TIMING SIDE CHANNELS FOR TRAFFIC ANALYSIS

Xun Gong[⋆†] *and Negar Kiyavash*[⋆‡]

[⋆] Coordinated Science Laboratory, UIUC
[†]Department of Electrical and Computer Engineering, UIUC
[‡]Department of Industrial and Enterprise Systems Engineering, UIUC
{xungong1, kiyavash}@illinois.edu

ABSTRACT

Traffic analysis often requires direct observations of network connections at local vantage points. In this work, we show that traffic analysis can be performed remotely by taking advantage of a timing side channel. The timing side channel results from a shared resource, namely, the scheduler between two users. Utilizing Shannon equivocation as a privacy metric, we prove that one user can learn the complete traffic pattern of the other user if the scheduler employs a first come first serve (FCFS) policy. Moreover, we show the feasibility of a real system attack exploiting the timing side channel inside a home digital subscriber line (DSL) router. This demonstrates the magnitude of the threat timing side channels pose for traffic analysis.

Index Terms— timing channel, timing side channel, traffic analysis, privacy

1. INTRODUCTION

1.1. Traffic Analysis

Traffic analysis is a crucial tool for inferring user behaviors in networks, since communications by and large are encrypted. Despite 'invisible' packet contents, an encrypted network connection often leaves noticeable 'footprints', such as packet timestamps and payload sizes. These communication patterns preserve to some extent properties of the upper-layer activities, and hence can be used to perform traffic analysis. For instance, one can classify application protocols of encrypted packet flows based on sizes, timings, and directions [1]. Moreover, statistical analysis of packet payloads can help detect abnormal activities such as network intrusions [2]. Traffic analysis may also be applied for malicious purposes, such as compromising user passwords [3], recovering spoken phrases in voice over IP (VoIP) conversations [4] and even identifying webpages a user is browsing [5].

This work was supported in part by National Science Foundation through the grant CCF 10-65022, CCF 10-54937 CAR, and in part by Air Force through the grant FA9550-11-1-0016, FA9550-10-1-0573.

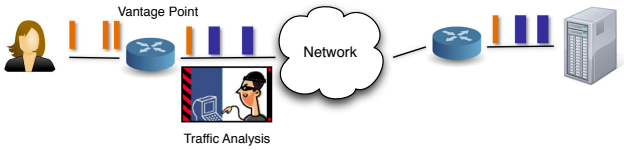

Fig. 1. A traffic analysis scenario. To gain information about the user's activity, the traffic analyzer monitors packet traces by having access to a router (vantage point) near the user.

Previous instances of traffic analysis often require direct observations of a network flow, which significantly restricts the practical application. For example in Figure 1, one has to stay close to a target user in order to capture useful packet traces for analysis. If such a local vantage point is not available, traffic analysis gets much harder. In this work, we show that even in such scenarios, it is still possible to learn traffic patterns by exploiting a timing side channel.

1.2. Timing Side Channels

Timing side channels exist in systems with common resources scheduled among multiple parties. For instance, in a single server queuing system, one user may infer the workloads of other users by noticing patterns of the server's busy periods. Although a timing side channel does not provide information as accurate as directly capturing traces, practical traffic analysis is still applicable. In [6], a timing side channel is discovered in cloud servers where jobs of several clients are assigned onto the same physical host. In that case, a cloud client would know the sizes of jobs issued by other clients if carefully measuring the time when the system hardwares are taken. Similar ideas apply to routers, where myriad packet flows meet and wait to be scheduled. By probing a router's buffer, one can estimate throughputs of connections belonging to others [7], which is shown to be useful for compromising an anonymity network [8].

In this paper, we study timing side channels for traffic analysis from both theoretical and practical viewpoints. This work differs from our previous work in [9], where the at-

tacker's probes are Bernoulli, i.e., the attacker samples the shared queue at random times according to a Bernoulli process. Given the same sampling rate, it can be proven that uniform sampling is optimal in terms of preserving the most information about the original arrival process. Therefore, the model in this paper considers a stronger attacker.

The rest of the paper is organized as follows. In §2, we introduce our timing side channel model arising from a FCFS scheduler. We characterize the information theoretical limit of this timing side channel in §3. A practical application of aforementioned timing side channel is presented in §4. We conclude in §5.

