Location Leaks on the GSM Air Interface

Denis Foo Kune, John Koelndorfer, Nicholas Hopper, Yongdae Kim
University of Minnesota
foo@cs.umn.edu, koeln005@umn.edu, hopper@cs.umn.edu, kyd@cs.umn.edu

Abstract

Cellular phones have become a ubiquitous means of communications with over 5 billion users worldwide in 2010, of which 80% are GSM subscribers. Due to their use of the wireless medium and their mobile nature, those phones listen to broadcast communications that could reveal their physical location to a passive adversary. In this paper, we investigate techniques to test if a user is present within a small area, or absent from a large area by simply listening on the broadcast GSM channels. With a combination of readily available hardware and open source software, we demonstrate practical location test attacks that include circumventing the temporary identifier designed to protect the identity of the end user. Finally we propose solutions that would improve the location privacy of users with low system impact.

1 Introduction

Wireless networks serving mobile subscribers with fixed base stations such as cellular networks have to track those subscribers to ensure adequate service delivery [3] and efficient utilization of limited radio resources. For example, an incoming voice call for a mobile station requires the network to locate that device and allocate the appropriate resources to handle the resulting bi-directional traffic [6]. The network thus has to at least loosely track the device within large regions in order to make the process of finding the device more efficient. This includes handling registration of mobile station in regions as well as hand off between towers within that region [3]. As part of its notification protocol, the network uses a broadcast medium to page mobile stations, notifying them that there is a message waiting for retrieval [6].

There are three main entities with intended differing access to location information of subscribers; the service provider that has access to all the location data of its users, law enforcement agencies that have the ability to subpoena that information, and external entities including other users with no explicit access to the location information. Location leaks in the communication protocol would mean that even entities with no access to the location database would be able to infer some location information from target users. In this study, we demonstrate that access to location information by the third group has a low entry barrier being attainable through open source projects running on commodity hardware.

The motivation for attackers to obtain pieces of location information of victims include anyone who would get an advantage from such data. For example, agents from an oppressive regime may no longer require cooperation from reluctant service providers to determine if dissidents are at a protest location. A second example could be the location test of a prominent figure by a group of insurgents with the intent to cause physical harm for political gain. Yet another example could be thieves testing if a user's cell phone is absent from a specific area and therefore deduce the risk level associated with a physical break-in of the victim's residence.

We focus on the common lower GSM stack layers at layers 2 and 3 and pay attention to the effects of the broadcast channels on the location privacy of users. We show that although GSM was designed to attempt to obfuscate the identity of the end device with temporary IDs, it is possible to map the phone number to its temporary ID. We also show that it is possible to determine if a user's device (and by extension, the actual user) is within an area of 100 km^2 with multiple towers by simply looking at the broadcast messages sent by the network. We also show that it is possible to test for a user on a single tower which could map to a relatively small geographic area of around 1 km^2 or less. In this work, we don't narrow it down to an exact building yet, but we can tell if the user is within a dozen city blocks.

Organization: We start with an overview of the 3GPP cellular network architecture in section 2 and describe the paging procedure that we use in our analysis. In section 3, we review other related works that have been focussed on the IP layer and above in the communication stack. In section 4, we charaterize the network and define primitives that we will use in our attack description in section 5. Then in section 6 we develop methods to deduce the geographic location of cellular base stations in the area and use it to map a region that we use for the evaluation of our attack. Finally, in section 7 we propose low-impact solutions that would prevent the current leak of location information on the lower layers of the GSM protocol stack.

