

# A New Cell-Counting-Based Attack Against Tor

Zhen Ling, Junzhou Luo, *Member, IEEE*, Wei Yu, Xinwen Fu, Dong Xuan, and Weijia Jia

**Abstract**—Various low-latency anonymous communication systems such as Tor and Anonymizer have been designed to provide anonymity service for users. In order to hide the communication of users, most of the anonymity systems pack the application data into equal-sized cells (e.g., 512 B for Tor, a known real-world, circuit-based, low-latency anonymous communication network). Via extensive experiments on Tor, we found that the size of IP packets in the Tor network can be very dynamic because a cell is an application concept and the IP layer may repack cells. Based on this finding, we investigate a new cell-counting-based attack against Tor, which allows the attacker to confirm anonymous communication relationship among users very quickly. In this attack, by marginally varying the number of cells in the target traffic at the malicious exit onion router, the attacker can embed a secret signal into the variation of cell counter of the target traffic. The embedded signal will be carried along with the target traffic and arrive at the malicious entry onion router. Then, an accomplice of the attacker at the malicious entry onion router will detect the embedded signal based on the received cells and confirm the communication relationship among users. We have implemented this attack against Tor, and our experimental data validate its feasibility and effectiveness. There are several unique features of this attack. First, this attack is highly efficient and can confirm very short communication sessions with only tens of cells. Second, this attack is effective, and its detection rate approaches 100% with a very low false positive rate. Third, it is possible to implement the attack in a way that

appears to be very difficult for honest participants to detect (e.g., using our hopping-based signal embedding).

**Index Terms**—Anonymity, cell counting, mix networks, signal, Tor.

## I. INTRODUCTION

CONCERNS about privacy and security have received greater attention with the rapid growth and public acceptance of the Internet, which has been used to create our global E-economy. Anonymity has become a necessary and legitimate aim in many applications, including anonymous Web browsing, location-based services (LBSs), and E-voting. In these applications, encryption alone cannot maintain the anonymity required by participants [1]–[3]. In the past, researchers have developed numerous anonymous communication systems. Generally speaking, mix techniques can be used for either message-based (high-latency) or flow-based (low-latency) anonymity applications. E-mail is a typical message-based anonymity application, which has been thoroughly investigated [4]. Research on flow-based anonymity applications has recently received great attention in order to preserve anonymity in low-latency applications, including Web browsing and peer-to-peer file sharing [5], [6].

To degrade the anonymity service provided by anonymous communication systems, traffic analysis attacks have been studied [3], [7]–[14]. Existing traffic analysis attacks can be categorized into two groups: passive traffic analysis and active watermarking techniques. Passive traffic analysis technique will record the traffic passively and identify the similarity between the sender's outbound traffic and the receiver's inbound traffic based on statistical measures [7]–[9], [15], [16]. Because this type of attack relies on correlating the timings of messages moving through the anonymous system and does not change the traffic characteristics, it is also a passive *timing* attack. For example, Serjantov *et al.* [7] proposed a passive packet-counting scheme to observe the number of packets of a connection that arrives at a mix node and leaves a node. However, they did not elaborate how packet counting could be done. To improve the accuracy of attacks, the active watermarking technique has recently received much attention. The idea of this technique is to actively introduce special signals (or marks) into the sender's outbound traffic with the intention of recognizing the embedded signal at the receiver's inbound traffic [13], [14], [17].

In this paper, we focus on the active watermarking technique, which has been active in the past few years. For example, Yu *et al.*

Manuscript received May 29, 2011; accepted November 05, 2011; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor M. Allman. Date of publication January 16, 2012; date of current version August 14, 2012. This work was supported in part by the National Key Basic Research Program of China (973 Program) under Grants 2010CB328104 and 2011CB302800; the National Science Foundation of China (NSFC) under Grants 60903162, 60903161, 61070158, 61070161, 61003257, 61070221, and 61070222/F020802; the US National Science Foundation (NSF) under Grants CNS0916584, CNS1065136, and CNS-1117175; CityU Applied R&D Funding (ARD) under Grants 9681001, 6351006, and 9667052; CityU Strategic Research Grant 7008110; ShenZhen-HK Innovation Cycle Grant ZYB200907080078A; the China Specialized Research Fund for the Doctoral Program of Higher Education under Grant 200802860031; Jiangsu Provincial Natural Science Foundation of China under Grant BK2008030; Jiangsu Provincial Key Laboratory of Network and Information Security under Grant BM2003201; and the Key Laboratory of Computer Network and Information Integration of Ministry of Education of China under Grant 93K-9. Any opinions, findings, conclusions, and recommendations in this paper are those of the authors and do not necessarily reflect the views of the funding agencies. The conference version of this paper was published in the Proceedings of the 16th ACM Conference on Computer and Communications Security (CCS), Chicago, IL, November 9–13, 2009.

Z. Ling and J. Luo are with the School of Computer Science and Engineering, Southeast University, Nanjing 210096, China (e-mail: zhenling@seu.edu.cn; jluo@seu.edu.cn).

W. Yu is with the Department of Computer and Information Sciences, Towson University, Towson, MD 21252 USA (e-mail: wyu@towson.edu).

X. Fu is with the Department of Computer Science, University of Massachusetts Lowell, Lowell, MA 01854 (e-mail: xinwenfu@cs.uml.edu).

D. Xuan is with the Department of Computer Science and Engineering, The Ohio State University, Columbus, OH 43210 USA (e-mail: xuan@cse.ohio-state.edu).

W. Jia is with the Department of Computer Science, City University of Hong Kong, Kowloon, Hong Kong (e-mail: wei.jia@cityu.edu.hk).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TNET.2011.2178036

can recognize the corresponding communication relationship, tracing the messages despite the use of anonymous networks. However, in order to accurately confirm the anonymous communication relationship of users, the flow-marking scheme needs to embed a signal modulated by a relatively long length of PN code, and also the signal is embedded into the traffic flow rate variation. Houmansadr *et al.* [14] proposed a nonblind network flow watermarking scheme called RAINBOW for stepping stone detection. Their approach records the traffic timing of the incoming flows and correlates them with the outgoing flows. This approach also embeds watermarks into the traffic by actively delaying some packets. The watermark detection problem was formalized as detecting a known spread-spectrum signal with noise caused by network dynamics. Normalized correlation is used as the detection scheme. Their approach can classify a typical SSH connection as a stepping stone connection in 3 min. As we can see, it is hard for the flow-marking technique to deal with the short communication sessions that may only last for a few seconds.

A successful attack against anonymous communication systems relies on accuracy, efficiency, and detectability of active watermarking techniques. Detectability refers to the difficulty of detecting the embedded signal by anyone other than the attackers. Efficiency refers to the quickness of confirming anonymous communication relationships among users. Although accuracy and/or detectability have received great attention [13], [14], [17], to the best of our knowledge, no existing work can meet all these three requirements simultaneously.

