

Studying Timing Analysis on the Internet with SubRosa

Hatim Daginawala and Matthew Wright

University of Texas at Arlington, Arlington, TX 76019, USA,
h.atim@hotmail.com, mwright@cse.uta.edu,
<http://isec.uta.edu/>

Abstract. Timing analysis poses a significant threat to anonymity systems that wish to support low-latency applications like Web browsing, instant messaging, and Voice over IP (VoIP). Research into timing analysis so far has been done through simulations or unrealistic local area networks. We developed SubRosa, an experimental platform for studying timing analysis attacks and defenses in low-latency anonymity systems. We present results of experiments on PlanetLab, a globally distributed network testbed. Our experiments validate the major conclusions, but not the detailed results, obtained by prior simulation studies. We also propose a new lightweight defense based on the principles of mix design called γ -buffering and show the limitations of this approach. Finally, motivated by our experimental results, we introduce spike analysis, a new timing analysis technique that takes advantage of unusual delays in a stream to substantially reduce errors over prior techniques.

1 Introduction

Low-latency anonymity systems provide network-level privacy for interactive applications such as Web browsing and secure shell (ssh). Such applications have timing constraints that prevent the use of techniques such as batching and reordering messages, as found in mixes [5]. Perhaps the most well-known and heavily-used low-latency anonymity system is Tor [10], a highly-distributed network of volunteer-run anonymizing proxies called Tor nodes. Users pass their TCP packets through a *circuit* of three Tor nodes and use a technique called *layered encryption* to ensure that the first Tor and second Tor nodes cannot see the contents and final destinations of the packets. In particular, this prevents the first Tor node from linking the user to her packets, which would compromise the user's privacy.

Ideally, an attacker that seeks to break the user's privacy would have to compromise or control all three Tor nodes on the user's circuit. However, timing analysis is a well-known threat to low-latency anonymity systems like Tor. In essence, an attacker that observes packets on both ends of a path can use the timings of those packets to confirm that the sender and receiver are communicating. Powerful global eavesdroppers can extend this attack to perform traffic analysis on the entire network [6]. More limited attackers, such as an eavesdropper that can see only parts of the network [17] or a malicious subset of the anonymizing proxies [14], can also perform traffic analysis.

In the case of a subset of malicious proxies, the first proxy and the last proxy on the circuit can use timing analysis to confirm that they share the same circuit. This allows

the attacker to link the user with her packets with only two proxies. Although this may not seem much harder than controlling all three proxies on a Tor circuit, prior work shows that it likely means an order of magnitude more effort for the attacker [30].

Several prior studies have shown that timing analysis can be highly effective against low-latency anonymity systems [6,14,31]. In fact, it seems to be effective despite expensive countermeasures such as circuits with more proxies and the use of *traffic shaping*, in which packets are sent only at specified times and *dummy packets* are sent when user packets are not available [14]. These prior studies, however, have either used simulation [6, 14] or small-scale experiments over local area networks [31].

Network timing over the Internet is difficult to model accurately for simulation or emulate correctly with delay generators. For example, one study uses exponentially-distributed delays [6], while another uses normally-distributed delays [28]. One study of Voice over IP (VoIP) network dynamics shows that the delay characteristics varied, such that they would best be modeled as gamma-distributed sometimes and exponentially-distributed other times [12]. Another study of VoIP dynamics suggests that a shifted gamma distribution is the best to model the network delay [15]. Choosing one of these may be sufficient to understand the performance of networks, but may not provide an adequate basis for the precise timings involved in timing analysis. Realistic models are important in understanding anonymity properties; in high-latency mixes, assumptions that the input traffic was Poisson distributed led to a design with poor anonymity properties [9].

To address the issue of realism in timing analysis experiments, we introduce SubRosa¹, an experimental platform for studying timing analysis and low-latency anonymity, designed to run on PlanetLab [1]. PlanetLab is a global overlay network that supports the development of new network services. SubRosa consists of server, client, and sink components that emulate the behavior of a low-latency anonymity system. We used SubRosa to evaluate a single timing analysis method and several defenses. In particular, we apply the *cross-correlation* method of timing analysis introduced by Levine, et al [14], and study *constant-rate cover traffic*, *defensive dropping*, and a new lightweight defense against timing analysis that we call *γ -buffering*. γ -buffering is designed to remove timing correlation from the network stream through limited delays at the proxies.

Our results show that defensive dropping is more effective than shown in simulation, even at lower drop rates than previously studied, and that γ -buffering is surprisingly ineffective unless buffering exceeds any burst of traffic. Deeper investigation into the causes of these results has led us to identify traffic patterns that are present even when defensive dropping makes typical statistical correlation ineffective. In particular, we have developed a new timing analysis technique that we call *spike analysis*, by which the attacker can achieve substantial improvements against all defenses, and defensive dropping in particular.

In Section 2, we discuss prior work in performing and defending against timing analysis in low-latency anonymity systems. Section 3 discusses the γ -buffering algorithm in depth. We then overview the SubRosa experimental platform, the PlanetLab

¹ The name is based on the Latin phrase, meaning “under the rose,” which has been used to signify secrecy, and was notably found in the recent novel [The Da Vinci Code](#). [4]

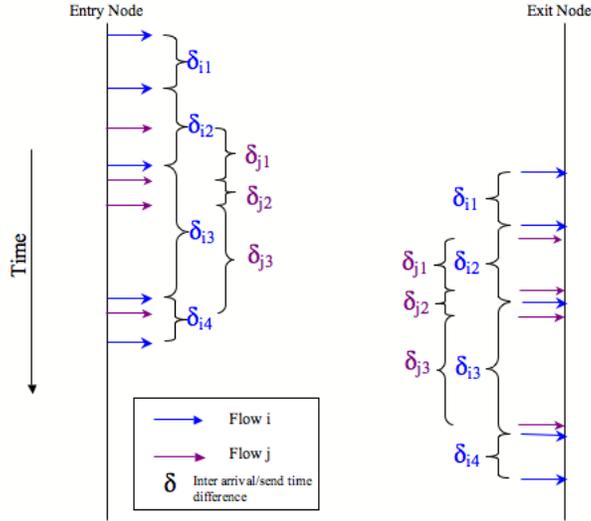


Fig. 1: Timing analysis based on correlating flows

testbed, and our experiment setup in Section 4. We present and discuss the results of our experiments, as well as the spike analysis algorithm, in Section 5 and conclude in Section 6.