2 Background

The original commercial cellular networks were deployed in the early 1980s based on analog voice, also known as *1G*. To better utilize the wireless radio resources and provide better scalability, protocols for digital voice were developed. While there were multiple standards available, the Global System for Mobile Communications (GSM) [1] was widely adopted as the de facto standard, now referred to as *2G*. With the technology boom in the late 1990s, there was an increased interest in carrying data on wireless cellular networks. General Packet Radio Service (GPRS) was designed to utilize existing GSM networks [2], and is sometimes referred to as *2.5G*. Another variant of GPRS, based on a different modulation technique from GSM was also designed around the same time and would produce the Enhanced Data rates for Global Evolution (EDGE) network with higher throughput than GPRS. In the mid 2000s, there were a number of new standards developed to offer better data rates including Wideband Code Division Multiple Access (W-CDMA), that gained wide adoption with the introduction of smartphones such as the iPhone 3G and HTC Dream (T-Mobile G1). Along with W-CDMA, High-Speed Packet Access (HSPA) protocols fall under a common umbrella named Universal Mobile Telecommunication Systems (UMTS) and are commonly referred to as *3G*. With increasing demand for higher bandwidth, two main standards have been developed as the next *4G* network, namely IEEE 802.16 (WiMax) [7] and Long Term Evolution (LTE) [8] with LTE quickly gaining popularity in 2011. LTE is the successor to UMTS, whereas WiMax was developed independently. In parallel to GSM in the *2G* networks, a service based on CDMA was built with the *3G* equivalent being CDMA-2000. Those networks have not gained wide adoption outside North America.

With the current focus on LTE and smart phones, there has not been much attention paid to the rest of the feature phones that make up 80% of the active subscription base for mobile phones [10]. The total number of estimated mobile phone subscribers worldwide was 5.3 billion in 2010 according to the United Nations' International Telecommunication Union (ITU) [10] and there is still a large subscriber base on GSM networks. In addition, with the recent pricing models based on data consumption, there has been an incentive for users to put their phones on networks with lower data rates to avoid accidental over-consumption. For example, the Apple iPhone allows the users to down-select their preferred network from the W-CDMA network to the EDGE network. In doing so, they also revert to the GSM network for voice communication, making them visible to GSM based attacks such as the one described in this work.

2.1 GSM infrastructure overview

A GSM cellular network is composed of 15 main logical entities [3]. The entities relevant to this work are as follows:

- The Visitor Location Register (VLR) is in charge of one or more areas that mobile stations may roam in and out of. This entity handles the temporary IDs (TMSIs) of the mobile stations.
- The Base Station System (BSS) is a network of base station transceivers (BTS) and controllers (BSC) responsible for communicating directly with the mobile station.
- The Mobile Station (MS) is the mobile device carried by the user. It is composed of the actual device and a Subscriber Identity Module (SIM).

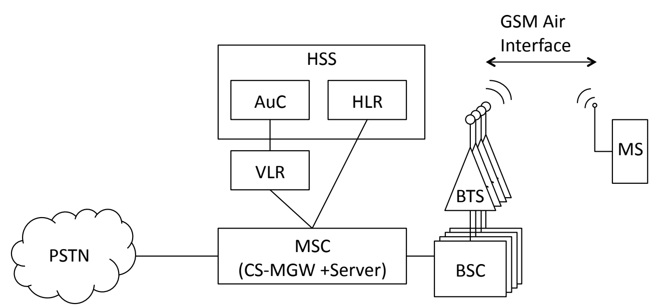

Figure 1. Simplified diagram of a cellular network connected to the PSTN. Only nodes relevant to this work are shown.

Figure 1 presents an overview of the architecture and the connections between the entities. The mobile station and the BTS talk over the wireless GSM protocol, also known as the air interface. Within the GSM air interface specification [4] there are multiple channels defined for the downlink and uplink communication between the MS and the BTS. The channels relevant to this work are as follows:

- *PCCH*: The broadcast downlink channel that all phones listen to. There are multiple frequencies identified by an absolute radio-frequency channel number (ARFCN) that can be used for this channel.
- *RACH*: The random access uplink channel available to any mobile station registered on the network.

- *SDCCH*: A specific uplink channel, assigned by the BTS.

In order to limit the consumption of radio resources, the cellular network will try to limit the traffic through its towers by only transmitting the messages required by the cell phones being served in that particular area. Thus, the paging request broadcast messages are sent through towers within a specific Location Area Code (LAC) serving the mobile station of interest. Those messages are sent over the broadcast PCCH downlink and are used to notify a mobile station that it needs to contact the BTS [6]. Mobile stations tune (or camp) on a particular frequency for their chosen BTS and are able to hear all the pages being issued. Each paging request message contains the unique identifier of the intended destination, either a globally unique International Mobile Station Identity (IMSI [5] clause 2.2) or a Temporary Mobile Subscriber Identity (TMSI [5] clause 2.4).