In this paper, we investigate a new cell-counting-based attack against Tor, a real-world, circuit-based low-latency anonymous communication network. This attack is a novel variation of the standard timing attack. It can confirm anonymous communication relationship among users accurately and quickly and is difficult to detect. In this attack, the attacker at the malicious exit router detects the data transmitted to a suspicious destination (e.g., server Bob). The attacker then determines whether the data is a *relay* cell or a *control* cell in Tor. After excluding the *control* cells, the attacker manipulates the number of *relay* cells in the circuit queue and flushes out all cells in the circuit queue. In this way, the attacker can embed a signal (a series of “1” or “0” bits) into the variation of the cell count during a short period in the target traffic. An accomplice of the attacker at the entry onion router detects and excludes the *control* cells, records the number of *relay* cells in the circuit queue, and recovers the embedded signal. The signal embedded in the target traffic might be distorted because the cells carrying the different bits (units) of the original signal might be combined or separated at middle onion routers. To address this problem, we develop the recovery algorithms to accurately recognize the embedded signal. Our theoretical analysis shows that the detection rate is a monotonously increasing function with respect to the delay interval and is a monotonously decreasing function of the variance of one way transmission delay along a circuit. In our real-world experiments, the experimental results match the theoretical results well. To be specific, our attack needs only 2 s to achieve a true positive rate of almost 100% and the false positive rate of almost 0%.

We have implemented the cell-counting-based attack against Tor and performed a set of real-world Internet experiments to validate the feasibility and effectiveness of the attack. The attack




Fig. 1. Tor network.

presented in this paper is one of the first to exploit the implementation of known anonymous communication systems such as Tor by exploiting its fundamental protocol design. There are several unique features for this attack. First, this attack is highly efficient and can quickly confirm very short anonymous communication sessions with tens of cells. Second, this attack is effective, and its detection rate approaches 100% with very low false positive rate. Third, the short and secret signal makes it difficult for others to detect the presence of the embedded signal. Our time-hopping-based signal embedding technique makes the attack even harder to detect. The attack poses a significant threat to the anonymity provided by Tor because the attack can confirm over half of communication sessions by injecting around 10% malicious onion routes on Tor [18], [19].

The remainder of this paper is organized as follows: We introduce the background in Section II. We present the cell-counting-based attack, including the basic idea, issues of the attack, and solutions, in Section III. In Section IV, we discuss various issues, including some extension, and the detectability and impact of the proposed attack. In Section V, we analyze the effectiveness of the attack. In Section VI, we show experimental results on Tor and validate our findings. We review related work in Section VII and conclude this paper in Section VIII.

## II. BACKGROUND

In this section, we first overview the components of Tor. We then present the procedures of how to create circuits and




Fig. 2. Cell format by Tor. (a) Tor cell format. (b) Tor relay cell format.

hard-coded into the Tor source code for a client to download the information of onion routers and build circuits through the Tor network.

Fig. 2 illustrates the cell format used by Tor. All cells have a 3-B header, which is not encrypted in the onion-like fashion so that the intermediate Tor routers can see this header. The other 509 B are encrypted in the onion-like fashion. There are two types of cells: *control* cell shown in Fig. 2(a) and *relay* cell shown in Fig. 2(b). The command field (*Command*) of a control cell can be: *CELL\_PADDING*, used for keepalive and optionally usable for link padding, although not used currently; *CELL\_CREATE* or *CELL\_CREATED*, used for setting up a new circuit; and *CELL\_DESTROY*, used for releasing a circuit. The command field (*Command*) of a relay cell is *CELL\_RELAY*. Note that relay cells are used to carry TCP stream data from Alice to Bob. The relay cell has an additional header, namely the relay header. There are numerous types of relay commands (*Relay Command*), including *RELAY\_COMMAND\_BEGIN*, *RELAY\_COMMAND\_DATA*, *RELAY\_COMMAND\_END*, *RELAY\_COMMAND\_SENDME*, *RELAY\_COMMAND\_EXTEND*, *RELAY\_COMMAND\_DROP*, and *RELAY\_COMMAND\_RESOLVE*. Note that all these can be found in *or.h* in released source code package by Tor.

### B. Circuit Creation and Data Transmission

In Tor, an *OR* maintains a TLS connection to other *ORs* or *OPs* on demand. The *OP* uses a way of source routing and chooses several *ORs* (preferably ones with high bandwidth and high uptime) from the locally cached directory, downloaded from the directory caches. The number of the selected *ORs* is referred as the path length. We use the default path length of three as an example. The *OP* iteratively establishes circuits across the Tor network and negotiates a symmetric key with each *OR*, one hop at a time, as well as handles the TCP streams from client applications. The *OR* on the other side of the circuit connects to the requested destinations and relays the data.

We now illustrate the procedure that the *OP* establishes a circuit and downloads a file from the server. *OP* first sets up a TLS connection with *OR1* using the TLS protocol. Then, tunneling through this connection, *OP* sends a *CELL\_CREATE* cell and uses the *Diffie-Hellman* (DH) handshake protocol to negotiate a base key  $K_1 = g^{xy}$  with *OR1*, which responds with a *CELL\_CREATED* cell. From this base key material, a forward symmetric key  $k_{f1}$  and a backward symmetric key  $k_{b1}$  are produced [21]. In this way, a 1-hop circuit *C1* is created. Similarly, *OP* extends the circuit to a 2-hop circuit and 3-hop circuit. After the circuit is set up between the *OP* and *OR3*, *OP* sends a *RELAY\_COMMAND\_BEGIN* cell to the exit onion router, and the cell is encrypted as  $\{\{\{Begin < IP, Port >\}_{k_{f3}}\}_{k_{f2}}\}_{k_{f1}}$ , where the subscript refers to the key used for encryption of one

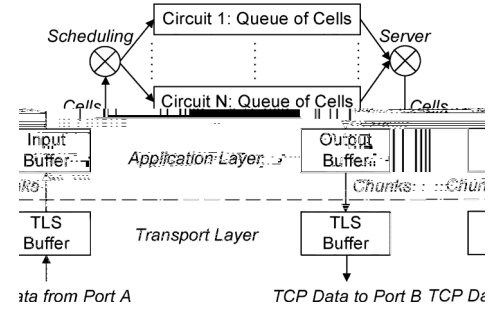


Fig. 3. Processing the cells at onion routers.

onion skin. The three layers of onion skin are removed one by one each time the cell traverses an onion router through the circuit. When *OR3* removes the last onion skin by decryption, it recognizes that the request intends to open a TCP stream to a port at the destination IP, which belongs to Bob. Therefore, *OR3* acts as a proxy, sets up a TCP connection with Bob, and sends a *RELAY\_COMMAND\_CONNECTED* cell back to Alice's *OP*. Then, Alice can download the file.