2 Related Work

Many systems have been proposed and developed for low-latency anonymous communications (e.g., [2, 10, 11, 18, 22, 23]). Various studies have investigated how to break the anonymity provided by these systems and developing further defenses against their attacks. This study focuses on timing analysis attacks: attacks designed to link users with their messages based on the timings of packets in the network. Timing analysis attacks can be classified into two categories: passive and active. A passive adversary attempts to perform timing analysis based only on observations of the timings of packets. An active attack assumes that the adversary not only can observe traffic but can also delay and inject packets. Active attackers certainly are more powerful, as they can overcome limited random perturbations in the packet timings [28, 29]. However, these attacks require the ability to delay packets with very precise timings, and this can be more difficult for large-scale attacks than simply eavesdropping and recording the delays.

Other active attacks, such as the congestion-based attacks of Murdoch and Danezis [16], use a different attacker model and seek a different result. For example, recent work by Hopper et al. combines latency measurements with a congestion attack to see if two Tor paths are the same, thereby linking ongoing connections [13]. We are specifically seeking to de-anonymize the connections, i.e. to identify the IP address of the sender.

We believe that both active and passive approaches to timing analysis should be examined. For this study’s threat model we use a passive adversary who controls a subset of the mixes. A nearly equivalent model is that of a passive eavesdropper that can monitor the timings of packets entering and exiting a subset of the mixes. Our results also can be helpful for understanding the capabilities of a global passive eavesdropper. We discuss this distinction in more detail in Section 5.

Figure 1 shows the basis of timing analysis; the relative time difference (δ) between two consecutive packets remains roughly the same as flows traverse the overlay network. In other words, δ between two packets leaving the entry node will be approximately the same as that of the packets entering the exit node of the circuit. This property of the network flow is used for timing analysis. Statistical correlation can be found between various distinct streams to determine the most likely sender and receiver of the stream and compromise the anonymity [14]. A dropped packet in the stream, whether intentional or due to network congestion, will cause the timings of the packets to be off by one. As a result, the correlation will be calculated between packets that do not match. To avoid this effect, the attacker can count of the number of packets received during a time window instead of simply matching the timings of packets [14].

2.1 Defenses Against Timing Analysis

Constant rate cover traffic along the entire path is a known defense against timing analysis attacks. When all the participating nodes send data at the same constant rate, cover traffic makes all the streams look the same and makes it difficult to find correlations based on the time difference to isolate the streams. Pipenet [7] and ISDN-Mix [21] use end-to-end cover traffic whereas Tarzan [11] uses hop-by-hop cover traffic. Constant rate cover traffic, however, adds tremendous overhead to the network. For the cover traffic to be foolproof, all the nodes must be synchronized and transmit the packets at the same constant rate. Even with synchronized constant rate cover traffic vulnerability to an active adversary still remains. For these reasons, and under the assumption that the routing infrastructure is uniformly busy, systems like Tor [10] and Crowds [22] do not use any cover traffic.

Based on the idea of cover traffic, to reduce the timing correlations, various defense mechanisms have been proposed. In partial-route padding [26], all the cover traffic is dropped at a designated intermediate mix. Defensive dropping [14] generalizes the idea of partial-route padding, where the initiator creates a dummy packet and marks it to be dropped at any intermediate mix at random. When the packets are dropped at random at a sufficiently large frequency, the timing correlation is reduced [14]. Adaptive padding [25] is designed to reduce timing correlation by inserting dummy packets in the network stream, instead of dropping the packets. It is used to fill statistically unlikely gaps in the packet flow without adding latency. We note that it is highly effective. However, it requires that inserted dummy packets be sent to the responder, which may not be practical in many scenarios. This is especially true when end-to-end traffic is not encrypted, as packets with random or computer-generated content should be easy for the attacker to identify and remove from streams.

3 A New Approach to Buffering for Low-Latency Anonymity

In the original mix design, each user must send one packet in each batching period. This has previously been extended to low-latency mixes by the use of constant rate traffic, in which each user emits packets at a constant rate, using dummy packets when real user traffic is not available. Constant rate traffic reduces the correlation between different flows since all the flows receive constant number of packets with the same time difference. For practical mix designs, various batching approaches have been proposed [8]. Relatively few designs that include buffering, however, have been proposed for low-latency anonymity systems. In this section, we introduce γ -buffering, a new lightweight defense that extends the idea of constant rate traffic to better prevent timing analysis attacks.

3.1 γ -Buffering

γ -Buffering is a technique for buffering traffic that can be used to undermine traffic analysis attacks in low-latency anonymous communications systems. This technique is designed with the aim to maintain low latencies at the cost of some cover traffic and can be adapted for different levels of allowable latency and bandwidth use.

The main insight behind this technique is that, with sufficiently high traffic rates, standard batching techniques from mixes can be used without slowing down traffic excessively. While low-latency mixes setup and utilize paths for streams of packets, packet-level batching can effectively intertwine the streams' timing characteristics and destroy many of the patterns that attackers would seek to use in timing analysis. Furthermore, batching creates a variable intermediate delay that may be able to remove watermarking introduced by an active attacker [27].

One obvious concern with batching in low-latency anonymity systems, however, is that users would have different activity levels and some streams would be greatly delayed while buffers waited for activity on the other streams. If, however, the system uses a high rate of end-to-end cover traffic, there will typically be enough traffic on all streams. The amount of traffic entering a proxy at a given time, however, can still depend on the number of connections entering the proxy and the network conditions for each connection. We could require that each path send a packet, as proposed in PIPENET [7], but that could lead to long delays and denial-of-service (DoS) attacks. An adversary could prevent a proxy from sending messages by preventing an initiator or number of initiators from sending messages to the mix.

In an attempt to remove enough timing patterns without introducing such risks, we have designed γ -buffering to be more flexible. The algorithm is simple: if there are p incoming connections to a proxy, then the proxy buffers at least $\gamma * p$ packets before sending the batch, where γ is a fraction that can vary depending on the system needs. This effectively turns each proxy into a threshold mix, with a threshold of $\gamma * p$ messages [24].

