On the Analysis of the *Predecessor Attack* on Anonymity Systems

Daniel R. Figueiredo[1] Philippe Nain[2] Don Towsley[1] *

[1]Dept. of Computer Science
University of Massachusetts
Amherst, MA 01003
{ratton, towsley}@cs.umass.edu

[2]INRIA
BP 93
06902 Sophia Antipolis, France
Philippe.Nain@inria.fr

Computer Science Technical Report 04-65

July 29, 2004

Abstract

Systems that allow users to communicate anonymously with a destination have received increasingly more attention since users of network applications became more concerned with their privacy. Unfortunately, anonymity systems are vulnerable to attacks that attempt to reveal the identity of nodes that communicate anonymously. Moreover, the distributed nature of such systems facilitates certain types of attacks. In this paper we focus on the *predecessor attack*, a robust traffic analysis attack that targets nodes that communicate with the same destination repeatedly over time. In particular, we perform a qualitative analysis of the attack using a generalized model for anonymous communication. We establish the necessary and sufficient conditions for the attack to succeed and also determine the effort required by the attacker. We consider different situations and investigate the scenario where multiple nodes communicate with the same destination. Our results show that for a common class of protocols, where paths are constructed uniformly at random, the attack always succeeds and the effort required is proportional to the number of initiators and the number of nodes in the system. Moreover, knowing the number of initiators present in the system does not reduce the effort required by the attacker. Understanding the capabilities and limitations of this attack is an important step toward designing more secure anonymity systems.

Keywords: anonymous protocols, predecessor attack, statistical analysis

1 Introduction

In recent years, users of network applications such as the World Wide Web have become increasingly more concerned with their privacy. Such concerns have led to an increased interest in practical interactive anony-

*This research has been supported in part by the NSF under grant awards EIA-0080119, ANI-0070067 and ANI-0085848, and by CAPES (Brazil). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

mous communication protocols, which aim at disguising the identity of the parties involved in low-latency bi-directional communication.

Several interactive anonymous communication systems have been proposed in the recent literature [10, 4, 11, 2]. In such systems, a group of nodes collectively provide anonymity by forwarding a message among themselves an arbitrary number of times before sending it to its intended destination, rendering the identity of the node that originated the message anonymous within the group of collaborating nodes. That is, from the destination's perspective, all group members are equally likely to have generated the message. Furthermore, a node that forwards messages cannot determine if its predecessor is the node that originated the message or an intermediate node along the forwarding path. The node originating a message is known as the *initiator*. Protocols that operate in this general manner are known as path-based anonymous protocols, as the initiator constructs a random path through the set of nodes participating in the system to reach the intended destination.

Unfortunately, anonymity systems are vulnerable to various attacks that attempt to reveal the identity of a user that is communicating anonymously. Traffic analysis is a class of attacks that has been shown to be powerful in breaking anonymity [1]. Defeating traffic analysis is not a trivial task and system designers usually target defenses against specific attacks. Although various types of traffic analysis attacks can be successfully defeated, a particular attack, known as the *predecessor attack*, has been shown to be robust in many proposed anonymity systems [13]. The predecessor attack allows an attacker with limited resources to learn the identity of an initiator that communicates repeatedly with a given destination for long periods of time.

The distributed nature of interactive path-based anonymity systems, many of which are based on the peer-to-peer (P2P) paradigm and can be deployed with little or no infrastructure, facilitates launching a predecessor attack. In such an environment, where nodes have no prior trust relationships with each other, an attacker can easily deploy multiple corrupted nodes, or simply impersonate multiple nodes using a single physical computer; recently, this is known as a Sybil attack [3]. By controlling more nodes in the system the attacker can gather more information in a shorter period of time, increasing the effectiveness of the predecessor attack. Even more troublesome is the fact that by complying with the anonymity protocol, it is very difficult, if not impossible, to detect attacker nodes. The threat posed by the predecessor attack on real path-based anonymous protocols is very realistic when initiators communicate repeatedly with the same destination.

In this paper we perform a qualitative analysis of the predecessor attack on generalized path-based anonymity protocols. The goal is to characterize the capabilities of this attack, establishing the necessary and sufficient conditions for the attack to succeed in different scenarios, and determining the effort required by the attacker to reveal the identity of the initiators. In particular, we consider the case where multiple initiators communicate with the same destination and establish the conditions under which attackers can correctly identify this set of nodes. Our results show that for a common class of protocols, where paths are constructed uniformly at random, the effort required by an attacker to reveal the set of initiators grows linear in the number of initiators present in the system (Section 5). Moreover, this effort is the same in both the cases where the number of initiators in the system is known to the attacker or not.

This remainder of this paper is organized as follows. In Section 2 we present previous related work. Section 3 describes the generalized model for anonymous communication used in the analysis and a concrete example of an anonymity system. The results for the single initiator case are presented in Section 4, while the analysis for multiple initiators is presented in Section 5. Section 6 presents a numerical analysis of the attack illustrating existing tradeoffs. Finally, Section 7 summarizes the paper and presents some discussion.

2 Related Work

Much work has been done evaluating anonymous protocols with respect to the quality of anonymity they can provide to their users. In particular, a number of metrics have been used to assess the quality of anonymous protocols, all related to the probability of correctly identifying a message initiator [10, 8, 5, 12]. Most metrics share the property that their value is proportional to the number of collaborating nodes in the system.

The predecessor attack was first pointed out in [10], where the authors provide an analysis of a specific anonymous protocol known as Crowds. More recently, [13] used a probabilistic framework to evaluate the predecessor attack in various anonymity systems. In particular, they prove that under the assumption that paths are constructed uniformly at random, the attacker always succeeds in correctly revealing the initiator, independent of the protocol being used. They also provide explicit upper bounds for the effort required by an attacker in several concrete protocols.

Our work differs from [13] in several aspects. First, we consider a more general formulation of the attack which is independent of the underlying protocol. Our formalization maps directly to the criteria used by an attacker to conduct the attack. Second, we derive asymptotic results, and not only upper bounds, for the effort required by attackers. Third, we investigate the more interesting case where multiple initiators communicate with the same destination. Within this scenario, we consider a several criteria for identifying the various initiators present in the system.

3 Model for Anonymous Communication and Predecessor Attack

In this section, we introduce a model for anonymous communication that is independent of specific protocols proposed and allows for a general analysis of the predecessor attack. We assume that the anonymity system is composed of a fixed set of nodes and that each node is uniquely identifiable (e.g., the IP address of each node will serve as their identity). The attacker controls (e.g., operates) a fixed set of attacker nodes, A. The remainder of the nodes in the system are honest, and are denoted by the set N. Let $n = |N|$ and $c = |A|$ be the number of honest and attacker nodes in each set, respectively (the total number of nodes in the system is $n + c$). A subset of the honest nodes, $I \subset N$, communicates anonymously with a fixed destination D outside the system. From the perspective of an initiator, all nodes in the system (honest and attackers) behave exactly the same and are indistinguishable. We assume all nodes in I communicate repeatedly and indefinitely with D. The goal of the attacker is to identify the members of I. Note that honest nodes may also be communicating anonymously with destinations different from D. However, we assume the attacker is only interested in revealing the identity of the nodes that are communicating with D. Of course, the attacker may also launch an attack to determine the set of initiators that are communicating with some other destination D', but this would be completely orthogonal.

