An Analysis of DoS Attack Strategies Against the LTE RAN

Jill Jermyn[1], Gabriel Salles-Loustau[2], and Saman Zonouz[2]

[1] *Department of Computer Science, Columbia University New York, NY jill@cs.columbia.edu*
[2] *Department of Electrical and Computer Engineering, University of Miami, Miami, FL g.sallesloustau@umiami.edu, s.zonouz@miami.edu*

Received 1 March 2014; Accepted 15 April 2014;
Publication 2 July 2014

Abstract

Long Term Evolution (LTE) is the latest 3GPP mobile network standard, offering an all-IP network with higher efficiency and up to ten times the data rates of its predecessors. Due to an increase in cyber crime and the proliferation of mobile computing, attacks stemming from mobile devices are becoming more frequent and complex. Mobile malware can create smart-phone botnets in which a large number of mobile devices conspire to perform malicious activities on the cellular network. It has been shown that such botnets can cause a denial of service (DoS) by exhausting user traffic capacity over the air interface. Through simulation and with studies in a real-world deployment, this paper examines the impact of a botnet of devices seeking to attack the LTE network using different types of strategies. We quantify the adverse effects on legitimate users as the size of the botnet scales up in both sparsely and densely-populated cells for varying traffic Quality of Service (QoS) requirements. Our results show that a single attacker can drastically reduce the QoS of legitimate devices in the same cell. Furthermore, we prove that the impact of the attack can be optimized by tuning the attack strategy, leveraging the LTE uplink MAC scheduler.

Keywords: LTE, DoS, security, mobile malware, botnets.

Journal of Cyber Security, Vol. 3 No. 2, 159–180.
doi: 10.13052/jcsm2245-1439.323

1 Introduction

Smartphones are becoming increasingly popular as multipurpose portable computing devices that run a complete software stack from the operating system to the user-level applications. Based on reports by ComScore Inc. [30], 110 million Americans used smartphones in 2012, and smartphones constitute 47% of the total mobile communication devices. Smartphone applications serve various sensitive and critical functionalities. At the same time, they are often developed by possibly untrusted and inexperienced third-party developers that may introduce new attack vectors and exploitable vulnerabilities. The increasing popularity of smartphones along with emerging possible attack vectors and vulnerabilities has turned them into appealing targets for malicious adversarial parties. Real-world solutions need to be designed to provide smartphone platforms with effective and practical security services. However, existing smartphone platforms generally face computational and storage limitations that hinder permanent deployment of comprehensive and heavyweight intrusion prevention and detection solutions.

As shown by many past research projects, smartphones remain vulnerable to various exploitation techniques. Intrusions on smartphones occur in one or more forms of the following three categories [14]. First, adversaries may violate user data confidentiality by exfiltrating device-resident, sensitive data to malicious end points. Second, attackers may cause data integrity violation via malicious unauthorized data modification within the user smartphones. Finally, the adversaries may target the availability of the system services provided by the smartphone platform. In this paper, we concentrate on infected devices targeting availability in the cellular network.

Long Term Evolution (LTE) is the latest cellular network standard for high-speed mobile devices. In 2013 there were 200 million devices connected over LTE [3], and this number is expected to surge to 1 billion by 2016 [22]. Not only are people reliant on LTE for their voice and data services, but with the rise of the Internet of Things (IoT), LTE has become an important resource for Machine-to-Machine (M2M) communication. According to Gartner there will be 26 billion devices on the IoT by 2020 [1]. Clearly, any impact on the availability of LTE services could cause catastrophic repercussions for the great number of devices that rely on them.

In this paper, we introduce a new set of denial of service (DoS) attack strategies that handheld devices could potentially use against LTE technologies. Malware executing these strategies would be neither particularly designed against a specific executable nor architectured against generic

computing devices. Instead, it would make use of how the LTE access networks manage a large-scale networked system consisting of thousands of end-user devices. Consequently, all devices utilizing LTE solutions become potentially vulnerable. Due to the frequency spectrums available to LTE networks, there exists a limited capacity that the air interface can handle, making the whole infrastructure vulnerable to physical layer DoS intrusions.

We investigate uplink data traffic and examine how the LTE Medium Access Control (MAC) scheduler allocates resources for devices during both normal benign and adversarial scenarios. Furthermore, we determine the optimal size of a botnet in a single cell that is needed to significantly hinder availability of service/degrade legitimate customer Quality of Service (QoS).

We have implemented and evaluated a real working prototype of our proposed attack strategies on the most recent Android platforms. Our experimental results empirically prove the feasibility of the attacks against LTE networks. They also show that the usability of smartphone devices could be significantly affected through installation of malicious legitimate-looking applications on the device. Our simulation results from a large-scale network show that deployment of these strategies could potentially impact the data network efficiency for all legitimate end-user devices sharing the same cell.

In summary, the contributions of this paper are the following:

- We introduce novel LTE-specific denial of service attack strategies against smartphone devices that make use of LTE networks. The attacks are not detectable by traditional signature-based detection solutions.
- We demonstrate the attack effectiveness through simulation and determine the optimal botnet size and type of flooding traffic that would be necessary to severely impact legitimate users in the same cell.
- We implement a working prototype for the most recent Android platform and evaluate its efficiency and practical deployability in a real-world setting.

This paper is organized as follows. Section 2 describes the related literature and real-world intrusions against resource-limited smartphone devices, followed by an explanation of how cellular botnets can be formed and a discussion on instances of them in the wild. Section 3 gives a high-level background on state-of-the-art LTE access networks and describes uplink scheduling in the LTE Radio Access Network (RAN). Our simulation experimental setup and results from a set of attack strategies are presented in Section 4, whereas Section 5 talks about practical applications and implementation in a real-world setting. Finally, Section 6 concludes the paper.