When $\gamma = 1$, each proxy will get an average of one packet from each incoming connection before sending the batch. For $\gamma > 1$, larger batch sizes help ensure security at the cost of higher average latencies. For $\gamma < 1$, we can ensure a low latency when no more than $(1 - \gamma) * p$ connections are blocked or delayed. The buffering parameter γ can

be controlled at the system level or by individual proxies. When controlled by proxies, different proxies may offer different γ values, allowing users to select paths with higher or lower amounts of buffering to suit their needs. If the ratio of users to proxies grows, meaning that there are more connections per proxy, then γ may be lowered due to greater cover traffic. This kind of non-uniform path selection, however, may have other consequences for user anonymity [3], and further investigation is beyond the scope of this work.

The other primary benefit of γ -buffering would be its resilience to changing network conditions and DoS attacks. When other users' connections fail or are delayed, this creates delays for the user. Some delay is good; if the user sends traffic too aggressively, then she will be subject to easier traffic analysis. The delay is bounded, however, as long as at least the user's own traffic continues to reach the proxy. End-to-end attacks will become more difficult, even when conducted by a global adversary, as long as some other paths continue to operate.

Another way in which γ -buffering would be resilient to active attacks is that delays introduced along a user's path will multiply to show delay on many other paths in the system. The increased delay, introduced early in the path, will likely delay an entire batch of packets at the next proxy. These delayed packets cause further upstream delays, with an effect that is exponential in the length of the path. An attacker may observe a delay that has propagated along either the original path or one that has been introduced by the buffering, making it difficult to distinguish false positives from correct matches. Unlike the addition of random delays along the path, γ -buffering introduces delays simultaneous with other delays in the system, so that timing effects occur together and are much less useful to the attacker for differentiating between paths.

We note that an attacker-controlled node could choose to not buffer packets properly. However, an attacker controlling the first and last proxies would still need to contend with buffering at intermediate nodes. If the attacker needs to control the intermediate nodes to be successful, then we have made a substantial improvement in the system's defense. If the exit and intermediate nodes are compromised, but the first node in the circuit is not, then the attacker will not only need to eavesdrop on the initiator or the first node, but also contend with buffering at the first node. The attacker has similar problems when the last node is not compromised.

3.2 Partial Buffering

Due to early findings that γ -buffering was not as effective as we had expected, we also developed a slight variation that we call *partial buffering*. Partial buffering implements γ -buffering at the intermediate proxies only, not at the first or last proxies. In particular, buffering at the first proxy creates variation in the traffic patterns that could then be observed at the last proxy. Intermediate buffering may only remove some of this variation. With partial buffering, these variations are not introduced by the first proxy, but we still get the benefits of buffering at the intermediate proxy. We evaluate both γ -buffering and partial buffering in our experiments.

4 Experimental Design

In this section, we describe SubRosa, an experimental platform for investigating timing analysis on the Internet. We also describe our experiments using SubRosa, which were conducted on PlanetLab.

4.1 PlanetLab

PlanetLab (see <http://www.planet-lab.org>) is a global research network that supports the development of new network services. It is an overlay network consisting of computers distributed over six continents, with the highest concentration of nodes in North America, Europe, and East Asia. Most of the machines are hosted by participating research institutions; all of the machines are connected to the Internet, and some have Internet2 connections. All the computers on PlanetLab run a Linux-based operating system from a read-only media. The key objective of PlanetLab's software is to support distributed virtualization – the ability to allocate a slice of PlanetLab's network-wide hardware resources to an application. This allows an application to run across all (or some) of the machines distributed over the globe, where at any given time, multiple applications may be running in different slices of PlanetLab.

One effect of virtualization is that our testing process may not get access to the system for relatively long periods of time. The advantage of this is that the system does not behave like a set of dedicated servers for individual streams — performance varies depending on the load from other experiments. The disadvantage is that the platform is less consistent than we might expect anonymizing servers with many simultaneous connections to be. Following the guidance of Peterson et al., this does not prevent us from getting accurate latency measurements [20]. To provide consistent expected performance and dedicated resources on a community shared network, PlanetLab allows for reservation of resources through the Sirius Calendar Service. Reservation entitles the slice a dedicated 25% of CPU capacity and 2.0 Mbps of the system's bandwidth. All other active slices share the remaining resources with equal priority. Unfortunately, Sirius is still under development and does not reliably schedule and allocate resources. PlanetLab has over 800 computers located in over 400 locations as of October 2007.

A distributed and geographically dispersed Linux-based network testbed was an ideal platform for collecting data for this research. The global distribution of PlanetLab nodes allows us to test the effects of real network latency. We created a slice with over 300 nodes for our experiments on PlanetLab. Sites with explicitly identified Internet2 connection were excluded. PlanetLab has various deployment and monitoring tools available, but we developed our own to meet our needs as the existing tools were lacking features needed to manage our experiments.

4.2 SubRosa: An Experimental Platform for Studying Timing Analysis

We developed SubRosa to run on PlanetLab and collect data for this study. We designed SubRosa to emulate the behavior of a Tor-like network over the unreliable UDP transport. It is a simple application for collecting timing data and does not use encryption. Only information visible to an adversary in the presence of encryption has been used to

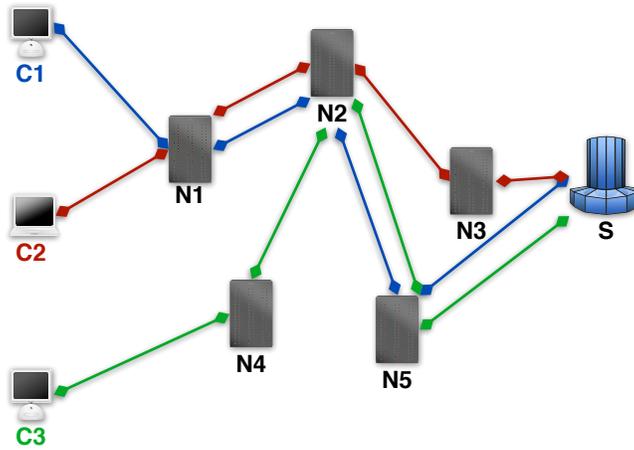


Fig. 2: Data flow in an anonymizing network

perform timing analysis in our experiments. This means that SubRosa does not capture delays due to encryption and decryption at the proxy. However, we found that the delays due to virtualization on PlanetLab nodes were occasionally quite high (more than one second) and likely suffice to emulate high-load scenarios. SubRosa is written in C and consists of three components: controller, srsrserver, srclient.