2. SYSTEM MODEL

We consider a router scheduling packets from two users, as depicted in Figure 2(a). In every time slot, each user either issues one packet or stays idle. All packets are buffered in a single FCFS queue before they are served. The router serves one packet in each time slot.

Assume one of the users is a malicious attacker who wants to learn the other user's traffic patterns. Since the service tokens are allocated in FCFS order, the attacker's departure process conveys information about the queue status, and hence the other user's arrival pattern. We use the following notation in the rest of the paper.

- We assume the user's packets are generated from a *Bernoulli process* of rate λ. Clearly, the difficulty of the attacker inferring the arrival pattern depends on the user's arrival model. It is easy to guess the arrival timings of the user if packets are issued in a predictable or regular manner; e.g., ON/OFF with a fixed period. On the contrary, as the *Bernoulli process* has the maximum entropy rate [10], hence our model sets a fairy difficult task for the attacker;
- ω is the arrival rate of the attacker. Considering that the router's service rate is 1 packet per time slot, we require $\omega \in [0, 1-\lambda)$ for the stability of the queue;
- t_i is the time slot when the attacker sends the i^{th} packet, and t'_i is its departure time;
- x_i denotes the number of arrivals from the other user between consecutive packets of the attacker (see Figure 2(b)). The attacker's goal is to learn the arrival pattern $\{x_i\}$ for all i.

2.1. Attack Strategy

We consider a periodic-sampling attacker where the attacker issues packets at *evenly-spaced* time slots to sample the user's arrival process (see Figure 2(b)), i.e., $t_{i+1} - t_i = \frac{1}{\omega}$.[1] Given the user's arrival process is Bernoulli, $x_i \sim Bernoulli(\frac{1}{\omega}, \lambda)$.

[1]More precisely, $t_{i+1} - t_i = \lfloor \frac{1}{\omega} \rfloor$ or $\lceil \frac{1}{\omega} \rceil$. For the convenience of analysis, we assume $\frac{1}{\omega}$ is an integer. The methodology and result of our analysis still are true without this assumption.

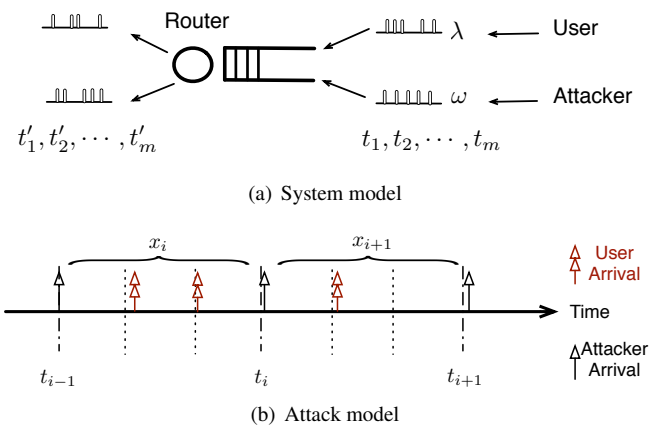

Fig. 2. A router scheduling packets for a user and an attacker. The attacker creates a timing side channel by issuing packets periodically to sample the queue length at the router buffer.

3. ANALYSIS OF INFORMATION LEAKAGE

In this section, we study how much information regarding the user's arrival pattern the attacker can learn using the timing side channel. We measure the information leakage of this timing side channel with a Shannon equivocation metric [11], as defined below.

Definition 1. *The equivocation rate, i.e., the conditioned randomness in user's arrival pattern given the observations of the attacker is given by*

$$\mathcal{P} = \lim_{n\to\infty} \frac{H\left(x_1, \cdots, x_n | t_0, \cdots, t_n, t'_0, \cdots, t'_n\right)}{n}, \qquad (1)$$

where H denotes the entropy function.

Metric $\mathcal{P}$ characterizes the remaining uncertainty of the target traffic pattern given the attacker's observations of the queue. A smaller value of $\mathcal{P}$ implies that the attack is more successful in learning the pattern; or equivalently the side channel is more suited for traffic analysis.