2 Related Work

In this section we review the past work on adversarial remote hacking techniques against smartphones as well as smartphone-originating attacks against the cellular and data network infrastructures.

Mylonas et al. [27] perform a fairly thorough empirical proof of how feasible and simple it is for average programmers to develop smartphone specific malware. The authors continue with an in-depth investigation of the deployed security mechanisms within various smartphone platforms and conduct a comparison among those platforms. A similar study was conducted by Jeon et al. [18], who performed a security analysis of smartphones and proposed potential countermeasures. Marforio et al. [24] design a new covert channel and application collusion attack that was previously unknown against the Android platforms. Using their proposed attack vector, the attackers could penetrate into the smartphone devices and steal sensitive user data through development of several individually legitimate-looking applications that co-operatively gain a sufficient set of permissions for user privacy violation. Aviv et al. [8] design and develop a novel intrusion against handheld tools that make use of touch screens using the oily residues on those screens. The authors show that it is possible to infer the user passwords if he/she enters it using the touch screen.

There have also been attacks proposed that could originate from smartphone devices against the cellular and data network infrastructures. Ricciato et al. [31] review potential and previously proposed denial of service attack models specifically designed against cellular and data networks. The authors discuss the trade-offs between optimality and robustness of the designed cellular networks. They determine that a single attacker can create traffic profiles that are capable of straining the entire network infrastructure. Traynor et al. [34] characterize the impact of the large-scale compromise and coordination of mobile phones in attacks against the core of cellular networks. The authors demonstrate that a botnet comprised of 11K phones could potentially degrade service to area-code sized regions by 93%. In another similar effort, Enck et al. [13] empirically show the feasibility of SMS-based attacks against the cellular networks to induce a denial of voice service to cities the size of Washington D.C. and Manhattan.

2.1 Botnet Threats in Cellular Networks

The vast expansion and increasing popularity of highly-capable but largely insecure smartphone devices that are often interconnected with the Internet is

a significant threat against large-scale cellular networks. In this section, we describe the well-known botnet intrusions in such cellular infrastructures and review the results from the past related literature.

Generally, botnets are defined as a collection of Internet-connected compromised computing devices, so-called zombies, communicating with other compromised systems in order to perform (often malicious) tasks. Unlike network worms, zombies are not autonomous and need to be ordered regarding what to do at each time instant. Such control orders may come from various communication channels such as an Internet Relay Chat (IRC) channel where the master sends control commands to be executed by the distributed zombies. Of the typical botnet tasks, one could mention large-scale spamming where all the zombies are ordered to send spam emails to the same target address, potentially causing a denial of service. Other similar intrusions may target Internet core services such as DNS [11] and BGP [21]. Traditionally, botnets have targeted desktop computer systems; however, with the increasing popularity of vulnerable and capable smartphone devices, the number of smartphone-specific botnets have recently risen by a significant factor.

In particular, smartphone botnets could potentially be more damaging to the underlying network infrastructures because cellular networks have more rigid hierarchical dependencies and hence are less likely to withstand similar misbehaviors. The past academic efforts investigating the possibility of a large-scale botnet attack against cellular networks have mainly consisted of two major categories [34]. First, researchers have attempted to explore whether the lack of authentication for signaling traffic in the wired network would enable an attacker with a physical connection to cause significant damage [20]. Second, there have been efforts to determine whether the same amount of damage is feasible by gathering a large set of compromised wireless devices and trying to either saturate the cellular network [35] or make use of the compromised smartphones as a spam generator to attack Internetbased resources [26]. The authors of [25] show that the threats described in [34] are concrete, and they demonstrate the ease of creating a mobile botnet on popular smartphones models.

There have been several instances of mobile botnets seen in the wild, most recently targeting the Android platform. MisoSMS, uncovered in December 2013, is one of the largest mobile botnets yet seen, stealing SMS messages and emailing them to a command-and-control infrastructure located in China [15]. One of the first Android malware to exhibit botnet-like capabilities was Geinimi, discovered in 2011, where a remote server had the ability to control and send instructions to infected devices [17]. However, Android is

not the only platform vulnerable to such malware. iKee.B, released throughout several European countries in 2009, spread to jailbroken iPhones by using the default SSH password. The malware stole sensitive information and performed malicious activities on infected devices by controlling them through a Lithuanian botnet server [28].

On a more comprehensive investigation of the smartphone botnets, Traynor et al. [34] evaluated the impact of the large-scale infection and co-ordination of mobile phones in attacks against the core of a cellular network, namely denial of service attacks using selected service requests on the central repository of user location and profile information in the network. According to their results, such attacks could degrade the core services even in cellular networks with capable databases to the extent of approximately 75% when the size of the botnet reaches around 140K zombies.

Cambiaso et al. [10] present a comprehensive survey of DoS intrusions targeting the data service network. The authors classify those intrusions into two main categories. First, attacks that involve high-bandwidth, flood-based approaches exploiting vulnerabilities of networking and transport protocol layers. Second, slow-rate attacks that exploit vulnerabilities of application layer protocols to accomplish DoS objectives. Specifically, there have been several real-world and experimental malware samples performing denial of service intrusions. Dondyk [12] presents a DoS attack against smartphones that prevent non-technical smartphone users from utilizing data services by exploiting the connectivity management protocol when encountered with a WiFi access point. Gobbo et al. [16] describe a DoS attack against the data provider network via an unauthenticated injection of malicious traffic in the mobile operator's infrastructure that causes significant service degradation.