- The *controller* component starts, stops and checks the status of the other components. We used it to deploy the rest of SubRosa onto PlanetLab and keep track of the versions of the components while conducting experiments.
- The *srsrserver* component acts as the Mix. It can act as the first mix, the last mix, or as any intermediary mix on multiple paths simultaneously and it also captures timing information of the packets it sends and receives.
- The *srclient* component represents the client and is responsible for generating data on the network.

SubRosa is designed to collect timing data for various defense algorithms against timing analysis and, hence, it is versatile. Hooks and exits are designed to collect timing data as well as easily implement different algorithms for collecting data. All the variable parameters are read from the configuration file.

Data flow in SubRosa is shown in Figure 2; circuit-building proceeds as in Tor [10]. $N1$ through $N5$ are nodes running srsrserver, $C1$ through $C3$ are nodes running srclient, and S is the responder, or sink. We did not utilize a sink in this study.

4.3 Methodology

In our experiments, we fixed the path length to three. Since we did not use a sink, the exit node on the path acts as a sink and generates the response. All the traffic generated

for the experiments was constant bit rate. Five PlanetLab nodes were selected based on the load and the available bandwidth to act as servers for our experiments. Due to the nature of the test bed, if and when a node was not available or was heavily loaded, that node was replaced by another carefully chosen node. For use as servers, we only selected nodes with uptime of more than 3 days, no bandwidth restrictions, and less than 0.5 seconds response time. Servers were selected by manual use of the CoMon monitoring infrastructure for PlanetLab. All the servers were closely monitored during the experiments using CoTop slice-based top for PlanetLab. Results were not considered in the analysis if the node did not perform consistently for the duration of the experiment.

More than 300 nodes were used as clients. For each experiment, clients were chosen at random from this set nodes. At startup, clients randomly choose three out of the five server nodes. Circuits are then established with those three nodes on the path. The duration of our experiments was approximately 15 minutes each. Logs were collected for analysis after each experiment, and timing data was extracted from the logs using Perl scripts. Timing data was converted to zero base time to avoid time synchronization issues and to accommodate the scripts for analyzing the data which were provided by [14]. Experiments were conducted using 25 clients and packets were generated every 300ms and 100ms. We also collected data with larger numbers of clients, up to 100, but saw no trends in the data as the number of clients grew.

Timing analysis For the timing analysis, we use the methodology and programs used for the work of Levine et al. [14]. We describe the basic idea here, but refer the reader to that paper for more details.

The main observation of our timing analysis method is that the timings seen by the first proxy on a given path and the timings seen by the last proxy should be statistically correlated. Thus, standard methods for statistical correlation should be effective in determining whether the two proxies are observing the same path. There is a significant caveat, however, in that a single dropped packet could cause the wrong packet times to be compared. Finding the right match between sent and received packets could be computationally expensive and is likely to only be an approximation. Levine et al. propose to divide time into *windows* and count the number of packets that arrive for each proxy during a given window. They use the timings of early packets to line up the windows. Finally, a statistical *cross-correlation* is taken on the packet counts. A number of these correlations is taken for each pair of first and last proxy, using a reasonable range of alignments for the windows, and the best such correlation for the pair is chosen.

The cross-correlation values provide a statistical diagnostic test that tells the attacker how closely the streams seen by the attacker are correlated. These values can be compared for streams that are and are not on the same path. We set a *threshold* correlation value and say that the attacker thinks that all pairs of streams with correlation values above the threshold are indeed on the same path.

We also introduce a new timing analysis technique called spike analysis. This technique is described in Section 5.2.

γ -Buffering γ -Buffering is implemented on the servers. Multiplier γ is configured at the start of the application. γ is multiplied with the number of active circuits on the node to obtain the number of packets to buffer, before queuing it on the send queue. During the buffering period all the packets received are stored in a list in random order. Once the desired number of packets is buffered, the packets from the list are put on the send queue and sent out using a pool of threads. The thread pool is created at the startup and is kept alive for the life of the application to avoid overhead and improve the response time. A pool of five threads was used for γ -buffering. γ -buffering was implemented with a delay of 180 seconds after the server started, to avoid lockups during the circuit building process. The packets sent during the first four minutes were discarded to compensate for the startup delay. γ values of 0.5, 1.0 and 1.5 were used for the experiments.

Defensive Dropping Defensive dropping is initiated by the client. The client, at random, selects the packets to be dropped and marks it to be dropped at the intermediary node. The drop command is set in the header for the intermediary node and, hence, no other nodes on the path are aware of the dropped packets. The server, upon receiving a packet with a drop command, logs the packet and stops further processing of the packet by discarding it from the receive queue. We used drop rates of 20% and 50%.

Constant-rate cover Baseline data using simple constant rate cover traffic, and no other defenses against timing analysis, were also collected for experiments with 25, 50, 75 and 100 nodes. Packets were generated at every 100 milliseconds and 300 milliseconds; other parameters were kept in line with other experiments.

5 Results and Discussion

In this section, we present the results of our experiments on timing analysis of a Tor-like network. We begin by showing how effective each defense is against statistical analysis. We then then examine timing results in more depth, describe a new timing analysis technique called *spike analysis*, and show the effectiveness of this technique.

5.1 Effectiveness of Defenses

We now show relative effectiveness of various defenses against statistical cross-correlation analysis. The attacker's correlation-based diagnostic test is subject to two types of errors: *false negatives*, in which the attacker fails to identify two streams as belonging to the same path and *false positives*, in which the attacker incorrectly identifies two streams from different paths as being from the same path. False negatives occur when the threshold value is too high, i.e. the attacker expects the correlation to be higher than it is. Similarly, false positives occur when the threshold value is too low. Thus, as the threshold rises, the false positives drop off at the expense of a rise in false negatives.

A Receiver Operator Characteristic (ROC) curve is a graphical representation of the trade off between the false negative and false positive rates for every possible threshold.

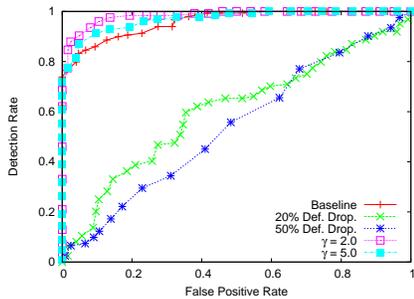


Fig. 3: ROC curves for detection against defensive dropping and γ -buffering.

