

On the Detection of Signaling DoS Attacks on 3G/WiMax Wireless Networks

Patrick P. C. Lee^{a,*}, Tian Bu^b, Thomas Woo^b

^a*Department of Computer Science and Engineering, The Chinese University of Hong Kong*

^b*Bell Labs, Alcatel-Lucent*

Abstract

Third Generation (3G) wireless networks based on the CDMA2000 and UMTS standards are now increasingly being deployed throughout the world. Because of their complex signaling and relatively limited bandwidth, these 3G networks are generally more vulnerable than their wireline counterparts, thus making them fertile ground for new attacks. In this paper, we identify and study a novel Denial of Service (DoS) attack, called *signaling attack*, that exploits the unique vulnerabilities of the signaling/control plane in 3G wireless networks. Using simulations driven by real traces, we are able to demonstrate the impact of a signaling attack. Specifically, we show how a well-timed low-volume signaling attack can potentially overload the control plane and detrimentally affect the key elements in a 3G wireless infrastructure. The low-volume nature of the signaling attack allows it to avoid detection by existing intrusion detection algorithms, which are often signature or volume-based. As a counter-measure, we present and evaluate an online early detection algorithm based on the statistical CUSUM method. Through the use of extensive trace-driven simulations, we demonstrate that the algorithm is robust and can identify an attack in its inception, before significant damage is done. Apart from 3G networks, we also show that many emerging wide-area networks such as 802.16/WiMax share the same vulnerability and our solution can also apply.

Key words: 3G wireless, security, DoS attacks.

Note: An earlier and shorter conference version of this paper appeared in IEEE INFOCOM '07 [19]. This paper additionally includes a discussion of the signaling attack in WiMax/802.16, provides evaluation on different dimensions of attacks,

*Corresponding author: Patrick P. C. Lee.

Email addresses: pclee@cs.columbia.edu (Patrick P. C. Lee),
tbu@research.bell-labs.com (Tian Bu), woo@research.bell-labs.com (Thomas Woo)

and presents more rigorous arguments in some of our analysis.

1. Introduction

The targets of most denial-of-service (DoS) attacks so far are wireline end-points, whose prevalence provides vast opportunities for an attacker to explore and launch new attacks. As the roll-out of nation-wide wireless data networks continues, we expect more types of DoS attacks will start targeting wireless networks. Currently, third generation (3G) wide-area wireless networks based on the CDMA2000 [28] and UMTS [30] standards are widely deployed. As of December 2005, there were over 300 million CDMA subscribers worldwide[6]. Emerging 3G data standards, such as EV-DO and HSDPA, promise to deliver broadband mobile Internet services with peak rates of 2.4 Mbps and 14.4 Mbps, respectively. The number of data subscribers is projected to reach a billion before 2010 [6]. As the number of data-capable wireless end-points escalates, they will become susceptible targets of new DoS attacks in the near future.

Apart from the sheer number of mobile endpoints, a multitude of other factors also contribute to making a 3G wireless network more vulnerable to attacks. These include:

- *Limited wireless link bandwidth:* As opposed to most wireline links, 3G wireless links tend to have much lower capacity thus it takes significantly less traffic to overload the link.
- *High signaling overhead:* With existing 3G standards, to transfer a similar amount of data, a lot more signaling messages/handshakes is needed for a wireless network than that in a wireline network. For instance, in order to improve the utilization of limited radio resources, a radio channel is only allocated to a mobile when there is data to transfer, and it will be revoked after an inactivity timeout. Such dynamic channel allocation and revocation procedures introduce lots of signaling operations.
- *Heavy control processing:* The hierarchical nature of current 3G (CDMA2000 or UMTS) networks places certain critical system functions such as power control, resource allocation, paging, etc. on a few infrastructure elements. The Radio Network Controller (RNC) and the base stations (BS) are involved in these activities for each mobile. By necessity, the engineering of these network elements is typically based on a certain load profile that is derived from the projected traffic patterns and behaviors of mobiles. Any operational deviation from these design assumptions can cause significant overload condition, and potentially non-graceful degradation.

In a nutshell, 3G wireless networks are significantly more fragile than wireline networks. To begin, most of the wireline DoS attacks would still apply to a wireless network. In addition, the above unique vulnerabilities of 3G networks can be exploited by new forms of wireless-specific DoS attacks.

In this paper, we introduce a novel DoS attack termed the *signaling attack*, which seeks to overload the control plane of a 3G wireless network using *low-rate, low-volume* attack traffic, based on some of the aforementioned 3G-specific vulnerabilities. Unlike conventional DoS attacks that focus on the data plane, the signaling attack creates havoc in the signaling plane of a 3G network by repeatedly triggering radio channel allocations and revocations. To accomplish this, an attacker first sends a low-volume packet burst to a mobile. If the mobile does not currently have a radio channel, the network will allocate a new one to complete the data transfer. After an inactivity timeout, the radio channel is torn down to recycle it back for others' use and help preserve the mobile energy that will otherwise be wasted on maintaining the channel. Immediately after the channel release, the attacker sends another low-volume packet burst to the mobile so as to trigger another radio channel establishment. By repeatedly doing so at appropriately timed periods, this can generate a considerable number of signaling operations. As detailed in Section 2, each channel establishment/release requires the RNC and BS to process more than 20 signaling messages. Launching this against large number of mobiles can easily introduce an excessive amount of signaling messages. The potential damage includes (1) overloading of RNC and BS, leading to reduced system performance, (2) denial of service to legitimate signaling messages due to congestion in the signaling paths, and (3) shortening of the mobile battery life.

As opposed to current DoS attacks that generate aggressive traffic, the signaling attack can be achieved with low-rate, low-volume traffic. Thus, the signaling attack can effectively evade detection by today's intrusion detection/prevention systems, which are effective mostly against flooding-based DoS attacks.

To understand the damage caused by the signaling attack, suppose that a 3G wireless network has inactivity timeout set to 5 s¹ and that an attacker generates a 40-byte packet burst (e.g., a TCP/IP packet with zero payload). By sending packet bursts periodically at a time slightly larger than 5 s, the attacker generates only 64bps attack traffic, which is invisible to volume-based detection systems. If the attacker is using a cable modem with 1.5Mbps uplink bandwidth, then it can simultaneously attack approximately 24K mobiles, a number that is sufficient to bring down a wireless network infrastructure (see Section 2). Note that this signaling attack can also be mounted to other emerging wide-area networks such as

¹It is the default value in many deployed networks.

802.16/WiMAX that share the same vulnerability.

We propose an online early detection mechanism for signaling attack by formulating it as a change-point detection problem, which aims to identify the sources of signaling attacks by monitoring any abnormal behavior against profiled benign behavior. Our detection algorithm is based on the *non-parametric CUSUM test*, which does not require the parameterization of the network traffic distribution of interest and has been shown to be asymptotically optimal for change-point problems [4, 5, 23, 32]. To make the algorithm robust against intelligent attackers that attempt to increase the detection delay and hence aggravate the damage, our detection mechanism is derived in a way that the detection delay depends only on the amount of additional signaling load due to the attack but not on anything else, such as the attack strategy.

To summarize, this paper makes the following contributions. We identify a novel signaling DoS attack specific for 3G wireless networks and demonstrate its severity through simulation driven by synthetic and real traces. The attack is unique in the sense that it overloads the control plane with only low-rate, low-volume attack traffic, and hence makes conventional signature or volume-based detection schemes ineffective. To combat this, we present a novel online early detection algorithm that can identify the attackers before they cause any major damage. Through extensive trace-driven evaluation, we demonstrate that our detection algorithm has a short detection delay while rarely generating false alarms.

