities of A .

Lemma 16 *Let $A = (M, I, a, B, p)$ be an anonymity system, $n = |I|$ and O its set of observable events. A^m satisfies probable innocence if and only if:*

$$(n - 1)p^m(o | a(i)) \geq p^m(o | a(j)) \quad \forall o \in O, \forall i, j \in I \quad (19)$$

Proof (only if) We can use equation (18) with $o_1 = \dots = o_m = o$ and then (17) to obtain (19). (if) We can write (19) as $\sqrt[n-1]{p(o | a(i))} \geq p(o | a(j))$. Let o_1, \dots, o_m be events, by applying this inequality to all of them we have:

$$\begin{aligned} \sqrt[n-1]{p(o_1 | a(i))} &\geq p(o_1 | a(j)) \\ &\vdots \\ \sqrt[n-1]{p(o_m | a(i))} &\geq p(o_m | a(j)) \end{aligned}$$

Then by multiplying these inequalities we obtain (18). □

Lemma 16 explains our previous example. The probability $p(o | a(2)) = 1/4$ was smaller than $p(o | a(1)) = 1/2$ but sufficient to provide probable innocence. But when we raised these probabilities to the power of two, $1/16$ was too small so the event oo would expose user 1. In fact, if we allow an arbitrary number of

repetitions equation (19) can never hold, unless the probability of all events under any user is the same, that is if the system is strongly anonymous.

Proposition 17 *Let A be an anonymity system. A^m satisfies probable innocence for all m if and only if A is strongly anonymous.*

Proof We rewrite equation (19) as ⁵ :

$$n - 1 \geq \left(\frac{p(o | a(j))}{p(o | a(i))} \right)^m \quad \forall o \in O, \forall i, j \in I \quad (20)$$

If A is strongly anonymous then by Definition 3: $p(o | a(i)) = p(o | a(j))$ for all o, i, j so the right side of inequality 20 is 1 thus it always holds (for $n \geq 2$). Otherwise there exist o, i, j such that $p(o | a(j)) > p(o | a(i))$. So (20) cannot hold for all m since $\alpha^m \rightarrow \infty$ when $m \rightarrow \infty$ for $\alpha > 1$. \square

5.1 Multiple paths attack

As stated in the original paper of Crowds, after creating a random path to a server, a user should use the same path for all the future requests to the same server. However there is a chance that some node in the path leaves the network, in that case the user has to create a new path using the same procedure. In theory the two paths cannot be linked together, that is the attacker cannot know that it is the same user who created the two paths. In practice, however, such a link could be achieved by means unrelated to the protocol such as the url of the server, the data of the request etc. By linking the two requests the attacker obtains more observables that he can use to track down the originator. Since the attacker also participates in the protocol he could voluntarily break existing paths that pass through him in order to force the users to recreate them.

If C is an anonymity system that models Crowds, then the m -paths version corresponds to the m -repetition of C , which repeats the protocol m times without re-selecting a user. From proposition 17 and since Crowds is not strongly anonymous, we have that probable innocence cannot be satisfied if we allow an arbitrary number of paths. Intuitively this is justified. Even if the attacker sees the event o_1 meaning that user 1 was detected (was right before a corrupted user in the path) it could be the case (with non-trivial probability) that user 2 was the real originator, he sent the message to user 1 and he sent it to the attacker. However, if there are ten paths and the attacker sees $o_1 \dots o_1$ (ten times) then it is much more improbable

⁵ Note that in order to have probable innocence (or strong anonymity) $p(o | a(i))$ should be non-zero for all o and i except from trivial systems where all observables have zero probabilities. Thus, we consider only non-zero values for $p(o | a(i))$.

that all of the ten times user 2 sent the message to user 1 and user 1 to the attacker. It appears much more likely that user 1 was indeed the originator.

This attack had been foreseen in the original paper of Crowds and further analysis was presented in [24,20]. However our result is more general since we prove that probable innocence is impossible for any protocol that allows “multiple paths”, in other words that can be modeled as an m -repetition, unless the original protocol is strongly anonymous. Also our analysis is simpler since we did not need to calculate the actual probabilities of any observables in a specific protocol.

6 Related Work

Anonymity and privacy have been an area of research for over two decades now, with an increasing interest on the subject during the last five years, resulting in a great number of publications. The most related work to ours, as we already discussed in the introduction and section 3, is the one of Reiter and Rubin ([15]) and the one of Halpern and O’Neill ([10]).