In order to communicate anonymously with a destination, the initiator first constructs a path through a sequence of nodes in the system that terminates at the intended destination. The path construction mechanism installs state in each intermediary node along the path such that subsequent messages and responses can be sent through this bidirectional path. Each node on the path knows the identity (e.g., IP address) of its immediate predecessor and successor nodes and these identities cannot be forged, as direct communication between them will take place. After some time, the initiator destroys the existing path to the destination and constructs a new one in order to continue communication. Although this is a practical concern of real

protocols, we simply assume that new paths must be reconstructed periodically. [1]

Attacker nodes are passive, never generating messages into the system, and fully-compliant with the anonymous protocol. They simply wait to be selected by some initiator to be on an anonymous path. Thus, each path constructed by an initiator can have zero, one, or more attacker nodes. We assume that the attacker can detect if more than one of their nodes lies on the same path. This assumption is reasonable and has been shown to be feasible [7]. We also assume that the attacker can observe the address of the intended destination when an attacker node is chosen to be on a path. Note that this information must appear in plain text (at least to the last node on the path) in order for the destination to be contacted.

The predecessor attack works as follows. The attacker nodes collectively maintain a single predecessor counter for each honest node in the system. Initially, all counters are set to zero. When an attacker node is selected to be on a new anonymous path, it first verifies if this path is intended for destination D. If so, the attacker increments the shared counter for its predecessor node in this path. These counters represent the number of times that each node in the system was observed as a predecessor to an attacker node on paths toward destination D. The attacker will then use the value of the predecessor counters to determine the set of nodes that are initiators.

3.1 Probabilistic Modeling

In order to analyze the predecessor attack within the framework described above, we must formalize how the protocol and the attack operates. We assume the anonymity system evolves in a discrete time fashion. At each instance of time, an initiator in the set I constructs a new anonymous path to destination D. The path constructed may contain zero or one attacker node (if it contains more than one, consider only the first attacker node). Let X_t denote the node that preceded an attacker node in the t-th path constructed by the initiators. $X_t \in X = \{e_1, e_2, \ldots, e_n, e_{n+1}\}$, where e_i is the unit vector in $\mathbb{N}^{n+1}$ with all components equal to zero except the i-th component that is equal to one. If node i is the predecessor to an attacker node in the t-th path constructed, then $X_t(i) = 1$ and $X_t(j) = 0$, for all $j \neq i$. $X_t(n+1) = 1$ if the attacker does not appear on the t-th path constructed by the initiators. Now consider the sequence $X_1, X_2, \ldots, X_T$ of vectors. Define $N_i(T) = \sum_{t=1}^{T} X_t(i)$, and note that $N_i(T)$ indicates the number of times that node i appeared as a predecessor to some attacker node after T paths were constructed, with $1 \leq i \leq n$. Moreover, $N_{n+1}(T)$ is the number of times that the attacker did not appear on a path after a total of T paths were constructed. The value of this last counter is not known by the attacker and thus cannot be during the attack. As an observation, note that $T = \sum_{i=1}^{n+1} N_i(T)$. Using these predecessor counters, the attacker will classify each node in the system as either an initiator or a non-initiator. The specific criteria for classification will be presented in the following sections.

To continue the analysis we must define how anonymous paths are constructed in the system. However, this depends on specific protocols and their parameterization. Thus, instead of dealing with a specific protocol instance, we will consider basic events that are capable of capturing any protocol. The set of probabilities associated with these events will be the sole parameter of our analysis.

We assume that nodes along an anonymous path from an initiator to destination D are chosen probabilistically. Moreover, paths constructed by the initiators are statistically independent from one another. In

[1] There are several reasons why paths must be periodically reconstructed, such as coping with changes in group membership [13], and reducing the probability of success of other types of attacks [11].

particular, $\{X_t\}$ forms a sequence of statistically independent and identically distributed random variables. Under these assumptions, we consider the probability that an attacker node is preceded by a given node $i, 1 \leq i \leq n$, when an anonymous path to destination D is constructed by some initiator. We denote this probability by b_i. Figure 1 illustrates this event. Note that any of the attacker nodes can appear on the path.

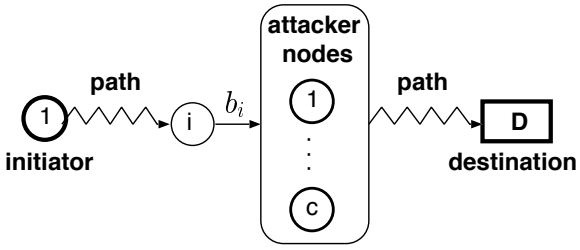

Figure 1: Idealized model showing node i appearing as the predecessor of an attacker node (which occurs with probability b_i)

The values of these probabilities depend on the path construction mechanism and its parameterization. Although we will work with an arbitrary set $\{b_i\}_{1\leq i\leq n}$ to determine the conditions under which the predecessor attack succeeds as well as the effort required by the attacker, in the next section we present a specific anonymity protocol and derive values for $\{b_i\}_{1\leq i\leq n}$.

3.2 A Path Construction Mechanism

In order to illustrate how $\{b_i\}_{1\leq i\leq n}$ can be computed, we introduce a concrete path construction mechanism that is a generalization of the approach used in the Crowds protocol [10]. The path construction starts with one of the initiators randomly choosing a node in the system, potentially itself or even the intended destination, to be next hop of the anonymous path. Each subsequent node in the path performs the same randomized choice, selecting either some node in the system to be the next hop on the path or terminating the path by selecting the destination node as the next hop. Let p_i be the probability of choosing node $i, 1 \leq i \leq n+c$ to be the next hop of a path. Moreover, let p_r be the probability of a node terminating the anonymous path and choosing the destination as the next hop.

Recall that the attacker controls a fixed set of nodes, A, in the system, with $|A| = c$. Thus, the probability an attacker nodes is chosen as the next hop in the path is given by $p_a = \sum_{i\in A} p_i$. If an attacker node is selected to be on the path, then it terminates the path by always choosing the destination as the next hop. Note that $\sum_{i\notin A} p_i + p_a + p_r = 1$. Figure 2 illustrates this path construction mechanism for system composed of 4 honest nodes and a set of attacker nodes.

Using this path construction mechanism, we can calculate $b_i, 1 \leq i \leq n$, the probability that the attacker will observe node i as its predecessor. By first conditioning on a given initiator and on the number of intermediary nodes visited between that initiator and node i, and then applying the law of total probability, we obtain:

$$b_j = \frac{p_j p_a}{p_r + p_a} \text{ if } j \notin I \tag{1}$$

$$b_i = \frac{p_i p_a}{p_r + p_a} + q_i p_a \text{ if } i \in I \tag{2}$$

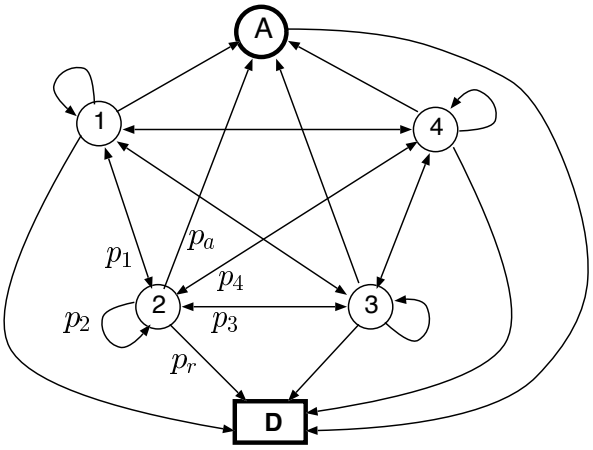

Figure 2: Path construction mechanism with 4 honest nodes and a set of attacker nodes.

where q_i is the probability that initiator i is the node constructing the path. Note that $b_i, i \in I$ contains the additional disjoint event of the initiator directly choosing the attacker as the first hop in the path.