### C. Processing Cells at Onion Routers

Fig. 3 illustrates the procedure of processing cells at onion routers. Note that the cells mentioned below are all *CELL\_RELAY\_DATA* cells, which are used to carry end-to-end stream data between Alice and Bob. To begin with, the onion router receives the TCP data from the connection on the given port A. After the data is processed by TCP and TLS protocols, the data will be delivered into the TLS buffer of the connection. When there is pending data in the TLS buffer, the read event of this connection will be called to read and process the data. The connection read event will pull the data from the TLS buffer into the connection input buffer. Each connection input buffer is implemented as a linked list with small chunks. The data is fetched from the head of the list and added to the tail. After the data in the TLS buffer is pulled into the connection input buffer, the connection read event will process the cells from the connection input buffer one by one. As stated earlier, the cell size is 512 B. Thus, 512-B data will be pulled out from the input buffer every time until the data remaining in the connection input buffer is smaller than 512 B. Since each onion router has a routing table that maintains the map from source connection and circuit ID to destination connection and circuit ID, the read event can determine that the transmission direction of the cell is either in the forward or backward direction. Then, the corresponding symmetric key is used to decrypt/encrypt the payload of the cell, replace the present circuit ID with the destination circuit ID, and append the cell to the destination circuit queue. If it is the first cell added to this circuit queue, the circuit will be made active by being added into a double-linked ring of circuits with queued cells waiting for a room to free up on the output buffer of the destination connection. Then, if there is no data waiting in the output buffer for the destination connection, the cell will be written into the output buffer directly, and then the write event of this circuit is added to the event queue. Subsequent incoming cells are queued in the circuit queue.

When the write event of the circuit is called, the data in the output buffer is pushed to the TLS buffer of the destination connection. Then, the write event will pull as many cells as possible from the circuit queue of the currently active circuit to the output




Fig. 4. Packet sequence versus packet size.




Fig. 5. Number of packets versus packet size.

buffer and add the write event of this circuit to the event queue. The next write event can carry on pushing data to the output buffer and pull the cells to the output buffer. In other words, the cells queued in the circuit queue can be delivered to the network via port *B* by calling the write event twice.

### III. CELL-COUNTING-BASED ATTACK

In this section, we first show that the size of IP packets in the Tor network is very dynamic. Based on this finding, we then introduce the basic idea of the cell-counting-based attack and list some challenging issues related to the attack and present solutions to resolve those issues.

#### A. Dynamic IP Packet Size of Traffic Over Tor

In Tor, the application data will be packed into equal-sized cells (e.g., 512 B). Nonetheless, via extensive experiments over the Tor network, we found that the size of IP packets transmitted over Tor is dynamic. Fig. 4 shows the size of received IP packets at the client over time, and Fig. 5 shows the frequency of the IP packet size. It can be observed that the size of packets from the sender to the receiver is random over time, and a large number of packets have varied sizes, other than the cell size or maximum transmission unit (MTU) size.

These observations can be reasoned as follows.

- 1) The varied performance of onion routers may cause cells not to be promptly processed. According to cell processing in Fig. 3, if an onion router is overloaded, unprocessed cells will be queued. Therefore, cells will be merged at the IP layer and sent out together. Those merged cells may be split into multiple MTU-sized packets and one non-MTU-sized packet.




Fig. 6. Cell-counting-based attack.

- 2) Tor network dynamics may incur those non-MTU-sized IP packets as well. If the network between onion routers is congested, cells will not be delivered on time. When this happens, cells will merge, and non-MTU-sized IP packets will show up.

#### B. Basic Idea of Cell-Counting-Based Attack

As we stated above, the packet size observed at the client shows a high probability to be random because of the performance of onion routers and Internet traffic dynamics. Motivated by this finding, we investigate a new cell-counting-based attack against Tor, which allows the attacker to confirm anonymous communication relationship among users very quickly. In addition, it will be hard for the client to detect our developed attack described in what follows.

As we mentioned before, this attack intends to confirm that Alice (client) communicates with Bob (server) over Tor. In order to do so, we assume that the attacker controls a small percentage of exit and entry onion routers by donating computers to Tor. This assumption is also used in other studies [3], [10], [18], [19]. The assumption is valid since Tor is operated in a voluntary manner [21]. For example, attackers may purchase Amazon EC2 virtual machines, which can be put into Tor. The attack can be initiated at either the malicious entry onion router or exit onion router, up to the interest of the attacker. In the rest of the paper, we assume that the attack is initiated at an exit onion router connected to server Bob and intends to confirm that Alice communicates with a known server Bob.

The basic idea is as follows. An attacker at the exit onion router first selects the target traffic flow between Alice and Bob. The attacker then selects a random signal (e.g., a sequence of binary bits), chooses an appropriate time, and changes the cell count of target traffic based on the selected random signal. In this way, the attacker is able to embed a signal into the target traffic from Bob. The signal will be carried along with the target traffic to the entry onion router connecting to Alice. An accomplice of the attacker at the entry onion router will record the variation of the received cells and recognize the embedded signal. If the same pattern of the signal is recognized, the attacker confirms the communication relationship between Alice and Bob.

As shown in Fig. 6, the workflow of the cell-counting-based attack is illustrated as follows.

*Step 1: Selecting the Target:* At a malicious exit onion router connected to the server Bob, the attacker will log the information, including server Bob's host IP address and port used for a given circuit, as well as the circuit ID. The attacker uses *CELL\_RELAY\_DATA* cells since those cells transmit the data stream. According to the description of Tor in Section II, we know that the attacker is able to obtain the

rst cell backward to the client, which is a *CELL\_CREATED* cell and is used to negotiate a symmetric key with the middle onion router. The second cell backward to the client will be a *CELL\_RELAY\_CONNECTED* cell. All sequential cells will be *CELL\_RELAY\_DATA* cell, and the attacker starts the encoding process shown in Step 2.

*Step 2: Encoding the Signal:* In Section II, we introduced the procedure of processing cells at the onion routers. The *CELL\_RELAY\_DATA* cells will be waiting in the circuit queue of the onion router until the write event is called. Then, the cells in the circuit queue are all

output buffer if the output buffer is empty. The write event is added to the event queue, and the cell waits to be written to the network by the write event. Since the interval is small, the three cells for the second bit “1” and the cell for the third bit “0” also arrive at the middle onion router and stay in the circuit queue. When the write event is called, the first cell for carrying the first bit “0” will be written to the network, while the following three cells for carrying the second bit of the signal and one cell for carrying the third bit of the signal will be written to the output buffer all together. When this happens, the original signal will be distorted (i.e., the third bit “0” of the signal will be lost). Therefore, the attacker needs to choose the proper delay interval for transmitting cells. In addition, we will discuss the types of the division and combination of the cells with details in Section III-C.2.

We now check conditions that preserve units of the signal during transmission. Let  $S = \{S_0,$








Fig. 9. Probability that a circuit chooses the malicious routers as entry and exit routers versus number of malicious Tor routers [19].




Fig. 10. Packet format. (a) Packet format of 1 cell. (b) Packet format of 2 cell. (c) Packet format of 3 cell.

1	2	2
Content Types	Versions	Length

Fig. 11. TLS header.

that it selects Tor routers with low bandwidth for the circuits is around 90%. Obviously, the low-bandwidth Tor router will be the bottleneck of the circuit.