2.2 Incoming Call protocol

The logical flow for the radio interface in a GSM network during an incoming call works as follows [6]:

1. The BSS attempts to find the mobile station. BTSs within the last LAC known to have seen the device send a paging request with the mobile station's IMSI or TMSI over the PCCH downlink.

2. Upon reception of the paging request, a mobile station will determine if the IMSI or TMSI matches its own. If it does, the mobile station will request radio resources from its BTS with a channel request message sent over the RACH uplink.

3. The BTS will indicate the details of the SDCCH intended for the mobile station in an immediate assignment message sent over the same PCCH downlink.

4. The mobile station responds with a paging reply over the allocated SDCCH uplink. The rest of the protocol allows the mobile station and the BTS to negotiate the security level and other setup parameters before data (text or voice) are transmitted. The initial protocol described is summarized in Figure 2.

Paging request messages: Paging request messages can be of type 1, 2 and 3 that simultaneously carries up to 2, 3 and 4 identities respectively. Paging requests are issued for every call or text message being sent to a mobile station within an LAC. The response to the paging request is the paging response that is sent from the mobile station to the BTS of its choice over the SDCCH uplink. Typically a mobile station will select the BTS with the strongest signal.

Immediate assignment messages: The immediate assignment is a response to the channel request message from the mobile station. It contains an identifier matching that from the channel request message. Since the immediate assignment messages are sent over a broadcast channel, the identifier allows the mobile station to discriminate between a message intended for it and messages intended to other mobile stations. While easy to match up with the requesting device if the uplink channel request message is heard, determining if an immediate assignment message was the result of a trigger from a paging request message by simply observing the PCCH downlink is more tricky. The ability to determine if an observed immediate assignment message is indeed intended for a mobile station of interest would indicate that the device is camped on the ARFCN of the PCCH carrying that message. Therefore, it would be an indication that the device is on the same tower as the one we are listening to.

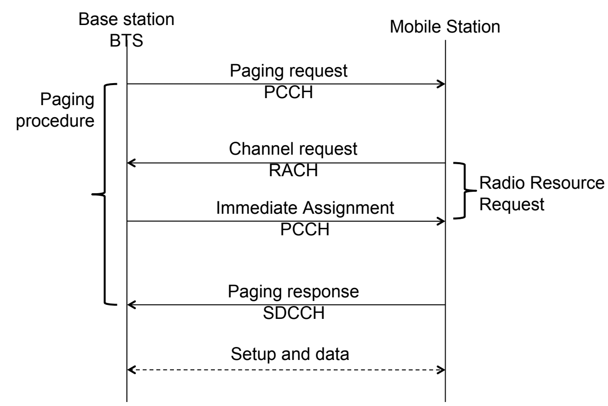

Figure 2. Sequence diagram for the first 4 messages in the paging procedure over the air interface between the BTS and the MS.

3 Related works

Location privacy has been studied previously [17, 22] in the context of Location Based Services, but those works looked at smart platforms such as Android [11] and Apple's iOS [12]. On those smart phones, the location information is mostly acquired, stored and transferred as application layer (OSI layer 7) data. Location inference based on lower layers of the GSM communication stack can be performed by the cell phone towers themselves using trilateration or triangulation [14], but that information requires collaboration with the service provider and it is typically reported, used and carried at the application layer as well. Location leaks from the lower IP layer on cellular networks have been investigated by Krishnanmurthi, Chaskar and Siren [23], but leaks due to the broadcast messaging at the bottom of the GSM protocol stack that generates location specific traffic have not yet been carefully studied.

Chen et al. have explored mapping techniques based on logging signal strengths of GSM cell towers [16]. They surveyed a 18.6 km by 25 km area in the downtown Seattle region and mapped cell towers to regions of maximal signal strength. Their goal was to assess the performance of three positioning algorithms given the GSM signal strength traces. That positioning calculation is carried out based on measurements made by the mobile station itself. What we are investigating is if we can locate a mobile station based on side effects of the communication protocol.

Husted and Myers propose a different system [21] where multiple devices attempt to listen to the IEEE 802.11 unique BSSID of the victim's smart device. With a sufficiently dense population of observers, the physical location of the victim can be tracked closely, without their knowledge. Our approach looks at a similar problem but on the GSM communication stack. Instead of searching for the BSSID, we look at the temporary identifiers of the victim's device. Our technique does not require a WiFi enabled phone; it works for any device that talks over the GSM network.

