Low-Cost Side Channel Remote Traffic Analysis Attack in Packet Networks

Sachin Kadloor[†], Xun Gong[†], Negar Kiyavash[‡], Tolga Tezcan[§], Nikita Borisov[†]
[†] ECE Department and Coordinated Science Lab.
[‡] IESE Department and Coordinated Science Lab.
[§] IESE Department
University of Illinois at Urbana-Champaign
{kadloor1,xungong1,kiyavash, ttezcan, nikita@illinois.edu}

***Abstract*—This paper presents a dangerous low-cost traffic analysis attack in packet-based networks, such as the Internet. The attack is mountable in any scenario where a shared routing resource exists among users. A real-world attack successfully compromised the privacy of a user without requiring significant resources in terms of access, memory, or computational power. The effectiveness of our attack is demonstrated in a scenario where the user's DSL router uses FCFS scheduling policy. Specifically, we show that by using a low-rate sequence of probes, a remote attacker can obtain significant traffic-timing and volume information about a particular user, just by observing the round trip time of the probes. We also observe that even when the scheduling policy is changed to round-robin, while the correlation reduces significantly, the attacker can still reliably deduce user's traffic pattern. Most of the router scheduling policies designed to date are evaluated mostly on the metrics of throughput, delay and fairness. Our work is aimed to demonstrate a need for considering an additional metric that quantifies the information leak between the individual traffic flows through the router.**

I. Introduction

As more of our daily activities are carried out over the Internet, the concern for their secrecy becomes increasingly urgent. This paper presents a straightforward, yet powerful traffic analysis attack that takes advantage of the side channel that is introduced by shared routing resources within packet-based networks, such as the Internet. To describe the side channel in more detail, we will consider an example of a user, Alice, using the Internet from home. She is connected by a DSL connection to her Internet Service Provider (ISP). The largest threats to the privacy of Alice's communication come either from a compromise of her home computer, or from an unscrupulous employee at her ISP. However, The ISP is bound by anti-wiretapping laws [1] to not invade into her privacy. In any case, Alice uses cryptographic protection, such as Transport Layer Security (TLS) [2] to protect the most sensitive data, such as passwords and credit card numbers.

An attacker, Bob, is curious about Alice's activities. He does not have access to Alice's computer or her ISP. In fact, Bob can possibly be on another ISP, possibly in another country. However, Bob has access to a resource in use by Alice, namely the packet queue inside her DSL router. This resource allows him to create a side channel that will leak information about Alice's traffic. This scenario is illustrated in Figure 1. Bob sends a low-bandwidth, but high frequency probe to the router and measures the round-trip time (RTT) of his probe packets. The DSL has an incoming and an outgoing port, for the traffics addressed to, and originating from Alice's computer. Bob's probe responses and Alice's incoming traffic share the same queue, hence the delay that Bob observes would vary based on the pattern of traffic addressed to Alice.

Although the pings travel through various intermediate routers, their roundtrip time is primarily affected by Alice's traffic. This is because of two reasons. One, the intermediate routers have significantly higher bandwidth compared to the volume of the traffic flowing through them. Also, the intermediate routers carry multiple traffic flows. These two imply that the delays incurred at these routers does not change with time. Hence, the delay induced by the users' traffic is well preserved.

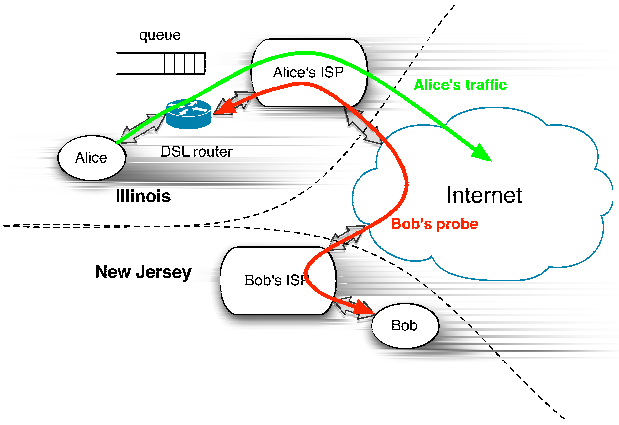


Fig. 1. Queueing Side Channel

We will demonstrate that, in fact, the delay in Bob's traffic pattern is strongly correlated with Alice's traffic pattern such that Bob can infer the activity Alice engages in. *The strength of this attack lies in the fact that it does not require special access or privileges for the attacker.* The attacker does not even need high computational power. Note that our attack can only reveal the timing and the volume of the traffic, and not the actual contents of the packets. However, recent research has shown that significant inferences can be drawn by observing just the traffic pattern of a user.

