Probabilistic Model Checking of an Anonymity System

Vitaly Shmatikov *

SRI International
333 Ravenswood Avenue
Menlo Park, CA 94025 U.S.A.

shmat@csl.sri.com

Abstract

We use the probabilistic model checker PRISM to analyze the Crowds system for anonymous Web browsing. This case study demonstrates how probabilistic model checking techniques can be used to formally analyze security properties of a peer-to-peer group communication system based on random message routing among members. The behavior of group members and the adversary is modeled as a discrete-time Markov chain, and the desired security properties are expressed as PCTL formulas. The PRISM model checker is used to perform automated analysis of the system and verify anonymity guarantees it provides. Our main result is a demonstration of how certain forms of probabilistic anonymity degrade when group size increases or random routing paths are rebuilt, assuming that the corrupt group members are able to identify and/or correlate multiple routing paths originating from the same sender.

1 Introduction

Formal analysis of security protocols is a well-established field. Model checking and theorem proving techniques [Low96, MMS97, Pau98, CJM00] have been extensively used to analyze secrecy, authentication and other security properties of

* Supported in part by DARPA contract N66001-00-C-8015 "Agile Management of Dynamic Collaboration."

protocols and systems that employ cryptographic primitives such as public-key encryption, digital signatures, etc. Typically, the protocol is modeled at a highly abstract level and the underlying cryptographic primitives are treated as secure "black boxes" to simplify the model. This approach discovers attacks that would succeed even if all cryptographic functions were perfectly secure.

Conventional formal analysis of security is mainly concerned with security against the so called *Dolev-Yao attacks*, following [DY83]. A Dolev-Yao attacker is a non-deterministic process that has complete control over the communication network and can perform any combination of a given set of attacker operations, such as intercepting any message, splitting messages into parts, decrypting if it knows the correct decryption key, assembling fragments of messages into new messages and replaying them out of context, etc.

Many proposed systems for anonymous communication aim to provide strong, non-probabilistic anonymity guarantees. This includes proxy-based approaches to anonymity such as the Anonymizer [Ano], which hide the sender's identity for each message by forwarding all communication through a special server, and MIX-based anonymity systems [Cha81] that blend communication between different senders and recipients, thus preventing a global eavesdropper from linking sender-recipient pairs. Non-probabilistic anonymity systems are amenable to formal analysis in the same non-deterministic Dolev-Yao model as used for verification of secrecy and authentication protocols. Existing techniques for the formal analysis of anonymity in the *non-deterministic* model include traditional process formalisms such as CSP [SS96] and a special-purpose logic of knowledge [SS99].

In this paper, we use *probabilistic* model checking to analyze anonymity properties of a gossip-based system. Such systems fundamentally rely on probabilistic message routing to guarantee anonymity. The main representative of this class of anonymity systems is Crowds [RR98]. Instead of protecting the user's identity against a global eavesdropper, Crowds provides protection against collaborating local eavesdroppers. All communication is routed randomly through a group of peers, so that even if some of the group members collaborate and share collected *local* information with the adversary, the latter is not likely to distinguish true senders of the observed messages from randomly selected forwarders.

Conventional formal analysis techniques that assume a non-deterministic attacker in full control of the communication channels are not applicable in this case. Security properties of gossip-based systems depend solely on the *probabilistic* behavior of protocol participants, and can be formally expressed only in terms of relative probabilities of certain observations by the adversary. The system must be modeled as a probabilistic process in order to capture its properties faithfully.

Using the analysis technique developed in this paper—namely, formalization of the system as a discrete-time Markov chain and probabilistic model checking of

this chain with PRISM—we uncovered two subtle properties of Crowds that cause degradation of the level of anonymity provided by the system to the users. First, if corrupt group members are able to detect that messages along different routing paths originate from the same (unknown) sender, the probability of identifying that sender increases as the number of observed paths grows (the number of paths must grow with time since paths are rebuilt when crowd membership changes). Second, the confidence of the corrupt members that they detected the correct sender increases with the size of the group. The first flaw was reported independently by Malkhi [Mal01] and Wright *et al.* [WALS02], while the second, to the best of our knowledge, was reported for the first time in the conference version of this paper [Shm02]. In contrast to the analysis by Wright *et al.* that relies on manual probability calculations, we discovered both potential vulnerabilities of Crowds by automated probabilistic model checking.

Previous research on probabilistic formal models for security focused on (i) probabilistic characterization of non-interference [Gra92, SG95, VS98], and (ii) process formalisms that aim to faithfully model probabilistic properties of cryptographic primitives [LMMS99, Can00]. This paper attempts to directly model and analyze security properties based on discrete probabilities, as opposed to asymptotic probabilities in the conventional cryptographic sense. Our analysis method is applicable to other probabilistic anonymity systems such as Freenet [CSWH01] and onion routing [SGR97]. Note that the potential vulnerabilities we discovered in the formal model of Crowds may not manifest themselves in the implementations of Crowds or other, similar systems that take measures to prevent corrupt routers from correlating multiple paths originating from the same sender.

2 Markov Chain Model Checking

We model the probabilistic behavior of a peer-to-peer communication system as a discrete-time Markov chain (DTMC), which is a standard approach in probabilistic verification [LS82, HS84, Var85, HJ94]. Formally, a *Markov chain* can be defined as consisting in a finite set of states S, the initial state s_0, the transition relation $T : S \times S \to [0, 1]$ such that $\forall s \in S \ \sum_{s' \in S} T(s, s') = 1$, and a labeling function from states to a finite set of propositions $L : S \to 2^{AP}$.

In our model, the states of the Markov chain will represent different stages of routing path construction. As usual, a state is defined by the values of all system variables. For each state s, the corresponding row of the transition matrix T_s defines the probability distributions which govern the behavior of group members once the system reaches that state.

2.1 Overview of PCTL

We use the temporal probabilistic logic PCTL [HJ94] to formally specify properties of the system to be checked. PCTL can express properties of the form "under any scheduling of processes, the probability that event E occurs is at least p."

First, define *state formulas* inductively as follows:

$$\Phi \ ::= \ true \mid false \mid a \mid \Phi \wedge \Phi \mid \Phi \vee \Phi \mid \neg\Phi \mid \mathcal{P}_{>p}[\Psi]$$

where atomic propositions a are predicates over state variables. State formulas of the form $\mathcal{P}_{>p}[\Psi]$ are explained below.

Define *path formulas* as follows:

$$\Psi \ ::= \ X \ \Phi \mid \Phi \, \mathcal{U}^{\leq k} \ \Phi \mid \Phi \, \mathcal{U} \ \Phi$$

Unlike state formulas, which are simply first-order propositions over a single state, path formulas represent properties of a chain of states (here *path* refers to a sequence of state space transitions rather than a routing path in the Crowds specification). In particular, $X\,\phi$ is true *iff* ϕ is true for every state in the chain; $\phi_1 \, \mathcal{U} \ \phi_2$ is true *iff* ϕ_1 is true for all states in the chain until ϕ_2 becomes true, and ϕ_2 is true for all subsequent states; $\phi_1 \, \mathcal{U}^{\leq k} \ \phi_2$ is true *iff* $\phi_1 \, \mathcal{U} \ \phi_2$ and there are no more than k states before ϕ_2 becomes true.