De Mulder et al. have been looking at methods of identifying users from their trace as they use parts of the GSM protocol [19]. Their method involves analyzing the database of registrations requests by the mobile station as it moves between cells. Our method does not require cooperation from the service provider. We simply listen on broadcast messages to deduce if a user is in a location of interest.

During a talk about sniffing the GSM communication at the 27th Chaos Communication Congress (27c3) [24], Nohl and Munaut introduced hints of how to watermark the paging channel. Their technique used malformed SMS messages to remain silent, which is different from our aborted call method discussed in section 6.

4 Preliminary measurements

The GSM specifications give us an overview of the mandatory behavior of the network but provide little insight on the behavior of an actual deployed network. To this end, we performed preliminary measurements on the T-Mobile and AT&T networks in a major metropolitan area to characterize actual deployments.

4.1 Measurement platform

Our measurement system as displayed in figure 3 is based on the Osmocom baseband platform [13] coupled with a land line phone capable of making outgoing calls. On the bottom right is the Motorola C118 connected via a serial reprogrammer cable to a Serial-to-USB converter which is in turn connected to a laptop. Using tools from

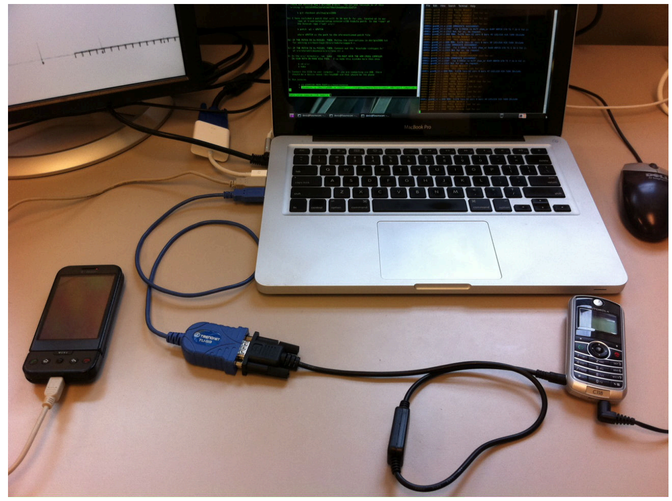

Figure 3. Our testbed showing the C118 running the GSM layer 1 (bottom right), the laptop running the GSM layer 2 and 3, and a T-Mobile G1 (bottom left) used as our mapping tool and later our victim.

the Osmocom project, we replaced the C118 firmware with the Osmocom GSM Layer 1 firmware. The new firmware tunes to the requested ARFCN from the laptop and relays all the layer 1 messages over its serial port back to the laptop. Running on the laptop is the `osmocon` tool that flashes the replacement firmware to the C118 and then turns into a relay that forwards packets from the serial port to a socket to be used by other applications from the Osmocom project. Running on the same laptop is the `mobile` application that connects to the previously opened socket and implements the GSM Layers 2 and 3. We only had to make minor modifications to the `mobile` application to interface with our set intersection tool.

The T-Mobile G1 (US) on the left in figure 3 was used as our mapping tool. It is running the Android 1.5 (downgraded TC4-RC29) operating system with a custom patched kernel based on 2.6.25 and the Cyanogen mod firmware. The custom kernel patches we added were to log messages between the application and the baseband. The phone's baseband version is the stock 62.33.20.08H. The baseband firmware responds to Hayes AT commands, with an expansion for GSM messages following the 3GPP TS 27.007 specifications. Of particular interest was the `AT+GSM` command that queries the baseband for the current information about the current TMSI, LAC, ARFCN as well as a list of neighboring ARFCNs. Our patch logged those messages when queried by the field test application at 1Hz, allowing us to match those with GPS coordinates from a separate device during our mapping studies. That same phone was later

Table 1. General observations on the GSM PCCH

	T-Mobile LAC `747b`	AT&T LAC `7d11`
Paging Requests - IMSI	27,120	8,897
Paging Requests - TMSI	257,159	84,526
Paging Requests Type 1	284,279	91,539
Paging Requests Type 2	1635	26
Paging Requests Type 3	0	1
Immediate Assignments	207,991	10,962
Observation period	24 hours	24 hours

used as our victim, as it allowed us to check our answers with the actual TMSI of the device.