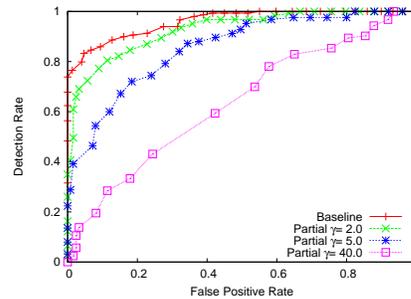


Fig. 4: ROC curves for detection against partial buffering with varying values for gamma.

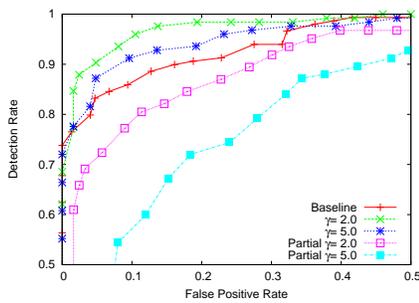


Fig. 5: Comparison of detection against partial buffering and γ -buffering.

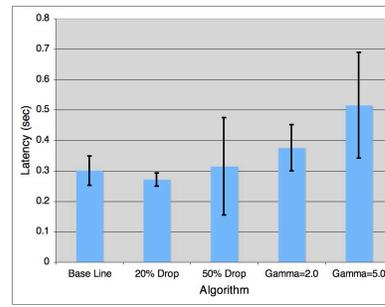


Fig. 6: Latency (RTT) for each defense. Error bars show the average standard deviation.

Conventionally, ROC curves show the false positive rate on the x-axis and the *detection rate*, the attacker’s chance of correctly linking the two streams, on the y-axis. Note that the detection rate is one minus the false negative rate. We plot different ROC curves on the same graph to visualize the relative comparison. Curves that are closer to the upper left-hand corner are better, in our case, better for the adversary. The worst case for an adversary would be a 45 degree diagonal, i.e. $x = y$, which is the same as random guessing.

We also estimate the area under the curve (AUC) for each ROC curve. The AUC gives us a single numeric value for each test that can be used for quantitative comparisons. The maximum AUC is 1.0, representing error-free detection, and the practical minimum is 0.5, which is given by the ROC curve $x = y$. We calculate our estimate of the AUC using a trapezoidal Riemann sum with the ROC points in our curves, i.e., the area under the curve between two points (x_1, y_1) and (x_2, y_2) is estimated as $A = \frac{1}{2}(y_1 + y_2) \times (x_2 - x_1)$, where $x_1 < x_2$. We have approximately five times as many points in our AUC calculation than shown in our graphs. For each set of results, we also calculate the *equal error rates* (EERs) — the point in the ROC curve where

Defense	Baseline	DD 20%	DD 50%	PB $\gamma = 2$	PB $\gamma = 5$
AUC	0.959	0.602	0.539	0.926	0.856
EER	12%	38%	48%	17%	24%
Latency (sec.)	0.300	0.271	0.314	0.375	0.515

Table 1: Area under the curve (AUC), equal error rate (EER), and latency values for defensive dropping (DD) and partial buffering (PB).

the false positive and false negative rates are equal. Table 1 shows the AUC, EER, and latency values for several settings.

We present a comparison of our baseline traffic (constant rate cover traffic), defensive dropping, and γ -buffering in Figure 3. First, we see that the attacker does quite well when only constant rate cover traffic is used. The ROC curve approaches the upper left hand corner, depicting a high success rate with relatively few errors. Table 1 shows that the attacker gets an AUC of 0.959. This appears to validate the simulation results of Levine et al. [14]. In particular, when the latency averages 50 ms and there are 1% drop rates on each application layer hop in their simulations, they report equal error rate of 8.1%. We show a higher rate of 12%, still validating the main point that statistical correlation-based traffic analysis is effective.

We see that defensive dropping is very effective against our timing analysis — both 20% and 50% dropping rates result in ROC curves close to a straight line $x = y$. Against defensive dropping of 20% and 50%, the attacker achieves AUCs of only 0.602 and 0.539, respectively. This is a significant improvement for defensive dropping in contrast with the simulation results of [14]. The equal error rates of 38% and 48% for defensive dropping rates of 20% and 50%, respectively, both greatly exceed any equal error rates achieved in their study. We explore the reasons for this difference in Section 5.2.

Perhaps more surprising is the relatively poor performance of γ -buffering. With values of γ between one and five, the attacker is as successful, if not more, than for the baseline constant rate traffic. For $\gamma = 2$ and $\gamma = 5$, we get AUCs of 0.983 and 0.970, respectively; both are larger than for the baseline traffic. We speculate that this is due to buffering at the first proxy in the circuit. This buffering can cause patterns in the traffic that may be found at the last proxy, despite further delays at intermediate nodes.

Our speculation about the failure of γ -buffering appears to be validated by relative success of partial buffering, in which the first proxy does not buffer packets. We show ROC curves for detection against partial buffering in Figure 4 and in comparison with γ -buffering in Figure 5. The latter figure only show the upper left hand quadrant of the curve for better viewing. As the curves show, partial buffering provides a better defense than the baseline constant rate traffic and much better than γ -buffering. Table 1 shows that for $\gamma = 2$ and $\gamma = 5$, we get AUCs of 0.926 and 0.856, respectively. For $\gamma = 5$, the equal error rate is 24%, which is a significant improvement over the baseline. For $\gamma = 2$, the equal error rate is 17%. Nevertheless, defensive dropping is much more effective than either setting. To get higher error rates, we set $\gamma = 40$ for a set of experiments. Even at this setting, which is impractical due to substantial delays for every packet, we

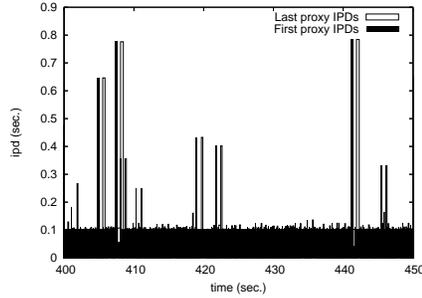


Fig. 7: Partial view of IPDs for a Baseline flow.