For any state s and path formula ψ, $\mathcal{P}_{>p}[\psi]$ is a state formula which is true *iff* state space paths starting from s satisfy path formula ψ with probability greater than p.

For the purposes of this paper, we will be interested in formulas of the form $\mathcal{P}_{>p}[true \ \mathcal{U} \ \phi]$, evaluated in the initial state s_0. Here ϕ specifies a system configuration of interest, typically representing a particular observation by the adversary that satisfies the definition of a successful attack on the protocol. Property $\mathcal{P}_{>p}[true \ \mathcal{U} \ \phi]$ is a liveness property: it holds in s_0 *iff* ϕ will eventually hold with greater than p probability. For instance, if `observe3` is a state variable representing the number of times one of the corrupt members received a message from the honest member no. 3, then $\mathcal{P}_{>0.5}[true \ \mathcal{U} \ \texttt{observe3} > 1]$ holds in s_0 *iff* the probability of corrupt members eventually observing member no. 3 twice or more is greater than 50%.

Expressing properties of the system in PCTL allows us to reason formally about the probability of corrupt group members collecting enough evidence to successfully attack anonymity. We use model checking techniques developed for verification of discrete-time Markov chains to compute this probability automatically.

2.2 PRISM model checker

The automated analyses described in this paper were performed using PRISM, a probabilistic model checker developed by Kwiatkowska *et al.* [KNP01]. The tool supports both discrete- and continuous-time Markov chains, and Markov decision processes. As described in section 4, we model probabilistic peer-to-peer communication systems such as Crowds simply as discrete-time Markov chains, and formalize their properties in PCTL.

The behavior of the system processes is specified using a simple module-based language inspired by Reactive Modules [AH96]. State variables are declared in the standard way. For example, the following declaration

```
deliver: bool init false;
```

declares a boolean state variable `deliver`, initialized to *false*, while the following declaration

```
const TotalRuns = 4;
...
observe1: [0..TotalRuns] init 0;
```

declares a constant `TotalRuns` equal to 4, and then an integer array of size 5, indexed from 0 to `TotalRuns`, with all elements initialized to 0.

State transition rules are specified using guarded commands of the form

```
[] <guard> -> <command>;
```

where `<guard>` is a predicate over system variables, and `<command>` is the transition executed by the system if the guard condition evaluates to *true*. Command often has the form $X_1' =$`<expression>`$_1 \wedge$ `...` $\wedge X_n' =$`<expression>`$_n$, which means that in the next state (*i.e.*, that obtained after the transition has been executed), state variable X_i is assigned the result of evaluating arithmetic expression `<expression>`$_i$

If the transition must be chosen probabilistically, the discrete probability distribution is specified as

```
[] <guard> -> <prob1>:<command1> +
              ... +
              <probN>:<commandN>;
```

Transition represented by `command`$_i$ is executed with probability `prob`$_i$, and Σ_i `prob`$_i = 1$. Security properties to be checked are stated as PCTL formulas (see section 2.1).

Given a formal system specification, PRISM constructs the Markov chain and determines the set of reachable states, using MTBDDs and BDDs, respectively. Model checking a PCTL formula reduces to a combination of reachability-based computation and solving a system of linear equations to determine the probability of satisfying the formula in each reachable state. The model checking algorithms employed by PRISM include [BdA95, BK98, Bai98]. More details about the implementation and operation of PRISM can be found at `http://www.cs.bham.ac.uk/~dxp/prism/` and in [KNP01].

Since PRISM only supports model checking of finite DTMC, in our case study of Crowds we only analyze anonymity properties of *finite* instances of the system. By changing parameters of the model, we demonstrate how anonymity properties evolve with changes in the system configuration. Wright *et al.* [WALS02] investigated related properties of the Crowds system in the general case, but they do not rely on tool support and their analyses are manual rather than automated.

3 Crowds Anonymity System

Providing an anonymous communication service on the Internet is a challenging task. While conventional security mechanisms such as encryption can be used to protect the content of messages and transactions, eavesdroppers can still observe the IP addresses of communicating computers, timing and frequency of communication, etc. A Web server can trace the source of the incoming connection, further compromising anonymity. The Crowds system was developed by Reiter and Rubin [RR98] for protecting users' anonymity on the Web.

The main idea behind gossip-based approaches to anonymity such as Crowds is to hide each user's communications by routing them randomly within a crowd of similar users. Even if an eavesdropper observes a message being sent by a particular user, it can never be sure whether the user is the actual sender, or is simply routing another user's message.

3.1 Path setup protocol

A *crowd* is a collection of users, each of whom is running a special process called a *jondo* which acts as the user's proxy. Some of the jondos may be corrupt and/or controlled by the adversary. Corrupt jondos may collaborate and share their observations in an attempt to compromise the honest users' anonymity. Note, however, that all observations by corrupt group members are *local*. Each corrupt member may observe messages sent to it, but not messages transmitted on the links between honest jondos. An honest crowd member has no way of determining whether

a particular jondo is honest or corrupt. The parameters of the system are the total number of members N, the number of corrupt members C, and the *forwarding probability* p_f which is explained below.

To participate in communication, all jondos must register with a special server which maintains membership information. Therefore, every member of the crowd knows identities of all other members. As part of the join procedure, the members establish pairwise encryption keys which are used to encrypt pairwise communication, so the contents of the messages are secret from an external eavesdropper.

Anonymity guarantees provided by Crowds are based on the path setup protocol, which is described in the rest of this section. The path setup protocol is executed each time one of the crowd members wants to establish an anonymous connection to a Web server. Once a routing path through the crowd is established, all subsequent communication between the member and the Web server is routed along it. We will call one run of the path setup protocol a *session*. When crowd membership changes, the existing paths must be scrapped and a new protocol session must be executed in order to create a new random routing path through the crowd to the destination. Therefore, we'll use terms *path reformulation* and *protocol session* interchangeably.

When a user wants to establish a connection with a Web server, its browser sends a request to the jondo running locally on her computer (we will call this jondo the *initiator*). Each request contains information about the intended destination. Since the objective of Crowds is to protect the *sender's* identity, it is not problematic that a corrupt router can learn the recipient's identity. The initiator starts the process of creating a random path to the destination as follows:

- The initiator selects a crowd member at random (possibly itself), and forwards the request to it, encrypted by the corresponding pairwise key. We'll call the selected member the *forwarder*.

- The forwarder flips a biased coin. With probability $1 - p_f$, it delivers the request directly to the destination. With probability p_f, it selects a crowd member at random (possibly itself) as the next forwarder in the path, and forwards the request to it, re-encrypted with the appropriate pairwise key. The next forwarder then repeats this step.