4.2 General observations

To understand the general trend of the PCCH traffic, we performed a series of captures on the T-Mobile network and the AT&T network. The results for a 24 hour period are summarized in table 1. Next, we measured the time delay between a call initiation on the PSTN and a paging channel request being issued on the PCCH. For our sample of 40 calls, we observed a mean delay of 2.96 seconds, with a standard deviation of 0.152. The observations are summarized in figure 4. During that same experiment, we measured the mean time delay between the PSTN call initiation and the actual ring on the mobile station to be 8.8 seconds with a standard deviation of 4.5. The median was 7.0 seconds. We also observed that calls aborted before 5.0 seconds following the PSTN call initiation would result in no rings or missed calls on the device, but by that time, the paging request would have already been sent. These initial measurements will help us quantify parameters used later in the section 5.

4.3 Observed messages on the PCCH

Paging Request of Type 1 that allow a single or two mobile identities to be paged per message [6] (clause 9.1.22) were the most commonly observed. See table 1. In our captures, we observed that over 90% of the Paging Requests had no identities; these were dismissed from the following plots since they would not trigger devices. A summary of the paging requests over 48 hours captured on the AT&T network is shown in figure 5. We observe general trends in local human activity with high traffic rates of 150 pages per minute in the middle of the day and low traffic rates of about 10 pages per minute between 00:00 (midnight) and 06:00 (6am) local time.

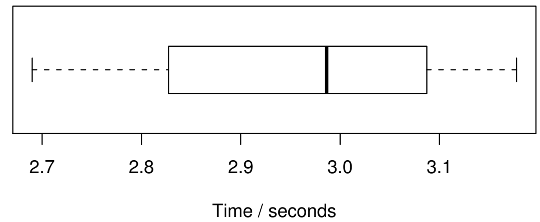

Figure 4. Delay between the PSTN call initiation and the corresponding paging request message broadcast on the PCCH.

From a different 21 hour capture on T-Mobile, we looked at the patterns of the time difference between two pages intended for the same TMSI. We observed that there appears to be a sharp decrease until around a time difference of 200 seconds, where the distribution levels off before dropping after a time difference of 600 seconds, beyond which two pages for the same TMSI becomes unlikely. See figure 6.

During the 51 hour AT&T experiment, we noticed that the traffic rate for the paging request messages (2.32/second) far exceeded the traffic rate for immediate assignment messages (0.554/second). Upon review of the GSM specification [6] it became apparent that the immediate assignment messages are limited to the ARFCN (and therefore to the cell phone tower) that a mobile station requested the radio resources from. Repeating those messages to all the BTSs in the region would be a poor utilization of the downlink bandwidth since those resources are local to a single BTS.

5 Attack description

5.1 Threat model

The attacker we consider in this work is anyone who stands to gain from obtaining the location information of a victim as outlined in the introduction. To be successful, the attacker first requires the ability to actively introduce a PCCH paging request for the victim, which can be achieved with a text message or a call initiation using any phone line. A regular PSTN land line is preferred for better timing accuracy and better control over the time to abort the call protocol to avoid notifications at the application layer on the victim's phone. The attacker also requires the ability to listen passively on the broadcast GSM PCCH paging channel. We describe the attacker's capability as follows:

- **PSTN active:** The attacker causes paging request messages to appear on the GSM PCCH by dialing the vic-

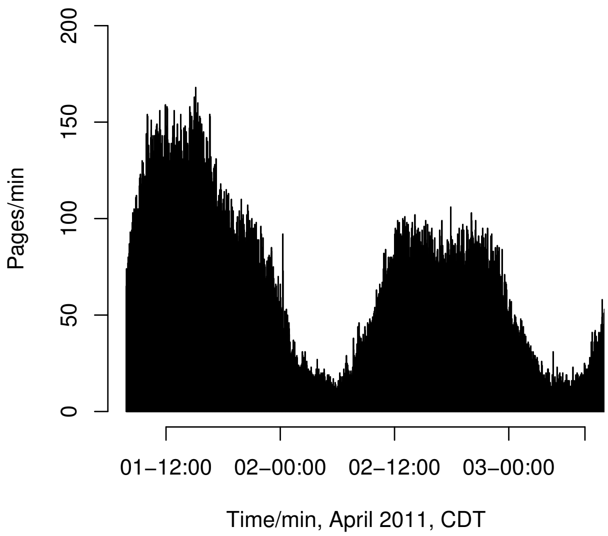

Figure 5. Observed Paging Requests per minute over a 48 hour period.