Apart from the above two, there are many papers in the anonymity bibliography in which formal definitions of various notions of anonymity are given. Schneider and Sidiropoulos ([17]) propose a definition of anonymity based on CSP. Hughes and Shmatikov ([12]) developed a modular framework to formalize a range of properties (including numerous flavors of anonymity and privacy) using the notion of *function views* to represent a mathematical abstraction of partial knowledge of a function. Syverson and Stubblebine ([22]) introduce the notion of *group principals* and an associated epistemic logic to axiomatize anonymity. In these papers, possibilistic frameworks are used and it is not clear how the definitions could be extended in a probabilistic setting.

On the other hand, Bhargava and Palamidessi ([2]) propose a probabilistic definition of strong anonymity using the same framework as this paper. The resulting definition can be seen as the strong variant of Definition 8 (in fact, it implies Definition 8 as shown in section 4.1.3). Serjantov and Danezis ([19]) and Diaz et al ([6]) take an information theoretical approach by considering the *entropy* of the probability distribution that the attacker assigns to the anonymous agents after observing the system.

Finally, we should mention an interesting work by Evfimievski et al ([8]) on the field of privacy preserving data mining. Their definition requires that the probability of a private value x_1 producing an output y should be at most γ times the corresponding probability of a different value x_2 . This is very close in spirit to our definition of probable innocence.

7 Conclusion

In this paper we have considered probable innocence, a weak notion of anonymity provided by real-world systems such as Crowds. We have analyzed the definitions of probable innocence existing in literature, in particular: the one by Reiter and Rubin which is suitable for systems which, like Crowds, satisfy certain symmetries, and the one given by Halpern and O’Neill, which expresses a condition on the probability of the users.

Our first contribution is a definition of probable innocence which is (intuitively) adequate for a general class of protocols, abstracts from the probabilities of the users and involves only the probabilities that depend solely on the system. We have shown that the new definition is equivalent to the existing ones under symmetry conditions (Reiter and Rubin) or uniform distribution of the users (Halpern and O’Neill).

A second contribution is the extension of the definition of probable innocence to the case of protocol repetition, which is induced by multiple paths attacks. We have shown a general negative result, namely that no protocol can ensure probable innocence under an arbitrary number of repetitions.

References

- [1] Oliver Berthold, Hannes Federrath, and Stefan Köpsell. Web mixes: A system for anonymous and unobservable internet access. In *Designing Privacy Enhancing Technologies, International Workshop on Design Issues in Anonymity and Unobservability*, volume 2009 of *Lecture Notes in Computer Science*, pages 115–129. Springer, 2000.
- [2] Mohit Bhargava and Catuscia Palamidessi. Probabilistic anonymity. In Martín Abadi and Luca de Alfaro, editors, *Proceedings of CONCUR*, volume 3653 of *Lecture Notes in Computer Science*, pages 171–185. Springer, 2005. <http://www.lix.polytechnique.fr/~catuscia/papers/Anonymity/concur.pdf>.
- [3] David Chaum. The dining cryptographers problem: Unconditional sender and recipient untraceability. *Journal of Cryptology*, 1:65–75, 1988.
- [4] Ian Clarke, Oskar Sandberg, Brandon Wiley, and Theodore W. Hong. Freenet: A distributed anonymous information storage and retrieval system. In *Designing Privacy Enhancing Technologies, International Workshop on Design Issues in Anonymity and Unobservability*, volume 2009 of *Lecture Notes in Computer Science*, pages 44–66. Springer, 2000.
- [5] Yuxin Deng, Catuscia Palamidessi, and Jun Pang. Weak probabilistic anonymity. In *Proceedings of the 3rd International Workshop on Security Issues in Concurrency*