Each forwarder maintains an identifier for the created path. If the same jondo appears in different positions on the same path, identifiers are different to avoid infinite loops. Each subsequent message from the initiator to the destination is routed along this path, *i.e.*, the paths are *static*—once established, they are not altered often. This is necessary to hinder corrupt members from linking multiple

paths originating from the same initiator, and using this information to compromise the initiator's anonymity as described in section 3.2.3.

3.2 Anonymity properties of Crowds

The Crowds paper [RR98] describes several degrees of anonymity that may be provided by a communication system. Without using anonymizing techniques, none of the following properties are guaranteed on the Web since browser requests contain information about their source and destination in the clear.

Beyond suspicion Even if the adversary can see evidence of a sent message, the real sender appears to be no more likely to have originated it than any other potential sender in the system.

Probable innocence The real sender appears no more likely to be the originator of the message than to not be the originator, *i.e.*, the probability that the adversary observes the real sender as the source of the message is less than $\frac{1}{2}$.

Possible innocence It appears to the adversary that there is a nontrivial probability that the message was originated by someone other than the real sender.

Probable innocence can be interpreted as *plausible deniability*. A system that guarantees the probable innocence property for message senders does not necessarily hide the sender's identity from the adversary. It merely puts the $\frac{1}{2}$ upper bound on the probability of detection. If the sender is observed by the adversary, she can then plausibly argue that she has been routing someone else's messages.

The Crowds paper focuses on providing anonymity against local, possibly cooperating eavesdroppers, who can share their observations of communication in which they are involved as forwarders, but cannot observe communication involving only honest members. We also limit our analysis to this case.

3.2.1 Anonymity for a single route

It is proved in [RR98] that, for any given routing path, the path initiator in a crowd of n members with forwarding probability p_f has *probable innocence* against c collaborating crowd members if the following inequality holds:

$$n \geq \frac{p_f}{p_f - \frac{1}{2}}(c+1) \tag{1}$$

More formally, let H_{1+} be the event that at least one of the corrupt crowd members is selected for the path, and I be the event that the path initiator appears in

the path immediately before a corrupt crowd member (*i.e.*, the adversary observes the real sender as the source of the messages routed along the path). Condition 1 guarantees that $P(I|H_{1+}) \leq \frac{1}{2}$. Note that this does not preclude the adversary from observing the path initiator more often than any other crowd member, *i.e.*, *probable innocence* is a weaker anonymity property than *beyond suspicion*.

3.2.2 Linkability of multiple routes

To maintain an anonymous connection when crowd membership changes, each initiator must rebuild its routing path to the destination through the new crowd. As a result of random forwarder selection, it is possible that both the old and the new path include corrupt forwarders. In general, it will not be immediately obvious to the adversary who controls both corrupt forwarders that the two paths originate from the same member. Session-specific information contained in the message may, however, provide clues that help the adversary link the paths. For example, if the initiator visits the same set of websites and/or its browsing patterns persist from session to session, it is relatively easy for the adversary to guess that messages observed along two different paths originate from the same place. Linking is even easier in the case of anonymous Web browsing since browser requests may contain cookies or other persistent data, relating sessions by the same (anonymous) user. Except cautioning the users "from continuing to browse the content related to what she was browsing prior to [path reformulation], lest collaborators are attempting to link paths based on that content" [RR98], the Crowds system *per se* does not provide protection against path linkage. Therefore, we assume in our analysis that attacks based on multiple-path observations are feasible. Other gossip-based anonymity systems such as onion routing [SGR97] may provide stronger protection against path linkability (*e.g.*, by inserting decoy traffic), making path linking attacks less feasible.

3.2.3 Anonymity for multiple routes

To prevent corrupt crowd members from linking multiple paths and using this information to infer the initiator's identity, the Crowds paper [RR98] suggests that paths should be static. Crowd membership, however, must change over time: new members join and some of the existing members fail, invalidating all paths in which they were involved as forwarders. Even if joins are batched, all paths must be scrapped and new paths built periodically. We demonstrate in section 6.1 that anonymity guarantees provided by Crowds degrade significantly if the adversary links only a relatively small (3-6) number of paths originating from the same member.

Malkhi [Mal01] and Wright *et al.* [WALS02] have made a similar observation,

proving that, given multiple linked paths, the initiator appears more often as a suspect than a random crowd member. The automated analysis described in section 6.1 confirms and quantifies this result. (The technical results of [Shm02] on which this paper is based had been developed independently of [Mal01] and [WALS02], before the latter was published). In general, [Mal01] and [WALS02] conjecture that there can be no reliable anonymity method for peer-to-peer communication if in order to start a new communication session, the initiator must originate the first connection before any processing of the session commences. This implies that anonymity is impossible in a gossip-based system with corrupt routers in the absence of decoy traffic.

In section 6.3, we show that, for any given number of observed paths, the adversary's confidence in its observations increases with the size of the crowd. This result contradicts the intuitive notion that bigger crowds provide better anonymity guarantees. It was discovered by automated analysis.

4 Formal Model of Crowds

In this section, we describe our probabilistic formal model of the Crowds system. Since there is no non-determinism in the protocol specification (see section 3.1), the model is a simple discrete-time Markov chain as opposed to a Markov decision process. In addition to modeling the behavior of the honest crowd members, we also formalize the adversary. The protocol does not aim to provide anonymity against global eavesdroppers. Therefore, it is sufficient to model the adversary as a coalition of corrupt crowd members who only have access to local communication channels, *i.e.*, they can only make observations about a path if one of them is selected as a forwarder. By the same token, it is not necessary to model cryptographic functions, since corrupt members know the keys used to encrypt peer-to-peer links in which they are one of the endpoints, and have no access to links that involve only honest members.

The modeling technique presented in this section is applicable with minor modifications to any probabilistic routing system. In each state of routing path construction, the discrete probability distribution given by the protocol specification is used directly to define the probabilistic transition rule for choosing the next forwarder on the path, if any. If the protocol prescribes an upper bound on the length of the path (*e.g.*, Freenet [CSWH01]), the bound can be introduced as a system parameter as described in section 4.2.3, with the corresponding increase in the size of the state space but no conceptual problems. Probabilistic model checking can then be used to check the validity of PCTL formulas representing properties of the system.

In the general case, forwarder selection may be governed by non-deterministic

`runCount`	Number of paths constructed so far ($\leq$`TotalRuns`).
`good`	The selected forwarder is honest.
`bad`	The selected forwarder is corrupt.
`lastSeen`	Identity of the preceding forwarder on the path.
`observe`$_i$	Number of times corrupt members observed member i.
Auxiliary flags	
`launch`	Holds only in the initial state s_0.
`new`	Ready to construct another path.
`start`	Beginning of new path construction.
`run`	Continue path construction.
`deliver`	Terminate the path.
`recordLast`	Record the identity of the preceding forwarder.
`badObserve`	A corrupt member is recording its observations.