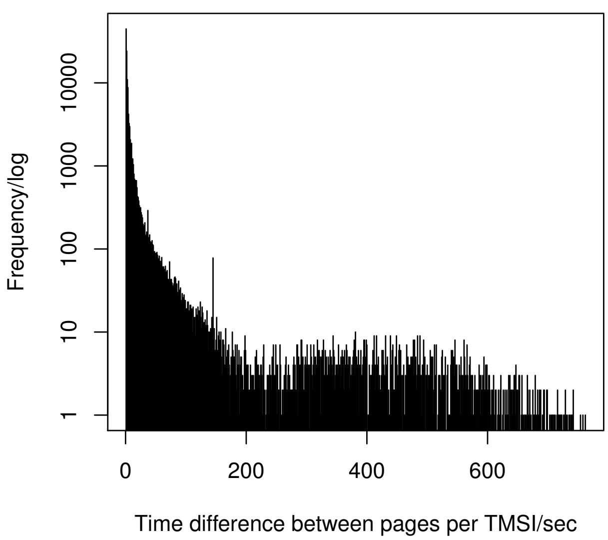

Figure 6. The time difference between paging requests for the same TMSI, observed over a 21 hour period.

tim's phone number or sending the victim a text message.

- **GSM passive:** The attacker is a passive listener on PCCH broadcast plaintext. There is no need to crack the encryption algorithm since we only require the beginning sequences of 4 messages in the radio resource (RR) setup phase of the GSM protocol.

It is possible to combine both capabilities on a single system. The PSTN active component can be composed of a dial up modem to which `ATDT` commands are issued to trigger paging requests on the PCCH. The GSM passive component can be composed of the system we described above. The set intersection tool can be started automatically after the PSTN active step. We also note that our attack would work on any system with an unencrypted broadcast paging channel with long-lived identifiers. Thus, the UMTS and LTE paging procedures could also be vulnerable if deployed in the same manner as observed GSM networks, provided basebands reporting all the messages from the paging channel are available.

5.2 Temporary IDs in local areas

A cellular network provider has to track the location of mobile users, at least to a coarse grain level in order to make efficient use of limited radio resources. Separating a large area into n smaller geographic regions such as Location Area Code (LAC) and making broadcast messages (including paging requests) local to those smaller regions on average would reduce the paging traffic down to $\frac{1}{n}$ times the original number of messages, on average. By observing the paging requests, we can tell if a victim is within that area if we know their unique ID. However, that ID (an IMSI or TMSI) is the only identifier visible on the GSM PCCH and the internal system mapping of telephone number to IMSI or TMSI is not known a priori.

5.3 Revealing identities

The TMSI has a meaning only within the LAC in which the device is located. In order to carry out our location test, we need to reveal the mapping between the PSTN phone number and the TMSI visible on the GSM PCCH. We focus on the TMSI since our observations in section 4.2 showed that over 90% of the paging requests contain TMSIs. The technique described below will work just as well with IMSIs, but those are omitted for clarity.

We first define the possible sets of candidate TMSIs after a call initiation by limiting the identifiers within a time window defined by $t_{min} \leq t \leq t_{max}$, where the time delays defined are taken relative to the completion of the dialed phone number on a PSTN land line. From our initial measurements in section 4, we empirically determined t_{min} to be 2.50 seconds and t_{max} to be 3.42 seconds to cover cases up to 3 standard deviations from the mean, assuming a Gaussian distribution for the delay between the PSTN call completion and the Paging Request Message on the PCCH. Specifically, we define the set of possible TMSIs in a given time window to be as follows.