From the equations above, we observe that if $p_i, 1 \leq i \leq n+c$ is publicly known, then the attacker can trivially compute $b_j, j \notin I$, and use this information during the predecessor attack. However, note that the attacker cannot compute $b_i, i \in I$, since it does not know $q_i, i \in I$, a priori.

3.2.1 Uniform path construction

Consider the special case of the above path construction mechanism where nodes select the next hop of the path uniformly at random among all nodes in the system. In this case, every node has the same probability as being selected as the next hop, thus, $p = p_1 = p_2 = \cdots = p_{n+c}$. Since the total number of nodes in the system is given by $n+c$, we have that $p = (1-p_r)/(n+c)$, and since the attacker controls c nodes, we have that $p_a = (1-p_r)c/(n+c)$. This uniform path construction scenario has been suggested in different anonymous protocols, and in particular, in the Crowds protocol [10]. Moreover, assume that all initiators communicate at the same rate with the given destination, such that $q_i = 1/|I|$ for $i \in I$. Thus, the basic probabilities are:

$$b_j = \frac{c(1-p_r)}{n+c} \frac{1-p_r}{np_r+c} \text{ if } j \notin I \tag{3}$$

$$b_i = \frac{c(1-p_r)}{n+c} \left(\frac{1-p_r}{np_r+c} + \frac{1}{|I|} \right) \text{ if } i \in I \tag{4}$$

An important observation from these equations is that all non-initiators have exactly the same probability of being a predecessor to an attacker node, while an initiator has a higher probability than a non-initiator. From this observation, one can understand why an attacker succeeds in identifying the set of initiators. In what follows, we will refer back to this uniform path construction mechanism, showing more formally that this is indeed the case.

Although the above derivation of $\{b_i\}_{1 \leq i \leq n}$ is pertinent to an anonymity system that operates a path construction mechanism corresponding to the Crowds protocol [10], it should be noted that a similar derivation can be carried out for generalized models of other existing protocols, such as [4, 11, 9].

4 Single Initiator Case

In this section we consider the case where a single node in the system is communicating anonymously with destination D, thus $|I| = 1$. The goal of the attacker is to identify this sole initiator among the set of honest nodes using the system.

The criterion used by the attacker to classify a node as the initiator is simple. The attacker chooses as initiator the node which has the highest predecessor counter. Assume without loss of generality, that node 1 is the initiator node. Then, the attacker correctly reveals the identity of the initiator if $N_1(T) > \max_{1<j\leq n}\{N_j(T)\}$ after T paths have been constructed by the sole initiator.

Given this criterion to identify the initiator, we can determine conditions under which the attack will succeed. In particular the following condition is necessary and sufficient:

$$b_1 > \max_{1<i\leq n}\{b_i\} \tag{5}$$

Intuitively, the probability that the attacker has the initiator node as a predecessor when a path is constructed has to be higher than that for any other node in the network. If this is the case, the attack will succeed given a sufficiently large number of paths constructed, as stated in the following theorem.

Theorem 1 *If* $b_1 > \max_{1<i\leq n}\{b_i\}$ *then* $\Pr[\lim_{T\to\infty} N_1(T) > \max_{1<j\leq n}\{N_j(T)\}] = 1$.

Proof: Without loss of generality, let b_2 be the maximum value among $b_i, 1 < i \leq n$, that is $b_2 = \max_{1<j\leq n}\{b_j\}$. From the strong law of large numbers we have that for all i, $\Pr[\lim_{T\to\infty} N_i(T)/T = b_i] = 1$. Thus, we have that $\Pr[\lim_{T\to\infty} N_1(T)/T = b_1] = 1$ and $\Pr[\lim_{T\to\infty} N_2(T)/T = b_2] = 1$. Since $b_1 > b_2$ by assumption, we conclude that $\Pr[\lim_{T\to\infty} N_1(T)/T > N_2(T)/T] = 1$. □

Thus, for any protocol for which condition (5) holds, the probability the attacker correctly identifies the initiator converges to one as the number of paths constructed by the initiator increases. In particular, the uniform path construction protocol presented in Section 3.2.1 satisfies condition (5) and therefore is subject to attack. The result in the theorem above was also established in [13] under the assumption that $b_2 = b_3 = \cdots = b_n$.

Note that if condition (5) does not hold for some node i, then the probability the attacker will correctly identify the initiator converges to zero. In this case, the node violating condition (5) will always be taken by the attacker to be initiator. For example, if the vast majority of the paths constructed by an initiator has node k as the first hop, and node k is not an attacker node, then an attacker conducting a predecessor attack will classify node k as the initiator, as b_k will be greater than b_1. A similar observation was made in [13] and this was referred to as a "setup attack".

The previous theorem only states that the predecessor attack will succeed with probability one when the number of paths constructed by the initiator approaches infinity. However, in practice we would like to characterize the number of paths that must be constructed by the initiator in order for the attacker to succeed with high probability. In particular, we would like to obtain the value for T for which $\Pr[N_1(T) > \max_{1<j\leq n}\{N_j(T)\}]$ is "sufficiently large" or "close enough" to one. This value reflects the effort required by the attacker in order to succeed, meaning that the attacker will need to wait for the initiator to construct this

many paths in order to identify it with high probability. Thus, we need to determine T such that following holds:

$$\Pr[N_1(T) > \max_{1<j\leq n}\{N_j(T)\}] = 1 - 1/f(n) \tag{6}$$

where $f(n)$ is any positive polynomial function in n. Note that the set of counters $\{N_i(T)\}_i$ form a multinomial random variable with $n+1$ possible outcomes with each outcome having probability b_i. Hoeffding has computed the asymptotic tail for a general set of events that have a multinomial distribution [6]. Using this result we can obtain an analytical expression for the probability defined in equation (6). This latter expression together with the assumption that n is large, allow us to prove the following theorem:

Theorem 2 *If* $T = \Theta(\log n/b_1)$ *then* $\Pr[N_1(T) > \max_{1<j\leq n}\{N_j(T)\}] = 1 - 1/f(n)$.

The proof is found in the appendix.

Thus, the attacker can reveal the identity of the initiator with high probability if the initiator constructs $T = \Theta(\log n/b_1)$ anonymous paths. Note that T depends on n, the number of honest nodes in the system, and on b_1, the probability that the initiator precedes an attacker node. A system with a larger number of honest nodes increases the effort required by the attacker, while increasing b_1 reduces the required effort. The value of b_1 depends on the actual protocol being used, however, one would expect that by increasing the number of attacker nodes in the system, c, the value of b_1 should increase. In fact, this is the case for the uniform path construction mechanism presented in Section 3.2.1. In this case, $b_1 = \Theta(c/n)$, such that the effort required to succeed with high probability is given by $T = \Theta(n \log n/c)$. Note that T decreases as the number of attacker nodes in the system increases. An upper bound for the effort in systems that use an uniform path construction mechanism was derived in [13]. The above asymptotic expression shows that their upper bound was indeed tight.

