Abstract

Design and Analysis of Efficient Anonymous-Communication Protocols

Aaron Johnson

2009

Research into protocols for anonymous communication in networks has yielded many designs and a wide range of functionality and performance options. Of these, only a few have been successful in practice. The experience in fielded systems emphasizes the following criteria for designing anonymous-communication protocols: communication functionality, anonymity, latency, and message complexity. Therefore, in this thesis we explore understanding and improving the performance of anonymous-communication protocols according to these criteria. To provide high confidence in our analysis against a malicious adversary, we formally model such protocols and rigorously analyze them according to precise definitions.

The first subject of our analysis is the onion-routing protocol. Onion routing is a successful protocol for anonymous communication in widespread public use. However, its anonymity has not received much rigorous analysis. Therefore, in the first half of the thesis, we model the protocol and its environment, and we analyze the resulting anonymity properties.

Our results show that onion routing provides unsatisfactory anonymity against a realistic adversary that controls some fraction of the network. In particular, the protocol faces an upper limit on anonymity that depends only on the size of the adversary. Therefore, in the last half of the thesis, we consider two ways to get past this limit: trust information and a new protocol design.

We first suppose that users have external trust information about the network that helps them avoid parts that are controlled by the adversary. Then we consider using this information to improve the anonymity provided by onion routing. Under a model of trust, we come up with practical and provable improvements.

Next we consider a new protocol that avoids a major weakness in onion routing: timing attacks. The protocol uses explicit timing and redundancy as its key techniques, and we prove that it can provide arbitrarily high anonymity. Finally, this protocol requires adding delays to smooth out variable latency in the underlying network. We estimate the magnitudes of these delays by performing measurements on a live anonymity network. Our results suggest that the added delays likely result in a small constant-factor increase over onion routing.

Design and Analysis of Efficient

Anonymous-Communication Protocols

A Dissertation
Presented to the Faculty of the Graduate School
of
Yale University
in Candidacy for the Degree of
Doctor of Philosophy

by
Aaron Johnson

Dissertation Director: Joan Feigenbaum

December, 2009

Contents

List of Figures

Acknowledgements

Many thanks to my advisor, Joan Feigenbaum, for giving me the freedom to choose my own route and helping me focus and follow it to the end. Thanks to my colleague Paul Syverson. Your ideas are the inspiration behind much of this work, and your experience and energy have been invaluable. Thanks to my friends for the pleasure of sharing this experience together. Finally, thanks to my family for always having a place for me at home.

Chapter 1

Introduction

Every method for human communication seems to be accompanied by a way to use it to communicate anonymously. Callers can block their phone numbers from being sent to the recipients, or they can use public telephones to hide their identity even from the telephone company. Anonymous letters are easily sent by omitting or falsifying the return address. Anonymous messages can even be sent in person using disguises or third parties.

The Internet is no exception. Indeed, the Internet by nature shields users behind their computers and allows them to communicate regardless of location, only helping users to communicate without revealing an identity. And users take advantage of this. People use the Internet to create alternate identities in online communities. They use the Internet to anonymously publish complaints about their governments and their jobs. Cybercriminals anonymously deface websites and exchange copyrighted media.

But the Internet is not as anonymous as it might appear. Every message sent directly from a computer is tagged with an address that allows the recipient to send a response. This address, perhaps combined with some information from the service provider, can be used to identify the source. Communication that happens over higher-level protocols, such as email, typically contains information that allows a message to be tracked back to the sender.

This opens up a basic question in computer networking: how can we provide anonymous communication over the network? Research into this question has identified a wide range of solutions. One one end of the range, there is the simple solution of redirecting messages through a proxy. This provides only very weak anonymity in that it requires users to trust the proxy and assumes that communication over the network

generally cannot be observed. It has the advantage of being simple and efficient. At the other end of the range is the "dining cryptographers" protocol [13]. This protocol provides perfect anonymity to senders and recipients. However, it requires every pair of users to communicate, which results in a total of $\Omega(n^2)$ messages, where n is the number of users. This overhead is very high when used on a large network such as the Internet.

There have been several deployments of anonymity systems for general use that provide information about how well these solutions work in practice [36]. Single-hop anonymizing proxies provide real-time relaying of general Internet traffic. Lists of free single-hop anonymizing proxies are easily available online, and there are several companies that provide that service for a fee. A more sophisticated class of system uses chains of relays to provide better anonymity, while still maintaining latencies that are low enough to support many types of Internet protocols [9, 5, 23]. These systems provide also real-time anonymous communication of Internet traffic, although with some performance loss. A third class of system is *anonymous remailers* [61, 56]. Anonymous remailers allow users to send email anonymously. They can achieve high guarantees of anonymity but require higher delays that make them unsuitable for other applications.

A lesson from these systems is that efficiency matters [22]. The popularity of these systems decreases as their resource requirements increase. Single-hop proxies require the network to send one extra message and add the time it takes to route traffic through the proxy to the latency. Their use is widespread and are the underlying technology for the few businesses that are successful in providing this service (for example, The Anonymizer [2]). Relay chains add messages and latency that is proportional to the length of the chain. These resource requirements were, for example, a main factor in the failure of the for-profit model of the Freedom Network [9]. As mentioned, anonymous remailers add significant delay, and this has limited their popularity.

Efficiency is generally gained at the expense of anonymity. For example, when there aren't too many extra messages nor much added latency, the users and destinations that are active at the same time are likely to be communicating with one another. Therefore, if the adversary can observe communication activity, he can use it to infer who is talking to whom. As another example, the adversary can slow down processing in parts of the network by overloading them with messages, and then he can observe how that affects the timing of messages in the protocol. In efficient protocols, users don't interact with every other agent, and so this may affect some users but not others, which can allow the adversary to figure out more about the execution of the protocol.

As is evident in these examples, a common weakness in efficient protocols is that the timing of messages (as opposed to, say, their number or content) can reveal private information about communication. This complicates the analysis of these protocols, as it tends to fall outside the models considered in more well-studied cryptographic protocols such as those for encrypted or authenticated communication. Due in part to this, the anonymity these protocols provide is often not well-understood.

In this thesis we analyze these timing issues, among others, as we evaluate the performance of efficient protocols. We first consider the anonymity provided by onion routing [40], which is a very popular protocol that uses chains of relays to provide anonymous communication. Our results show that, due primarily to the ability of the adversary to correlate timing patterns in traffic coming in from users with such patterns going out to destinations, the anonymity provided by onion routing has a worst case that is significantly worse than less rigorous analysis has suggested (for example, [80]). Even in the best case, there is a hard upper limit, depending on the size of the adversary, to how much anonymity is provided.

Therefore, we consider designing protocols with better anonymity that maintain good efficiency. Specifically, we would like to maintain performance competitive with onion routing, because it has demonstrated that its performance is acceptable in practice. First, we consider modifying onion routing to use outside trust information about which agents are likely to be adversarial. If an agent has a good idea about which routers are trustworthy, he can prefer those routers when forwarding his traffic through the network. We formalize a notion of trust and consider how use it to optimally choose those routers. Second, we consider ways to avoid entirely the timing attacks of the adversary. Such attacks rely on timing patterns in the protocol traffic created either by the users or the adversary. We enable the users to avoid creating such patterns, and give a protocol that prevents the adversary from creating them.

Our design and analysis is done rigorously within precise mathematical models. Analysis of anonymous-communication protocols, especially efficient protocols that are designed to be useful in practice, is often done informally. Such analyses can leave important parts of the network environment, such as the user behavior or adversary capabilities, somewhat ambiguous. We provide precise modeling of all aspects of the protocols we consider, and evaluate their performance according to clear definitions of our criteria. The basic model and definitions are given in Chapter 2; extensions are given in Chapters 6, 7, and 8.

1.1   Analyzing Onion routing

We start by modeling and analyzing *onion routing*, which uses chains of relays and layered encryption to anonymously deliver messages. Onion routing is widely used, and there are several similar protocols in use as well. Despite this, recent research demonstrates how much we have to learn about its properties [55, 10, 64, 8, 57, 43, 28]. The anonymity of the protocol is affected by the timing properties of the network, including timing delays on communication links, congestion at routers, and inaccuracies in local clocks. It also depends on user behavior, such as how users choose destinations and the pattern of their traffic. The nature of the adversary is a factor as well. He may control small parts of the network, may be able to make observations at many points in the network, or may be able to move around the network during the operation of the protocol. We construct a model of the protocol that makes some concrete choices on these questions, which we justify, when possible, by experience in practice. Then we obtain results about the properties of onion routing in this model, in particular about the anonymity it provides.

In onion routing, agents are divided into *users*, *routers*, and *destinations*. Users use the protocol to open connections for two-way communication with destinations. For a given connection, the user selects a sequence of routers, known as a *circuit*, that will be used to forward the user's traffic. The user establishes a circuit by first directly opening a circuit with the first router, and then iteratively extending the circuit by sending message over the existing circuit. Messages are encrypted with the key of each router in the circuit in the reverse order that the routers appear. For a message that is being sent from the user to the destination, after the message arrives at the ith router on a circuit, the router decrypts the message and sends the result to the $(i{+}1)$st router. For a message that is being sent from the destination to the user, after the message arrives at the ith router, the router encrypts the message and sends the result to the $(i-1)$st router. When the user is done with the circuit, he sends a command to the first router to destroy it, which gets relayed down the circuit. This scheme keeps a router from learning the identities of the agents in the circuit except for the one before and the one after. Or, at least, the message contents alone do not provide this information, as we have already mentioned several attacks that can allow the adversary to learn more.

This protocol is a practical scheme that provides some anonymity with reasonable performance. It adds latency and extra messages proportionally to the length of the path, which we can hold constant even as the network grows. It provides two-way, connection-based communication and does not require that the destination participate in the anonymity-network protocol. These features make it useful for anonymizing much of the communication that takes place over the Internet today, such as web browsing, chatting, and

remote login.

The basic ideas behind onion routing can be found in many fielded systems [9, 5, 23], although the term "onion routing" is often reserved exclusively for the original protocol and its direct successors [40, 79, 23]. It has achieved considerable popularity as the Tor network [23, 83]. Tor develops and distributes open-source client and server software for users and routers, respectively. The routers are run by volunteers. As of May 2009, there were over 1500 active routers in the network [85]. The usage level of the network is hard to measure exactly, because, to protect user privacy, Tor routers collect very little network data. It is estimated by Mittal and Borisov [60] and by Goldberg [38] that the network has hundreds of thousands of users and carries terabytes of traffic per day, and these estimates are substantiated by the data published by McCoy et al.[57].[1].

We formalize onion routing, using I/O automata, in Chapter 4. Then, in Chapter 5, we characterize the anonymity of the resulting system in a possibilistic sense, that is, we determine when it is possible, as far as the adversary can tell, that a user didn't perform a given action. Our results show that the adversary can only determine the identity of the agents on a circuit that he controls or is adjacent to. This implies that anonymity depends on the set of users for which the adversary controls the first router on the path and the set for which the adversary controls the last router, because controlling these positions are the only ways that the adversary can determine the user and destination of a circuit, respectively.

This is a weak notion of anonymity, in that, in a more realistic model in which the system operates probabilistically, certain outcomes may be more likely than others. Therefore, in Chapter 6, we add a probability distribution to executions in our model and analyze the *relationship anonymity* (that is, how well the adversary can tell that a user and destination are talking) in a probabilistic sense. Specifically, the metric that we use is the expected probability that the adversary assigns to the actual destination of a given user, and the lower this value is the more anonymous a user and destination are. The relationship anonymity of a given user u and destination d depends on the behavior of other users; if the other users are more likely to choose different destinations than u, for example, then their actions are less likely to be confused with the actions of u.

In fact, we show that there are two possible worst cases: all users except u always select the destination

---

[1]They observed 7571 unique users in one day at an entry node. As of May 2009, constituted approximately 20% of routers [85]. Every user chooses entry nodes randomly from a set of default size 3[84]. Therefore we can estimate the total number of unique users for the day as $7571 * 250/3 \approx 630917$. Also, they forwarded 709GB of traffic over four days, and report that the reported bandwidth of their server was in the top 5%. If the rest of the top 5% allowed exit traffic, that yields $709 * 5/4 = 886.25$GB/day

u is least likely to select (other than d), and all users except u always select d. Then we show that the worst cases are significantly worse, in that the adversary must control many fewer routers in the worst case to maintain the same anonymity as in a "typical" case. When there are many users and destinations, the anonymity achieved in a typical case (previously proposed by Shmatikov and Wang [76]) against an adversary that controls a fraction b of the routers is only achieved in the worst case when the adversary controls the smaller fraction b^2. We also show a lower bound on our metric that is constant in the number of agents, given that the relative size b of the adversary is constant. This result is consistent with the less formal analysis of Syverson et al. [80].

1.2 Improved Efficient Protocols

The limit on relationship anonymity in onion routing is not very satisfactory. Against an adversary controlling a fraction b of the routers, there is a non-negligible probability b^2 of having no anonymity, and this does not change with the size of the network nor are we able to trade some performance for anonymity. To understand what this barrier means in practice, consider Tor. Under the current size and usage of the Tor network, with only thirty routers we expect the adversary to compromise over three users per day and for a given user to be compromised at least once a month[2]. Even worse, these rates of compromise can be obtained with only two routers if the adversary can run the two highest-bandwidth routers[3]. Thus, the system provides almost no anonymity against a wide variety of realistic opponents, such as governments, ISPs, and criminals willing to purchase the use of botnets.

However, onion routing does provide desirable communication functionality with good efficiency, and this has contributed to its popularity over more secure protocols. Therefore, in Chapters 7 and 8 we consider new protocols to provide better anonymity while maintaining the desirable properties of onion routing.

In Chapter 7, we consider the case that users have some outside knowledge about which routers in the network are less likely to try to compromise anonymity. There are several possible sources of such *trust* knowledge, such as personal relationships with the operators of certain routers, knowledge about the security

---

[2]Assume the Tor network consists of 1500 routers [85]. An adversary that runs 30 routers has over a 50% chance of compromising 1 in 1750 *circuits* (the persistent connections through the network that users set up to route traffic). McCoy et al. [57] observed 7571 unique clients while running a router for one day. At this usage rate, our adversary, with 30 routers, would expect to compromise over 3 users per day. Also, currently in Tor users choose new circuits every 10 minutes. Therefore, any given user has a 50% chance of being compromised at least once in about 2 weeks of continuous use.

[3]As observed by Bauer et al. [4], Tor weights router selection by bandwidth to provide good performance. Thus an adversary need only control 2% of the bandwidth to achieve these success rates. McCoy et al. observed in [57] that the top 2% of routers transported about 50% of the traffic. In such a situation, supplying 2% of the bandwidth in Tor can be achieved by supplying just the two routers with the highest bandwidth.

of different platforms, and information about surveillance programs by governments. We formalize a notion of trust, and then analyze how users can improve the anonymity of onion routing by taking into account how trusted each router is when selecting circuits. Our results include a general path-selection algorithms, analysis of a natural approximation algorithm, and optimal strategies in the special case that all the routers are either "trusted" or "untrusted."

Using trust knowledge can help users avoid adversarial routers, but the extent to which it helps is limited by the amount of trust users have in the routers. In Chapter 8, we describe a protocol that provides good anonymity even without trust assumptions. The anonymity our protocol provides increases without bound as the number of users increases, and it has efficiency that is comparable to that of onion routing. The protocol deals with the fundamental problem in onion routing - the adversary can use timing to determine which routers are part of the same circuit. To prevent the adversary from manipulating the timing of the protocol, we include explicit timing commands with the messages and use redundancy to help prevent the adversary from blocking messages. Redundancy opens up more opportunities for the adversary to observe the communication, though, and balancing these two threats is the key contribution of our design. The explicit timing requires adding some delays that depend on the network, and so we also perform a measurement study of the timing properties of the Tor network to evaluate what kind of performance the system might achieve in practice. The results suggest that, for the majority of messages in the system, latency would be multiplied by at most a factor of two or three.

Chapter 2

Models and Definitions

2.1 Introduction

The first step towards rigorously analyzing and designing anonymity protocols is to build an appropriate model of the environment in which they will operate. Our environment consists of the network, users, and the adversary. We make our modeling decisions in order to best capture the relevant features of the environment as they have been revealed in the research and practice of low-latency anonymous communication. The second step is to choose the criteria with which to evaluate the protocols and make precise definitions of them. The criteria that we will consider are functionality, anonymity, latency, and message complexity.

2.2 Models

We give a basic model here that will serve as a basis for all of the protocols we describe and analyze. Features will be added to this model in later chapters as we seek better understanding of and designs for anonymous communication.

2.2.1 Input/Output Automata

We express our model and protocols using input/output automata [53]. Several different formalisms have been used to formalize anonymous-communication protocols, including an ACP-style process algebra [55], the universally composable framework [10], and a run-and-systems logical semantics [41]. The main obstacle

to useful anonymous communication is preventing *traffic analysis*, that is, hiding traffic patterns of traffic from the adversary. I/O automata allow us to focus on this problem rather than the cryptographic aspects of the protocols. Also, I/O automata are quite general and allow us to model asynchronous computation and communication.

2.2.2   Executions

We use standard notions of *execution* and *fairness*. An execution is a possible run of the network given its initial state. Fairness for us means that any message an automaton wants to send will eventually be sent and every sent message is eventually received.

We introduce the notion of a *cryptographic* execution. This is an execution in which no agent sends a control message encrypted with active keys it doesn't possess before it receives that message. We restrict our attention to such executions, and must require our encryption operation to prevent an attacker from outputting a control message, with more than negligible probability, when it is encrypted with keys he doesn't possess. This is reasonable because we can easily create a ciphertext space that is much larger than a limited control message space P. Note that this precludes the use of public-key encryption to encrypt the packets because such messages can easily be constructed with the public keys of the routers.

Definition 1. *An* execution *is a sequence of states of an I/O automaton alternating with actions of the automaton. It begins with an initial state, and two consecutive states are related by the automaton transition function and the action between them. Every action must be* enabled, *meaning that the acting automaton must be in a state in which the action is possible at the point the action occurs.*

Because there are no internal actions in the automata, all actions are message sends or message receives. Frequently, we will treat an execution as a sequence of actions, because the states are implicit from these and the initial states.

Definition 2. *A finite execution is* fair *if there are no actions enabled in the final state. Call an infinite execution* fair *if every output action that is enabled in infinitely many states occurs infinitely often.*

Definition 3. *An execution is* cryptographic *if an agent sends a message containing $\{p\}_{k_1,\ldots,k_i}$ only when it possesses all keys $k_1,\ldots,k_i$, or when for the largest j such that the agent does not possess k_j, $1 \leq j \leq i$, the agent has already received a message containing $\{p\}_{k_1,\ldots,k_j}$.*

2.2.3 Network

We model onion routing as a fully-connected asynchronous network of I/O automata. The network is composed of FIFO channels. There is a set of users U, a set of routers R, and a set of destinations D. Let $n = |U|$ be the number of users and m be the number of routers. Let $N = U \cup R \cup D$. The term *agent* refers to any element of N.

The network delivers messages in the message space $\mathcal{M}$. $\mathcal{M}$ must satisfy the following properties:

1. It contains the integers, that is, $\mathbb{Z} \subseteq \mathcal{M}$.

2. Message tuples are valid messages, that is, $\mathcal{M}^k = \mathcal{M}$, $k \in \mathbb{N}$.

Our network model provides asynchronous communication between every pair of agents. Asynchrony is an important feature, because synchronizing the actions of different users can be used to make them indistinguishable to an adversary and is a key technique used by anonymous-communication protocols [24]. The network model does not include, for now, any more explicit notion of time than the ordering on actions.

2.2.4 Users

We model user behavior very simply. At the start of the protocol each user wants to send a message to one destination and to receive one message back.

User behavior is an important feature of our model because it has a big effect on anonymity. When there are very few users that communicate, for example, good anonymity is simply not achievable. As another example, if each user only communicates with a single unique destination, and the adversary knows this, then the user cannot communicate anonymously unless the adversary can't tell if the user is communicating at all.

Our model of user behavior is an intentionally simple example of bidirectional communication. We will show that many of the common anonymous-communication protocols do not have satisfactory performance even for such basic user behavior. We will then focus on protocols with better performance, which will then allow us to consider the effect of more complicated, and more realistic, users.

There are many features of user behavior that are not addressed by our basic model. Users may choose destinations in a predictable way. They may contact multiple destinations simultaneously. Users and destinations may participate in bidirectional streams of messages rather than single message/response pairs. Instead of communicating just once, users may participate in many communication sessions over time.

These ways in which plausible user behavior differs from our model may be significant in an anonymous-communication protocol. An adversary might be able to learn a user's frequent communication partners if he chooses them somewhat predictably over time. Moving to stream communication may open up additional opportunities for an adversary to determine the destination. The simple user behavior model already leads to interesting problems in evaluating and designing anonymous-communication protocols. Also, we do examine more complicated user behavior in Chapters 6 and 8. These show that our simple model is robust in some ways.

2.2.5 Adversary

The adversary in our system is a set of users and routers $A \subseteq U \cup R$. The adversary is *active* in the sense that the automata running on members of A are completely arbitrary. We call an agent a *compromised* if $a \in A$. We assume that the adversary knows the number of users, routers, and destinations.

Note that we do not allow the adversary to compromise any of the destinations. It does seem possible in several situations this could be the case - for example, communication with a known hostile destination or a destination that is under surveillance. However, the case that this doesn't happen generally captures protocol performance and protocol design challenges, and therefore that is the case we focus on. Because our design criteria will require that all protocols use just one router to communicate with the destination, the case that this router is compromised presents the same situation as the case that the destination itself is compromised. Therefore many of our analyses extend naturally to the case that the adversary can compromise destinations.

We choose to make the adversary *local*, in that he controls a subset of the network. In particular, this limits him to only a partial observation of network activity. This contrasts with a *global* adversary, popular in analyzing mix networks, that can observe the entire network. That adversary is too strong for a large network such as the Internet. This is a very helpful restriction for designing effective protocols; we can confuse a local adversary just by making sure that some network events occur outside of his view. A global adversary precludes this, and seems to require protocols that have high resource requirements.

We also make the adversary active in the network. This is opposed to a *passive* adversary, that is, an adversary that obeys the protocol and just makes inferences from his observations, which is considered often in the literature. A passive adversary is interesting for several reasons: it is a trivial behavior for the adversary to implement, it is hard to detect and protect against because it is behaving the same as a normal user, and there are mechanisms (for example, zero-knowledge proofs) that help enforce correct adversary

behavior. However, an active adversary is much more realistic, and ideally a protocol should work against such an adversary.

Because we are trying to keep some knowledge hidden from the adversary, it is important to specify the knowledge that the adversary starts with. Our basic knowledge assumption is that the adversary knows the number of users, routers, and destinations that exist in the network. This is already probably too strong, in that it gives the adversary more information than he likely has. However, in the protocols we consider, the adversary can make good statistical guesses about the number of users, routers, and, in some cases, destinations, just by observing the operation of the protocol at the router he controls.

2.2.6 Weaknesses

Our model is necessarily simple, and it omits many features of actual networks that are relevant to anonymous communication. We mention a few.

First, the network does not include any notion of time. We consider this somewhat in Chapter 8, but many aspects remain absent. In particular, we do not consider congestion in the network, in which the time to process a message at a routers depends on the number of messages the have received. In real networks, congestion can be used to incapacitate parts of the network or create patterns of artificial that can reveal the patterns of communication [63]. Also, there may be different delays between different hosts in the network, and this can reveal information about where messages are likely to have originated [43].

Also, we have just one round of communication involving ane message and a response. Multiple connections over time give the adversary an opportunity to learn from what he sees and adapt. The adversary might, for example, seek to compromise machines that are of interest after observing some messages being sent [80]. Also, when not all users participate in the same number of connections at a given time, the adversary may be able to perform an *intersection attack* [48, 18], that is, to observe which users and destinations are simultaneously active and infer the common destinations of the users. Multiple messages allow the adversary to perform *destination fingerprinting*, that is, to infer the destination given the volume and timing of return messages [42]. This is particularly an issue when there is a small set of popular destinations, such as on the Web.

2.3 Criteria

The criteria that we will use to evaluate anonymous-communication protocols are functionality, anonymity, latency, and message complexity. The popularity of onion routing can be attributed in part to the fact that it has reasonable performance according to all of these criteria, while the many competing protocols usually perform badly in at least one. Note, however, that good performance on each of these may not be sufficient for a useful protocol in many of the situations in which anonymous communication might be applied. For example, one might want the user to be *unobservable*, that is, that the adversary is unable to determinine whether or not the user is communicating with anyone at all. More weakly, one might want the protocol to be hard to distinguish from other popular protocols on the same network. However, it does seem that for many of the Internet applications that onion routing is used to anonymize - for example, Web browsing, file sharing, chat, and email - any protocol that succeeds against these four criteria is likely to be useful.

2.3.1 Functionality

The communication functionality that we want our protocol to provide is to, first, deliver a message from a user to a given destination and, second, to deliver a possible response message from the destination to the user. Formally, the protocol must implement the operation AnonymousMessage(d,m). It takes as inputs a destination $d \in D$ and a message $m \in \mathcal{M}$, and it returns a number $n \in \mathbb{N}$ that identifies that use of the operation. After a user u calls AnonymousMessage, some router $r \in R$ in the network must eventually deliver the message $[g(n), m]$ to d, where $g(n) : \mathbb{N} \to \mathbb{N}$ is an injective map. Then, if at some point d sends the message $[g(n), m']$ to r, $m' \in \mathcal{M}$, the network must eventually deliver the message $[n, m']$ to u. Note that the protocol is only required to implement these actions when all the agents correctly execute it.

Our functionality requirement is deliberately simple to avoid complicating the analysis of protocols according to all of our criteria. Also, it allows us to directly compare the protocols for anonymous communication in the literature, which generally implement this basic operation.

We require that only one router communicate with the destination for a given communication from the user. This requirement is motivated by the observation in practice that it is difficult to get hosts to support a special protocol for anonymous communication. Therefore, the destination communicates using standard point-to-point protocols with one router that acts as an anonymous proxy for the user. Removing this requirement would make it easier to design effective protocols, but it is a real-world constraint that we feel is important to adopt.

2.3.2 Anonymity

There are multiple types of anonymity involved in anonymous communication and several plausible ways of measuring each one. Indeed, there is a growing body of research solely concerned with identifying and comparing various notions of anonymity. Our approach is to pick an important type of anonymity and study how to implement it very well under our chosen criteria, evaluating anonymity with some reasonable metric. We are mostly not concerned with examining diferent types of anonymity and determining the best metrics.

Therefore we emphasize one type of anonymity evaluated according to one metric. We will focus primarily on *relationship anonymity*, that is, on how well the adversary can determine the communication partners of a user. To measure this, we primarily use the probability that the adversary assigns to the user's destination after observing the protocol, under probabilistic extensions our model.

Types

In general, any action that a user can perform can be performed anonymously. This means that if the adversary can tell that a certain action was performed, then he cannot tell which user was the actor. We will also consider the user anonymous, in the case that adversary cannot tell whether or not the action was performed.

There are several types of anonymity associated with network communication that may be of interest:

- *Sender anonymity*: A message has sender anonymity if the adversary cannot tell which user sent that message.

- *Receiver anonymity*: A message has receiver anonymity if the adversary cannot tell which destination a specific message is sent to.

- *Relationship anonymity*: A user and destination have relationship anonymity if the adversary cannot tell which destination a user communicates with.

In general, there is a weak logical relationship between these three notions - if user u and destination d have relationship anonymity, then, for every message m sent by u, m has sender anonymity or receiver anonymity. Otherwise, for some message the adversary would be able to determine that it was sent by u to d. If u and d do not have relationship anonymity, we cannot say anything about the anonymity of their messages. It is possible, for example, that every message has both sender and receiver anonymity, and yet the adversary can still determine that u sent at least one of them to d.

We focus on relationship anonymity because it has been the primary goal of fielded anonymous communication systems. Also, providing receiver anonymity is made fundamentally difficult by our functionality requirement that one final router send a user's message to its destination. This final router may be compromised and will observe the message and its recipient. Conversely, providing sender anonymity seems very close to providing relationship anonymity, because in our model every user communicates with exactly one destination. Therefore for u to have relationship anonymity, we only need there to be one message with sender anonymity for which u is among the possible senders and that is to a destination other than d.

Metrics

There are several reasonable metrics with which to measure anonymity. Unlike the primitives in cryptography, which have similar information-hiding goals to anonymous communication and are in several cases closely related [45], the proper definitions for anonymity have not been settled. This seems to be because the kinds of very strong guarantees that cryptographic constructions can satisfy have not been achieved by useful constructions for anonymous communication. The way to measure the anonymity that existing constructions do achieve is then less clear.

In an anonymous-communication protocol, the adversary makes a series of observations of the network. There is a set of possible executions, including the execution that actually occurred, that are consistent with his observations. Metrics that are functions on this set are *possibilistic* in that they only consider which situations are possible. For a given agent and action, they include

- *Boolean*: whether it is possible that the agent didn't perform the action

- *Anonymity set*: the number of actors that may have performed the action

Possibilistic anonymity is a very weak notion of anonymity, as the adversary may have some external information that allows him to conclude that certain executions are very unlikely. To model this uncertainty, we can add a probability measure to executions. Then the adversary can infer a conditional probability distribution on executions given his adversary. Metrics that take into account the probability measure are *probabilistic*. For a given agent and action, the probabilistic metrics that have been proposed include

- *Agent probability*: the conditional probability that the adversary assigns to the event that the correct agent performed the action

- *Entropy*: the entropy of the conditional distribution over agents

15

- *Min-entropy*: the min-entropy of the conditional distribution over agents

- *Variation distance*: the variation distance between a prior distribution over agents for a given action and a posterior distribution (that is, the conditional distribution)

Probabilistic metrics take into account both the randomness that is a part of many anonymous-communication protocols and predictable behavior of users that is reasonably modeled as probabilistic. From the metrics that have been proposed, we select agent probability for most of the thesis. It emphasizes how well a protocol hides specific actions by a specific user, and it is relatively easy to work with mathematically.

We use different definitions of anonymity for the different models and protocols that we consider. Therefore will defer for now giving precise definitions of anonymity.

2.3.3 Latency

The latency of a protocol measures how long it takes for a message to be received by its intended recipient after it is sent. Because our basic model lacks an explicit notion of time, we give here a definition of latency based on the longest sequence of messages that must be sent and received after a message is sent until the destination finally receives it.

For a given execution, the actions form a partial order in which, for actions a_1 and a_2, $a_1 \prec a_2$ if a_1 and a_2 are actions of the same agent's automaton and a_1 was performed by that automaton before a_2. We define the latency of an anonymous message in a given protocol execution to be the length of the longest chain in the partial order from the action in which a user sends the anonymous message (that is, calls AnonymousMessage) to the action in which the message is sent. If no such chain exists, we define the latency to be infinity. Then we can define the general latency of a protocol.

Definition 4. *Let $\mathcal{E}$ be the set of fair, cryptographic executions when all agents run the given protocol. The latency of a protocol is the minimum over all executions $\alpha \in \mathcal{E}$ of the protocol of the latency of α.*

We take the smallest latency over all executions because it capture the latency inherent in the protocol. If we were to take the largest possible latency, executions that include long sequences of messages unrelated to delivering the anonymous message could yield unrealistically large values.

Our definition of latency assumes that all agents faithfully run the protocol. Our model of the adversary allows him to run abitrary programs, however. This is an important issue, because a protocol in which the adversary can easily delay or deny the delivery of messages is not very usable. If we assume that the

adversary's goal is to deanonymize communication, however, assuming that he employs strategies that are most disruptive to message delivery may overstate the problem. Therefore we separately consider latency under adversarial misbehavior.

Definition 5. *Let $\mathcal{E}_A$ be the set of fair, cryptographic executions when the honest agents run the given protocol and the adversary runs A. The latency against A is the minimum over all executions $\alpha \in \mathcal{E}_A$ of the latency of α. The* adversarial latency *is the maximum over all adversaries A of the latency against A.*

It is true that delaying message delivery may help the adversary deanonymize users. If users repeatedly attempt to send an anonymous message until it is successfully delivered, the adversary can try to prevent delivery until the user is successfully deanonymized. This issue has been explored by Borisov et al. [8]. We do not consider anonymity under this kind of user behavior.

2.3.4   Message Complexity

In general, we care about the total communication requirements of a protocol. Important metrics include the number of messages, message sizes, and the message contention [53]. These metrics can be considered for the system overall or for individual agents. We want our protocols to have reasonable communication requirements for each agent, and so we are interested in the per-agent metrics. Moreover, for the protocols we consider, message size and contention are not a problem in theory or practice. Therefore we will only consider the number of messages sent by the agents in the protocol.

Definition 6. *Let $\mathcal{E}$ be the set of fair, cryptographic executions of a protocol when all agents run the protocol. Let $m(\alpha)$ be the number of messages sent in execution α. Let $a(\alpha)$ be the number of calls to* AnonymousMessage. *The* message complexity *of the protocol is the minimum over all $\alpha \in \mathcal{E}$ of $m(\alpha)/a(\alpha)$.*

Similar to defining latency, we choose the minimum over executions because we want the message complexity to represent the fundamental resource requirements of the protocol.

As before, we cannot assume that the adversary faithfully runs the protocol, and we want to consider by how much his malicious behavior can increase the message complexity of the protocol. Because the adversary can send an arbitrary number of messages, we normalize by the number of messages that he send.

Definition 7. *Let $\mathcal{E}_A$ be the set of fair, cryptographic executions when the honest agents run the given protocol and the adversary runs A. Let $m(\alpha)$ be the number of messages sent in execution α, let $m_A(\alpha)$ be*

the number of messages sent by the adversary A, and let $a(\alpha)$ be the number of calls to *AnonymousMessage*. *The message complexity against A is the minimum over all executions $\alpha \in \mathcal{E}_A$ of $m(\alpha)/(m_A(\alpha)a(\alpha))$. The* adversarial message complexity *is the maximum over all adversaries A of the message complexity against A.*

Chapter 3

Related Work

Our contributions to anonymous-communication protocols take place in a large body of related work. Indeed, in addition to the most direct connections to other work on anonymous communication, the whole area has much in common with a variety of other areas in security in privacy, such as anonymous data storage [21, 14, 86], anonymous data publishing [26], private information retrieval [35], secure multiparty function evaluation [89, 39], and cryptography in general [45]. We survey here the work that most directly addresses either the questions in anonymous communication that we consider or the concepts and techniques we use to come up with solutions.

3.1 Anonymous Communication Protocols

A large number of protocols for anonymous communication have been designed. They vary along several dimensions, including the communication functionality they provide, such as one-way message delivery or two-way delivery/response; the network model they require, such as a synchronous or asynchronous network, and a wired or wireless network; the adversary model they are designed against, such as a passive or active adversary, and local or global adversary; the latency they add; and their message complexity. These systems are described under varying models and assumptions. We have one specific set of models, definitions, and criteria that we are interested in, which is given in Chapter 2. Therefore, we describe and evaluate many of the existing protocols under these choices. This shows how, in the situations we consider important, they compare to one another, and highlights the contributions of our analysis of onion routing and of our new

protocol designs. Good surveys of existing protocols for anonymous communication are available [47, 16, 27], and thus our descriptions will be brief.

In the following, let $b = |A \cap R|/|R|$ be the fraction of routers that the adversary controls. Let $d = |D|$ be the number of destinations.

3.1.1 Single-hop Proxy

Users simply randomly choose a router to forward their message.

Anonymity:
$$\begin{cases} 1 & \text{with probability } b \\ 0 & \text{with probability } 1 - b \end{cases}$$

Latency: 2

Adversarial Latency:
$$\begin{cases} \infty & \text{with probability } b \\ 2 & \text{with probability } 1 - b \end{cases}$$

Message complexity: 2

Adversarial message complexity: 2

3.1.2 Mix Networks

In a mix network, users choose a common sequence of mixes (that is, a *cascade* of routers), and, in reverse order of the sequence, encrypt the message with the identity of the next mix [12]. A reply block is contained with the encrypted identity of the user. When a router receives a message it decrypts the message and forwards the result to the next router in the chain. There are many improvements and variations on this idea (for example, [49, 17, 69, 25, 73]). We just consider the basic scheme here, which exhibits the problem, common in mix networks, that the mixes can selectively forward messages. Let l be the length of the cascade.

Anonymity: If the adversary controls the first and last routers, he can just forward a message from one user and observe its final destination. Therefore the anonymity is

$$\begin{cases} 0 & \text{with probability } b^2 \\ 1 & \text{with probability } 1 - b^2 \end{cases}$$

Latency: l

Adversarial Latency: The adversary can drop messages on any path in which he appears. Therefore the adversarial latency is

$$
\begin{cases}
l+1 & \text{with probability } (1-b)^l \\
\infty & \text{with probability } 1-(1-b)^l
\end{cases}
$$

Message complexity: $l+1$

Adversarial message complexity: $l+1$

3.1.3   Crowds

In Crowds [68], users randomly forward messages among themselves. When they join, users send a request to create a persistent path to a random user. That user extends the path with probability p, and otherwise the paths ends. Messages are sent along the path between the first user and a destination he specifies. Anonymity is provided because the adversary cannot determine when the user he received the message from was the source.

Anonymity: All members on a path can see the destination, and so the only relevant event is whether or not the adversary observes the user. Therefore the anonymity is

$$
\begin{cases}
1 - ((1-b) - 1/n)p & \text{with probability } b \\
(1-b)p/n & \text{with probability } 1-b
\end{cases}
$$

Latency:

$$
1+i \qquad \text{with probability } p^{i-1}(1-p), i \geq 1
$$
$$
1/(1-p) \quad \text{in expectation}
$$

Adversarial Latency: The adversary can drop messages on any path in which he appears. Therefore the adversarial latency is

$$
1+i \quad \text{with probability } p^{i-1}(1-p)(1-b)^i, i \geq 1
$$
$$
\infty \qquad \text{with probability } b/(1-(1-b)p)
$$

Message complexity: The number of messages sent per anonymous message is the same as the latency,

because every message corresponds to a hop on the path. Therefore the message complexity is

$$1 + i \qquad \text{with probability } p^{i-1}(1-p), i \geq 1$$
$$1/(1-p) \quad \text{in expectation}$$

Adversarial message complexity: The best that the adversary can do is not to extend further any paths that he is asked to be a part of, because the complexity measure charges the adversary for every such message that he sends. Thus the adversarial message complexity is only smaller than message complexity.

3.1.4 Dining Cryptographers networks

In Dining Cryptographers networks [13], every pair of users shares a private key. In the ith round, a user that doesn't want to send a bit sends the XOR of the ith bit of all of his keys, and a user that does want to send a bit sends the inversion of that XOR. Repeating this protocols allows users to send full messages. To avoid collisions, we assume that time is divided into blocks. The first block is a reservation block of size kn^2, where n is the number of users. Each bit potentially reserves a time slot for sending. Users randomly choose a bit to flip to reserve a slot, and there are no collisions in the block with probability at least $1 - e^{-(1-1/n)/2k}$. The reservation block is followed by send blocks. We can save time by scheduling the user that flipped the i *reserved* bit in the ith slot, rather than just scheduling the ith bit in the ith block. Send blocks are followod by an equal number of response blocks, where the ith response block is intended for the sender of the ith send block. Then a new reservation block follows the last response block. We can modify the scheme to allow the point-to-point communication required by our scheme by requiring users to designate the endpoint and destination of a message in a header before the message. We assume that users encrypt the message and destination with the public key of the endpoint. This protocol implements our required functionality assuming that every destination send its response before their reponse round starts.

Anonymity: $1/n$ (perfect)

Latency: Users have to wait for their time slot to send. If r users reserved slots, the latency is $r/2 + 1$ in expectation.