Note that the attacker can calculate the queue length at the buffer from his packet timings. Define q_i to be the queue length at the beginning of time t_i (before the i^{th} attack job arrives), then the attacker knows

$$q_i = t'_i - t_i - 1. \qquad (2)$$

Lemma 3.1. *When the total arrival rate satisfies $\lambda + \omega < 1$, the user's privacy is given by*

$$\mathcal{P} = H(X|Q_1, Q_2), \qquad (3)$$

where $Q_2 = \left(Q_1 + X + 1 - \frac{1}{\omega}\right)_+$, $X \sim Bernoulli(\frac{1}{\omega}, \lambda)$, and Q_1, Q_2 have identical marginal distributions.[2]

[2]The function $(a)_+ = \max\{a, 0\}$.

Proof. Applying the entropy chain rule,

$$H\left(x_1,\cdots,x_n|t_0,\cdots,t_n,t_0',\cdots,t_n'\right) = \sum_{i=1}^{n} H\left(x_i|x_1,\cdots,x_{i-1},t_0,\cdots,t_n,t_0',\cdots,t_n'\right). \quad (4)$$

Notice that the queue length seen by the attacker updates according to

$$q_i = \left(q_{i-1}+1+x_i-\frac{1}{\omega}\right)_+, \quad i=1,2,\cdots,n \quad (5)$$

which indicates that the pair (q_i, q_{i+1}) is a sufficient statistic of the pair $(t_i, t_i')'s$ in estimating x_i. Moreover, from (2) it is easy to check that $(t_i, t_i')'s$ is a sufficient statistic of (q_i, q_{i+1}) as well. Therefore, each term of the summation in (4) can be rewritten as

$$H\left(x_i|x_1\cdots x_{i-1},t_0\cdots t_n,t_0'\cdots t_n'\right) = H\left(x_i|q_i,q_{i+1}\right). \quad (6)$$

Notice that $\{(q_i,q_{i+1}), i=0,1,\cdots\}$ forms a Markov chain. Define a Lyapunov function as $V(q_i,q_{i+1}) = q_{i+1}$, it can be shown that this chain is stable if $\lambda+\omega<1$. This implies that the limit of $H(x_i|q_i,q_{i+1})$ for $i\to\infty$ exists. Denote the stationary states of (x_i,q_i,q_{i+1}) by random variables (X,Q_1,Q_2). Thus, limit of (6) can be represented by $H(X|Q_1,Q_2)$. Substituting this limit into (4) and (1), concludes the proof. Note that X has the same distribution as x_i, $Bernoulli(\frac{1}{\omega},\lambda)$, and $Q_2 = \left(Q_1+X+1-\frac{1}{\omega}\right)_+$. □

Theorem 3.2. *When the attacker issues packets at the maximum available rate, $\omega\to 1-\lambda$, user's traffic pattern is entirely leaked through the side channel, i.e.,*

$$\lim_{\omega\to 1-\lambda}\mathcal{P}=0. \quad (7)$$

Proof. Because $Q_2=\left(Q_1+X+1-\frac{1}{\omega}\right)_+$, when $Q_2>0$, X is fixed by $Q_2+\frac{1}{\omega}-1-Q_1$. Hence, we have $H(X|Q_1,Q_2>0)=0$, and can rewrite (3) as

$$\mathcal{P}=P(Q_2=0)H(X|Q_1,Q_2=0). \quad (8)$$

Derive the z-transform of Q_2 (or Q_1) as

$$E\left[z^{Q_2}\right]=\frac{\sum_{i=0}^{\frac{1}{\omega}-2}P(Q_2=i)\sum_{j=0}^{\frac{1}{\omega}-2-i}P(X=j)(z^{\frac{1}{\omega}-1}-z^{i+j})}{z^{\frac{1}{\omega}-1}-(1-\lambda+\lambda z)^{\frac{1}{\omega}}} \quad (9)$$

and take $z\to 1$ on both sides. We get

$$\sum_{i=0}^{\frac{1}{\omega}-2}P(Q_2=i)\sum_{j=0}^{\frac{1}{\omega}-2-i}P(X=j)(\frac{1}{\omega}-1-i-j)=\frac{1}{\omega}(1-\omega-\lambda). \quad (10)$$

In (10), when $\omega\to 1-\lambda$, $P(Q_2=i)=0, \forall\, i\le\frac{1}{\omega}$. Applying this result to (8), proves (7). □

4. EXPERIMENTS

We already showed that an FCFS scheduler creates a timing side channel that allows one user to learn the complete traffic pattern of another. Now we provide an evidence of such channels in real-world systems.