The remainder of the paper proceeds as follows. Section 2 presents the vulnerabilities of the signaling plane in a general 3G wireless network and explains how they are susceptible to the signaling attack. Based on trace-driven simulation, Section 3 demonstrates the severe damage resulting from the signaling attack. In Section 4, we present a CUSUM-based detection mechanism to identify the source of a signaling attack. In Section 5, we evaluate the effectiveness of our detection mechanism. Section 6 addresses possible enhancements to our current work and the reaction mechanism. Section 7 reviews related work on different forms of DoS attacks in wireline and wireless networks, and we conclude in Section 8.

2. Signaling attacks on 3G Wireless Networks

In this section, we first overview the network elements of a 3G network architecture and their inter-connections. Here, we use the UMTS network architecture as our example. We focus on the signaling procedures among the network elements for radio channel establishment and release. We then demonstrate how an attacker may exploit these signaling procedures to overload the control plane. Finally, we address how the signaling attack can be launched to IEEE 802.16 wireless metropolitan-area networks as well.




Figure 1: UMTS Network Architecture.

2.1. UMTS Network Architecture

Figure 1 shows the typical architecture of a UMTS wireless network. We first describe two of its main components: the Gateway GPRS Support Node (GGSN) and the Serving GPRS Support Node (SGSN). The GGSN is a GPRS network entity that serves as the mobile wireless gateway between an SGSN and the Internet. When a mobile successfully authenticates and registers with the network, a Point-to-Point (PPP) link is set up between the GGSN and the mobile. On the other hand, the SGSN is responsible for sending data to and from mobile stations, in addition to maintaining information about the location of a mobile and performing authentication for the mobile. Typically, there are multiple SGSNs, each of which serves the GPRS users physically located in its serving area.

Another key component of a UMTS network is the Radio Network Controller (RNC), which is the point where wireless link layer protocols terminate. The RNC provides the interface between a mobile communicating through a Base Station (BS) and the network edge. This includes management of radio transceivers in BS equipment (radio resource control), admission control, channel allocation, as well as management tasks such as handoffs between BSs and deciding power control parameters. The functionalities of a BS include wireless link transmission/reception, modulation/demodulation, physical channel coding, error handling, and power control. In this hierarchical architecture, multiple mobiles communicate with a BS, and multiple BSs communicate with an RNC, and multiple RNCs talk to the SGSN/GGSN.

2.2. Signaling for Radio Resource Control

The data service session provided by the network to mobile is often referred to as a *data call* that starts from the mobile connects to the network for data service and stops when the mobile disconnects from the network. Figure 2 shows the steps of a UMTS data call. The paging stage is not presented as it is necessary only when the network originates a data call. We also skip the authentication and security

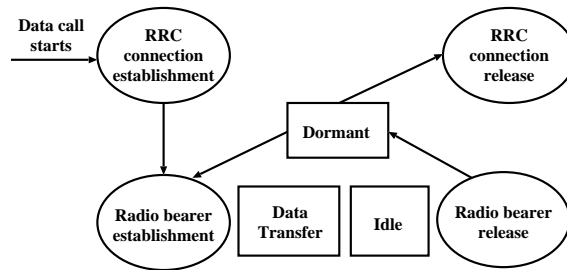


Figure 2: UMTS data call. Multiple RABs may be established within one RRC connection.

events that are not relevant in the context. Detailed specification of a UMTS data call can be found in [1].

A Radio Resource Control (RRC) connection is first created. While there is only one RRC connection during a data call, the network may establish one or more Radio Access Bearers (RABs) within the single RRC connection in an on-demand fashion. RABs are the actual radio resources for data communication. They are released after a timeout period for inactivity so that they can be reused by other mobiles. In addition, the mobiles may extend their battery life by releasing the idle RAB because maintaining a RAB requires periodic channel condition updates that consume much energy.

When new data arrives for a mobile without an RAB, an RAB establishment procedure is invoked to allocate radio resources for data delivery. The procedure varies depending on the type of RAB being established or released. Figure 3 shows the procedure for establishing a synchronized RAB. In a nutshell, the RNC first requests the base station to prepare resources for a new RAB through the appropriate Radio Link Reconfiguration messages. The RNC and the base station also establish uplink and downlink synchronism for the RAB. The RNC and the mobile then agree upon the transport parameters through the base station (which acts as a relay) and finally set up an RAB. There are about a total of 15 signaling messages being processed by the RNC. The processing load would be even higher in order to support fast handoff.

The allocated RAB is released after an inactivity timeout. Similarly, the release procedure will vary depending on the type of the RAB. Figure 4 shows the procedure for the release of a synchronized RAB. The RNC first requests the base station to prepare to release the resources of the RAB. Then the RNC and the mobile release the RAB by exchanging control messages via the base station. Finally, the base station releases its resources. There are a total of 12 signaling messages being processed by the RNC. Similarly, the support of fast handoff would result in even more signaling steps.

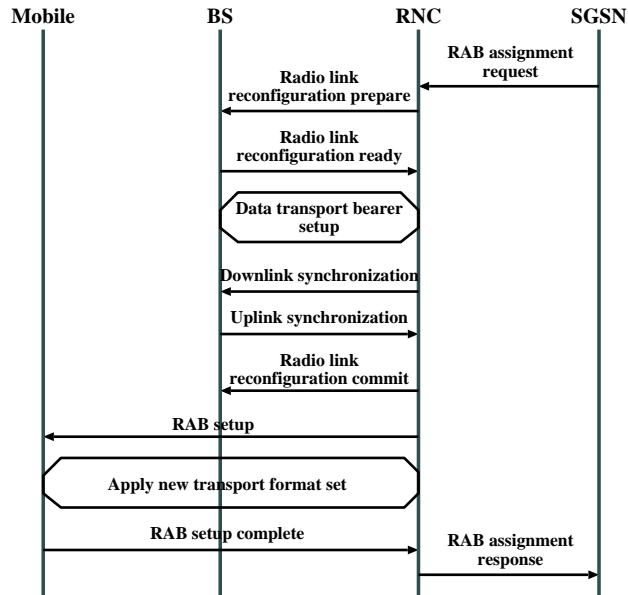


Figure 3: UMTS Radio Bearer Establishment

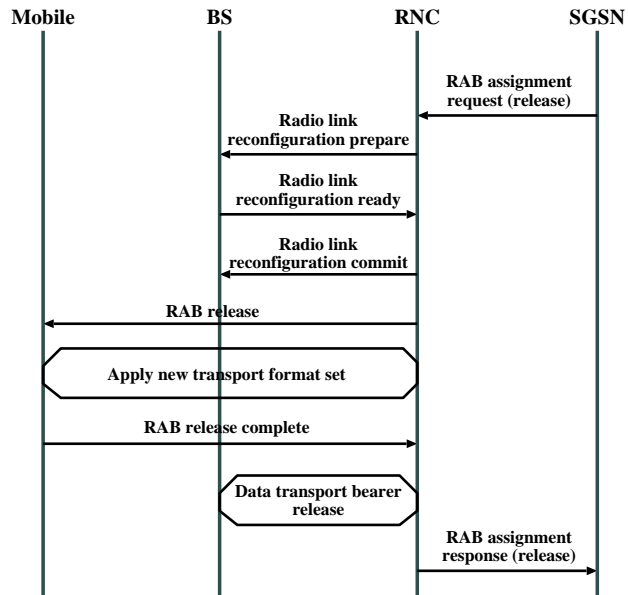


Figure 4: UMTS Radio Bearer Release

2.3. Signaling Attacks on UMTS Networks

Due to the signaling overhead required for RAB setup/release, an attacker may seek to trigger excessive amount of signaling messages in order to overload an RNC and potentially the BSs. This can be done by *regularly sending low-volume bursts at appropriately timed periods such that immediately after a RAB is torn down due to inactivity, a burst arriving from the attacker will trigger a new RAB establishment*. It should be noted that the attack traffic can simply be composed of regular data packets (e.g., TCP/IP packets) and a personal computer with Internet connectivity is sufficient for launching the attack.