Adversarial latency: The adversary can easily jam the service by constantly sending. Therefore the adversarial latency is ∞.

Message complexity: The reservation block requires all n users to send kn^2 bits to $n-1$ other users, for a total of $kn^3(n-1)$ bits. Users send $n(n-1)$ messages total in each of n send blocks, for a total of $n^2(n-1)$ messages. Therefore the message complexity is $(kn^3(n-1) + n^2(n-1))/n = O(n^3)$.

Adversarial message complexity: The adversary must reserve a bit to cause $O(n^2)$ extra messages to be sent. Therefore the adversarial message complexity is the same as the message complexity: $O(n^3)$.

3.1.5 Onion Routing

Onion routing is very similar to protocols using mixes. The key difference for our analysis is that users choose their routes independently of each other. Each user constructs a persistent routes, or *circuit*, of length l by recursively constructing a circuit of length $l-1$ and sending a message down it to extend the route by one. Messages down the circuit from the user are encrypted by the user once for each router on the circuit in reverse order. A router on the circuit takes an incoming message, decrypts it, and forwards the result to the last router. Messages up the circuit, going in the direction of the destination towards the user, follow the reverse process and are decrypted as they are passed from router to router. The message to the destination includes the identity of the destination.

Anonymity: We analyze the anonymity of onion routing in Chapters 5 and 6. The results are that for reasonable user behavior and adversary sizes, the worst case expected anonymity is about $b + (1-b)p_d^u + O(\sqrt{\log n/n})$. More "typical" parameter values yield expected anonymity of $b^2 + (1-b^2)p_d^u + O(1/n)$.

Latency: l

Adversarial latency The adversary can block messages on circuits of which it is a part. Therefore the adversarial latency is

$$\begin{cases} l & \text{with probability } (1-b)^l \\ \infty & \text{with probability } 1 - (1-b)^l \end{cases}$$

Message complexity: The number of messages during the circuit-building process is $\sum_{i=1}^{l} 2i = l(l+1)$, and sending a message adds an additional $2l+1$ messages. Therefore the message complexity is $O(l^2)$.

Adversarial message complexity: Assuming there is no limit to the length of circuits, the adversary can build arbitrarily long circuits and thereby generate a total amount of traffic that is an arbitrarily high multiple of the amount of his own traffic. Therefore the adversarial message complexity is ∞.

3.2 Analysis of Anonymous Communication Protocols

Numerous papers have informally discussed the security of the design of onion routing and related systems, as well as theoretical and experimentally demonstrated attacks. There have also been numerous formalizations of anonymous communication. However, formal analyses have primarily been of systems other than onion routing, for example, Dining Cryptographers networks and Crowds.

There has been work formalizing systems similar to onion routing but without persistent circuits. Camenisch and Lysyanskaya [10] prove that the cryptography their protocol uses doesn't leak any information to nodes in the path other than the previous and next nodes, but leave open what anonymity it provides. This question is answered in part by Mauw et al. [55], who formalize a similar connectionless protocol in an ACP-style process algebra and, under a possibilistic definition of anonymity, show that it provides sender and receiver anonymity against a global passive adversary. Cryptography is dealt with using high level assumptions, similar to our work. Their model and analysis has much in common with ours, but it does differ in several important ways. First, the protocol they investigate is connectionless: each data "onion" stores the identities of the routers it will pass through. This is significantly different from onion routing, which is circuit-based. Second, the analysis is done with respect to a passive adversary, which exhibits only a subset of the behavior of an active adversary. Third, in their model agents choose destinations asynchronously and the adversary must take into account every onion he has seen when trying to break anonymity. In our model all agents choose a single destination, which gives the adversary more power. In some ways, our work extends theirs, and several of the differences noted here appear in [55] as suggestions for further research.

Ours is not the first formalization of anonymous communication. Earlier formalizations used CSP [71], graph theory and possible worlds [44], and epistemic logic [82, 41]. These earlier works focused primarily on formalizing the high-level concept of anonymity in communication. For this reason, they applied their formalisms to toy examples or systems that are of limited practical application and can only provide very strong forms of anonymity, for example, dining-cryptographers networks. Also, with the exception of [41], they have at most a limited ability to represent probability and probabilistic reasoning. Our goal in Chapter 4 is on the formalization of a widely deployed and used, practical, low-latency system.

Halpern and O'Neill [41] give a general formulation of anonymity in systems that applies to our model. They describe a "runs-and-systems" framework that provides semantics for logical statements about systems. They then give several logical definitions for varieties of anonymity. It is straightforward to apply this framework to the network model and protocol that we give in [31]. Our possibilistic definitions of sender

anonymity, receiver anonymity, and relationship anonymity then correspond to the notion of "minimal anonymity" as defined in their paper. The other notions of anonymity they give are generally too strong and are not achieved in our model of onion routing.

Other works have presented probabilistic analysis of anonymous communication [68, 74, 88, 15, 18, 54, 49] and even of onion routing [80]. The work of Shmatikov and Wang [76], in particularl, is similar to ours. It calculates relationship anonymity in mix networks and incorporates user distributions for selecting destinations. However, with the exception of [74], these have not been formal analyses. Also, whether for high-latency systems such as mix networks, or low-latency systems, such as Crowds and onion routing, many of the attacks in these papers rely on observing user behavior over time, what is known as an *intersection attack*. Several of these papers set out frameworks for making that more precise. In particular, [15], [18], and [54] constitute a progression towards quantifying how long it takes (in practice) to reveal traffic patterns in realistic settings.

Our model does not consider observing user behavior over time. Our probabilistic model effectively assumes that the intersection attack is done. The adversary already has a correct distribution of a user's communication partners. We are investigating the anonymity of a communication in which a user communicates with one of those partners in the distribution. This follows the anonymity analyses performed in much of the literature [49, 55, 68, 80], which focus on finding the source and destination of an individual communication. Our analysis differs in that we take into account the probabilistic nature of the users' behavior.

3.3   Improved Anonymous-Communication Protocols

In Chapters 7 and 8, we present two new protocols for anonymous communication. Our goal in each case is to improve anonymity while maintaining efficiency in latency and message complexity.

3.3.1   Trust

Chapter 7 explores improving onion routing when users have outside knowledge of which routers to trust. Trust has many meanings and applications in computer security [29, 78, 7, 1, 11, 75, 70, 19]. Much of the literature is concerned in one way or another with propagation or transfer of trust from where it is to where it needs to be. Our concern is not with the transfer of trust information but with what it means in the context

of onion routing and how to make use of it. We consider how trust associated with network nodes or links might be used to protect (or reveal) information that would undermine the anonymity of communicants.

3.3.2 Preventing Timing Attacks

Timing attacks are a major challenge in low-latency anonymous communication [50]. They have been observed in some of the earliest low-latency systems [3], including initial versions of onion routing [40]. These attacks are also closely related to traffic analysis in mix networks [67].

In a passive timing attack, the adversary observes timing patterns in a network flow, and then correlates them with patterns in other traffic that it observes. If the adversary is able to observe both the user and the destination, he can thereby link the two. The ability of the adversary to perform this correlation has been experimentally demonstrated several times [91, 50, 66, 64, 4].

A solution to passive timing attacks is to get rid of identifying patterns in the traffic by padding and delaying it. The drawbacks to such an approach are added latency and bandwidth overhead. Our protocol relies on the existence of some acceptable and effective padding scheme. Constant-rate padding, in which traffic is sent at a constant rate by filling in the gaps with dummy packets, is probably the most obvious such scheme. It has appeared multiple times in the literature [80, 33, 50]. Levine et al. [50] propose a "defensive dropping" mechanism which adds dummy packets at the start of the circuit and drops them at various routers before the end. This reduces the correlation between any patterns in the incoming streams and patterns in the outgoing streams. Shmatikov and Wang [77] propose a variable-rate padding scheme. In their scheme, packets from the user are forwarded with little delay, and dummy packets are added by the intermediate routers according to a probability distribution on the packet inter-arrival times. Wang et al. [87] describe a link-padding scheme for low-latency systems, but their system is designed for a situation in which the adversary is not active and the destination participates in the protocol. This situation does not reflect the kind of Internet communications that have proven useful and that we target.

All of these schemes are vulnerable to an adversary that actively delays packets from the user. Yu et al. [90] show that this can be done in a way that makes it difficult for the user or the system to detect that the attack is occurring. According to the results of Shmatikov and Wang, their variable-rate scheme of inserting padding packets throughout the network provides the most protection against such an attack. However, in the scenario we are designing for, the final destination does not participate in the protocol. Therefore there must be one last router in the system that provides the point of contact to the destination. This final router

26

can identify the dummy packets, and therefore can observe timing patterns that the adversary created earlier in the stream.

Simply delaying packets that pass directly through adversarial routers isn't the only active timing attack that has been demonstrated. Murdoch and Danezis [63] show that in onion routing the adversary can actively add delay patterns to the data by sending bursts of traffic through a router. This can be used to determine the routers on a given circuit. A queuing scheme has been suggested by McLachlan and Hopper [58] to mitigate this problem. Fu et al. [32] describe how the presence of a flow on a router can also be determined by ping times to the router, and suggests the simple countermeasure of treating ping packets equally. We do not address such congestion attacks in this thesis, and it is outside of our model. Borisov et al. [8] look at the case that the adversary doesn't just delay, but drops packets in a denial-of-service (DoS) attack aimed at forcing users to move to circuits that the adversary can deanonymize. This is an easy behavior to observe, although proving the culprit may be difficult, and the authors suggest that reputation mechanisms can be used to limit such attacks. Such an attack was also discussed by Dingledine et al. [24]. Mittal and Borisov [60] examine the use of active attacks to subvert the host lookup mechanism in the Salsa and AP3 peer-to-peer protocols. This isn't a timing attack, but they uncover a tradeoff offered by the protocol between preventing active and passive attacks in the lookup that is very similar to the tradeoff we examine in our use of redundancy.

Hopper et al. [43] explore several attacks based on the network latency between hosts. In particular, they show the latencies from multiple hosts to a user can be very identifying. The user in our protocol communicates with several routers as a first hop, and in light of this attack we took care not to allow the adversary to infer these latencies.

One key feature of our protocol is the use of a layered mesh to provide redundancy. The use of redundancy for several purposes has been also explored in previous protocols. Syverson [81] suggests using router "twins," pairs of routers that share the same key, to provide a backup in case a router fails. Two redundancy schemes to manage failures, K-Onion and Hydra-Onion, are proposed by Iwanik et al. [46]. In the K-Onion scheme, onions are encoded in a way that provides more than one option for the next hop to route around host failures. In the Hydra-Onion scheme, onions get forwarded to two routers at every hop instead of one. The purpose of this scheme is to prevent an active attack in which the adversary kills messages from a known user and observes which destination stops receiving traffic. Redundancy to support failure is not our goal, and such schemes are in some ways complementary to our own. However, the redundancy in our protocol

does protect against honest node failures as well as malicious ones. Our analysis, however, only considers malicious node failures. The aim of the Hydra-Onion scheme is similar to ours, although they only analyze the ability of the adversary to kill a connection. Morever, the scheme is designed for an adversary that controls links, and not one that controls routers.

Another critical feature of our protocol is the use of explicit timing to coordinate the users and routers. This is similar to the timing instructions of Stop-and-Go mixes [49]. Such mixes are given a time window within which the packets must arrive, and they delay forwarding by an exponentially-distributed amount of time. Although the techniques are similar, this scheme is designed for mix networks and not stream communication, and this scheme does give the adversary some slack time within which a timing signature could possibly be placed. Moreover, the lack of any redundancy means that any slowdowns within the network, even of benign origin, can quite easily kill a connection. The protocol is also susceptible to DoS attacks.

The timing approach we take is also similar to the use of synchronicity in mix networks by Dingledine et al. [24]. They describe synchronous batching on free routes and show that it typically provides better anonymity than cascades. Our scheme can be viewed as low-latency synchronous batching. In their case the synchronicity is provided by long delays in the mix network, while in our case synchronicity is provided by the padding schemes. They are able to disregard active attacks in which messages are dropped or delayed, because each mix is able to wait for enough time that any missing packets must be maliciously withheld. They rely on reputation systems [25] to eventually find and punish such action. With low-latency communication, on the other hand, we cannot assume that such delays will not happen. Moreover, our solution does not rely on the use of an external reputation mechanism. This is an advantage, because reputation mechanisms in general only detect large-scale misbehavior and must tolerate smaller-scale targeting of users.

Chapter 4

Formalizing Onion Routing

4.1 Introduction

Low latency and other performance characteristics of onion routing can be demonstrated experimentally; anonymity-preserving properties cannot. Also, even with careful design, vulnerabilities can persist. The initial authentication protocol for Tor, for example had a cryptographic-assumption flaw that left it open to man-in-the-middle attacks. The revised authentication protocol was then proven to be free of such flaws [37]. As anonymous-communication networks are increasingly relied upon for sensitive personal and business transactions, it is increasingly important to assure their users that anonymity will be preserved. Long-established components of such assurance in system security include a formal model, proving security guarantees in that model, and arguing that the model captures essential features of the deployed system.

An onion-routing network consists of a set of onion routers and clients. To send data, a client chooses a sequence of routers, called a *circuit*, and constructs the circuit using the routers' public keys. During construction, a shared symmetric key is agreed upon with each router. Before sending data, these keys are used to encrypt each packet once for each router in the circuit and in the reverse of the order that the routers appear in the circuit. Each router uses its shared key to decrypt the data as it is sent down the circuit so it is fully decrypted at the end. Data flowing up to the client has a layer of encryption added by each onion router, all of which are removed by the client. The layered encryption helps hide the data contents and destination from all but the last router and the source from all but the first router. The multiple encryptions and decryptions also make it harder for an observer to follow the path the data takes through the network.

In this chapter we formalize this protocol in order to rigorously analyze the anonymity it provides.

4.2 Automata

We give the automata descriptions for the users and routers that are based on the Tor protocol [23]. We have simplified the protocol in several ways. In particular we do not perform key exchange, do not use a stream cipher, have each user construct exactly one circuit to one destination, do not include circuit teardowns, eliminate the final unencrypted message forward, and omit stream management and congestion control. Section 5.3.7 discusses the effects of changing some of these features of our protocol.

We use circuit identifiers to simplify the definitions and proofs of anonymity in our protocol. Circuit identifiers mimic the effects of a timing attack in onion routing. They have the added advantages of making it clear when this power is used. Theorem 5 shows that they are easily simulated by an adversary, and therefore they do not affect the anonymity of the protocol.

We assume that each router-and-user pair shares a set of secret keys; however, the router does not know which of its keys belong to which user. This separates, for now, key distribution from the rest of the protocol. We assume that all keys in the system are distinct. Let K be the keyspace. The triple (u, r, i) will refer to the ith key shared by user u and router r.

Let P be the set of control messages that are used in the protocol, and $\bar{P}$ be the extension of P by encryption with up to l keys. We assume that $\bar{P} \subseteq \mathcal{M}$. The control messages will be tagged with a link identifier and circuit identifier when sent, so let the protocol message space be $M = \mathbb{N} \times \mathbb{N} \times \bar{P}$. We denote the encryption of $p \in P$ using key k with $\{p\}_k$, and the decryption with $\{p\}_{-k}$. For brevity, the multiply encrypted message $\{\{p\}_{k_1}\}_{k_2}$ will be denoted $\{p\}_{k_1, k_2}$. Brackets will be used to indicate the list structure of a message (*i.e.* $[p_1, p_2, \ldots]$).

During the protocol each user u iteratively constructs a circuit to his destination. u begins by sending the message $\{\text{CREATE}\}_{k_1}$ to the first router, r_1, on his circuit. The message is encrypted with a key, k_1, shared between u and r_1. r_1 identifies k_1 by repeatedly trying to decrypt the message with each one of its keys until the result is a valid control message. It responds with the message CREATED. (Note that this is different from the implemented Tor protocol, in which the CREATE message would be encrypted with the public key for r_1 rather than one of the shared keys it holds.) Given a partially-constructed circuit, u adds another router, r_i, to the end by sending the message $\{[\text{EXTEND}, r_i, \{\text{CREATE}\}_{k_i}]\}_{k_{i-1}, \ldots, k_1}$ down

the circuit. As the message gets forwarded down the circuit, each router decrypts it. r_{i-1} performs the CREATE steps described above, and then returns the message $\{\text{EXTENDED}\}_{k_{i-1}}$. Each router encrypts this message as it is sent back up the circuit.

Link identifiers are used by adjacent routers on a circuit to differentiate messages on different circuits. They are unique to the adjacent pair. Circuit identifiers are also included with each message and identify the circuit it is traveling on. They are unique among all circuits.

The user automaton's state consists of the sequence of routers in its circuit, the destination, a number that identifies its circuit, and a number that indicates the state of its circuit. The user automaton runs two threads, one to extend a circuit that is called upon receipt of a message and the other to start circuit creation that is called at the beginning of execution. We express these in pseudocode rather than I/O automata, but note that the state changes in a particular branch occur simultaneously in the automaton. $k(u, c, b)$ refers to the key used by user u with router c_b in the bth position in circuit c. The automaton for user u appears in Automaton 1.

The router automaton's state is a set of keys and a table, T, with a row for each position the router holds in a circuit. Each row stores the previous and next hops in the circuit, identifying numbers for the incoming and outgoing links, and the associated key. There is only one thread and it is called upon receipt of a message. In the automaton for router r, we denote the smallest positive integer that is not being used on a link from r to q or from q to r as $minid(T, q)$. The automaton for router r appears in Automaton 2.

Automaton 1 User u

1: $c \in \{(r_1, \ldots, r_l) \in R^l | \forall_i r_i \neq r_{i+1}\}$; init: arbitrary   ▷ User's circuit
2: $d \in D$; init: arbitrary   ▷ User's destination
3: $i \in \mathbb{N}$; init: random   ▷ Circuit identifier
4: $b \in \mathbb{N}$; init: 0   ▷ Next hop to build
5: **procedure** START
6:     SEND($c_1, [i, 0, \{\text{CREATE}\}_{k(u,c,1)}]$)
7:     $b = 1$
8: **end procedure**
9: **procedure** MESSAGE(msg,j)   ▷ $msg \in M$ received from $j \in N$
10:     **if** $j = c_1$ **then**
11:         **if** $b = 1$ **then**
12:             **if** $msg = [i, 0, \text{CREATED}]$ **then**
13:                 $b++$
14:                 SEND($c_1, [i, 0, \{[\text{EXTEND}, c_b, \{\text{CREATE}\}_{k(u,c,b)}]\}_{k(u,c,b-1),\ldots,k(u,c,1)}]$)
15:             **end if**
16:         **else if** $b < l$ **then**
17:             **if** $msg = [i, 0, \{\text{EXTENDED}\}_{k(u,c,b-1),\ldots,k(u,c,1)}]$ **then**
18:                 $b++$
19:                 SEND($c_1, [i, 0, \{[\text{EXTEND}, c_b, \{\text{CREATE}\}_{k(u,c,b)}]\}_{k(u,c,b-1),\ldots,k(u,c,1)}]$)
20:             **end if**
21:         **else if** $b = l$ **then**
22:             **if** $msg = [i, 0, \{\text{EXTENDED}\}_{k(u,c,b-1),\ldots,k(u,c,1)}]$ **then**
23:                 $b++$
24:                 SEND($c_1, [i, 0, \{[\text{DEST}, d, HELLO]\}_{k(u,c,b-1),\ldots,k(u,c,1)}]$)
25:             **end if**
26:         **end if**
27:     **end if**
28: **end procedure**

Automaton 2 Router r

1: $keys \subseteq K$, where $|keys| \geq |U| \cdot \lceil \frac{l}{2} \rceil$; init: arbitrary   ▷ Private keys
2: $T \subset N \times \mathbb{N} \times N \times \mathbb{Z} \times keys$; init: $\emptyset$ ▷ Routing table
3: **procedure** MESSAGE($[i, n, p], q$)   ▷ $[i, n, p] \in M$ received from $q \in N$
4:      **if** $[q, n, \emptyset, -1, k] \in T$ **then**   ▷ In link created, out link absent
5:          **if** $\exists_{s \in R-r, b \in P} p = \{[\text{EXTEND}, s, b]\}_k$ **then**
6:              SEND($s, [minid(T, s), b]$)
7:              $T = T - [q, n, \emptyset, -1, k] + [q, n, s, -minid(T, s), k]$
8:          **else if** $\exists_{d \in D} p = \{[DEST, d, b]\}_k$ **then**
9:              SEND($d, [minid(T, d), b]$)
10:              $T = T - [q, n, \emptyset, -1, k] + [q, n, d, minid(t, d), k]$
11:          **end if**
12:      **else if** $[s, m, q, -n, k] \in T$ **then**   ▷ In link created, out link initiated
13:          **if** $p = \text{CREATED}$ **then**
14:              $T = T - [s, m, q, -n, k] + [s, m, q, n, k]$
15:              SEND($s, [i, m, \{\text{EXTENDED}\}_k]$)
16:          **end if**
17:      **else if** $\exists_{m>0}[q, n, s, m, k] \in T$ **then**   ▷ In and out links created
18:          SEND($s, [i, m, \{p\}_{-k}]$)   ▷ Forward message down the circuit
19:      **else if** $[s, m, q, n, k] \in T$) **then**   ▷ In and out links created
20:          SEND($s, [i, m, \{p\}_k]$)   ▷ Forward message up the circuit
21:      **else**
22:          **if** $\exists_{k \in keys} p = \{\text{CREATE}\}_k$ **then**   ▷ New link
23:              $T = T + [q, n, \emptyset, -1, k]$
24:              SEND($q, [i, n, \text{CREATED}]$)
25:          **end if**
26:      **end if**
27: **end procedure**

Chapter 5

Possibilistic Anonymity in Onion Routing

5.1 Introduction

We give precise definitions of anonymity of onion routing as formalized in Chapter 4. Using these, we provide necessary and sufficient conditions for anonymity to be provided to a user in our model. It should be noted that we do not analyze the validity of the cryptography used in the protocol and instead base our proofs on some reasonable assumptions about the cryptosystem.

We only consider possibilistic anonymity here. An action by user u is considered to be anonymous when there exists some system in which u doesn't perform the action, and that system has an execution that is consistent with what the adversary sees. The actions for which we consider providing anonymity are sending messages, receiving messages, and communicating with a specific destination. More refined definitions of anonymity, [72, 20], incorporate probability. We apply such a definition to our system in Chapter 6, by defining a probability measure over executions or initial states.

The main result we show is that the adversary can determine a router in a given user's circuit if and only if it controls an adjacent router, with some other minor conditions. In particular, the adversary can determine which user owns a circuit only when the adversary controls the first hop. The set of users which have an uncompromised first hop form a sender "anonymity set," among which the adversary cannot distinguish. Similarly, the adversary can determine the last router of a circuit only when it controls the last router or

the penultimate router. Such circuits provide receiver anonymity. Also, a user and his destination have relationship anonymity when he has receiver anonymity or his sender anonymity set includes another sender with a destination that is different or unknown to the adversary.

The first-hop/last-hop attack is well-known [80], but we state it in full detail and show that, in a reasonable formal model, it is the only attack that an adversary can mount. Also, our results persist with or without some of the nontrivial design choices, such as multiple encryption and stream ciphers. This doesn't imply that these features are unnecessary – they may make attacks more difficult in practice and meet other goals such as data integrity, but it does illuminate their effect on the security of the protocol.

5.2   Anonymity Analysis

Definition 8. *A* protocol configuration $C : U \rightarrow \{(r_1, \ldots, r_l) \in R^l \times D \times \mathbb{N}_+ | \forall_i r_i \neq r_{i+1}\}$ *maps each user to the circuit, destination, and circuit identifier in his automaton state.*

For the remainder of this chapter, we will refer to a protocol configuration simply as a *configuration*.

The actions that we want to be performed anonymously are closely related to the circuits that the users try to construct during an execution and their destinations. Therefore, in order to prove that certain actions are performed anonymously in the network, we show that the adversary can never determine this circuit and destination information. This is a possibilistic notion of of anonymity. We do this by identifying classes of adversary-indistinguishable configurations.

Because $i \in N$ only sees those messages sent to and from i, an execution of a configuration C may appear the same to i as a similar execution of another configuration D that only differs from C in the circuit positions and destinations that are not adjacent to i and in circuit identifiers that i never sees. To be assured that i will never notice a difference, we would like this to be true for all possible executions of C. These are the fair cryptographic executions of C, and likewise the executions of D should be fair and cryptographic. We will say that these configurations are indistinguishable if, for any fair cryptographic execution of C, there exists a fair cryptographic execution of D that appears identical to i, *i.e.* in which i sends and receives what appear to be the same messages in the same order.

Agent i's power to distinguish among executions is weakened by encryption in two ways. First, we allow a permutation on keys to be applied to the keys of encrypted or decrypted messages in an execution. This permutation can map a key from any router other than i to any other key of any other router other than i,

because i can only tell that it doesn't hold these keys. It can map any key of i to any other key of i, because i doesn't know for which users and circuit positions its keys will be used. Second, i cannot distinguish among messages encrypted with a key he does not possess, so we allow a permutation to be applied to control messages that are encrypted with a key that is not shared with i. This second requirement must be justified by the computational intractability of distinguishing between encrypted messages with more than a negligible probability in our cryptosystem.

Definition 9. *Let D_A be a relation over configurations indicating which configurations are indistinguishable to $A \subseteq N$. For configurations C and C', $C \sim_{D_A} C'$ if for every fair cryptographic execution α of C, there exists some action sequence β such that the following conditions hold with C' as the initial state:*

1. *Every action of β is enabled, except possibly actions done by members of A.*

2. *β is fair for all agents, except possibly those in A.*

3. *β is cryptographic for all agents.*

4. *Let Ξ be the subset of permutations on the active keyspace $U \times R \times \lceil \frac{l}{2} \rceil$ such that each element restricted to keys involving $a \in A$ is a permutation on those keys. We apply $\xi \in \Xi$ to the encryption of a message sequence by changing every list component $\{p\}_{(u,r,i)}$ in the sequence to $\{p\}_{\xi(u,r,i)}$.*

   Let Π be the subset of permutations on $\bar{P}$ such that, for all $\pi \in \Pi$,

   (a) *π is a permutation on the set $\{\{p\}_{k_1,\ldots,k_i}\}_{p \in P}$*

   (b) *$\pi(\{p\}_{k_1,\ldots,k_i,k_a}) = \pi(\{p\}_{k_1,\ldots,k_i})$ when k_a is shared by the adversary*

   We apply $\pi \in \Pi$ to a message sequence by changing every message $\{p\}_{k_1,\ldots,k_i}$ in the message sequence to $\pi(\{p\}_{k_1,\ldots,k_i})$.

   Then there must exist $\xi \in \Xi$ and $\pi \in \Pi$ such that applying ξ and π to the subsequence of α corresponding to actions of A yields the subsequence of β corresponding to actions of A.

If $C \sim_{D_A} C'$, we say that C is *indistinguishable* from C' to A. It is clear that an indistinguishability relation is reflexive and transitive.

5.2.1   Anonymity

The sender in this model corresponds to the user of a circuit, the receiver to the destination of the circuit, and the messages we wish to communicate anonymously are just the circuit control messages. The circuit

identifiers allow the adversary to link together all the messages initiated by a user and attribute them to a single source. (Recall that in our model, users open a single circuit to a unique destination at one time.) Therefore sender anonymity is provided to u if the adversary can't determine which circuit identifier u is using. Similarly, receiver anonymity is provided to d for messages from u if the adversary can't determine the destination of the circuit with u's identifier. Also, relationship anonymity is provided to u and d if the adversary can't determine u's destination.

Definition 10. *User u has* sender anonymity *in configuration C with respect to adversary A if there exists some indistinguishable configuration C' in which u uses a different circuit identifier.*

Definition 11. *Destination d has* receiver anonymity *on user u's circuit, in configuration C, and with respect to adversary A, if there exists some indistinguishable configuration C' in which a user with u's circuit identifier, if one exists, has a destination other than d.*

Definition 12. *User u and destination d have* relationship anonymity *in configuration C if there is some indistinguishable configuration C' in which the destination of u is not d.*

5.3   Indistinguishable Configurations

Now we will show that sometimes the adversary cannot determine the circuit, destination, or circuit identifier. More specifically, an adversary can only determine which user, router, or destination occupies a given position in a circuit when the adversary controls it or a router adjacent to it on that circuit. Also, when the adversary controls no part of a circuit it cannot determine its identifier.

In order to do this, we must show that, given a pair of configurations (C, C') that are indistinguishable by these criteria, for every execution of C there exists an execution of C' that appears identical to the adversary. To do this we will start with a fair cryptographic execution of C, describe how to transform it, and prove that this transformed sequence forms a fair, cryptographic, and indistinguishable execution of C'. We will also show that a pair of configurations that are distinguishable by the described criteria allow no such transformation.

5.3.1   Message Sequences

To start, we observe that, in spite of the arbitrary actions of the adversary, the actions of the uncompromised users and routers are very structured. The protocol followed by the user and router automata defines a simple

sequence of message sends and receives for every circuit. A user or router will only send messages from the part of such a sequence consisting of its actions.

The user automaton gives this subsequence for users. It consists of messages between the user and the first router on its circuit, and is parameterized by the user u, the circuit c, the destination d, and the circuit identifier. We will refer to this sequence as $\sigma_U(u, c, d, j)$. Because the user automaton ignores the circuit identifier on received messages, we use an asterisk to indicate that any value is valid. Let $\sigma_U(u, c, d, j)$ be:

Step	From	To	Message
1	u	c_1	$[j, 0, \{\text{CREATE}\}_{k(u,c,1)}]$
2	c_1	u	$[*, 0, \text{CREATED}]$
3	u	c_1	$[j, 0, \{[\text{EXTEND}, c_2, \{\text{CREATE}\}_{k(u,c,2)}]\}_{k(u,c,1)}]$
4	c_1	u	$[*, 0, \{\text{EXTENDED}\}_{k(u,c,1)}]$
$1 + 2i$	u	c_1	$[j, 0, \{[\text{EXTEND}, c_{i+1}, \{\text{CREATE}\}_{k(u,c,i+1)}]\}_{k(u,c,i),...,k(u,c,1)}]$
$2 + 2i$	c_1	u	$[*, 0, \{\text{EXTENDED}\}_{k(u,c,i),...,k(u,c,1)}]$
$2 \leq i < l$			
1+2l	u	c_1	$[j, 0, \{[\text{DEST}, d, \text{HELLO}]\}_{k(u,c,l),...,k(u,c,1)}]$

Lemma 1. *For user u to be enabled to send a message in an action sequence under configuration C, the the following conditions must be satisfied:*

1. *The send appears in $\sigma_U(u, C(u))$.*

2. *The prefix of $\sigma_U(u, C(u))$ ending before the send has appeared in the sequence.*

3. *This prefix is the longest such prefix to appear.*

Proof. Every message sent by Automaton 1 appears in $\sigma_U(u, C(u))$. This proves that (1) is satisfied.

We show inductively that (2) is satisfied. Assume that (1)-(3) hold for the sends that appear in step $1 + 2(i-1)$, $0 < i \leq l$. Given some message that u is enabled to send, by (1) it appears in some step $1 + 2i$ in $\sigma_U(u, C(u))$. Examining where this message is sent in Automaton 1, we observe that step $2i$ must have occurred. Also, we observe that the automaton state variable b must equal i. For b to be set to i, the message in step $1 + 2(i-1)$ must have been sent, because b is only incremented after that event. By our inductive assumption, when that send occurred, steps 1 to $2(i-1)$ had occurred. Putting these together, we get the prefix in $\sigma_U(u, C(u))$ before step $1 + 2i$, showing (2). The inductive base case is trivially true, because the prefix of step 1 is empty.

Now, to show (3), we just observe that once step $1 + 2i$ occurs, b is incremented. Because b is never decremented, the condition guarding the send message of step $1 + 2i$ is never again satisfied. $\qquad \square$

Similarly, the router automaton defines the action sequence that a router performs during the creation of a circuit. A different sequence exists for every router r, user u, circuit position $1 \le i \le l$, circuit c, destination d, and link identifiers $m, n, p \in \mathbb{N}$. We will denote a particular sequence $\sigma_R(r, c, d, u, i, m, n)$. Frequently we will drop parameters that we don't care about, for example, referring to $\sigma_R(r, c, d, u, i)$ when the specific link identifiers don't matter, and may abuse this notation by treating it as one sequence rather than a family of sequences. We use $k(u, c, i)$ as before.

If $i < l$, the sequence $\sigma_R(r, c, d, u, i, m, n)$ is:

Step	From	To	Message
1	c_{i-1}	r	$[j_1, n, \{\text{CREATE}\}_{k(u,c,i)}]$
2	r	c_{i-1}	$[j_1, n, \text{CREATED}]$
3	c_{i-1}	r	$[j_2, n, \{[\text{EXTEND}, c_{i+1}, s]\}_{k(u,c,i)}]$
4	r	c_{i+1}	$[j_2, m, s]$
5	c_{i+1}	r	$[j_3, \text{CREATED}]$
6	r	c_{i-1}	$[j_3, \{\text{EXTENDED}\}_{k(u,c,i)}]$

If $i = l$, the sequence $\sigma_R(r, c, d, u, i, m, n)$ is:

Step	From	To	Message
1	c_{i-1}	r	$[j_1, n, \{\text{CREATE}\}_{k(u,c,i)}]$
2	r	c_{i-1}	$[j_1, n, \text{CREATED}]$
3	c_{i-1}	r	$[j_2, n, \{[\text{DEST}, d, s]\}_{k(u,c,l)}]$
4	r	c_{i+1}	$[j_2, m, s]$

Using the σ_R sequences, we can characterize which messages a router can send at any point in an action sequence. Let α be a finite execution, and τ_i be the length i prefix of some sequence $\sigma \in \sigma_R(r)$. We say that τ_i has *occurred* in α when τ_i is the longest prefix of σ that appears as a subsequence of α, and, at the point at which step 1 (4) occurs in α, n (m) must be the smallest number not yet in r's table as an entry to or from c_{i-1} (c_{i+1}), the agent that sent (received) the message in step 1 (4).

Lemma 2. *For r to be enabled to send a message μ at the end of α, one of three cases must apply for some $\sigma \in \sigma_R(r)$:*

1. *The message is part of the circuit-building protocol. Sending μ is the $(2i)^{th}$ step in σ, $1 \le i \le 3$. Then*

τ_{2i-1} must occur in α and τ_{2i} must occur in the execution created when μ is appended to α.

2. *The message is a forward down the circuit. $\mu = [m, p]$. r is to be sending μ to c_{i+1}. σ has occurred in α. The link identifiers to c_{i-1} and c_{i+1} are n and m, respectively. α contains the action of the message $[n, p]$ being sent to r from c_{i-1} after σ occurs. Also, the number of such receives from c_{i-1} is greater than the number of sends of μ to c_{i+1} that happen after σ occurs.*

3. *The message is a forward up the circuit. $\mu = [n, p]$. r is to be sending μ to c_{i-1}. σ has occurred in α. The link identifiers to c_{i-1} and c_{i+1} are n and m, respectively. α contains the action of the message $[m, p]$ being sent to r from c_{i+1} after σ occurs. Also, the number of such receives from c_{i+1} is greater than the number of sends of μ to c_{i-1} that happen after σ occurs.*

Proof. Suppose r is enabled to send μ at the end of α. This must occur as one of the send operations in Automaton 2, which occur in steps 9, 15, 18, 20, and 24.

If the send is enabled in line 6,9, 15, or 24, then it appears as Step 4, 4, 6, or 2, respectively, of a $\sigma_R(r)$ sequence. Then the message is part of a circuit-building protocol, and the case (1) will apply. For each step $2i$ in $\sigma_R(r)$, $1 \le i \le 3$, the argument is the same. The condition guarding the send in Automaton 2 includes the receipt of the message in step $2i - 1$ of $\sigma_R(r)$. It also requires the table to have an entry that is only made when the message in step $2i - 2$ is sent. Thus τ_{2i-1} must be a prefix in α for r to be enabled to send μ. Moreover, once μ is sent, the table row is changed, preventing μ from being sent again. Therefore τ_{2i-1} is the longest prefix of τ_i. Lines 6 and 22 of Automaton 2 ensure that the minimum-identifier condition was satisfied when 4 and 1 of $\sigma_R(r)$ were executed, respectively. Therefore τ_{2i-1} occurred in α. Sending μ executes the next step of $\sigma_R(r)$. It respects the minimal-identifier condition as part of the automaton specification, and because τ_{2i} becomes the new longest prefix, τ_{2i} occurs.

If the send is enabled in line 18 Automaton 2, it is a forward down the circuit, and case (2) will apply. Take m, n, p, c_{i-1}, and c_{i+1} in the statement of the case to be n, m, p, r, and s in Automaton 2, respectively. In order for the send to be enabled in line 18 of Automaton 2, there must exist the entry in the table T that is inserted, in line 10 or line 14, when the message in step 4 or step 6 of $\sigma_R(r)$ is sent, respectively. Thus, by case (1), which we have already shown, a σ_R sequence must have occurred in α. The MESSAGE procedure gets called once for every time a message $[n, p]$ is sent to r from c_{i-1}, and therefore line 18 must have been executed fewer times than the number of receives from c_{i-1} of μ after σ occurs.

The case for a send enabled by line 20 of Automaton 2 is a forward up the circuit, and case (3) applies.

The argument is exactly the same as for case (2), with c_{i-1} and c_{i+1} switched. □

We can use these lemmas to partition the actions performed by an agent in an execution of configuration C. We will use these partitions to construct executions of indistinguishable configurations and prove that they satisfy the requirements of Definition 9.

For a user u we create two partitions. The first is formed by the maximal prefix of $\sigma_U(u, C(u))$ such that each receive in the partition causes the b variable of u's state to be incremented. The condition on the receives is required for a unique maximal prefix to deal with the case that an adversary sends sequence responses multiple times. The second partition is formed from all of u's other actions. By Lemma 1 this is composed of receiving unnecessary messages due to adversarial actions, and we will call this the "junk" partition.

For a router r, we create a partition for each entry in its routing table at any point in the execution and an extra junk partition. For a given routing table entry we create a partition out of the maximum-length subsequence of some $\sigma_R(r)$ sequence, say, σ, such that each receive modifies the same entry in the routing table. We also include every send and receive of a forward performed using that entry. This partition is said to be associated with σ. Every other action done by the router is put in a junk partition, and, by Lemma 2, this partition is composed only of receives.

5.3.2 Indistinguishable Users

Now we prove that an active adversary cannot determine which user creates a given circuit unless the first router on that circuit is controlled by the adversary or the owners of all the other circuits have been determined. That is, an adversary cannot distinguish between a configuration C and the configuration C' that is identical to C except for two circuits with uncompromised first routers that are switched between their owners. In order to do so, we must show that, for any fair, cryptographic execution of C, there exists some action sequence of C' satisfying the indistinguishability requirements of Definition 9. To do so, we simply swap between the switched users the messages that pass between them and the first routers on their circuits and switch the encryption keys of these messages.

Theorem 1. *Let u, v be two distinct users such that neither they nor the first routers in their circuits are compromised (that is, are in A). Let C' be identical to C except the circuits of users u and v are switched. C is indistinguishable from C' to A.*

Proof. Let α be a fair, cryptographic execution of C. To create a possible execution of C', first construct α' by replacing any message sent or received between u (v) and $C_1(u)$ $(C_1(v))$ in α with a message sent or received between v (u) and $C_1(u)$ $(C_1(v))$. Then let ξ be the permutation that sends u to v and v to u and other users to themselves. Create β by applying ξ to the encryption keys of α'.

1. *Every action by an agent in $N \setminus A$ in β is enabled.* All receives are enabled in β, because sends and corresponding receives are modified together.