5 Multiple Initiator Case

In this section we consider the more general setting where multiple initiators in the system communicate repeatedly with the same destination. This situation maybe expected for example, when the destination in under consideration is a popular site, when a group of users coordinate their communication, or simply by chance in an anonymity system with a large and diverse user population.

There is a fundamental difference between the single initiator case and the multiple initiator case, namely that in the latter the attacker cannot link a specific communication with the destination to a given initiator. The attacker can only identify the set of initiators, but cannot determine which initiator is indeed communicating when a given path is constructed.

Consider the set I of nodes in the system that communicate repeatedly and indefinitely with destination D. Each initiator in set I may communicate at a different rate with the destination. In particular, let $r_i > 0,\ 1 \leq i \leq |I|$ be the rate at which initiator $i \in I$ establishes anonymous paths with destination D. We will assume that initiators communicate uniformly over time with the destination, such that communication rates can be mapped into Bernoulli probabilities. Thus, $q_j = r_j / \sum_{i\in I} r_i$ is the probability that a given path is constructed by initiator j.

As before, we will use the basic event that an attacker node will be preceded by some given node j when a path to destination D is constructed. Let $b_j^{(i)}$ denote the probability node j precedes an attacker node when initiator i is the node establishing the path. In general, this probability may depend on the initiator. By the law of total probability, we have

$$b_j = \sum_{i \in I} b_j^{(i)} q_i$$

The attacker will again maintain the predecessor counters and use them to identify the set of initiators. As before, $N_j(T)$ indicates the number of times that node j appeared as a predecessor to an attacker node after a total of T paths were constructed by the initiators, with $1 \leq j \leq n$.

In a first analysis, let's assume the attacker knows the number of initiators in the system, but not their identities. In particular, let $m = |I|$ be the number of initiators. In this case, the attacker can proceed in a similar fashion as in the single initiator case, and classify as initiators the m nodes which have the highest predecessor counters. The attacker will correctly identify the set I if $\min_{i \in I}\{N_i(T)\} > \max_{j \notin I}\{N_j(T)\}$ after a total of T paths have been constructed.

Using this criterion, we can determine the conditions under which the attack will succeed. In particular the following condition is necessary and sufficient for a successful attack:

$$\min_{i \in I}\{b_i\} > \max_{j \notin I}\{b_j\} \tag{7}$$

Intuitively, the probability any initiator is a predecessor to the an attacker node has to be higher than that of any other node in the system. Note that this condition is similar to the single initiator case. The attack will succeed in correctly identifying the set I given a sufficient number of paths constructed, as stated in the following theorem.

Theorem 3 *If condition (7) holds, then* $\Pr[\lim_{T \to \infty} \min_{i \in I}\{N_i(T)\} > \max_{j \notin I}\{N_j(T)\}] = 1$.

Proof: Without loss of generality, let b_1 be the minimum value among $b_i, i \in I$, that is, $b_1 = \min_{i \in I}\{b_i\}$. Without loss of generality, let b_2 be the maximum value among $b_j, j \notin I$, that is $b_2 = \max_{j \notin I}\{b_j\}$. From the strong law of large numbers we have that for all j, $\Pr[\lim_{T \to \infty} N_j(T)/T = b_j] = 1$. Since $b_1 > b_2$ by assumption, we conclude that $\Pr[\lim_{T \to \infty} \min_{i \in I}\{N_i(T)\} > \max_{j \notin I}\{N_j(T)\}] = 1$. □

It is also interesting to characterize the effort required by the attacker in order to identify the set of initiators with high probability. Following the derivation for the single initiator case, we determine T such that $\Pr[\min_{i \in I}\{N_i(T)\} > \max_{j \notin I}\{N_j(T)\}]$ is sufficiently large. The derivation of this result is identical to the single initiator case, where $b_1 = \min_{i \in I}\{b_i\}$. By again assuming that n is large, we can state the following theorem:

Theorem 4 *If* $T = \Theta(\log n / b_1)$ *then* $\Pr[\min_{i \in I}\{N_i(T)\} > \max_{j \notin I}\{N_j(T)\}] = 1 - 1/f(n)$.

where $f(n)$ is any positive polynomial in n.

Despite the similarities between the above result and Theorem 2, the effort required in the multiple initiator case is potentially much larger than in the single initiator case. Note that b_1 in the theorem above

represents the least active initiator, while all paths in the single initiator case are constructed by the same node. For example, if all m initiators in I were identical then b_1 in the theorem above would be about m times smaller than in the single initiator case. Thus, the attacker would need around m times more paths (in total) in order to identify the set I with high probability. This is precisely what occurs in the case of the uniform path construction mechanism described in Section 3.2.1. In that system, $b_1 = \Theta(c/(mn))$, and thus, $T = \Theta(mn \log n/c)$. Note that this is m times larger than the single initiator case.

It may be unreasonable to assume that the attacker knows the number of initiators in the system. Identifying the set I in this case requires a different criterion. As before, we would like to establish a criterion that will succeed in identifying I given a large enough number of paths constructed.

Before we describe a new criterion, we introduce a new event of interest to the attacker. Let $\tilde{b_i}$ denote the probability node $i, 1 \leq i \leq n$ is a predecessor to an attacker assuming node i is not an initiator. Thus, $\tilde{b_i} = b_i$ if $i \notin I$. However, if node i is an initiator then $\tilde{b_i}$ is the probability that node i is a predecessor to an attacker assuming $r_i = 0$. Thus, $\tilde{b_i}$ for $i \in I$ is the probability node i is a predecessor if i was not an initiator. Note that $\tilde{b_i} < b_i$ if $i \in I$. We will assume that the attacker knows $\tilde{b_i}$ for all $1 \leq i \leq n$. This is reasonable for some anonymity systems where the attacker can trivially compute $\tilde{b_i}$. In fact, this is indeed the case for the generalized path construction mechanism presented in Section 3.2.

The attacker will classify node i as an initiator if the relative frequency observed by its counter deviates too much from $\tilde{b_i}$. More formally, the attacker classifies as initiators all nodes for which $N_i(T)/T - \tilde{b_i} > \varepsilon$, where ε is a small positive constant.

The identification of set I using this criterion is correct if all initiators are classified as initiators and no non-initiator is classified as an initiator. More formally, the attacker is correct if $\forall_{i \in I}\, N_i(T)/T - \tilde{b_i} > \varepsilon$ and $\forall_{j \notin I}\, N_j(T)/T - \tilde{b_j} < \varepsilon$, after a total of T paths have been constructed by the initiators.

Given this formulation, it is necessary and sufficient for the following condition to hold in order for the attack to succeed:

$$\forall_{i \in I}\, b_i - \tilde{b_i} > \varepsilon \tag{8}$$

Note that ε can always be chosen small enough such that this condition is satisfied. In particular, if $\varepsilon < \min_{i \in I}\{b_i - \tilde{b_i}\}$ then the condition above always holds. Although this maximum value for ε is not known a priori, the attacker can simply choose an arbitrarily small value. Choosing any ε within this range ensures that all initiators are correctly classified. However, as we will shortly demonstrate, the effort required to correctly classify the set I will strongly depend on the choice of ε.