### B. Controlling Exit Onion Routers Only

If the attacker does not control entry onion routers, the cell-counting-based attack can still be successful. An attacker can sniff the packets transmitted between an entry onion router and a client. The attacker may recover the embedded signal based on the size of the packet. In this way, the number of required malicious routers in Tor can also be reduced while the attack still has a desired impact.

We now introduce the structure of the IP packet that envelops the cell(s) and passes along the network. Without loss of generality, we assume that MTU is 1500 B. Fig. 10(a) illustrates the structure of IP packet that envelops one cell, including an IP header, a TCP header, an empty TLS application record, and a TLS application record of enveloping one cell. The TLS record packet incorporates a TLS header (5 B), a TLS message (not to exceed  $2^{14}$  B), a MAC (Message Authentication Code, 20 B), and a TLS padding (12 B). Fig. 11 illustrates the header of the TLS packet, with the length of 5 B. The field of content type identifies the record-layer protocol type contained in this record, with the length of 1 B. In our case, we are concerned with the




Fig. 12. Time-hopping technique.

TLS application record, with content type of 23. The field of version identifies the major or minor version of TLS for the contained message, with the length of 2 B. The field of length identifies the length of protocol message(s), not to exceed  $2^{14}$  B.

Fig. 10(b) illustrates the structure of IP packet that envelops two cells and has a length of 1150 B. Because an IP packet that envelops three cells exceeds the MTU (1500 B), this IP packet will be segmented; one segment has the packet of 1500 B, and the other segment has the packet of 214 B. Fig. 10(c) illustrates the structure of IP packet that envelops three cells and is segmented. Hence, the attacker can map “0” bit of the signal to one IP packet, with the length of 638 B. By appropriately choosing a delay interval at the exit onion router, the “1” bit of the signal will have two cases: two IP packets with one cell [shown in Fig. 10(a)] and two cells [shown in Fig. 10(b)], i.e., the signal is divided as Type I with  $k = 1$ , as well as two IP packets with three cells [shown in Fig. 10(c)], which is neither divided nor combined. Therefore, from packet size pattern, the attacker is still able to recognize the signal embedded in the IP packet stream by using our signal detection mechanism. Actually, the fact that multiple cells can be packed into a packet guarantees the correct signal encoding via the variation of the cell count. When such a packet arrives at the TLS buffer, those cells form a group, which is read into the circuit queue. This is our mechanism that generates a signal bit “1” or “0.”

### C. Attack Detectability

The proposed cell-counting-based attack is difficult to detect. As we know, the attack transmits a short and secret random signal known only to the attackers. It is difficult to detect within the target traffic. Based on the evaluation data shown in Section VI, the success of this attack requires only a short secret signal—such as 5 b—while achieving a detection rate of almost 100% and a false positive rate of  $1/10^5$ . It would be hard to classify such a short sequence of random signals as the attack sequence in bursty network traffic.

To further improve the attack invisibility, we adopt the time-hopping-based signal embedding technique, which can greatly reduce the probability of interception and recognition [27]. Fig. 12 illustrates the principle of the time-hopping technique. For the time hopping, there exist random intervals between signal bits. At the exit onion router, the duration of those intervals are varied according to a pseudorandom control code, which is known to only the attackers. To recover the signal at the entry onion router, an accomplice of the attacker can use the same secret control code to help position the signal bits and recover the whole signal. Intuitively, if the interval between the bits is large enough, the inserted signal bits appear sparse within the target traffic, and it is difficult to determine whether groups of cells are caused by network dynamics or by intention. Therefore, the secret signal embedded into the target traffic is no different than the noise. In addition, when a malicious entry node has confirmed the communication relationship, it can separate the group of cells by adding delay between the cells so that not even the client can observe the



embedded signal. In Section VI, we demonstrate the effectiveness of this time-hopping-based technique, and the detailed approach is shown in Algorithm 2 in Appendix A.

In our proposed attack, a secret signal is embedded into the target traffic, which implies a secret sequence of groups of one and three cells. One may be concerned that if the sequence of groups of one and three cells is unnatural and the entry node is honest and aware of the attack, it will detect the sequence and thus distinguish the traffic flow with an embedded signal from a flow without a signal. However, with the time-hopping technique, groups of one and three cells are separated by random intervals, and it is hard to differentiate them from those caused by network dynamics. As a side note, the false positives in detecting signal bits in Section VI's figures imply that normal network traffic does have groups of one and three cells caused by network dynamics. In addition, since the embedded signal is very short and only known to attackers, we conjecture that it is very difficult to distinguish traffic with embedded signals from normal traffic based on this very short secret sequence of cell groups.

#### D. Difference From Existing Attacks

The proposed cell-counting-based attack may dramatically degrade anonymity that Tor maintains. Different from other existing attacks, the cell-counting-based attack is accurate, efficient, and difficult to detect. This attack requires much fewer packets and incurs little overhead while achieving a higher detection rate than most traffic analysis attacks, including traffic correlation attacks in [10], [13], [17], [28], and [29]. Since this attack utilizes the atomic unit of a traffic flow, i.e., cells/packets (and their size), this attack is highly efficient and can confirm very short communication sessions with only tens of cells. Although the tagging-based attack [19], [30] may require few packets, it tears down the Tor circuits and is relatively easy to detect. A simple passive cell-counting attack may count the cells at points of exit and entry onion routers and correlate the counting. However, there is no guarantee of detection rate and false positive rate because of the large number of connections running through Tor. In addition, our attack achieves a low false positive rate with a very small amount of target traffic as demonstrated in Section VI. Therefore, as a powerful traffic correlation attack, the proposed attack poses a great challenge against Tor.

#### E. Countermeasures

We now discuss possible countermeasures. It is also difficult for Tor to defeat the cell-counting attack. One possible countermeasure is that Tor routers add delay between cells in order to disrupt malicious cell groups. However, choosing such a delay will be very challenging. A too short delay cannot separate cells (at the network layer), while a long delay may dramatically degrade Tor's performance, which is already the biggest bottleneck of using Tor [22], [26], [31]. A second way to reduce the impact of the proposed attack is to use purely random routing algorithms and reduce the chance of traffic flows passing malicious Tor onion routers. However, such a random routing algorithm will also degrade Tor performance. Its effect is also very limited since the attacker can inject more malicious routers into Tor to increase the impact.

Dummy traffic may be used to distort the timing of the signal. A constant rate padding along a circuit may incur too

much overhead. Levine *et al.* [8] investigated a defensive dropping scheme, in which dummy traffic can be randomly dropped at the intermediated routers. An end-to-end defensive dropping cannot be applied to Tor directly. Tor adopts Advanced Encryption Standard Counter Mode (AES-CTR) to encrypt the cells. The AES counter at each onion router and onion proxy is synchronized. Defensive dropping will disrupt this AES counter and cause decryption errors at the onion proxy or the exit routers [30]. These errors will tear down the circuits. Shmatikov *et al.* [15] proposed an adaptive padding scheme by injecting dummy packets into statistically unlikely gaps in the packet flow, destroying timing fingerprints without adding any latency to the application traffic. However, in our case, the attacker controls the exit router, and the signal can be embedded in the dummy traffic as well.