[29] exploits opportunistic scheduling in 3G networks by showing that devices can report false channel condition reports. The authors demonstrate that only five malicious devices in a 50-user cell can consume up to 95% of timeslots, in effect causing 2 second end-to-end packet delay on VoIP applications that make them virtually useless. There has been some recent research on DoS attacks against the LTE air interface. For example, the authors of [19] use simulations to determine the number of attackers needed to degrade service of legitimate VoIP users. Although peripherally similar to our work in scope, this research does not examine the impact of a botnet due to various QoS requirements nor does it provide results from a real-world environment.

Bassil et al. [9] simulate an attack against LTE where malicious devices request high-bandwidth GBR bearers while having a low Modulation and Coding Scheme (MCS) index. They demonstrate denial of service for

legitimate TCP-based applications with two malicious devices and show that high priority voice bearers are not affected by attacks since they preempt video bearers. Their results are insightful and support many of our claims in this paper, yet their experiments do not show results for variable sized botnets nor a variable number of legitimate users. Additionally, the authors only consider scenarios where malicious devices request bandwidth for high QoS applications, such as conversational video, and fail to demonstrate that botnets demanding low QoS applications can severely impact resource availability as well. In our paper, however, we quantify the impact of a botnet on legitimate users as the number of malicious devices scales for both lightly-used and densely-populated cells, and for both low QoS and high QoS traffic.

3 LTE Radio Access Network

The LTE radio interface, also known as the Uu interface, lies between the eNodeB and the User Equipment (UE). Offering high peak transmission rates, the physical layer implements Orthogonal Frequency-Division Multiple Access (OFDMA) in the downlink and Single-Carrier FDMA (SC-FDMA) in the uplink. Data and Signaling Radio Bearers (DRB and SRB) transmit user-plane and control-plane traffic on the air interface. DRBs support many different types of service requirements, for example voice call and mobile broadband access, through their QoS configuration [32]. QoS describes the combination of requirements for categories of data traffic, including latency, guaranteed bit rate (GBR), jitter, and error rate. Classes of traffic can be grouped by a QoS Class Identifier (QCI) which indicates the traffic's requirements for delay, priority, and error rate. Figure 1 shows a list of standard QCI profiles as defined by 3GPP [7]. End-to-end QoS is important for distinguishing certain real-time services, such as voice, from basic broadband access and providing them with reliable delivery.

QoS scheduling over the air interface in LTE in both the uplink and downlink direction is handled by the MAC scheduler in the eNodeB [23]. Non-GBR bearers, which do not guarantee minimum resources, are allocated bandwidth depending on the QCI of the service and the current cell utilization. They are typically provided with service according to fairness criteria. Examples of applications using Non-GBR bearers are web browsing and FTP. GBR bearers, on the other hand, are granted throughput up to an agreed-on guaranteed bit rate depending on the current cell utilization and in some cases may force a lower priority user to relinquish services. Real time applications, voice, and streaming video typically use GBR bearers. The job

QCI	Resource Type	Priority	Packet Delay Budget	Packet Error Loss Rate	Example Services
1	GBR	2	100 ms	10^{-2}	Conversational Voice
2		4	150 ms	10^{-3}	Conversational Video (Live Streaming)
3		3	50 ms	10^{-3}	Real Time Gaming
4		5	300 ms	10^{-6}	Non-Conversational Video (Buffered Streaming)
5	Non-GBR	1	100 ms	10^{-6}	IMS Signalling
6		6	300 ms	10^{-6}	Video (Buffered Streaming) TCP-based (e.g., www, e-mail, chat, ftp, p2p file sharing, progressive video, etc.)
7		7	100 ms	10^{-3}	Voice, Video (Live Streaming) Interactive Gaming
8		8	300 ms	10^{-6}	Video (Buffered Streaming) TCP-based (e.g., www, e-mail, chat, ftp, p2p file
9		9			sharing, progressive video, etc.)

Figure 1 3GPP standardized QCI Characteristics

of the MAC scheduler is to allocate the air interface resources so that bearer QoS requirements are met and priorities among different QCIs are sustained [6]. Although the exact algorithms used in scheduler implementation are vendor specific, they must balance fairness and QoS when making allocation decisions.

Studying uplink scheduling related to LTE network availability is particularly important with the rise of M2M communications using 4G networks, as most M2M applications are uplink dominant and will therefore make high demands on uplink bandwidth [33]. The manner in which the MAC scheduler handles such voluminous requests for uplink resources will play an important role in how customers are affected by high service demands, whether they be of legitimate or malicious origin. In addition, QoS distinctions in LTE will become increasingly important for M2M applications that have strict requirements for delay and reliability of service, such as medical devices and smart grid.

4 Simulation Experiments and Results

This section describes our set of simulation experiments for testing multiple DoS attack strategies against the LTE RAN. As described in Section 3, the LTE MAC scheduler is responsible for making bandwidth allocation decisions and is influenced by the particular traffic QoS requested for a bearer. Our experiments illustrate that these scheduling decisions can be exploited to optimize certain DoS attack strategies. We show that the network fails to recognize a botnet as malicious and consequently tries to furnish it with resources by reducing legitimate device throughput.