The attack may cause severe damage to the control plane: (a) the signaling path between the RNC and the BS is congested with the setup/release messages associated with each RAB; (b) the RNC processor is tied up with maintaining states and processing signaling messages and is thus overloaded. The consequence of (a) is that valid signaling messages may not receive allocation of resources, causing it to be dropped by the RNC due to insufficient buffers and/or excessive delay that leads to timeout. The consequence of (b) is that overloading an RNC can effectively deny legitimate services to the mobiles being served by the RNC. Another side effect of the signaling attack is the potential draining of the mobile battery. Normally, for power conservation, a mobile switches to a low-power idle or dormant state when there is no packet to be sent or received. Since low-volume bursts are sent regularly, mobiles would stay active longer than necessary.

To quantify the damage due to the signaling attack, we conduct simple calculation. The transmissions of signaling messages are based on Asynchronous Transfer Mode (ATM) [1]. As ATM encodes traffic into 53-byte fixed-sized cells, each RAB establishment/release procedure contains at least $(27 \text{ signaling messages}) \times (53 \text{ bytes}) \approx 10 \text{ Kbits}$. Suppose that the total Iub interface capacity (i.e., the capacity between base stations and an RNC) is in the range of 150-200 Mbps [26], and that the RAB inactivity timeout is 5 s, which is the default value in many operational 3G networks. To launch the signaling attack, an attacker only needs to generate 40-byte packets (e.g., TCP/IP packets with zero payload) to 20K mobiles periodically at a time slightly larger than 5 s, or equivalently 1.28Mbps attack traffic. This will result in at least 200 Mbps signaling traffic, and hence overload an RNC. Note that the amount of attack traffic is small enough to be generated using a cable modem with 1.5Mbps uplink bandwidth. The low-volume nature of the signaling attack makes it hard to be detected by conventional intrusion-detection mechanisms that are designed to defend against flooding-based DoS attacks.

2.4. Signaling Attacks on CDMA2000

CDMA2000 also has a hierarchical architecture where PDSN and MobileIP Home Agents handle similar packet functionality as GGSN and SGSN do in UMTS




Figure 5: 802.16 data transport.

networks. In addition, CDMA2000 follows similar radio resource allocation and revocation procedures as does UMTS. Thus, the signaling protocols of CDMA2000 can also be exploited by attackers to overload the RNC.

2.5. Signaling Attacks on IEEE 802.16 Networks

We close this section by briefly describing the signaling protocols of data transport in IEEE 802.16 networks and showing how they can be vulnerable to signaling attacks. Our description of the signaling protocols is based on the 802.16e-2005 standard [12], which extends and updates the previous 802.16-2004 standard to include mobile subscriber stations. We refer readers there for detailed discussion.

In IEEE 802.16 wireless metropolitan-area networks, the data service for a mobile is based on the concept of *service flow*, a MAC-layer transport service that describes the uni-directional flow of either uplink or downlink data. The establishment of a service flow uses a two-phase model: a service flow is first *admitted* with provisioned resources, and then the service flow is then *activated* to have the resources committed on an on-demand basis. The service flow may be de-activated later to conserve network resources.

Figure 5 outlines the procedures of re-activating a service flow for data transport. Initially, when a mobile has no downlink or uplink data, it will enter either Sleep Mode or Idle Mode, both of which aim to reduce the power consumption of the mobile. Upon the availability of data, the serving BS will awaken the mobile. The mobile then establishes a connection with the BS via *initial ranging*. Ranging parameters are then adjusted for the connection. Finally, the service flow is re-activated for data transfer, and the mobile returns to the normal operation stage. Depending on whether the serving BS has the necessary information, the mobile may need to carry out more signaling operations (not shown in Figure 5), such as basic capability negotiation, authentication and key management, re-registration, as well as IP connectivity re-establishment. Given the above signaling procedures, attackers may also launch similar signaling attacks to an 802.16 BS by triggering unnecessary state transitions that overload the BS with signaling processing.

3. Impact of the Signaling Attack

To demonstrate the impact of the signaling attack, the ideal approach is to attack an operational 3G wireless network and observe performance degradation or even network breakdown. However, the experiments on a real network are not feasible due to both economical and legal reasons. We may also collect data traces and corresponding signaling traces from 3G operational networks and simulate the signaling attack by inserting hypothetical attack traffic into the traces. However, data traces from live 3G networks are not currently available due to sensitivities around privacy and competition. Therefore, we consider two types of traces for the purpose of demonstration: (1) synthetic traces derived from the 3G traffic models presented in the literature and (2) real traces collected from a campus-wide wireless network.

3.1. Simulation with Synthetic Traces

Based on the 3G traffic models described in [2, 9, 20], we construct a discrete-event traffic simulator that generates synthetic traces for a UMTS network. The trace generation is composed of three levels. In the *session level*, the simulator creates a new user session according to a Poisson arrival process with mean 1 s, such that each session is of type voice or data with probabilities 0.4 and 0.6, respectively. We assume that a voice session has an exponentially distributed duration with mean 120 s, while a data session lasts until all session data has been transferred, where the data volume follows a lognormal distribution with parameters $\mu=11.1170$ s and $\sigma=1.3818$ s. In the *burst level*, each (voice or data) session consists of alternating ON/OFF periods such that packets are generated during the ON period only. Both ON and OFF durations are exponentially distributed, with the same mean 3 s for voice sessions and with means 0.2075 s and 12 s, respectively, for data sessions. In the *packet level*, packets are generated based on the distributions of inter-arrival time and packet size. Since the packet-level time variations are so small that they are insignificant to the inactivity timer of a radio channel, we assume that the packet-level parameters are constant. Thus, we fix the packet inter-arrival time and packet size to be 0.02 s and 31 bytes, respectively, for voice sessions (based on the AMR 12.2 mode) and 8.3 ms and 480 bytes, respectively, for data sessions. We assume that each mobile is associated with a single user session, and that after the session is completed, the mobile remains connected to the wireless network for 30 minutes based on the observation in [17]. In addition, we assume that all mobiles within the network are served by a single RNC. Our goal is to evaluate the number of signaling events triggered on this specific RNC.

After generating the synthetic traces, we add an attacker that periodically injects additional low-volume packet bursts at every *attack interval*, which is set to

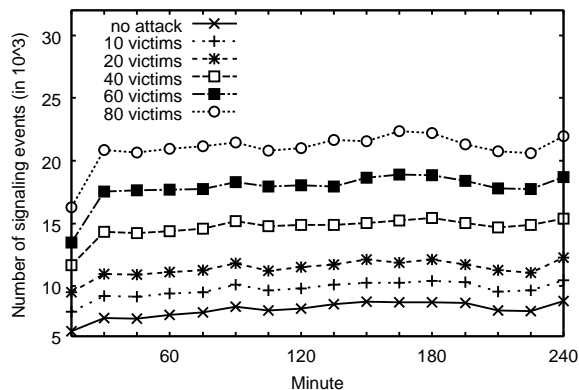


Figure 6: Signaling load under a signaling attack using a synthetic trace.

be slightly larger than the inactivity timeout of a radio channel so as to maximize the chance of introducing signaling events. Specifically, the attacker randomly selects a subset of mobiles that are connected to the network. It then sends a 40-byte packet (e.g., a TCP packet with zero payload) to each of the attacked mobiles. Any attacked mobile that has been idle over the channel inactivity timeout will invoke a signaling event at the RNC for establishing a new radio channel. Here, we focus on a UMTS network that has inactivity timeout set to 5 s. It should be noted that this inactivity timeout is a tunable parameter that is set up by network operators once and rarely changes.

Figure 6 illustrates, over a 4-hour-long synthetic trace, the signaling load recorded every 15 minutes when different numbers of mobiles are attacked at each attack interval. It clearly shows the severe damage resulting from the signaling attack. For example, when 80 mobiles are attacked at each attack interval, the signaling load increases by 2.5 times as compared to the no-attack case.