- (SecCo),
Electronic Notes in Theoretical Computer Science. Elsevier Science Publishers, 2005.
http://www.lix.polytechnique.fr/~catuscia/papers/Anonymity/report_wa.pdf.
- [6] Claudia Díaz, Stefaan Seys, Joris Claessens, and Bart Preneel. Towards measuring anonymity. In Roger Dingledine and Paul F. Syverson, editors, *Proceedings of the workshop on Privacy Enhancing Technologies (PET) 2002*, volume 2482 of *Lecture Notes in Computer Science*, pages 54–68. Springer, 2002.
- [7] Roger Dingledine, Michael J. Freedman, and David Molnar. The free haven project: Distributed anonymous storage service. In *Designing Privacy Enhancing Technologies, International Workshop on Design Issues in Anonymity and Unobservability*, volume 2009 of *Lecture Notes in Computer Science*, pages 67–95. Springer, 2000.
- [8] Alexandre Evfimievski, Johannes Gehrke, and Ramakrishnan Srikant. Limiting privacy breaches in privacy preserving data mining. In ACM Press, editor, *Proceedings of the Twenty-Second ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems (PODS)*, pages 211–222. ACM Press, 2003.
- [9] Joseph Y. Halpern and Kevin R. O’Neill. Anonymity and information hiding in multiagent systems. In *Proc. of the 16th IEEE Computer Security Foundations Workshop*, pages 75–88, 2003.
- [10] Joseph Y. Halpern and Kevin R. O’Neill. Anonymity and information hiding in multiagent systems. *Journal of Computer Security*, 13(3):483–512, 2005.
- [11] Oltea Mihaela Herescu and Catuscia Palamidessi. Probabilistic asynchronous π -calculus. In Jerzy Tiuryn, editor, *Proceedings of FOSSACS 2000 (Part of ETAPS 2000)*, volume 1784 of *Lecture Notes in Computer Science*, pages 146–160. Springer, 2000. http://www.lix.polytechnique.fr/~catuscia/papers/Prob_asy_pi/fossacs.ps.
- [12] Dominic Hughes and Vitaly Shmatikov. Information hiding, anonymity and privacy: a modular approach. *Journal of Computer Security*, 12(1):3–36, 2004.
- [13] Marta Z. Kwiatkowska, Gethin Norman, and David Parker. PRISM 2.0: A tool for probabilistic model checking. In *Proceedings of the First International Conference on Quantitative Evaluation of Systems (QEST) 2004*, pages 322–323. IEEE Computer Society, 2004.
- [14] Catuscia Palamidessi and Oltea M. Herescu. A randomized encoding of the π -calculus with mixed choice. *Theoretical Computer Science*, 335(2-3):73–404, 2005. http://www.lix.polytechnique.fr/~catuscia/papers/prob_enc/report.pdf.
- [15] Michael K. Reiter and Aviel D. Rubin. Crowds: anonymity for Web transactions. *ACM Transactions on Information and System Security*, 1(1):66–92, 1998.
- [16] Peter Y. Ryan and Steve Schneider. *Modelling and Analysis of Security Protocols*. Addison-Wesley, 2001.
- [17] Steve Schneider and Abraham Sidiropoulos. CSP and anonymity. In *Proc. of the European Symposium on Research in Computer Security (ESORICS)*, volume 1146 of *Lecture Notes in Computer Science*, pages 198–218. Springer, 1996.

- [18] Roberto Segala and Nancy Lynch. Probabilistic simulations for probabilistic processes. *Nordic Journal of Computing*, 2(2):250–273, 1995. An extended abstract appeared in *Proceedings of CONCUR '94*, LNCS 836: 481-496.
- [19] Andrei Serjantov and George Danezis. Towards an information theoretic metric for anonymity. In Roger Dingleline and Paul F. Syverson, editors, *Proceedings of the workshop on Privacy Enhancing Technologies (PET) 2002*, volume 2482 of *Lecture Notes in Computer Science*, pages 41–53. Springer, 2002.
- [20] V. Shmatikov. Probabilistic model checking of an anonymity system. *Journal of Computer Security*, 12(3/4):355–377, 2004.
- [21] Vitaly Shmatikov. Probabilistic analysis of anonymity. In *15th IEEE Computer Security Foundations Workshop (CSFW)*, pages 119–128, 2002.
- [22] Paul F. Syverson and Stuart G. Stubblebine. Group principals and the formalization of anonymity. In *World Congress on Formal Methods (1)*, pages 814–833, 1999.
- [23] P.F. Syverson, D.M. Goldschlag, and M.G. Reed. Anonymous connections and onion routing. In *IEEE Symposium on Security and Privacy*, pages 44–54, Oakland, California, 1997.
- [24] M. Wright, M. Adler, B. Levine, and C. Shields. An analysis of the degradation of anonymous protocols. In *ISOC Network and Distributed System Security Symposium (NDSS)*, 2002.