This paper provides guidance to anonymous protocol design and implementation. To design an anonymous communication system, we have to consider the impact of the design on all protocol layers. For example, Tor implements an overlay protocol and preserves equal-sized cells on the application layer. However, the equal-sized cells on the application layer cannot guarantee that packets on the network layer are also equal-sized. Hence, the equal-sized cells on the application layer cannot guarantee the anonymity provided by Tor. Indeed, our attack exploits the Tor protocol's impact on the network layer.

### V. ANALYSIS

In this section, we show the analytical results for the accuracy and efficiency of the cell-counting-based attack. For attack accuracy, we derive closed formulas for detection rate and false positive rate. Our theoretical analysis shows that the detection rate is a monotonously increasing function with respect to the delay interval and is a monotonously decreasing function of the variance of one way transmission delay along a circuit. Our experimental results in Section VI match the theoretical results well.

#### A. Detection Rate

We view that the major factor causing detection error is network dynamics, which leads to combination and division of cell groups. Our analysis is based on the network configuration described in the second paragraph of Section II-B. The round-trip delay between two onion routers can be modeled by a log-normal distribution [32]. We first investigate the properties of the log-normal distribution and then use the delay model to derive detection rate analytically.

A log-normal random variable has the property that its logarithm has a Gaussian distribution. Let  $X_i$  be a Gaussian random variable with the probability density function (PDF), we have(9)where

$$f_{x_i}(x) = \frac{1}{\sigma_{x_i} \sqrt{2\pi}} e^{-\frac{(x - \mu_{x_i})^2}{2\sigma_{x_i}^2}}$$




Fig. 13. One-way trip time probability density function. (a) Germany.

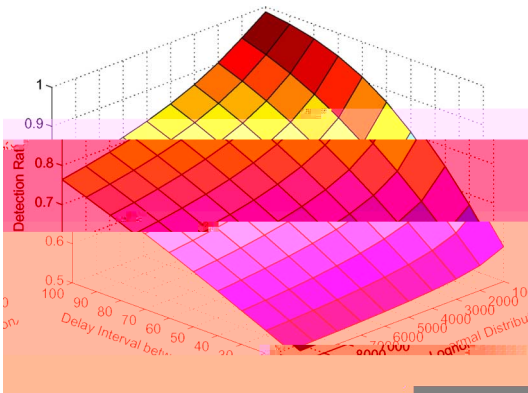


Fig. 14. Theoretical detection rate versus delay interval and variance of the log-normal distribution.

Detection rate  $P_{D,n}$  is defined as the detection rate for an  $n$ -bit original signal. Given




Fig. 16. Channel model.

Based on the channel model shown in Fig. 16, we know that  $\max I(X, Y)$  can be derived by (28)

$$\begin{aligned} \max_p I(X, Y) \\ &= \max_p H(X) - H(X|Y) \end{aligned}$$




Fig. 18. Detection rate versus delay interval (Note: The rate is for detecting one bit).

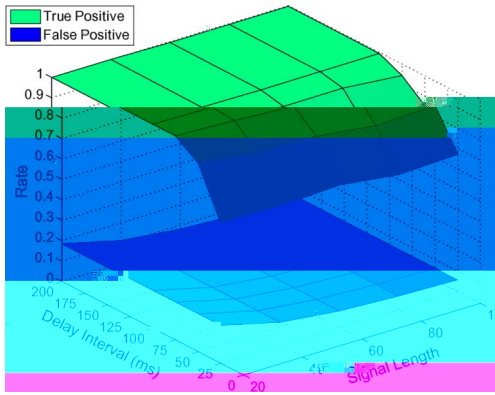


Fig. 19. Detection rate versus delay interval and signal length with detection scheme 1 (Note: The rate is for detecting one bit).

as detection scheme 1. Moreover, we chose to recover all types with  $k = 2$  as detection scheme 2.

When we evaluate the false positive rate, the client downloads 30 files via Tor again. However, no signal is embedded into the traffic at the exit onion router. Denote the traffic with no signal as clean traffic. We generate a 100-b random signal and apply detection schemes 1 and 2 to the clean traffic collected at the entry onion router. By checking how many bits of this signal show up in the clean traffic, we can calculate the false positive rate.

We conduct the above experiment to evaluate the true positive and false positive by using a 100-b random signal. Fig. 18 illustrates the correlation between the detection rate (true positive) and the delay interval for transmitting cells associated to different units of the signal. As we can see from this figure, the detection rate will increase dramatically when the delay interval is slightly increased in two detection schemes. As expected, the detection rate of scheme 2 is higher than scheme 1 with a slightly increasing false positive rate, while the overall false positive rate for each scheme is a fixed value. When the delay interval approaches 100 ms, the detection rate of two schemes approaches 100%. All these findings validate that our investigated attack can significantly degrade the anonymity service provided by Tor.

Fig. 19 illustrates the detection rate in terms of signal length and the delay interval for scheme 1. Note that the detection rate

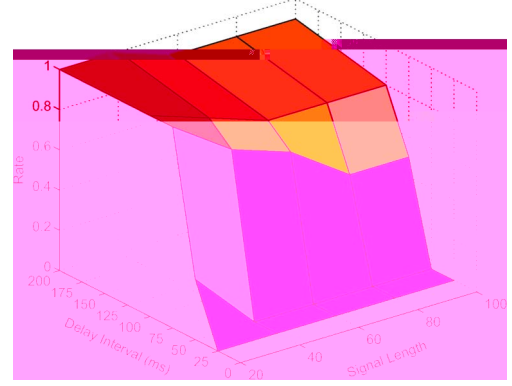


Fig. 20. Detection rate versus delay interval and signal length with detection scheme 1.

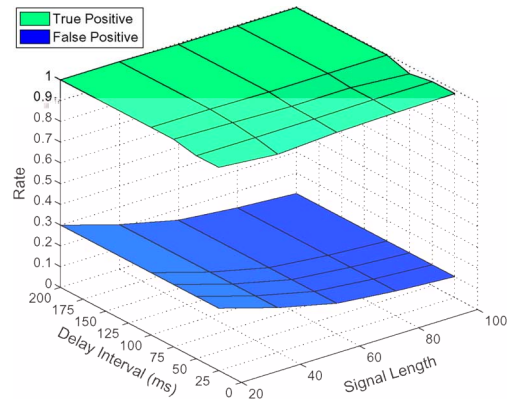


Fig. 21. Detection rate versus delay interval and signal length with detection scheme 2 (Note: The rate is for detecting one bit).

in Fig. 19 is for detecting one bit. As we can see from this figure, when we increase the signal length from 20 to 100, the detection rate will be slightly decreased, and the false positive rate will be constantly very low (less than 5%). When the signal length is 20, and the delay interval between signals is 100 ms, 100% detection rate can be achieved. In addition, Fig. 20 illustrates the detection rate for detecting the whole signal in terms of signal length and the delay interval for scheme 1. When the signal length is 20, and the delay interval between the signals is 100 ms, a detection rate of 100% can be achieved. In Fig. 20, the false positive approaches 0%, and this matches our theoretical analysis in Section V-B. This validates that the investigated




Fig. 22. Detection rate versus delay interval and signal length with detection scheme 2.