Also, the 4-hour-long synthetic trace contains 2.4 GB of normal traffic, while the additional traffic for attacking 80 mobiles accounts for $(14400 / 5 \times 80 \times 40 \text{ bytes}) / 2.4\text{GB} = 0.4\%$ of the total traffic only. Among the attack traffic, about 90% of which triggers a signaling event at an RNC. This shows that an RNC can be heavily overloaded with low-volume traffic.

3.2. Simulation with Real Traces

As noted in Section 3.1, the packet-level time scales are much smaller than the inactivity timer of a radio channel. Therefore, although the wireless access medium may disturb the packet transmission delay, the channel inactivity and hence the RAB setup/release events are less influenced by link delay variance of the access medium. Instead, they are mainly influenced by the session arrival pattern resulting

from user and application behaviors as well as network transport layer protocols. In particular, we observe that when a mobile completes its session and remains connected to the wireless network, an attacker can send TCP/IP packets to the mobile to trigger signaling events. Thus, regardless of the access medium, the impact of a signaling attack that we observe should be similar to that in the previous synthetic-trace simulation.

Due to the unavailability of real traces from a 3G operational network, we instead take the traces collected from the IEEE 802.11b wireless local area network (LAN) at Dartmouth College in Fall 2003 [17] for our evaluation purpose. We emphasize that by no means do we suggest that these traces are ideal for a 3G wireless network. Nevertheless, we envision that both IEEE 802.11b and 3G/WiMax networks have similar user and application behaviors in Internet access. Also, note that the 802.11b traces that we are using are collected at the TCP/IP layer. In Section 4.2, we detail how we can capture IP packets in a 3G/WiMax network. Therefore, our analysis is equally applicable if we are given the 3G/WiMax traces at the TCP/IP layer.

In order to use the traces to drive the simulation of a 3G network, we first assume that all hosts associated with the access points are mobiles in a 3G network that is served by a single RNC. We then simulate the signaling load at this specific RNC before and after the signaling attack.

We use the same setup as in Section 3.1 for injecting attack traffic, i.e., at each attack interval that is slightly larger than 5 s, we have an artificial attacker pick a random set of mobiles that are connected to the wireless network and send a 40-byte TCP packet to each of the mobiles. Our goal is to infer the number of signaling events from attack traffic at the TCP/IP layer.

Figure 7 illustrates the signaling load when different numbers of mobiles are attacked at each attack interval, using the traces collected on November 3, 2003. As in Section 3.1, the signaling load is significantly increased due to the signaling attack. For instance, by periodically attacking 80 mobiles, the signaling load of the RNC increases more than five times, while (not shown in the figure) only less than 0.6% additional traffic is introduced. We also repeat the simulation using traces collected from different days and the results are similar.

3.3. Discussion

When a malicious host launches signaling attacks, it might not know the exact set of IP addresses that have been assigned to *active mobiles* (i.e., the mobiles being connected to the wireless network). However, the IP address segments of major wireless service providers are public information [11]. Recent studies [7, 24] also discuss how to identify active mobiles in a cellular data network, such as creating




Figure 7: Signaling load under a signaling attack using the real traces collected from Dartmouth College.

hit-lists for targeted mobiles. To simplify our simulation without compromising our key results, we only insert attack traffic for active mobiles.

Since the bandwidth required for attacking a single mobile is very small, the attacker may choose to send packets to all or a large portion of the IP addresses of a wireless network. Any packet targeting an inactive mobile is dropped by GGSNs (UMTS) or PDSNs (CDMA2000), whereas any packet reaching an active mobile that has its radio channel released would trigger a signaling event for a channel setup. In this case, the signaling attack can still be achieved with low-volume attack traffic even the active mobiles only occupy a small portion of the IP address segment. For example, in the real-trace simulation in Section 3.2, if we assume that the active mobiles only occupy 10% of the IP address segment, the attacker only needs 6% additional traffic to increase the signaling load of an RNC by five times.

Our trace-driven simulation results show that an attacker can increase the signaling load of an RNC by a multiplicative factor by sending a small amount of attack traffic to a small number of mobiles (e.g., 80 victims in our simulation). As shown in Section 2.3, to actually bring down an RNC, an attacker may need to attack more mobiles (e.g., 20K victims). Nevertheless, the low-volume nature of the signaling attack still holds, since the amount of attack traffic for such a practical attack remains significantly small (see Section 2.3 for details).

4. Detecting Signaling Attacks

In order to defend against signaling attacks, we propose to monitor all data traffic that enters and leaves a 3G wireless network for detecting any ongoing sig-

nalng attacks and their originators. Our *online* detection algorithm is based on the statistical *cumulative sum (CUSUM)* test, whose objective is to identify any anomaly that deviates from normal behavior as early as possible.

In this section, we first describe what data samples should be used as the input to the detection algorithm and where they should be collected. We then overview the CUSUM test and demonstrate how to apply the CUSUM test for robust detection of signaling attack with the appropriate choices of parameters.

4.1. Data Sample Collection

We now explain how data samples are collected for the purpose of attack detection. We define a *remote host* to be a node that originates downlink traffic to a mobile within a wireless network. Note that a remote host can be either a node in the Internet or another mobile in the same wireless network. We uniquely identify a *flow* (a.k.a. *session*) by the address pair of a mobile and a remote host. We refer a packet to be *inbound* (resp. *outbound*) if it is transmitted from (resp. to) a remote host to (resp. from) a mobile. We treat a packet as inbound if the remote host is a mobile itself. We say a remote host *triggers* a *virtual setup* if it initiates a packet to a flow that has not transmitted any inbound or outbound packet over the last inactivity timeout.

It is important to note that virtual setups are measured at the flow level. They do not necessarily lead to any actual radio channel setup (and hence any signaling event) since a mobile does not release its radio channel if there have been data transfers between the mobile and other remote hosts in the last inactivity timeout. However, the virtual setup provides a measure of the intention of a remote host to launch a signaling attack. If the remote host triggers a significant number of virtual setups over a short period of time, then there is a strong evidence that the remote host attempts to mount a signaling attack, regardless of how many signaling operations are actually caused by these virtual setups.

In our analysis, we only focus on inbound attacks, i.e., the signaling attack is always originated from remote hosts. Although a mobile may frequently establish and release radio channels with other non-mobiles, the signaling events are associated with a single mobile only. Hence, the increase in the RNC load is very limited.

For a given remote host, we define the *inter-setup time* as the time interval between two adjacent virtual setups (which can be on the same flow or on two different flows) that are triggered by the remote host. The smaller the inter-setup time is, the higher likelihood that the corresponding remote host is mounting a signaling attack. Therefore, for each remote host, we monitor a sequence of its inter-setup time samples $\{t_1, t_2, \dots\}$ in the CUSUM test (see Section 4.3). We then use the abrupt decrease of the inter-setup time as an indicator of a signaling attack.

4.2. Placement of Detection Points

In order to collect input data for the detection algorithm, we have to monitor all IP packets entering and leaving a wireless network, decode their source and destination addresses, and compute the inter-setup times. Although we may collocate the detection elements with an RNC/SGSN for monitoring all mobiles served by the RNC/SGSN, neither the RNC nor the SGSN works at the IP layer, leading to extra efforts for extracting IP packets from either radio frames (in the RNC) or encapsulated packets (in the SGSN). An ideal place for data collection would be between the GGSN and the Internet where all traffic is composed of IP packets. However, the inter-mobile communication always turns around at the GGSN and is thus not available for data collection in this setting. To tackle this problem, we suggest two solutions. The first one is to add a packet sniffer that passively captures all packets entering and leaving the GGSN and pass them to the detection unit. Another solution is to configure the GGSN so that all inter-mobile traffic is routed to the detection unit and then back to the GGSN. This loop-back feature is supported by the GGSN and being used for firewalling inter-mobile traffic today.