4.1 Experimental Setup

All of our experiments were performed on the OPNET Modeler's LTE model [4], a comprehensive platform that is standards compliant at Release 8. The model consists of several network elements, including the UEs, eNodeB (LTE base station), Evolved Packet Core (EPC), and IP servers. Traffic sent from UEs is customizable with regard to QoS, intensity, start time, and duration. For our experiments, we selected a 3MHz LTE deployment with a single cell and designated a subset of UEs as attackers that attempt to saturate the RAN with large amounts of uplink traffic at a particular time during the simulation. The remaining UEs depict legitimate devices, as their traffic profiles are set within reasonable typical usage patterns. Figure 2 shows

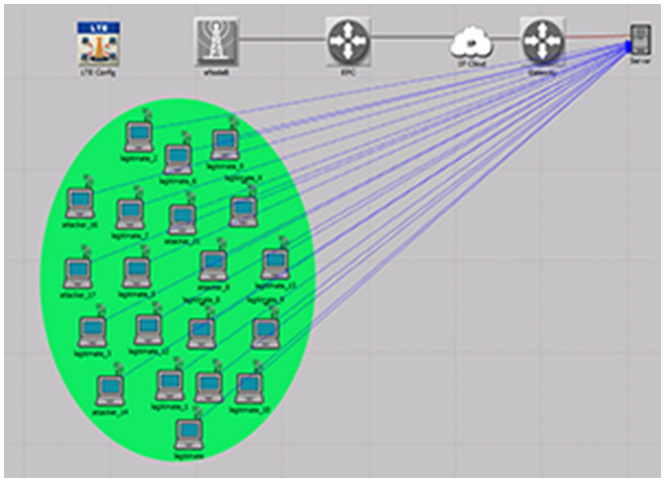

Figure 2 Simulation screenshot of legitimate and malicious UEs in a single LTE cell

a screenshot of an OPNET simulation scenario containing both legitimate and malicious UEs in a single cell.

We simulated several attack scenarios to determine the impact of various types of botnets on legitimate user QoS. In each experiment we modified the size of the botnet, the number of legitimate users in the cell, and the QoS requirements of the traffic sent by both malicious and legitimate devices. Video conferencing is indicated by traffic with QCI 2, while QCI 9 traffic represents web browsing or file transfer.

4.2 Simulation Experiments and Results

Our first experiment deployed a lightly-used cell of 20 legitimate users each sending traffic to a remote IP server at 100,000 b/s for 20 minutes. During the traffic period a botnet of malicious devices attempts to flood the network by sending 2 Mbps per device of uplink traffic. Table 1 (a) and (b) summarize the impact of the botnet on the cell's legitimate devices for an increasing number of malicious devices and for traffic with varying QCI. Clearly the impact of the attack fluctuates based on the QCI of the traffic sent by both the legitimate and malicious devices. In general, when legitimate devices send traffic with a low QCI, a botnet can cause a complete denial of service for those devices. For example, Table 1 (a) shows that a botnet of only one device precipitates zero throughput for 80% of the legitimate users. We see a similar impact even when the attacker's traffic has a high QCI. The reason for the complete rejection of allocated bandwidth is that the scheduler is unable to furnish the guaranteed bit rate of the legitimate UEs and therefore denies them completely. It is also interesting to see that when attackers send high QCI traffic, increasing the botnet size by 20 times doesn't exacerbate the impact any more than with a single malicious device. However, when the attackers send low QCI traffic, we see a further deterioration of legitimate throughput when the botnet reaches 15 devices.

When legitimate traffic has a high QCI, shown in Table 1 (b), the botnet attack inflicts different consequences. We see a drastic reduction in legitimate UE QoS when the attackers send low QCI traffic. Although all legitimate devices are allocated some uplink bandwidth, their throughput declines so significantly during the attack that their applications would be virtually unusable. For example, the throughput of legitimate users is reduced to less than half when there are only two malicious devices sharing the same cell. The impact of the attack proves significantly worse as the size of the botnet increases. Table 1 (b) shows that when the botnet

Table 1 Uplink throughput for legitimate UEs in a lightly-used cell for different QoS traffic and varying botnet size, (a) Legitimate traffic: QCI 2, (b) Legitimate traffic: QCI 9

Size of Botnet [number of devices]	**Attacker Traffic: QCI 9 [percentage of legitimate devices: uplink throughput]**	**Attacker Traffic: QCI 2**
1	20%: 100,000 b/s 80%: 0 b/s	20%: 100,000 b/s 80 %: 0 b/s
2	20%: 100,000 b/s 80%: 0 b/s	20%: 100,000 b/s 80 %: 0 b/s
5	20%: 100,000 b/s 80%: 0 b/s	20%: 100,000 b/s 80 %: 0 b/s
10	20%: 100,000 b/s 80%: 0 b/s	20%: 100,000 b/s 80 %: 0 b/s
15	20%: 100,000 b/s 80%: 0 b/s	10%: 100,000 b/s 10 %: 28,000 b/s 60%: 0 b/s
20	20%: 100,000 b/s 80%: 0 b/s	5%: 31,000 b/s 15 %: 13,000 b/s 80%: 0 b/s

(a)

Size of Botnet [number of devices]	**Attacker Traffic: QCI 9 [legitimate uplink throughput]**	**Attacker Traffic: QCI 2**
1	71,500 b/s	62,500 b/s
2	65,700 b/s	43,000 b/s
5	58,000 b/s	15,000 b/s
10	46,000 b/s	4,100 b/s
15	40,000 b/s	190 b/s
20	35,300 b/s	190 b/s

(b)

scales from 1 to 5 devices, the bit rate of each legitimate user is reduced by 76%. When it reaches 15 malicious devices, the legitimate throughput becomes less than 0.2% the traffic rate actually sent by each legitimate device.