4.1. A Timing Side Channel in DSL

We examine a typical home digital subscriber line (DSL) environment. In Figure 3, *Alice* subscribes to a DSL Internet service. Although she applies encryptions to protect the packet contents, her traffic patterns are still vulnerable to traffic analysis because of existence of a timing side channel that we describe in the following. All packets destined for *Alice* have to pass through a queue inside her Internet Service Provider (ISP) before they are scheduled to *Alice*'s computer. If an attacker can join this queue, he gets the chance to learn Alice's traffic patterns by analyzing the queuing delays he experiences. Therefore, as long as *Bob* knows *Alice*'s IP address, he can issue Internet control message protocol (ICMP) requests to ping *Alice*, and observes the round trip times (RTT).[3]

Besides Alice's traffic, delays on intermediate routers on Bob's path to Alice's DSL router affect his RTTs. Therefore, RTT of the i^{th} ping packet, denoted by τ_i, can be expressed as

$$\tau_i=\sum_{l\in\text{links on the path}}\left(p_i^l+t_i^l+w_i^l\right), \quad (11)$$

where p_i^l is the propagation delay, t_i^l is the transmission delay, and w_i^l is the queueing delay experienced by the i^{th} ping packet on link l. Notice that p_i^l and t_i^l are determined by the length and bandwidth of the underlying physical link, so they do not change during the short period of the attack. We substitute this portion of the delay with the smallest RTT seen in the entire ping sequence to get an approximation to $\tau_i\approx\min_j\tau_j+\sum_{l\in links}w_i^l$. Moreover, since *Alice*'s DSL link has a bandwidth of only several Mbps, much lower than the rest of the links on the path (most are at backbone nodes), the queuing delay for the most part results at Alice's DSL router. Ignoring congestions on other links, we get $\tau_i\approx\min_j\tau_j+(t_i'-t_i)$. Recall that t_i and t_i' denote the arrival and departure time of Bob's i^{th} ping packet, respectively. Thus, the RTTs seen by *Bob* capture the queueing effect of *Alice*'s traffic pattern.

4.2. Measurements

We verified the above analysis by setting up an experiment. During the experiment, the host initiated HTTP connections from a DSL line with download bandwidth of 3 Mbps. At the same time, we scheduled ping requests to this host from a remote server. The pings were sent every 10 ms (i.e., $t_i - t_{i-1} = 10$ ms).

Figure 4(a) shows the volume of HTTP packets downloaded within every interval of 10 ms. Figure 4(b) depicts the observed RTTs in response to *Bob*'s pings. A quick inspection of these figures reveals a strong correlation between the

[3] We use ping packets as they have low bandwidths and thus hard to notice. Other protocols (e.g., TCP) can be considered when *Alice*'s DSL router disables the ping functionality.

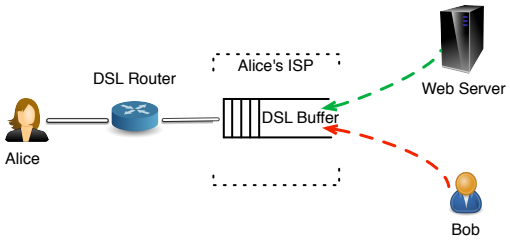

Fig. 3. A timing side channel in home DSLs. *Alice* is surfing the Internet through a DSL line, while *Bob* issues ping packets to sample the queue length at *Alice*'s DSL buffer. This creates a timing side channel, allowing *Bob* to learn *Alice*'s arrival pattern.

download traffic pattern of *Alice* and *Bob*'s RTTs. Finally, applying the processing described in §4.1 to the received RTTs, we obtain the sequence in Figure 4(c), which resembles the traffic pattern in Figure 4(a) even more closely.

The fact that the estimated RTTs in Figure 4(c) reveal features of the original traffic pattern, enables the attacker to perform traffic analysis by exploiting the aforementioned timing side channel. It is noteworthy, that what sets this attack from previous traffic analysis threats apart is that it does not require access to a local vantage point. In fact, in [12], we showed an attacker was able to remotely fingerprint the webpage visited by Alice using a home DSL. The interested reader is referred to [12] for more details.