Table 1: State variables.

rules. Non-deterministic transitions would give rise to a Markov decision process. In the case of Crowds, however, forward selection is probabilistic rather than non-deterministic. Therefore, there is no need to model the system as a Markov decision process.

4.1 Overview of the model

We model crowd members' behavior only in the path setup protocol, ignoring all subsequent communication conducted along an established static path. Once a path is set up, every forwarder on the path receives messages from the same member and cannot gain any additional information about the true originator of the messages.

Since paths must be rebuilt on a regular basis, we introduce the number of path reformulations (*i.e.*, number of times the path construction protocol is executed) as a parameter of the model (`TotalRuns`) and allow the adversary to accumulate observations over time in order to try to infer the identity of the path initiator. This assumes that a corrupt crowd member is capable of determining whether two paths originate from the same initiator, without necessarily knowing that initiator's identity (see section 3.2.2).

Each state of our model represents a particular stage of routing path construction. In the multiple-path case, we distinguish different paths. A state is completely defined by the values of state variables listed in table 1.

4.2 Model of honest members

4.2.1 Initiation

Path construction is initiated as follows (syntax of PRISM is described in section 2.2):

```
[] launch ->
     runCount'=TotalRuns &
     new'=true & launch'=false;
[] new & (runCount>0) ->
     (runCount'=runCount-1) &
     new'=false & start'=true;
[] start ->
     lastSeen'=0 & deliver'=false &
     run'=true & start'=false;
```

4.2.2 Forwarder selection

The initiator (*i.e.*, the first crowd member on the path, the one whose identity must be protected) randomly chooses the first forwarder from among all N group members. We assume that all group members have an equal probability of being chosen, but the technique can support any discrete probability distribution for choosing forwarders.

Forwarder selection is a single step of the protocol, but we model it as two probabilistic state transitions. The first determines whether the selected forwarder is honest or corrupt, the second determines the forwarder's identity. The randomly selected forwarder is corrupt with probability `badC`$= \frac{C}{N}$, and honest with probability `goodC`$= 1-$`badC`, where N is the size of the crowd, and C is the number of corrupt members.

```
[] (!good & !bad & !deliver & run) ->
     goodC: good'=true & run'=false &
            recordLast'=true   +
      badC: bad'=true & run'=false &
            badObserve'=true;
```

4.2.3 Path construction

If the selected forwarder is honest, its identity is recorded in `lastSeen`. Recording the forwarder's identity models the fact that the source IP addresses of requests routed by honest forwarders can be observed by a corrupt member if it happens to

be next on the path. Any of the $N - C$ honest crowd members can be selected as the forwarder with equal probability. To illustrate, for a crowd with 10 honest members, the following transition models the second step of forwarder selection:

```
[] recordLast & CrowdSize=10 ->
      0.1: lastSeen'=0 & run'=true &
           recordLast'=false   +
      0.1: lastSeen'=1 & run'=true &
           recordLast'=false   +
      ...
      0.1: lastSeen'=9 & run'=true &
           recordLast'=false;
```

According to the protocol, each honest crowd member must decide whether to continue building the path by flipping a biased coin. With probability p_f, the forwarder selection transition is enabled again and path construction continues, and with probability $1 - p_f$ the path is terminated at the current forwarder, and all requests arriving from the initiator along the path will be delivered directly to the recipient.

```
  [] (good & !deliver & run) ->
  // Continue path construction
           PF: good'=false +
  // Terminate path construction
        notPF: deliver'=true;
```

The specification of the Crowds system imposes no upper bound on the length of the path. Moreover, the forwarders are not permitted to know their relative position on the path. Note, however, that the amount of information about the initiator that can be extracted by the adversary from any path, or any finite number of paths, is finite (see sections 4.3 and 4.5).

In systems such as Freenet [CSWH01], requests have a *hops-to-live* counter to prevent infinite paths, except with very small probability. To model this counter, we may introduce an additional state variable `pIndex` that keeps track of the length of the path constructed so far. The path construction transition is then coded as follows:

```
// Example with Hops-To-Live
// (NOT CROWDS)
//
// Forward with prob. PF, else deliver
```

```
[] (good & !deliver & run &
    pIndex<MaxPath) ->
         PF: good'=false & pIndex'=pIndex+1  +
       notPF: deliver'=true;
// Terminate if reached MaxPath,
// but sometimes not
// (to confuse adversary)
[] (good & !deliver & run &
    pIndex=MaxPath) ->
     smallP: good'=false  +
     largeP: deliver'=true;
```

Introduction of `pIndex` obviously results in exponential state space explosion, decreasing the maximum system size for which model checking is feasible.

4.2.4 Transition matrix for honest members

To summarize the state space of the discrete-time Markov chain representing correct behavior of protocol participants (*i.e.*, the state space induced by the above transitions), let $s^{(j)}_{i_1 \ldots i_k}$ be the state in which k links of the jth routing path from the initiator have already been constructed, and assume that $i_1 \ldots i_k$ are the honest forwarders selected for the path. Let $\mathbf{s}^{(j)}_{i_1 \ldots i_k}$ be the state in which path construction has terminated with $i_1 \ldots i_k$ as the final path, and let $\tilde{s}^{(j)}_{\ldots}$ be an auxiliary state. Then, given the set of honest crowd members $\mathcal{H}$ s.t. $|\mathcal{H}| = N - C$, the transition matrix T is such that $T(s^{(j)}_{i_1 \ldots i_k}, \mathbf{s}^{(j)}_{i_1 \ldots i_k}) = 1 - p_f$, $T(s^{(j)}_{i_1 \ldots i_k}, \tilde{s}^{(j)}_{i_1 \ldots i_k}) = p_f$, $\forall i \in \mathcal{H}\, T(\tilde{s}^{(j)}_{i_1 \ldots i_k}, s^{(j)}_{i_1 \ldots i_k, i}) = \frac{1}{N-C}$. Since there is no *a priori* upper bound on the length of the path, the state space of the honest members is infinite.

4.3 Model of corrupt members

Following the standard approach in security analysis, we are interested in evaluating security of the Crowds system against the *strongest possible adversary*, *i.e.*, the adversary who combines the capabilities of all hostile agents present in the system. In the worst case, a single adversary controls all corrupt crowd members and is able to correlate information obtained from different members. To model the worst-case adversary, we collapse all corrupt members into a single agent. In our formal model, this is implemented by selecting the single-agent adversary as a forwarder with probability $\frac{C}{N}$ (see section 4.2.2), *i.e.*, the probability of selecting the adversary is equal to the cumulative probability of selecting *some* corrupt member.

This abstraction does not limit the class of attacks that can be discovered using the approach proposed in this paper. Any attack found in the model where individual corrupt members are kept separate will be found in the model where their capabilities are combined in a single worst-case adversary. The reason for this is that every observation made by one of the corrupt members in the model with separate corrupt members will be made by the adversary in the model where their capabilities are combined. The amount of information available to the worst-case adversary and, consequently, the inferences that can be made from it are at least as large as those available to any individual corrupt member or a subset thereof.