Our second set of simulation experiments sought to disclose the effect of a botnet on a single legitimate device that attempts to transmit data during an ongoing attack, again for a lightly-used cell of 20 legitimate devices. The device sends 200,000 b/s uplink traffic once all malicious devices initiate the attack (the remaining legitimate devices still send 100,000 b/s uplink traffic each). Table 2 summarizes the results of these experiments. In all cases where the attackers send low QCI traffic, the single legitimate UE is completely

Table 2 Uplink throughput for a legitimate UE in a lightly-used cell that attempts to connect during an ongoing attack for different QoS traffic and varying botnet size, (a) Legitimate traffic: QCI 2, (b) Legitimate traffic: QCI 9

Size of Botnet [number of devices]	Attacker Traffic: QCI 9 [legitimate uplink throughput]	Attacker Traffic: QCI 2 [legitimate uplink throughput]
1	0 b/s	0 b/s
2	0 b/s	0 b/s
5	0 b/s	0 b/s
10	0 b/s	0 b/s
15	0 b/s	0 b/s
20	0 b/s	0 b/s

(a)

Size of Botnet [number of devices]	Attacker Traffic: QCI 9 [legitimate uplink throughput]	Attacker Traffic: QCI 2 [legitimate uplink throughput]
1	120,000 b/s	96,000 b/s
2	120,000 b/s	71,000 b/s
5	66,000 b/s	20,000 b/s
10	46,000 b/s	4,000 b/s
15	48,000 b/s	190 b/s
20	40,890 b/s	190 b/s

(b)

denied bandwidth, even with only one attacker in the cell. When sending traffic requiring high QCI, the UE is also more severely impacted by the botnet than are the other legitimate devices that start transmitting before the attack. With a single attacker sending low QCI traffic, the legitimate device's throughput is cut to 48% of its requested rate, whereas the other legitimate UEs are able to achieve 62.5% of their desired throughput. A botnet of 15 attackers reduces the UE's throughput to 0.1%, at 190 b/s. Although the impact of the botnet is worse when the attackers send low QCI traffic while legitimate devices send high QCI traffic, everyone sending high QCI traffic generates significantly adverse consequences for the single legitimate device as well. When in the same cell as a botnet of 20 devices, the single legitimate user is granted only 34% of the bit rate it can send when there is a botnet of one device. This single client is therefore penalized more than the bulk of legitimate devices in the same cell that are allocated about 49% of their bit rate during an attack with one malicious device. Additionally, when the size of the botnet reaches 10 attackers, the single legitimate device's throughput reaches that of the other legitimate clients, even though it is requesting twice the bit rate. These results indicate that the network views the botnet as legitimate and therefore doesn't

reduce the resources allocated to the botnet when the legitimate user attempts to transmit during the attack.

In our third set of simulation experiments, we examined the impact of a botnet on a densely-populated cell made up of 200 legitimate devices sending 20,000 b/s of uplink traffic and a varying number of attackers transmitting 2Mbps each. Clearly the results in Table 3 (a) and (b) indicate that a much smaller botnet can cause even more drastic ramifications than in a lightly used cell. This phenomenon is due to the eNB trying to accommodate the malicious UEs by reducing the throughput of devices already present in the cell. For example, it takes only one attacker sending low QCI traffic to spark a complete DoS for 98% of the legitimate population, with the remaining 2% of devices able to throughput only 0.5% of their requested bit rate.

Table 3 Uplink throughput for legitimate UEs in a densely-populated cell for different QoS traffic and varying botnet size, (a) Legitimate traffic: QCI 2, (b) Legitimate traffic: QCI 9

Size of Botnet [number of devices]	**Attacker Traffic: QCI 9** [percentage of legitimate devices: uplink throughput]	**Attacker Traffic: QCI 2** [percentage of legitimate devices: uplink throughput]
1	2%: 105 b/s	2%: 104 b/s
	98%: 0 b/s	98%: 0 b/s
2	2%: 105 b/s	2%: 104 b/s
	98%: 0 b/s	98%: 0 b/s
5	2%: 105 b/s	2%: 104 b/s
	98%: 0 b/s	98%: 0 b/s
10	2%: 105 b/s	2%: 104 b/s
	98%: 0 b/s	98%: 0 b/s
15	2%: 104 b/s	2%: 104 b/s
	98%: 0 b/s	98%: 0 b/s
20	2%: 100 b/s	2%: 100 b/s
	98%: 0 b/s	98%: 0 b/s

(a)

Size of Botnet [number of devices]	**Attacker Traffic: QCI 9** [percentage of legitimate devices: uplink throughput]	**Attacker Traffic: QCI 2** [percentage of legitimate devices: uplink throughput]
1	100%: 3,800 b/s	100%: 3,100 b/s
2	100%: 3,800 b/s	100%: 2,400 b/s
5	100%: 3,750 b/s	100%: 821 b/s
10	100%: 3,660 b/s	100%: 268 b/s
15	100%: 3,567 b/s	66%: 103 b/s
		34%: 0 b/s
20	100%: 3,500 b/s	57%: 103 b/s
		43%: 0 b/s

(b)

Similarly, when legitimate devices send high QCI traffic and attackers low QCI traffic (Table 3 (b)), it takes only one attacker to reduce legitimate throughput by 85%, yet in a lightly used cell, the same reduction requires 15 attackers. It is interesting to observe that when legitimate clients send low QCI traffic, the impact of the botnet is the same regardless of the QoS requirements of the attack traffic, as shown in Table 3 (a). However, when legitimate users send high QCI traffic, a botnet sending low QCI traffic causes a significantly worse impact as its scale increases than one sending high QCI traffic. In both cases, nevertheless, it takes only one attacker to reduce the legitimate throughput by 81–84%.

Figure 3 gives some insight on how legitimate user QoS is impacted by a botnet as it scales up, measured in terms of uplink packet latency between the device and eNodeB. In this experiment, legitimate users in both lightlyused (20 UEs) and densely-populated (200 UEs) cells attempt to send QCI 9 uplink traffic while a botnet of varying size sends QCI 2 traffic. As the results indicate, the larger the botnet, the greater the impact. It is interesting to note that it takes a smaller botnet to adversely affect legitimate users in a denselypopulated cell than it does in a lightly-used cell. Additionally, a botnet with only one device can produce such a high packet latency for legitimate clients sending QCI 9 traffic that it would result in their applications, such as web browsing or FTP, appearing unresponsive or even being unusable.