   For any user $w \notin \{u, v\}$, all messages in $\sigma_U(w, C(w))$ go to or from w, so none are added or removed from α in α'. Also none of the messages in this sequence would be modified by ξ because they are encrypted with a key of w, and ξ doesn't convert messages in α' to be messages of $\sigma_U(w, C(w))$. Therefore if a message is enabled to be sent from w in β it was enabled in α.

   For user u, when u sends a message to $C_1(v)$ in β, it corresponds to v sending a message to $C_1(v)$ in α. v is enabled to do so in α so at that point it has sent and received exactly those messages of $\sigma_U(v, C(v))$ necessary to enable that send. In β we have changed those messages to be messages between u and $C_1(v)$ while modifying the encryption keys, so the necessary $\sigma_U(u, C'(u))$ messages have appeared to enable the send. No additional messages in $\sigma_U(u, C'(u))$ could have appeared in β, because u and v do not communicate on link identifier 0 in α. Therefore u is enabled to send the message. A similar argument works for v.

   For a router $r \notin A \cup \{C_1(u), C_1(v)\}$, the only change in messages to or from r between α and β is from the permutation ξ applied to the encryption keys of the messages. Applying ξ preserves the occurrence of some prefix of $\sigma_R(r, C'(w), w)$ at any point in the execution, for any $w \notin \{u, v\}$. For $w = u$, applying ξ turns the occurrence of some $\sigma_R(r, C(u), u)$ into an occurrence of $\sigma_R(r, C'(u), v)$, and vice versa for $w = v$. It also preserves the appearance of messages to forward and the actual forwarding. Thus any action performed by r in β is enabled because it corresponds to a similar enabled action in α.

   Now consider a message μ sent from $C_1(u)$ in β.

   It may be that μ is part of a $\sigma_R(C_1(u), C(w), w)$ sequence for some $w \notin \{u, v\}$ in α. Then μ is enabled in β since none of the messages in $\sigma_R(C_1(u), C(w), w)$ come from u or v and none involve u or v in the encryption keys, so all exist in β that did in α and no additional ones do. It could also be that, in α, μ is a forward in some circuit not belonging to u or v. Then μ is still enabled in β for a similar reason, recognizing that although it might involve the encryption keys of u or v, the content of messages is

42

ignored in forwards.

Another case is that, in α, μ is part of some $\sigma_R(C_1(u), C(u), u, i)$. Then in β, μ is part of $\sigma_R(C_1(u), C'(v), v, i)$. This is because every message from u to $C_1(u)$ is changed to a message from v to $C_1(u)$ and every encryption key involving u changes to one involving v. It can be seen by inspecting the messages in a σ_R sequence that consistently replacing the user in the encryption keys in a σ_R sequence and (when $i = 1$) the previous hop from u to v, as is done to create β, transforms a $\sigma_R(C_1(u), C(u), u, i)$ sequence into a $\sigma_R(C_1(u), C'(v), v, i)$ sequence. No additional messages can enter into this sequence in β because they must be encrypted with a key of v, and any such message will have appeared with a key of u in α and will have filled the same spot in the $\sigma_R(C_1(u), C(u), u, i)$ sequence there. Thus μ is enabled in β. Also, for similar reasons, if μ is a forward in u's circuit in α, then it is a forward for v's circuit in β.

The final case is when, in α, μ is in a $\sigma_R(C_1(u), C(v), v)$ sequence or is a forward on v's circuit. Since v does not communicate directly with $C_1(u)$ as a user in α, it must be that $C_1(u)$ is some intermediate router. Then the only changes to the $\sigma_R(C_1(u), C(v), v)$ messages in α are the encryption keys, which are applied consistently to all the sequence messages and are not interfered with by messages in α already using the target keys since they are also modified. Therefore if μ corresponds to a $\sigma_R(C_1(u), C(v), v)$ send in α, it is a $\sigma_R(C_1(u), C'(u), u)$ message enabled in β. Also, for similar reasons, if μ was a foward in v's circuit in α, it is an enabled forward on u's circuit in β.

The analogous argument works for $C_1(v)$.

2. *β is fair for agents in $N \setminus A$.*

   For any user $w \notin \{u, v\}$, every $\sigma_U(w, C'(w))$ message received in β in its non-junk partition is the same message in α. Therefore every send w is enabled to perform in β it is enabled to perform in α. Because α is fair for w so is β.

   Now consider u. The transformation properly changes the messages from $C_1(v)$ to v in $\sigma_U(v, C(v))$ to messages sent to u that are in the same position in the $\sigma_U(u, C'(u))$ sequence. No extra messages can appear since they must be encrypted using u's keys, and then they would have been encoded with v's keys in α and been part of the $\sigma_U(v, C(v))$ sequence there. Therefore every send that u is enabled to perform in β, v is enabled to perform in α. Since α is fair for v, then β is fair for u. The analogous argument works for v.

For router $r \notin A \cup \{C_1(u), C_1(v)\}$ to be enabled to perform a send in β but not α, there must be a message in some sequence $\sigma_R(r, C'(w), w, i)$ that r receives in β but doesn't in the corresponding sequence in α. This cannot be for any $w \notin \{u, v\}$, because the messages in this σ_R are not modified, except possibly the content of forwards which doesn't affect their validity. All messages in β that are to r and are in some $\sigma_R(r, C'(u), u)$ are also are sent to r in α and are part of some $\sigma_R(r, C'(v), v)$. Therefore if such a message enables r to send something in β there exists a similar message enabling r to send something in α. Also forwards along u's circuit in β exist as forwards along v's circuit in α. The analogous argument works for messages of some sequence $\sigma_R(r, C', v)$.

For $C_1(u)$ to be enabled to perform a send in β but not α, there must be a message in some sequence $\sigma_R(C_1(u), C'(w), w)$ or forward that $C_1(u)$ receives in β but doesn't in α. There can not be such a message in the $\sigma_R(C_1(u), C'(w), w)$ sequence for any $w \notin \{u, v\}$, because the messages in this sequence are not modified in the transformation and no new messages encrypted with w's key are created. Also the sender and recipient of forwards from w aren't modified, and the content of forwards which doesn't affect their validity. Now suppose $w = v$. For any message that appears at the end of some $\sigma_R(C_1(u), C'(v), v, i)$ in β that $C_1(u)$ doesn't respond to there must not be an analogous message in $\sigma_R(C_1(u), C(u), u, i)$ in α or $C_1(u)$ would be enabled at that point as well. But this message must be encrypted with v's keys and would be modified by the ξ permutation and thus play the same role for $C_1(u)$ in α. Again, the content of forwards doesn't matter and any forward on v's circuit in β corresponds to a forward on u's circuit in α. The analogous argument works for the case $w = u$. Therefore every send enabled for $C_1(u)$ in β is enabled in α, and β is fair for $C_1(u)$. The analogous argument works for $C_1(v)$.

3. *β is cryptographic.*

We've already shown that uncompromised routers and users perform enabled actions in β. Since the automatons only allow sending messages encrypted with keys the agent doesn't possess after receiving them, the actions of these agents do not prevent β from being cryptographic. For a compromised user or router, let's say that a control message encrypted with a foreign key is sent before being received at least once. If the encryption key doesn't involve u or v, then the same message gets sent in α before being received, contradicting the fact that α is cryptographic. If the key does involve u, then in α it involves v, in which case if the message is received in α beforehand, it must have received it in β since the key permutation takes v to u. Likewise for messages encrypted with one of v's keys. The fact that

α is cryptographic then implies that β is cryptographic.

4. *We can find a $\xi \in \Xi$ and $\pi \in \Pi$ that turn α into a sequence that agrees with β in all the adversary actions.*

   ξ is simply the user permutation used to create β, transposing users u and v, and π is the identity on all messages. Applying these to α yields a sequence that agrees with β everywhere except for messages between u (v) and $C_1(u)$ $(C_1(v))$, which we assumed are not adversarial.

$\square$

5.3.3 Indistinguishable Routers and Destinations

Now we prove that an adversary cannot determine an uncompromised router or destination on a given circuit unless it controls the previous or next router on that circuit. The proof is similar to that of Theorem 1, although it is complicated by the fact that the identities of routers in a circuit are included in multiple ways in the circuit creation protocol. Given an execution of C, we identify those messages that are part of the circuit creation sequence of the modified circuit and then change them to change the router or destination. Then we show that, in the sense of Definition 9, from the adversary's perspective this sequence is indistinguishable from the original and could be an execution of C'.

Theorem 2. *Say there is some user $u \notin A$ such that u's circuit in C contains three consecutive routers, $r_{i-1}, r_i, r_{i+1} \notin A$. Let C' be equal to C, except r_i is replaced with r_i' in u's circuit, where $r_i' \notin A \cup \{r_{i-1}, r_{i+1}\}$. C' is indistinguishable from C to A. The same holds for uncompromised routers (r_i, r_{i+1}) if they begin u's circuit and are replaced with (r_i', r_{i+1}), uncompromised routers (r_{i-1}, r_i), if they end u's circuit and are replaced with (r_{i-1}, r_i'), or uncompromised router and destination (r_{i-1}, r_i) if they are replaced with (r_{i-1}, r_i').*

Proof. Let α be some fair cryptographic execution of C, and let $h(C(u), i)$ denote the number of occurrences of the ith router in the circuit $C(u)$ among the first i routers. We modify α in steps to create an indistinguishable sequence β:

1. If r_i is a router, replace all message components of the form
   $[\text{EXTEND}, r_i, \{\text{CREATE}\}_{(u,r_i,h(C(u),i))}]$ with $[\text{EXTEND}, r_i', \{\text{CREATE}\}_{(u,r_i',h(C'(u),i))}]$. If r_i is a destination, replace all message components of the form $[\text{DEST}, r_i, b]$ with $[\text{DEST}, r_i', b]$.

2. Consider the partition of router r_{i-1}'s actions that are associated with a $\sigma_R(r_{i-1}, C(u), u, i-1)$ sequence. Replace all messages in this partition that are to and from r_i with the same messages to and from r_i'. Modify the link identifiers on these messages so that they are the smallest identifiers in use between r_{i-1} and r_i' at that point in α. Increase link identifiers that are in use between r_{i-1} and r_i' to make room for these new connections and decrease link identifiers that are in use between r_{i-1} and r_i to fill in the holes created by the removed connections. Perform similar modifications for routers r_i and r_{i+1}.

3. Replace all encryption keys of the form $(u, r_i, h(C(u), i))$ with $(u, r_i', h(C'(u), i))$. Increment as necessary the third component of the encryption keys used between u and r_i' to take into account that r_i' appears once more in $C'(u)$ than it does in $C(u)$. Also decrement as necessary the third component of the keys used between u and r_i to take into account that r_i appears once less in $C'(u)$ than it does in $C(u)$.

Now we show that the action sequence thus created, β, is a fair cryptographic execution of C':

1. *Every action by an agent in $N \setminus A$ in β is enabled.*

   It is easy to see that all receives are enabled in β since sends and corresponding receives are modified together.

   Our strategy to show that all sends in β are enabled will be to consider the separate non-junk partitions of α after the transformation. First we will show that no sends from uncompromised agents appear in β outside of these transformed partitions. Then we show that any given non-junk partition of α is transformed into a subsequence that is "locally" enabled under C'. A user (router) action sequence is locally enabled if each send satisfies the conditions of Lemma 1 (2) applied just to that sequence. Then we show that it is "globally" enabled in the sense that the sends in the transformed user (router) partition continue to satisfy Lemma 1 (2) when considered over the entire sequence β.

   It is easier to proceed this way since going from a locally to globally-enabled sequence just requires that certain actions *don't* exist in the larger sequence. For users, none of the sends in the transformed non-junk partition can appear again in the larger sequence between being enabled and being sent in the partition. This must also be true for transformed router partitions, and additionally the link identifiers used must be unique and minimal at the point of link creation. Via Lemmas 1 and 2 a locally-enabled action sequence satisfying these global conditions contains only enabled sends in β.

The fact that there are no sends from uncompromised agents in β outside of the transformed non-junk α partitions helps us prove that actions are globally enabled. By inspecting the three changes made to α, it is clear that no actions are added or deleted from α, and that sends (receives) in α are sends (receives) in β. Since every send in α from an agent $a \in N \setminus A$ is part of one of its non-junk partitions, every send by an uncompromised agent in β is part of one of the transformed partitions.

We can use this to show that a given sequence is globally enabled. If the sequence is a locally-enabled transformed user partition, it is automatically globally enabled because there are no sends outside the partition to interfere with it. Similarly, for locally-enabled transformed router partitions, we automatically satisfy the send non-inteference property in β as long as we satisfy the second requirement on the link identifiers.

This second requirement for routers is slightly simpler to achieve by noting that all CREATE messages in β were transformed from CREATE messages in α. Therefore we only need to show that the link IDs used in β are unique and minimal among the link creations in α after transformation.

For user $v \neq u$ the non-junk partition has not been modified, and therefore, by Lemma 1, every send is locally enabled in β. Therefore every action by v is enabled.

Now consider the user u's non-junk partition in α. We've modified steps $2i-1,2i,2i+1$, and $2i+2$ as necessary to change the $\sigma_U(u, C(u))$ prefix to a $\sigma_U(u, C'(u))$ prefix. All these are enabled by Lemma 1, and so this is locally enabled. No sends appear outside of this transformed partition in β. Thus the partition is globally enabled.

Now consider a router $r \notin \{r_{i-1}, r_i, r'_i, r_{i+1}\}$ and a partition of r in α. The partition consists of a prefix of a $\sigma_R(r, C(w))$ sequence, for some user w, and possibly some forwards. The only changes made to the partition are key relabelings and some modification of the messages of forwards. The relabeling turns the $\sigma_R(r, C(w))$ prefix into some $\sigma_R(r, C'(w))$ prefix of the same length, and so sends in this sequence are locally enabled. Forwards are enabled regardless of content, and so they are also locally enabled. No link identifiers of r have changed, and so they are still unique and minimal. Therefore the whole partition is globally enabled.

Now consider r_{i-1}. Take some non-junk partition of α that does not contain a sequence in which r_{i-1} is the $(i-1)$th circuit router of u, that is, that does contain a $\sigma_R(r_{i-1}, C(w), w, j)$ sequence, $w \neq u$ or $j \neq i-1$. For the same reasons as the preceding case, it is transformed into a sequence

associated with a $\sigma_R(r_{i-1}, C'(w), w)$ sequence. Thus it is locally enabled. The partition that is a prefix of $\sigma_R(r_{i-1}, C(u), u, i-1)$ can be seen by inspection to be modified to be a locally-enabled sequence associated with $\sigma_R(r_{i-1}, C'(u), u, i-1)$. The link identifiers used in every transformed partition of r_{i-1} are unique and minimal in β because the original partitions had unique and minimal IDs in α, we haven't changed the IDs or neighbors of any partitions not connecting with r'_i or r_i, and we have changed the ID in partitions connecting with r'_i or r_i to make IDs unique and minimal after changing a partition to connect with r'_i instead of r_i. Thus the whole sequence is globally enabled. The analogous arguments work for r_{i+1}, r_i and r'_i.

2. *β is fair for agents in $N \setminus A$.*

   To show this we will again consider the transformed partitions of α. We have shown that they form enabled sequences, and now need to show that no messages from the transformed junk partition belong in these sequences. For users, this means that the next step in a transformed σ_U partition isn't received. For routers, it means that the next step in a transformed σ_R partition isn't received, no new forwards on a created circuit are received, and that no new valid CREATE messages are received.

   Consider a user $w \neq u$. Every receive action by w in β is from a receive action by w in α. The messages received by w are never modified to use one of w's keys, so a message encrypted with w's keys in β uses the same keys in α. Also the content of an encrypted message is never changed to be a message that appears in $\sigma_U(w, C'(w))$. Therefore any receive that is a step of $\sigma_U(w, C'(w))$ in β is the same step in $\sigma_U(w, C(w))$ in α. Therefore β is fair for w.

   Now consider user u. As shown, the transformed non-junk partition in α is a locally-enabled sequence in β. For u to have an unperformed enabled action in β, the next message in the $\sigma_U(u, C'(u))$ sequence must come from the junk sequence and be unanswered. $\sigma_U(u, C(u))$ and $\sigma_U(u, C'(u))$ differ in steps $(2 + 2j)$, $i \leq j \leq l$. These only differ in the encryption keys, and these have been renumbered in such a way that the $(2 + 2j)$th step in $\sigma_U(u, C'(u))$ appears if and only if the $(2 + 2j)$th step appears in $\sigma_U(u, C(u))$. Thus an enabling receive in β implies that such a receive exists in α. Therefore β is fair for u.

   Now consider a router $r \notin \{r_{i-1}, r_i, r'_i, r_{i+1}\}$. For a given transformed partition, no new messages of the associated $\sigma_R(r)$ sequence can appear in β. This is because $\sigma_R(r)$ messages are all encrypted for r and we have created no such messages, nor have we modified such messages so as to create a new message

in the $\sigma_R(r)$ sequence. For forwards, first we recognize that the transformation maintains source and link ID consistency for r in the sense that if we were to group r's receives in α by their source and link ID the transformed groups would be the same as the same groups created in β. Therefore for an incoming message to be transformed into a forward on a created circuit, it must previously be sent with the link identifiers of the incoming link, but since content in forwards doesn't matter, this would be a forward in α as well. Finally there are no valid CREATE messages received in β that aren't received in α. No new messages are sent to r, no messages are transformed into a CREATE, no keys have been modified to belong to r, and r's link IDs have been consistently changed. Therefore β is fair for r.

Now consider r_{i-1}. For a transformed partition of r_{i-1}, suppose that some receive extends the associated $\sigma_R(r_{i-1})$ or acts as a forward on the created circuit. This receive must be from the junk partition of r_{i-1} since we have shown its that non-junk partitions form enabled sequences. This message can't be from a router not in $\{r_i, r_i'\}$ because all such messages existed in α with the same link ID, source, and destination, and the content is either the same or is a forward, which would still be a forward in α. It can't come from r_i. This router is uncompromised and therefore properly uses link identifiers. If the message were part of the associated $\sigma_R(r_{i-1})$ sequence in β, it would exist in α with the ID in use by r_{i-1} and r_i in the sequence in α and with the same content, so this can't be the case. If the message were a forward, again it would exist in α with the link ID in use between r_{i-1} and r_i in the partition, and would therefore be a forward in α as well. Similar arguments work for messages from r_i'. Finally, no new partitions can be created, because CREATE messages to r_{i-1} on unique link IDs in β are the same in α. Therefore β is fair for r_{i-1}. The analogous arguments work for r_{i+1}.

Now consider r_i. Because only link identifiers to r_{i-1} and r_{i+1} have been changed, and those routers are uncompromised, all messages in β to r_i from a given router and with a given link ID are transformed from all messages in α from the same router and of a given (possibly different) link ID. Suppose a message receive in β enables a send in the associated $\sigma_R(r_i)$ that is not enabled in the related sequence in α, or acts as a new forward. It must have been sent in α on the link ID in α of that partition. Because messages aren't redirected to r_i and senders aren't changed, it must have been sent in α by the same sender. Because content doesn't change in the $\sigma_R(r_i)$ messages and doesn't matter in forwards this message would perform the same function in α, contradicting the fairness of α. No new partitions can be created because new CREATE messages aren't made by the transformation, senders are the same, and link identifiers are renumbered in such a way that distinct link IDs from a router in α are

distinct in β. Therefore β is fair for r_i.

A similar argument works for r_i' over its partitions in α, but we do reassign a partition of r_i to r_i', which we must also consider. Notice that the messages redirected to r_i' exist on unique link IDs in β with r_{i-1} and r_{i+1} in β. Therefore these cannot enable actions on the other transformed partitions, and vice versa. Also no junk messages of r_i' can enable actions in this transformed partition because the connecting routers, r_{i-1} and r_{i+1}, are uncompromised and will have sent these messages on a link ID that is different from the ID of the transformed new partition in β. Finally, we show that the only valid CREATE messages received by r_i' in β are those in transformed partitions of α. Every CREATE in β is a CREATE in α. Every valid CREATE to r_i' in α becomes a valid CREATE in β because it is part of a transformed partition and we have shown that these become enabled sequences. The only messages redirected to r_i' belong to r_i's reassigned partition, which forms a fair sequence in α and maintains this after transformation. The final possibility for a new CREATE is a CREATE message from r_i''s junk partition that was sent to r_i' in α but was encrypted with a key of r_i, which then gets changed in the transformation. The only such message is $\{\text{CREATE}\}_{(u, r_i, h(C(u),i))}$. Only r_{i-1} could send this message in α because α is cryptographic, and u only produces such a message encrypted with the key of r_{i-1}'. It would only send this message if it were to receive the message $\{\text{EXTEND}, r_i', \{\text{CREATE}\}_{(u, r_i, h(C(u),i))}\}_{(u, r_{i-1}, h(C(u),i-1))}$, which u never sends in α. Again by the cryptographic property of α r_{i-1} never sends this CREATE to r_i' in α. Thus no valid CREATE messages are received by r_i' in β that are not transformed from partitions in α, and we have shown that these transformed partitions are fair. Therefore β is fair for r_i'.

3. *β is cryptographic.*

For uncompromised routers, the fact that all sends are enabled in β guarantees cryptographic sends, because the protocol ensures this property. Compromised routers send and receive all the same messages, but to which the transformation function has been applied. Therefore because α is cryptographic, β is.

4. *We can find key and message permutations that turn β into a sequence that agrees with α in all adversary actions.*

No messages are redirected towards or away from $a \in A$ when constructing β. We apply the message permutation to β of transposing

$\{[\text{EXTEND}, r'_i, \{\text{CREATE}\}_{(u,r'_i,h(C'(u),i))}]\}_{k_1,...,k_j}$ and

$\{[\text{EXTEND}, r_i, \{\text{CREATE}\}_{(u,r_i,h(C(u),i))}]\}_{k_1,...,k_j}$, $1 \le j \le l$, where k_j isn't shared by the adversary. (If we replaced a destination we instead transpose $\{[DEST, r'_i, b]\}_{k_1...,k_j}$ and $\{[DEST, r_i, b]\}_{k_1...,k_j}$.) We also apply the key permutation that sends $(u, r'_i, h(C'(u), i))$ to $(u, r_i, h(C(u), i))$ and undoes the renumbering of the r'_i and r_i keys. Then the subsequence of actions by a in β is identical to the subsequence in α.

$\square$

5.3.4 Indistinguishable Identifiers

Theorem 3. *Say there is some uncompromised user u such that all routers and the destination in $C(u)$ are uncompromised. Then let C' be a configuration that is identical to C, except that u uses a different circuit identifier. C' is indistinguishable from C to A.*

Proof. Let α be a fair, cryptographic execution of C. To create β, simply change every occurrence of u's circuit identifier in C ($C_{l+2}(u)$) to its identifier in C'. β is enabled, fair, and cryptographic for C' because no message containing $C_{l+2}(u)$ gets sent to the adversary in α and the protocol itself ignores circuit identifiers except to forward them on. It appears the same to A for the same reason. $\square$

5.3.5 Indistinguishable Configurations

Let D'_A be the transitive closure of pairs that are indistinguishable by Theorems 1, 2, and 3. Theorems 1, 2, and 3 will allow us to characterize exactly which configurations are indistinguishable.

Lemma 3. *D'_A is an equivalence relation on indistinguishable configurations.*

Proof. If $C_1 \sim_{D_A} C_2$, for some configurations C_1 and C_2, then, for every execution α of configuration C_1, there exists an indistinguishable execution β of C_2. Then if $C_2 \sim_{D_A} C_3$, for some configuration C_3, there must exist an execution γ of C_3 that is indistinguishable from β. α and γ are then indistinguishable, and thus indistinguishability of configurations is transitive. Therefore taking a transitive closure of indistinguishable pairs, as we do to create D'_A, preserves the indistinguishability of all pairs.

Theorems 1, 2, and 3 are all symmetric, in that if, for configurations C_1 and C_2, $C_1 \sim C_2$ by one of the theorems, then $C_2 \sim C_1$ by the same theorem. Therefore D'_A is symmetric. D'_A is transitive by

construction. D'_A is reflexive because a configuration is trivially indistinguishable from itself. Therefore D'_A is an equivalence relation. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\square$

We introduce some notation to conveniently refer to the configurations related by D'_A.

Definition 13. *For configurations C and D, we say that $C \approx_{D_A} D$ if C and D are related by a chain of configurations that are indistinguishable by Theorems 1, 2, and 3.*

We can easily tell which configurations are in the same equivalence class using the following function. It reduces a circuit to an identifier, the compromised positions, and the positions adjacent to compromised positions.

Definition 14. *Let $\rho : U \times N^l \times D \times \mathbb{N}_+ \times \mathcal{P}(N) \to \mathbb{N} \times \mathcal{P}(N \times \mathbb{N}_+)$ be:*

$$\rho(u, c, d, j, A) = \begin{cases} (j, \{(r, i) \in N \times \mathbb{N}_+ | c_i = r \wedge (c_{i-1} \in A \vee c_i \in A \vee c_{i+1} \in A)\}) & \text{if } c_i \in A \text{ for some } i \\ (0, \emptyset) & \text{otherwise} \end{cases}$$

In the preceding let c_0 refer to u and c_{l+1} refer to d.

We overload this notation and use $\rho(C)$ to refer to the multiset formed from the circuits of configuration C adjoined with their user and reduced by ρ. That is, $\rho(C) = \{\rho(u, C(u), A) | u \in U\}$. The following lemma shows that ρ captures the indistinguishable features of a configuration according to Theorems 1, 2, and 3.

Lemma 4. *Let C and D be configurations. $C \approx_{D_A} D$ if and only if $\rho(C) = \rho(D)$.*

Proof. First we show that if $\rho(C) = \rho(D)$, then $C \approx_{D_A} D$. Because $\rho(C) = \rho(D)$, there is a permutation π of the users such that $\rho(C(u)) = \rho(D(\pi(u)))$. Let C' be the configuration such that, for every user u, we replace the routers in $C(u)$ that are not compromised or adjacent to compromised routers by the routers in the same circuit position in $D(\pi(u))$, and such that we do the same for an uncompromised destination. C' is indistinguishable from C by Theorem 2. Let C'' be the configurations that results when we permute the users of C' by π. π must act as the identity on all routers that are compromised or have compromised first routers, because otherwise $(u, 0)$ would exist in $\rho(C(u))$ and would not exist in $\rho(D(\pi(u)))$, contradicting their equality. Therefore, C'' is indistinguishable from C' by Theorem 1. Finally, let C''' be the result of changing, for all users such that $\rho(u) = (0, \emptyset)$, the circuit identifiers $C''_{l+2}(u)$ to $D_{l+2}(u)$. C''' is indistinguishable from C'' by Theorem 3, because $\rho(C(u)) = (0, \emptyset)$ implies that u and $C_i(u)$ are uncompromised, $1 \leq i \leq l$.

For all users u, if $C_i'''(u)$, $C_{i-1}'''(u)$, and $C_{i+1}'''(u)$ are uncompromised, $1 \leq i \leq l+1$, then

$$C_i'''(u) = C_i'(\pi^{-1}(u)) \tag{5.1}$$

$$= D_i(\pi(\pi^{-1}(u))) \text{ by construction.} \tag{5.2}$$

If $C_i'''(u)$, $C_{i-1}'''(u)$, or $C_{i+1}'''(u)$ is compromised, $1 \leq i \leq l+1$, then

$$C_i'''(u) = C_i'(\pi^{-1}(u)) \tag{5.3}$$

$$= C_i(\pi^{-1}(u)) \tag{5.4}$$

$$= D_i(\pi(\pi^{-1}(u))) \text{ because } \rho(C(u)) = \rho(D(\pi(u))). \tag{5.5}$$

If $C'''(u) = (0, \emptyset)$, then the circuit identifier $C_{l+2}'''(u) = D_{l+2}(u)$ by construction. If $C'''(u) \neq (0, \emptyset)$, then $C_{l+2}'''(u) = D_{l+2}(u)$, because

$$\rho(C'''(u)) = \rho(C(\pi^{-1}(u))) \tag{5.6}$$

$$= \rho(D(\pi(\pi^{-1}(u)))), \tag{5.7}$$

and the circuit identifier is included in $\rho(c)$ when c contains compromised agents. Therefore $C''' = D$, and thus $C \approx_{D_A} D$.

Now we show that if $C \approx_{D_a} D$, then $\rho(C) = \rho(D)$. Because $C \approx_{D_A} D$, there must be a chain of indistuinguishable configurations relating C to D such that each adjacent pair in the chain is indistinguishable by Theorems 1, 2, and 3. For any pair of configurations C_1 and C_2 that are indistinguishable by one of these theorems, $\rho(C_1) = \rho(C_2)$:

1. If C_1 and C_2 are related by Theorem 1, a pair of uncompromised users with uncompromised first routers is swapped. ρ is invariant under this operation.

2. If C_1 and C_2 are related by Theorem 2, then a router or destination r_i such that r_i, r_{i-1}, and r_{i+1} are uncompromised is replaced by r_i'. ρ is invariant under this operation.

3. If C_1 and C_2 are related by Theorem 3, the circuit identifier of an uncompromised circuit is modified. ρ is invariant under this operation.

Therefore $\rho(C) = \rho(D)$. □

Now we show that the equivalence relation is in fact the entire indistinguishability relation and that Theorems 1, 2, and 3 characterize which configurations are indistinguishable. The reason for this is that an adversary can easily determine which entries in the compromised routers belong to the same circuits and what positions they hold in those circuits. The adversary links together entries in its routers by using the circuit identifiers that are uniquely associated with each circuit. And because circuits have a fixed length compromised routers can determine their position in the circuit by counting the number of messages received after the circuit entry is made.

Theorem 4. *Configurations C and D are indistinguishable only if $C \approx_{D_A} D$.*

Proof. Suppose that C and D are not in the same equivalence class. Let the adversary run the automata prescribed by the protocol on the agents it controls. Let α be a fair, cryptographic execution of C and β be a fair, cryptographic execution of D.

Partition the adversary actions of α into subsequences that share the same circuit identifier. There is at most one such partition for each circuit. Circuit positions that are created in the same partition belong to the same circuit. In each partition the adversary can determine the absolute location of a circuit position filled by a given compromised agent a by counting the total number of messages it sees after the initial CREATE. A can also determine the agents that precede and succeed a on the circuit and the circuit identifier itself. Therefore A can determine the reduced circuit structure $\rho(C)$ from α.

The adversary can use β in the same way to determine $\rho(D)$. By Lemma 4, $\rho(C) \neq \rho(D)$, and so A can always distinguish between C and D. $\qquad\square$

5.3.6 Anonymity

The configurations that provide sender anonymity, receiver anonymity, and relationship anonymity follow easily from Theorems 1, 2, 3, and 4.

Corollary 1. *User u has sender anonymity in configuration C with respect to adversary A if and only if at least one of the following cases is true:*

1. *u and $C_1(u)$ are uncompromised, and there exists another user $v \neq u$ such that v and $C_1(v)$ are uncompromised.*

2. *u and $C_i(u)$ are uncompromised, for all i.*

Proof. Suppose u has sender anonymity in configuration C. Then there must be an indistinguishable configuration D in which u uses a different circuit identifier. By Theorem 4, $C \approx_{D_A} D$, and therefore, by definition, there must be a chain of configurations from C to D such that consecutive pairs are indistinguishable by Theorems 1, 2, and 3. At some point in this chain, the circuit identifier of u must change. If, at this point, the consecutive configurations are indistinguishable by Theorem 1, u must swap circuits with some v, and the conditions for this are exactly those of (1). It is not possible for the consecutive configurations at this point to be indistinguishable by Theorem 2, because it cannot affect the circuit identifier of u. If, at this point, the configurations are indistinguishable by Theorem 3, the conditions of (2) must hold.

Conversely, if (1) holds we can apply Theorem 1 to C to yield an indistinguishable configuration in which u has a different circuit identifier. If (2) holds we can apply Theorem 3 to C to change the circuit identifier of u entirely. $\square$

Corollary 2. *Destination d has receiver anonymity on u's circuit, in configuration C, and with respect to adversary A if and only if at least one of the following cases is true:*

1. *u and $C_l(u)$ are uncompromised, and there exists more than one destination.*

2. *u and $C_i(u)$ are uncompromised, for all i.*

Proof. We use the assumption that all destinations are uncompromised. Suppose d has receiver anonymity in configuration C. Then there must be an indistinguishable configuration D in which d is not the destination on a circuit with the circuit identifier of u in C. By Theorem 4, $C \approx_{D_A} D$, and therefore, by definition, there must be a chain of configurations from C to D such that consecutive pairs are indistinguishable by Theorems 1, 2, and 3. At some point in this chain, d must no longer be the destination on a circuit with the circuit identifier of u in C. Theorem 1 only affects users, and thus cannot apply at this point. If Theorem 2 applies, the conditions of (1) must be satisfied. If Theorem 3 applies, the conditions of (2) must be satisfied.

Conversely, if (1) holds we can apply Theorem 2 to C to yield an indistinguishable configuration in which d is not the destination of a circuit with the circuit identifier of u in C. If (2) holds, we can apply Theorem 3 to C to change the circuit identifier of u entirely. $\square$

Corollary 3. *User u and destination d have relationship anonymity in configuration C with respect to adversary A if and only if at least one of the following cases apply:*

1. *u and $C_l(u)$ are uncompromised, and there exists more than one destination.*

2. u and $C_1(u)$ are uncompromised. There exists another user $v \neq u$ such that v and $C_1(v)$ are un-compromised. $C_{l+1}(v) \neq d$, or $C_l(v)$ is uncompromised and there exists more than one destination.

Proof. Suppose u and d have relationship anonymity in configuration C. By Theorem 4, $C \approx_{D_A} D$. Therefore, by definition, there must be a chain of configurations $C = C^0, C^1, \ldots, C^m = D$ such that each consecutive pair, (C^i, C^{i+1}), is indistinguishable by Theorems 1, 2, and 3. At some point i in the chain, d must cease to be the destination of u. Suppose that, at this point, the configurations are indistinguishable by Theorem 1. Then u and $C_1^i(u)$ are uncompromised, and, for some $v \neq u$, v and $C_1^i(u)$ are uncompromised. Also, the destination of v is not d. Then either $C_{l+1}(v) \neq d$, or v and d have relationship anonymity in C. We inductively assume (performing induction on the length of the chain) that Corollary 3 holds. Then if (1) holds for v, (2) is satisfied for u. If (2) holds for v, then, for some w, v and w can be swapped to give v relationship anonymity. w can therefore play the role of v in the conditions for (2). In the base case, swapping u and v happens at $i = 0$, and it must be the case that $C_{l+1}(v) = d$, satisfying (1).

Conversely, if (1) holds, we can apply Theorem 2 to replace d with another destination, yielding an indistinguishable configuration in which the destination of u is not d. Now suppose (2) holds. If $C_{l+1}(v) \neq d$, then we can apply Theorem 1 to swap u and v in C, yielding an indistinguishable configuration in which the destination of u is not d. Otherwise, we can swap u and v by Theorem 1, and then we are able to apply Theorem 2 to replace d as the destination of u. $\square$

5.3.7 Model Changes

We chose the described protocol to balance two goals. The first was to accurately model the Tor protocol. The second was to make it simple enough to be analyzed, and also so that the main ideas of the analysis weren't unnecessarily complicated. Our results are robust to changes of the protocol, however. We can make the protocol simpler by removing circuit identifiers and multiple encryption without weakening the indistinguishability results. In the other direction, we can make it more complicated with a stream cipher and multiple circuits per user without weakening the distinguishability results.

Circuit identifiers can be simulated by an adversary performing a timing attack, and therefore are just a convenience that helps us reason about the anonymity of the protocol. Consider modifying the protocol so that the messages sent by users do not include circuit identifiers. That is, remove the leading number from all messages received or sent in Automaton 1. Call this protocol "identifier-free onion routing." The

following theorem shows that an adversary can simulate the use of circuit identifiers.

Theorem 5. *There exists an adversary that can calculate, for an execution of identifier-free onion routing, a sequence of messages he could have received had the users used circuit identifiers.*

Proof. The adversary must be able to link together the different parts of the circuit-building process that he can observe to add circuit identifiers consistently to the messages he receives. To do so, the adversary gives a unique number n_i to each router a_i that he controls, where $n_i > 2l$ and $|n_i - n_j| > 2l$ for all i, j. Then, he runs the protocol normally, except that, after sending a CREATED message from his router a_i, he sends n_i dummy messages up the same link. These messages get forwarded up the circuit. Then, if the adversary observes between n_i and $2l + n_i$ messages coming up a circuit at some other compromised router a_j, he knows that a_i and a_j are on the same circuit.

Then the adversary can pick any set of n numbers as circuit identifiers, assign those to circuits, and for each circuit add its identifier to the beginning of those messages that are sent on that circuit. □

Theorem 5 shows that circuit identifiers do not affect anonymity, because sender anonymity, receiver anonymity, and relationship anonymity are invariant under the actual numbers used as identifiers. The protocol automata only test identifiers for equality, and therefore permuting circuit identifiers doesn't affect which executions are indistinguishable. Thus if the adversary can determine a set of circuit identifiers which could have been used in the protocol, the sender, receiver, and relationship anonymity is the same as if those identifiers or any permutation of them had been used.

Note that we could destroy the ability for the adversary to send dummy messages along a circuit by including message authentication codes. This is a standard addition to onion routing. However, by including some notion of time, the adversary simply delay messages, release them at specific times, and observe where messages are later observed. The protocol we give in Chapter 8 is designed to deal with this problem, and it does so using some non-cryptographic ideas.

Multiple encryption does not appear to be necessary for the distinguishability theorems, and therefore the anonymity results. Consider a single-encryption protocol in which the user only encrypts each message with the key of the last router added to the circuit. Messages aren't encrypted or decrypted as they pass up and down a circuit. The adversary still is not able to determine parts of a circuit that aren't adjacent to a compromised agent. The proof of this under multiple encryption did not use the changing representation of messages going along a circuit, and only relied on the last key of the multiple encryption to hide the content

of messages. Single encryption does allow the adversary to easily link entries in his routers by sending messages along the circuit. This power is already available in our model from circuit identifiers, though.

Stream ciphers are used in the Tor protocol and prevent signaling along a circuit using dummy messages. Sending such messages will throw off the counter by some routers on the circuit and the circuit will stop working. We can model a stream cipher by expressing the encryption of the ith message p with key k as $\{p\}_{(k,i)}$, and allowing a different permutation to be applied for every pair (k,i). This can only increase the size of the configuration indistinguishability relation. However, the proof for the distinguishability of configurations only relies on the ability of the adversary to decrypt using his keys, count messages, and recognize the circuit identifier. Therefore it should still hold when the model uses a stream cipher. Also, with a stream cipher the circuit identifier is still not necessary for our results. The adversary can again use the process described above to link entries in compromised routers, since although it involves sending dummy messages, they are sent after the circuit creation is finished and therefore do not interfere with it.

Allowing users to create multiple circuits doesn't weaken the adversary's power to link together its circuit positions and determine their position, but the number of configurations that are consistent with this view does in some cases increase. Let users create an arbitrary number of circuits, where a different key is used between a user and router for every position-and-circuit pair. Then the protocol executes as if each circuit is created by a different user. Therefore the adversary should be able to determine the reduced circuit structure $\rho(C)$ of a configuration C. Moreover, the adversary should not be able to determine any additional router information, which does not depend on the identity of the user. However, the function mapping circuits to users is no longer a permutation, and therefore there should only be more possibilities for the user of a given circuit. This would only improve anonymity.