In the adversary model of [RR98], each corrupt member can only observe its local network. Therefore, it only learns the identity of the crowd member immediately preceding it on the path. We model this by having the corrupt member read the value of the `lastSeen` variable, and record its observations. This corresponds to reading the source IP address of the messages arriving along the path. For example, for a crowd of size 10, the transition is as follows:

```
[] lastSeen=0 & badObserve ->
      observe0'=observe0 + 1 &
      deliver'=true & run'=true &
      badObserve'=false;
...
[] lastSeen=9 & badObserve ->
      observe9'=observe9 + 1 &
      deliver'=true & run'=true &
      badObserve'=false;
```

The counters $\texttt{observe}_i$ are persistent, *i.e.*, they are not reset for each session of the path setup protocol. This allows the adversary to accumulate observations over several path reformulations. We assume that the adversary can detect when two paths originate from the same member whose identity is unknown (see section 3.2.2).

The adversary is only interested in learning the identity of the *first* crowd member in the path. Continuing path construction after one of the corrupt members has been selected as a forwarder does not provide the adversary with any new information. This is a very important property since it helps keep the model of the adversary finite. Even though there is no bound on the length of the path, at most *one* observation per path is useful to the adversary. To simplify the model, we assume that the path terminates as soon as it reaches a corrupt member (modeled by `deliver'=true` in the transition above). This is done to shorten the average path length without decreasing the power of the adversary.

Each forwarder is supposed to flip a biased coin to decide whether to terminate the path, but the coin flips are local to the forwarder and cannot be observed by other members. Therefore, honest members cannot detect without cooperation that corrupt members always terminate paths. In any case, corrupt members can make their observable behavior indistinguishable from that of the honest members by continuing the path with probability p_f as described in section 4.2.3, even though this yields no additional information to the adversary.

4.4 Multiple paths

The discrete-time Markov chain defined in sections 4.2 and 4.3 models construction of a single path through the crowd. As explained in section 3.2.2, paths have to be reformulated periodically. The decision to rebuild the path is typically made according to a pre-determined schedule, *e.g.*, hourly, daily, or once enough new members have asked to join the crowd. For the purposes of our analysis, we simply assume that paths are reformulated some finite number of times (determined by the system parameter T=`TotalRuns`).

We analyze anonymity properties provided by Crowds after T successive path reformulations by considering the state space produced by T successive executions of the path construction protocol described in section 4.2. As explained in section 4.3, the adversary is permitted to combine its observations of some or all of the T paths that have been constructed (the adversary only observes the paths for which some corrupt member was selected as one of the forwarders). The adversary may then use this information to infer the path initiator's identity. Because forwarder selection is probabilistic, the adversary's ability to collect enough information to successfully identify the initiator can only be characterized probabilistically, as explained in section 5.

4.5 Finiteness of the adversary's state space

The state space of the honest members defined by the transition matrix of section 4.2.4 is infinite since there is no *a priori* upper bound on the length of each path. Corrupt members, however, even if they collaborate, can make at most one observation per path, as explained in section 4.3. As long as the number of path reformulations is bounded (see section 4.4), only a finite number of paths will be constructed and the adversary will be able to make only a finite number of observations. Therefore, the adversary only needs finite memory and the adversary's state space is finite.

In general, anonymity is violated if the adversary has a high probability of making a certain observation (see section 5). To find out whether Crowds satisfies

Crowd	Path reformulations			
	3	4	5	6
5 honest members	1,198	3,515	8,653	18,817
10 honest members	6,563	30,070	111,294	352,535
15 honest members	19,228	119,800	592,060	2,464,167
20 honest members	42,318	333,455	2,061,951	10,633,591

Table 2: Size of state space.

a particular anonymity property, it is thus sufficient to look *only* at the adversary's state space. We can safely ignore the (infinite) state space of the honest members, because only a finite subset thereof yields observations that can be used by the adversary to infer the path initiator's identity (see section 4.3). Because the state space of the adversary's observations is finite, the problem of finding anonymity violations for a fixed number of path reformulations is simply the problem of computing the probability of reaching some state in a finite state space, and can be handled by probabilistic model checking.

5 Formalization of Anonymity Properties

For certain values of system parameters, Crowds ensures that the originator of any path enjoys *probable innocence* against corrupt forwarders on that path (see section 3.2.1). Suppose, however, that corrupt, collaborating crowd members are able to link several paths originating from the same initiator as described in section 3.2.2. What is the likelihood that the corrupt members will be able to observe the initiator with significantly higher probability than any other member? What is their confidence in their observations? In this section, we formalize these questions as PCTL formulas over the Markov chain representing the Crowds system. In section 6, we use the PRISM model checker to answer them.

The properties we analyze are somewhat different from those considered in the original Crowds paper [RR98]. While Crowds may be "anonymous" in the probable innocence sense of section 3.2, we believe that a user who employs Crowds to hide her identity over multiple sessions with the same destination may want to know what are the chances of detection even if such detection is not, technically, a violation of probable innocence *for any given path*. Even though probable innocence provides the user with plausible deniability for each session, if the user is detected over multiple sessions, she will not be able to plausibly deny that she is the real sender.

Let K_i be the number of times the adversary observes a crowd member i, *i.e.*, there are K_i paths in which i selected a corrupt crowd member as the next forwarder, thus permitting the adversary to record i's identity. Let K_0 be the number of times the path initiator is observed—either because a corrupt crowd member was selected as the first forwarder, or because the initiator itself was selected as one of the forwarders on its own path, and is followed by a corrupt member.

We consider two notions of what it means for a crowd member to be *detected*. With metric A, a member is detected if it is observed more often than any other member, *i.e.*, $\forall j \neq i \; K_i > K_j$. With metric B, a member is detected if it is observed at least twice, *i.e.*, $K_i > 1$. The difference between these notions of detection is discussed in section 6.2.

Define events $E^a_{\text{det}}, E^b_{\text{det}}, E^b_{\text{fpos}}, E^b_{\text{nofpos}}$ as follows:

$$
\begin{array}{lcl}
E^a_{\text{det}} & = & K_0 > K_j \quad \forall j \neq 0 \\
 & & \text{(initiator observed more often than anybody else)} \\
E^b_{\text{det}} & = & K_0 > 1 \\
 & & \text{(initiator observed twice or more)} \\
E^b_{\text{fpos}} & = & K_j > 1 \quad \text{for some } j \neq 0 \\
 & & \text{(\textit{false positive}: non-initiator observed twice or more)} \\
E^b_{\text{nofpos}} & = & K_j \leq 1 \quad \forall j \neq 0 \\
 & & \text{(complement of false positive)}
\end{array}
$$

We are interested in the following probabilities:

$$
\begin{array}{lcl}
P_{\text{a}} & = & P(E^a_{\text{det}}) \\
 & & \text{(detection of the true path initiator — metric A)} \\
P_{\text{b}} & = & P(E^b_{\text{det}}) \\
 & & \text{(detection of the true path initiator — metric B)} \\
P_{\text{conf}} & = & P(E^b_{\text{nofpos}} | E^b_{\text{det}}) \\
 & & \text{(detection of \textit{only} the true initiator — metric B)}
\end{array}
$$

These probabilities are *not* conditional on selection of at least one corrupt member among the forwarders. In this setting, we analyze anonymity properties simply as a function of the total number of path reformulations without concern for whether the adversary had a chance to observe all the reformulations.