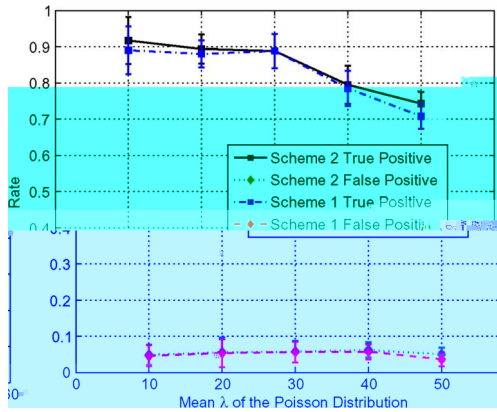


Fig. 23. Correlation between detection rate and mean  $\lambda$  of the Poisson distribution (Note: The rate is for detecting one bit).

embeds units of a signal into the target traffic, as introduced in Section IV-B. For this encoding scheme, we generate an array  $Q[1 \times n]$  by using a Poisson distribution with a mean  $\lambda$ . We first send the number of cells  $Q[i]$  without embedding signals, and then embed a signal bit. In this set of experiments, we also chose a signal length of 100. Since the units of the signal are embedded randomly in a hopping fashion in the time domain, it is hard for the multi-flow attack [37] to detect the embedded signal in the traffic. Fig. 23 illustrates the relationship between detection rate (true positive) and the mean  $\lambda$  of nonwatermarked cells (which corresponds to the random time interval. No signal is embedded into those cells). From Fig. 23, we can see that this improved encoding scheme can still achieve very high detection rates along with a very low false positive rate. Since this new encoding scheme does not embed the signal into all *CELL\_RELAY\_DATA* cells, the attack will require more cells in order to be successful. Additionally, based on Algorithm 1, we use Algorithm 2 to recognize the signal embedded in the sparsely encoded cells.

We also use *tcpdump* to capture the IP packets transmitted between the entry node and the client and demonstrate that an attacker may also use packet size to recognize the embedded signal. Fig. 24 illustrates the variance of IP packet size. As we can see, there are three types of IP packet sizes, and the corresponding packet structures are shown in Fig. 10. According to detection scheme 1 (blind detection approach), we can map bit “0” of the signal to the IP packet size of 638 B. Bit “1”




Fig. 24. Variance of IP packet size.

of the signal has two cases: IP packet of 638 B followed by IP packet of 1150 B, as well as one IP packet of 1500 B and of 214 B, as we discussed in Section IV-B. Therefore, we can decode the signal between 0 and 1 s as “0010100.” Note that since the delay interval is very small among the second (638 B), the third (1500 B), and the fourth (214 B) IP packets, they are mostly overlapped in the figure. As we know, the packet drops will incur TCP retransmissions, and it may result in a distorted signal. One way to address the packet drop issue is to use *tcpdump* and check the TCP sequence numbers for assisting the signal recovery. A second way is to increase the delay interval between the signal bits and reduce the impact of packet drops. As illustrated in Fig. 24, we know that the attacker is able to recognize the signal based on the size of sniffed IP packets using the signal detection mechanism discussed in Section IV-B in addition to using the cell count. In our recent work, we proposed a packet-size-based attack [38] that compromises Tor’s communication anonymity with no need of controlling Tor routers. An attacker can manipulate size of packets between a Web site and an exit onion router and embeds a signal into the target traffic. An accomplice at the user side can sniff the traffic and recognize this signal. If the victim traffic is marked by our signal four times, the detection rate approaches over 90% with the delay interval of 400 ms, and the false positive rate can be suppressed to less than 4%.

## VII. RELATED WORK

A good review of mix systems can be found in [4] and [5]. There has been much research on degrading anonymous communication through mix networks. Existing traffic analysis attacks against anonymous communication can largely be categorized into two groups: passive traffic analysis and active watermarking techniques. Passive traffic analysis techniques have shown that the attacks record the traffic passively and identify the similarity between server’s outbound traffic and client’s inbound traffic [8], [9]. Other recent research works have shown that the attackers can infer sensitive information from the encrypted network traffic by examining patterns in terms of the sizes of packet and its timing [1], [39]–[41]. For example, Liberatore and Levine [40] examined the packet sizes of HTTP traffic transmitted over persistent connection or tunneled via SSH port forwarding can statistically identify the Web pages. Wright *et al.* [41] investigated the statistical distribution of packet sizes in encrypted Voice over IP (VoIP) connections



and identified the language spoken based on the distribution in each conversation. Later work of Wright *et al.* [42] also investigated how an eavesdropper could identify spoken phrases in encrypted VoIP.

The active watermarking techniques intend to embed specific secret signal (or marks) into the target traffic [10], [13], [17], [43]. Such techniques can reduce the false positive rate significantly if the signal is long enough and does not require massive training study of traffic cross correlation as required in passive traffic analysis. For example, Yu *et al.* [13] proposed a flow-marking scheme based on the DSSS technique. This approach could be used by attackers to secretly confirm the communication relationship via mix networks. Øverlier *et al.* [3] studied a scheme using one compromised mix router to identify the “hidden server” anonymized by Tor. Wang *et al.* [17] also investigated the feasibility of a timing-based watermarking scheme in identifying the encrypted peer-to-peer VoIP calls. Peng *et al.* [44] analyzed the secrecy of timing-based watermarking traceback proposed in [43], based on the distribution of traffic timing. Kiyavash *et al.* [37] proposed a multi flow approach detecting the interval-based watermarks [12], [45] and DSSS-based watermarks [13]. This multi flow-based approach intends to average the rate of multiple synchronized watermarked flows and expects to observe a unusual long silence period without packets or a unusual long period of low-rate traffic.

## VIII. CONCLUSION

In this paper, we introduced a novel cell-counting-based attack against Tor. This attack is difficult to detect and is able to quickly and accurately confirm the anonymous communication relationship among users on Tor. An attacker at the malicious exit onion router slightly manipulates the transmission of cells from a target TCP stream and embeds a secret signal (a series of binary bits) into the cell counter variation of the TCP stream. An accomplice of the attacker at the entry onion router recognizes the embedded signal using our developed recovery algorithms and links the communication relationship among users. Our theoretical analysis shows that the detection rate is a monotonously increasing function with respect to the delay interval and is a monotonously decreasing function of the variance of one way transmission delay along a circuit. Via extensive real-world experiments on Tor, the effectiveness and feasibility of the attack is validated. Our data showed that this attack could drastically and quickly degrade the anonymity service that Tor provides. Due to Tor’s fundamental design, defending against this attack remains a very challenging task that we will investigate in our future research.