Based on his observations, Bob can perform *traffic analysis*: recovery of secret data based on packet counts, timings, and sizes. Recent research has demonstrated that traffic analysis can be used for a number of attacks: identification of web sites visited [3], [4], guesses at passwords typed [5], and recovery of phrases spoken over VoIP [6]. However, it was previously assumed that such attacks required access to one of the routers along the path of Alice's traffic. We will show that anyone who can send probe traffic to one of these routers can use the shared router queue as a side channel to carry out traffic analysis. This dramatically increases the attack surface available for traffic analysis. Even internal communication wholly contained within a network in an enterprise or military base may be subject to traffic analysis, if the network is connected to the public Internet. Likewise, even on a closed network, a low-clearance insider may be able to carry out such probes from a remote location. Remote traffic analysis can be applied to a wide variety of contexts, so we begin by discussing a few important attack scenarios.

1) Cyberstalking: An attacker may put a particular user under surveillance in order to learn more about that user's behavior on the Internet. The motives for such individual surveillance can vary widely. An attacker may want to determine whether a user is active or not, which will help him determine whether someone is likely to be home at a given time. An attacker may also wish to construct a profile of the user: Is the user likely to be affluent? What is the user's gender? Are there children in the home?

Another possibility for attack is to de-anonymize a user pseudonym, used on a blog or a discussion board. By correlating user activity with the times that updates are posted on the blog/board, it will eventually be possible to obtain statistical evidence that implicates the user.

Finally, as stated before, some encrypted traffic analysis techniques reveal the actual contents of a user's communication: statistical models can recover characters typed based on keystroke timings, while others can recover features from encrypted VoIP calls, such as the language being spoken, or test for the presence of specific phrases.

2) Unmasking Relationships: Remote traffic analysis can also be used to detect whether two users are communicating, thus inferring social or professional relationships between them. Such communications may take the form of a direct network connection between them via TCP or VoIP, or indirect conversation by instant message. In both cases, correlation of messages leaving the computer of one user and arriving at the computer of the other can give evidence of communication.

3) Monitoring at Scale: Because each probe requires only a low volume of traffic, it is possible for attackers to monitor multiple routers at once. In addition to being able to monitor multiple people, more pervasive surveillance can provide new opportunities for monitoring. For example, when de-anonymizing a pseudonym or finding connections between users, cyberstalking allows only confirmation attacks, where an existing suspicion can be confirmed or disproved, whereas more pervasive monitoring can discover previously unknown relationships.

Having discussed how our attack can be used to invade into a user's privacy on the Internet, we now discuss some of the other attacks proposed in the literature. Following that, in Section III, we formulate mathematically the problem at hand. We then present the results of our simulation and the experiment in Section IV. We conclude by describing a list of potential research problems in Section V, that arise out of the observations made in this paper.

II. Related Work

Perhaps the most related work to our proposed attack is the "clogging" attack proposed by Back et al. [7], which was proposed as an attack against low-latency anonymity schemes. As the name suggests, anonymous communication schemes allow their users to communicate with one another while masking the relationships between the users, i.e. concealing who communicates with whom. The majority of these schemes forward end-user traffic though redirecting relays, using multi-layered encryption between the source and every intermediate relay. This is done to rebuff traffic analysis and guarantee end-to-end privacy. The layered encryption is implemented through a series of relay nodes creating a virtual "anonymous tunnel". In the clogging attack proposed in [7], an adversary at one end of an anonymous tunnel can determine which nodes participate in the tunnel by sequentially clogging each of the possible relays and looking for a corresponding drop in the throughput across the tunnel. However, given the current architecture of the Internet, launching such an attack is infeasible because it would require very high bandwidth and time.