Chapter 6

Probabilistic Anonymity in Onion Routing

6.1   Introduction

In Chapters 4 and 5, we present a formal I/O-automata model of onion routing and proved anonymity guarantees. In this chapter, we construct an abstraction of that model and treat the network simply as a black box to which users connect and through which they communicate with destinations. The abstraction captures the relevant properties of a protocol execution that the adversary can infer from his observations - namely, the observed users, the observed destinations, and the possible connections between the two. We give a map between the models and show that they posses the same anonymity properties. In this way, we show that one can abstract away from much of the design specific to onion routing so that our results may apply both to onion routing and to other low-latency anonymous-communication designs. The executions described in Chapter 2 disappear from the analysis done herein.

Our previous analysis in the I/O-automata model was possibilistic, a notion of anonymity that is simply not sensitive enough. It makes no distinction between communication that is equally likely to be from any one of a hundred senders and communication that came from one sender with probability .99 and from each of the other 99 senders with probability .000101. An adversary in the real world is likely to have information about which scenarios are more realistic than others. In particular, users' communication patterns are not totally random. When the adversary can determine with high probability, for example, the sender of a

message, that sender is not anonymous in a meaningful way.

Using this intuition, we add a probability distribution to the I/O-automata model of onion routing given in Chapter 4. For any set of actual circuit sources and destinations, there is a larger set that is consistent with the observations made by an adversary. The adversary can then infer conditional probabilities on this larger set using the distribution. This gives the adversary probabilistic information about the facts we want the network to hide, such as the initiator of a communication.

The probability distribution that we use models heterogeneous user behavior. In our onion-routing protocol, each user chooses a circuit to a destination. We make this choice probabilistic and have each user choose a destination according to some probability distribution, allowing this distribution to be different for different users. We assume that the users choose their circuits by selecting the routers on it independently and at random.

After observing the protocol, the adversary can in principle infer some distribution on circuit source and destination. He may not actually know the underlying probability distribution, however. In particular, it doesn't seem likely that the adversary would know how every user selects destinations. In our analysis, we take a worst-case view and assume that the adversary knows the distribution exactly.

This assumption may not be too pessimistic, because over time the adversary might learn a good approximation of user behavior via the long-term intersection attack [18]. In an intersection attack, one watches repeated communication events for patterns of senders and receivers over time. Unless all senders are on and sending all the time (in a way not selectively blockable by an adversary) and/or all receivers receiving all the time, if different senders have different receiving partners, there will be patterns that arise and eventually differentiate the communication partners. It has long been recognized that no system design is secure against a longterm intersection attack.

If the adversary does know the distribution over destinations for every user, it may seem as though anonymity has been essentially lost anyway. However, even when the adversary knows how a user generally behaves, the anonymity network may make it hard for him to determine who is responsible for any specific action, and the anonymity of a specific action is what we are interested in.

We analyze relationship anonymity in our probabilistic onion-routing model. Our definition of relationship anonymity in Chapter 5 simply requires that, as far as the adversary can tell, there is more than one possibility for a given user's destination. The probability distribution on executions we add in this chapter yields a conditional distribution on these possible destinations. Therefore we can use a quantitative metric

for anonymity, as discussed in Chapter 2. We will use the probability assigned to the correct destination as our metric. In part, this is because it is the simplest metric. Also, any statements about entropy and maximum probability metrics only make loose guarantees about the probability assigned to the actual subject, a quantity that clearly seems important to the individual users.

We look at the value of this anonymity metric for a choice of destination by a user. Fixing a destination by just one user, say u, does not determine what the adversary sees, however. The adversary's observations are also affected by the destinations chosen by the other users and the circuits chosen by everybody. Because those variables are chosen probabilistically under the distribution we added, the anonymity metric will have its own distribution. Several statistics about this distribution might be interesting; in this chapter, we look at its expectation.

The distribution of the anonymity metric for a given user and destination depends on the other users' destination distributions. If their distributions are very different, the adversary may have an easy time separating out the actions of the user. If they are similar, the user may more effectively hide in the crowd.

We begin by providing a kind of worst-case guarantee to a user with a given destination distribution by finding the maximum expectation over the possible destination distributions of the other users. Our results show that the worst case is when every other user either always visits the destinations the user is otherwise least likely to visit or always visits his actual destination. Which one is worse depends on how likely he was to visit his destination in the first place. If he is unlikely to visit it, it is worse when everybody else always visits his otherwise least-likely destination, because the adversary can generally infer that he is not responsible for communication to that destination. When he is likely to visit it, the adversary considers him likely to be the culprit whenever the destination is observed, and so observing that destination often causes the adversary to suspect the truth. We give an approximation to the user's anonymity in these worst cases for large user populations that shows that on average it decreases by about the fraction of the network the adversary controls.

We then consider anonymity in a more typical set of user distributions. In the model suggested by Shmatikov and Wang [76], each user selects a destination from a common Zipfian distribution. Because the users are identical, every user hides well among the others. As the user population grows, the anonymity loss in this case tends to the square of the fraction of the network that is compromised.

We expect this to have potential practical applications. For example, designs for shared security-alert repositories to facilitate both forensic analysis for improved security design and quicker responses to widescale

attacks have been proposed [52]. A participant in a shared security-alert repository might expect to be known to communicate with it on a regular basis. Assuming reports of intrusions, etc., are adequately sanitized, the concern of the participant should be to hide when it is that updates from that participant arrive at the repository, that is, which updates are likely to be from that participant as opposed to others.

6.2 Technical Preliminaries

6.2.1 Model

We describe our analysis of onion routing in terms of a black-box model of anonymous communication. We are using a black-box model for two reasons: First, it abstracts away the nonessential details, and second, its generality immediately suggests ways to perform similar analyses of other anonymity networks. It models a round of anonymous communication as a set of inputs owned by users and a set of outputs owned by destinations. The adversary observes the source of some of the inputs and the destination of some of the outputs. This captures the basic capabilities of an adversary in an onion-routing network that controls some of the routers. In this situation, the adversary can determine the source of messages when he controls the first router on the source's circuit and the destination of messages when he controls the last router. In order for the adversary always to be able to recognize when it controls the onion router adjacent to the circuit source, we assume that the initiating client is not located at an onion-routing network node. This is the case for the vast majority of circuits in Tor and in all significant deployments of onion routing and similar systems to date. The black box system can also model a mix network under attack by a global, passive adversary. Such a model was used by Kesdogan et al. [48] in their analysis of an intersection attack.

We add two assumptions to specialize this model to onion routing. First, we assume that every user owns exactly one input and is responsible for exactly one output in a round. Certainly users can communicate with multiple destinations simultaneously in actual onion-routing systems. However, it seems likely that in practice most users have at most some small constant number of active connections at any time, and the smaller this constant is the fewer possibilities there are that are consistent with the adversary's observations. Therefore, this assumption is a conservative one that gives the adversary as much power to break anonymity as the limited number of user circuits can provide. Second, we assume the adversary can link together an input and output from the same user when he observes them both. This is another conservative assumption that is motivated by the ability of an active adversary to use the timing and volume of messages on a circuit

to recognize that the observations he makes at different locations on the circuit are part of the same circuit. These timing and volume patterns can be created by the user or they can be artifically created by the adversary. An example of the latter strategy is described in the proof of Theorem 5, and the purpose is indeed to allow the adversary to link together different parts of the same circuit.

Let U be the set of users with $|U| = n$. Let Δ be the set of destinations. A round of communication is defined by the *black-box configuration* that includes the destinations chosen by the users and the observations made by the adversary.

Definition 15. *A* black-box configuration C *in a black-box system consists of:*

1. $C_D : U \rightarrow \Delta$, *which indicates the selection of a destination by each user.*

2. $C_I : U \rightarrow \{0, 1\}$, *which indicates the set of users whose inputs are observed.*

3. $C_O : U \rightarrow \{0, 1\}$, *which indicates the set of users whose outputs are observed.*

For the remainder of the chapter, the term *configuration* will refer to a black-box configuration. A user's input, output, and destination will be called its *circuit*.

We include the probabilistic behavior of users by adding a probability distribution over configurations. Let each user u select a destination d from a distribution p^u over Δ, where we denote the probability that u chooses d as p_d^u. Every input and output is independently observed with probability b. This reflects the probability that the first or last router of a user's circuit is compromised when the user selects the circuit's routers independently and at random, and the adversary controls a fraction b of the routers. The probability of a configuration C is the joint probability of its events:

$$Pr[C] = \prod_{u \in U} \left(p_{C_D(u)}^u \right) \left(b^{C_I(u)} (1-b)^{1-C_I(u)} \right) \left(b^{C_O(u)} (1-b)^{1-C_O(u)} \right). \tag{6.1}$$

For any configuration, there is a larger set of configurations that are consistent with the inputs and outputs that the adversary sees. We will call two configurations *indistinguishable* if the sets of inputs, outputs, and links between them that the adversary observes are the same.

Definition 16. *Configurations C and $\overline{C}$ are* indistinguishable *if there exists a permutation $\pi : U \rightarrow U$ such that for all $u \in U$:*

1. $C_I(u) = 1 \wedge C_O(u) = 1 \Rightarrow \overline{C}_I(u) = 1 \wedge \overline{C}_O(u) = 1 \wedge C_D(u) = \overline{C}_D(u)$

63

2. $C_I(u) = 1 \wedge C_O(u) = 0 \Rightarrow \overline{C}_I(u) = 1 \wedge \overline{C}_O(u) = 0$

3. $C_I(u) = 0 \wedge C_O(u) = 1 \Rightarrow \overline{C}_I(\pi(u)) = 0 \wedge \overline{C}_O(\pi(u)) = 1 \wedge C_D(u) = \overline{C}_D(\pi(u))$

4. $C_I(u) = 0 \wedge C_O(u) = 0 \Rightarrow \overline{C}_I(\pi(u)) = 0 \wedge \overline{C}_O(\pi(u)) = 0$

Thus, two configurations are indistinguishable if they have the same pattern of observed inputs, outputs, and destinations, while allowing the identities of users with unobserved inputs to be permuted. The indistinguishability relation is an equivalence relation, and, in particular, is symmetric, because, if C and $\overline{C}$ are indistinguishable under π, then $\overline{C}$ and C are indistinguishable under π^{-1}. Therefore we use the notation $C \approx \overline{C}$ to indicate that configurations C and $\overline{C}$ are indistinguishable.

6.2.2 Relationship Anonymity

We analyze the relationship anonymity of users and destinations in our model. As described in Chapter 2, relationship anonymity refers to the adversary's ability to determine if a user and destination communicate with each other.

A user has relationship anonymity in a possibilistic sense if there is an indistinguishable configuration in which the user does not communicate with his destination. For example, under this definition a user with observed output but unobserved input sends that output anonymously if there exists another user with unobserved input and a different destination.

The probability distribution we have added to configurations allows us to use a probabilistic notion of anonymity and incorporate the degree of certainty that the adversary has about communication between users and destinations. After making observations in the actual configuration, the adversary can infer a conditional probability distribution on configurations. We measure the relationship anonymity of user u and destination d by the posterior probability that u chooses d as his destination. The lower this is, the more anonymous we consider their relationship.

The relationship anonymity of u and d varies with the destination choices of the other users and the observations of the adversary. If, for example, u's output is observed, and the inputs of all other users are observed, then the adversary assigns u's destination a probability of 1. Because we want to examine the relationship anonymity of u conditioned only on his destination, we end up with a distribution on the anonymity metric. We look at the expectation of this distribution.

This expectation depends on the destination distributions of all of the users. If the other users behave

similarly to u, for example, it might be hard to determine which observed destination belongs to u, while if they rarely choose the same destinations as u, it might be easy to figure out u's destination. We consider how this expectation varies with the other users' destination distributions. In particular, we examine the maximum expectation for a given user u over the p^v, $v \neq u$, as this provides a lower bound to the user on his anonymity. We continue by examining the expectation in a more likely situation.

6.2.3 Model Equivalence

What makes the black-box model abstraction useful is that relationship anonymity is preserved under a natural map from the protocol configurations of Definition 8 to black-box configurations. Let $\mathcal{C}_p$ be the set of protocol configurations and $\mathcal{C}_b$ be the set of black-box configurations. Let $\phi : \mathcal{C}_p \rightarrow \mathcal{C}_b$ be the function such that $\phi(C^p) = (C_D^b, C_I^b, C_O^b)$, where for all $u \in U$

$$
\begin{aligned}
C_D^b(u) &= C_{l+1}^p(u) \\
C_I^b(u) &= \begin{cases} 1 & \text{if } u \in A \vee C_1^p(u) \in A \\ 0 & \text{otherwise} \end{cases} \\
C_O^b(u) &= \begin{cases} 1 & \text{if } C_l^p(u) \in A \\ 0 & \text{otherwise} \end{cases}
\end{aligned}
$$

The black-box model essentially discards all information in a protocol configuration except the user destinations, whether or not the adversary observes a user, and whether or not the adversary observes a destination. By Corollary 3, these properties of a configuration determine relationship anonymity of the configuration. Indeed, user u has possibilistic relationship anonymity in protocol configuration C_1^p if and only if u has possibilistic relationship anonymity in black-box configuration $\phi(C_1^p)$. In fact, with an appropriate probability distribution over executions of the protocol configurations, ϕ preserves probabilistic anonymity.

Let $C^p(\alpha)$ be the configuration of execution α. Extend the notion of indistinguishability to executions such that, for executions α and β, $\alpha \approx_{D_A} \beta$ if α and β satisfy the four conditions of Definition 9. Let X be a random execution, and Y be a random black-box configuration.

Theorem 6. *There exists a distribution μ on fair cryptographic executions such that, for any fair crypto-*

graphic execution α and black-box configuration C^b,

$$Pr_\mu[\phi(C^p(X)) = C^b | X \approx_{D_A} \alpha] = Pr[Y = C^b | Y \approx \phi(C^p(\alpha))]. \tag{6.2}$$

Proof. We describe the distribution on fair cryptographic executions μ via distributions conditional on protocol configurations. Let C be a protocol configuration. At the end of a finite cryptographic execution α of C there is a finite set E of enabled actions. Let the probability under μ that an enabled action is the next action in all fair cryptographic executions that contain α as a prefix be $1/|E|$. That is, at any point in a fair cryptographic execution, every enabled action is equally likely under μ. To make μ a full distribution on configurations, we let each user u choose routers uniformly at random and a destination according to the distributions p^u.

We first observe that a bijection exists between the fair cryptographic executions of two indistinguishable protocol configurations, and that bijection is between indistinguishable executions. Let C and D be protocol configurations, $C \sim_{D_A} D$. The proofs of Theorems 1, 2, and 3 transform executions from one configuration into the other in a way that is invertible and pairs indistinguishable executions. By the Bernstein-Schroeder theorem, then, there exists a bijection between the fair cryptographic executions of C and D the continues to pair indistinguishable executions. Now let C and D be any pair of indistinguishable protocol configurations, that is, $C \approx_{D_A} D$. By Theorem 4, they are related by a chain of executions $C = C_0 \sim_{D_A} C_1 \sim_{D_A} \ldots \sim_{D_A} C_k = D$. Then there must also be a bijection between the executions of C and the executions of D, and it pairs indistinguishable executions. Call this bijection ψ_{CD}

ψ preserves the conditional distribution on executions between indistinguishable protocol configurations. Let C and D be protocol configurations such that $C \approx_{D_A} D$. Given a finite execution α of C, let $\mathcal{E}_\alpha$ be the set of fair cryptographic executions of C with α as a prefix, and let $\psi_{CD}(\mathcal{E}_\alpha)$ be their image under ψ_{CD}. If C and D are related by Theorem 1, then the number of enabled actions at any point in α is the same as the set at the same point in $\psi_{CD}(\alpha)$ (by Theorem 1 they are identical after swapping two users). If C and D are related by Theorem 2, the number of enabled actions at every point in α are the same for all users and the same for all routers except the pair r_i and r_i' given in the statement of Theorem 2. For that pair, the actions related to the circuit that is different between C and D are enabled in $\psi(\alpha)$ for r_i' instead of r_i. This maintains the same number of enabled actions at every point in α and $\psi(\alpha)$. If C and D are related by Theorem 3, then every agent has the same number of enabled actions at every point in α and $\psi(\alpha)$. If C and

D are indistinguishable by a chain of configurations related by Theorems 1, 2, and 3, then we have shown that this chain preserves the number of enabled actions at every step in the chain and therefore between α and $\psi(\alpha)$. Because the probability of α is defined by the number of enabled actions at every step, then, for random executions X of C and Y of D, $Pr[X \in \mathcal{E}_\alpha | C] = Pr[Y \in \psi_{CD}(\mathcal{E}_\alpha)|D]$.

Let Z be a random protocol configuration. Now we can say that, given an execution α of C, $Pr[Z = z | X \approx_{D_A} \alpha] = Pr[Z = z | Z \approx_{D_A} C]$. Thus executions have disappeared from the picture, and we just need to relate the probability of protocol configurations that map to a black-box configuration under ϕ to the distribution on black-box configurations already defined.

Let $\mathcal{C}$ be an equivalence class of protocol configurations under $\approx_{D_A}$. Given, $C_1, C_2 \in \mathcal{C}$, let $b_i = |\{C \in \mathcal{C} : \phi(C) = \phi(C_i)\}|, i \in \{1, 2\}$. We show that the number of indistinguishable protocol configurations that map to each black-box configuration is the same, that is, that $b_1 = b_2$. Because $C_1 \approx_{D_A} C_2$, $\rho(C_1) = \rho(C_2)$, where ρ, as given in Definition 14, is the function that reduces circuits to the components adjacent to the adversary. Add to each reduced circuit in $\rho(C_i)$ the user and destination of the full circuit in C_i, $i \in \{1, 2\}$, and call this ρ_i'. There potentially exist protocol configurations consistent with ρ_1' and ρ_2' that map to different black-box configurations; in particular, it may be that $\phi(C_1) \neq \phi(C_2)$. However, further assignments of routers and circuit identifiers to ρ_i' do not change the black-box functions C_D, C_I, and C_O, and therefore $\rho(C') = \rho(C_i)$ for all C' consistent with ρ_i'. The number of such assignments only depends on the open circuit positions and not on the identities of the users and destinations. Thus the number of ways to make these assignments is the same for both ρ_1' and ρ_2'. Moreover, any $C' \in \mathcal{C}$ that is not consistent with ρ_1' or ρ_2' does not map to $\phi(C_1)$ or $\phi(C_2)$ because C' must assign some user to a different destination. The number of possible assignments consistent with ρ_i' is b_i, and so $b_1 = b_2$.

The probability of a protocol configuration C is $(1/m)^{ln} \prod_{u \in U} p_{C_{l+1}}^u$. Therefore, because the number of protocol configurations that map to a black-box configuration C^b under ϕ is the same for all C^b in the same indistinguishable equivalence class, the conditional probability of C^b depends only on the assignment of users to destinations. More specifically, for an execution α of C, X a random execution, and a black-box configuration $C^b \approx \phi(C)$,

$$
\begin{aligned}
Pr_\mu[\phi(C^p(X)) = C^b | X \approx_{D_A} \alpha] &= \prod_{u \in U} p_{C_D^b(u)}^u \quad &(6.3)\\
&= Pr[Y = C^b | Y \approx \phi(C^p(\alpha))]. \quad &(6.4)
\end{aligned}
$$

Equation 6.4 follows from Equation 6.1, and it gives us the theorem. □

6.3 Expected Anonymity

Let the set $\mathcal{C}$ of all configurations be the sample space and X be a random configuration. X is then distributed according to Equation 6.1. Let Y be the posterior probability of the event that u chooses d as a destination, that is, $Y(C) = Pr[X_D(u) = d | X \approx C]$. Y is our metric for the relationship anonymity of u and d.

6.3.1 Calculation and Bounds

Let $\mathbb{N}^\Delta$ represent the set of multisets over Δ. Let $\rho(\Delta^0)$ be the number of permutations of $\Delta^0 \in \mathbb{N}^\Delta$ that only permute elements of the same type:

$$\rho(\Delta^0) = \prod_{\delta \in \Delta} |\{\delta \in \Delta^0\}|!.$$

Let $\Pi(A, B)$ be the set of all injective maps $A \to B$. The following theorem gives an exact expression for the conditional expectation of Y in terms of the underlying parameters U, Δ, p, and b:

Theorem 7.

$$
\begin{aligned}
E[Y|X_D(u) = d] = \;& b(1-b)p_d^u + b^2 \\
& + \sum_{S \subseteq U : u \in S} \sum_{\Delta^0 \in \mathbb{N}^\Delta : |\Delta^0| \leq S} b^{n - |S| + |\Delta^0|}(1-b)^{2|S| - |\Delta^0|} \\
& \left(\begin{aligned} & \sum_{T \subseteq S - u : |T| = |\Delta^0| - 1} \sum_{\pi \in \Pi(T+u, \Delta^0) : \pi(u) = d} p_d^u \prod_{v \in T} p_{\pi(v)}^v \\ & + \sum_{T \subseteq S - u : |T| = |\Delta^0|} \sum_{\pi \in \Pi(T, \Delta^0)} p_d^u \prod_{v \in T} p_{\pi(v)}^v \end{aligned} \right)^2 \\
& [\rho(\Delta^0)]^{-1} (p_d^u)^{-1} \left(\sum_{T \subseteq S : |T| = |\Delta^0|} \sum_{\pi \in \Pi(T, \Delta^0)} \prod_{v \in T} p_{\pi(v)}^v \right)^{-1}
\end{aligned}
$$

(6.5)

Proof. At a high level, the conditional expectation of Y can be expressed as

$$\mathrm{E}[Y|X_D(u) = d] = \sum_{C \in \mathcal{C}} \Pr[X = C | X_D(u) = d] Y(C)$$

We calculate Y for a configuration C by finding the relative weight of indistinguishable configurations in which u selects d. The adversary observes some subset of the circuits. If we match the users to circuits in some way that sends users with observed inputs to their own circuits, the result is an indistinguishable

configuration. Similarly, we can match circuits to destinations in any way that sends circuits on which the output has been observed to their actual destinations in C.

The value of $Y(C)$ is especially simple if u's input has been observed. If the output has not also been observed, then $Y(C) = p_d^u$. If the output has also been observed, then $Y(C) = 1$.

For the case in which u's input has not been observed, we have to take into account the destinations of and observations on the other users. Let $S \subseteq U$ be the set of users s such that $C_I(s) = 0$. Note that $u \in S$. Let Δ^0 be the multiset of the destinations of circuits in C on which the input has not been observed, but the output has.

Let $f_0(S, \Delta^0)$ be the probability that in a random configuration the set of unobserved inputs is S and the set of observed destinations with no corresponding observed input is Δ^0:

$$f_0(S, \Delta^0) = b^{n-|S|+|\Delta^0|}(1-b)^{2|S|-|\Delta^0|}[\rho(\Delta^0)]^{-1} \sum_{T \subseteq S: |T|=|\Delta^0|} \sum_{\pi \in \Pi(T, \Delta^0)} \prod_{v \in T} p_{\pi(v)}^v.$$

Let $f_1(S, \Delta^0)$ be the probability that in a random configuration the set of unobserved inputs is S, the set of observed destinations with no corresponding observed input is Δ^0, the output of u is observed, and the destination of u is d:

$$f_1(S, \Delta^0) = b^{n-|S|+|\Delta^0|}(1-b)^{2|S|-|\Delta^0|}[\rho(\Delta^0)]^{-1} p_d^u \sum_{T \subseteq S-u: |T|=|\Delta^0|-1} \sum_{\pi \in \Pi(T+u, \Delta^0): \pi(u)=d} \prod_{v \in T} p_{\pi(v)}^v.$$

Let $f_2(S, \Delta^0)$ be the probability that in a random configuration the set of unobserved inputs is S, the set of observed destinations with no corresponding observed input is Δ^0, the output of u is unobserved, and the destination of u is d:

$$f_2(S, \Delta^0) = b^{n-|S|+|\Delta^0|}(1-b)^{2|S|-|\Delta^0|}[\rho(\Delta^0)]^{-1} p_d^u \sum_{T \subseteq S-u: |T|=|\Delta^0|} \sum_{\pi \in \Pi(T, \Delta^0)} \prod_{v \in T} p_{\pi(v)}^v.$$

Now we can express the posterior probability $Y(C)$ as

$$Y(C) = \frac{f_1(S, \Delta^0) + f_2(S, \Delta^0)}{f_0(S, \Delta^0)}. \tag{6.6}$$

The expectation of Y is a sum of the above posterior probabilities weighted by their probability. The probability that the input of u has been observed but the output hasn't is $b(1-b)$. The probability that

both the input and output of u have been observed is b^2. These cases are represented by the first two terms in Equation 6.5.

When the input of u has not been observed, we have an expression of the posterior in terms of sets S and Δ^0. The numerator $(f_1 + f_2)$ of Equation 6.6 itself sums the weight of every configuration that is consistent with S, Δ^0, and the fact that the destination of u is d. However, we must divide by p_d^u, because we condition on the event $\{X_D(u) = d\}$.

These observations give us the final summation in Equation 6.5. $\qquad\qquad\square$

The expression for the conditional expectation of Y in Equation 6.5 is not easy to interpret. It would be nice if we could find a simple approximation. The probabilistic analysis in [80] proposes just such a simplification by reducing it to only two cases: $i)$ the adversary observes the user's input and output and therefore identifies his destination and $ii)$ the adversary doesn't observe these and cannot improve his a $priori$ knowledge. The corresponding simplified expression for the expection is

$$E[Y|X_D(u) = d] \approx b^2 + (1 - b^2)p_d^u. \tag{6.7}$$

This is a reasonable approximation if the final summation in Equation 6.5 is about $(1-b)p_d^u$. This summation counts the case in which u's input is not observed, and to achieve a good approximation the adversary must experience no significant advantage or disadvantage from comparing the users with unobserved inputs (S) with the discovered destinations (Δ^0).

The quantity $(1 - b)p_d^u$ does provide a lower bound on the final summation. It may seem obvious that considering the destinations in Δ^0 can only improve the accuracy of adversary's prior guess about u's destination. However, in some situations the posterior probability for the correct destination may actually be smaller than the prior probability. This may happen, for example, when some user v, $v \neq u$, communicates with a destination e, $e \neq d$, and only u is a $priori$ likely to communicate with e. If the adversary observes the communication to e, it may infer that it is likely that u was responsible and therefore didn't choose d.

It is true, however, that in expectation this probability can only increase. Therefore Equation 6.7 provides a lower bound on the expected anonymity.

The proof of this fact relies on the following lemma. Let $\mathcal{E}$ be an event in some finite sample space Ω. Let $\mathcal{A}_1, \ldots, \mathcal{A}_n$ be a set of disjoint events such that $\mathcal{E} \subseteq \bigcup_i \mathcal{A}_i$, and let $\mathcal{A}^j = \bigcup_{i=1}^j \mathcal{A}_i$. Let $\mathcal{E}_i = \mathcal{E} \cap \mathcal{A}_i$. Finally, let $Z(\omega) = \sum_i 1_{\mathcal{E}_i} Pr[\mathcal{E}_i]/Pr[\mathcal{A}_i]$ (where $1_{\mathcal{E}_i}$ is the characteristic function for $\mathcal{E}_i$). $Z(\omega)$ is thus the

conditional probability $Pr[\mathcal{E}|\mathcal{A}_i]$, where $\omega \in \mathcal{E}_i$.

Lemma 5. $Pr[\mathcal{E}|\mathcal{A}^n] \leq E[Z|\mathcal{E}]$

Proof.

$$
\begin{aligned}
Pr[\mathcal{E}|\mathcal{A}^n] \quad &= \frac{Pr[\mathcal{E}]}{Pr[\mathcal{A}^n]} \\
&= \frac{\left(\sum_i \frac{Pr[\mathcal{E}_i]\sqrt{Pr[\mathcal{A}_i]}}{\sqrt{Pr[\mathcal{A}_i]}}\right)^2}{Pr[\mathcal{A}^n]Pr[\mathcal{E}]} \qquad \text{by a simple rewriting} \\
&\leq \frac{\left(\sqrt{\sum_i \frac{(Pr[\mathcal{E}_i])^2}{Pr[\mathcal{A}_i]}}\sqrt{\sum_i Pr[\mathcal{A}_i]}\right)^2}{Pr[\mathcal{A}^n]Pr[\mathcal{E}]} \qquad \text{by the Cauchy-Schwartz inequality} \\
&= \sum_i \frac{(Pr[\mathcal{E}_i])^2}{Pr[\mathcal{A}_i]Pr[\mathcal{E}]} \\
&= E[Z|\mathcal{E}]
\end{aligned}
$$

$\square$

Theorem 8. $E[Y|X_D(u) = d] \geq b^2 + (1 - b^2)p_d^u$

Proof. As described in the proof of Theorem 7

$$
E[Y|X_D(u) = d] = b^2 + b(1 - b)p_d^u + (1 - b)E[Y|X_D(u) = d \wedge X_I(u) = 0]
$$

To apply Lemma 5, take the set of configurations $\mathcal{C}$ to be the sample space Ω. Take $\{X_D(u) = d\}$ to be the event $\mathcal{E}$. Take the indistinguishability equivalence relation to be the sets $\mathcal{A}_i$. Finally, take Y to be Z. Then the lemma shows that $E[Y|X_D(u) = d \wedge X_I(u) = 0] \geq p_d^u$. $\square$

To examine the accuracy of our approximation, we look at how large the final summation in Equation 6.5 can get as the users' destination distributions vary. Because this is the only term that varies with the other user distributions, this will also provide a worst-case guarantee on expected anonymity. Our results will show that the worst case can occur when the users other than u act as differently from u as possible by always visiting the destination u is otherwise least likely to visit. Less obviously, we show that the maximum can also occur when the users other than u always visit d. This happens because it makes the adversary observe destination d often, causing him to suspect that u chose d. Our results also show that the worst-case expectation is about $b + (1 - b)p_d^u$, which is significantly worse than the simple approximation above.

As the first step in finding the maximum of Equation 6.5 over $(p^v)_{v \neq u}$, we observe that it is obtained when every user $v \neq u$ chooses only one destination d_v, that is, $p_{d_v}^v = 1$ for some $d_v \in \Delta$.

Lemma 6. *A maximum of $E[Y|X_D(u) = d]$ over $(p^v)_{v \neq u}$ must occur when, for all $v \neq u$, there exists some $d_v \in \Delta$ such that $p_{d_v}^v = 1$.*

Proof. Take some user $v \neq u$ and two destinations $e, f \in \Delta$. Assign arbitrary probabilities in p^v to all destinations except for f, and let $\zeta = 1 - \sum_{\delta \neq e,f} p_\delta^v$. Then $p_f^v = \zeta - p_e^v$. Consider $E[Y|X_D(u) = d]$ as a function of p_e^v. The terms t_i of $E[Y|X_D(u) = d]$ that correspond to any fixed S and Δ^0 are of the following general form, where $\alpha_i, \beta_i, \gamma_i, \delta_i, \epsilon_i, \eta_i \geq 0$:

$$t_i = \frac{(\alpha_i p_e^v + \beta_i(\zeta - p_e^v) + \gamma_i)^2}{\delta_i p_e^v + \epsilon_i(\zeta - p_e^v) + \eta_i}.$$

This is a convex function of p_e^v:

$$t_i'' = \frac{2(\gamma_i(\delta_i - \epsilon_i) + \beta_i(\delta_i\zeta + \eta_i) - \alpha_i(\epsilon_i\zeta + \eta_i))^2}{(\epsilon_i(\zeta - p_e^v) + \delta_i p_e^v + \eta_i)^3} \geq 0$$

The leading two terms of $E[Y|X_D(u) = d]$ are constant in p^v, and the sum of convex functions is a convex function, and so $E[Y|X_D(u) = d]$ is convex in p_e^v. Therefore, a maximum of $E[Y|X_D(u) = d]$ must occur when $p_e^v \in \{0, 1\}$. $\qquad \square$

The following lemma shows that we can further restrict ourselves to distribution vectors in which, for every user except u, the user either always chooses d or always chooses the destination that u is otherwise least likely to visit.

Lemma 7. *Order the destinations $d = d_1, \ldots, d_{|\Delta|}$ such that $p_{d_i}^u \geq p_{d_{i+1}}^u$ for $i > 1$. Then a maximum of $E[Y|X_D(u) = d]$ must occur when, for all users v, either $p_{d_1}^v = 1$ or $p_{d_{|\Delta|}}^v = 1$.*

Proof. Assume, following Lemma 6, that $(p^v)_{v \neq u}$ is an extreme point of the set of possible distribution vectors.

Equation 6.5 groups configurations first by the set S with unobserved inputs and second by the observed destinations Δ^0. Instead, group configurations first by S and second by the set $T \subseteq S$ with observed outputs. Because every user except u chooses a destination deterministically, Y only depends on the sets S and T. Therefore we can use the notation $Y(S, T)$ without ambiguity.

$$
\begin{aligned}
E[Y|X_D(u) = d] = \quad & b(1-b)p_d^u + b^2 + \\
& \sum_{S:u \in S} \sum_{T:T \subseteq S} b^{n-|S|+|T|}(1-b)^{2|S|-|T|} Y(S, T)
\end{aligned}
\tag{6.8}
$$

72

Select two destinations $d_i, d_j, 1 < i < j$. We break up the sum in Equation 6.8 and show that, for every piece, the sum can only be increased by changing $(p^v)_v$ so that any user that always chooses d_i always chooses d_j instead.

Fix $S \subseteq U, u \in S$. Let $S_i, S_j \subseteq S$ be such that $p^s_{d_i} = 1$ if and only if $s \in S_i$, and $p^s_{d_j} = 1$ if and only if $s \in S_j$. Fix $T' \subseteq S \backslash S_i \backslash S_j$ and some $l \geq |T'|$. Let s_{d_k} be the number of users in $S - u$ that always visit d_k and t_{d_k} be the number of users in $T' - u$ that always visit d_k. Let $f(S, T')$ be the sum of terms in Equation 6.8 that are indexed by S and some T such that $|T| = l$ and $T \supseteq T'$. Let $m = l - |T'|$.

$$
\begin{aligned}
f(S, T') &= b^{n-|S|+l}(1-b)^{2|S|-l} \sum_{k=0}^{m} \binom{s_{d_i}}{k}\binom{s_{d_j}}{m-k} \prod_{e \in \Delta \backslash \{d, d_i, d_j\}} \binom{s_e}{t_e} \\
&\quad p^u_d \left(\binom{s_d}{t_d - 1}\binom{s_{d_i}}{k}\binom{s_{d_j}}{m-k} + \binom{s_d}{t_d}\binom{s_{d_i}}{k}\binom{s_{d_j}}{m-k} \right) \\
&\quad \left(\begin{array}{l} \prod_{e \in \Delta \backslash \{d_i, d_j\}} \binom{s_e}{t_e}\binom{s_{d_i}}{k-1}\binom{s_{d_j}}{m-k}p^u_{d_i} + \prod_{e \in \Delta \backslash \{d_i, d_j\}} \binom{s_e}{t_e}\binom{s_{d_i}}{k}\binom{s_{d_j}}{m-k-1}p^u_{d_j} \\ + \sum_{f \in \Delta \backslash \{d_i, d_j\}} \prod_{e \in \Delta \backslash \{f, d_i, d_j\}} \binom{s_e}{t_e}\binom{s_f}{t_f-1}\binom{s_{d_i}}{k}\binom{s_{d_j}}{m-k}p^u_f + \prod_{e \in \Delta \backslash \{d_i, d_j\}} \binom{s_e}{t_e}\binom{s_{d_i}}{k}\binom{s_{d_j}}{m-k} \end{array} \right)^{-1} \\
&= b^{n-|S|+l}(1-b)^{2|S|-l} \prod_{e \in \Delta \backslash \{d_i, d_j\}} \binom{s_e}{t_e} p^u_d \left(\frac{t_d}{s_d + 1 - t_d} + 1 \right) \\
&\quad \sum_{k=0}^{m} \binom{s_{d_i}}{k}\binom{s_{d_j}}{m-k} \frac{(s_{d_i}+1-k)(s_{d_j}+1-m+k)}{\left(\begin{array}{l} p^u_{d_i}k(s_{d_j}+1-m+k) + p^u_{d_j}(m-k)(s_{d_i}+1-k) \\ + (s_{d_i}+1-k)(s_{d_j}+1-m+k)\left(\sum_{f \in \Delta \backslash \{d_i, d_j\}} p^u_f \frac{t_f}{s_f+1-t_f} + 1 \right) \end{array} \right)} \\
&= \alpha \sum_{k=0}^{m} \binom{s_{d_i}}{k}\binom{s_{d_j}}{m-k} \frac{(s_{d_i}+1-k)(s_{d_j}+1-m+k)}{\left(\begin{array}{l} p^u_{d_i}(s_{d_i}+1)(s_{d_j}+1-m+k) + p^u_{d_j}(s_{d_j}+1)(s_{d_i}+1-k) \\ + (s_{d_i}+1-k)(s_{d_j}+1-m+k)\beta \end{array} \right)} \\
\alpha, \beta &\geq 0
\end{aligned}
$$

This can be viewed as the weighted convolution of binomial coefficients. Unfortunately, there is no obvious way to simplify the expression any further to find the maximum as we trade off s_{d_i} and s_{d_j}. There is a closed-form sum if the coefficient of the binomial product is a fixed-degree polynomial, however. Looking

at the coefficient, we can see that it is concave.

$$c_k = \frac{(s_{d_i}+1-k)(s_{d_j}+1-m+k)}{p^u_{d_i}(s_{d_i}+1)(s_{d_j}+1-m+k)+p^u_{d_j}(s_{d_j}+1)(s_{d_i}+1-k)+(s_{d_i}+1-k)(s_{d_j}+1-m+k)\beta}$$

$$\frac{d^2 c_k}{dk^2} = -\frac{\left(\begin{array}{c} 2((s_{d_i}+1)(s_{d_j}+1)(2+s_{d_i}+s_{d_j}-m)^2 p^u_{d_i} p^u_{d_j}+ \\[4pt] b((s_{d_i}+1)(s_{d_j}+1+k-m)^3 p^u_{d_i}+(s_{d_j}+1)(s_{d_i}+1-k)^3 p^u_{d_j})) \end{array}\right)}{((s_{d_j}+1+k-m)(b(s_{d_i}+1-k)+p^u_{d_i}(s_{d_i}+1))+(s_{d_j}+1)(s_{d_i}+1-k)p^u_{d_j})^3} \le 0$$

We can use this fact to bound the sum above by replacing c_k with a line tangent at some point k_0. Call this approximation $\tilde{f}$. Holding $s_{d_i}+s_{d_j}$ constant, this approximation is in fact equal at $s_{d_i}=0$ because the sum has only one term. Then if $s_{d_i}=0$ still maximizes the sum, the theorem is proved.

$$\begin{aligned} f(S,T') &\le \sum_{k=0}^{m}\binom{s_{d_i}}{k}\binom{s_{d_j}}{m-k}(c'_{k_0}(k-k_0)+c_{k_0}) \\[6pt] &= \binom{s_{d_i}+s_{d_j}}{m}\left(c_{k_0}+c'_{k_0}\frac{m\cdot s_{d_i}}{s_{d_i}+s_{d_j}}-c'_{k_0}k_0\right) \\[6pt] &= \tilde{f}(S,T') \end{aligned}$$

The linear approximation will be done around the point $k_0 = m\cdot s_{d_i}/(s_{d_i}+s_{d_j})$. This results in a simple form for the resulting approximation, and also the mass of the product of binomial coefficients concentrates around this point. Set $\nu = s_{d_i}+s_{d_j}$ to examine the tradeoff between s_{d_i} and s_{d_j}.

$$\begin{aligned} \tilde{f}(S,T') &= \binom{\nu}{m}(c_{\frac{m\cdot s_{d_i}}{\nu}}) \\[6pt] &= \binom{\nu}{m}\frac{((\nu-s_{d_i})(\nu-m)+\nu)((s_{d_i}+1)\nu-m\cdot s_{d_i})}{\left(\begin{array}{l} p^u_{d_i}\nu(s_{d_i}+1)((\nu-s_{d_i})(\nu-m)+\nu)+ \\[4pt] p^u_{d_j}\nu(\nu-s_{d_i}+1)(\nu+s_{d_i}(\nu-m))+ \\[4pt] \beta((s_{d_i}+1)\nu-m\cdot s_{d_i})((\nu-s_{d_i})(\nu-m)+\nu) \end{array}\right)} \end{aligned}$$

Direct calculation can easily show that $\frac{d^2\tilde{f}}{d(s_{d_i})^2} \ge 0$. Then it is straightforward to show that $f(s_{d_i}=0) \ge f(s_{d_j}=0)$.