It is also possible to define the goal of the attacker differently. Since the range of valid ε is unknown to the attacker and choosing an arbitrarily small ε may require too much effort (as we will soon observe), an attacker may find it sufficient to identify only a subset of I. In this case, the attacker initially chooses a value for ε. All initiators for which $b_i - \tilde{b_i} > \varepsilon$ can potentially be correctly identified. We consider the attack successful even if some initiators, for which $b_i - \tilde{b_i} < \varepsilon$, are not identified. Note that under this definition, the success of an attack no longer depends on the value of ε, as a successful attack is possible under any choice of ε (if ε is too large, then no initiator will be identified). Although this interpretation is advantageous from the perspective of the attacker, we will focus on criteria that attempt to identify the entire set I.

Condition (8) says nothing about the probabilities concerning the non-initiators. In fact, an explicit condition is not necessary, as a non-initiator node will not be incorrectly classified given a sufficient number

of paths. This occurs because $\forall_{j\notin I}$, $N_j(T)/T$ converges to $\tilde{b_j}$ as T becomes large. The following theorem states that condition (8) is necessary and sufficient for correct classification of the set I.

Theorem 5 *If condition (8) holds and* $\varepsilon < \min_{i\in I}\{b_i - \tilde{b_i}\}$, *then* $\Pr[\lim_{T\to\infty} \forall_{i\in I}\, N_i(T)/T - \tilde{b_i} > \varepsilon] = 1$ *and* $\Pr[\lim_{T\to\infty} \forall_{j\notin I}\, N_j(T)/T - \tilde{b_j} > \varepsilon] = 0$

Proof: From the strong law of large numbers we have that for all i, $\Pr[\lim_{T\to\infty} N_i(T)/T = b_i] = 1$. Condition (8) states that $\forall_{i\in I}\, b_i - \tilde{b_i} > \varepsilon$. From this, we conclude that $\Pr[\lim_{T\to\infty} \forall_{i\in I}\, N_i(T)/T - \tilde{b_i} > \varepsilon] = 1$. Now since $\forall_{j\notin I}\, b_j = \tilde{b_j}$, we can conclude that for any positive ε, $\Pr[\lim_{T\to\infty} \forall_{j\notin I}\, N_j(T)/T - \tilde{b_j} > \varepsilon] = 0$ □

We can also characterize the effort required by the attacker to identify the set I with high probability under this criterion. In particular, we would like obtain T such that $\Pr[\forall_{i\in I}\, \{N_i(T)/T - \tilde{b_i} > \varepsilon]$ is sufficiently large and $\Pr[\forall_{j\notin I}\, N_j(T)/T - \tilde{b_j} > \varepsilon]$ is sufficiently small. By applying Hoeffding's result we can obtain an analytical expression for each of these probabilities.

We start by evaluating the first probability, which concerns identifying the initiators. Let $k = \arg\min_{i\in I}\{b_i - \tilde{b_i}\}$. Thus, initiator k is the most likely initiator to be misclassified by the attacker, given that T paths were constructed. Under the assumption that n is large, we have the following theorem.

Theorem 6 *If* $T = \Theta(\log n/(b_k - \tilde{b_k} - \varepsilon)$ *then* $\Pr[\forall_{i\in I}\, N_i(T)/T - \tilde{b_i} > \varepsilon] = 1 - 1/f(n)$.

where $f(n)$ is any positive polynomial in n. The proof is found in the appendix.

The above result depends on b_k, $\tilde{b_k}$ and ε. Note that by decreasing ε, the effort required by the attacker decreases. Although this may appear counter-intuitive, this is actually intuitive if one carefully considers the probability in question. Thus, from the perspective of this probability, the least effort for an attacker is achieved when ε is set to zero. The effort required also depends on the difference $b_k - \tilde{b_k}$, with a larger difference requiring less effort.

We continue to evaluate the effort related to the second probability, which concerns the non-initiator nodes. Let $l = \arg\max_{j\notin I}\{b_j\}$. Thus, node l is the most likely non-initiator to be classified as an initiator by the attacker. If we assume n is large we can state the following theorem.

Theorem 7 *If* $T = \Theta(b_l\ \log n/\varepsilon^2$ *then* $\Pr[\forall_{j\notin I}\, N_j(T)/T - b_j < \varepsilon] = 1 - 1/f(n)$.

where $f(n)$ is any positive polynomial in n. The proof is found in the appendix.

Note that, contrary to Theorem 6, the effort required by Theorem 7 increases as ε decreases. In fact, at ε equal to zero the effort blows up to infinity. Thus, from the perspective of not misclassifying a non-initiator, an attacker would prefer a larger ε. The effort required by Theorem 7 is also proportional to b_l. Intuitively, if a non-initiator node precedes an attacker with smaller probability, then effort required by the attacker is reduced.

The above two theorems characterize the effort required such that the attacker can identify the initiator set with high probability. Since an attacker should correctly classify all initiators and not misclassify any

non-initiator, the effort required is the maximum between the result given by Theorems 6 and 7. Thus, $T = \Theta(\max(\log n/(b_k - \tilde{b_k} - \varepsilon), b_l \log n/\varepsilon^2))$ in order for this attack to succeed with high probability.

We can obtain the optimal value for ε, for which the effort required is minimized. Since the results for the effort in Theorems 6 and 7 are strictly monotonic in ε, we can obtain the optimal by equating these two equations and solving for ε. Thus, we have:

$$\varepsilon^* = \Theta(\sqrt{b_k b_l}) \tag{9}$$

The tradeoff between ε and the required effort will be illustrated in the numerical analysis in Section 6.

Although for some systems it may be possible for the attacker to compute $\tilde{b_i}$ a priori for all i, in general, this information may not be available to the attacker. In this case, can the attacker always correctly reveal the set of initiators, given a large enough number of paths constructed? As we show next, this is still possible under some circumstances.

Consider the following criteria used by the attacker to determine the set of initiators. The attacker chooses the minimum counter value (divided by T) as the base for comparison. Any node with a counter (divided by T) above the minimum plus a fixed quantity is classified as an initiator. More formally, let $x = \min_{1 \le i \le n}\{N_i(T)/T\}$ and ε a small positive constant. Node i is classified as an initiator if $N_i(T)/T - x > \varepsilon$. Note that the attacker relies solely on its counters and no other information is known. This criteria will correctly classify the set I under the following necessary and sufficient conditions:

$$\min_{i \in I}\{b_i\} - \min_{j \notin I}\{b_j\} > \varepsilon \tag{10}$$

$$\max_{j \notin I}\{b_j\} - \min_{j \notin I}\{b_j\} < \varepsilon \tag{11}$$

Thus, in an anonymity systems where the above two conditions hold, the attacker will correctly identify the set I given a sufficiently large number of paths constructed. This is stated in the stated in the following theorem.

Theorem 8 *If conditions (10) and (11) hold true and* ε *is chosen accordingly then* $\Pr[\lim_{T\to\infty} \min_{i\in I}\{\frac{N_i(T)}{T}\} - \min_{j\notin I}\{\frac{N_j(T)}{T}\} > \varepsilon)] = 1$ *and* $\Pr[\lim_{T\to\infty} \max_{j\notin I}\{\frac{N_j(T)}{T}\} - \min_{j\notin I}\{\frac{N_j(T)}{T}\} > \varepsilon] = 0$.