## APPENDIX A

Algorithm 1 shows the signal recovery mechanism with continuously embedded bits at a malicious Tor entry node. Algorithm 2 gives the signal recovery mechanism at a malicious Tor entry node when the time-hopping-based approach is used for embedding a signal into the target traffic.

---

### Algorithm 1: Recovery Mechanism for Continuously Embedded Bits

---

#### Require:

(a)

```

6:   end while
7:   if  $Q[j] == C[i]$  then
8:      $i = i + 1$ ;  $Q[j]$  is removed.
9:     Detect  $S[j]$  with  $C[i]$  by using Algorithm 1
10:  else if  $Q[j] < C[i]$  then
11:    The signal  $S[j]$  is combined with  $Q[j]$ .
12:     $C[i] = C[i] - Q[j]$ 
13:    Detect  $S[j]$  with  $C[i]$  by using Algorithm 1
14:  end if
15:   $i = i + 1$ ;  $j = j + 1$ 
16: end while

```

## APPENDIX B

We now prove the fact that if  $L$  is a random variable with a log-normal distribution,  $M = L/2$  also has a log-normal distribution. Let  $M$  be the random variable of one way delay through the circuit and  $M$  is approximately  $L/2$ , where  $L$  is a round-trip delay along the circuit. Since the CDF of one-way delay  $M$  is derived by

$$F_M(m) = P(M < m) \quad (34)$$

$$= P(L/2 < m) \quad (35)$$

$$= P(L < 2m) \quad (36)$$

$$= F_L(2m) \quad (37)$$

$$= \int_0^{2m} f(l) dl \quad (38)$$

the PDF of  $M$  can be derived by

$$f_M(m) = F'_M(m) \quad (39)$$

$$= F'_L(2m) \quad (40)$$

$$= 2f_L(2m) \quad (41)$$

$$= 2 \frac{1}{2m\sigma_{x_i}\sqrt{2\pi}} e^{-\frac{(ln(2m) - \mu_{x_i})^2}{2\sigma_{x_i}^2}} \quad (42)$$


$$= \frac{1}{m\sigma_{x_i}\sqrt{2\pi}} e^{-\frac{(ln(m) - (\mu_{x_i} - ln 2))^2}{2\sigma_{x_i}^2}}. \quad (43)$$

As we can see, the PDF of  $M$  follows a log-normal distribution as well.

## REFERENCES

- [1] Q. X. Sun, D. R. Simon, Y. Wang, W. Russell, V. N. Padmanabhan, and L. L. Qiu, "Statistical identification of encrypted Web browsing traffic," in *Proc. IEEE S&P*, May 2002, pp. 19–30.
- [2] X. Fu, Y. Zhu, B. Graham, R. Bettati, and W. Zhao, "On flow marking attacks in wireless anonymous communication networks," in *Proc. IEEE ICDCS*, Apr. 2005, pp. 493–503.
- [3] L. Øverlier and P. Syverson, "Locating hidden servers," in *Proc. IEEE S&P*, May 2006, pp. 100–114.
- [4] G. Danezis, R. Dingleline, and N. Mathewson, "Mixminion: Design of a type III anonymous remailer protocol," in *Proc. IEEE S&P*, May 2003, pp. 2–15.
- [5] R. Dingleline, N. Mathewson, and P. Syverson, "Tor: The second-generation onion router," in *Proc. 13th USENIX Security Symp.*, Aug. 2004, p. 21.
- [6] "Anonymizer, Inc.," 2009 [Online]. Available: <http://www.anonymizer.com/>
- [7] A. Serjantov and P. Sewell, "Passive attack analysis for connection-based anonymity systems," in *Proc. ESORICS*, Oct. 2003, pp. 116–131.
- [8] B. N. Levine, M. K. Reiter, C. Wang, and M. Wright, "Timing attacks in low-latency MIX systems," in *Proc. FC*, Feb. 2004, pp. 251–265.
- [9] Y. Zhu, X. Fu, B. Graham, R. Bettati, and W. Zhao, "On flow correlation attacks and countermeasures in Mix networks," in *Proc. PET*, May 2004, pp. 735–742.
- [10] S. J. Murdoch and G. Danezis, "Low-cost traffic analysis of Tor," in *Proc. IEEE S&P*, May 2006, pp. 183–195.
- [11] K. Bauer, D. McCoy, D. Grunwald, T. Kohno, and D. Sicker, "Low-resource routing attacks against anonymous systems," in *Proc. ACM WPES*, Oct. 2007, pp. 11–20.
- [12] X. Wang, S. Chen, and S. Jajodia, "Network flow watermarking attack on low-latency anonymous communication systems," in *Proc. IEEE S&P*, May 2007, pp. 116–130.
- [13] W. Yu, X. Fu, S. Graham, D. Xuan, and W. Zhao, "DSSS-based flow marking technique for invisible traceback," in *Proc. IEEE S&P*, May 2007, pp. 18–32.
- [14] N. B. Amir Houmansadr and N. Kiyavash, "RAINBOW: A robust and invisible non-blind watermark for network flows," in *Proc. 16th NDSS*, Feb. 2009, pp. 1–13.
- [15] V. Shmatikov and M.-H. Wang, "Timing analysis in low-latency MIX networks: Attacks and defenses," in *Proc. ESORICS*, 2006, pp. 18–31.
- [16] V. Fussenig, E. Staab, U. Sorger, and T. Engel, "Slotted packet counting attacks on anonymity protocols," in *Proc. AISL*, 2009, pp. 53–60.
- [17] X. Wang, S. Chen, and S. Jajodia, "Tracking anonymous peer-to-peer VoIP calls on the internet," in *Proc. 12th ACM CCS*, Nov. 2005, pp. 81–91.
- [18] K. Bauer, D. McCoy, D. Grunwald, T. Kohno, and D. Sicker, "Low-resource routing attacks against anonymous systems," Univ. Colorado Boulder, Boulder, CO, Tech. Rep., Aug. 2007.
- [19] X. Fu, Z. Ling, J. Luo, W. Yu, W. Jia, and W. Zhao, "One cell is enough to break Tor's anonymity," in *Proc. Black Hat DC*, Feb. 2009 [Online]. Available: <http://www.blackhat.com/presentations/bh-dc-09/Fu/BlackHat-DC-09-Fu-Break-Tors-Anonymity.pdf>
- [20] R. Dingleline, N. Mathewson, and P. Syverson, "Tor: Anonymity online," 2008 [Online]. Available: <http://tor.eff.org/index.html.en>
- [21] R. Dingleline and N. Mathewson, "Tor protocol specification," 2008 [Online]. Available: [https://gitweb.torproject.org/torspec.git?a=blob\\_plain;hb=HEAD;f=torspec.txt](https://gitweb.torproject.org/torspec.git?a=blob_plain;hb=HEAD;f=torspec.txt)
- [22] J. Reardon, "Improving Tor using a TCP-over-DTLS tunnel," Master's thesis, University of Waterloo, Waterloo, ON, Canada, Sep. 2008.
- [23] R. Dingleline and N. Mathewson, "Tor path specification," 2008 [Online]. Available: [https://gitweb.torproject.org/torspec.git?a=blob\\_plain;hb=HEAD;f=path-spec.txt](https://gitweb.torproject.org/torspec.git?a=blob_plain;hb=HEAD;f=path-spec.txt)
- [24] X. Fu, Z. Ling, W. Yu, and J. Luo, "Network forensics through cloud computing," in *Proc. 1st ICDCS-SPCC*, Jun. 2010, pp. 26–31.
- [25] M. Perry, "TorFlow: Tor network analysis," in *Proc. 2nd HotPETs*, 2009, pp. 1–14.
- [26] R. Pries, W. Yu, S. Graham, and X. Fu, "On performance bottleneck of anonymous communication networks," in *Proc. 22nd IEEE IPDPS*, Apr. 14–28, 2008, pp. 1–11.
- [27] G. Smillie, *Analogue, Digital Communication Techniques*. London, U.K.: Butterworth-Heinemann, 1999.
- [28] N. S. Evans, R. Dingleline, and C. Grothoff, "A practical congestion attack on Tor using long paths," in *Proc. 18th USENIX Security Symp.*, Aug. 10–14, 2009, pp. 33–50.
- [29] S. J. Murdoch, "Hot or not: Revealing hidden services by their clock skew," in *Proc. 13th ACM CCS*, Nov. 2006, pp. 27–36.
- [30] R. Pries, W. Yu, X. Fu, and W. Zhao, "A new replay attack against anonymous communication networks," in *Proc. IEEE ICC*, May 19–23, 2008, pp. 1578–1582.
- [31] D. McCoy, K. Bmmoad and Wmt