$\square$

74

Therefore, in looking for a maximum we can assume that every user except u either always visits d or always visits $d_{|\Delta|}$. We can use the same idea with d and $d_{|\Delta|}$ as was used with d_i and d_j and consider $E[Y|X_D(u) = d]$ as we trade off the number of users visiting them. Doing this shows that a maximum is obtained either when all users but u always visit d or when they always visit $d_{|\Delta|}$. This is the worst-case expected anonymity.

Theorem 9. *A maximum of $E[Y|X_D(u) = d]$ occurs when either $p_d^v = 1$ for all $v \neq u$ or when $p_{d_{|\Delta|}}^v = 1$ for all $v \neq u$.*

Proof. Assume, following Lemma 7, that $(p^v)_{v \neq u}$ is such that $p_d^v = 1$ or $p_{d_{|\Delta|}}^v = 1$ for all $v \neq u$.

The expected posterior probability can be given in the following variation on Equation 6.8:

$$E[Y|X_D(u) = d] = b(1-b)p_d^u + b^2 + \sum_{S:u \in S} \sum_{t=0}^{|S|} b^{n-|S|+t}(1-b)^{2|S|-t} Y(S, t)$$

Here $Y(S, t)$ is the expectation of Y conditional on the event that S is the set of users with unobserved inputs and t is the size of the subset of S with observed outputs.

Take some arbitrary S containing u. Let $s = |S|$, and let $d = |\{s \in S - u : p_d^s = 1\}|$. $Y(S, t)$ can be expressed as follows, where k represents the number of observed outputs of users in S with destination d:

$$Y(S, t) = \sum_{k=0}^{t} \binom{d+1}{k}\binom{s-d-1}{t-k}\left[\frac{p_d^u\binom{d}{k-1}\binom{s-d-1}{t-k} + p_d^u\binom{d}{k}\binom{s-d-1}{t-k}}{p_d^u\binom{d}{k-1}\binom{s-d-1}{t-k} + p_{d_{|\Delta|}}^u\binom{d}{k}\binom{s-d-1}{t-k-1} + \binom{d}{k}\binom{s-d-1}{t-k}}\right]$$

$$= \sum_{k=0}^{t} \binom{d+1}{k}\binom{s-d-1}{t-k}\left[\frac{p_d^u(d+1)(s-d-t+k)}{\begin{array}{l} p_d^u(d+1)(s-d-t+k) + p_{d_{|\Delta|}}^u(d+1-k)(s-d) \\ + (1-p_d^u - p_{d_{|\Delta|}}^u)(d+1-k)(s-d-t+k) \end{array}}\right]$$

Because $Y(S, t)$ only depends on s, d, and t, we overload the notation and let $Y(s, t, d) = Y(S, t)$. The makes the dependence on d explicit. We will show that $Y(s, t, d)$ achieves a maximum over d at either $d = 0$ or $d = s - 1$, and that the maximizing value of d doesn't depend on s or t. These facts imply the theorem.

As in Theorem 7, we can view the expression for $Y(s, t, d)$ as the weighted convolution of binomial coefficients:

$$Y(s, t, d) = \sum_{k=0}^{t} \binom{d+1}{k}\binom{s-d-1}{t-k} a(d, k) \tag{6.9}$$

75

To make the analysis of this sum more tractable, we will approximate $a(d, k)$ with a function $\tilde{a}(d, k)$ that is linear in k. In the following analysis, we are only concerned with values k such that the binomial coefficients $\binom{d+1}{k}$ and $\binom{s-d-1}{t-k}$ are nonzero, that is, that satisfy $\max(0, t-s+d+1) \leq k \leq \min(d+1, t)$.

Observe that $\frac{\partial^2 a}{\partial k^2} \neq 0$. This is clear from a direct calculation of the second derivative. Therefore it is of constant sign. We will use this fact and linearly approximate a with both a tangent line and a line through the endpoints. One of the two approximations will overestimate $a(d, k)$ for all k in range for a fixed d. Therefore one of the lines will overestimate $Y(s, t, d)$ at that value d.

The first approximation overestimates a when it is concave. Let $k_0 = t(d+1)/s$. Also let a' refer to $\frac{\partial a}{\partial k}$. The linear approximation $\tilde{a}$ is the tangent at k_0:

$$\tilde{a}(d, k) = (k - k_0)a'(d, k_0) + a(d, k_0)$$

Then there is a closed-form for the sum using this approximation:

$$\tilde{Y}(s, t, d) = \sum_{k=0}^{t} \binom{d+1}{k}\binom{s-d-1}{t-k}\tilde{a}(d, k)$$
$$= \binom{s}{t}a(d, k_0)$$

A calculation of $\frac{\partial \tilde{Y}}{\partial d}$ shows that it is nonnegative. Therefore a maximum of $\tilde{Y}$ is found at $d = n - 1$.

The second approximation overestimates a when it is convex. The endpoints of the range of k are $k_0 = \max(0, t-s+d+1)$ and $k_1 = \min(d+1, t)$. Let $\tilde{a}$ be the line going through the function at these points:

$$\tilde{a}(d, k) = \frac{(a(d, k_1) - a(d, k_0))(k - k_0)}{k_1 - k_0} + a(d, k_0)$$

Substituting $\tilde{a}$ for a in the expression for $Y(s, t, d)$ we get:

$$\tilde{Y}(s, t, d) = \binom{s}{t}a(d, k_0) + \binom{s-1}{t-1}\frac{(d+1)(a(d, k_1) - a(d, k_0))}{k_1 - k_0}$$

There are four possible cases for the values of k_1 and k_0.

1. $t \leq s - d - 1 \wedge t \leq d + 1$

   In this case $k_0 = 0$, $k_1 = t$, and $d \in [t-1, s-1-t]$. Then it can be shown that $\tilde{Y}(s, t, t-1) \geq \tilde{Y}(s, t, d)$. Therefore the maximum of $\tilde{Y}$ is obtained at $d = t - 1$.

2. $t \le s - d - 1 \wedge t \ge d + 1$

   In this case $k_0 = 0$, $k_1 = d + 1$, and $d \in [0, \min(s - t - 1, t - 1)]$. It can be easily shown that $\frac{\partial \tilde{Y}}{\partial d} \le 0$. Therefore the maximum of $\tilde{Y}$ is obtained at $d = 0$.

3. $t \ge s - d - 1 \wedge t \le d + 1$

   In this case $k_0 = t - s + d + 1$, $k_1 = t$, and $d \in [\max(s - t - 1, t - 1), s - 1]$. It can be shown that if $\frac{\partial^2 \tilde{Y}}{\partial d^2} \le 0$ then $\frac{\partial \tilde{Y}}{\partial d} \ge 0$. Therefore the maximum of $\tilde{Y}$ is obtained at $d = \max(s - t - 1, t - 1)$ or $d = s - 1$.

4. $t \ge s - d - 1 \wedge t \ge d + 1$

   In this case, $k_0 = t - s + d + 1$, $k_1 = d + 1$, and $d \in [s - t - 1, t - 1]$. It is easily shown that $\frac{\partial^2 \tilde{Y}}{\partial d^2} \ge 0$. Therefore the maximum of $\tilde{Y}$ is obtained at $d = s - t - 1$ or $d = t - 1$.

Thus we have the following set of points as candidate maxima: $\{0, s - t - 1, t - 1, s - 1\}$.

When $t \le s - t$, the intervals of d with the same endpoints are $[0, t-1]$, $[t-1, s-1-t]$, and $[s-1-t, s-1]$. We just showed that the maximum in these intervals is at 0, $t-1$, and $s-1-t$ or $s-1$, respectively. Therefore $\tilde{Y}(s, t, 0) \ge \tilde{Y}(s, t, t - 1) \ge \tilde{Y}(s, t, s - 1 - t)$, and so the maximum over the whole interval is at $d = 0$ or $d = s - 1$.

When $t \ge s - t$, the intervals of d with the same endpoints are $[0, s - 1 - t]$, $[s - 1 - t, t - 1]$, and $[t - 1, s - 1]$. We showed above that the maximum in these intervals is 0, $s - 1 - t$ or $t - 1$, and $t - 1$ or $s - 1$, respectively. Therefore $\tilde{Y}(s, t, 0) \ge \tilde{Y}(s, t, s - 1 - t)$, and so the maximum over the whole interval is at one of $\{0, t - 1, s - 1\}$. Examination of the expressions for $\tilde{Y}$ at these values shows that if $\tilde{Y}(s, t, 0) \le \tilde{Y}(s, t, t - 1)$, then $\tilde{Y}(s, t, t - 1) \le \tilde{Y}(s, t, s - 1)$. Therefore the maximum is always found at $d = 0$ or $d = s - 1$.

The sum in Equation 6.9 for $Y(s, t, d)$ has only one term at $d = 0$ and $d = s - 1$, and therefore $\tilde{Y}(s, t, d) = Y(s, t, d)$ at these points. $\tilde{Y}$ is an overestimate of Y, and thus $Y(s, t, d)$ is also a maximum at $d = 0$ or $d = s - 1$. $\qquad \square$

6.3.2 Asymptotic Anonymity

The exact value of the maximum of $E[Y|X_D(u) = d]$ is not simple to express, but we can give a straightforward approximation for large user populations n. We focus on large n, because anonymity networks, and onion routing in particular, are understood to have the best chance at providing anonymity when they have many users. Furthermore, Tor is currently used by an estimated 200,000 people.

Theorem 10. *When $p^v_{d_{|\Delta|}} = 1$, for all $v \neq u$,*

$$E[Y|X_D(u) = d] = b + b(1-b)p^u_d + (1-b)^2 p^u_d \left(\frac{1-b}{1 - (1 - p^u_{d_{|\Delta|}})b} + O\left(\sqrt{\frac{\log n}{n}} \right) \right). \qquad (6.10)$$

Proof. Let $f(n)$ represent the conditional expectation of the posterior probability in the case that u's input and output are not observed. Observe that in the case that u's output is observed $Y = 1$, because u is the only user that visits d. Thus $E[Y|X_D(u) = d] = b + b(1-b)p^u_d + (1-b)^2 f(n)$.

When u's input and output are unobserved, the value of Y depends on the number of other users with unobserved inputs, s, and the number of those s users with observed outputs, t. We can then express f as:

$$
\begin{aligned}
f(n) &= E[Y|X_D(u) = d \wedge X_I(u) = 0 \wedge X_O(u) = 0] \\
&= \sum_{s=0}^{n-1} (1-b)^s b^{n-1-s} \binom{n-1}{s} \sum_{t=0}^{s} b^t (1-b)^{s-t} \binom{s}{t} \frac{p^u_d \binom{s}{t}}{p^u_{d_{|\Delta|}} \binom{s}{t-1} + \binom{s}{t}} \\
&= \sum_{s=0}^{n-1} (1-b)^s b^{n-1-s} \binom{n-1}{s} \sum_{t=0}^{s} b^t (1-b)^{s-t} \binom{s}{t} \frac{p^u_d (s - t + 1)}{p^u_{d_{|\Delta|}} t + s - t + 1}
\end{aligned}
$$

Consider the inside sum. It is equal to the expected value of $g_0(s, T) = \frac{p^u_d (s - T + 1)}{p^u_{d_{|\Delta|}} T + s - T + 1}$, where $T \sim Bin(s, b)$. As s gets large, Chernoff bounds on the tails show that they contribute little to the expectation. Let $\varepsilon_0(s)$ be the terms of the sum for which t is more than $\sqrt{s \log s}$ from its expectation, $\mu_0 = bs$:

$$
\begin{aligned}
\varepsilon_0(s) &= \sum_{t: |t - \mu_0| > \sqrt{s \log s}} b^t (1-b)^{s-t} \binom{s}{t} g_0(s, t) \\
&\leq \sum_{t: |t - \mu_0| > \sqrt{s \log s}} b^t (1-b)^{s-t} \binom{s}{t} && \text{because } g_0 \leq 1 \\
&\leq 2e^{-c_0 \log s / b} && \text{for any } c_0 < 1/2 \text{ and large } s \\
& && \text{by Chernoff's inequality} \\
&= s^{-c_1} && \text{for some } c_1 > 1/2
\end{aligned}
$$

Now we look at how much the values of g_0 inside the tail differ from the value of g_0 at its expectation. Let $\varepsilon_1(s, t) = g_0(s, t) - g_0(s, \mu_0)$ be this difference. It can be shown through direct calculation that $\frac{\partial \varepsilon_1}{\partial t} \leq 0$

and $\frac{\partial^2 \varepsilon_1}{\partial t^2} \leq 0$. Therefore, for values of t that are within $\sqrt{s \log s}$ of μ_0,

$$
\begin{aligned}
|\varepsilon_1(s,t)| &\leq \left| \varepsilon_1\left(s, \mu_0 + \sqrt{s \log s} \right) \right| \\
&= \frac{p_d^u p_{d_{|\Delta|}}^u (1 + 1/s)}{\left(1 - \left(1 - p_{d_{|\Delta|}}^u \right) b + 1/s \right) \left(\sqrt{\frac{s}{\log s}} \left(1 - \left(1 - p_{d_{|\Delta|}}^u \right) b \right) - (1 - p) + (s \log s)^{-1/2} \right)} \\
&= O\left(\sqrt{\frac{\log s}{s}} \right)
\end{aligned}
$$

Looking now at the outside sum we have

$$
f(n) = \sum_{s=0}^{n-1} (1-b)^s b^{n-1-s} \binom{n-1}{s} \left[\frac{p_d^u(s(1-b)+1)}{s\left(1 - \left(1 - p_{d_{|\Delta|}}^u \right) b \right) + 1} + O(s^{-c_1}) + O\left(\sqrt{\frac{\log s}{s}} \right) \right].
$$

This is the expectation of a function of a binomially distributed random variable with mean $\mu_1 = (1-b)(n-1)$. Let $\varepsilon_2(n)$ be the parts of this sum that are greater than $k(n-1)$ from μ_1, $k < \min(b, 1-b)$:

$$
\begin{aligned}
\varepsilon_2(n) &\leq \sum_{s:|s-\mu_1|>k(n-1)} (1-b)^s b^{n-1-s} \binom{n-1}{s} && \text{because inner sum is at most 1} \\
&\leq 2e^{-c_2 k^2 (n-1)/(1-b)} && \text{for some } c_2 > 0 \text{ by Chernoff's inequality} \\
&= O(e^{-c_3 n}) && c_3 = c_2 k^2/(1-b)
\end{aligned}
$$

Let g_1 be the non-vanishing inner term of $f(n)$:

$$
g_1(s) = \frac{p_d^u(s(1-b)+1)}{s\left(1 - \left(1 - p_{d_{|\Delta|}}^u \right) b \right) + 1}.
$$

As s grows it approaches a limit of

$$
g_1^* = \frac{p_d^u(1-b)}{1 - \left(1 - p_{d_{|\Delta|}}^u \right) b}.
$$

Let $\varepsilon_3(s) = g_1(s) - g_1^*$ be the difference from this limit. Direct calculation shows that $\frac{d\varepsilon_3}{ds} \leq 0$ and $\frac{d^2\varepsilon_3}{ds^2} \geq 0$.

Therefore, for values of s within $k(n-1)$ of μ_1,

$$|\varepsilon_3(s)| \leq \varepsilon_3(\mu_1 - k(n-1))$$

$$= bp_d^u p_{d|\Delta|}^u / \left[\left(1 - \left(1 - p_{d|\Delta|}^u\right)b\right)\left((n-1)(1-b-k)\left(1 - \left(1 - p_{d|\Delta|}^u\right)b\right) + 1\right)\right]$$

$$= O(1/n)$$

Now we can show the desired asymptotic expression for the entire sum:

$$f(n) = O(e^{-c_3 n}) + \sum_{s=\mu_1 - k(n-1)}^{\mu_1 + k(n-1)} (1-b)^s b^{n-1-s} \binom{n-1}{s} \left[g_1^* + \varepsilon_3(s) + O(s^{-c_1}) + O\left(\sqrt{\frac{\log s}{s}}\right)\right]$$

$$= g_1^* + O(e^{-c_3 n}) + O(1/n) + O(n^{-c_1}) + O\left(\sqrt{\frac{\log n}{n}}\right)$$

$$= \frac{p_d^u(1-b)}{1 - \left(1 - p_{d|\Delta|}^u\right)b} + O\left(\sqrt{\frac{\log n}{n}}\right)$$

$\square$

Theorem 11. *When $p_d^v = 1$, for all $v \neq u$,*

$$E[Y|X_D(u) = d] = b^2 + b(1-b)p_d^u + (1-b)\frac{p_d^u}{1 - (1-p_d^u)b} + O\left(\sqrt{\frac{\log n}{n}}\right). \tag{6.11}$$

Proof. The proof of this theorem is similar to that of Theorem 10.

Let $f(n)$ represent the conditional expectation of the posterior probability in the case that u's input is not observed. Thus $E[Y|X_D(u) = d] = b^2 + b(1-b)p_d^u + (1-b)f(n)$.

We can express f as

$$f(n) = E[Y|X_D(u) = d \wedge X_I(u) = 0]$$

$$= \sum_{s=0}^{n-1}(1-b)^s b^{n-1-s}\binom{n-1}{s}\sum_{t=0}^{s+1} b^t(1-b)^{s+1-t}\binom{s+1}{t}\frac{p_d^u\binom{s}{t} + p_d^u\binom{s}{t-1}}{\binom{s}{t} + p_d^u\binom{s}{t-1}}$$

$$= \sum_{s=0}^{n-1}(1-b)^s b^{n-1-s}\binom{n-1}{s}\sum_{t=0}^{s+1} b^t(1-b)^{s+1-t}\binom{s+1}{t}\frac{p_d^u(s+1)}{s - (1-p_d^u)t + 1}.$$

As s gets large, Chernoff bounds on the tails of the inner sum show that they contribute little to the expectation. Let $g_0(s,t) = \frac{p_d^u(s+1)}{s-(1-p_d^u)t+1}$. Let $\varepsilon_0(s)$ be the terms of the inner sum for which t is more than

$\sqrt{s \log s}$ from $\mu_0 = b(s+1)$:

$$\varepsilon_0(s) = \sum_{t:|t-\mu_0|>\sqrt{(s+1)\log(s+1)}} b^t(1-b)^{s+1-t}\binom{s+1}{t}g_0(s,t)$$

$$\leq \sum_{t:|t-\mu_0|>\sqrt{(s+1)\log(s+1)}} b^t(1-b)^{s+1-t}\binom{s+1}{t} \qquad \text{because } g_0 \leq 1$$

$$\leq 2e^{-c_0\log(s+1)/b} \qquad \text{for any } c_0 < 1/2 \text{ and large enough } s$$

$$\qquad \text{by Chernoff's inequality}$$

$$= (s+1)^{-c_1} \qquad c_1 > 1/2$$

Let $\varepsilon_1(s,t) = g_0(s,t) - g_0(s,\mu_0)$. It can be shown through direct calculation that $\frac{\partial \varepsilon_1}{\partial t} \geq 0$ and $\frac{\partial^2 \varepsilon_1}{\partial t^2} \geq 0$. Therefore, for values of t that are within $\sqrt{s \log s}$ of μ_0,

$$|\varepsilon_1(s,t)| \leq \varepsilon_1\left(s, \mu_0 + \sqrt{(s+1)\log(s+1)}\right)$$

$$= \frac{p_d^u(1-p_d^u)}{(1-(1-p_d^u)b)\left(\sqrt{\frac{s+1}{\log(s+1)}}(1-(1-p_d^u)b)-(1-p)\right)}$$

$$= O\left(\sqrt{\frac{\log s}{s}}\right)$$

Looking now at the outside sum we have

$$f(n) = \sum_{s=0}^{n-1}(1-b)^s b^{n-1-s}\binom{n-1}{s}\left[\frac{p_d^u}{1-(1-p_d^u)b}+O(s^{-c_1})+O\left(\sqrt{\frac{\log s}{s}}\right)\right].$$

Let $\varepsilon_2(n)$ be the parts of this sum that are greater than $k(n-1)$ from $\mu_1 = (n-1)(1-b)$, $k < \min(b, 1-b)$:

$$\varepsilon_2(n) \leq \sum_{s:|s-\mu_1|>k(n-1)} (1-b)^s b^{n-1-s}\binom{n-1}{s} \qquad \text{because inner sum is at most 1}$$

$$\leq 2e^{-c_2 k^2(n-1)/(1-b)} \qquad \text{for some } c_2 > 0 \text{ by Chernoff's inequality}$$

$$= O(e^{-c_3 n}) \qquad c_3 = c_2 k^2/(1-b)$$

Now we can show the desired asymptotic expression for the entire sum:

$$f(n) = O(e^{-c_3 n}) + \sum_{s=(1-b-k)(n-1)}^{(1-b+k)(n-1)} (1-b)^s b^{n-1-s} \binom{n-1}{s} \left[\frac{p_d^u}{1 - (1-p_d^u)b} + O(s^{-c_1}) + O\left(\sqrt{\frac{\log s}{s}}\right) \right]$$

$$= \frac{p_d^u}{1 - (1-p_d^u)b} + O(e^{-c_3 n}) + O(n^{-c_1}) + O\left(\sqrt{\frac{\log n}{n}}\right)$$

$$= \frac{p_d^u}{1 - (1-p_d^u)b} + O\left(\sqrt{\frac{\log n}{n}}\right)$$

$\square$

To determine which distribution is the worst case for large n, simply examine the difference between the limits of the expressions in Theorems 10 and 11. It is clear from this that the worst-case distribution is $p_d^v = 1, \forall v \neq u$, only when $p_{d_{|\Delta|}}^u \geq \frac{(1-b)(1-p_d^u)^2}{p_d^u(1+b)-b}$. This happens when $p_d^u \geq 1/2$ and $p_{d_{|\Delta|}}^u$ is near $1-p_d^u$. For $p_{d_{|\Delta|}}^u$ small, which we would expect as it must be less than $1/|\Delta|$, the worst-case distribution is $p_{d_{|\Delta|}}^v = 1, \forall v \neq u$. In this case the expected assigned probability is about $b + (1-b)p_d^u$. This can be viewed as decreasing the "innocence" of u from $1 - p_d^u$ to $(1-b)(1-p_d^u)$. It is also equal to the lower bound on anonymity in onion routing when the adversary controls a fraction $\sqrt{b}$ of the network.

6.4 Typical Distributions

It is unlikely that users of onion routing will ever find themselves in the worst-case situation. The necessary distributions just do not resemble what we expect user behavior to be like in any realistic use of onion routing. Our worst-case analysis may therefore be overly pessimistic. To get some insight into the anonymity that a typical user of onion routing can expect, we consider a more realistic set of users' destination distributions in which each user selects a destination from a common Zipfian distribution. This model of user behavior is used by Shmatikov and Wang [76] to analyze relationship anonymity in mix networks and is motivated by observations that the popularity of sites on the web follows a Zipfian distribution. Our results show that a user's expected assigned probability is close to $b^2 + (1-b^2)p_d^u$ for large populations, which is the best that can be expected when there is a b^2 probability of total compromise.

Let each user select his destination from a common Zipfian distribution p: $p_{d_i} = 1/(\mu i^s)$, where $s > 0$

and $\mu = \sum_{i=1}^{|\Delta|} 1/i^s$. It turns out that the exact form of the distribution doesn't matter as much as the fact that it is common among users.

Theorem 12. *When $p^v = p^w$, for all $v, w \in U$,*

$$E[Y|X_D(u) = d] = b^2 + (1 - b^2)p_d^u + O(1/n)$$

Proof. Let p be the common destination distribution. The expected assigned probability can be expressed as

$$E[Y|X_D(u) = d] = b^2 + b(1-b)p_d^u + (1-b)\sum_{s=1}^{n} b^{n-s}(1-b)^{s-1}\binom{n-1}{s-1}$$

$$\sum_{t=0}^{s}(1-b)^{s-t}b^t\left[\binom{s-1}{t-1}\sum_{\Delta \in D^t : \Delta_1 = d}\prod_{i=2}^{t}p_{\Delta_i}\psi(s,\Delta) + \binom{s-1}{t}\sum_{\Delta \in D^t}\prod_{i=1}^{t}p_{\Delta_i}\psi(s,\Delta)\right].$$

Here, s represents the size of the set of users with unobserved inputs, t represents the size of the subset of those s users that also have observed outputs, Δ represents the t observed destinations, and $\psi(s,\Delta)$ is the posterior probability.

Let $\Delta_d = |\{i : \Delta_i = d\}|$. Let $x^{\underline{n}}$ denote the falling power $x(x-1)\ldots(x-n+1)$. The posterior probability ψ can be expressed simply as

$$\psi(s,\Delta) = \frac{\Delta_d(s-1)^{\underline{t-1}} + p_d(s-1)^{\underline{t}}}{s^{\underline{t}}}$$

$$= (\Delta_d + p_d(s-t))/s$$

The sum $\sum_{\Delta \in D^t : \Delta_1 = d}\prod_{i=2}^{t}p_{\Delta_i}\psi(s,\Delta)$ calculates the expectation for ψ conditioned on s, t, and the destination of u being observed. The expression for ψ shows that this depends linearly on the expected value of Δ_d. This expectation is simply $1 + p_d(t-1)$, because one destination in this case is always d, and each of the other $t-1$ is d with probability p_d. The sum $\sum_{\Delta \in D^t}\prod_{i=1}^{t}p_{\Delta_i}\psi(s,\Delta)$ similarly depends linearly on the expectation of Δ_d, which in this case is $p_d t$.

With this observation, it is a straightforward calculation to show that the inner sum over t is simply:

$$b\frac{p_d(s-1)+1}{s} + (1-b)p_d$$

We insert this into the larger sum and simplify:

$$E[Y|X_D(u) = d] = b^2 + b(1-b)p_d^u + (1-b)\sum_{s=1}^{n} b^{n-s}(1-b)^{s-1}\binom{n-1}{s-1}\left[b\frac{p_d(s-1)+1}{s} + (1-b)p_d\right]$$

$$= b^2 + b(1-b)p_d^u + (1-b)\left[b\left(p_d + \frac{(1-p_d)(1-(1-b)^{n+1})}{b(n+1)}\right) + (1-b)p_d\right]$$

$$= b^2 + (1-b^2)p_d^u + O(1/n)$$

$\square$

Chapter 7

Improving Onion Routing Through Trust

7.1 Introduction

When designing or analyzing anonymous communication networks, researchers generally assume that all nodes routing traffic are equally trusted. But this typically is incorrect. There is much information available to those selecting routes that can affect trust: information about who runs some components of the infrastructure, what computing platforms are used, how long and how reliably some components have been running, etc. And if anonymous-communication protocols were to begin taking trust into account, then even more extensive and diverse sources of trust information might be available.

We consider using trust information to improve the relationship anonymity of onion routing by choosing circuits that are less likely to contain adversarial routers. We first set out a simple model that should facilitate reasoning about using trust in routing. We define trust simply to be the probability that an attempt by the adversary to control a node fails. This captures both the uncertainty in the adversary's ability to subvert a given router and the uncertainty in a user's knowledge about a giver router. Fully-trusted nodes, such as those run by the user, can be modeled as having a trust value of one, and untrusted nodes, such as those run by the adversary, can be modeled as having a trust value of zero.

The adversary cannot be expected to attempt to compromise every router on a large network such as the Internet. Therefore we consider an adversary with a bound on the number of routers he can attack. Also,

the adversary can be expected to know or learn something about how users choose routes, and therefore the adversary is given this information before deciding which routers to attempt to compromise. This sets up a game between the user and adversary, in which the user chooses a distribution over circuits and the adversary decides which routers to attack to minimize the user's anonymity.

The problem we then focus on is to maximize over distributions on circuits the minimum anonymity over the possible routers the adversary can attack. The analyses in Chapters 5 and 6 show that the first and last routers of the users' circuits determine the users' anonymity, and therefore we can focus just on how they are chosen. We look at the problem of minimizing the chance that the adversary controls the first and last routers. This is primarily done for mathematical convenience, as opposed to optimizing the more complicated probabilistic metric studied in Chapter 6. However, while this approach does not always optimize anonymity, we can show that for the case in which users choose destinations identically, examined in Section 6.4, it obtains anonymity that is within a small additive error of optimal.

The essential problem for the user is to balance between choosing routers that are hard for the adversary to compromise and choosing from a large number of routers to prevent the adversary from being able to focus his attacks. We first analyze the general case of this problem, in which there is an arbitrary number of trust levels. We observe that a straightforward algorithm to calculate an optimal distribution runs in time exponential in the size of the adversary. We consider a natural simplification of looking at distributions on individual nodes rather than pairs of nodes and considering the product distribution as an approximation of the joint distribution on pairs. We find two optimal distributions over single nodes, but we then show that optimal distributions on pairs are arbitrarily better than products of those optimal distributions on single nodes.

In practice, it is unlikely that one can realistically assign many different levels of trust, and so we next consider restricting to the case where there are only two trust levels for nodes in the network. We describe three distributions and prove that in every case one of them must be optimal. Lastly, we discuss determining in practice when one of the three distributions is optimal based on the values of the system variables: trust values, size of the trusted and untrusted sets, and the size of the adversary.

In addition to the difficulty of accurately assigning trust levels, we acknowledge several other weaknesses in our modeling and analysis of trust. Our model does not explicitly consider users with different trust levels. We expect that in practice users have different opinions about which routers are trustworthy, and even who the adversary is. Our model can be applied by each user individually, but the recommendations for router

selection are then likely to be different, and this opens the opportunity for the adversary to infer the source of a circuit by its routers. Also, the trust information reflected in circuit selection may itself be sensitive. For example, users may wish to run a set of honest nodes secretly so that they aren't noticed and attacked by the adversary. We do not consider to what extent the adversary can learn this information.

Despite these weaknesses, we intend that our results have some useful application in the practice of anonymous communication. Indeed, the success of onion routing in practice is a motivation for studying it and trying to improve it. Therefore, it is useful to consider how our modeling and analysis translates to Tor [83], the widely-deployed and used public onion-routing network. Because Tor nodes are run by volunteers, it is easy for the adversary to simply set up routers and try to get them selected as the first and last positions on users' circuits. When successful, such an attack completely deanonymizes a user. In our model, these routers would be assigned a trust level of zero, and our objective function is exactly to avoid choosing such routers as the first and last routers on a circuit. Thus our recommendations for using trust may seem to translate well into avoiding a very plausible attack on a deployed system.

The Tor system also gives us a suggestion for a potential improvement to our recommendations for using trust to improve anonymity. One way that Tor reduces the threat of linking destinations to sources is by the use of *entry guard nodes*, a small number of routers that a single client uses persistently to connect to the Tor network. The selection of the set of guard nodes happens infrequently, and circuits are built by selecting the first router on a circuit randomly from the set of entry guards. If a client has chosen guard nodes that are not compromised, it can never be linked to its destination by a pair of compromised first/last routers. Our analysis throughout the thesis is just with respect to one round of communication, while guard nodes are an attempt to deal with the problem of making connections over time. As part of a larger effort to understand and deal with the problems that arise with connections over time, it seems like the selection of entry guards could benefit from using trust.

7.2   Trust model

A user wants to use a network of onion routers for anonymous communication. He trusts some onion routers more than others in the sense that he trusts that they are less likely to attempt to compromise his anonymity. How should he take this trust into account when he selects his paths?

7.2.1 The model

To make this question concrete, we need to make the notions of trust, anonymity, and an adversary precise. We use the network and user models described in Chapter 2 and the onion-routing protocol described in Chapter 4.

We modify the adversary that is trying to compromise the user's anonymity. The adversary selects k routers in R that he will attempt to compromise and use for deanonymization. If a router is not selected, it cannot be used by the adversary in an attack.

When an onion router i is selected, the adversary fails to compromise it with probability t_i. This represents the user's trust in the router. It will be convenient to define $c_i = 1 - t_i$, the probability that the adversary *does* successfully compromise router i when he attempts to do so.

A user selects a path for a circuit from some probability distribution. If the adversary has selected and successfully compromised the first and last nodes on the chosen path, he can correlate timing patterns between the two, and the user has no anonymity. We seek to improve anonymity by minimizing the probability of this event. To calculate this probability, we need only look at the user's distribution over entry-and-exit-node pairs.

We would like to find the probability distribution over pairs of routers that minimizes the maximum chance that the adversary can obtain of selecting both members of the pair and successfully compromising them. More precisely, we want to find $p \in \Delta_{n(n-1)/2}$, that is, a probability distribution p over pairs in R, that minimizes

$$c(p) = \max_{K \subseteq R : |K| = k} \sum_{\{r,s\} \in \binom{K}{2}} p(r,s) c_r c_s.$$

For a set S and $j \leq |S|$, we use $\binom{S}{j}$ to represent the collection of all subsets of S of size j. Also, for convenience, we write $p(\{r, s\})$ as $p(r, s)$.

This problem complements our results in Chapter 6. There, we analyze anonymity under uniformly-random path selection. Here, we consider improving anonymity by optimizing our path-selection distribution.

7.2.2 The adversary

The adversary's size represents a limit of his resources. While caution might suggest designing against an adversary that can compromise the entire network as a worst case, the optimal strategy in this case would be to select the two most-trusted routers. This strategy fails our intuition for good strategies because a realistic

adversary does have limited resources, but in such a strategy is not hindered much by this as he can focus all his efforts on the trusted routers. Also, especially for large diverse networks, it is typically unrealistic to assume that the adversary has the capacity to attack the entire network. System and protocol designs have been shown to provide a guarantee against various types of failure or compromise as long as no more than some fixed threshold of nodes is compromised at any time, for example, Byzantine fault-tolerance.

The particular partial-network adversary from which our work derives is the roving adversary of Ostrovsky and Yung [65]. They introduced and were motivated by the concept of proactive security, in which an adversary could compromise arbitrary optimal sets of nodes given his current information. The roving adversary can potentially compromise every node in the network, but it can compromise no more than a fixed maximum number of nodes at any one time. Proactive security is concerned with properties that are resilient to such attacks. This can be useful for secret sharing and other distributed applications. The adversary model was applied to onion routing by Syverson et al. [80].

We alter the basic roving adversary model in two ways. First, to incorporate trust we add the idea that an adversary does not always succeed when attempting to compromise a node. Second, the adversary selects only one set to attack—there is only one move by the adversary. It may be useful to bring multiple rounds back in for future work. Though it is unlikely that the adversary could target the first and routers of an individual circuit (given the short duration of typical onion-routing circuits), roving could allow the adversary to learn various communication and trust properties of the network and its users.

The adversary is assumed to have prior knowledge of the distribution that is used to pick a route, and he uses this knowledge to pick the set of nodes that he will attempt to compromise. It is realistic in many settings to assume the adversary has such knowledge. For example, the probability distributions may be set in some software or common system parameters given to a wide group in which there is at least one compromised member. The adversary may also be able to infer trust information from outside knowledge about the user.

7.2.3 Trust

Trust is captured in our model with the probability t_i that the adversary's attempt to compromise a node fails. This notion accommodates several different means by which users in the real world might trust an onion router.

The probability might represent the user's estimate of how likely it is that the operator of a given node

is trying to provide, rather than break, anonymity. It might represent the user's faith in the security of a given node against outside attack.

To arrive at such conclusions, the users must rely on some outside knowledge. This might include knowledge of the organizations or individuals who run nodes, both knowledge of their technical competence and the likelihood of themselves harboring ill intent. It might also include knowledge of computing platforms on which a network node is running, geopolitical information about the node, knowledge about the hosting facility where a node might be housed or the service provider(s) for its access to the underlying communications network, etc.

Admittedly, it may not be the case that one can realistically assign specific probabilities to each node in the network separately. It is for this reason that we consider in sections 7.5 and 7.6 restriction to just two trust levels. Even if one cannot be certain of the probability of compromise to assign at one level or another, one may be in a position to know the divergence of those levels. This is particularly the case if one is considering nodes run by, for example, security or law-enforcement agencies of friendly governments or their contractors vs. the rest of the nodes on the network. Alternatively, one can imagine sets of nodes run by reputable human rights groups, NGOs, or human rights agencies of friendly governments.

Unlike many other areas, network performance or reliability reputation are not good bases for trust for anonymous communication. That is because an adversary that is focused on learning as much as possible about communication patterns has incentive to run the highest performing, most reliable nodes in the network. Thus, many of the usual metrics do not apply.

7.2.4   Anonymity

We use trust information to improve anonymity by minimizing the probability that the adversary controls the first and last routers in a user's circuit. If the adversary controls these routers, then he knows the destination of the user with certainty. Therefore it makes sense to make this as unlikely as possible.

Ideally, however, we would use trust information to directly maximize the metric for relationship anonymity used in Chapter 6. And using this metric may indeed lead to a different optimal algorithm for selecting circuits. For example, suppose that every user except u always visits a destination that u never visits. In this case, u should just minimize the chance that the adversary controls his last router, because if the adversary observes the destination of u, then he knows that u must be the source, and if he does not observe u's destination, then he cannot improve his prior knowledge. Minimizing the chance that the adversary con-

trols just the last router is in general incompatible with minimizing the chance that the adversary controls the first and last routers. For instance, suppose that the routers only have two trust levels, $t_1 \geq t_2$, with $U = \{r \in R : t_r = t_1\}$ and $V = \{r \in R : t_r = t_2\}$. In Section 7.5, we show that, in this case, for large enough $|V|$ the optimal strategy to avoid choosing compromised first and last routers is to choose them either both from U or both from V but never one from U and one from V. With large enough $|V|$ this strategy makes any given $v \in V$ arbitrarily more likely to be chosen as the last router than any $u \in U$. However, if we apply Theorem 14 to this case, then we can see that for large enough n the unique optimum to avoid choosing a compromised last router is to choose $u \in U$ with probability proportional to $1/c_1$ and $v \in V$ with probability proportional to $1/c_2$. This is therefore incompatible with the optimal distribution to avoid compromised pairs.

However, our intuition to minimize the chance of choosing compromised first and last routers may be justified in more likely scenarios of user behavior. In Section 6.4, we analyzed the case in which all users share a common destination distribution and discovered that the expected anonymity is dominated by the probability that the first and last routers are compromised. In fact, this observation still holds in our new model with trust. Let the common probability of choosing destination d be p_d. Suppose the adversary attempts to compromise some subset of routers K. Let users choose pairs of routers to be their first and last routers according to some distribution p. Let the following be the probabilities of the possible outcomes of choosing a circuit:

$p_{11} :$ probability that the first and last routers are compromised

$p_{01} :$ probability that the first router is uncompromised and the last router is compromised

$p_{10} :$ probability that the first first router is compromised and the last router is uncompromised

$p_{00} :$ probability that the first and last routers are uncompromised.

The following theorem generalizes Theorem 12.