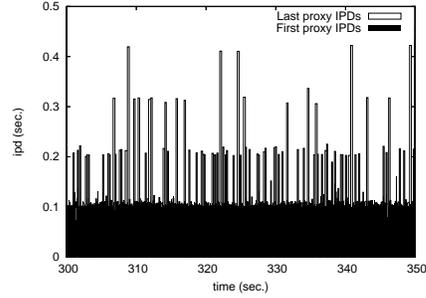


Fig. 8: Partial view of IPDs for a flow with 20% defensive dropping.

only get an AUC of 0.635 and an equal error rate of 41% — approximately equivalent to 20% defensive dropping.

We note that the difference between γ -buffering and partial buffering depends on the attacker model in an important way that we have ignored to this point. If the attacker controls the first and last proxies on the path, then our analysis holds. If the attacker can eavesdrop on all of the traffic going into and exiting these same two nodes, but doesn't control the first proxy, it is somewhat more difficult for the attacker. In this case, γ -buffering may be better than partial buffering to help hide the link between the user and the traffic exiting the first proxy. Exploring the implications of this attacker model in greater depth is beyond the scope of this work.

We also examine the latency from each of our schemes. As we see in Figure 6, the latency is roughly the same for the baseline and defensive dropping cases, between 271 ms and 314 ms round trip time (RTT). We see an increase in RTT as we introduce partial buffering. For $\gamma = 2$, we get a 25% increase in the average RTT over the baseline data. For $\gamma = 5$, this rises to 72%. This suggests that small values of γ are likely to be practical, while values over $\gamma = 5$ would increase latencies too much for general use. The error bars show the average of the standard deviations of the RTTs within each run. The large standard deviation values for 50% defensive dropping and $\gamma = 5$ appear to reflect experimental variation. The maximum RTT was over seven seconds.

5.2 Taking Advantage of Bursty Traffic

We now take a deeper look at the traffic patterns that lead to the results described above. To more carefully study the variability of the timings, we present representative graphs showing inter-packet delays (IPDs), which are the gaps between packets. In Figures 7, 8, 11, 12, 9, 10, we show the IPDs as sent by the first proxy in the circuit (dark bars) and as received by the last proxy in circuit (light bars). In other words, they represent the two timings used in our timing analysis techniques. We show 50 second intervals and the two sets of times have been offset slightly to make patterns visually recognizable. We say that *spikes* in these graphs represent high IPDs.

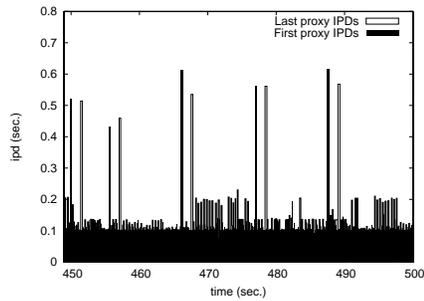


Fig. 9: Partial view of IPDs for a flow with $\gamma = 2$.

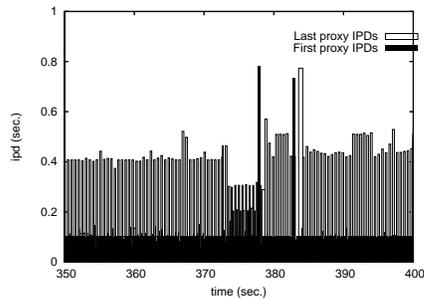


Fig. 10: Partial view of IPDs for a flow with $\gamma = 5$.

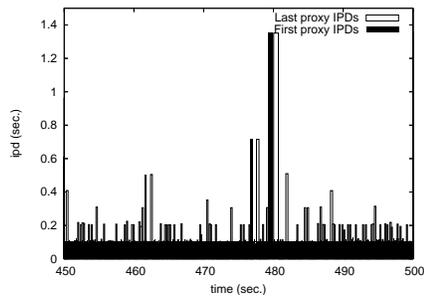


Fig. 11: Partial view of IPDs for a flow with 20% defensive dropping.

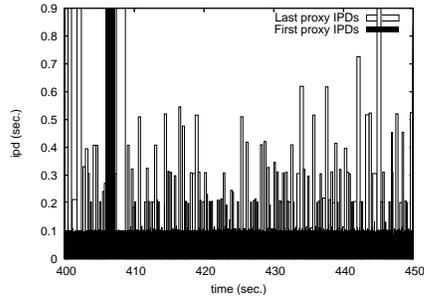


Fig. 12: Partial view of IPDs for a flow with 50% defensive dropping.

We see in Figure 7 a good match between the first proxy and last proxy IPDs for our baseline constant rate streams with no other defense. In particular, we see that every spike in the first proxy IPDs is matched with a nearly identical spike in the last proxy IPDs. Some additional spikes are present in the last proxy's IPDs, as would be expected due to additional delays and possible variability between the two measurement points. Given the similarity in the patterns, the success of statistical methods is not surprising. In contrast, we can look at a sample of IPDs from experiments with 20% defensive dropping in Figure 8. We see that there is relatively little connection between the two streams' timings. The two largest spikes in the first proxy's measurement may or may not be represented among the other spikes in the last proxy's measurement; matching the timings with confidence is difficult. Essentially, the spikes due to defensive dropping are larger than the spikes due to normal variations between the client and the first proxy.

An interesting question to consider is why we get higher error rates for the attacker than were reported in simulations by Levine et al. [14]. We speculate that perhaps in simulation, one can get better and more consistent matches in timings between the first and last proxy. If there is, e.g., clock skew, this could cause problems for the attacker.

Algorithm 1 Spike Analysis

```
TopFirstIPDs = GetTop5(FirstProxyIPDs)
TopLastIPDs = GetTop25(LastProxyIPDs)
MinError =  $\infty$ 
for Offset = -100.0 sec. to 100.0 sec. by 0.01 sec. do
  Error = 0;
  for all fIPD in TopFirstIPDs do
    Match = GetClosestInTime(fIPD, TopLastIPDs)
    if TimeDifference(fIPD, Match)  $\geq$  1 sec. then
      Error += 1 sec.
    else
      Error += ( $|fIPD - Match|$ )2
    end if
  end for
  Error = SquareRoot(Error / 5)
  if Error < MinError then
    MinError = Error
  end if
end for
```

Such an effect would be less problematic for attacking constant rate streams, in which there may be enough matching data early in the stream. Compensating for clock skew is possible [19], but is beyond the scope of this work.