Note also that while multiple agents may be "detected" according to metric B (more than one agent may be observed at least twice by the adversary), at most one agent may be "detected" according to metric A. Therefore, for metric A, P_{conf} is always equal to 1.

Event probabilities defined above are expressed as PCTL formulas and stated in PRISM syntax. Since conditional probabilities are not supported in PRISM, P_{conf}

is computed as $\frac{P(E^b_{\text{nofpos}} \wedge E^b_{\text{det}})}{P(E^b_{\text{det}})} = P(E^b_{\text{nofpos}} | E^b_{\text{det}})$. In PRISM syntax, `[true U F] > p` stands for the $\mathcal{P}_{>p}[true\ \mathcal{U}\ \phi]$ (see section 2.1). Anonymity properties are formalized as follows (for a crowd with 10 honest members):

```
// Detection (metric A)
launch ->
    [true U (new & runCount=0 &
    observe0 > observe1 &
    observe0 > observe2 &
    ...
    observe0 > observe9)] > 0.2

// Detection (metric B)
launch ->
    [true U (new & runCount=0 &
    observe0 > 1)] > 0.2

// False positive (metric B)
launch ->
    [true U (new & runCount=0 &
    observe0 <= 1 & (
    observe1 > 1 |
    ...
    observe9 > 1)] > 0.2
```

Recall that `launch` is the flag which is true only in the initial state, whereas `new & runCount=0` is true only after all path reformulations have completed, and the adversary has collected all available observations.

6 Analysis Results

After modeling the behavior of crowd members as described in section 4, and specifying anonymity properties as described in section 5, we used PRISM to perform probabilistic model checking of different system configurations and compute the relevant probabilities. Table 2 describes the size of the state space for models of different size. The number of corrupt crowd members does not affect the size of the state space since all corrupt members are modeled as a single process (see section 4.3). The only parameter affected by the number of corrupt members is the probability of selecting a corrupt member as one of the forwarders.

Crowd		Path reformulations			
		3	4	5	6
5 honest, 1 corrupt	P_{a}	31.3%	34.5%	38.5%	42.5%
	P_{b}	13.8%	23.5%	33.3%	42.7%
	P_{conf}	100.0%	97.4%	93.1%	86.9%
10 honest, 2 corrupt	P_{a}	25.4%	27.9%	31.6%	36.1%
	P_{b}	10.4%	18.1%	26.3%	34.6%
	P_{conf}	100.0%	98.9%	96.2%	92.5%
15 honest, 3 corrupt	P_{a}	23.6%	25.8%	29.4%	34.0%
	P_{b}	9.4%	16.5%	24.1%	31.8%
	P_{conf}	100.0%	98.9%	97.5%	95.0%
20 honest, 4 corrupt	P_{a}	22.6%	24.7%	28.2%	32.8%
	P_{b}	8.9%	15.6%	23.0%	30.5%
	P_{conf}	100.0%	99.4%	97.8%	96.1%
10 honest, 1 corrupt	P_{a}	19.0%	20.4%	21.7%	23.2%
	P_{b}	3.7%	6.8%	10.5%	14.5%
	P_{conf}	100.0%	99.6%	98.1%	96.6%
20 honest, 2 corrupt	P_{a}	16.7%	17.7%	18.7%	20.0%
	P_{b}	3.0%	5.5%	8.6%	12.0%
	P_{conf}	100.0%	99.6%	98.8%	98.3%

Table 3: Probabilities of observations by the adversary.

As in most approaches based on model checking, the size of the state space to be explored increases exponentially with the size of the system, making analysis of large systems infeasible. In the Crowds case, the model has relatively few dynamic parameters and it is possible to analyze realistic system configurations with a few dozen members, similar in size to the implementations of Crowds that have actually been deployed. For example, the biggest configuration we analyzed involves 20 honest members, $\frac{1}{6}$ probability of selecting a corrupt member as a forwarder, and 6 path reformulations. Assuming that paths are rebuilt daily, as recommended by the original Crowds paper [RR98, section 8.2], this roughly models a crowd of 24 members running for a week.

The state explosion problem is significantly worse for systems with parameters whose value changes at each stage of routing path construction, *e.g.*, the *hops-to-live* counter (see section 4.2.3). For such systems, only fairly small configurations (up to 10 members) can be feasibly analyzed with PRISM.

Table 3 lists computed event probabilities. In all of the experiments, forwarding

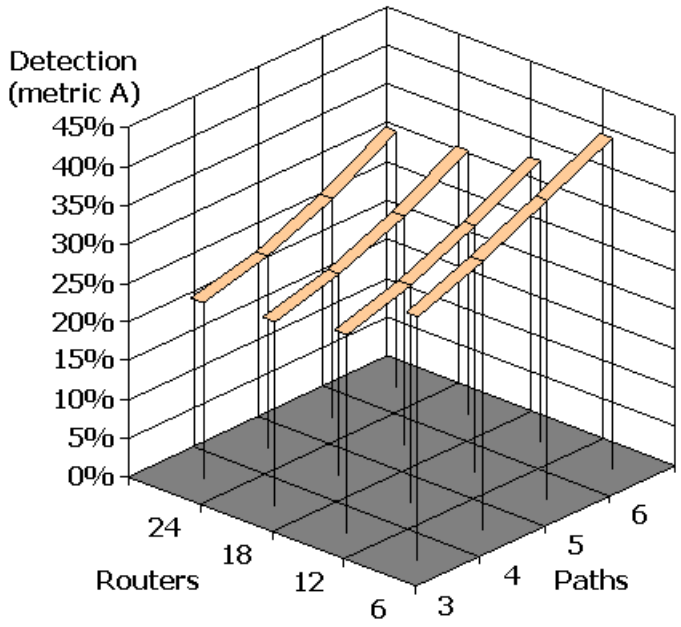

Figure 1: Metric A: probability of observing the true initiator more often than any other member ($\frac{1}{6}$ of routers are corrupt)

probability $p_f = 0.8$, and c, n and p_f satisfy condition 1. Therefore, for any given single path, the initiator enjoys probable innocence.

Recall that P_{a} and P_{b} are the probabilities of, respectively, observing the true path initiator more often than any other crowd member and observing the initiator twice or more, while $P_{\text{conf}} = P(E^b_{\text{nofpos}} | E^b_{\text{det}})$ is the probability of observing *only* the initiator twice or more. P_{conf} can be interpreted as the adversary's "confidence." If P_{conf} is high, as soon as the corrupt members observe the same honest member twice, they can be confident that the member is indeed the path initiator.