Proof: From the strong law of large numbers we have that for all i, $\Pr[\lim_{T\to\infty} N_i(T)/T = b_i] = 1$. Condition (10) provides that $\min_{i\in I}\{b_i\} - \min_{j\notin I}\{b_j\} > \varepsilon$. Thus, we can conclude that $\Pr[\lim_{T\to\infty} \min_{i\in I}\{\frac{N_i(T)}{T}\} - \min_{j\notin I}\{\frac{N_j(T)}{T}\} > \varepsilon] = 1$. Now condition (11) provides that $\max_{j\notin I}\{b_j\} - \min_{j\notin I}\{b_j\} < \varepsilon$. Thus, we can conclude that $\Pr[\lim_{T\to\infty} \max_{j\notin I}\{\frac{N_j(T)}{T}\} - \min_{j\notin I}\{\frac{N_j(T)}{T}\} > \varepsilon] = 0$ □

Note that conditions (10) and (11) are more stringent than the previously defined conditions. In fact, unlike condition (8), ε above cannot be chosen arbitrarily small, since doing so will violate condition (11). Even in a system where the above two conditions hold for some value of ε, the attacker would have to know this value in order to correctly classify the set I. This requirement seems unreasonable without any prior knowledge of the system.

However, if a uniform path construction mechanism is used, such as the one described in Section 3.2.1, the above two conditions are satisfied. Recall that in such systems, all non-initiator nodes have the same

probability of preceding an attacker attacker node, and thus, $\max_{j \notin I}\{b_j\} = \min_{j \notin I}\{b_j\}$. In this case, the attacker can choose any small positive value for ε without violating the condition (11) while ensuring condition (10) holds. In fact, condition (11) is no longer needed, as it is always satisfied by any $\varepsilon > 0$, and condition (10) becomes identical to condition (8), since $\tilde{b_j} = b_j$ for $j \notin I$.

Thus, in this uniform scenario, the effort required by the attacker in order to succeed with high probability is the same as the result provided by Theorems 6 and 7. Using equation (4), which gives the values for $\{b_i\}$ in the uniform case, we have that $\min_{i \in I}\{b_i - \tilde{b_i}\} = \Theta(c/(mn))$, where m is the number of initiators in the system, and $\max_{j \notin I}\{b_j\} = \Theta(c/n^2)$, assuming n is large and $c < n$. Thus, the effort required is given by $T = \Theta(\max(mn\log n/(c - mn\varepsilon), c\log n/(n\varepsilon)^2))$.

The optimal value for ε, which is given by equation (9), in this case is given by $\Theta(c/n\sqrt{1/(nm)})$. If we assume ε to be equal to the optimal, then the required effort is simply given by $T = \Theta(mn\log n/c)$. We observe that the effort required increases linearly in the number of initiators, while decreases linearly with the number of attacker nodes. Also, if $m = 1$ the effort required to identify the sole initiator using this criteria is the identical to the effort required under the single initiator criteria presented in Section 4. Interestingly, the effort required using this criteria is also equal to the case where the attacker knows the number of initiators in the system. Thus, asymptotically, the attacker does not need less effort to identify the set I if it has knowledge of the number of initiators present in the system.

6 Numerical Case Study

We conduct a numerical evaluation of the predecessor attack to illustrate that an attacker can correctly identify the initiator set. We investigate two of the criteria that were presented in the previous section and discuss some tradeoffs. In the evaluation that follows, we consider the uniform path construction mechanism described in Section 3.2.1 with the following parameters: $n = 80, c = 20, p_r = 0.25$. The numerical results were obtained through repeated simulations of the anonymity system.

We first assume that the number of initiators is known to the attacker. Thus, the criteria for determining the initiator set is simply to select the $m = |I|$ larger counter values. In the case of ties between the m-th and $(m+1)$-th larger counter values, we stipulate that the attacker fails to identify the set correctly. Figure 3 illustrates the probability of correctly identifying the initiator set as a function of the total number of paths created. As expected, the attacker can correctly classify the initiator set given sufficient rounds. Note that as the number of initiators increases, more effort is required from the attacker.

If the attacker does not know the number of initiators in the system, then a different criteria to determine the initiator set must be used. Let's also assume the attacker does not know $\{\tilde{b_i}\}$. In this case, we use the criteria that chooses the minimum counter value as the base of comparison, as discussed in Section 5. Recall that using this criteria, a node i is classified as an initiator if $N_i(T)/T - x > \varepsilon$, where $x = \min_{1 \le i \le n}\{N_i(T)/T\}$ and ε is a small positive constant chosen by the attacker. Figure 4 illustrates the probability of correctly identifying the initiator under this criteria. Again, in all cases the attacker is capable of correctly identifying the initiator set.

Interestingly, note that the difference between the effort required in the two criteria considered, as illustrated by Figures 3 and 4, does not differ much, although in the first case the attacker has much more information about the system. In fact, the asymptotic effort required for a successful attack in both cases are

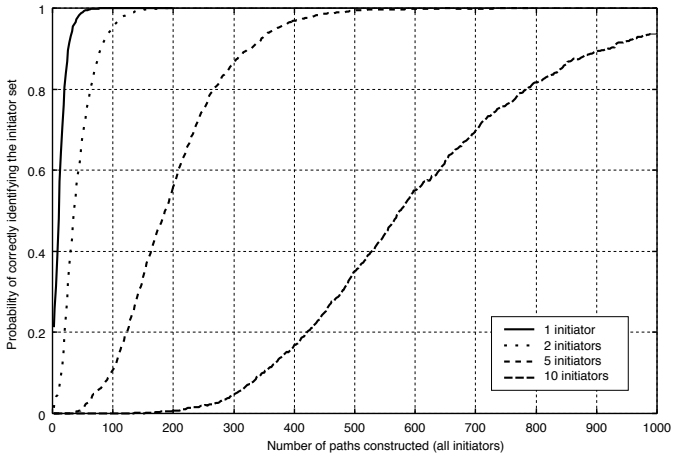

Figure 3: Probability of correctly identifying the initiator set – attacker knows the number of initiators in the system.

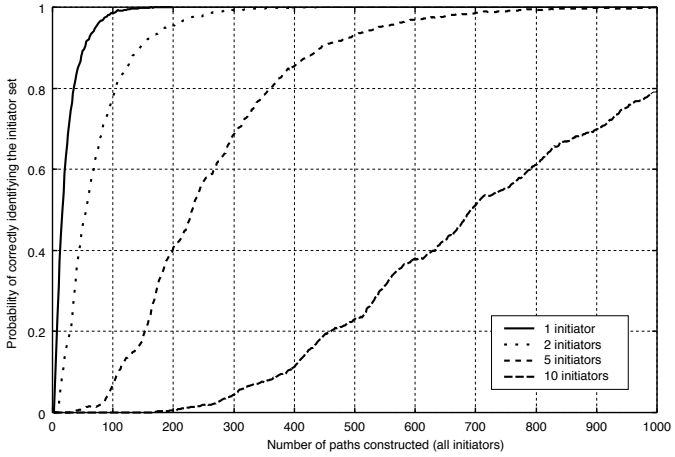

Figure 4: Probability of correctly identifying the initiator set – attacker uses minimum counter value as basis for comparison.

identical, as established in Section 5. In practice, the actual extra effort required by the attacker when the number of initiators in the system is not known is quite small.