Some security measures such as firewalls and intrusion detection systems may have been placed between the GGSN and the Internet for protecting a wireless network from other general attacks. We may then implement our detection algorithm as an additional detection module in an existing detection system and take advantage of the traffic policies that have already been built into the system for reacting to the detected attacks.

4.3. Overview of the CUSUM Test

We first overview the CUSUM test and motivate how it is well-suited for intrusion detection. Here, we consider its *non-parametric version* that does not assume any a priori distribution of the data samples being considered. Thus, as long as the data samples are not extremely dependent, we are guaranteed the *asymptotic optimality*, i.e., the detection time is minimized among all possible detection schemes subject to a fixed worst-case expected false alarm rate [5, 32]. Detailed discussion of the CUSUM test and its non-parametric version can be found in [4, 5, 23, 32].

In the context of the signaling attack detection, for each remote host, the CUSUM test monitors a set of n inter-setup time samples $\{t_1, t_2, \dots, t_n\}$. Each inter-setup time sample t_n is assigned a score $z(t_n)$. When a sample t_n is available, we update the *CUSUM statistic* q_n as follows:

$$q_n = \max(q_{n-1} + z(t_n), 0),$$
$$\text{Take action if } q_n \geq h,$$
(1)

where $h > 0$ is the pre-specified *CUSUM threshold*. Note that if an inter-setup time sample t follows a malicious behavior, then the expected score $E(z(t))$ should be

positive so that q_n will eventually rise above threshold h . On the other hand, $E(z(t))$ should be negative when the data samples follow a benign behavior. We justify the choice of $z(t)$ in Section 4.4.

The CUSUM test is adequate for identifying any abrupt change of a benign behavior to a malicious behavior. To understand this, note that if a host is benign, $z(t_n)$ is negative (in the expected sense) and the corresponding q_n will stay around the zero value, regardless of how long the benign behavior has been observed. However, when the benign behavior turns to a malicious one, q_n increases and eventually surpasses the threshold. Therefore, the CUSUM test prevents an attacker from suppressing q_n with a long history of benign behavior. This ensures that the CUSUM test detects a malicious behavior in a timely manner.

4.4. Choice of Score $z(t)$

In this subsection, we justify the choice of $z(t)$, the score being added to the cumulative score S_n in the CUSUM test for a given inter-setup time t .

If $z(t)$ was not defined appropriately, an attacker could *trick* the CUSUM test by choosing a “more intelligent” malicious behavior in order to avoid being detected. To illustrate this, suppose that an inter-setup time t is regarded as benign when $t < \alpha$, or as malicious when $t \geq \alpha$, where α is some pre-specified parameter. We then define $z(t) = 1$ if $0 \leq t < \alpha$, and $z(t) = -1$ if $t \geq \alpha$. We assume that the threshold h is greater than 1. Suppose that a remote host called Attacker 1 triggers RAB setups every time $t = \alpha/2$. Since $z(t) > 0$ for every observed inter-setup time, Attacker 1 will eventually be detected. Now, suppose that another remote host termed Attacker 2 triggers virtual setups whose inter-setup time samples alternate between 0 and α . In this case, the cumulative score S_n will alternate between 0 and 1 as well. Since we assume $h > 1$, Attacker 2 will never be detected, although both Attacker 1 and Attacker 2 launch the same rate of signaling attacks (in the expected sense).

As a result, it is important to prevent an attacker from tricking the CUSUM test with an intelligent strategy. We seek to define $z(t)$ such that no attacker can evade detection as it causes extra signaling overhead, i.e., if any two attackers introduces the same expected rate of signaling attacks, then the expected delay to detect both attackers should be the same, regardless of their attack strategies. In other words, if any two attackers introduce the same number of virtual setups over time T , then they should have the same cumulative score $\sum z(t)$, which captures the behavior of all inter-setup time samples, when the CUSUM test is carried out at time T .

With this objective, we can formally prove that $z(t)$ must be a *linear function of t* . The detailed proof is shown in Appendix.

The formula for $z(t)$ can now be derived as follows. We consider a tunable parameter α , defined as the cut-off point such that if inter-setup time $t > \alpha$, it is

Algorithm 1 Detection algorithm

```
1: for each arrival of packet  $P$  belonging to  $F = (M, R)$  do
2:   if  $R$  has not been marked “malicious” then
3:     if  $P$  is inbound & virtual setup triggered then
4:       if the setup is NOT the first one triggered by  $R$  then
5:         Set  $t = \text{now} - R$ ’s last setup time
6:         Set  $R.q = \max\{R.q + (-t/\alpha + 1), 0\}$ 
7:         if  $R.q > h$  then
8:           Mark the remote host as “malicious”
9:           Set  $R$ ’s last setup time = now
```

likely to be a benign sample, while if $t < \alpha$, it is likely to be a malicious sample. Since $z(t)$ is linear, it attains maximum when $t = 0$. By setting $z(0) = 1$, we have $z(t) \leq 1$ for all t , and hence the ceiling of the CUSUM threshold $\lceil h \rceil$ denotes the minimum number of inter-setup time samples required to decide if a remote host is malicious. As a result, we can write $z(t)$ as

$$z(t) = \frac{-t}{\alpha} + 1. \quad (2)$$

4.5. Choice of h

For a given choice of $z(t)$ with $E(z(t)) < 0$, the probability that the CUSUM statistic q_n surpasses a positive value h is small but not zero either. However, the probability decreases as the value of h increases. Therefore, the false positive ratio is a decreasing function of h . On the other hand, q_n takes longer to surpass a large h even when $E(z(t)) > 0$ and hence the detection delay increases. The choice of h is to trade off between the false positive rate and the detection delay. Note that for some specific distributions (e.g., exponential) of data samples, h is often derived from the *average running length (ARL)* between two false positives that measures the false alarm tolerance. In the non-parametric setting (i.e., no a priori distribution of data samples is known), we will show how the empirical value of h can be obtained from the real traces in Section 3.2.

4.6. Detection Algorithm

We now present the detection algorithm that monitors if any remote host is mounting a signaling attack, as shown in Algorithm 1. The detection algorithm is applied only to the remote hosts that are not marked as malicious attackers. Upon the arrival of a packet P (either inbound or outbound), the detection algorithm first determines the corresponding flow F , which identifies mobile M and remote host R . An inbound packet from R only triggers a virtual setup for flow F (Line 3) when no packet has been observed in flow F over last inactivity timeout. If the virtual




Figure 8: Distributions of the auto-correlation of the inter-setup time per remote host.

setup is not the first one due to packets from R (Line 4), a new inter-setup time is computed (Line 5) and the CUSUM values associated with R are updated (Line 6) using Equations 1 and 2. R is marked malicious when the associated CUSUM statistic crosses the pre-selected threshold h (Lines 7 and 8). The most recent time that R triggers a virtual setup is tracked for computing the next inter-setup time (Line 9). Note that the initial values of the CUSUM parameters for all remote hosts are always set to zero.

4.7. Weak Dependency Condition for the CUSUM Test

In the non-parametric CUSUM test, we assume that the data samples are not heavily dependent [5, 32]. To validate this assumption, we compute the auto-correlation of the inter-setup time samples for each remote host, using the traces described in Section 3.2. Figure 8 illustrates the distribution of the resulting auto-correlations across all remote hosts. We observe that the dependency is not strong at all, as the auto-correlation of the inter-setup time samples is in $[-0.2, 0.2]$ for more than 80% of the remote hosts. To further eliminate the dependency among the inter-setup time samples, we may adopt the *batch-mean method* [14], in which the input of the CUSUM test is the mean of a batch of inter-setup time samples, or random sampling, in which the CUSUM test takes an inter-setup time sample as input with a fixed probability.