6.1 Increasing path reformulations

As conjectured by the original Crowds paper [RR98] and independently predicted by Malkhi [Mal01] and Wright *et al.* [WALS02], anonymity guarantees provided by the system degrade with the increase in the number of different paths that may be observed by the adversary and linked as initiated by the same crowd member. This holds for both detection metrics considered in this paper. After relatively few path reformulations—even if not all of the paths involve corrupt members—

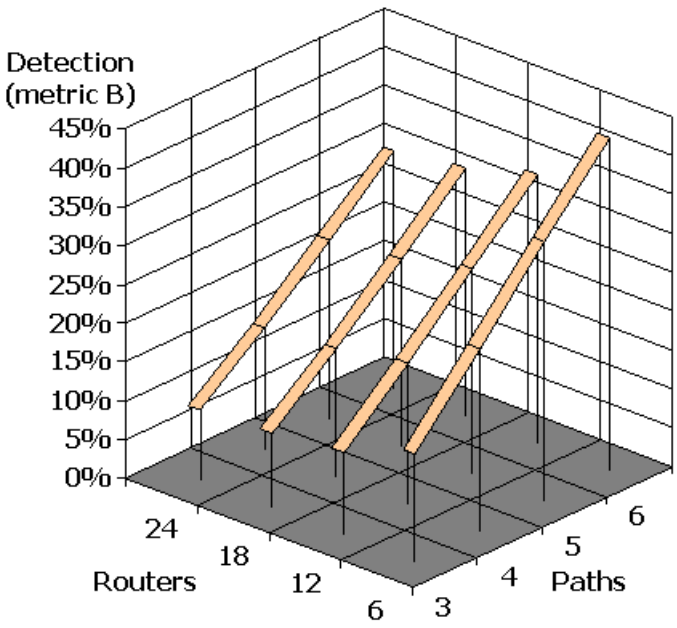

Figure 2: Metric B: probability of observing the true initiator at least twice ($\frac{1}{6}$ of routers are corrupt)

the probability of observing the path initiator more often than any other member grows significantly (figure 1), and so does the probability of observing the path initiator more than twice (figure 2). This means that even with static paths and the corresponding reduction in the frequency of path reformulation (see section 3.2.3), the system could be vulnerable. For example, in a crowd of 6 members, only 1 of whom is corrupt, the single corrupt member has a better than 30% chance of detecting the true path initiator (E^a_{det} and E^b_{det} events) after 5 path reformulations *without* assuming that it is selected as one of the forwarders in every path.

6.2 Comparison of detection metrics

In our analysis, we consider two notions of what it means to "detect" a crowd member. With metric A, a member is detected if it is observed by the adversary more frequently than any other member. With metric B, a member is detected if it is observed at least twice. Direct comparison between the two notions is not straightforward, and depends on non-technical factors outside the protocol specification, such as the purpose of the adversary's observations.

Metric A has the benefit of being unambiguous: no more than one crowd mem-

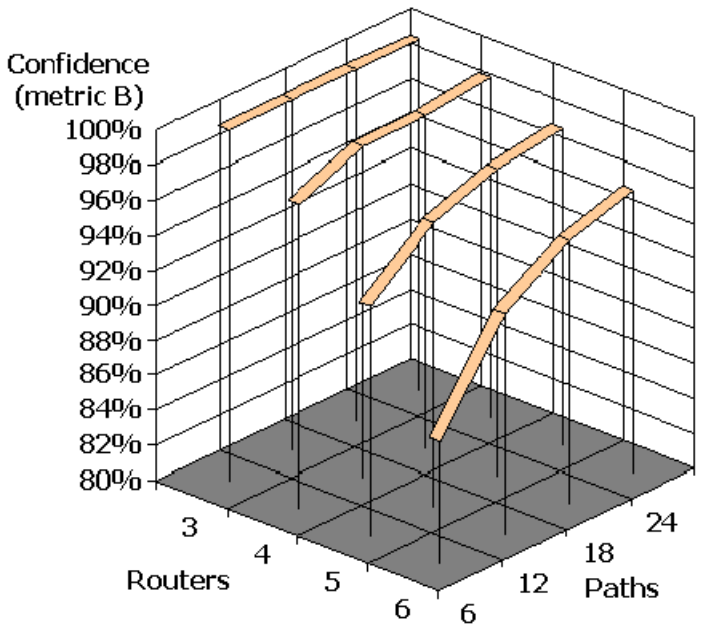

Figure 3: Metric B: probability of observing only the true initiator at least twice ($\frac{1}{6}$ of routers are corrupt)

ber can possibly be detected. Therefore, the adversary's "confidence" P_{conf} is always 100%. Metric B, on the other hand, provides stronger evidence (*e.g.*, for investigative purposes), at least for configurations where P_{conf} is high, since it always requires multiple observations of the same agent. For example, suppose there have been 3 path reformulations, and a corrupt member was selected as a forwarder in only 1 of the paths. Whichever honest member happened to precede the corrupt member on that path will be considered "detected" according to metric A, since it has been observed more often than any other member ($1 > 0$). In this case, metric B would provide higher assurance that the true initiator has been detected.

Since more than one member can be detected according to metric B, it is most useful when the adversary's "confidence" P_{conf} is high, *i.e.*, for a small number of paths or a large crowd (see section 6.3 for the explanation of the latter point). As can be seen in table 3, P_{conf} decreases in each row with the increase in the number of path reformulations. The reason for this is that as more paths are constructed, the chances of a random honest member appearing twice or more before a corrupt member and thus being mistaken for the initiator (E^b_{fpos} event) increase.

6.3 Increasing crowd size

A somewhat surprising result, uncovered by automated analysis with PRISM, is the change in P_{conf} for metric B with the increase in the size of the crowd as long as the proportion of corrupt members remains constant. As the crowd grows, P_{conf} actually increases for any given number of path reformulations (see figure 3). This implies that the larger the crowd, the more confidence the adversary has that if it observes the same honest member at least twice, that member is the true initiator. Since the probability of detection P_{b} decreases only slightly with the increase in the size of the crowd, increased confidence of the adversary in its observations can be interpreted as a degradation of anonymity.

An intuitive explanation of this result is that in a sufficiently large crowd, a random honest member has only a negligible chance of being selected for more than one path (in the extreme case of an infinite crowd, the probability that a forwarder who is not the initiator appears in two or more different paths is 0). The only member that has a non-negligible probability of appearing in multiple paths is the path initiator. Therefore, assuming detection (E^b_{det}) occurs, the adversary's confidence that the true initiator was detected grows with the size of the crowd.