Spikes are quite prominent when γ -buffering is used. For $\gamma = 2$, shown in Figure 9, we see that large spikes transfer easily from the first proxy to the last proxy. In Figure 10, we see that a few spikes carry through despite the larger variation in IPD. These results help to illustrate how statistical correlation can still find timing correlations despite buffering.

Although much of the information in a stream is lost when defensive dropping is used, defensive dropping can still leave spikes in traffic, especially larger spikes. In Figures 11 and 12, we see a number of large spikes that have apparently transferred from the first proxy through to the last proxy despite defensive dropping. This motivated the development of a simple timing analysis technique called spike analysis that is designed specifically for when there are spikes in the IPDs, e.g. due to slow response time on the proxies or by the user node.

The spike analysis algorithm is given as tested in Algorithm 1. The essential idea is to match the top f largest spikes (we use $f = 5$) in the first proxy's IPDs with f out of the top l largest spikes (we use $l = 25$) in the last proxy's IPDs. We use $l > f$, as we expect additional spikes in the last proxy's IPDs. We try a range of offsets, up to 100 seconds on either side of our best time synchronization estimate. The IPD values in Algorithm 1 include both the time of occurrence and the IPD. Matching an IPD from the first proxy's data with an IPD from the last proxy's data means finding an IPD that is close in the time of occurrence. We choose the closest matches given the offset and calculate the error as the standard deviation. If no match is within 1.0 seconds, we add an error of 1.0 second (prior to dividing by the number of matches sought and taking

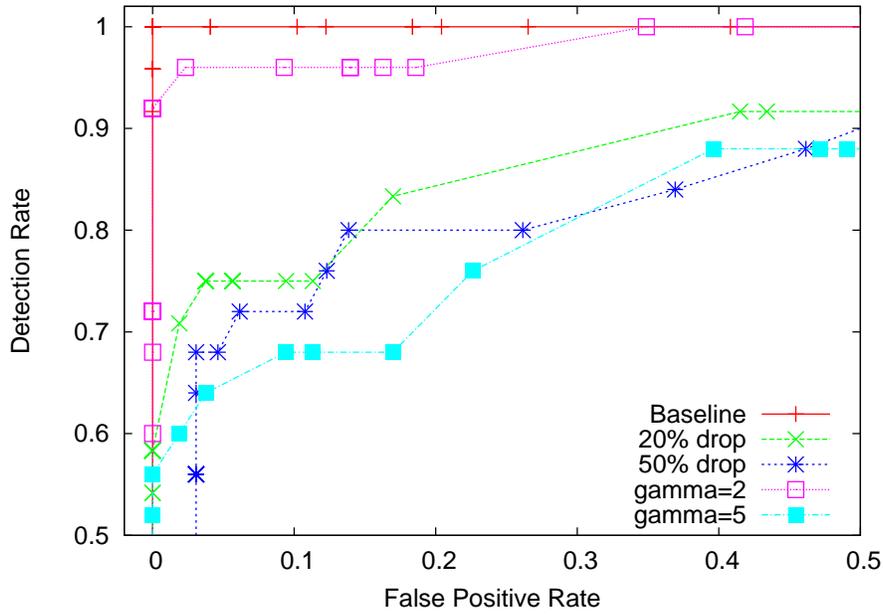


Fig. 13: Spike Analysis: ROC curves for different defenses.

Defense	Baseline	DD 20%	DD 50%	PB $\gamma = 2$	PB $\gamma = 5$
AUC	1.0	0.913	0.871	0.989	0.864
EER	0%	17%	20%	4%	24%
Latency (sec.)	0.300	0.271	0.314	0.375	0.515

Table 2: Spike Analysis: Area under the curve (AUC), equal error rate (EER), and latency values for defensive dropping (DD) and partial buffering (PB).

the square root). These are arbitrary values that worked well in our tests. Finally, we select the offset that gives the lowest error.

In Figure 13, we see the effectiveness of spike analysis on our data in the form of ROC curves. Table 2 provides AUC and equal error rates. The main result is that defensive dropping can be attacked much more successfully with spike analysis than with statistical correlation. In particular, the AUC for 20% defensive dropping is 0.913 and the equal error rate is 17% when using spike analysis, as compared with 0.602 AUC and 38% equal error rate for statistical correlation. Note that the figure shows only the upper left hand quadrant and that the axes have been shifted slightly to show the lines with zero error rate. In particular, the attacker gets zero error for baseline traffic — he can get 100% detection with no false positives as shown by the line touching the point (0, 1) in the graph. The AUC for baseline traffic is the maximum of 1.0.

We also see that partial buffering with $\gamma = 2$ is much less effective against spike analysis than statistical analysis, with an AUC of 0.989 and an equal error rate of only

4%. Perhaps more interestingly, we see that partial buffering with $\gamma = 5$ is only a little bit less effective against spike analysis. The AUC is slightly higher at 0.864 and the equal error rate of 24% is the same as with statistical analysis. In fact, this is the most effective defense of the settings we tried against spike analysis. As we see with Figure 10, the IPD graphs when $\gamma = 5$ have many spikes at the last proxy. From this observation, we speculate that spike analysis with two mismatched streams is finding spikes in the last proxy’s IPDs that fit the profile of the first proxy’s IPDs. This leads to false positives. We have observed that tightening the time requirements for more precise matches can lead to more false negatives. Thus, with enough buffering, we seem to be able to maintain a defense against these passive timing analysis techniques.

6 Conclusion

To facilitate the research of low-latency anonymous systems in general and timing analysis on low latency systems in particular, we developed the SubRosa experimental platform. We modeled SubRosa on Tor-like systems but choose UDP as the transport protocol to best study systems in which packet timing can be dictated by design choices rather than TCP behavior. We ran experiments on the PlanetLab overlay network and collected network timing data. Using this data, we performed successful timing analysis and evaluated several known defenses to avoid the attack. We introduced a light-weight defense against timing analysis attack, γ -buffering, based on threshold mixes. Our results show that defensive dropping provides the best defense among those we tested against statistical correlation-based timing analysis attacks by a passive adversary.