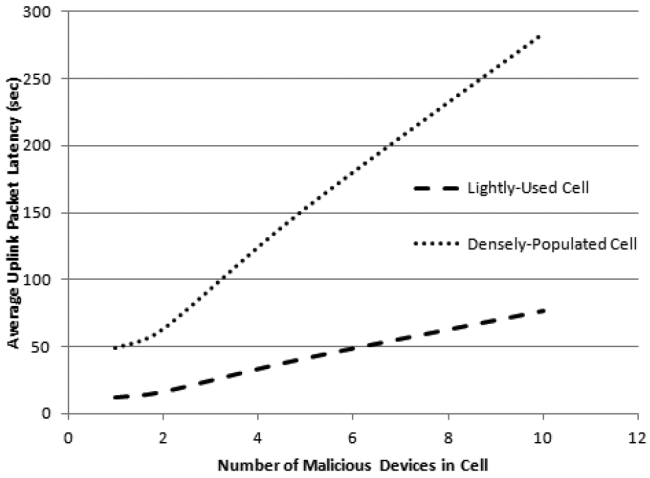

Figure 3 Average uplink packet latency over the air interface during an attack of variable sized botnets on densely-populated and lightly-used cells

5 Mobile Application Implementation and Results

This section describes the implementation of an Android application and a set of tools we have built to assess the effects of such attacks on existing infrastructure, along with an analysis of our experimental results. We used multiple smartphones connected to the same 4G LTE antenna and observed the transmission rate of a given file sent from a legitimate device to a server as we modified the traffic generated by the other surrounding devices. For every scenario, we recorded the data transmitted by the legitimate phone and used Speedtest.net App [5] to perform allocated bandwidth tests.

We implemented and deployed an application on a legitimate device that uploads a file to a server through a TCP connection and used a network capture tool (*tcpdump*) on the device to record the rate of traffic leaving the phone. Two separate devices acted as network flooders that utilized a UDP network traffic generator [2] and were configured to send traffic to a destination different than that of the smartphone. We captured the upload bit rate originating from the legitimate device and recorded the elapsed transmission time on each attempt. We also ran bandwidth tests for different network load on the antenna.

Table 4 shows the legitimate device file transfer rate we obtained during an attack with one and two malicious flooding devices as compared to during normal operations. We observe a degradation of the service between the non-flooded antenna test and the flooded antenna tests both in terms of bandwidth and response time. We observe an increase up to 19ms of response time (e.g., Ping test) and a diminution of upload bandwidth up to 3.64Mbps.

Next, we ran our flooder during an ongoing legitimate file transfer to observe the impact of a suddenly loaded antenna. Figure 4 shows the upload throughput for a file transfer on the legitimate device under normal circumstances (from 0s to 53s), while the antenna is flooded by one device

Table 4 Legitimate device file transfer rate results with one and two malicious flooding devices compared to during normal usage

	Average upload rate for a file transfer	**Speedtest.net App Results**
Normal Usage	5.12Mbps	Upload: 6.43Mbps Ping: 74ms
One Flooder	3.39Mbps	Upload: 3.63Mbps Ping: 93ms
Two Flooders	2.32Mbps	Upload: 2.79Mbps Ping: 84ms

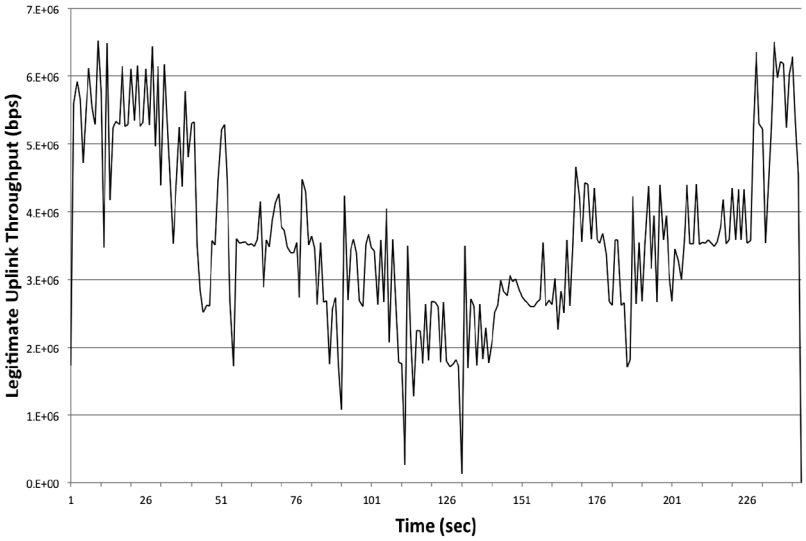

Figure 4 Legitimate throughput for a file upload through an antenna with different load over time

(from 53s to 109s), two devices (from 109s to 161s), again with one device (from 161s to 228s), and again under normal load (from 228s to 244s). We can observe a drop in the upload throughput while the flooders are transmitting, followed by a return to normal when they are stopped. In the worst case, we see normal throughput reduced up to 12,000 times when there are 2 flooders transmitting.

Although our results using existing infrastructure are small-scale, they prove the feasibility of our proposed attack strategies and show that a single malicious device has the potential to degrade legitimate user operations.