Murdoch and Danezis [8] proposed a dramatically lower-cost form of this attack and validated in the Tor network [9]. In their attack, a malicious server, through an on-off traffic pattern, systematically clogs an anonymous tunnel, while measuring the latency of the connections that run across all relays. The relays with latency functions that are most correlated to the clogging periods are identified as the members of the tunnel. Evans et al. [10] extended this attack to require less bandwidth by using long paths in the Tor network. Finally, Chakravarty et al. [11] showed that similar attacks can be used to detect normal Internet routers that forward Tor paths, as opposed to just Tor nodes, by using available bandwidth estimation techniques [12]. These attacks are structurally similar to our attack, but the probes are used to detect the presence of an attacker-controlled flow in order to violate anonymity guarantees. In our attack, the flow of interest is not under the attackers control, thus our problem exploits a side channel, rather than a covert one. In particular, previous work has used coarse-grained on-off patterns to detect flow presence, whereas we show that it is possible to infer fine-grained information about a user flow.

III. Mathematical formulation

We will consider a time slotted system. In each time slot, we will assume that there can be at most one packet arrival from each stream. This assumption is valid if we sample the

time fine enough. Let $x_i^{(k)} \in \{0,1\}$ denote the number of incoming packets from stream $k \in \{1,2\}$ in time slot i. Let $y_i^{(k)} \in \{0,1\}$ denote the number of packets of stream k served in the time slot i. Let $Q_i^{(k)}$ represent the number of packets of stream k which have not been served until time i. It follows the recursion

$$Q_i^{(k)} = Q_{i-1}^{(k)} + x_i^{(k)} - y_i^{(k)} \tag{1}$$

Let stream 1 denote the string of pings, and let stream 2 denote the traffic pattern of the other user. Define the arriving time of the n^{th} packet in the k^{th} queue as

$$T_n^{(k)} = inf\left\{ i : \sum_{j=1}^{i} x_j^{(k)} \geq n \right\} \tag{2}$$

and define the departure time of the n^{th} packet from the k^{th} queue as

$$T'^{(k)}_n = inf\left\{ i : \sum_{j=1}^{i} y_j^{(k)} \geq n \right\} \tag{3}$$

Then the roundtrip time of the n^{th} packet is

$$\tilde{T}_n^{(k)} = T'^{(k)}_n - T_n^{(k)} \tag{4}$$

Also, we define the number of arriving packets in one queue between two consecutive arrivals in the other queue, $\tilde{x}_n^k$.

$$\tilde{x}_n^{(1)} = \sum_{j=T_{n-1}^{(2)}}^{T_n^{(2)}} x_j^{(1)} \tag{5}$$

The correlation between the roundtrip times of the pings and the traffic pattern in the other stream is

$$\frac{1}{N}\sum_{i=1}^{N} \tilde{x}_i^{(1)}\tilde{T}_i^{(2)} - \frac{1}{N}\sum_{i=1}^{N} \tilde{x}_i^{(1)} \cdot \frac{1}{N}\sum_{i=1}^{N} \tilde{T}_i^{(2)} \tag{6}$$

We will analyze the performance of first come first served, and the round-robin policy against our attack in terms of the correlation metric in the following section.

IV. Performance Results

To estimate the effectiveness of our attack, we first ran a series of simulation experiments. We captured both the time of arrival of the packets and the packet sizes of real traffic arriving to our computer using Wireshark. These captured packets represent the user's traffic, about which the attacker wants to learn. The packets were captured in two scenarios, one while browsing the Internet, and second while playing a video on Youtube. The average traffic data rate is vastly different in these two scenarios, it is 21.8 KBps while browsing and 105 KBps while watching the video.

We then generate a stream of equally spaced pings, each of size 74 Bytes, and compute the roundtrip time of these pings as seen by the attacker. Roundtrip times are calculated both for the case when FCFS policy is implemented at the router and when the policy used is round-robin. We vary the frequency of these pings: 5, 50, and 500 pings/sec, which correspond to a ping data rate of 0.37, 3.7, and 37 kBps, respectively. These rates are too small for the router to suspect the presence of an attack.