5. Performance Evaluation

In this section, we evaluate our detection mechanism against the signaling attack based on trace-driven simulation. From Section 3, we observe similar damage brought by the signaling attack using both synthetic and real traces. Thus, our

following evaluation focuses on the real traces collected from Dartmouth College. While the wireless LAN traces are not ideal, our evaluation methodology can be directly applied when real traces from a 3G operational network are available.

Here, we consider a UMTS system whose inactivity timeout for a RAB is set to be 5 s.

5.1. Metrics

In evaluating an intrusion-detection system, there are three fundamental metrics that we can consider: (1) *false positive ratio*, the fraction of benign remote hosts that are mistakenly identified as malicious over all observed remote hosts, (2) *false negative ratio*, the fraction of malicious remote hosts that are missed over all observed remote hosts, and (3) *detection time*, the delay that the detection algorithm needs to identify a malicious remote host since it starts the signaling attack.

Here, we do not consider the false negative ratio. Because of the linearity of $z(t)$, the CUSUM decision depends only on the additional signaling load due to the signaling attack, regardless of what the attack strategy is. If there exists a false negative, then it implies that the signaling load generated by this missed attacker is within an acceptable level. In this case, the malicious hosts that are missed by our detection mechanism actually have little impact on the total signaling load. As a result, our following analysis focuses on the false positive ratio and the detection time.

The above metrics are quantified based on how we classify a remote host as benign or malicious. Here, we assume that all remote hosts in the original traces are benign, and that the remote hosts that we inject to the traces for generating attack traffic are malicious.

The false positive ratio and the detection time are intrinsically conflicting metrics that depend on how we choose the tunable parameters α and h for our detection algorithm (see Section 4). Intuitively, by choosing a small value of α and a large value of h , we reduce the false positive ratio, but in the meantime lengthen the detection time. We investigate the trade-off in the following subsections.

5.2. False Positive Ratio vs. Different Tunable Parameters

In this subsection, we analyze the false positive ratio of our detection mechanism with different values of α and h .

For a given a 24-hour trace, we set α to be the p -th percentile of inter-setup times of all remote hosts. In other words, we bound the proportion of inter-setup time samples that have positive $z(t)$ to be no more than $p\%$. Here, we consider the cases where $p = 1, 5, 10, 15,$ and 20 . We then compute the false positive ratio versus different values of h using our detection mechanism from the simulation.




Figure 9: False positive ratio vs. h .

Figure 9 illustrates the false positive ratio versus different values of h , together with the percentile values of the observed inter-setup time samples. In general, most remote hosts only trigger a few virtual setups with large inter-setup times. Thus, the false positive ratio in most cases is very small. For instance, when α is no higher than the 10th-percentile, the false positive ratio is less than 0.06% for $h \geq 3$.

5.3. Detection Time vs. Different Tunable Parameters

To assess the detection time, we add to the traces a malicious remote host that generates low-rate, low-volume packets to a number of active mobiles every attack interval (see Section 3). We assume that the packets are randomly spaced over the attack interval. Also, we randomly choose the active mobiles to be attacked. We repeat the experiment 10 times with different seeds, and the measured detection time is averaged over the experimental instances.

Similar to the setting in Section 3, we fix the attack interval to be slightly larger than the inactivity timeout (which is 5 s according to our assumption). We assume that the attacker targets 5 mobiles at each attack interval.

Figure 10 illustrates the detection time versus different values of h , where α is set at different percentiles of inter-setup time samples. For smaller values of α and larger values of h , the detection mechanism requires longer detection time to identify the attacker, but in the meantime it produces fewer false positives (see Figure 9). This suggests the trade-off between different combinations of α and h .

5.4. Evaluation for the Signaling Attack Defense

We now evaluate our detection mechanism for defending against the signaling attack. Since the false positive ratio is low according to our previous evaluation, we fix $\alpha = 7.1138$ s, which is the 10th-percentile inter-setup time, and $h = 3$.




Figure 10: Detection time vs. h .




Figure 11: Detection time vs. no. of attacked mobiles

Our setup is similar to that in Section 5.3. In addition to the average detection time taken from 10 experimental instances, we also obtain the minimum and maximum detection times among those instances, and these bounds are shown as the endpoints of the vertical lines in the plots.

In Section 3, we demonstrate how a single remote host can overload a RNC via the signaling attack. Figure 11 shows the detection time versus different numbers of victims. In most cases, the single attacker can be identified in less than 10 s since the signaling attack starts. As the attacker becomes more aggressive by attacking more mobiles, we need less detection time to identify this attacker. For instance, as shown in Section 3, the attacker can double the signaling load of an RNC by periodically attacking 20 mobiles. However, such a signaling attack can be detected in less than one second. This shows our detection mechanism can quickly identify an attacker before it causes significant damage.




Figure 12: Detection time vs. attack interval.

5.5. Evaluations for Different Dimensions of Signaling Attacks

We now evaluate how our detection mechanism responds to different dimensions of signaling attacks.

Experiment 1 (Detecting a single attacker with different attack intervals):

In addition to the number of mobiles being attacked, the intensity of the signaling attack is also determined by the frequency that an attacker generates attack packets. In this experiment, we evaluate the impact of the attack interval, at which an attacker regularly attacks mobiles. We focus on a single attacker that attacks 10 mobiles at each attack interval.

Figure 12 plots the detection time at different attack intervals. Not surprisingly, the detection mechanism takes longer to identify the attacker as the attack interval increases (i.e., the damage of the signaling attack decreases). However, given that 10 mobiles are attacked, the signaling attack can be identified within a single attack interval in all cases.

Experiment 2 (Detecting multiple attackers with different attacking strategies): One of the major objectives in deriving our detection mechanism is to ensure that if any two attackers have the same expected attack rate, then they should be identified with the same expected detection time (see Section 4.4). In this experiment, we study how well our detection mechanism achieves this objective.

We consider two attackers termed Attacker 1 and Attacker 2, both of which periodically attack the same number of active mobiles based on an attack interval of length L s. For Attacker 1, we set L to be a constant L_c for all attack intervals, while for Attacker 2, we update L based on a uniform distribution $U(L_c - 2.5, L_c + 2.5)$. Thus, both attackers have the same expected attack rate but different arrival patterns of the attack traffic. We then measure the detection time for different values of L_c . To see how the results are affected when the detection time spans different numbers of attack intervals, we consider the cases where both attackers target 2 mobiles and




Figure 13: Detection time for different attack strategies: (a) 2 victims and (b) 5 victims.

5 mobiles at each attack interval. Experiment 1 shows that when the attack interval is slightly larger than 5 s, the detection times for attacking 2 mobiles and 5 mobiles respectively span more than one attack interval and less than one attack interval.

Figure 13 plots the (expected) detection time when L_c is varied. While Attacker 2 may result in significantly high detection time for some cases, both Attackers 1 and 2 have very close detection time in the expected sense.

Summary: We show that our detection mechanism requires small detection time to identify the source of a signaling attack, and hence allows early defense before any significant damage. Also, it maintains a very small false positive ratio. Furthermore, our detection mechanism remains robust toward different attack strategies that have the same attack intensity.

6. Enhancements to the Detection Mechanism and the Reaction

In this section, we address some enhancements to our detection mechanism so as to improve its scalability and enable it to handle more challenging attacks. We also describe a graceful approach for reacting to attacks that minimizes impact of false alarm yet maintain load at sustainable level.

6.1. Scaling up the Detection

Our detection mechanism keeps track of every remote host, and thus the storage requirement is proportional to the number of remote hosts, which can be up to the total number of hosts in the Internet. Although the amount of memory consumed by our simulation is not a vital concern due to the relatively small number of remote hosts in our traces as compared to the upper bound, our detection mechanism should be able to survive the abrupt surge of the number of remote hosts without exhausting the memory storage.