6 Conclusion

The rising popularity of smartphones has proliferated an abundance of mobile malware that could potentially perform large-scale coordinated attacks on communication infrastructure. Due to the limited frequency spectrum available to cellular networks, the physical layer is potentially vulnerable to denial of service attacks that can impact the bandwidth availability to all users in a cell. In this paper we examine several DoS attack strategies against the LTE RAN and study how the MAC uplink scheduler enables certain flavors of the attack, depending on the requested QoS of clients and the population

of a cell, to be more effective than others. We study a variety of attack strategies in which we vary the traffic QoS requirements of legitimate and malicious devices in lightly-used and densely-populated cells for increasing botnet sizes. Our simulation results indicate that a single attacker is capable of significantly reducing the QoS experienced by legitimate devices in the same cell and, using certain strategies, inducing a complete denial of service for those clients. Since the network views the malicious UEs as benign, it tries to accommodate them as much as possible by assigning resources and reducing the throughput of legitimate devices present in the cell. We follow our sim-ulation experiments with a real working prototype on the Android platform that proves the feasibility of our proposed attack strategies and demonstrates the impact legitimate users can experience during the attacks.

References

[1] Gartner Says the Internet of Things Installed Base Will Grow to 26 Billion Units By 2020. http://www.gartner.com/newsroom/id/2636073.

[2] gen-send: A Simple UDP Traffic Generater Application. http://www.citi.umich.edu/projects/qbone/generator.html.

[3] Global LTE Subscription Growth. http://www.4gamericas.org/index.cfm?useaction=page&pageid=2197.

[4] OPNET Modeler. http://www.opnet.com/solutions/network_rd/modeler.html.

[5] Speedtest.net App. http://www.speedtest.net/mobile/.

[6] LTE eNodeB MAC Scheduler Interface. White paper, Roke, 2009. http://www.roke.co.uk/resources/datasheets/108-lte-mac-scheduler-interface.pdf.

[7] 3rd Generation Partnership Project; LTE; Technical Specification Group Services and System Aspects. Policy and charging control architecture; 3gpp ts 23.203. v12.3.0, 2012.

[8] Adam J Aviv, Katherine Gibson, Evan Mossop, Matt Blaze, and Jonathan M Smith. Smudge attacks on smartphone touch screens. In *Proceedings of the 4th USENIX conference on Offensive technologies*, pages 1–7. USENIX Association, 2010.

[9] R. Bassil, I.H. Elhajj, A. Chehab, and A. Kayssi. A resource reservation attack against lte networks. In *Communications and Information Technology (ICCIT), 2013 Third International Conference on*, pages 262–268, June 2013.

[10] Enrico Cambiaso, Gianluca Papaleo, Giovanni Chiola, and Maurizio Aiello. Slow dos attacks: definition and categorisation. *International Journal of Trust Management in Computing and Communications*, 1(3): 300–319, 2013.

[11] David Dagon, Manos Antonakakis, Kevin Day, Xiapu Luo, Christopher P Lee, and Wenke Lee. Recursive dns architectures and vulnerability implications. In *NDSS*, 2009.

[12] E. Dondyk and C.C. Zou. Denial of convenience attack to smartphones using a fake wi-fi access point. In *Consumer Communications and Networking Conference (CCNC), 2013 IEEE*, pages 164–170, 2013.

[13] William Enck, Patrick Traynor, Patrick McDaniel, and Thomas La Porta. Exploiting open functionality in sms-capable cellular networks. In *Proceedings of the 12th ACM Conference on Computer and Communications Security*, CCS '05, pages 393–404, New York, NY, USA, 2005. ACM.

[14] Adrienne Porter Felt, Matthew Finifter, Erika Chin, Steve Hanna, and David Wagner. A survey of mobile malware in the wild. In *Proceedings of the 1st ACM Workshop on Security and Privacy in Smartphones and Mobile Devices*, SPSM '11, pages 3–14, New York, NY, USA, 2011. ACM.

[15] Anthony Freed. Misosms malware sends your text messages to attackers in china, 2013. http://www.tripwire.com/state-of-security/top-security-stories/misosms-malware-sends-text-messages-china/.

[16] Nicola Gobbo, Alessio Merlo, and Mauro Migliardi. A denial of service attack to gsm networks via attach procedure. In *Security Engineering and Intelligence Informatics*, pages 361–376. Springer, 2013.

[17] George Hulme. Geinimi android malware has 'botnet-like' capabilities, 2011. http://www.csoonline.com/article/650866/geinimi-android-malware-has-botnet-like-capabilities?source=rss_cso_exclude_net_net.

[18] Woongryul Jeon, Jeeyeon Kim, Youngsook Lee, and Dongho Won. A practical analysis of smartphone security. In *Human Interface and the Management of Information. Interacting with Information*, pages 311–320. Springer, 2011.

[19] M. Khosroshahy, Dongyu Qiu, and M.K. Mehmet Ali. Botnets in 4g cellular networks: Platforms to launch ddos attacks against the air interface. In *Mobile and Wireless Networking (MoWNeT), 2013 International Conference on Selected Topics in*, pages 30–35, 2013.

[20] Kameswari Kotapati, Peng Liu, and Thomas F LaPorta. Cata practical graph & sdl based toolkit for vulnerability assessment of 3g networks.

In *Security and Privacy in Dynamic Environments*, pages 158–170. Springer, 2006.

[21] Mohit Lad, Ricardo Oliveira, Beichuan Zhang, and Lixia Zhang. Understanding resiliency of internet topology against prefix hijack attacks. In *Dependable Systems and Networks, 2007. DSN'07. 37th Annual IEEE/IFIP International Conference on*, pages 368–377. IEEE, 2007.

[22] Lam, Wayne. Wireless Communication Report-4G-LTE Landscape. https://technology.ihs.com/413870/wireless-communications-report-4g-lte-landscape-2012.

[23] LTE; Evolved Universal Terrestrial Radio Access (E-UTRA). Medium access control (mac) protocol specification. 3gpp ts 36.321. v12.0, 2013.