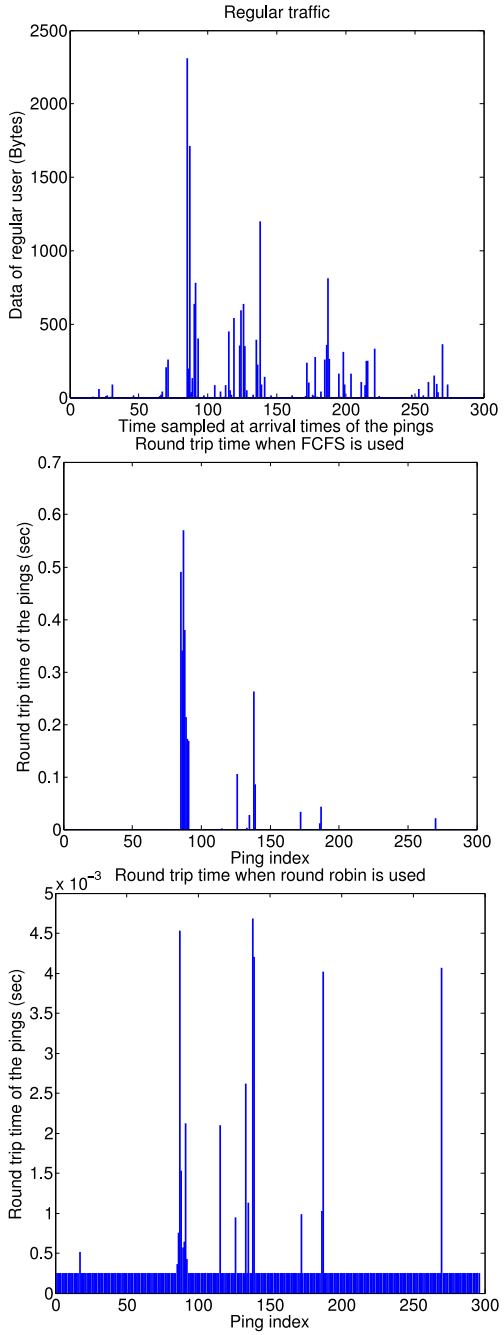


Fig. 2. Ping rate = 5 pings/sec

The router's service rate was fixed at 296 kBps. Figures 2, 3, and 4 plot the round trip time of each ping frequency for FCFS and round-robin policy. The top plot in each figure depicts the user's traffic, sampled at the arrival times of the pings. The middle plot is the round trip time of attacker's pings when FCFS policy is used and the bottom plot the round trip