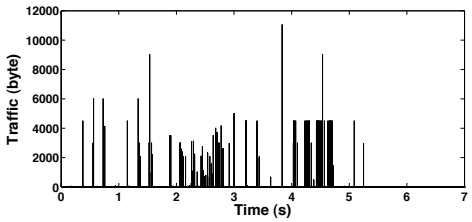

(a) HTTP bytes downloaded in each interval of 10 ms

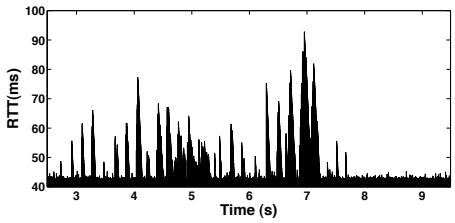

(b) RTTs of pings sent every 10 ms

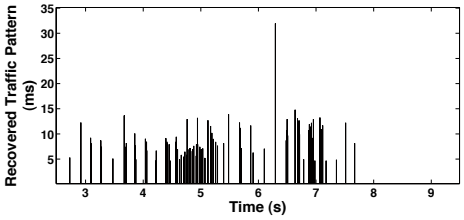

(c) RTTs after processing

Fig. 4. The measurements of timing side channel in DSL.

5. CONCLUSION

We studied the information leakage of the timing side channel that arises from an FCFS scheduler servicing two users. Such a timing channel enables one user to learn the traffic pattern of the other through the queueing delay he experiences. We prove that there exists an attack strategy, i.e, a sequence of jobs that the attacker issues, which leads to learning the other user's traffic pattern without ambiguity. In other words, the equivocation rate, i.e., the conditional randomness in the arrival process of the user given the observations of the attacker goes to zero. Moreover, we give evidence of the existence of such a channel in home DSL routers that demonstrates the threat timing side channel pose as they enable remote traffic analysis from non-local vantage points.

6. REFERENCES

[1] C. V. Wright, F. Monrose, and G. M. Masson, "On inferring application protocol behaviors in encrypted network traffic," *Journal of Machine Learning Research*, vol. 6, pp. 2745–2769, 2006.

[2] K. Wang and S. J. Stolfo, *Anomalous payload-based network intrusion detection*, 2004.

[3] D. Song, D. Wagner, and X. Tian, "Timing analysis of keystrokes and SSH timing attacks," in *10th USENIX Security Symposium*, Dan S. Wallach, Ed. Aug. 2001, USENIX Association.

[4] Charles V. Wright, Lucas Ballard, Scott E. Coull, Fabian Monrose, and Gerald M. Masson, "Spot me if you can: Uncovering spoken phrases in encrypted VoIP conversations," in *IEEE Symposium on Security and Privacy*, Washington, DC, USA, May 2008, pp. 35–49, IEEE Computer Society Press.

[5] Dominik Herrmann, Rolf Wendolsky, and Hannes Federrath, "Website fingerprinting: attacking popular privacy enhancing technologies with the multinomial naïve-Bayes classifier," in *ACM Workshop on Cloud Computing Security*, Chicago, Illinois, USA, Oct. 2009, pp. 31–42, ACM.

[6] T. Ristenpart, E. Tromer, H. Shacham, and S. Savage, "Hey, you, get off of my cloud: exploring information leakage in third-party compute clouds," in *ACM Conference on Computer and Communications Security*, 2009, pp. 199–212.

[7] S. Kadloor, X. Gong, N. Kiyavash, T. Tezcan, and N. Borisov, "Low-cost side channel remote traffic analysis attack in packet networks," in *IEEE International Conference on Communications (ICC)*. IEEE, 2010, pp. 1–5.

[8] Steven J. Murdoch and George Danezis, "Low-cost traffic analysis of Tor," in *IEEE Symposium on Security and Privacy*, 2005, pp. 183–195.

[9] X. Gong, N. Kiyavash, and P. Venkitasubramaniam, "Information Theoretic Analysis of Side Channel Information Leakage in FCFS Schedulers," in *IEEE International Symposium on Information Theory*, 2011.

[10] J. A. McFadden, "The entropy of a point process," *Siam Journal on Applied Mathematics*, vol. 13, 1965.

[11] R. B. Ash, *Information Theory*, 1990.

[12] X. Gong, N. Borisov, N. Kiyavash, and N. Schear, "Website detection using remote traffic analysis," in *Privacy Enhancing Technologies Symposium*, 2012.