We also introduced spike analysis, a new timing analysis approach motivated by our observations of inter-packet delays. Spike analysis takes advantage of the presence of large spikes in the IPD traffic pattern. If the largest spikes are smaller than the changes in IPD caused by perturbations like defensive dropping and γ -buffering, the attack will not work well. However, an attacker who controls the first proxy in the circuit could induce the necessary spikes by delaying packets. This active attack requires much lower timing precision than the watermarking methods of [28] and we have shown here that delaying only five packets can be enough to get good results. A reputation system or other mechanism to ensure that attackers do not manipulate traffic in this way would have to be very sensitive to relatively small changes in the traffic pattern. The adaptive padding approach of [25] may be a useful supplement to the defenses we have tested here, as it actively seeks to remove the spikes that form the basis of spike analysis. The practical issues involved in using this technique, however, would need to be further refined.

7 Acknowledgments

The work of Daginawala and Wright was supported in part by National Science Foundation award CNS-0549998. Thanks to George Danezis for valuable discussions during early phases of the work and Cyrus Bavarian for insight and suggestions into the design of SubRosa. Thanks also to the anonymous reviewers for their valuable comments and to Jaideep Padhye for a useful way to visualize our data.

References

1. Planetlab. <http://www.planetlab.org>.
2. A. Back, I. Goldberg, and A. Shostack. Freedom 2.0 security issues and analysis. Zero-Knowledge Systems, Inc. white paper, Nov 2000.
3. Kevin Bauer, Damon McCoy, Dirk Grunwald, Tadayoshi Kohno, and Douglas Sicker. Low-resource routing attacks against TOR. In *ACM WPES*, 2007.
4. D. Brown. *The Da Vinci Code*. Doubleday Press, 2003.
5. D. Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, 24(2):84–88, Feb 1981.
6. G. Danezis. The traffic analysis of continuous-time mixes. In *Proc. Privacy Enhancing Technologies Workshop (PET)*, May 2004.
7. W. Dei. Pipenet 1.1, August 1996. <http://www.eskimo.com/~weidai/pipenet.txt>.
8. C. Diaz and B. Preneel. Taxonomy of mixes and dummy traffic. In *Proc. Intl. Information Security Management, Education and Privacy (I-NetSec '04)*, Aug 2004.
9. Claudia Diaz, Len Sassaman, and Evelyne Dewitte. Comparison between two practical mix designs. In *Proceedings of ESORICS 2004*, September 2004.
10. R. Dingledine, N. Mathewson, and P. Syverson. Tor: The next-generation Onion Router. In *Proc. USENIX Security Symposium*, August 2004.
11. M. Freedman and R. Morris. Tarzan: A peer-to-peer anonymizing network layer. In *Proc. ACM CCS*, Nov. 2002.
12. R. Ganesh, B. Kaushik, and R. Sadhu. Modelling Delay Jitter in Voice over IP. *ArXiv Computer Science e-prints*, January 2003.
13. Nicholas Hopper, Eugene Y. Vasserman, and Eric Chan-Tin. How much anonymity does network latency leak? In *Proceedings of CCS 2007*, October 2007.
14. B. N. Levine, M. Reiter, C. Wang, and M. Wright. Timing analysis in low-latency mix systems. In *Proc. Financial Cryptography*, February 2004.
15. H. Li and L. Mason. Estimation and simulation of network delay traces for voip in service overlay network. In *Proc. Intl. Symposium on Signals, Systems and Electronics (ISSSE '07)*, July 2007.
16. Steven J. Murdoch and George Danezis. Low-cost traffic analysis of Tor. In *Proc. IEEE Symposium on Security and Privacy*, May 2005.
17. Steven J. Murdoch and Piotr Zieliński. Sampled traffic analysis by internet-exchange-level adversaries. In *Proceedings of the Seventh Workshop on Privacy Enhancing Technologies (PET 2007)*, June 2007.
18. A. Nambiar and M. Wright. Salsa: a structured approach to large-scale anonymity. In *Proc. ACM Conference on Computer and Communications Security (CCS '06)*, Oct 2006.
19. Vern Paxson. *Measurement and Analysis of End-to-End Internet Dynamics*. Berkeley, California, 1997. Ph.D Dissertation.
20. Larry Peterson, Vivek Pai, Neil Spring, and Andy Bavier. Using PlanetLab for Network Research: Myths, Realities, and Best Practices. Technical Report PDN-05-028, PlanetLab Consortium, June 2005.
21. A. Pfitzmann, B. Pfitzmann, and M. Waidner. ISDNMixes: Untraceable communication with very small bandwidth overhead. In *Proc. GI/ITG Communication in Distributed Systems*, Feb 1991.
22. M. K. Reiter and A. D. Rubin. Crowds: Anonymity for Web Transactions. *ACM TISSEC*, 1(1):66–92, Nov 1998.
23. M. Rennhard and B. Plattner. Practical anonymity for the masses with MorphMix. In *Proc. Financial Cryptography (FC '04)*, February 2004.

24. A. Serjantov, R. Dingledine, and P. Syverson. From a trickle to a flood: Active attacks on several mix types. In *Proc. Information Hiding Workshop (IH)*, October 2002.
25. V. Shmatikov and M. H. Wang. Timing analysis in low latency mix networks: Attacks and defenses. In *Proc. European Symposium on Research in Computer Security (ESORICS '06)*, Sep 2006.
26. P. Syverson, G. Tsudik, M. Reed, and C. Landwehr. Towards an analysis of Onion Routing security. In *Workshop on Design Issues in Anonymity and Unobservability*, July 2000.
27. M. Venkateshaiah and M. Wright. Evading stepping stone detection under the cloak of streaming media. Technical Report CSE-2007-6, Dept. of Computer Science and Engineering, U. Texas at Arlington, 2007.
28. Xinyuan Wang, Shiping Chen, and Sushil Jajodia. Tracking anonymous peer-to-peer VoIP calls on the Internet. In *Proceedings of the ACM Conference on Computer Communications Security (CCS)*, November 2005.
29. Xinyuan Wang, Shiping Chen, and Sushil Jajodia. Network flow watermarking attack on low-latency anonymous communication systems. In *Proceedings of the IEEE Symposium on Security and Privacy (S&P 2007)*, May 2007.
30. M. Wright, M. Adler, B.N. Levine, and C. Shields. The predecessor attack: An analysis of a threat to anonymous communications systems. *ACM Transactions on Information and Systems Security (TISSEC)*, 7(4), 2004.
31. Y. Zhu, X. Fu, B. Graham, R. Bettati, and W. Zhao. On flow correlation attacks and countermeasures in mix networks. In *Proc. Privacy Enhancing Technologies (PET)*, May 2004.