Theorem 13. $E[Y|X_D(u) = d] = p_{11} + p_d(1 - p_{11}) + \frac{(1-p_d)(p_{01})(1-(p_{11}+p_{10})^n)}{(p_{01}+p_{00})n}.$

Proof. The proof is similar to the proof of Theorem 12. We can express the expected assigned probability

as

$$E[Y|X_D(u) = d] = p_{11} + p_{10}p_d^u + \sum_{s=1}^{n} \binom{n-1}{s-1} (p_{11} + p_{10})^{n-s}.$$

$$\sum_{t=0}^{s} (p_{01})^t (p_{00})^{s-t} \left[\binom{s-1}{t-1} \sum_{\Delta \in D^t : \Delta_1 = d} \prod_{i=2}^{t} p_{\Delta_i} \psi(s, \Delta) + \binom{s-1}{t} \sum_{\Delta \in D^t} \prod_{i=1}^{t} p_{\Delta_i} \psi(s, \Delta) \right],$$

where, again, s represents the size of the set of users with unobserved inputs, t represents the size of the subset of those s users that also have observed outputs, Δ represents the t observed destinations, and $\psi(s, \Delta)$ is the posterior probability that the destination of u is d.

As before,

$$\psi(s, \Delta) = (\Delta_d + p_d(s - t))/s. \tag{7.1}$$

Also as before, the two sums over $\Delta \in D^t$ are expectations of ψ conditioned on s and t, which we can calculate using Equation 7.1 and the fact that all users choose d with probability p_d:

$$\sum_{\Delta \in D^t : \Delta_1 = d} \prod_{i=2}^{t} p_{\Delta_i} \psi(s, \Delta) = (1 + p_d(s - 1))/s \tag{7.2}$$

and

$$\sum_{\Delta \in D^t} \prod_{i=1}^{t} p_{\Delta_i} \psi(s, \Delta) = p_d \tag{7.3}$$

With this observation, we can apply the binomial theorem to show that the inner sum over t is

$$(p_{01} + p_{00})^{s-1} \left[(1 - p_d)(p_{01})/s + (p_d)(p_{01} + p_{00}) \right].$$

We insert this into the larger sum and simplify:

$$E[Y|X_D(u) = d] = \quad p_{11} + p_{10}p_d^u + \sum_{s=1}^n \binom{n-1}{s-1} \cdot \tag{7.4}$$

$$(p_{11} + p_{10})^{n-s}(p_{01} + p_{00})^{s-1}\left[(1 - p_d)(p_{01})/s + (p_d)(p_{01} + p_{00})\right]$$

$$= \quad p_{11} + p_{10}p_d^u + (1 - p_d)(p_{01})\frac{(p_{11} + p_{10} + p_{01} + p_{00})^n - (p_{11} + p_{10})^n}{(p_{01} + p_{00})^n} + \tag{7.5}$$

$$(p_d)(p_{01} + p_{00})(p_{11} + p_{10} + p_{01} + p_{00})^{n-1}$$

$$= \quad p_{11} + p_d(1 - p_{11}) + (1 - p_d)(p_{01})\frac{1 - (p_{11} + p_{10})^n}{(p_{01} + p_{00})^n}. \tag{7.6}$$

Equation 7.5 follows by the binomial theorem. Equation 7.6 is obtained by observing that

$p_{11} + p_{10} + p_{01} + p_{00} = 1$. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad$ □

Given some distribution over pairs of routers, we can view the adversary as choosing K to maximize $E[Y|X_D(u) = d]$. Theorem 13 shows that, when users use a common destination distribution, the resulting anonymity is close to the approximation that only considers whether or not the first and last routers are compromised:

$$\left| \max_{K \subseteq R:|K|=k} E[Y|X_D(u) = d] - \max_{K \subseteq R:|K|=k} p_{11} + p_d(1 - p_{11}) \right| \le 1/n.$$

Therefore we can approximately minimize the maximum of $E[Y|X_D(u) = d]$ by minimizing the maximum of $p_{11} + p_d(1 - p_{11})$, and this yields an additive error of at most $1/n$. The probability p_{11} is greater than zero, and therefore the distribution that achieves the minimax of p_{11} achieves the minimax of $p_{11} + p_d(1 - p_{11})$. Finding the minimax of p_{11} is exactly the objective function we set out in our formulation of the trust problem.

The model does not include some other methods the adversary can use, for example congestion attacks [63, 28], denial-of-service attacks [8], latency [43], or destination fingerprinting [42, 51]. The attacks on which we focus are conceptually much simpler than these others, but more importantly, as noted in Section 7.1, none of these other attacks succeeds with as much certainty using as few resources as does this one. Note that such entry-exit correlation attacks could also be done by the links from source to the entry onion router on the entry side and links from the exit onion router to the destination on the exit side (or by the destination itself). For example, an autonomous system or internet exchange on these links could participate in a correlation attack [30, 64]. We focus, however, on just the attack as it can be performed by network

nodes. Besides simplifying analysis, the ease with which nodes can be added to the network makes this a reasonable restriction.

Using this model, the user's selection of the pair constituting the first and last onion routers on his path is the only relevant factor in his anonymity. The user may make this selection using any probability distribution p over pairs of routers.

7.2.5 Objective function

We set as our objective function to find the distribution on pairs of routers that minimizes the probability of circuit compromise over all possible sets that the adversary could choose:

$$\min_{p \in \Delta_{n(n-1)/2}} \max_{K \subseteq R : |K| = k} \sum_{\{r,s\} \in \binom{K}{2}} p(r,s) c_r c_s.$$

This provides a worst-case guarantee, and if the user has a distribution with a low worst-case value, he is guaranteed anonymity with high probability regardless of the adversary's actions. As a worst-case criterion, however, it may direct the user to protect against adversarial actions that are unlikely. Indeed, while the adversary's goal is to find the subset $K \subseteq R$ that maximizes his chance of compromise, it is easy to see that this problem in general is equivalent to the **NP**-hard problem CLIQUE. Therefore the adversary may fail in many cases to actually select the worst-case set.

7.3   Strategies for the general case

Given arbitrary trust values $t_1, \ldots, t_n$, we would like to find a polynomial-time algorithm that takes as input the trust values and outputs an optimal or near-optimal distribution p^*.

7.3.1   Exact algorithm

There is a straightforward formulation of this problem as a linear program. Let the set of variables be p_{ij}, $i, j \in R$. The following constraints ensure that p is a probability distribution:

$$\sum_{\{r,s\} \in \binom{R}{2}} p_{rs} = 1$$
$$0 \le p_{rs} \le 1 \qquad \text{for all } \{r, s\} \in \binom{R}{2}.$$

We want to find the distribution that satisfies the minimax criterion

$$\min_{p} \max_{K \in \binom{R}{k}} \sum_{\{r,s\} \in \binom{K}{2}} c_r c_s p(r,s).$$

For any fixed K, the sum

$$c(p,K) = \sum_{\{r,s\} \in \binom{K}{2}} p(r,s) c_r c_s$$

is linear in p. Therefore the minimax criterion minimizes the maximum of linear functions. We can thus transform it into a simple minimization problem by adding a slack variable t and some linear constraints. We force t to be greater than the maximum of our linear functions:

$$t - c(p,K) \geq 0 \quad \text{for all } K \in \binom{R}{k}$$

Then the objective function is simply $\min t$. Unfortunately, this linear program is of exponential size $\left(O(n^k)\right)$ because of the constraints for each subset.

7.3.2 Choosing a simple distribution

A straightforward simplification is to consider restricting the output to be a distribution in which the first and last routers are chosen independently and identically at random and then minimizing the probability that they are individually compromised.

Let p_R be a distribution on R. We consider the distribution p_R^* that minimizes the probability that an adversary chooses and successfully compromises a single router:

$$
\begin{aligned}
c(p_R) &= \max_{K \in \binom{R}{k}} \sum_{r \in K} p_R(r) c_r \\
p_R^* &= \operatorname*{argmin}_{p_R} c(p_R)
\end{aligned}
$$

The following theorem states that it is always optimal either to put all the probability on the most trusted router or to set the probabilities such that the values $c_i p_R(r_i)$ are equal for all $r_i \in R$.

Theorem 14. *Let $c_\mu = \min_j c_j$. Let p_R^1 put all the probability on the most trusted router:*

$$
p_R^1(r) = \begin{cases} 1 & \text{if } r = r_\mu \\ 0 & \text{otherwise} \end{cases}
$$

Let p_R^2 set probability inversely proportional to c_i:

$$
p_R^2(r_i) = \alpha/c_i
$$

where $\alpha = \left(\sum_i 1/c_i\right)^{-1}$.

Then

$$
c(p_R^*) = \begin{cases} c(p_R^1) & \text{if } c_\mu \le k\alpha \\ c(p_R^2) & \text{otherwise} \end{cases}
$$

Proof. Suppose p_R is an optimal distribution. Sort the routers so that $c_1 p_R(r_1) \ge c_2 p_R(r_2) \ge \ldots \ge c_n p_R(r_n)$. The set K that maximizes $\sum_{r \in K} c_r p_R(r)$ is then $\{r_1, r_2, \ldots, r_k\}$, and the value of p_R is $c(p_R) = \sum_{i=1}^{k} c_i p_R(r_i)$.

Let l be the largest index such that $c_l p_R(r_l) = c_k p_R(r_k)$.

If $l < n$, we could decrease $c_i p_R(r_i), k \le i \le l$ by moving $\epsilon c_k/c_i$ probability from r_i to r_{l+1}. This decreases $c_i r_i$ by $c_k \epsilon$ and increases $c_{l+1} p_R(r_{l+1})$ by $\epsilon c_{l+1} c_k/c_i$. For small enough ϵ we maintain that if $i < j$ then $c_i p_R(r_i) \ge c_j p_R(r_j)$, and therefore we reduce the value $c(p_R)$. Therefore p_R cannot be optimal, contradicting our assumption.

Thus it must be that $l = n$. Let m be the smallest index such that $c_m p_R(r_m) = c_k p_R(r_k)$. Assume that p_R is an optimal distribution that has the smallest m possible.

If $m = 1$, we are in the case that $c_i p_R(r_i) = c_j p_R(r_j)$ for $1 \le i, j \le n$. This is the distribution p_R^2.

Suppose $m > 1$. If $p_R(r_m) = 0$, then $c(p_R) = \sum_{i=1}^{m-1} c_i p_R(r_i)$. Let $c_\mu = \min_i c_i$. Because all of the probability is contained in a set that the adversary can completely select, we do not increase $c(p_R)$ by

96

moving all the probability to r_μ:

$$
\begin{aligned}
c(p_R) &= \sum_{i=1}^{m-1} c_i p_R(r_i) \\
&\geq \sum_{i=1}^{m-1} c_\mu p_R(r_i) \\
&= c_\mu.
\end{aligned}
$$

c_μ is equal to $c(p_R^1)$.

Now consider the case that $p_R(r_m) > 0$. Recall that $c_i p_R(r_i) = c_j p_R(r_j)$ for all pairs r_i, r_j, in the set $S = \{r_i, m \leq i \leq n\}$. Consider moving probability between r_{m-1} and S in a way that maintains the equality of $c_i p_R(r_i)$ for $r_i \in S$. This can be achieved by setting the probability of r_{m-1} to

$$
p_R'(r_{m-1}, t) = p_R(r_{m-1}) + t
$$

and the probability of $r_i \in S$ to

$$
p_R'(r_i, t) = p_R(r_i) - \frac{t}{c_i(\sum_{r_j \in S} 1/c_j)}.
$$

For small enough values of t, this preserves the property that if $i > j$ then $c_i p_R'(r_i, t) \leq c_j p_R'(r_j, t)$. Therefore $c(p_R') = \sum_{i=1}^{k} c_i p_R'(r_i, t)$. The fact that p_R' is linear in t makes $c(p_R')$ also linear in t for small enough values of t.

If $D_t c(p_R')|_{t=0} \geq 0$, then for $t < 0$ large enough $c(p_R')$ doesn't increase. This corresponds to moving probability from r_{m-1} to S, and the smallest t that maintains the ordering by $c_i p_R'(r_i)$ results in $c_{m-1} p_R'(r_{m-1}) = c_m p_R'(r_m)$. This contradicts the assumption about the minimality of the index m.

If $D_t c(p_R')|_{t=0} \leq 0$, then for $t > 0$ small enough $c(p_R')$ doesn't increase. This corresponds to moving probability from S to r_{m-1}. In fact, no positive value of t increases $c(p_R')$. This is because setting $t > 0$ decreases the probability of all r_i, $i > k$, and only increases the probability of $r_{m-1}, m \leq k$, and thus preserves the fact that $c(p_R') = \sum_{i=1}^{k} c_i p_R'(r_i, t)$. Therefore we can increase t until $c_i p_R'(r_i) = 0$ for all $r_i \in S$. This puts us in the case where $p_R'(r_m) = 0$, which we have already shown implies that $c(p_R') \geq c(p_R^1)$.

Thus we have shown that either p_R^1 or p_R^2 is an optimal distribution. $c(p_R^1) = c_1$ and $c(p_R^2) = k\alpha$. Therefore, if $c_1 \leq k\alpha$, $c(p_R^*) = c(p_R^1)$, and otherwise $c(p_R^*) = c(p_R^2)$. □

We might hope that the product distributions $p_R^1 \times p_R^1$ and $p_R^2 \times p_R^2$ over $R \times R$ are good approximations to an optimal distribution p^*. However, this is not the case, and we can find inputs such that $c(p_R^i)/c(p^*)$, $i \in \{1, 2\}$, is arbitrarily high. In fact, we can show this for slightly improved distributions p^1 and p^2 over $\binom{R}{2}$.

Notice that $p_R^i \times p_R^i$, $i \in \{1, 2\}$, puts positive probability on the user choosing the same router twice. The problem as formulated in Section 7.2 allows distributions only over distinct pairs of routers (that is, elements of $\binom{R}{2}$). This doesn't affect the optimum, however. There is always an optimal distribution that puts zero probability on $(r, r) \in R \times R$. Let p be a distribution on $R \times R$. Then let

$$
p'(r, s) = \begin{cases} 0 & \text{if } r = s \\ p(r, s) + q(r, s) & \text{otherwise} \end{cases}
$$

where, for all $r \in R$, $\sum_{s \neq r} q(r, s) = p(r, r)$ and, for all $r, s \in R, q(r, s) \geq 0$.

Lemma 8. $c(p') \leq c(p)$

Proof.

$$
\begin{aligned}
c(p) &= max_{K \subseteq R : |K| = k} \sum_{r \in K} p(r, r) c_r + \sum_{r, s \in K : r \neq s} p(r, s) c_r c_s \\
&= max_{K \subseteq R : |K| = k} \sum_{r \in K} \sum_{s \neq r} q(r, s) c_r + \sum_{r, s \in K : r \neq s} p(r, s) c_r c_s \\
&\leq max_{K \subseteq R : |K| = k} \sum_{r \in K} \sum_{s \neq r} q(r, s) c_r c_s + \sum_{r, s \in K : r \neq s} p(r, s) c_r c_s \\
&= max_{K \subseteq R : |K| = k} \sum_{r, s \in K : r \neq s} (p(r, s) + q(r, s)) c_r c_s \\
&= c(p')
\end{aligned}
$$

$\square$

Now assume that $c_1 \leq c_2 \leq \ldots \leq c_n$ and consider two distributions over $\binom{R}{2}$:

$$
p^1(r, s) = \begin{cases} 1 & \text{if } r = c_1 \wedge s = c_2 \\ 0 & \text{otherwise} \end{cases}
$$

and

$$p^2(r, s) = \frac{\alpha}{c_r c_s}$$

where $\alpha = \left(\sum_{\{r,s\} \in \binom{R}{2}} 1/(c_r c_s) \right)^{-1}$. By Lemma 8 $c(p^1) \leq c(p_R^1)$ and $c(p^2) \leq c(p_R^2)$.

Now let $\mathcal{I}_n = (c_1, \ldots, c_n, k)$ be a problem instance that, as n grows, satisfies

1. $c_1 = O(1/n)$.

2. $c_2 > c$ for some constant $c \in (0, 1)$.

3. $k = o(n)$

4. $k = \omega(1)$

For large enough n, $\mathcal{I}_n$ has an optimal value that is arbitrarily smaller than the values achieved by p^1 and p^2. Let $c(\mathcal{I}_n, p)$ be the value of $\mathcal{I}_n$ under distribution p.

Theorem 15.

$$c(\mathcal{I}_n, p^1)/c(\mathcal{I}_n, p^*) = \Omega\left(\frac{n}{k}\right) \tag{7.7}$$

$$c(\mathcal{I}_n, p^2)/c(\mathcal{I}_n, p^*) = \Omega(k) \tag{7.8}$$

Proof. The following distribution achieves the ratios in Eqs. 7.7 and 7.8. Let

$$p^3(r, s) = \begin{cases} \frac{\alpha}{c_r c_s} & \text{if } r = r_1 \\ 0 & \text{otherwise} \end{cases}$$

where $\alpha = \left(\sum_{i>1} 1/(c_1 c_i) \right)^{-1}$. This distribution puts weight on all distinct pairs that include r_1. It represents a middle approach between putting all the probability on the lightest pair, as p^1 does, and spreading the probability over all pairs, as p^2 does. The optimal distribution for each $\mathcal{I}_n$ only has higher ratios with p^1 and p^2 than p^3 does.

The ratio between p^1 and p^3 is

$$\begin{aligned}
\frac{c(\mathcal{I}_n, p^1)}{c(\mathcal{I}_n, p^3)} &= \frac{c_1 c_2}{(k-1)/\left(\sum_{i=2}^n 1/(c_1 c_i)\right)} \\
&\geq (1 + c_2(n-2)/c_n)/(k-1) \\
&= \Omega\left(\frac{n}{k}\right).
\end{aligned}$$

99

The ratio between p^2 and p^3 is

$$\frac{c(\mathcal{I}_n, p^2)}{c(\mathcal{I}_n, p^3)} = \frac{\binom{k}{2} \left(\sum_{i \neq j} 1/(c_i c_j) \right)^{-1}}{(k-1)/ \left(\sum_{i=2}^{n} 1/(c_1 c_i) \right)} \tag{7.9}$$

$$= \frac{k}{2} \left(1 + c_1 \frac{\sum_{2 \leq i < j \leq n} 1/(c_i c_j)}{\sum_{i=2}^{n} 1/c_i} \right)^{-1} \tag{7.10}$$

$$\geq \frac{k}{2} \left(1 + \frac{c_1}{2} \left(\sum_{i=2}^{n} 1/c_i - 1 \right) \right)^{-1} \tag{7.11}$$

$$= \Omega(k). \tag{7.12}$$

In Equation 7.11, $\sum_{i=2}^{n} 1/c_i$ is bounded by n because $c_i > c$, $i > 1$. The last line then follows because $c_1 = O(1/n)$. $\qquad \square$

Intuitively, the reason p^1 does arbitrarily worse than p^3 is that it doesn't take advantage of an adversary of size $o(n)$ by putting probability on $\Omega(n)$ pairs, while p^2 does arbitrarily worse than p^3 because it puts probability on pairs $\{r_i, r_j\}$, $i, j > 1$, that have $\Omega(n)$ times higher probability of being successfully compromised than pairs including r_1.

7.4   When pairing off, trust is everything

Allowing arbitrary trust values may be unnecessarily general. Users are unlikely to have precise knowledge of the probability of compromise for each onion router in the network. Instead, they seem more likely to have a few classes of trust into which they can partition the routers, or to have detailed knowledge about only a small number of routers. This fact may help us deal with the apparent computational intractability of the general problem. Also, the potentially complicated optima that result from arbitrary trust values may not satisfy other criteria for path-selection strategies that are outside the given problem formulation. For example, we may want the number of possible optimal strategies to be small so users share their behavior with many others, or we may want the strategies to be robust to small changes in trust values.

Therefore, we now consider the case that there are only two trust values. We refer to the nodes with higher trust as the *trusted* set, and nodes with lower trust as the *untrusted* set. This case is simple yet results in non-obvious conclusions, and also still provides practical advice to users.

In Section 7.5 we show that, when there are only two trust values, there are three strategies that are potentially optimal. But first we give here a lemma that allows us to consider only distributions that treat

the routers within a trust set identically. Note that this lemma holds for general trust values.

Lemma 9. *Let U be a set of routers with identical trust values c, where $|U| = m$. Let V be the rest of the routers, where $|V| = n$. Then the set of routers is $R = U \cup V$. There exists an optimal distribution p in which the following hold:*

1. *For all $\{u, v\}, \{w, x\} \in \binom{U}{2}$, $p(u, v) = p(w, x)$.*

2. *For all $v \in V$, $u, w \in U$, $p(v, u) = p(v, w)$.*

Proof. Consider some distribution over pairs $p : \binom{R}{2} \to [0, 1]$, $\sum_{\{r,s\} \in \binom{R}{2}} p(r, s) = 1$. Consider any subset $S \subseteq V$. Let X_S be a subset chosen randomly from all subsets X of size k such that $X \cap V = S$. Let $j = k - |S|$ be the size of $X_S \cap U$. Let $c(p, K)$ be the probability of compromise under p, given that set K is chosen by the adversary. That is,

$$c(p, K) = \sum_{\{r,s\} \in \binom{K}{2}} p(r, s) c_r c_s$$

We can calculate the expected probability of compromise of X_S as follows:

$$\mathbb{E}\left[c(p, X_S)\right]$$

$$= \left\{\binom{m}{j}^{-1} \sum_{T \subseteq \binom{U}{j}} \left[\begin{array}{l} \displaystyle\sum_{\{t,u\}\in\binom{T}{2}} p(t,u)c^2 + \\[2ex] \displaystyle\sum_{u\in T, v\in S} p(u,v)c\cdot c_v + \\[2ex] \displaystyle\sum_{\{v,w\}\in\binom{S}{2}} p(v,w)c_v c_w \end{array} \right] \right. \tag{7.13}$$

$$= \left\{ \begin{array}{l} \displaystyle\binom{m}{j}^{-1}\binom{m-2}{j-2}c^2 \sum_{\{t,u\}\in\binom{U}{2}} p(t,u) + \\[3ex] \displaystyle\binom{m}{j}^{-1}\binom{m-1}{j-1}c \sum_{v\in S, u\in U} p(v,u)c_v + \\[3ex] \displaystyle\sum_{\{v,w\}\in\binom{S}{2}} p(v,w)c_v c_w \end{array} \right. \tag{7.14}$$

$$= \left\{ \begin{array}{l} \displaystyle\frac{j(j-1)c^2}{m(m-1)} \sum_{\{t,u\}\in\binom{U}{2}} p(t,u) + \\[3ex] \displaystyle\frac{j\cdot c}{m} \sum_{v\in S, u\in U} p(v,u)c_v + \\[3ex] \displaystyle\sum_{\{v,w\}\in\binom{S}{2}} p(v,w)c_v c_w \end{array} \right. \tag{7.15}$$

There must be some set $T \subseteq U$ of size j such that $c(p, S \cup T)$ is at least the expectation expressed in Equation 7.15. If we modify p to treat all nodes in U the same, and thus satisfy the conditions in the statement of the lemma, every such T achieves the value in Equation 7.15. Let p' be this modified distribution:

$$p'(r,s) = \left\{ \begin{array}{ll} \sum_{\{t,u\}\in\binom{U}{2}} p(t,u)/\binom{m}{2} & \text{if } \{r,s\}\in\binom{U}{2} \\[2ex] \sum_{u\in U} p(r,u)/m & \text{if } r\in V, s\in U \\[2ex] \sum_{u\in U} p(s,u)/m & \text{if } r\in U, s\in V \\[2ex] p(r,s) & \text{if } \{r,s\}\in\binom{V}{2} \end{array} \right.$$

The probability of compromise for any value $S \cup T$ of X_S is

$$
c(p', S \cup T) = \binom{j}{2}\binom{m}{2}^{-1} \sum_{\{t,u\} \in \binom{U}{2}} p(t,u)c^2 +
$$

$$
\frac{j}{m} \sum_{v \in S} \sum_{u \in U} p(v,u)c_v c + \tag{7.16}
$$

$$
\sum_{\{v,w\} \in \binom{S}{2}} p(v,w)c_v c_w.
$$

Equations 7.15 and 7.16 are equal, and therefore

$\max_{T:|T|=j} c(p', S \cup T) \leq \max_{T:|T|=j} c(p, S \cup T)$. Because this holds for all $S \subseteq V$, $\max_{K:|K|=k} c(p', K) \leq$

$\max_{K:|K|=k} c(p, K)$. $\qquad \square$

7.5 Choosing pairs to avoid compromise

Now we analyze optimal distributions for selecting pairs when there are two trust values in the network, c_1

and c_2, with $c_1 \leq c_2$. We show that, in this case, one of the following strategies is always optimal: (i) choose

pairs such that $p(r,s)c_r c_s$ is equal for all $\{r,s\} \in \binom{R}{2}$, (ii) choose a pair of trusted routers uniformly at

random, or (iii) choose only fully-trusted or fully-untrusted pairs such that the adversary has no advantage

in attacking either trusted or untrusted routers. Distribution (i) is analogous to distribution p^2 as described

in Section 7.3.2. Distribution (ii) is analogous to distribution p^1 of Section 7.3.2, with the difference that

(ii) spreads probability to all the most-trusted routers and not just two. Distribution (iii) shows that

non-obvious distributions can exist even when the trust values are very restricted.

Let U be the trusted set, with trust value c_1, $|U| = m$. Let V be the untrusted set, with trust value c_2,

$|V| = n$.

Theorem 16. *Let $v_0 = \max(k - m, 0)$ and $v_1 = \max(k - n, 0)$. Then let $g_0 = \frac{v_0(v_0-1)}{n(n-1)}$ and $g_1 = \frac{v_1(v_1-1)}{m(m-1)}$.*

One of the following is an optimal distribution:

$$p(r,s) = \begin{cases} \dfrac{(c_2)^2}{\binom{m}{2}(c_2)^2+(mn)(c_1c_2)+\binom{n}{2}(c_1)^2} \\ \qquad\qquad \textit{if } \{r,s\} \in \binom{U}{2} \\[2em] \dfrac{(c_1c_2)}{\binom{m}{2}(c_2)^2+(mn)(c_1c_2)+\binom{n}{2}(c_1)^2} \\ \qquad\qquad \textit{if } (r,s) \in U \times V \cup V \times U \\[2em] \dfrac{(c_1)^2}{\binom{m}{2}(c_2)^2+(mn)(c_1c_2)+\binom{n}{2}(c_1)^2} \\ \qquad\qquad \textit{if } \{r,s\} \in \binom{V}{2} \end{cases}$$

(7.17)

$$p(r,s) = \begin{cases} \binom{m}{2}^{-1} & \textit{if } \{r,s\} \in \binom{U}{2} \\ 0 & \textit{otherwise} \end{cases}$$

(7.18)

$$p(r,s) = \begin{cases} \binom{m}{2}^{-1} \dfrac{c_2^2(1-g_0)}{c_1^2(1-g_1)+c_2^2(1-g_0)} \\ \qquad\qquad \textit{if } \{r,s\} \in \binom{U}{2} \\[2em] \binom{n}{2}^{-1} \dfrac{c_1^2(1-g_1)}{c_1^2(1-g_1)+c_2^2(1-g_0)} \\ \qquad\qquad \textit{if } \{r,s\} \in \binom{V}{2} \\[2em] 0 \\ \textit{if } (r,s) \in U \times V \cup V \times U \end{cases}$$

(7.19)

Proof. Let p be some distribution on $\binom{R}{2}$. By Lemma 9, we can assume that $p(t,u) = p(x,y)$, if $t,u,x,y \in U$. Similarly, $p(v,w) = p(x,y)$, if $v,w,x,y \in V$. Again using Lemma 9, $p(u,v) = p(u,y) = p(x,y)$, if $u,x \in U$ and $v,y \in V$. This shows that all pairs intersecting both U and V have equal probability.

If $k \geq n+m$, the adversary can try to compromise all routers. Thus the best strategy is to only choose pairs from the trusted set U, as described in Equation 7.18. From now on, assume that $k < n+m$.

Let $K_j \subseteq R$ be of size k and have an intersection with U of size j. The value of j alone determines the probability of compromise for K_j, because it determines the number of pairs in $\binom{U}{2}$, $U \times V$, and $\binom{V}{2}$. As we have just shown, the exact pairs included do not matter because their probability is determined by their class. Let $p_1 = \sum_{\{t,u\} \in \binom{U}{2}} p(t,u)$, $p_2 = \sum_{(u,v) \in U \times V} p(u,v)$, and $p_3 = \sum_{\{v,w\} \in \binom{V}{2}} p(v,w)$. Then we can say

that

$$c(p, K_j) = \binom{j}{2}(c_1)^2 \frac{p_1}{\binom{m}{2}} + j(k-j)c_1 c_2 \frac{p_2}{mn} +$$
$$\binom{k-j}{2}(c_2)^2 \frac{p_3}{\binom{n}{2}}$$

(7.20)

To narrow the set of possible optimal assignments of p_1, p_2, and p_3, we will first consider the effect of varying p_2. The quantity we want to minimize is the maximum value of Equation 7.20. Equation 7.20 is a quadratic function of j. Assume that the second derivative is non-zero. If it is zero it is easy to show that the distribution p is the distribution described in Equation 7.17. Otherwise, we will show that we can improve the maximum by changing p_2. We can find the local extremum by taking the derivative of Equation 7.20 and setting it to zero. Solving for j gives

$$j^* = \frac{\begin{array}{c} n(n-1)p_1 c_1^2 - k(m-1)(n-1)p_2 c_1 c_2 + \\ (2k-1)m(m-1)p_3 c_2^2 \end{array}}{\begin{array}{c} 2(n(n-1)p_1 c_1^2 - (m-1)(n-1)p_2 c_1 c_2 + \\ m(m-1)p_3 c_2^2) \end{array}}.$$

(7.21)

Unfortunately, j^* must be integral to represent a worst-case subset, and therefore we cannot just substitute the expression in Equation 7.21 into Equation 7.20 and solve for the optimal value of p_2. There may in fact be two values of j that are maxima, and varying p_2 could possibly increase the value at one while decreasing the value at other. Therefore, while varying p_2, we simultaneously vary p_1 and p_3 to maintain the local extremum of Equation 7.20 at j^*. Then both possible maxima are changed in the same way.

By observing that $p_3 = 1 - p_1 - p_2$ in Equation 7.21 we can see that p_1 and p_2 are linearly related. Solve this for p_1 and call the expression p_1'. Now let $j' \in \mathbb{N}$, $0 \le j' \le k$, be any value that maximizes $c(p, K_{j'})$. j' is either an endpoint of $[0, k]$ or a closest integer to a local maximum. Substitute p_1' for p_1 in $c(p, K_{j'})$, and the result is a linear function of p_2. Therefore either increasing or decreasing p_2 does not increase $c(p, K_{j'})$. Suppose we move p_2 in the direction that decreases $c(p, K_{j'})$. Because we vary p_1' (and p_3) with p_2 in such a way as to maintain the extremum of the parabola at the same value j^*, j' is maintained as a maximum of $c(p, K_j)$ as long as the second derivative of $c(p, K_{j'})$ remains non-zero.

The process of changing p_2 stops when (i) the second derivative of $c(p, K_{j'})$ becomes zero, (ii) p_2 reaches zero, (iii) p_3 reaches zero, or (iv) p_1 reaches zero.

Case (i): In this case, all sets have the same value. This is only satisfied when the distribution is that of Equation 7.17.

Case (ii): In this case, all probability is in pairs of two trusted or two untrusted nodes. Therefore the maximizing value of j must be when it is as small as possible or as large as possible, that is, at $\max(0, k-n)$ or $\max(k, m)$. If the former case is strictly larger, we can reduce it by decreasing p_3 and increasing p_1. If the latter case is strictly larger, we can do the reverse. Therefore the value in these two cases must be equal. To find the probabilities p_1 and p_3 that satisfy this, let $p_3 = 1 - p_1$, $v_0 = \max(k-m, 0)$, and $v_1 = \max(k-n, 0)$. Then setting them equal and solving for p_1 yields the condition

$$p_1 = \frac{c_2^2 \left(1 - \frac{v_0(v_0-1)}{n(n-1)}\right)}{c_1^2 \left(1 - \frac{v_1(v_1-1)}{m(m-1)}\right) + c_2^2 \left(1 - \frac{v_0(v_0-1)}{n(n-1)}\right)}. \tag{7.22}$$

Equation 7.22 then gives us the probability for each pair in $\binom{U}{2}$ and $\binom{V}{2}$, and this is the same as the distribution in Equation 7.19.

Case (iii): In this case, $p_3 = 0$. Then if $p_2 = 0$ also, we put all probability in the trusted nodes, which is the distribution described in Equation 7.18.

Now suppose that $p_2 > 0$. We will consider moving probability between p_1, p_2, and p_3 to show that this case isn't possible. Let $p_2 = 1 - p_1$ in Equation 7.20 and call this $c_3(p, K_j)$. Then use this to consider trading off p_1 and p_2 to find the optimal assignment. As p_1 varies, the change in the value of the set K_j is

$$D_{p_1} c_3(p, K_j) = \frac{j c_1}{m} \left[\frac{(j-1)c_1}{m-1} - \frac{(k-j)c_2}{n}\right]. \tag{7.23}$$

Next, let $p_2 = 1 - p_1 - p_3$ in Equation 7.20 and call this $c_4(p, K_j)$. Moving p_2 to p_3 results in a change of

$$D_{p_3} c_4(p, K_j) = \frac{(k-j)c_2}{n} \left[\frac{(k-j-1)c_2}{n-1} - \frac{j c_1}{m}\right]. \tag{7.24}$$

Let $j^* \in \operatorname{argmax}_j c(p, K_j)$ be the largest integer that is a maximum of $c(p, K_j)$.

We observe that $D_j^2 c(p, K_j) \le 0$. If not, we would have $j^* = k$. Then Equation 7.23 shows that decreasing p_1 would decrease the value at j^*, and p_1 is non-zero so we could do this because, at $p_1 = 0$, $c(p, K_j)$ is largest at $j^* = \lceil k/2 \rceil \ne k$. Such a decrease would contradict the optimality of j^*.

Now, because $D_j^2 c(p, K_j) \le 0$, there may be some $\hat{j} \in \operatorname{argmax}_j c(p, K_j)$ such that $\hat{j} < j^*$. There are four cases to consider here: (1) $D_{p_1} c_3(p, K_{j^*})$, $D_{p_1} c_3(p, K_{\hat{j}}) \le 0$, (2) $D_{p_1} c_3(p, K_{j^*})$, $D_{p_1} c_3(p, K_{\hat{j}}) \ge 0$, (3)

$D_{p_1}c_3(p, K_{j^*}) \geq 0$ and $D_{p_1}c_3(p, K_{\hat{j}}) \leq 0$, and (4) $D_{p_1}c_3(p, K_{j^*}) \leq 0$ and $D_{p_1}c_3(p, K_{\hat{j}}) \geq 0$.

In case (1), we could decrease c at j^* and $\hat{j}$ by moving probability from p_2 to p_1. This would contradict the optimality of p.

For case (2), we use the fact that

$$0 \leq a < b \Rightarrow \frac{a-1}{b-1} < \frac{a}{b}. \tag{7.25}$$

Inequality 7.25 implies that if $D_{p_1}c_3(p, K_j) \geq 0$, then $D_{p_3}c_4(p, K_j) \leq 0$. Therefore we could decrease c at j^* and $\hat{j}$ by moving probability from p_2 to p_3, contradicting the optimality of p.

For case (3), we show that we can still decrease both j^* and $\hat{j}$ while maintaining their equality, and hence maximality, by moving some probability from p_2 to p_1 and p_3. Moving probability from p_2 to p_1 increases the value at j^* and decreases the value at $\hat{j}$. This implies, by Inequality 7.25, that moving probability from p_2 to p_3 decreases the value at j^*. Furthermore, can assume that it increases it at $\hat{j}$ because otherwise we could decrease both j^* and $\hat{j}$ by moving probability directly from p_2 to p_3.

For j^* and $\hat{j}$ to be integral maxima of Equation 7.20, it must be that $j^* - 1 = \hat{j}$. Also, solving $D_{p_1}c_3 = D_{p_3}c_4$ for j, we find that at this point, $D_{p_1}c_3 \leq 0$ and $D_{p_3}c_4 \leq 0$. Therefore, j^* is at most one more than this point. We can observe by calculation that within this range the ratio $|D_{p_1}c_3/D_{p_3}c_4|$ is less than one. Similarly, $\hat{j}$ is at most one less than this point, and within this range the ratio $|D_{p_1}c_3/D_{p_3}c_4|$ is greater than one.

This shows that we can move probability from p_2 to p_1 and p_3 at rates that decrease the value at both j^* and $\hat{j}$. Because they were maximum, we have lowered the value of the worst-case subset K_{j^*}, contradicting the optimality of p.

Case (4) is not possible because $D_j^2[D_{p_1}c_2] \geq 0$ and $D_{p_1}c_3(p, K_0) = 0$.

Case (iv): In this case, if $p_2 > 0$, the case is symmetric to the case of $p_1, p_2 > 0$ and we can apply the same argument. Therefore assume that $p_2 = 0$, which implies that $p_3 = 1$. It must be that $m < n$ because otherwise we could set $p_1 = 1$ and $p_3 = 0$ and improve the worst case. But now consider moving some probability from p_3 to p_1. Let $p_1 = 1 - p_3$ in Equation 7.20 and call this c_3. The change in the worst-case subset, K_n, is

$$D_{p_3}c_3(p, K_n) = c_2^2 - \frac{(k-n)(k-n-1)c_1^2}{m(m-1)}.$$

This must be greater than zero because $c_2 \geq c_1$ and $k - n < m$. Therefore decreasing p_3 decreases $c(p, K_n)$, contradicting the optimality of p. □

7.6   Choosing a distribution

We have shown that there are three possibilities for an optimal strategy in choosing nodes that will minimize the best chances a fixed size adversary has to compromise both endpoints of an onion-routing circuit when a trusted set is available. To choose a distribution, a user can simply calculate the probability of compromise in each case and use the distribution with the smallest result. The optimal distribution depends on all the variables in the system: the trust values, the size of the trusted set, the size of the untrusted set, and the size of the adversary.

In the first distribution, described in Equation 7.17, the user chooses pairs $\{i, j\}$ to make $p(i, j)c_i c_j$ equal for all i, j. This is a random choice of pairs weighted by the trust in the pair. The probability of compromise under this strategy is

$$C_1 = \frac{k(k-1)c_1^2 c_2^2}{m(m-1)c_2^2 + 2mnc_1 c_2 + n(n-1)c_1^2}. \tag{7.26}$$

This strategy is optimal when the network is large compared to the adversary, and so it benefits the user to spread out his distribution, even to less-trusted routers. It is also optimal when the trust values are close.

In the second distribution, described in Eq 7.18, the user randomly selects pairs from within the trusted set. This can only be optimal if the size k of the adversary is larger than the size m of the trusted set. Otherwise, the user could decrease the probability of compromise by putting some of the pair-selection distribution on pairs outside the trusted set. Doing so would not change the adversary's worst-case subset, which is entirely in the trusted set, but it would decrease the probability that those nodes are chosen by the user. The probability of compromise, assuming $k > m$, is simply

$$C_2 = c_1^2. \tag{7.27}$$

We can compare this to Equation 7.26 and observe that c_1 can always be made small enough to make this value less than the value of the first strategy. These equations also show that choosing only trusted nodes will be optimal when k is large relative to the network. When $k = m + n$, this case is always optimal.