7 Conclusions

Probabilistic model checking is a well-established technique for verification of hardware and concurrent protocols. The main contribution of this paper is to demonstrate how it can be applied to the analysis of security properties based on discrete probabilities. As a case study, we analyzed anonymity properties of the Crowds system, a "real-world" protocol for anonymous Web browsing.

Anonymity in Crowds is based on constructing a random routing path to the destination through a group of members, some of whom may be corrupt. The path construction protocol is purely probabilistic, therefore, we modeled it as a discrete-time Markov chain, without introducing non-determinism and thus avoiding the need for Markov decision processes. We considered the worst-case local adversary, who combines the capabilities of all corrupt crowd members, but can only make an observation if one of the corrupt members was selected as a forwarder. The adversary was permitted to combine its observations of a finite number of different paths, modeling the fact that paths in Crowds must be reformulated on a regular basis. Since the number of paths is finite, the state space of the adversary's observations is also finite. Therefore, the problem of analyzing anonymity — that is, computing the probability that the adversary will be able to successfully infer the identity of the path initiator — is amenable to automated probabilistic model checking.

In addition to proving feasibility of the model checking approach to verification of probabilistic security properties, we uncovered potential vulnerabilities of the Crowds system. These include the increase in the probability that the true path initiator will be detected as the number of path reformulations grows, and the increase in the adversary's confidence with the increase in crowd size. The former has been reported by other researchers (the model described in this paper had been constructed independently before the other results were published), while the latter was reported for the first time in the conference version of this paper. We also show that correctly stating the definition of a successful "attack" on anonymity is a non-trivial task. There are several possible definitions of what it means for the adversary to "detect" the path initiator, and direct comparison between them is not always possible.

Acknowledgements. This paper has greatly benefitted from the comments of the anonymous reviewers, and discussions with Dahlia Malkhi and Jon Millen.

References

[AH96] R. Alur and T. Henzinger. Reactive modules. In *Proc. 11th Annual IEEE Symposium on Logic in Computer Science (LICS)*, pages 207–218, 1996.

[Ano] `http://www.anonymizer.com`.

[Bai98] C. Baier. On algorithmic verification methods for probabilistic systems, 1998. Fakultät für Mathematik & Informatik, Universität Mannheim.

[BdA95] A. Bianco and L. de Alfaro. Model checking of probabilistic and nondeterministic systems. In *Proc. Foundations of Software Technology and Theoretical Computer Science (FST & TCS)*, volume 1026 of *LNCS*, pages 499–513. Springer-Verlag, 1995.

[BK98] C. Baier and M. Kwiatkowska. Model checking for a probabilistic branching time logic with fairness. *Distributed Computing*, 11(3):125–155, 1998.

[Can00] R. Canetti. A unified framework for analyzing security of protocols. IACR Cryptology ePrint Archive 2000/067 (`http://eprint.iacr.org`), December 2000.

[Cha81] D. Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, 24(2):84–88, 1981.

[CJM00] E.M. Clarke, S. Jha, and W. Marrero. Verifying security protocols with Brutus. *ACM Transactions in Software Engineering Methodology*, 9(4):443–487, 2000.

[CSWH01] I. Clarke, O. Sandberg, B. Wiley, and T.W. Hong. Freenet: A distributed anonymous information storage and retrieval system. In *Proc. International Workshop on Design Issues in Anonymity and Unobservability*, volume 2009 of *LNCS*, pages 46–66. Springer-Verlag, 2001.

[DY83] D. Dolev and A. Yao. On the security of public key protocols. *IEEE Transactions on Information Theory*, 29(2):198–208, 1983.

[Gra92] J.W. Gray. Toward a mathematical foundation for information flow security. *J. Computer Security*, 1(3):255–294, 1992.

[HJ94] H. Hansson and B. Jonsson. A logic for reasoning about time and probability. *Formal Aspects of Computing*, 6(5):512–535, 1994.

[HS84] S. Hart and M. Sharir. Probabilistic temporal logics for finite and bounded models. In *Proc. ACM Symposium on the Theory of Computing (STOC)*, pages 1–13, 1984.

[KNP01] M. Kwiatkowska, G. Norman, and D. Parker. PRISM: Probabilistic symbolic model checker. Technical Report 760/2001, University of Dortmund, September 2001. Also in Proc. PAPM/PROBMIV 2001 Tools Session.

[LMMS99] P. Lincoln, J.C. Mitchell, M. Mitchell, and A. Scedrov. Probabilistic polynomial-time equivalence and security analysis. In *Proc. World Congress on Formal Methods*, volume 1708 of *LNCS*, pages 776–793. Springer-Verlag, 1999.

[Low96] G. Lowe. Breaking and fixing the Needham-Schroeder public-key protocol using FDR. In *Proc. Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, volume 1055 of *LNCS*, pages 147–166. Springer-Verlag, 1996.

[LS82] D. Lehmann and S. Shelah. Reasoning with time and chance. *Information and Control*, 53(3):165–198, 1982.

[Mal01] D. Malkhi. Private communication, 2001.

[MMS97] J.C. Mitchell, M. Mitchell, and U. Stern. Automated analysis of cryptographic protocols using Murφ. In *Proc. IEEE Symposium on Security and Privacy*, pages 141–153, 1997.

[Pau98] L. Paulson. The inductive approach to verifying cryptographic protocols. *J. Computer Security*, 6(1):85–128, 1998.

[RR98] M. Reiter and A. Rubin. Crowds: Anonymity for web transactions. *ACM Transactions on Information and System Security*, 1(1):66–92, 1998.

[SG95] P. Syverson and J.W. Gray. The epistemic representation of information flow security in probabilistic systems. In *Proc. 8th IEEE Computer Security Foundations Workshop*, pages 152–166, 1995.

[SGR97] P. Syverson, D. Goldschlag, and M. Reed. Anonymous connections and onion routing. In *Proc. IEEE Symposium on Security and Privacy*, pages 44–54, 1997.

[Shm02] V. Shmatikov. Probabilistic analysis of anonymity. In *Proc. 15th IEEE Computer Security Foundations Workshop*, pages 119–128, 2002.

[SS96] S. Schneider and A. Sidiroupoulos. CSP and anonymity. In *Proc. ESORICS*, volume 1146 of *LNCS*, pages 198–218. Springer-Verlag, 1996.

[SS99] P. Syverson and S. Stubblebine. Group principals and the formalization of anonymity. In *Proc. World Congress on Formal Methods*, volume 1708 of *LNCS*, pages 814–833. Springer-Verlag, 1999.

[Var85] M. Vardi. Automatic verification of probabilistic concurrent finite-state programs. In *Proc. 26th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 327–338, 1985.

[VS98] D. Volpano and G. Smith. Probabilistic non-interference in a concurrent language. In *Proc. 11th IEEE Computer Security Foundations Workshop*, pages 34–43, 1998.

[WALS02] M. Wright, M. Adler, B.N. Levine, and C. Shields. An analysis of the degradation of anonymous protocols. In *Proc. ISOC Network and Distributed System Security Symposium (NDSS)*, 2002.