One approach is to bound the storage requirement. Suppose that we hash all remote hosts into a table of fixed size n . We define a *super node* that corresponds to the subset of remote hosts being hashed to the same entry. We then apply the detection algorithm to each super node rather than individual remote hosts. If a super node is not identified as malicious, obviously none of the remote hosts in the super node are malicious. On the other hand, if the super node is flagged malicious, we cannot tell if any of its associated remote hosts is malicious, because it is possible that some other remote hosts are malicious, or all remote hosts are benign but their aggregate behavior appears to be malicious.

In order to reduce false positives due to hash collisions, we introduce m parallel hash tables with independent hash functions (as in [8]). Each remote host is hashed to one entry in each of the m tables in parallel, i.e., any remote host is a member of m super nodes, each of which corresponds to a unique hash table (function). Now, a remote host is considered to be an attacker if its corresponding m super nodes are all marked malicious. In general, because the fraction of malicious remote hosts is small, a super node is less likely to have a malicious remote host. Moreover, the probability that a remote host collides with the same set of other hosts in all m hash tables decreases exponentially as m increases. Thus, the probability that a benign host has all its super nodes flagged malicious becomes very small. As a result, we can decrease (exponentially) the false positive ratio by increasing m .

By using m hash tables of size n , we can bound the total memory size to be $O(mn)$, while maintaining a small false positive ratio.

6.2. Coordinated Attacks

Our detection mechanism seeks to identify any remote host that triggers excessive amount of RAB setup/release. One way for an attacker to elude detection is to distribute the attack load over a large number of malicious remote hosts (e.g., by forming a botnet) and launch a coordinated attack. For instance, the malicious hosts may trigger virtual setups in a round-robin manner. Thus, each of these hosts results in benign inter-setup time samples that evade detection, but the aggregate damage can be highly severe.

One way to further enhance our current detection mechanism is to incorporate a *network telescope* [22]. We observe that although attackers may know the segment of IP addresses of a particular wireless network, they generally do not have the knowledge of active IPs unless they can monitor all traffic within the wireless network, which is not feasible in most cases. Most attack packets would eventually reach inactive IPs. Since we do not expect any legitimate traffic destined for inactive IPs, one way to deal with a coordinated attack is to treat all remotes hosts that send traffic to inactive IPs as one super remote host during the execution of our detection mechanism. By monitoring the inter-setup time samples between the

super remote host and the inactive IPs, we can detect the presence of any group of coordinated attackers.

6.3. Attacks using Spoofed Addresses

Although attacks from spoofed address have been shown to be unpopular [33], attackers may still send malicious traffic using spoofed addresses. In this case, they not only hide their real origins, but also trick the detection mechanism to make wrong decisions. For instance, if an attacker always spoofs the same address for all attack traffic, the detection mechanism will identify the spoofed address as malicious and then block traffic coming from the spoofed address. This is only a problem when the spoofed address belongs to an active remote host that has legitimate communication with the mobiles.

There have been studies (e.g., [15]) on identifying and filtering packets with spoofed addresses that we may employ before applying the detection algorithm. In future work, we analyze how to extend our detection mechanism to address spoofing attacks as well.

6.4. Reaction Mechanism

In response to an identified malicious remote host, one possibility is to filter the subsequent traffic originated from that remote host. On the other hand, we may choose not to take any action since we can endure the presence of a signaling attack as long as the signaling load is within a sustainable level. This reduces the chance of blocking legitimate hosts that are falsely marked malicious (see Section 5.2).

As the signaling load rises above a threshold, we may start blocking identified attackers that generate the most signaling load until the signaling load decreases to below the threshold. By doing so, we only take necessary actions to protect the control plane from being overloaded while minimizing the chance to penalize legitimate remote hosts.

7. Related Work

In this section, we review the related work on defending DoS attacks in both wireline and wireless networks.

Traditional DoS attacks are flooding-based such that an attacker generates *high-rate, high-volume* data traffic so as to deplete network resources using overwhelming data traffic. A number of defense mechanisms have been proposed for wireline networks, such as the pushback and traceback mechanisms [13, 25, 27]. Another class of DoS attacks is based on the *low-rate, high-volume* TCP attack [18], in which an attacker periodically generates high-volume packet bursts in order to force all TCP flows to repeatedly enter the retransmission timeout state. In this

paper, we describe a *low-rate, low-volume* signaling attack that targets 3G or equivalent wireless infrastructures. Unlike the low-rate TCP attack, the signaling attack does not necessarily have any periodic pattern.

The flooding-based DoS attack is a common threat to both wireline and wireless networks. Other forms of DoS attacks that specifically target wireless networks include packet-forwarding disruption [3, 10], base-station impersonation [21], control-channel congestion via a sufficient number of SMS messages [7], and depletion of mobile batteries [24]. In particular, the DoS attack in [7] saturates the control channels for SMS communication, while that in [24] keeps a mobile in a high-battery-consumption state. An independently published work that is closely related to ours is [29], which also identifies a DoS attack that exploits the expensive connection setup and release procedures. Our work is different from [29] in mainly two ways. First, we look at the problem from different perspectives: the discussion of [29] is based on the GSM (2G) architecture, while ours is based on the UMTS (3G) architecture. Second, and most importantly, we additionally propose a detection mechanism that is shown to be effectively detect the presence of the signaling attack.

Statistical online detection schemes have been studied by [16, 32] for countering DoS attacks. Specifically, [16] focuses on detecting malicious connection attempts based on Wald's test [31]. However, this detection scheme requires a priori probability distributions for the benign and malicious behaviors. In contrast, [32] propose a non-parametric CUSUM test to detect flooding-based DoS attacks based on periodic sampling. In this paper, we propose a different CUSUM-based method that is suitable for detecting the low-rate signaling attack and ensure that no attacker can intelligently escape from our detection mechanism.

8. Conclusions

We have presented a new DoS attack, called *signaling attack*, which targets 3G wireless networks. This attack works by exploiting the heavy-weight nature of signaling in 3G wireless networks. We have shown via trace-driven simulation that the signaling attack can substantially overload a wireless infrastructure using only minimal traffic. Because of its low-rate, low-volume property, the signaling attack can evade the detection of traditional counter-DoS systems. In view of this, we have proposed a statistical CUSUM-based detection mechanism to defend against the signaling attack. Using real-world traces, we have shown that our detection mechanism can identify the source of a signaling attack in a timely manner before the damage becomes aggravated, while producing very few false positives. In addition, our detection mechanism is robust as it depends solely on the additional signaling load and based on any assumed attack strategy.

Appendix

In Section 4.4, we claim that the score function $z(t)$ is a linear function of inter-setup time t . Here, we prove this claim formally. Recall that any two attackers that trigger the same number of virtual setups over time T should lead to the same cumulative score $\sum z(t)$ when the CUSUM test is conducted at time T . Formally, we present the following rule of defining $z(t)$ for this criterion:

Definition 1. (Rule of defining $z(t)$): Let $\{t_1, t_2, \dots, t_n\}$ and $\{t'_1, t'_2, \dots, t'_{n'}\}$ be any two sequences of inter-setup time samples of two different remote hosts. If $n = n'$ and $\sum_{i=1}^n t_i = \sum_{i=1}^{n'} t'_i$, then their respective cumulative scores should be the same, i.e., $\sum_{i=1}^n z(t_i) = \sum_{i=1}^{n'} z(t'_i)$.

To understand the above rule, we consider a UMTS network whose inactivity timeout for a RAB is set to be 5 s. Suppose that a remote host sends four packets to a mobile at times $T = 0, 6, 11.5,$ and 21 s, and hence introduces the inter-setup time samples are $t_1 = 6$ s, $t_2 = 5.5$ s, and $t_3 = 7$ s. Then the remote host will trigger three virtual setups within 21 s. The same effect is obtained if another remote host sends the four packets at times $T = 0, 7, 14,$ and 21 s. Therefore, both scenarios should lead to the same cumulative score $\sum_{i=1}^3 z(t_i)$.