The third distribution, given in Equation 7.19, is perhaps the least obvious one, and arises as a result of the fact that users choose their distribution over pairs, while the adversary attacks individual routers. This distribution is only optimal when $k > m$ or $k > n$. Otherwise, under this distribution, there is a worst-case set $K_1 \subseteq U$ and a worst-case set $K_2 \subseteq V$. Then the probability of compromise can be reduced by moving

probability from pairs in U and pairs in V to the cross-pairs between U and V while maintaining K_1 and K_2 as worst-case sets. Therefore assume that $k > m$ or $k > n$. Let $v_0 = \max(k - m, 0)$ and $v_1 = \max(k - n, 0)$. Then let $g_0 = v_0(v_0 - 1)/(n(n - 1))$ and $g_1 = v_1(v_1 - 1)/(m(m - 1))$. The probability of compromise under this distribution is

$$
C_3 \;=\;
\begin{cases}
\dfrac{c_1^2 c_2^2 (1 - g_0)}{c_1^2(1 - g_1) + c_2^2(1 - g_0)} + \\[2ex]
\dfrac{v_0(v_0 - 1)c_1^2 c_2^2 (1 - g_1)}{n(n - 1)(c_1^2(1 - g_1) + c_2^2(1 - g_0))} & \text{if } k > m \\[3ex]
\dfrac{v_1(v_1 - 1)c_1^2 c_2^2 (1 - g_0)}{m(m - 1)(c_1^2(1 - g_1) + c_2^2(1 - g_0))} + \\[2ex]
\dfrac{c_1^2 c_2^2 (1 - g_1)}{c_1^2(1 - g_1) + c_2^2(1 - g_0)} & \text{if } k > n
\end{cases}
\tag{7.28}
$$

To make some sense of this, it is helpful to consider some special cases. When $n > k$, $m < k$, a worst-case set is contained in V, and so the probability of compromise is

$$
C_3 = \frac{k(k - 1)c_1^2 c_2^2}{n(n - 1)(c_1^2 + c_2^2(1 - g_0))}
$$

We can see that there is some large n such that C_3 is less than C_2 and C_1. What happens in this case is that there is a large number of routers, and the user wants to spread his probability among them. However, because $k > m$, spreading the probability to all cross-pairs (pairs including one trusted and one untrusted router) means that an adversary selecting as many trusted routers as possible gains $(k - m)m/(mn) = (k - m)/n$ of the probability on such pairs. On the other hand, when spreading probability to untrusted pairs, $(k - m)(k - m - 1)/(n(n - 1))$ of the shifted probability is captured by the adversary. Therefore, moving probability from cross-pairs to trusted pairs and untrusted pairs reduces the amount captured by the adversary by a factor of $O(1/n)$.

Chapter 8

A Protocol for Highly-Anonymous Low-Latency Communication

8.1   Introduction

In this chapter, we present a low-latency protocol that provides arbitrarily good anonymity. The barrier to this in onion routing is the ability of the adversary to determine if two circuit locations that it observes belong to the same circuit. Then, if the adversary controls the first and last routers, he can link the user with his destination.

In the onion-routing protocol we give in Chapter 4, the adversary can link circuit locations by sending a specific number of onions on a circuit after it has been constructed. If the same number of onions is observed at another point, then the two locations are part of the same circuit. We can prevent this attack, and indeed prevent any attack that involves changing or adding onions on the circuit, by using authentication codes to verify the validity of onions. This is included in, for example, the Tor protocol [23]. However, the ability to link locations remains, as the adversary can delay a message on a circuit, allow the rest of the messages to be delivered, and then release the message. A router that receives protocol messages after that release must be part of the same circuit. Our protocol is designed to make it very unlikely that the adversary is in a position to control the timing of onion delivery in this way.

In order to capture the problem we add a notion of time to our model. This gives agents the power to reason about the ordering of events in the system. Then they can react to their observations to make

themselves act more like the other agents and thereby confuse the adversary. Also, our protocol will involve intentional delays, and explicit timing allows us to estimate the impact on latency.

The key ingredients of our solution are timestamping onions with their intended send time and sending multiple copies of the same onion over redundant paths. These ideas have appeared before in the literature (cf. [49, 24] for timing techniques and [81, 46] for redundancy). The essential aspects of the protocol include:

1. A mesh topology from the user to the destination that balances limiting view of the user's message to a small number of routers while providing redundancy against malicious delays

2. A predefined return time from the destination to the user

3. A method to select free routes in the mesh topology

There are several choices in our protocol that are not optimized, and it seems quite possible that substantial improvements to anonymity and efficiency may yet be gained - for example, in the forward topology and the use of cryptography. These were left open because the problem we investigated had a wide space of possible outcomes. Could we achieve arbitrarily good security with reasonable performance? Was using free routes possible once you leave the simple path topology? As a result, our aim was to show good *asymptotic* peformance in our design criteria. The theoretical results, however, suggest that the approach is indeed viable and that a promising next step is to optimize within the framework of the scheme we describe.

We are still, however, concerned with the eventual practicality of our protocol. Therefore we extend the basic protocol to handle two-way communication over persistent connections. These are the communication semantics of the TCP protocol and the kind of communication supported by Tor. Sending multiple onions in each direction can be handled directly by our protocol. However, we need to make sure that the pattern of traffic does not identify the users. To prevent this, we make black-box use of a *padding scheme* to prevent passive timing attacks. Padding schemes specify how to delay and add dummy traffic to an input stream to make different streams looks similar. Several padding schemes that defeat passive timing attacks have been proposed [80, 34, 77, 87], and furthermore we believe that there is still potential for substantial improvement.

To evaluate the likely performance of our protocol in practice, we measure the feature of the underlying network that affects the latency added by the protocol - delay variations in the host network. Specifically, we measured the likely latency costs of running our protocol on the existing Tor network. The Tor network consists of routers of widely different capacities located around the world and therefore is likely to exhibit a high variability in performance. This provides a strenuous, real-world scenario in which to evaluate our

protocol's performance. We examine this by taking measurements of the inter-host network latencies and the per-host processing delays of the Tor routers over the course of a month.

Our results are therefore both theoretical and experimental:

1. We show that the user can be made anonymous with arbitrarily high probability as long as b is less than $1/2$. The user is anonymous within the set of users with identical traffic patterns as produced by the input padding scheme.

2. We prove that our mesh topology in the forward direction is optimal for anonymity in the limit.

3. The latency of our protocol is proportional to the length of the mesh and path. We show that the probability of compromise decreases to zero polynomially in that length. This compares well with onion routing, which adds latency proportional to its path length.

4. The bandwidth used is $2w + (l-1)w^2 + 1$, where l is the mesh/path length, and $w = \Theta(\log l)$. This compares well asymptotically to the $l + 2$ copies that are sent in an onion-routing path of the same total length.

5. For most of our measurements, we observe that added onion delay would need to be less than a factor of two to achieve desirable reliability.

The results suggest that our approach indeed has the potential to mitigate active timing attacks. However, the case is not fully made. A truly suitable padding scheme, with low overhead and large *anonymity sets* (groups of users within which the adversary cannot distinguish users), does not currently exist. There still exists many possible improvements to the construction of our scheme, such as changing the mesh/path length or using persistant circuits and private keys. Moreover, most of the delays measured appear to be caused by congestion at the router. Tor was designed for optimal throughput rather than predictable and reliable latency, and therefore we expect that a usable network could be built with even better performance than our measurements suggest.

8.2   Model

We describe and analyze our protocol using the network, user, and adversary models of Chapter 2. However, we need to incorporate a more detailed model of timing and network delays. Agents need to be able to

determine the global timing of events to enforce timing patterns. Also, network and processing delays need to be somewhat predictable so that agents can take them into account.

8.2.1 Network

Each event in the network occurs at some global time. We assume that each user and router has a local clock that accurately measures time, and that these clocks have run some sort of synchronization protocol [59]. Let δ_{sync} be the largest difference between two synchronized clocks in the network.

There is some probabilistic network delay $d_{net}(r, s)$ between every pair (r, s) of routers. This models the unpredictable delay between routers due to factors such as congestion and route instability. There is also some probabilistic processing delay $d_{proc}(r)$ at every router r. This reflects changes in delay at a router due to timesharing with other processes. The delay of a message is the sum of delays drawn independently from these two distributions. We also let the delay of one message be independent of the delays of the other messages. We assume the distributions of both sources of delay is known to the system. In practice, this may be achievable via network measurements.

8.2.2 Adversary

It seems plausible that an adversary can run routers that are at least as fast as the other routers on the network, and that it may dedicate them to running the protocol. Therefore, in contrast to non-adversarial routers, we pessimistically assume that the adversary has complete control over the processing delay $d_{proc}(a)$ of routers $a \in A$.

8.3 A Time-stamping Solution

The adversary can take advantage of time in our model to deanonymize users in the onion-routing protocol when he controls the first and last routers. To do so, the adversary delays the message from any given user until the messages from all other users are highly likely to have been delivered. Then he forwards the delayed message and observes which destination receives a message afterwards. In general, the adversary can affect the arrival time of onions by delaying when he forwards them. To prevent this, for the message from the user we include the time every onion should be forwarded. Also, we use redundant routers to send multiple copies of a given onion. If at least one copy shows up in time, the onion is successfully forwarded. A reply sent

from the destination to the user must, by our requirements, pass initially through a single router. Because that last router may be compromised, we cannot trust it to send the onion at a specified time. However, the reply is destined for an anonymity protocol participant, the user; therefore, unlike traffic from the user, we can send a dummy reply if a real one doesn't arrive in time. Observe that this does not work for the message from the user, because the final router sees which onions actually contain a message.

8.3.1 From the user

First, consider what we could do if propagation and processing delays were deterministic. The user could send through a simple path in the network an onion that in the ith layer includes the latest time that the onion should arrive at the ith router. This would indicate the time that that router should process and forward the onion. Any delay must be due to the adversary.

Using this scheme, the adversary could delay onions to create a gap that propagates down the path. Our approach to this problem is to send multiple copies of an onion down redundant, intersecting paths. A router on a path needs only one copy of the onion to arrive in time from any incoming path in order to forward it by its send time.

This approach has limits, because each redundant router adds another chance for the adversary to observe an entire path from source to destination. For example, suppose that we simply send onions over k paths of length l that intersect at a final router, where every router is chosen uniformly at random. Let b be the fraction of routers that are controlled by the adversary. The probability that at least one path is entirely composed of compromised routers is $b(1 - (1 - b^{l-1})^k)$. This quickly goes to b as k increases. We use a layered-mesh topology to balance the ability of the adversary to passively observe a path with his ability to actively perform a timing attack and provide good total anonymity.

Unfortunately, in real networks, delays are somewhat unpredictable. For example, an onion might be delayed by congestion in the underlying network. However, if the distribution of delays is known, we know how long we need to wait at a router for onions to arrive with any fixed probability. We will set that probability to balance decreasing added latency with decreasing the chance that the onion arrives later than planned for, which creates the opportunity for a successful timing attack. Then we add in this buffer time to the send time.
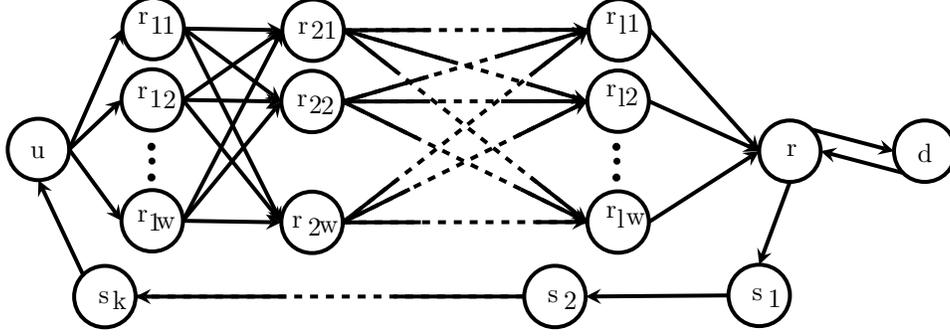
Figure 8.1: Layered-mesh topology

Topology

The layered-mesh topology we propose is pictured in Figure 8.1. For some length l and width w, user u sends a copy of the onion to the w members r_{1j} of the first layer. Then, in layer i, each router r_{ij} sends one copy of every onion it receives to each router $r_{(i+1)k}$ of the next layer. Finally, the routers r_{lj} in last layer send a copy of each onion received to a single router r, that finishes decrypting them and sends the data on to the destination. We call this structure the *layered mesh*.

Timestamps

As described, the user sets timestamps to instruct routers to maintain a specific timing of events. A user may have different latencies to the different routers in the first layer. If the time that one of these routers is instructed to forward the onion only depended on the network and processing delays of that router, the first-layer routers could send copies of the same onion at different times. This provides information to the next layer about the identity of the user. Similarly, the adversary could use these different times to link other layers in the mesh. Therefore, we set the send time for each layer to be the send time of the previous layer plus the longest delay at our chosen reliability level p. We must also add some extra delay to allow for clock skews.

Let $d^*(r, s)$ be the amount of delay we need to ensure that an onion from r is processed at s with success probability p:

$$p = Pr[d_{net}(r, s) + d_{proc}(s) \leq d^*(r, s)].$$

At time t, let user u send a message. The user chooses the same send time for all routers in the same

layer. The send time for routers r_{1j} in the first layer is

$$t_1 = t + \max_j d^*(u, r_{1j}).$$

The send time for routers r_{ij} in the ith layer is

$$t_i = t_{i-1} + \max_{j,k} d^*(r_{(i-1)j}, r_{i,k}) + \delta_{sync}.$$

If a router receives its first copy of an onion after the send time has passed, it immediately forwards the onion to the routers in the next layer. At worst, the delay is the result of attempts by the adversary to delay certain onions. Sending the onion later or not at all in that case would only make it easier for the adversary to observe its lateness later in the mesh. Forwarding it immediately might even allow the onion to eventually get back on schedule. At best, the delay is just a result of network delays and forwarding has no effect on anonymity.

Onions

Let M be the message to be sent. We will encrypt the message with a public key shared by all members of a layer. Given that the layers are set up in advance and known to all, such a key can be generated by a trusted third party or by electing a leader to do it. Let $\{M\}_{r_i}$ denote the encryption of M with the public key of layer i. Let $n_{r_i}, n_{s_1} \in \mathbb{N}$ be random integers. Then the onion that u sends to the routers in layer 1 is

$$\{n_{r_1}, t_1, \{n_{r_2}, t_2, \cdots \{n_r, d, n_{s_1}, k_r, M\}_r \cdots \}_{r_2}\}_{r_1}$$

For each layer i, a user generates the random number $n_{r_i} \in \mathbb{N}$ as an onion identifier. The routers keep track of the onion identifiers they have seen. When they receive an onion, they decrypt it and examine the identifier. Routers only forward an onion if its identifier n_i has not been seen before. n_{s_1} is the identifier that r should use with s_1 when sending back any reply, and k_r is a private key that will let r encrypt the return message for u.

The onion encoding and forwarding scheme should provide anonymity by preventing attacks such as determining the routing information in the encrypted onion, preventing forgery, and prohibiting replay attacks. For clarity of presentation, we have described a simple scheme that achieves this. We observe,

however, that several improvements to the protocol could be made. For example, the protocol could do a more explicit connection open and close to reduce the lists of identifiers that routers have to maintain. Also, symmetric keys could be exchanged to speed up onion processing. Another improvement that we could incorporate is forward secrecy. There are numerous cryptographic details that must be carefully set out (for example, as in [10]) for our protocol to have a cryptographically secure and efficient implementation. These are not our focus here.

8.3.2   To the user

A message returning to the user from the destination must first pass through the one router that is selected to communicate directly with the destination. This router may be compromised, and it may try to affect the timing of in the return message. We manage this by giving the intermediate routers a send time for the return message. They enforce it by delaying the return onion until that time, and by creating a dummy onion at that point if no return onion has arrived. We note that this doesn't work for the message from the user because a dummy packet can be recognized by the final router as part of our functionality requirement. We choose a simple path of length k for the return onion (Fig. 8.1), because there is no anonymity advantage to adding redundancy here. We call this structure the *return path*. Note that specifying a return time beforehand violates our functionality requirement in that it only delivers return messages that arrive before the send time. We therefore must slightly adjust our functionality and assume that replies from the destination are always sent within some bounded time. Then as long as the user chooses a longer return time, the protocol will implement this functionality.

To communicate the desired return time to the return path, we take advantage of the one-way communication protocol already developed. The user sends the return time via the layered mesh to every router in the return path. At the same time, the user generates the random numbers $n_{s_i} \in \mathbb{N}^+, 1 \leq i < k$, and sends two random numbers $n_{s_i}, n_{s_{i+1}}$ and a key k_{s_i} to each router s_i, again using the layered mesh,. These will be the incoming and outgoing identifiers for the onion. The user also sends n_k and u to s_k. Let M be the message to be sent back from d to u. The return onion sent to s_i is

$$O_i = < n_{s_i}, \{\cdots \{\{M\}_{k_r}\}_{k_{s_1}} \cdots \}_{k_{s_{i-1}}} > .$$

After r receives M from d, it will take n_{s_1} and k_r out of the inbound onion from the user and send O_1 to

s_1. When s_i receives O_i it looks for a matching n_{s_i} in the pairs of number it has received and then forms O_{i+1} to send when the previously-arranged time arrives.

8.3.3 Choosing the routes

A simple method to select the layered mesh and return path would be for it to be chosen by the network and shared among all users. This is analogous to cascades in mix networks [6]. A disadvantage of cascades is that number of users that can be confused with one another, that is, the size of the anonymity set, is at most the number of users that can simultaneously be handled by a router.

Therefore, we allow users to select their own routes, known as free routes. This can provide a much bigger anonymity set. A relevant comparison of the approaches in the context of mix networks is given in [24]. This procedure must be designed carefully because users are no longer guaranteed that anybody else chooses the same routers they do. We let the set of nodes be divided into groups of size w that form the layers described above. These layers are further stratified by the position they can serve in a mesh, for example, no second-layer group could serve in any other position in a route chosen by anyone else.

When the user first selects a layer that is completely uncompromised, his onion becomes indistinguishable from the other users that select that layer. His onion can also be confused with routes that have consecutive uncompromised layers covering the same mesh position because the adversary cannot tell what was selected in that position. This continues until the end of the mesh or until the user selects a completely compromised layer (after which the adversary can drop onions to create patterns that are detectable at the end). As the user population grows, the expected number of users sharing a completely uncompromised layer reaches the capacity of those routers. Therefore the expected size of the anonymity set grows with n and also depends on b.

Though outside of our threat model, one can consider an attacker that attempts to compromise specific nodes over an extended period in an attempt to increase the chances of owning an entire layer and an entire path from a client to that layer. Against such an attacker one would want to have a dynamic network structure in which the positions that layers can be in change periodically as does the composition of each layer group. These may or may not change at the same rate. There are numerous tradeoffs to be made in deciding how best to do this. It is beyond the scope of this work to do more than note the existence of such responses to a long-term strategically roaming attacker.

8.4 Anonymity

The purpose of our protocol is to provide better anonymity than onion routing at reasonable cost in latency and bandwidth. A major drawback to onion routing is that the probability of compromise is b^2 and cannot be improved by making tradeoffs elsewhere in the protocol, for example by choosing a longer path. We will show how in fact our scheme can provide arbitrarily good probability by increasing the length and width of the layered mesh.

First, the design of the protocol onions essentially limits the adversary to observing directly and linking the user and his destination. For the message from the user, the use of encryption and unique identifiers forces it to be passed through the required layers and limits it to being forwarded by an honest router at most once. It also hides the source and destination from all but the first and last routers, and it makes messages leaving one layer unlinkable to messages leaving another layer. For a reply to the user, the source and destination are not included in the onions, and encryption prevents an onion leaving one router from being linked with a message leaving another router.

For the adversary to break a user's anonymity, then, he will need to either observe the entire path between source to destination or link onions at different steps on that path. The latter depends on the adversary's ability to delay the onion. To evaluate this possibility, we will make the simplifying assumption that the network and processing delay between two routers r, s never falls above the time allowed for it $d^*(r, s)$. In our model, such failures can happen with probability $1 - p$, where p can be set arbitrarily close to 1. Certainly, if the adversary can deanonymize a user even under this assumption, he can do so with possible link failures. When such failures do occur, they open up the chance for an adversary to successfully delay onions. However, late onions are typically irrelevant because the same onion will have been forwarded by another router in the previous layer. Also, late onions that are the first of their type to arrive are immediately forwarded, thus benign delays are somewhat self-correcting, and we estimate that they do not open up a large opportunity for an active timing attack. However, if we wish to make a conservative estimate, we can expect that for any given onion a fraction p of the onion copies will fail to arrive in time. We can simply add p to the fraction of routers controlled by the adversary to obtain an effective fraction $b^* = b + p$ of compromised routers that can be used to determine the probability that the onion is successfully delayed or dropped.

Assuming no link failures, then, the anonymity of a user only depends on which routers the adversary controls. Among users starting a connection at the same time as u, that have free routers overlapping u's route such that the adversary cannot distinguish them, and that are not deanonymized by the adversary,

either u is anonymous or he can be fully identified via a delaying attack. Because the routers in the topology are selected randomly, there is some probability that the adversary can deanonymize a user. Let $\mathcal{C}$ be the event that the adversary can compromise anonymity. We can quantify the probability of $\mathcal{C}$.

Theorem 17.

$$Pr[\mathcal{C}] = b^{k+1} + b(1 - b^k)\left[b^w\frac{1 - (1 - (1 - b)^w - b^w)^l}{b^w + (1 - b)^w} + (1 - (1 - b)^w - b^w)^l\right] \tag{8.1}$$

Proof. To compromise anonymity, the adversary must control the router r that communicates with the destination. Therefore assume that $r \in A$. This happens with probability b.

Then he can link the destination via the return path only when he controls every router on it. This happens with probability b^{k+1}.

The adversary can deanonymize the user via the mesh in two ways. First, the adversary can control an entire path from the first layer in the mesh to r, but not any entire layer. This happens with probability $(1 - (1 - b)^w - b^w)^l$. Second, if the adversary controls a path to a completely controlled layer i, he can then use it to selectively delay onions from one user. He can identify the user to target because he controls a path from that user to layer i. This happens with probability $b^w(1 - (1 - b)^w - b^w)^{i-1}$. If neither of these is the case, then it must be that there is an honest path in the mesh leading to a completely honest layer. This layer receives all onions from all users via the honest path, and beyond it the adversary cannot distinguish between users. Therefore deanonymization is not possible.

Adding all of the probabilities, we get

$$Pr[\mathcal{C}] = b^{k+1} + b(1 - b^k)\left[b^w\sum_{i=0}^{l-1}(1 - (1 - b)^w - b^w)^i + (1 - (1 - b)^w - b^w)^l\right]$$

Simplifying the partial geometric sum gives the theorem. □

Theorem 17 shows that, when less than half of the routers are compromised, we can make the probability that a user is anonymous arbitrarily high by setting w, l, and k large enough.

Corollary 4.

$$\lim_{w,l,k\to\infty} Pr[\mathcal{C}] = \begin{cases} 0 & b \leq 1/2 \\ 1/4 & b = 1/2 \\ b & b > 1/2 \end{cases}$$

Proof. As $k \to \infty$, the first term in Equation 8.1 goes to zero. As $k, l \to \infty$, the second term goes to $b^{w+1}/(b^w + (1-b)^w)$. When $b < 1/2$, this goes to zero, when $b = 1/2$, this goes to $1/4$, and when $b > 1/2$, this goes to b. $\qquad\square$

Corollary 4 shows that anonymity can be made arbitrarily good when $b < 1/2$, but that it is worse than onion routing when $b \geq 1/2$. Therefore assume from now on that $b < 1/2$. We would like to determine how big the layered mesh and return path grow as we increase our desired level of anonymity. First, we will consider how wide the mesh must be for a given depth to achieve optimal anonymity. This affects the number of messages that need to be sent. Also, it will allow us to evaluate anonymity as a function of the lengths k and l, quantities which we would like to keep small to provide good latency. Luckily, it turns out that the optimal width w^* grows slowly as a function of l.

Theorem 18. $w^* = O(\log l)$

Proof. Let $\mathcal{A} = \mathcal{C}^c$ be the event that users are anonymous. From Thm. 17, we can determine that

$$Pr[\mathcal{A}] = 1 - b + b(1-b)^w \frac{1 - (1 - b^w - (1-b)^w)^l}{b^w + (1-b)^w}. \tag{8.2}$$

Let x be the third term in Equation 8.2. It is the only term that varies with w. Because $b < 1/2$, we can replace b with $1 - b$ at point in x to obtain a simpler upper bound:

$$x \leq b(1 - (1 - (1-b)^w)^l). \tag{8.3}$$

Similarly, we can replace $1 - b$ with b to obtain a lower bound:

$$x \geq b(1 - (1 - (1-b)^w)^l)/2. \tag{8.4}$$

Equation 8.3 is decreasing in w. Equation 8.4 is increasing for small w, achieves a maximum at $w = \log(1/2)/\log p$, and is decreasing for larger w. At its maximum, Equation 8.4 has value $b/2$. Therefore $Pr[A]$ must be maximized at a smaller w than that at which the upper bound in Equation 8.3 reaches $b/2$. We can

further approximate to find this w:

$$x \leq b(1 - (1 - (1 - b)^w)^l) \tag{8.5}$$

$$\leq bl(1 - b)^w \tag{8.6}$$

Expression 8.6 is $b/2$ when w is $\frac{\log(2l/b)}{\log(1/(1-b))}$. $\qquad\square$

Network users are particularly sensitive to latency, and each additional step in the layered mesh and return path represents a potentially global hop on our network. To keep the lengths l and k of the mesh and return path small, then, we would like for $Pr[\mathcal{C}]$ to converge quickly to its limit. As suggested by Thm. 18, let $w = \log l$, and consider the convergence of $Pr[\mathcal{C}]$. It is clear that the first term shrinks exponentially with k. The convergence time of the second term is polynomial, although it does improve as b gets smaller.

Theorem 19. *Let* $c_1 = \log b$ *and* $c_2 = \log(1 - b)$. *Then*

$$Pr[\mathcal{C}] = \Theta(l^{c_1 - c_2}).$$

Proof. Let us fix k and only consider how $Pr[\mathcal{C}]$ changes with increasing l. Let x be the part of Equation 8.1 that varies with l and substitute in $w = \log l$:

$$x = bl^{c_1} \frac{1 - (1 - l^{c_1} - l^{c_2})^l}{l^{c_1} + l^{c_2}} + b(1 - l^{c-1} - l^{c_2})^l.$$

The first term in x is less than l^{-c_2}. We can give an upper bound for the second term by using the inequality $1 - y \leq e^{-y}$:

$$b(1 - l^{c-1} - l^{c_2})^l \leq b(1 - l^{c_2})^l$$

$$\leq be^{-l^{1+c_2}}$$

$1 + c_2 > 0$ because $b < 1/2$. Therefore $x \leq b(l^{c_1 - c_2}) + e^{-l^{1+c_2}} = O(l^{c_1 - c_2})$.

We can give a lower bound for the first term in x by using the fact that $1 - y \geq e^{-2y}$ for $y \leq 1/2$:

$$bl^{c_1} \frac{1 - (1 - l^{c_1} - l^{c_2})^l}{l^{c_1} + l^{c_2}} \geq bl^{c_1} \frac{1 - e^{-2(l^{c_1} + l^{c_2})l}}{2l^{c_2}}$$

$$\geq \frac{b}{2}(1 - e^{-2l^{1+c_2}})l^{c_1 - c_2}$$

Therefore $x = \Omega(l^{c_1 - c_2})$. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad$ $\square$

Our analysis shows how our scheme can provide arbitrarily good probability for $b < 1/2$. Is it possible to improve this to include values of b greater than $1/2$? Also, the convergence as the topology grows is not as fast as we could hope. The return path converges exponentially to perfect anonymity, but the layered mesh does not, and so far we have not justified the specific choice of a mesh. Is it possible that we could design a topology that has improved convergence time?

First, we observe that some other plausible topologies do not perform as well. For example, suppose the user sends onions to k routers, each of which forwards it to the same next router, and then from then on there is a path to the destination. The probability that anonymity is compromised in this situation is $b(b(1 - (1 - b)^k + (1 - b)b^k))$. While for small values of k and b this is a small improvement over the simple path in onion routing, as k grows the anonymity goes to zero.

As another example, consider using a binary tree. The user sends to the leaves of the tree, each child forwards the onion to its parent, and the parent forwards one copy of the onions it receives. Set p so that the probability of a sent onion arriving late is arbitrarily small. There are three events at a router r that are relevant to onion loss propagating up the tree:

$\qquad\qquad C_r \qquad$ r is compromised.

$\qquad\qquad \mathcal{D}_r \qquad$ Onions from all users can be stopped from coming through r.

$\qquad\qquad \mathcal{E}_r \qquad$ Onions from a known user can be stopped from coming through r.

Let $c_1(r)$ and $c_2(r)$ be the children of r. These events at r can be described in terms of in the same events at its children.

$$\mathcal{D}_r = C_r \vee (\mathcal{D}_{c_1(r)} \wedge \mathcal{D}_{c_2(r)})$$

$$\mathcal{E}_r = C_r \wedge (\mathcal{E}_{c_1(r)} \vee \mathcal{E}_{c_2(r)}) \vee (\mathcal{E}_{c_1(r)} \wedge \mathcal{D}_{c_2(r)}) \vee (\mathcal{D}_{c_1(r)} \wedge \mathcal{E}_{c_2(r)})$$

We use these descriptions to give recurrence relations for the probabilities of the events. Let r be at

distance i from a leaf. Let D_i be the probability of $\mathcal{D}_r$.

$$
\begin{aligned}
D_0 &= b \\
D_i &= b + (1 - b)D_{i-1}^2
\end{aligned}
$$

Let E_i be the probability of $\mathcal{E}_r$.

$$
\begin{aligned}
E_0 &= b \\
E_i &= bE_{i-1}(2 - E_{i-1}) + (1 - b)E_{i-1}(2D_{i-1} - E_{i-1}) \\
&= E_{i-1}(2b + 2(1 - b)D_{i-1} - E_{i-1})
\end{aligned}
$$

The probability that the adversary can block all onions going through a router is always at least b, and increasing its height in the tree increases the probability up to a limit.

Theorem 20.

$$
\lim_{i \to \infty} D_i = \begin{cases} \frac{b}{1-b} & b \le \frac{1}{2} \\ 1 & b \ge \frac{1}{2} \end{cases}
$$

Proof. Let D^* be the fixpoint of D_i.

$$
\begin{aligned}
D^* &= b + (1 - b)(D^*)^2 \\
D^* &= \frac{1 - |1 - 2b|}{2(1 - b)}
\end{aligned}
$$

It is clear from the definition that D_i is strictly increasing for $b > 0$. Therefore D^* is achieved in the limit. $\square$

We can use this result to determine the limit of E_i.

Theorem 21.

$$
\lim_{i \to \infty} E_i = \begin{cases} 0 & b \le \frac{1}{4} \\ 4b - 1 & \frac{1}{4} \le b \le \frac{1}{2} \\ 1 & b \ge \frac{1}{2} \end{cases}
$$

Proof. Replace D_i with D^* in the definition of E_i, and call it E_i'. We can compute the fixpoints for E_i':

$$
E'^* = E'^*(2b + 2(1 - b)D^* - E'^*)
$$

124

$$E'^* = 2b - |1 - 2b| \vee 0$$

We will treat these as the fixpoints for E_i, although more careful arguments need to be made to show that we can treat D_i as its limit in this way.

Observe that when $b \le 1/4$, E_i is strictly decreasing.

$$2b + 2(1 - b)D_i \le 2b + 2(1 - b)D^* \le 1$$

Therefore, in the limit, E_i reaches the higher fixpoint for E'^*, which in this case is $E^* = 0$.

Now consider the case that $1/4 \le b \le 1/2$. When $E_i > 2b + 2(1 - b)D^* - 1$, E_i decreases with i. When $E_i < 2b + 2(1 - b)D_{i-1} - 1$, E_i is increasing. Careful limiting arguments show that as $D_i \to D^*$, E_i also approaches a limit. Because E_i increases when it is small, this limit must be the fixpoint $2b - |1 - 2b| = 4b - 1$.

Finally, let $b \ge 1/2$. E_i is multiplied each step by $2b + 2(1 - b)F_i - E_i$, which, in this case, approaches $2 - E_i$. Therefore, for large i, E_i is increasing, unless E_i is very close to 1. However, even then E_i cannot decrease too much. Choosing a convergent sequence of F_i yields a sequence of E_i that approaches 1. $\qquad\square$

For the adversary to link a user with his destination, the final router r_f must be compromised. Also, a signal must propagate either up the send tree or down the return path. The probability of at least one of these compromises occuring is at most $b(E_h + 1 - b^l)$. Therefore, by Theorem 21, we can push it arbitrarily low when $b \le 1/4$ by increasing h and l.

As Theorem 21 shows, anonymity in the binary tree is not as good as it is in the layered mesh. The following theorem shows that the layered mesh is optimal in the limit among all topologies.

Theorem 22. *Let $c(b)$ be the probability of anonymity compromise when the fraction of adversarial routers is b. Then, if $b < 1/2$, $c(b) < \epsilon$ implies that $c(1 - b) > 1 - b - \frac{1-b}{b}\epsilon$.*

Proof. Compromise happens only when the adversary controls the last router and onions from certain users can be selectively blocked before then. Even when the adversary controls an entire path from the user to the destination, this can be viewed as a case in which he can selectively block onions. Cycles in the topology do not affect which onions reach the destination. Therefore we can assume that the topology is acyclic.

We assume without loss of generality that there are two users, u and v, and the adversary is trying to block an onion from u and while at the same time allowing an onion from v. To do this, it is always an optimal strategy for the adversary to forward only v's onion if it can determine the source and otherwise

block all onions. Therefore we assume the adversary implements this strategy.

Assume that the router that communicates with the destination is adversarial. Then, for a given topology, let α be a Boolean function on the remaining router positions that indicates if that position is controlled by the adversary or not. If anonymity is not compromised under α, there must be a set of honest routers C such that C forms a cut between the users and destination and every router in C sends onions from u and v. To see this consider the induced subgraph of routers that send the onion from v but not from u. It cannot connect u and the last router or anonymity would be compromised. Therefore, there must be a set of routers that are outgoing neighbors to this subgraph. These routers and the honest first-hop routers constitute the desired cut.

Every router in C must have a path of honest routers from the user. Otherwise there would be a set of adversarial routers that cut a router in C from the users, and no onions from the users would reach that router. Now consider the assignment $\overline{\alpha} = 1 - \alpha$. If α provides anonymity, then in $\overline{\alpha}$ there must be a cut consisting of adversarial routers such that each router has a path of adversarial routers from the users. Thus anonymity is compromised in $\overline{\alpha}$.

Let $c'(b)$ be the probability of anonymity compromise given that the final router is compromised. The probability of assignment α with a fraction b of compromised routers is equal to the probability of assignment $\overline{\alpha}$ with a fraction $1 - b$ of compromised routers. Because the complement of an assignment providing anonymity is an assignment the compromises anonymity, we have

$$1 - c'(b) \leq c'(1 - b). \tag{8.7}$$

$c(b) = bc'(b)$, and, therefore, if $c(b) < \epsilon$, then $c'(b) < \epsilon/b$. Combining this with Inequality 8.7, we get that $c'(1 - b) > 1 - \epsilon/b$, which implies that $c(1 - b) > 1 - b - \frac{1-b}{b}\epsilon..$ $\qquad \square$

Theorem 22 implies that if the probability of compromise for a topology goes to zero for $b \in [0, \beta]$, then it must go to b for $b \in [1 - \beta, 1]$. This is achieved by the layered-mesh topology for the largest range of b possible, $[0, 1/2)$.

8.5   Extended scheme

Our functionality up to this point has been very simple - we just require one message and one response to be delivered. This requirement is enough to evaluate current schemes and our proposals, and its simplicity keeps the analysis simple. However, we would ultimately like to come up with protocols that are useful in practice. Therefore, we extend the protocol so that it can be used in a real system.

In practice, users and destination often wish to send multiple messages back and forth. In particular, it is very useful to support two-way *stream communication*, in which users open bidirectional connections with destinations, and messages inserted into the connection at any time by a communication partner are delivered to the other partner. To extend our protocol to support multiple messages, we run multiple times the protocol for a single message and response, with two small changes.

The first change modifies how messages are sent from the users to the destination. Our model was that a group of users all want send a single message at the same time. To extend that assumption to multiple messages, the entire group must send the same number of messages and must send them at the same time. It seems quite unlikely that very many uncoordinated users will happen to synchronize for very long. Therefore, we use a padding scheme to fit the messages from a user into a pattern that is more likely to be shared by other users. The padding scheme $\mathcal{P}$ is a black box that, for every time step, takes the presence of data from the user and returns whether or not an onion should be sent. If an onion is to be sent and there is no message from the user, then a dummy onion with garbage message content is sent.

The second change modifies how replies are sent from the destination to the user. Our assumption has been that the destination sends at most one reply after receiving the message. Again, it is unlikely that very many destinations send all their responses messages simultaneously. Therefore, as before, we use a padding scheme for the return traffic. However, we cannot rely on the final router to follow this scheme because it may be adversarial. Therefore the scheme must be determined in advance and sent to every router on the return path. Any router on the return path can therefore fill in missing onions. Unlike traffic from the user, filling in delayed or dropped onions works for traffic from the destination because the dummy onions get sent all the way to the user and aren't revealed to the final router.

Let S_u be the set of users with traffic patterns that are the same and are received by the network at the same time as the traffic pattern of user u. Our proposed protocol relies on the effectiveness of the padding scheme. At best, it makes u indistinguishable within the set S_u supplied by $\mathcal{P}$.

We will simply assume the existence of a suitable padding scheme. Several padding schemes that could

be used in our context have already been proposed [80, 34, 77, 87], and future improvements may be used as a component of our protocol. One existing padding scheme that could be used is constant-rate padding, in which the length of the connection is chosen in advance and onions are sent at a constant rate in each direction for the duration of the connection. This typically causes high added latency and/or message overhead, however. The padding scheme of Shmatikov and Wang [77] could be adapted for use in our protocol by fixing the return scheme (perhaps by sampling in advance from the delay distribution). It is not hard to conceive of novel padding schemes that might satisfy these requirements, although getting a good mix of anonymity and performance does not seem easy.

8.6   Latency Measurements

Our protocol requires routers to hold messages until a specified send time. Latency between routers varies, and therefore this send time must be set high enough to guarantee with sufficiently high probability that it occurs after the onion actually arrives. The amount of time that onions spend being delayed before the send time depends on the variability of latency. Because network performance is critical to realizing effective anonymous communication [22], and we wish to evaluate the latency in our protocol.

In order to do so, we have measured latency in the Tor [23] network. Performance data in the Tor network gives us reasonable estimates for the performance of our protocol for several reasons. First, the essential router operation in both protocols is decrypting and forwarding onions. Second, the network is globally distributed and therefore includes a wide variety of network and host conditions. Third, the volunteer-run network is the most successful example of anonymous communication in practice, and it seems likely that any future system for anonymous communication would include a volunteer component.

We can therefore estimate the added delays in our protocol by simply using the network-delay and processing-delay distributions we observe in the Tor network. In addition, we are able to examine the data more closely and understand the underlying factors affecting these distributions, such as Internet route congestion, bandwidth at Tor routers, and process load at Tor routers.

8.6.1   Methodology

During each experiment, we measured the latency of Tor routers in three ways:

1. For a given router, we first opened five TCP connections in a row and measured the time between the

TCP SYN request and the SYNACK response. This provides an estimate of the network delay. We refer to this delay as the *round-trip time (RTT)*.

2. We then created a circuit from a host on the Yale network to the router and opened a TCP stream over that circuit from the router back to the Yale host. We measured the time between sending the stream-open request to the router and receiving the TCP connection from the router. This time includes host-processing delay from decrypting the onion and responding to the command. We refer to this delay as the *connection delay*.

3. Finally, we sent five 400-byte segments of data up the established TCP stream. The segments were sent every 500ms in separate Tor cells (Tor's 512-byte packets). We measured the time between sending a given segment and receiving all 400 bytes. This time includes host processing delay from decrypting the onion and forwarding the data. We refer to this delay as the *packet delay*.