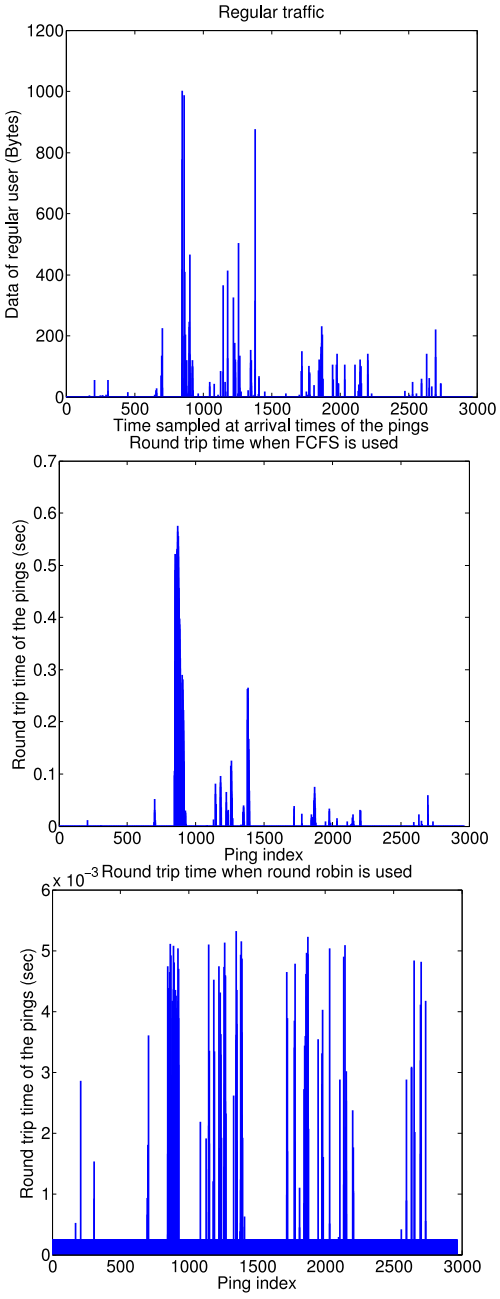


Fig. 3. Ping rate = 50 pings/sec

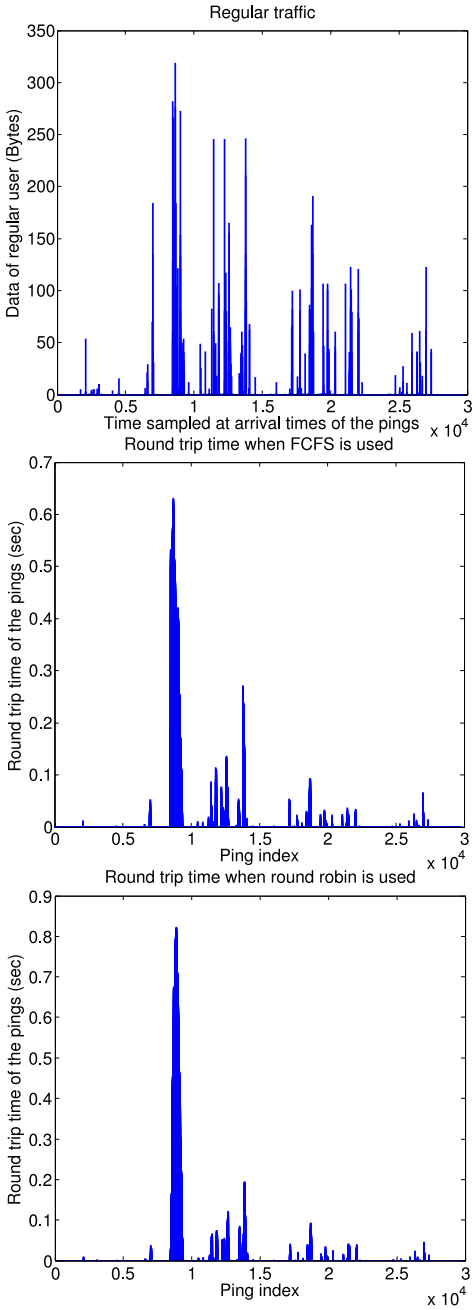


Fig. 4. Ping rate = 500 pings/sec

time when round-robin is used.

It is clearly seen that as long as the ping is not too infrequent, there is a high correlation between the input traffic and the round trip times. While the round-robin policy only reveals whether the user is active or not, the FCFS policy reveals the exact pattern. This is however not true at high ping frequencies when round-robin policy reveals as much information as the FCFS policy, and hence is vulnerable.

The simulations above demonstrate the potential of our proposed attack to uncover the traffic pattern of a remote user. *Can this attack be used in practice?* To answer this question, we setup a testbed. We observed traffic of a home DSL user in Illinois, while simultaneously sending a ping probe from a computer in New Jersey at a frequency of 100 per second. Round trip times were calculated and were later compared with the actual traffic pattern of the home user. Figure 5 shows the results of our experiment. The user's traffic is shown in green, and the round-trip times of the ping probe are shown in red. It is quite clear from the figure that there is a correlation between the two. Note that the data rate of the pings is less than 50 Kbps of bandwidth, and is unlikely to be noticed.