A much more serious concern for the attacker is the choice of ε. Recall that Theorem 5 states that any ε small enough will allow the attacker to correctly classify all initiators. However, as shown by the result of Theorems 6 and 7, the actual effort required is very sensitive to this parameter. We will illustrate the impact of ε by considering an attack a system with 2 initiators. Recall that ε should be chosen such that condition (10) is satisfied, that is, $\varepsilon < \min_{i \in I}\{b_i - \tilde{b}_i\}$, which in this example yields $\varepsilon < 0.2$. Figure 5 illustrates the effort required when $\varepsilon = 0.01, 0.1, 0.19$. As expected, the results show clearly that the effort required is not monotonic in ε. Intuitively, if ε is too small, then the attacker is likely to classify a non-initiator as being an initiator. Similarly, if ε is too large (but still within the minimum), then the attacker is likely to classify an initiator as a non-initiator. However, no matter the choice of ε the attacker always correctly identify the set I given that a sufficiently large number of paths are constructed. From our results, we observe that a moderate ε yields the best performance.

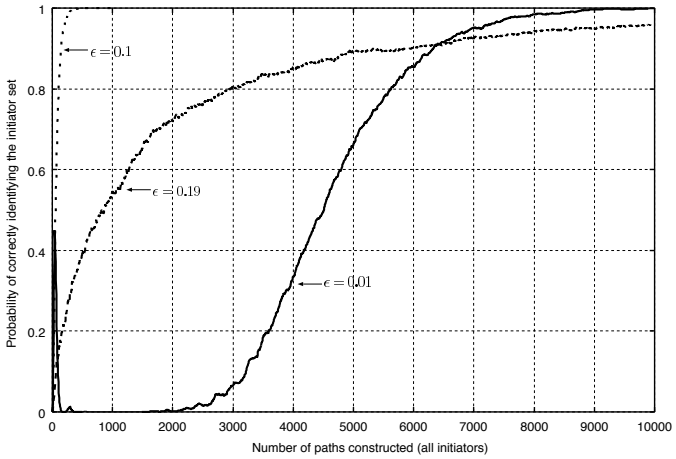

Figure 5: Impact of ε on the probability of correctly identifying the initiator.

The range of valid ε is most likely not known to the attacker. In particular, even in the uniform scenario this information is not available to the attacker. A good choice of ε seems critical for an efficient successful attack. The results shown in Figure 4 were obtained using ε equal to the midpoint in its valid range. Although this can be seen as a potential limitation, the attacker can attempt to dynamically select values for ε, refining the set of initiators as more paths are constructed.

7 Summary

Understanding the capabilities and limitations of an attack is an important step towards designing more secure anonymity systems. Our results show that the predecessor attack is very robust and can succeed even if the resources dispensed by the attacker are limited. In particular, we have shown that practical anonymity systems where uniform path construction mechanisms are used, are always subject to a successful attack. Moreover, our theoretical results show that in this case the effort required is $\Theta(mn \log n/c)$, which grows linearly with the number of initiators. In fact, the actual effort needed for a successful attack can be easily attained, as illustrated by the numerical analysis of a simulated anonymity system. An interesting observation, shown by the analysis and corroborated by the numerical analysis, is that knowing the number of initiators present in the system does not reduce the effort required by the attacker to identify them.

A natural question is how should we design a practical anonymity systems that is not subject to the predecessor attack? In general, this seems difficult. Clearly, the use of uniform path construction mechanisms gives the attacker an advantage, as this enables the correct classification of the set of initiators. Thus, anonymity systems should be designed such that the conditions required for success never hold. However, without prior knowledge of which nodes are controlled by the attacker, this seems hard to accomplish. A simple idea is to have initiators choose random path construction distributions that are not uniform. This would be an attempt to trick the attacker in thinking the system has a large number of initiators. However, future work is needed to understand this idea in detail and to consider other approaches in this direction.

8 Proof of Theorems Determining the Effort Required

All proofs presented below use Hoeffding's result for the asymptotic tail probability of a generic multinomial distribution [6]. Thus, it is helpful and informative that we first present a tailored version of this result.

Consider a set of $n+1$ distinct events each occurring probability $b_i, 1 \leq i \leq n+1$. Let $b = (b_1, \ldots, b_{n+1})$. Consider a sequence of T independent trials and random vector $N(T) = (N_1(T), N_2(T), \ldots, N_{n+1}(T))$, such that $N_i(T)$ is the number of trials that resulted in the i-th event. $N(T)$ is given by the multinomial distribution. Let $x = (x_1, \ldots, x_n, x_{n+1})$ such that $x_i = N_i(T)/T$.

Define a region of interest, M, over the space defined by all possible values for x. Equation (2.9) of [6] yields

$$\Pr[M|b] \quad = \quad \exp\{-TI(M,b) + O(\log T)\} \tag{12}$$

where

$$I(M,b) \quad = \quad \inf_{x \in M} \sum_{i=1}^{n+1} x_i \log(x_i/b_i) \tag{13}$$

Note that $I(M,b)$ is usually referred to as the rate function, as it provides the rate of the exponential decay of the probability, as T increases. For the cases we consider below, the infimum in equation (13) can be mapped into the following optimization problem:

$$\min_x \quad \sum_{i=1}^{n+1} x_i \log(x_i/b_i) \tag{14}$$

$$\text{subject to} \quad \sum_{i=1}^{n+1} x_i = 1 \tag{15}$$

$$x_i > 0, \quad 1 \leq i \leq n+1 \tag{16}$$

$$x \in M \tag{17}$$

Again for the cases we consider, this optimization problem can then be solved analytically through the use of Lagrange multipliers.

Let $f(n)$ be any positive polynomial function in n. The objective is to obtain T such that $\Pr[M|b] = 1/f(n)$. Applying equation (12), we have:

$$\Pr[N_1(T) < \max_{1<j\leq k} \{N_j(T)\}] = 1/f(n) \tag{18}$$

$$\Rightarrow \quad \exp\{-TI(M,b) + O(\log T)\} = 1/\log f(n) \tag{19}$$

$$\Rightarrow \quad -TI(M,b) + O(\log T) = -\log f(n) \tag{20}$$

$$\Rightarrow \quad T = \Theta(\log n / I(M,b)) \tag{21}$$

This asymptotic result above can be proved formally in the case T and n are both assumed to be large.

Theorem 2: If $T = \Theta(\log n/b_1)$ *then* $\Pr[N_1(T) > \max_{1<j\leq n}\{N_j(T)\}] = 1 - 1/f(n)$.

Proof: We will work with $\Pr[N_1(T) < \max_{1<j\leq n}\{N_j(T)\}]$, as this probability converges to zero as T approaches infinity (see Theorem 1). Define the region of interest $M = \{x | x_1 < \max_{1<j\leq n}\{x_j\}\}$.

Note that in order for a vector x to be included in M, it is sufficient for $x_1 < x_j$ for some $1 < j \leq n$. Moreover, due to the monotonicity of the objective function of $I(M,b)$ on both parameters x and b (see [6]), for a fixed counter j, the minimum will be achieved when $x_1 = x_j$. Considering the parameter vector b, the minimum over all possible values for j is achieved with the variable that has the largest value. Without loss of generality assume that $b_2 = max_{2<i\leq n}\{b_i\}$. Therefore, the minimum is obtained when $x_1 = x_2$. Note that if we assume b_2 is strictly greater than $b_i, 2 < i \leq n$, the optimization problem has unique solution. Otherwise, if k other values of b_i are identical to b_2, then the optimization problem has k solutions, each obtained when equating x_1 to a particular x_j. Thus, without loss of generality, we substitute constraint (17) with $x_1 = x_2$, which allows for an analytical solution through the use of Lagrange multipliers. In particular, we have

$$I(M,b) \quad = \quad -\log(1+2\sqrt{b_1 b_2} - (b_1+b_2)) \tag{22}$$

To continue with the analysis and simplify the above equation, we let $b_2 = o(b_1)$, which satisfies condition (5). Recall that $f(n) = o(g(n))$ iff $\lim_{n\to\infty} f(n)/g(n) = 0$. Thus, assuming n is large and $c < n$, $I(M,b) = \Theta(b_1)$. In the previous, the Taylor expansion of $\log(1-x)$ was used, which provides for $\log(1-x) = -x + \Theta(x^2)$ when x is close to zero. Applying this result to equation (21) we have that $T = \Theta(\log n/b_1)$. □

Theorem 6: If $T = \Theta(\log n/(b_k - \tilde{b_k} - \varepsilon)$ *then* $\Pr[\forall_{i\in I}\, N_i(T)/T - \tilde{b_i} > \varepsilon] = 1 - 1/f(n)$.