[24] Claudio Marforio, Aurélien Francillon, Srdjan Capkun, Srdjan Capkun, and Srdjan Capkun. *Application collusion attack on the permission-based security model and its implications for modern smartphone systems*. Department of Computer Science, ETH Zurich, 2011.

[25] Collin Mulliner and Jean-Pierre Seifert. Rise of the iBots: 0wning a telco network. In *Proceedings of the 5th IEEE International Conference on Malicious and Unwanted Software (Malware)*, 2010.

[26] Collin Mulliner and Giovanni Vigna. Vulnerability analysis of mms user agents. In *Computer Security Applications Conference, 2006. ACSAC'06. 22nd Annual*, pages 77–88. IEEE, 2006.

[27] Alexios Mylonas, Stelios Dritsas, Bill Tsoumas, and Dimitris Gritzalis. Smartphone security evaluation-the malware attack case. *SECRYPT*, 11: 25–36, 2011.

[28] Phillip Porras, Hassen Sadi, and Vinod Yegneswaran. An analysis of the ikee.b iphone botnet. In AndreasU. Schmidt, Giovanni Russello, Antonio Lioy, NeeliR. Prasad, and Shiguo Lian, editors, *Security and Privacy in Mobile Information and Communication Systems*, volume 47 of *Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering*, pages 141–152. Springer Berlin Heidelberg, 2010.

[29] R. Racic, D. Ma, Hao Chen, and Xin Liu. Exploiting and defending opportunistic scheduling in cellular data networks. *Mobile Computing, IEEE Transactions on*, 9(5): 609–620, 2010.

[30] ComScore reports June 2012 U.S. mobile subscriber market share. http://www.comscore.com/Insights/Press_Releases/2012/8/comScore_Reports_June_2012_U.S._Mobile_Subscriber_Market_Share.

[31] Fabio Ricciato, Angelo Coluccia, and Alessandro DAlconzo. A review of dos attack models for 3g cellular networks from a system-design perspective. *Computer Communications*, 33(5): 551–558, 2010.

[32] S. Sesia, M. Baker, and I. Toufik. *LTE, The UMTS Long Term Evolution: From Theory to Practice*. Wiley, 2009.

[33] Muhammad Zubair Shafiq, Lusheng Ji, Alex X. Liu, Jeffrey Pang, and Jia Wang. A first look at cellular machine-to-machine traffic: Large scale measurement and characterization. In *Proceedings of the 12th ACM SIGMETRICS/PERFORMANCE Joint International Conference on Measurement and Modeling of Computer Systems*, SIGMETRICS '12, pages 65–76, New York, NY, USA, 2012. ACM.

[34] Patrick Traynor, Michael Lin, Machigar Ongtang, Vikhyath Rao, Trent Jaeger, Patrick McDaniel, and Thomas La Porta. On cellular botnets: measuring the impact of malicious devices on a cellular network core. In *Proceedings of the 16th ACM conference on Computer and communications security*, pages 223–234. ACM, 2009.

[35] Patrick Traynor, Patrick McDaniel, Thomas La Porta, et al. On attack causality in internet-connected cellular networks. In *Proceedings of 16th USENIX Security Symposium on USENIX Security Symposium*, pages 1–16. USENIX Association, 2007.

Biographies

Jill Jermyn is a PhD student at Columbia University, where she works under Professor Salvatore Stolfo in the Intrusion Detection Systems Lab. Some of her research interests are wireless and cellular network security, mobile, and cloud computing. Part of her previous experience includes internships at AT&T Security Research Center and IBM Watson Research Center. Starting

Fall 2014 she will be Adjunct Professor of Computer Science at Purchase College. Jill has been granted numerous awards for her work, including several from Google, Facebook, Applied Computer Security Associates (ACSA), Brookhaven National Laboratory, and the National Physical Science Consortium. Prior to her career in computing, Jill pursued a career as a concert violinist. She has performed at venues such as Carnegie Hall, Lincoln Center, Kennedy Center, the Austrian Cultural Forum NY, and the Los Angeles County Museum of Art, to name a few.

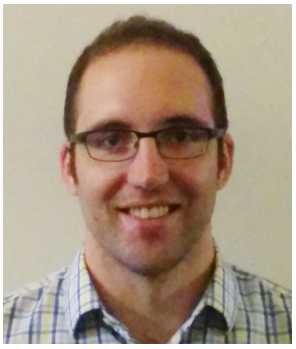

Gabriel Salles-Loustau is a PhD candidate in the 4N6 Cyber Security and Forensics Laboratory in the Electrical and Computer Engineering Department at the University of Miami. His research interests include systems and network security, mobile devices systems security and data privacy.

Saman Zonouz is an Assistant Professor in the Electrical and Computer Engineering Department at the University of Miami (UM) since August 2011, and the Director of the 4N6 Cyber Security and Forensics Laboratory.

He has been awarded the Faculty Fellowship Award by AFOSR in 2013, the Best Student Paper Award at IEEE SmartGridComm 2013, the EARLY CAREER Research award from the University of Miami in 2012 as well as the UM Provost Research award in 2011. The 4N6 research group consists of 1 post-doctoral associate and 8 Ph.D. students, and their research has been funded by grants from NSF, ONR, DOE/ARPA-E, and Fortinet Corporation. Saman's current research focuses on systems and smartphone security and privacy, trustworthy cyber-physical critical infrastructures, binary and malware analysis, as well as adaptive intrusion tolerance architectures. Saman has served as the chair, program committee member, and a reviewer for international conferences and journals. He obtained his Ph.D. in Computer Science, specifically, intrusion tolerance architectures, from the University of Illinois at Urbana-Champaign in 2011.