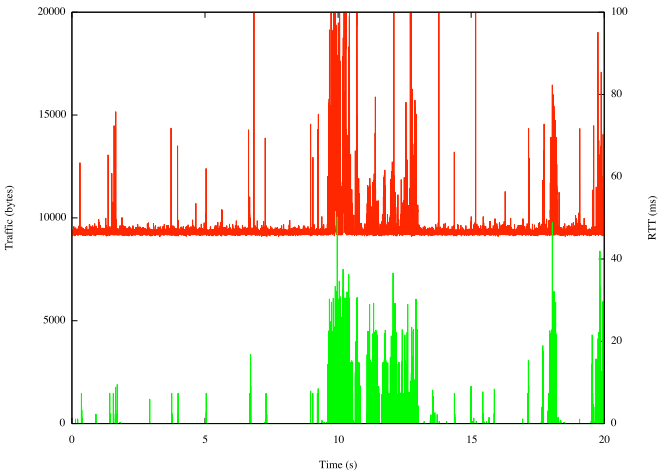


Fig. 5. Real traffic on a DSL line vs. observed probe RTTs

V. DISCUSSION

We have successfully demonstrated a nefarious attack that takes advantage of the side channel introduced by shared routing resources within packet-based networks, such as the Internet. The performance of the attack was shown using real traffic and simulation of FCFS and round-robin policies. This simple attack is so powerful because the queuing policies are not designed with purpose of preserving privacy of the users. Rather, they are designed having delay-throughput trade-offs in mind. This work illustrates the urgency for development of privacy preserving scheduling policies. The objective of such a policy will be to minimize the correlations between the traffic patterns of one user and the roundtrip times (or the waiting times) of the packets of the other user. Moreover, while correlation-based metrics such as (6) are commonly used for determining the amount of information an adversary can learn via the side channel, information-theoretically one would like to determine the maximum rate of information an attacker can reliably learn based on the observations available to her/him. Thus, a privacy-preserving scheduling policy must in fact increase the uncertainly of the adversary.

In the most general formulation of the problem, the goal of the attacker is to maximize the mutual information $I(Y;X|Z)$ where X denotes the traffic pattern of interest, Y is the observed waiting time of the probe signal, and Z denotes the probe signal. The queuing policy is modeled by the transition probability $p(Y|X,Z)$. Finding an optimal privacy-preserving policy in terms of either minimizing the correlation metric or the conditional mutual information will be the focus of our future work.

REFERENCES

[1] "Wire and electronic communications interception and interception of oral communications," United States Code, Title 18, Part I, Chapter 119.

[2] T. Dierks and C. Allen, "The TLS protocol version 1.0," RFC2246, Jan. 1999.

[3] G. Bissias, M. Liberatore, D. Jensen, and B. Levine, "Privacy vulnerabilities in encrypted HTTP streams," in *Privacy Enhancing Technologies*, 2006, pp. 1–11. [Online]. Available: http://dx.doi.org/10.1007/11767831_1

[4] M. Liberatore and B. N. Levine, "Inferring the source of encrypted HTTP connections," in *CCS '06: Proceedings of the 13th ACM conference on Computer and communications security*. New York, NY, USA: ACM Press, 2006, pp. 255–263. [Online]. Available: http://dx.doi.org/10.1145/1180405.1180437

[5] D. X. Song, D. Wagner, and X. Tian, "Timing analysis of keystrokes and SSH timing attacks," in *USENIX Security Symposium*, 2001.

[6] C. V. Wright, L. Ballard, S. E. Coull, F. Monrose, and G. M. Masson, "Spot me if you can: Uncovering spoken phrases in encrypted VoIP conversations," in *SP '08: Proceedings of the 2008 IEEE Symposium on Security and Privacy*. Washington, DC, USA: IEEE Computer Society, 2008, pp. 35–49.

[7] A. Back, U. Moller, and A. Stiglic, "Traffic analysis attacks and trade-offs in anonymity providing systems," *Lecture Notes in Computer Science*, vol. 2137, no. 245-257, p. 76, 2001.

[8] S. Murdoch and G. Danezis, "Low-cost traffic analysis of tor," in *2005 IEEE Symposium on Security and Privacy*, 2005, pp. 183–195.

[9] R. Dingledine, N. Mathewson, and P. Syverson, "Tor: The second-generation onion router," in *USENIX Security Symposium*, 2004, pp. 303–320.

[10] N. S. Evans, R. Dingledine, and C. Grothoff, "A practical congestion attack on Tor using long paths," in *USENIX Security Symposium*, Aug. 2009.

[11] S. Chakravarty, A. Stavrou, and A. D. Keromytis, "Identifying proxy nodes in a Tor anonymization circuit," in *Workshop on Security and Privacy in Telecommunications and Information Systems*, Dec. 2008, pp. 633–639.

[12] J. Strauss, D. Katabi, and F. Kaashoek, "A measurement study of available bandwidth estimation tools," in *IMC '03: Proceedings of the 3rd ACM SIGCOMM conference on Internet measurement*. New York, NY, USA: ACM, 2003, pp. 39–44.