Proof: We will work with $\Pr[\exists_{i\in I}\, N_i(T)/T - \tilde{b_i} < \varepsilon)]$, as this probability converges to zero as T approaches infinity (see Theorem 5). Define the region of interest $M = \{x|\exists_{i\in I}\, x_i - \tilde{b_i} < \varepsilon)\}$.

Note that in order for a vector x to be included in M, it is sufficient if some $x_i, i \in I$ is smaller than $\varepsilon + \tilde{b_i}$. Moreover, due to monotonicity of the objective function of $I(M,b)$ on both parameters x and b (see [6]), for a fixed choice of i, the minimum will be achieved when $x_i = \varepsilon + \tilde{b_i}$. Let $k = \arg\min_{i\in I}\{b_i - \tilde{b_i}\}$ Thus, the minimum over all possible $i \in I$, is achieved with node k, and in particular when $x_k = \varepsilon + \tilde{b_k}$. Thus, without loss of generality, we substitute constraint (17) with this previous condition, which then allow us to obtain an analytical solution through the use of Lagrange multipliers. In particular, we have

$$\begin{aligned} I(M,b) \quad &= \quad (\tilde{b_k} + \varepsilon)\log(\frac{\tilde{b_k}+\varepsilon}{b_k}) + \\ &\qquad (1-\tilde{b_k}-\varepsilon)\log(\frac{1-\tilde{b_k}-\varepsilon}{1-b_k}) \end{aligned} \tag{23}$$

To continue with the analysis and simplify the above equation we assume that b_k is small, which is expected when n is large. Thus, we have $I(M,b) = \Theta(b_k - \tilde{b_k} - \varepsilon)$. In this derivation we used the Taylor expansion of $\log(1-x)$, which provides for $\log(1-x) = -x + \Theta(x^2)$ when x is close to zero. Applying this result to equation (21) we have that $T = \Theta(\log n/(b_k - \tilde{b_k} - \varepsilon)$. □

Theorem 7: If $T = \Theta(b_l\ \log n/\varepsilon^2$ *then* $\Pr[\forall_{j\notin I}\, N_j(T)/T - b_j < \varepsilon] = 1 - 1/f(n)$.

Proof: Recall that $\tilde{b_j} = b_j$ for all $j \notin I$. Also, note that $\Pr[\exists_{j\notin I}\, N_j(T)/T - \tilde{b_j} > \varepsilon)]$ converges to zero as T approaches infinity (see Theorem 5). Define the region of interest $M = \{x|\exists_{j\notin I}\, x_j - b_j > \varepsilon\}$.

Note that in order for a vector x to be included in M, it is sufficient if some $x_j, j \notin I$ is larger than $\varepsilon + b_j$. Moreover, due to monotonicity of the objective function of $I(M,b)$ on both parameters x and b (see [6]),

for a fixed choice of j, the minimum will be achieved when $x_j = \varepsilon + b_j$. Let $l = \arg\max_{j \notin I}\{b_j\}$. Thus, the minimum over all possible $j \notin I$, is achieved with node l, and in particular, when $x_l = \varepsilon + b + l$. Thus, without loss of generality, we substitute constraint (17) with the previous condition, which then allow us to obtain an analytical solution through the use of Lagrange multipliers. In particular, we have

$$\begin{aligned} I(M,b) &= (b_l+\varepsilon)\log(1+\frac{\varepsilon}{b_l}) + \\ &\quad (1-b_l-\varepsilon)\log(1-\frac{\varepsilon}{1-b_l}) \end{aligned} \tag{24}$$

To continue with the analysis and simplify the above equation, we assume b_l is small which is likely to be the case when n is large. Thus, we obtain $I(M,b) = \Theta(\varepsilon^2/b_l)$. In the previous, the again used the Taylor expansion of $\log(1-x)$. Applying this result to equation (21) we have that $T = \Theta(b_l \log n/\varepsilon^2)$. □

References

[1] Adam Back, Ulf Mller, and Anton Stiglic. Traffic analysis attacks and trade-offs in anonymity providing systems. In *4th International Workshop on Information Hiding*, volume 2137 of *Lecture Notes in Computer Science*. Springer-Verlag, 2001.

[2] Krista Bennett and Christian Grothoff. Gap – practical anonymous networking. *Proc. Workshop on Privacy Enhancing Technologies (PET)*, 2003.

[3] John R. Douceur. The sybil attack. In *Proc. of the First International Workshop on Peer-to-Peer Systems (IPTPS '02)*, Cambridge, MA, March 2002.

[4] Michael J. Freedman and Robert Morris. Tarzan: A peer-to-peer anonymizing network layer. In *Proc. of the 9th ACM Conference on Computer and Communications Security (CCS 2002)*, Washington, DC, November 2002.

[5] Yong Guan, Xinwen Fu, Riccardo Bettati, and Wei Zhao. An optimal strategy for anonymous communication protocols. In *Proc. 22nd IEEE International Conference on Distributed Computing Systems (ICDCS 2002)*, Jul 2002.

[6] Wassily Hoeffding. Asymptotically optimal tests for multinomial distributions. *Annals of Mathematical Statistics*, 36(2):369 – 401, Apr 1965.

[7] Brian Levine, Michael Reiter, Chenxi Wang, and Matthew Wright. Timing analysis in low-latency mix systems. In *Proc. of the 8th International Conference on Financial Cryptography*, Feb 2004.

[8] Brian N. Levine and Clay Shields. Hordes: A protocol for anonymous communication over the internet. *ACM Journal of Computer Security*, 10(3), 2002.

[9] Michael G. Reed, Paul F. Syverson, and David M. Goldschlag. Anonymous connections and onion routing. *IEEE Journal on Selected Areas in Communication Special Issue on Copyright and Privacy Protection*, 16(4):482–494, 1998.

[10] Michael K. Reiter and Aviel D. Rubin. Crowds: anonymity for Web transactions. *ACM Transactions on Information and System Security*, 1(1):66–92, 1998.

[11] Marc Rennhard and Bernhard Plattner. Introducing morphmix: Peer-to-peer based anonymous internet usage with collusion detection. In *Proc. of the Workshop on Privacy in the Electronic Society (WPES)*, Washington, DC, November 2002.

[12] Andrei Serjantov and George Danezis. Towards an information theoretic metric for anonymity. *Proc. Workshop on Privacy Enhancing Technologies (PET)*, 2482, 2002.

[13] Matt Wright, Micah Adler, Brian N. Levine, and Clay Shields. An analysis of the degradation of anonymous protocols. In *Proc. ISOC Network and Distributed System Security Symposium (NDSS 2002)*, Feb 2002.