Based on Definition 1, we can prove the following theorem.

Theorem 1. *Definition 1 holds if and only if $z(t)$ is a linear function of the inter-setup time t .*

Only if part: Let δt be some small positive value. Since $\sum_{i=1}^n t_i + \delta t = \sum_{i=1, i \neq j}^n t_i + (t_j + \delta t) = \sum_{i=1, i \neq k}^n t_i + (t_k + \delta t)$, we have $\sum_{i=1, i \neq j}^n z(t_i) + z(t_j + \delta t) = \sum_{i=1, i \neq k}^n z(t_i) + z(t_k + \delta t)$. This implies $\frac{z(t_j + \delta t) - z(t_j)}{\delta t} = \frac{z(t_k + \delta t) - z(t_k)}{\delta t}$. As δt tends to zero, the derivatives of $z(t_i)$ and $z(t_j)$ are the same. Since t_i and t_j can be any possible inter-setup time samples, the derivative of $z(t)$ is a constant, implying that $z(t)$ is linear.

If part: Let $z(t) = c_1 t + c_2$, where c_1 and c_2 are some constants. Given the inter-setup time samples $\{t_1, t_2, \dots, t_n\}$ and $\{t'_1, t'_2, \dots, t'_{n'}\}$, we have $\sum_{i=1}^n z(t_i) = \sum_{i=1}^n (c_1 t_i + c_2) = c_1 (\sum_{i=1}^n t_i) + n c_2 = c_1 (\sum_{i=1}^{n'} t'_i) + n' c_2 = \sum_{i=1}^{n'} (c_1 t'_i + c_2) = \sum_{i=1}^{n'} z(t'_i)$.

References

- [1] 3GPP. UTRAN Functions, Examples on Signaling Procedures (Release 1999), Jun 2002. TR 25.931 v3.7.0.
- [2] 3GPP. IP Transport in UTRAN, Dec 2003. TR 25.933 v5.4.0.

- [3] I. Aad, J.-P. Hubaux, and E. W. Knightly. Denial of Service Resilience in Ad Hoc Networks. In *Proc. of ACM MOBICOM*, Sep 2004.
- [4] M. Basseville and I. V. Nikiforov. *Detection of Abrupt Changes: Theory and Application*. Prentice Hall, 1993.
- [5] B. E. Brodsky and B. S. Darkhovsky. *Non-parametric methods in Change-Point Problems*. Kluwer Academic Publishers, 1993.
- [6] CDMA Development Group. <http://www.cdg.org>.
- [7] W. Enck, P. Traynor, P. McDaniel, and T. L. Porta. Exploiting Open Functionality in SMS-Capable Cellular Networks. In *Proc. of ACM CCS*, 2005.
- [8] C. Estan and G. Varghese. New Directions in Traffic Measurement and Accounting: Focusing on the Elephants, Ignoring the Mice. *ACM Trans. on Computer Systems*, 21(3):270–313, Aug 2003.
- [9] A.-B. Garcia, E. Garcia, M. Alvarez-Campana, J. Berrocal, and E. Vazquez. A Simulation Tool for Dimensioning and Performance Evaluation of the UMTS Terrestrial Radio Access Network. *Lecture Notes in Computer Science*, 2515, Nov 2002.
- [10] Y.-C. Hu, A. Perrig, and D. B. Johnson. Ariadne: A Secure On-Demand Routing Protocol for Ad Hoc Networks. In *Proc. of ACM MOBICOM*, Sep 2002.
- [11] IANA. IP Address Services. <http://www.iana.org/ipaddress/ip-addresses.htm>.
- [12] IEEE 802.16e Task Group. IEEE Standard for Local and metropolitan area networks Part 16: Air Interface for Fixed and Mobile Broadband Wireless Access Systems Amendment 2: Physical and Medium Access Control Layers for Combined Fixed and Mobile Operation in Licensed Bands and Corrigendum 1, Feb 2006. 802.16E-2005 & 802.16/COR1.
- [13] J. Ioannidis and S. M. Bellovin. Implementing Pushback: Router-Based Defense Against DDoS Attacks. In *Proc. of the Network and Distributed System Security (NDSS)*, Feb 2002.
- [14] R. Jain. *The Art of Computer Systems Performance Analysis: Techniques for Experimental Design, Measurement, Simulation, and Modeling*. John Wiley & Sons, New York, 1991.

- [15] C. Jin, H. Wang, and K. Shin. Hop-Count Filtering: An Effective Defense Against Spoofed DoS Traffic. In *ACM International Conference on Computer and Communications Security (CCS)*, Oct 2003.
- [16] J. Jung, V. Paxson, A. W. Berger, and H. Balakrishnan. Fast Portscan Detection Using Sequential Hypothesis Testing. In *IEEE Symposium on Security and Privacy 2004*, May 2004.
- [17] D. Kotz and K. Essien. Analysis of a Campus-Wide Network. In *Proc. of ACM MOBICOM*, Sep 2002.
- [18] A. Kuzmanovic and E. Knightly. Low-Rate TCP-Targeted Denial of Service Attacks and Counter Strategies. *IEEE/ACM Trans. on Networking*, 14(5), Oct 2006.
- [19] P. P. C. Lee, T. Bu, and T. Woo. On the Detection of Signaling DoS Attacks on 3G Wireless Networks. In *Proc. of IEEE INFOCOM*, May 2007.
- [20] C. Lindemann, M. Lohmann, and A. Thümmler. Adaptive Performance Management for Universal Mobile Telecommunications System Networks. *Computer Networks*, 38(4), Mar 2002.
- [21] U. Meyer and S. Wetzel. A Man-in-the-middle Attack on UMTS. In *Proc. of the ACM Workshop on Wireless Security (WiSe)*, Oct 2004.
- [22] D. Moore, C. Shannon, G. M. Voelker, and S. Savage. Network Telescopes: Technical Report. Technical Report CS2004-0795, UCSD, Jul 2004.
- [23] E. S. Page. Continuous Inspection Schemes. *Biometrika*, 41(1/2):100–115, Jun 1954.
- [24] R. Racic, D. Ma, and H. Chen. Exploiting MMS Vulnerabilities to Stealthily Exhaust Mobile Phone’s Battery. In *Proc. of SecureComm*, Aug 2006.
- [25] S. Savage, D. Wetherall, A. Karlin, and T. Anderson. Network Support for IP Traceback. *IEEE/ACM Trans. on Networking*, 9(3), Jun 2001.
- [26] C. SMith and D. Collins. *3G Wireless Networks*. McGraw-Hill, 2002.
- [27] A. C. Snoeren, C. Partridge, L. A. Sanchez, C. E. Jones, F. Tchakountio, B. Schwartz, S. T. Kent, and W. T. Strayer. Single-Packet IP Traceback. *IEEE/ACM Trans. on Networking*, 10(6), Dec 2002.

- [28] TIA/EIA/cdma2000. *Mobile Station - Base Station Compatibility Standard for Dual-Mode Wideband Spread Spectrum Cellular Systems*. Washington: Telecommunication Industry Association, 1999.
- [29] P. Traynor, P. McDaniel, and T. L. Porta. On Attack Causality in Internet-Connected Cellular Networks. In *USENIX Security Symposium (SECURITY)*, Aug 2007.
- [30] UMTS. *Release 5*. 3G Partnership Project.
- [31] A. Wald. *Sequential Analysis*. New York, Dover Publications, 1973.
- [32] H. Wang, D. Zhang, and K. G. Shin. Change-Point Monitoring for Detection of DoS Attacks. *IEEE Trans. on Dependable and Secure Computing*, 1(4), Dec 2004.
- [33] V. Yegneswaran, P. Barford, and J. Ullrich. Internet Intrusions: Global Characteristics and Prevalence. In *ACM SIGMETRICS*, Jun 2003.