- [38] Z. Ling, J. Luo, W. Yu, and X. Fu, "Equal-sized cells mean equal-sized packets in Tor?," in *Proc. IEEE ICC*, Jun. 2011, pp. 1–6.
- [39] D. X. Song, D. Wagner, and X. Tian, "Timing analysis of keystrokes and timing attacks on SSH," in *Proc. 10th USENIX Security Symp.*, Aug. 2001, p. 25.
- [40] M. Liberatore and B. N. Levine, "Inferring the source of encrypted HTTP connections," in *Proc. ACM CCS*, Oct. 2006, pp. 255–263.
- [41] C. V. Wright, L. Ballard, F. Monrose, and G. M. Masson, "Language identification of encrypted VoIP traffic: Alejandra y Roberto or Alice and Bob?," in *Proc. 16th Annu. USENIX Security Symp.*, Aug. 2007, pp. 43–54.
- [42] 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 conversation," in *Proc. IEEE S&P*, May 2008, pp. 35–49.
- [43] X. Wang and D. S. Reeves, "Robust correlation of encrypted attack traffic through stepping stones by manipulation of inter-packet delays," in *Proc. ACM CCS*, Nov. 2003, pp. 20–29.
- [44] P. Peng, P. Ning, and D. S. Reeves, "On the secrecy of timing-based active watermarking trace-back techniques," in *Proc. IEEE S&P*, May 2006, pp. 335–349.
- [45] Y. J. Pyun, Y. H. Park, X. Wang, D. S. Reeves, and P. Ning, "Tracing traffic through intermediate hosts that repacketize flows," in *Proc. IEEE INFOCOM*, May 2007, pp. 634–642.



**Zhen Ling** received the B.S. degree in computer science from Nanjing Institute of Technology, Nanjing, China, in 2005, and is currently pursuing the Ph.D. degree in computer science and engineering at Southeast University, Nanjing, China.

He joined Department of Computer Science, City University of Hong Kong, Hong Kong, from 2008 to 2009 as a Research Associate, and then joined the Department of Computer Science, University of Victoria, Victoria, BC, Canada, in 2011 as a visiting scholar. His research interests include network

security, privacy, and forensics.



**Junzhou Luo** (M'10) received the B.S. degree in applied mathematics and M.S. and Ph.D. degrees in computer network from Southeast University, Nanjing, China, in 1982, 1992, and in 2000, respectively.

He is a Full Professor with the School of Computer Science and Engineering, Southeast University. His research interests are next-generation network, protocol engineering, network security and management, grid and cloud computing, and wireless LAN.

Prof. Luo is Co-Chair of the IEEE SMC Technical Committee on Computer Supported Cooperative

Work in Design.



**Wei Yu** received the B.S. degree in electrical engineering from Nanjing University of Technology, Nanjing, China, in 1992, the M.S. degree in electrical engineering from Tongji University, Shanghai, China, in 1995, and the Ph.D. degree in computer engineering from Texas A&M University, College Station, in 2008.

He is an Assistant Professor with the Department of Computer and Information Sciences, Towson University, Towson, MD. Before that, he worked for Cisco Systems, Inc., San Jose, CA, for almost nine years. His research interests include cyberspace security, computer network, and distributed systems.



**Xinwen Fu** received the B.S. degree in electrical engineering from Xi'an Jiaotong University, Xi'an, China, in 1995, the M.S. degree in electrical engineering from the University of Science and Technology of China, Hefei, China, in 1998, and the Ph.D. degree in computer engineering from Texas A&M University, College Station, in 2005.


He is an Assistant Professor with the Department of Computer Science, University of Massachusetts Lowell, Lowell, which he joined in the summer of 2008 as a faculty member. From 2005 to 2008, he was an Assistant Professor with the College of Business and Information Systems, Dakota State University, Madison, SD. His current research interests are in network security and privacy.



**Dong Xuan** received the B.S. and M.S. degrees in electronic engineering from Shanghai Jiao Tong University (SJTU), Shanghai, China, in 1990 and 1993, respectively, and the Ph.D. degree in computer engineering from Texas A&M University, College Station, in 2001.

Currently, he is an Associate Professor with the Department of Computer Science and Engineering, The Ohio State University (OSU), Columbus. He was on the faculty of Electronic Engineering at SJTU from 1993 to 1998. His research interests include distributed computing, computer networks, and cyberspace security.

Dr. Xuan received the NSF CAREER Award in 2005 and the Lumley Research Award from the College of Engineering, OSU, in 2009.



**Weijia Jia** received the B.Sc. and M.Sc. degrees from Center South University, Changsha, China, in 1982 and 1984, respectively, and the Master of Applied Science and Ph.D. degrees from the Polytechnic Faculty of Mons, Mons, Belgium, in 1992 and 1993, respectively, all in computer science.

He is currently a Full Professor with the Department of Computer Science and the Director of Future Networking Center, ShenZhen Research Institute, City University of Hong Kong (CityU), Hong Kong. He joined the German National Research Center for Information Science (GMD), Bonn (St. Augustine), Germany, from 1993 to 1995 as a Research Fellow. In 1995, he joined the Department of Computer Science, CityU, as an Assistant Professor. His research interests include next-generation wireless communication, protocols and heterogeneous networks; distributed systems, and multicast and anycast QoS routing protocols.