We took measurements in the Tor network from March 13, 2009, to March 30, 2009. Every hour, we downloaded the latest router directories and used them to identify routers to test during that hour. Following the Tor-specification instruction for choosing routers, we only measured routers with the modifiers "Exit", "Fast", "Running", "Stable", and "Valid". The routers also had to be non-hibernating and could not have exit policies that disallowed connections to our host IP address and port. For every router we tested during the hour, we started an experiment after an exponentially-distributed random delay. After each experiment finished, we started another one after another exponentially-distributed random delay. There was an average number of five experiments per router per hour.

The Yale host was located at 192.31.236.5 and the server port we used was 80. The Tor client was a custom client written in Java. Packet traces were recorded using `tcpdump`. The timestamps on these were used to determine the timing of events during the measurement process.

8.6.2 Results

Over the course of the month, 745 routers appeared that met our selection criteria. The average number of experiments per router was 399, the median was 89, the minimum was 1, and the maximum was 1776. These numbers reflect the uptime of Tor routers. A histogram of the number of measurement experiments performed on the routers is given in Figure 8.2.

We estimate network delay by looking at the round-trip time of a TCP SYN/SYNACK pair. Round-trip
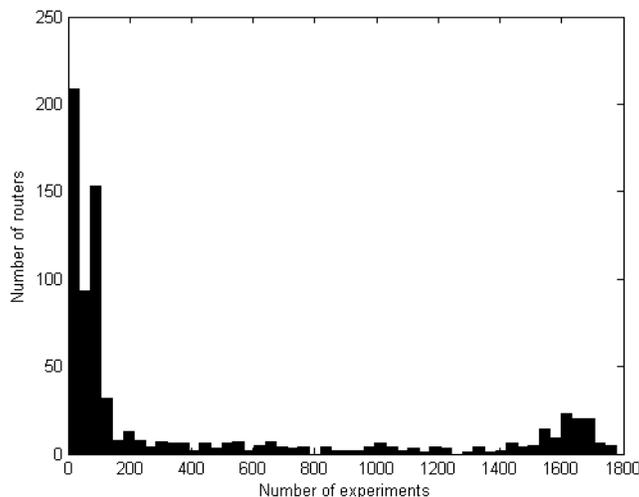
Figure 8.2: Histogram of experiments performed on routers

times for all experiments are shown in Figure 8.3. The mean of these times is 161ms and the median is 110ms. Similar to the connection delays, we see a peak bin centered at 100ms. These data are consistent with the measurements of Mukherjee [62], which, are shown to be well-modeled by a gamma distribution. In particular, these delays are unimodal.

A histogram of all the connection delays measured is shown in Figure 8.4(a). The top 5% of delays have been removed to show the rest in greater detail. It shows that nearly all delays are less than 1s. Also, we can see that the distribution is bimodal, with peaks at about 120ms and 860ms.

This pattern is present when considering the connection delays per router. We group the delays up to 1s into bins of size 20ms. Then we examine the top three most-populated bins for evidence of bimodality. We find 100 routers for which there exist, among these three bins, both one centered below 300ms and one centered above 700ms. The connection delays over time for a one such router - *b0b3r* (193.221.122.229) - is shown in Figure 8.4(b). The clear timing stratification suggests a cause other than varying Internet-route congestion or host-resource contention. Furthermore, the RTTs are unimodal, and so the cause is likely to be in the Tor routers. We believe that this is due to read/write rate-limiting that Tor servers manage by periodically filling global token buckets.

We can combine the RTT and connection delay measurements to estimate the distribution of host processing delays. For each router, we estimate the RTT and connection-delay distributions by partitioning delay times from 0 to 1000ms into 50 bins and treating the normalized frequency counts as probabilities.
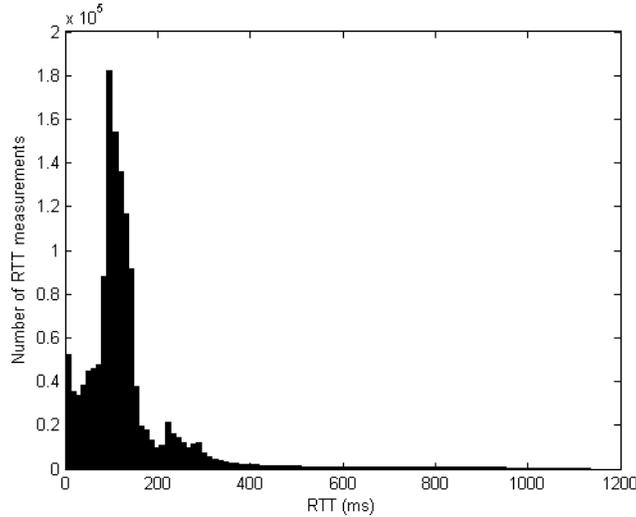
Figure 8.3: Round-trip times (top 1% removed)

The connection delay is the sum of the RTT and the processing delay, and therefore, if we assume independence of the RTT and processing-delay distributions, we can derive the processing delay distribution for each router. Considering the resulting distribution over all routers, as shown in Figure 8.5, there is almost a 40% probability of having a processing delay of nearly zero. Thus processing delays are not due to the inherent cost of computation but to limited resources at the routers.

The delays of all 400-byte packets that were successfully forwarded is shown in Figure 8.6(a). We again see times that cluster around certain levels. Here, there are six obvious levels, approximately centered at 125ms, 860ms, 1350ms, 1840ms, 2330ms, and 2820ms. To examine any such clustering per router, we group the packet delays for each router by 20ms intervals from 0-5000ms and look at the largest six such groups. There are 136 routers with two groups of these six within 100ms of the observed cluster peaks. There are 14 routers with three such groups. The delay time-series of the router $b0b3r$ (193.221.122.229), which exhibits clustering, is shown in Figure 8.6(b). We can see that the different levels are interleaved. Thus phenomenon is undoubtedly due to the same mechanism underlying the similar pattern in connection delays.

From the delay measurements we can estimate the added delay that would have resulted from our protocol. The protocol chooses a time t that must elapse from the time the packet is originally sent before the packet is forwarded. It is set such that with some success probability p the packet arrives at the forwarding router in less than t time. In our data, we examine the tradeoff that varying t sets up between the fraction p of packets that arrive in time and the forwarding delay that gets added to them.

131

(a) All delays (top 5% removed) (b) Delays for router *b0b3r*
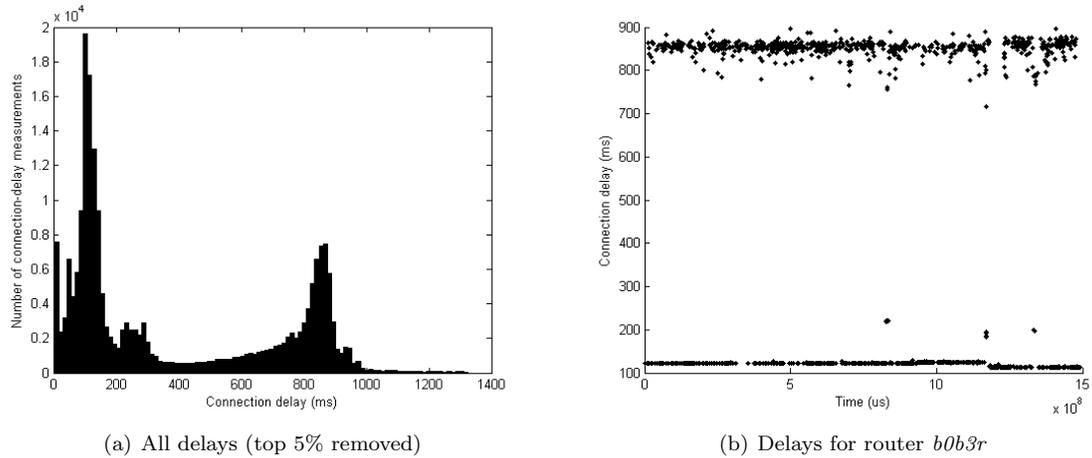
Figure 8.4: Connection delays

We divide the delay measurements by router and into periods of 6 hours. This period length is large enough to have an expected 30 experiments per period per router but is still short enough to adjust for the changes in average delay that occurs in many routers over time. Within each period, to achieve success probability p we set the send time t to be the smallest value that is larger than at least a fraction p of delays. For those delays smaller than t, using the send time t adds some delay. We look at the relative increase, that is, the total new delay divided by the original delay.

The distribution over all successfully-opened streams of relative connection delays to achieve a success probability of 0.95 is shown in Figure 8.7(a). At the 50th percentile, the relative connection delay is less than 1.48. We also must consider the failure rate of opened streams. We could simply consider failures to be connection with infinite delay, but failures often occur for reasons that have nothing to do with congestion. Therefore it is useful to keep them separate. These failures occur after a Tor circuit has been successfully created with the destination router. It thus seems unlikely that most of these failures are due to resource constraints. The distribution of connection failures over all periods is shown in Figure 8.7(b). At the 50th percentile, the failure rate is less than 0.005.

The data for relative packet delays appears in Figure 8.8(a). At the 50th percentile, the relative packet delay is just over 2.95. We also considered the distribution over periods of failure rates for forwarding packets, shown in Figure 8.8(b). The rate stays below 0.005 until the 99th percentile. The reason for this is that packet sends are not even attempted when the connection fails. It shows that once a router successfully opens a connection, it is quite likely to successfully forward all the data that it receives.
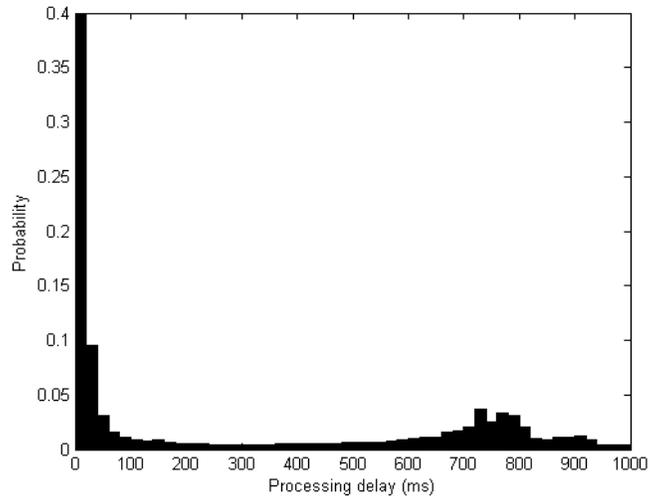
Figure 8.5: Processing distribution
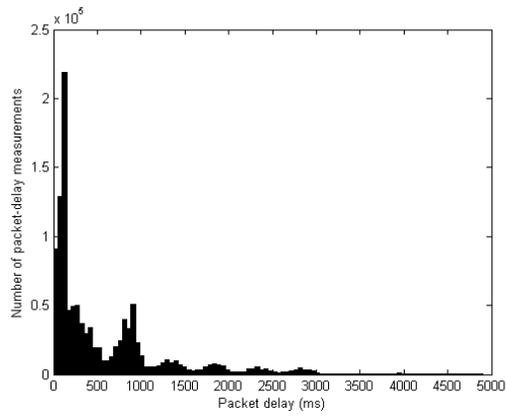
(a) Packet delays (top 1% removed)
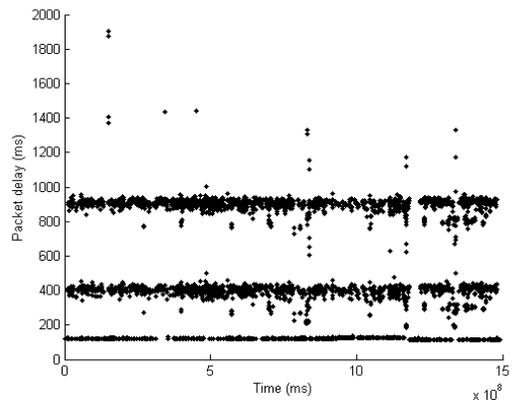
(b) Packet delays for router *b0b3r*

Figure 8.6: Packet delays

133
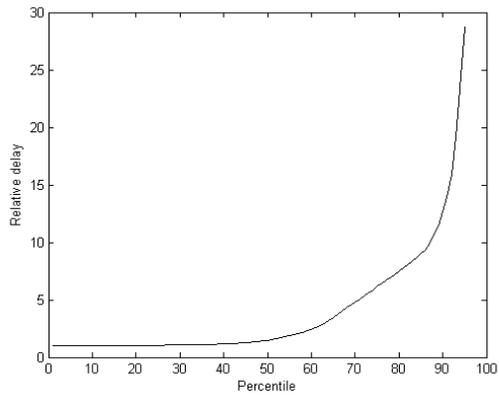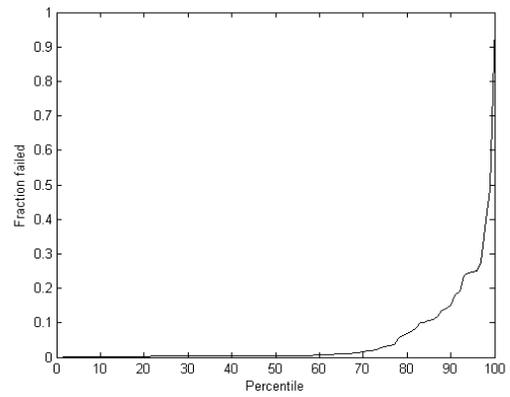
(a) Relative delays at p=0.95   (b) Failures

Figure 8.7: Connection delays and failures
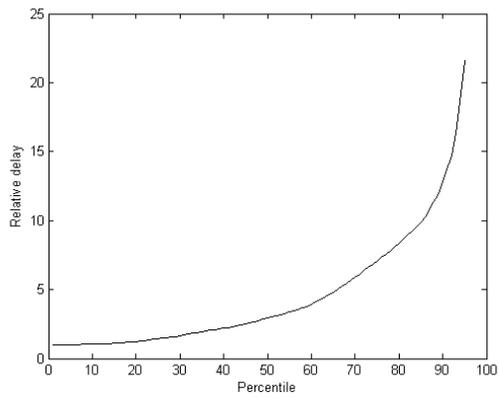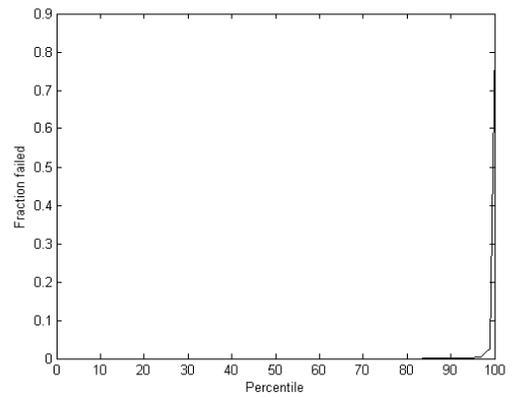
(a) Relative delays at p=0.95   (b) Failures

Figure 8.8: Packet delays and failures

Chapter 9

Conclusion and Future Work

We have been motivated by a gap between provably and highly-anonymous communication protocols such as Dining Cryptographers networks and practical protocols such as Crowds. Onion routing, in particular, has been intriguing because of its popularity, our incomplete understanding of its properties, and its potential as a base for designing new protocols. In our attempt to analyze and improve upon it, rather than focus on developing a set of tools and applying them to a set of technically-similar problems, we chose to follow the topic from basic modeling and low-level analysis to new designs and network measurement.

Our models borrow heavily from existing literature, but do adhere to a standard of fully specifying all components - the users, the network, the adversary, and the criteria for evaluation. They include several features that we hope to see developed in future work on anonymous communication. One such feature is a precise user model. Our user model, in which the population comes together at once to choose destinations probabilistically and send messages to them, is incomplete in many ways. However, it shows that the performance of anonymous communication networks depends closely on the behavior of users. This is in contrast to traditional cryptographic operations in which the adversary is the main actor. How many users there are, how active they are, and how they relate to each other is an important component of anonymous communication. Another feature is the local, active adversary model. This model certainly does not capture everything an adversary might plausibly do - for example, he might also add congestion into the network or move around during protocol operation. However, the local, active adversary is realistic enough that restricting ourselves to his behavior has and may continue to give insight into ways to protect against his actions. A third key feature is the use of a precise probabilistic anonymity metric with which to perform

analysis. Although much has already been written about different metrics, their relationship to one another is not well-understood, and there does not seem to be consensus on their appropriateness for mathematical analysis. We have seen the value of such precise definitions, though. In our probabilistic analysis of onion routing, for example, we prove upper and lower bounds on user anonymity and isolate the effect of users on each other's anonymity. Finally, we explicitly consider latency and message complexity as important criteria for protocol design.

We analyze the anonymity of onion routing possibilistically at a low level and the probabilistically at a higher level. Our analysis of onion routing is unsurprising in many ways. The possibilistic analysis shows that the adversary could only determine those parts of the circuit he directly observes. The probabilistic analysis shows that, in at least one typical case, the dominating weakness in anonymity is the chance that the adversary controls the first and last routers on a circuit. However, these calculations can be surprisingly intricate, and by following a formal model and definitions, we were forced to consider the implications of the behavior of the population on the individual.

We develop two proposals for improved efficient anonymous communication. The first considers using trust to improve onion routing, and it uncovers a key tension between working with trusted components and presenting a moving target to the adversary. The second considers the careful use of redundancy and the explicit timing of events to defeat the first/last correlation attack and achieve anonymity below the barrier faced by onion routing. At the same time, the protocol achieves reasonable efficiency. To evaluate this last point, we performed several weeks of measurement on Tor, possibly the largest anonymous-communication network publicly available. While the protocol design isn't mature enough for implementation and testing, the results of the study show that it has promise as a practical tool. Both new protocols offer perhaps the most potential for future work.

There are several interesting and important questions on trust to examine: What happens when a network is shared between entities that do not share trust levels placed on the nodes? What is the impact of trust on profiling in this case? What is the effect of learning if we add time to the model and allow the adversary to rove rather than conducting a one-off attack?

Though our motivation in our trust model is onion routing, our analysis applies to any network where it would be beneficial to reduce the chance of circuit-endpoint threats by choosing circuits with less vulnerable endpoints. It clearly generalizes to other low-latency anonymity designs, such as Crowds [68]. It also applies beyond networks for anonymity to other concerns. For example, network endpoints may be able

to collaborate to cover up checksum or other errors that might flag data-integrity attacks. And, capturing internet traffic for any kind of analysis (cryptanalysis, textual analysis, traffic analysis, etc.) may be easier to do or harder to detect or both if pairs of nodes are collaborating for route capture. Alternatively they might collaborate for unfair resource sharing. Similar observations apply to ad-hoc and peer-to-peer networks and to sensor networks, for which vulnerability of cheap, low-power, and physically accessible nodes is a known concern. Going further, our results are not restricted in applicability to path endpoints. In any setting in which sets of principals can collaborate so that a successfully compromised pair can conduct an attack our results are potentially applicable. Examining larger numbers of nodes being attacked than just pairs is one possible generalization of this work that should apply in many settings.

The protocol described in Chapter 8 would benefit from some developments that have to potential to make it truly useful and effective.

Foremost among these is to design and evaluate a usable padding scheme. We would like it to provide large anonymity sets and low overhead. It should also allow the return padding scheme to be predetermined. It could be possible to weaken the total predictability of return traffic by allowing the user to send updates of the return padding scheme to the return path. One would have to be careful to maintain a large set of other users providing the same updates, however.

Also, we have avoided for now optimizing the efficiency of processing at the routers. Onion routing, in particular, has developed good techniques to make this aspect of the protocol fast. For example, we could use persistant circuits to speed up detecting duplicate packets. We could also switch to private keys to speed up decryption.

Our analysis could be improved in some areas as well. First, we have considered benign link delays to be failures to evaluate their effect on anonymity. However, the protocol calls for such late arrivals to be immediately forwarded. The result of such forwarding seems likely to, at some point, lead to the packet catching up to one of its send times. Understanding this process better could improve the expected anonymity of the protocol. Also, we have observed that most of the delays measured appear to be due to congestion at the router. Tor is not optimized for latency, and therefore we should seek to understand the underlying resource issues in such network to accurately determine the added latencies of our protocol.

Bibliography

[1] Martín Abadi. On SDSIs linked local name spaces. *Journal of Computer Security*, 6(1/2):2–22, 1998.

[2] Anonymous web surfing by anonymizer. `https://www.anonymizer.com`.

[3] Adam Back, Ulf Möller, and Anton Stiglic. Traffic analysis attacks and trade-offs in anonymity providing systems. In *Proceedings of Information Hiding Workshop (IH 2001)*, pages 245–257, 2001.

[4] Kevin Bauer, Damon McCoy, Dirk Grunwald, Tadayoshi Kohno, and Douglas Sicker. Low-resource routing attacks against tor. In Ting Yu, editor, *WPES'07: Proceedings of the Workshop on Privacy in the Electronic Society*, pages 11–20. ACM Press, October 2007.

[5] Oliver Berthold, Hannes Federrath, and Marit Köhntopp. Project anonymity and unobservability in the internet. In *Proceedings of the tenth conference on Computers, freedom and privacy (CFP 2000)*, pages 57–65, 2000.

[6] Oliver Berthold, Andreas Pfitzmann, and Ronny Standtke. The disadvantages of free MIX routes and how to overcome them. In H. Federrath, editor, *Proceedings of Designing Privacy Enhancing Technologies: Workshop on Design Issues in Anonymity and Unobservability*, pages 30–45. Springer-Verlag, LNCS 2009, July 2000.

[7] Matt Blaze, Joan Feigenbaum, and Jack Lacy. Decentralized trust management. In *In Proceedings of the 1996 IEEE Symposium on Security and Privacy*, pages 164–173. IEEE Computer Society Press, 1996.

[8] Nikita Borisov, George Danezis, Prateek Mittal, and Parisa Tabriz. Denial of service or denial of security? How attacks on reliability can compromise anonymity. In Sabrina De Capitani di Vimercati, Paul

Syverson, and David Evans, editors, *CCS'07: Proceedings of the 14th ACM Conference on Computer and Communications Security*, pages 92–102, October 2007.

[9] Philippe Boucher, Adam Shostack, and Ian Goldberg. Freedom systems 2.0 architecture. White paper, Zero Knowledge Systems, Inc., 2000.

[10] Jan Camenisch and Anna Lysyanskaya. A formal treatment of onion routing. In *Proceedings of CRYPTO 2005*, pages 169–187, 2005.

[11] Ajay Chander, John C. Mitchell, and Drew Dean. A state-transition model of trust management and access control. In *Proceedings of the 14th IEEE Computer Security Foundations Workshop, CSFW '01*, pages 27–43. IEEE Computer Society, 2001.

[12] David Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, 4(2), 1981.

[13] David Chaum. The dining cryptographers problem: Unconditional sender and recipient untraceability. *Journal of Cryptology*, 1:65–75, 1988.

[14] Ian Clarke, Oskar Sandberg, Brandon Wiley, and Theodore W. Hong. Freenet: A distributed anonymous information storage and retrieval system. In *Proceedings of Designing Privacy Enhancing Technologies: Workshop on Design Issues in Anonymity and Unobservability*, pages 46–66, July 2000.

[15] George Danezis. Statistical disclosure attacks: Traffic confirmation in open environments. In *Proceedings of Security and Privacy in the Age of Uncertainty, (SEC2003)*, pages 421–426, 2003.

[16] George Danezis and Claudia Diaz. A survey of anonymous communication channels. Technical Report MSR-TR-2008-35, Microsoft Research, January 2008.

[17] George Danezis, Roger Dingledine, and Nick Mathewson. Mixminion: Design of a Type III Anonymous Remailer Protocol. In *Proceedings of the 2003 IEEE Symposium on Security and Privacy*, pages 2–15, 2003.

[18] George Danezis and Andrei Serjantov. Statistical disclosure or intersection attacks on anonymity systems. In *Proceedings of 6th Information Hiding Workshop (IH 2004)*, pages 293–308, 2004.

[19] Sabrina De Capitani di Vimercati, Sushil Jajodia, Stefano Paraboschi, and Pierangela Samarati. Trust management services in relational databases. In *ASIACCS '07: Proceedings of the 2nd ACM symposium on Information, computer and communications security*, pages 149–160, New York, NY, USA, 2007. ACM.

[20] Claudia Díaz, Stefaan Seys, Joris Claessens, and Bart Preneel. Towards measuring anonymity. In *Proceedings of Privacy Enhancing Technologies Workshop (PET 2002)*, pages 54–68, 2002.

[21] Roger Dingledine, Michael J. Freedman, and David Molnar. The free haven project: Distributed anonymous storage service. In H. Federrath, editor, *Proceedings of Designing Privacy Enhancing Technologies: Workshop on Design Issues in Anonymity and Unobservability*. Springer-Verlag, LNCS 2009, July 2000.

[22] Roger Dingledine and Nick Mathewson. Anonymity loves company: Usability and the network effect. In Ross Anderson, editor, *Proceedings of the Fifth Workshop on the Economics of Information Security (WEIS 2006)*, Cambridge, UK, June 2006.

[23] Roger Dingledine, Nick Mathewson, and Paul Syverson. Tor: The second-generation onion router. In *Proceedings of the 13th USENIX Security Symposium*, pages 303–319. USENIX Association, August 2004.

[24] Roger Dingledine, Vitaly Shmatikov, and Paul Syverson. Synchronous batching: From cascades to free routes. In *Proceedings of Privacy Enhancing Technologies workshop (PET 2004)*, volume 3424 of *LNCS*, pages 186–206, May 2004.

[25] Roger Dingledine and Paul Syverson. Reliable MIX Cascade Networks through Reputation. In Matt Blaze, editor, *Proceedings of Financial Cryptography (FC '02)*. Springer-Verlag, LNCS 2357, March 2002.

[26] Cynthia Dwork. Differential privacy: A survey of results. In *Proceeding of the 5th International Conference on Theory and Applications of Models of Computation (TAMC 2008)*, pages 1–19, 2008.

[27] Matthew Edman and Bülent Yener. On anonymity in an electronic society: A survey of anonymous communication systems.

[28] Nathan S. Evans, Christian Grothoff, and Roger Dingledine. A practical congestion attack on Tor using long paths, 2009. Manuscript.

[29] Ronald Fagin and Joseph Y. Halpern. I'm OK if you're OK: On the notion of trusting communication. *Journal of Philosophical Logic*, 17(4):329–354, November 1988.

[30] Nick Feamster and Roger Dingledine. Location diversity in anonymity networks. In Sabrina De Capitani di Vimercati and Paul Syverson, editors, *Proceedings of the Workshop on Privacy in the Electronic Society (WPES 2004)*, pages 66–76, 2004.

[31] Joan Feigenbaum, Aaron Johnson, and Paul Syverson. A model of onion routing with provable anonymity. In *Proceedings of Financial Cryptography and Data Security (FC 2007)*, pages 57–71, 2007.

[32] Xinwen Fu, Bryan Graham, Riccardo Bettati, and Wei Zhao. Active traffic analysis attacks and countermeasures. In *Proceedings of the 2003 International Conference on Computer Networks and Mobile Computing*, pages 31–39, 2003.

[33] Xinwen Fu, Bryan Graham, Riccardo Bettati, and Wei Zhao. Analytical and empirical analysis of coutermeasures to traffic analysis attacks. In *Proceedings of the 2003 International Conference on Parallel Processing*, pages 483–492, 2003.

[34] Xinwen Fu, Bryan Graham, Riccardo Bettati, and Wei Zhao. On effectiveness of link padding for statistical traffic analysis attacks. In *ICDCS '03: Proceedings of the 23rd International Conference on Distributed Computing Systems*, page 340, Washington, DC, USA, 2003. IEEE Computer Society.

[35] William I. Gasarch. A survey on private information retrieval (column: Computational complexity). *Bulletin of the EATCS*, 82:72–107, 2004.

[36] Ian Goldberg. Privacy-enhancing technologies for the Internet, II: Five years later. In Roger Dingledine and Paul Syverson, editors, *Proceedings of Privacy Enhancing Technologies workshop (PET 2002)*. Springer-Verlag, LNCS 2482, April 2002.

[37] Ian Goldberg. On the security of the Tor authentication protocol. In George Danezis and Philippe Golle, editors, *Proceedings of the Sixth Workshop on Privacy Enhancing Technologies (PET 2006)*, pages 316–331, Cambridge, UK, June 2006. Springer.

[38] Ian Goldberg. *Digital Privacy: Theory, Technologies, and Practices*, chapter 1. Auerbach Publications, December 2007.

[39] O. Goldreich, S. Micali, and A. Wigderson. How to play any mental game. In *STOC '87: Proceedings of the nineteenth annual ACM symposium on Theory of computing*, pages 218–229, New York, NY, USA, 1987. ACM.

[40] David M. Goldschlag, Michael G. Reed, and Paul F. Syverson. Hiding routing information. In *Information Hiding: First International Workshop, Proceedings*, pages 137–150, 1996.

[41] Joseph Y. Halpern and Kevin R. O'Neill. Anonymity and information hiding in multiagent systems. *Journal of Computer Security*, 13(3):483–514, 2005.

[42] Andrew Hintz. Fingerprinting websites using traffic analysis. In Roger Dingledine and Paul Syverson, editors, *Privacy Enhancing Technologies: Second International Workshop, PET 2002*, pages 171–178, San Francisco, CA, USA, April 2002. Springer-Verlag, LNCS 2482.

[43] Nicholas Hopper, Eugene Y. Vasserman, and Eric Chan-Tin. How much anonymity does network latency leak? In Sabrina De Capitani di Vimercati, Paul Syverson, and David Evans, editors, *CCS'07: Proceedings of the 14th ACM Conference on Computer and Communications Security*, pages 82–91. ACM Press, 2007.

[44] Dominic Hughes and Vitaly Shmatikov. Information hiding, anonymity and privacy: A modular approach. *Journal of Computer Security*, 12(1):3–36, 2004.

[45] Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. Cryptography from anonymity. In *FOCS '06: Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science*, pages 239–248, Washington, DC, USA, 2006. IEEE Computer Society.

[46] Jan Iwanik, Marek Klonowski, and Miroslaw Kutylowski. Duo–onions and hydra–onions – failure and adversary resistant onion protocols. In *Proceedings of the IFIP TC-6 TC-11 Conference on Communications and Multimedia Security*, pages 1–15, September 2004.

[47] Douglas Kelly. *A taxonomy for and analysis of anonymous communications networks*. PhD thesis, Air Force Institute of Technology, March 2009.

[48] Dogan Kesdogan, Dakshi Agrawal, and Stefan Penz. Limits of anonymity in open environments. In *Proceedings of Information Hiding Workshop (IH 2002)*, 2002.

[49] Dogan Kesdogan, Jan Egner, and Roland Büschkes. Stop-and-go MIXes: Providing probabilistic anonymity in an open system. In *Proceedings of Information Hiding Workshop (IH 1998)*, 1998.

[50] Brian N. Levine, Michael K. Reiter, Chenxi Wang, and Matthew K. Wright. Timing attacks in low-latency mix-based systems (extended abstract). In *Proceedings of Financial Cryptography (FC '04)*, pages 251–265, 2004.

[51] Marc Liberatore and Brian Neil Levine. Inferring the source of encrypted HTTP connections. In Rebecca N. Wright, Sabrina De Capitani di Vimercati, and Vitaly Shmatikov, editors, *CCS'06: Proceedings of the 13th ACM Conference on Computer and Communications Security*, pages 255–263. ACM Press, 2006.

[52] Patrick Lincoln, Phillip Porras, and Vitaly Shmatikov. Privacy-preserving sharing and correlation of security alerts. In *Proceedings of the 13th USENIX Security Symposium*, pages 239–254, 2004.

[53] Nancy A. Lynch. *Distributed Algorithms*. Morgan Kaufmann Publishers Inc., 1996.

[54] Nick Mathewson and Roger Dingledine. Practical traffic analysis: Extending and resisting statistical disclosure. In *Proceedings of Privacy Enhancing Technologies workshop (PET 2004)*, pages 17–34, 2004.

[55] Sjouke Mauw, Jan Verschuren, and Erik de Vink. A formalization of anonymity and onion routing. In *Proceedings of the European Symposium on Research in Computer Security (ESORICS 2004)*, pages 109–124, 2004.

[56] David Mazières and M. Frans Kaashoek. The Design, Implementation and Operation of an Email Pseudonym Server. In *Proceedings of the 5th ACM Conference on Computer and Communications Security (CCS 1998)*. ACM Press, November 1998.

[57] Damon McCoy, Kevin Bauer, Dirk Grunwald, Tadayoshi Kohno, and Douglas Sicker. Shining light in dark places: Understanding the Tor network. In Nikita Borisov and Ian Goldberg, editors, *Proceedings of the Eighth International Symposium on Privacy Enhancing Technologies (PETS 2008)*, pages 63–76, Leuven, Belgium, July 2008. Springer.

[58] Jon McLachlan and Nicholas Hopper. Don't clog the queue: Circuit clogging and mitigation in P2P anonymity schemes. In *Proceedings of Financial Cryptography (FC '08)*, January 2008.

[59] David Mills. Network time protocol (version 3) specification, implementation. RFC 1305, Internet Engineering Task Force, March 1992.

[60] Prateek Mittal and Nikita Borisov. Information leaks in structured peer-to-peer anonymous communication systems. In *Proceedings of the 15th ACM conference on Computer and communications security (CCS 2008)*, pages 267–278, October 2008.

[61] Ulf Möller, Lance Cottrell, Peter Palfrader, and Len Sassaman. Mixmaster Protocol — Version 2. Draft, 2003.

[62] Amarnath Mukherjee. On the dynamics and significance of low frequency components of internet load. *Internetworking: Research and Experience*, 5:163–205, 1992.

[63] Steven J. Murdoch and George Danezis. Low-cost traffic analysis of Tor. In *Proceedings of the 2005 IEEE Symposium on Security and Privacy*, pages 183–195. IEEE CS, May 2005.

[64] Steven J. Murdoch and Piotr Zieliński. Sampled traffic analysis by internet-exchange-level adversaries. In Nikita Borisov and Philippe Golle, editors, *Proceedings of the Seventh Workshop on Privacy Enhancing Technologies (PET 2007)*, Ottawa, Canada, June 2007. Springer-Verlag, LNCS 4776.

[65] Rafail Ostrovsky and Moti Yung. How to withstand mobile virus attacks. In *Proceedings of the Tenth ACM Symposium on Principles of Distributed Computing (PODC '91)*, pages 51–59. ACM Press, 1991.

[66] Lasse Øverlier and Paul Syverson. Locating hidden servers. In *Proceedings of the 2006 IEEE Symposium on Security and Privacy*, pages 100–114. IEEE CS, May 2006.

[67] Jean-François Raymond. Traffic analysis: Protocols, attacks, design issues, and open problems. In Hannes Federrath, editor, *Designing Privacy Enhancing Technologies: International Workshop on Design Issues in Anonymity and Unobservability*, pages 96–114. Springer-Verlag, LNCS 2009, July 2000.

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

[69] Marc Rennhard and Bernhard Plattner. Introducing MorphMix: Peer-to-peer based anonymous internet usage with collusion detection. In Sabrina De Capitani di Vimercati and Pierangela Samarati, editors, *Proceedings of the ACM Workshop on Privacy in the Electronic Society, WPES 2002*, pages 91–102. ACM Press, 2002.

[70] Audun Jøsang, Roslan Ismail, and Colin Boyd. A survey of trust and reputation systems for online service provision. *Decision Support Systems*, 43(2):618 – 644, 2007.

[71] Steve Schneider and Abraham Sidiropoulos. CSP and anonymity. In *Proceedings of the European Symposium on Research in Computer Security (ESORICS 1996)*, pages 198–218, 1996.

[72] Andrei Serjantov and George Danezis. Towards an information theoretic metric for anonymity. In *Proceedings of Privacy Enhancing Technologies Workshop (PET 2002)*, pages 41–53, 2002.

[73] Andrei Serjantov and Richard E. Newman. On the anonymity of timed pool mixes. In *Proceedings of the Workshop on Privacy and Anonymity Issues in Networked and Distributed Systems*, pages 427–434, Athens, Greece, May 2003. Kluwer.

[74] Vitaly Shmatikov. Probabilistic model checking of an anonymity system. *Journal of Computer Security*, 12(3-4):355–377, 2004.

[75] Vitaly Shmatikov and Carolyn Talcott. Reputation-based trust management. *Journal of Computer Security*, 13(1):167–190, 2005.

[76] Vitaly Shmatikov and Ming-Hsui Wang. Measuring relationship anonymity in mix networks. In *Proceedings of the Workshop on Privacy in the Electronic Society (WPES 2006)*, October 2006.

[77] Vitaly Shmatikov and Ming-Hsui Wang. Timing analysis in low-latency mix networks: Attacks and defenses. In *Proceedings of the European Symposium on Research in Computer Security (ESORICS 2006)*, pages 18–33, 2006.

[78] G. J. Simmons and Catherine Meadows. The role of trust in information integrity protocols. *Journal of Computer Security*, 3(1):71–84, 1994/1995.

[79] Paul Syverson, Michael Reed, and David Goldschlag. Onion routing access configurations. In *Proceedings of the DARPA Information Survivability Conference and Exposition (DISCEX 2000)*, pages 34–40, 2000.

[80] Paul Syverson, Gene Tsudik, Michael Reed, and Carl Landwehr. Towards an analysis of onion routing security. In Hannes Federrath, editor, *Designing Privacy Enhancing Technologies: International Workshop on Design Issues in Anonymity and Unobservability*, pages 96–114. Springer-Verlag, LNCS 2009, July 2000.

[81] Paul F. Syverson. Onion routing for resistance to traffic analysis. In *Proceedings of the 3rd DARPA Information Survivability Conference and Exposition (DISCEX-III)*, volume 2, pages 108–110, 2003.

[82] Paul F. Syverson and Stuart G. Stubblebine. Group principals and the formalization of anonymity. In *Proceedings of the World Congress on Formal Methods (FM'99), Vol. I*, pages 814–833, 1999.

[83] The Tor project home page. `https://www.torproject.org/` .

[84] Tor path specification. `http://www.torproject.org/svn/trunk/doc/spec/path-spec.txt`.

[85] Torstatus - tor network status. http://torstatus.kgprog.com/, April 2008.

[86] Marc Waldman, Aviel Rubin, and Lorrie Cranor. Publius: A robust, tamper-evident, censorship-resistant and source-anonymous web publishing system. In *Proceedings of the 9th USENIX Security Symposium*, pages 59–72, August 2000.

[87] Mehul Motani Wei Wang and Vikram Srinivasan. Dependent link padding algorithms for low latency anonymity systems. In *Proceedings of the 15th ACM Conference on Computer and Communications Security (CCS 2008)*, pages 323–332, October 2008.

[88] Michael K. Wright, Micah Adler, Brian N. Levine, and Clay Shields. The predecessor attack: An analysis of a threat to anonymous communications systems. *ACM Trans. Inf. Syst. Secur.*, 7(4):489–522, 2004.

[89] Andrew C. Yao. Protocols for secure computations. In *SFCS '82: Proceedings of the 23rd Annual Symposium on Foundations of Computer Science*, pages 160–164, Washington, DC, USA, 1982. IEEE Computer Society.

[90] Wei Yu, Xinwen Fu, Steve Graham, Dong Xuan, and Wei Zhao. Dsss-based flow marking technique for invisible traceback. In *Proceedings of IEEE Symposium on Security and Privacy*, 2007.

[91] Ye Zhu, Xinwen Fu, Bryan Graham, Riccardo Bettati, and Wei Zhao. On flow correlation attacks and countermeasures in mix networks. In *Proceedings of Privacy Enhancing Technologies workshop (PET 2004)*, pages 207–225, 2004.