$$I_j = \begin{cases} \mathrm{TMSI}_t & t_{\min} \leq t \leq t_{\max} \\ \emptyset & \text{otherwise,} \end{cases}$$

where $1 \leq j < n$, $t_{min} = \mu - 3\sigma$, $t_{max} = \mu + 3\sigma$ and μ, σ are the mean and standard deviation of the PSTN to paging request delays measured during the calibration phase 4.2. We repeat this process n times, collecting n possible sets of unique TMSIs $I_1, ..., I_n$. We wait at least $t_d > t_{max}$ seconds between each call to give the system a chance to reset. We then compute the intersection of all the sets to extract a very small number of possible TMSIs, $I_1 \cap I_2 \cap ... \cap I_n$. We note that depending on the paging channel traffic patterns and TMSI reassignment policies, it may be possible to get 0 or more TMSIs. We have observed that the TMSI assignments tend to last for several hours, making it probable that we will obtain at least one TMSI for a short t_d which is on the order of 200 seconds as determined by prior observations (section 4.2). From our experiments, we determined that we only need a small number for n, typically 2 or 3 to narrow down the exact TMSI. Finally, we note that for this test to be successful, the attacker has to be able to hear the paging request for the victim device. Empirically, the T-Mobile LAC `747d` covers an area in excess of 100km^2, and repeats all paging requests at each BTS within that area. It thus makes the attack to reveal the identity of the mobile station very practical.

5.4 Absence testing

A natural extension of the previous finding is an absence test to determine if a mobile station is not in that region. The resulting intersection would yield a null set, even for a large number, n, of PSTN calls. This information can be useful for an attacker if the absence of the mobile device is indicative of the absence of the user as well. Starting with the method to reveal the TMSI as outlined above, and with the assumption that the TMSI does not change for the duration of the attack measurements, we run the attack n times with a delay of t_d interval between each reading and recover the TMSI sets $I_1, ..., I_n$. In this case, if $I_1 \cap I_2 \cap ... \cap I_n = \emptyset$, we can reasonably conclude that the mobile device is not registered in the current LAC, and if powered up, the device is outside of the region. A quick test can be done by making a call and letting it complete. If the device is turned off, the call will tend to go to voicemail faster than if the device is made to ring.

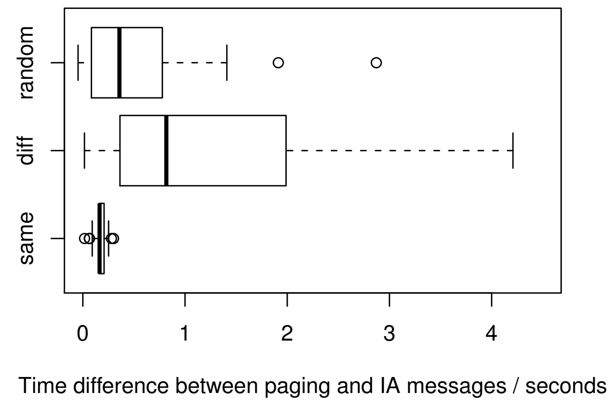

Figure 7. Time difference between the paging request and the very next immediate assignment message under our 3 test conditions.

5.5 Presence testing on the same BTS

Once an attacker determines that a target device is in the same LAC as he is, the next step is to determine if the device is listening on the same BTS. Recall that the PCCH downlink carries the paging request and immediate assignment messages, but the identifier in the immediate assignment is chosen by the mobile station and communicated in the channel request which is unknown to us. We thus developed a technique to determine if we are on the same tower as the victim by looking at the time difference between an immediate assignment and its triggering paging request.

We listened for the paging request for the victim's phone using the TMSI, and measured the time delay before we observe the very next immediate assignment message on the PCCH. We examined 3 conditions with our testbed being 10m from the victim:

1. Camped on the same ARFCN as the victim, and triggered paging requests using PSTN calls
2. Camped on a different ARFCN as the victim and triggered paging request using PSTN calls
3. Camped on an arbitrary ARFCN, and sampled the PCCH starting at random times.

We denote t_p as the time stamp of the paging request for our target device and t_a for the time stamp of the very next immediate assignment message. We want to compare the time difference $\delta t = t_a - t_p$ in the 3 test cases above.

For each test, we had a sample size of at least 40 readings in order to obtain a power of over 80% for the Welch two-sample t-test used in our analysis. In the first case, with a sample size of 46, we observed the mean time difference $\mu_{\delta t} = 0.177$ seconds with a standard deviation of $\sigma_{\delta t} = 0.0572$. In the second case with a sample size of 43 where we were listening to another ARFCN, we observed $\mu_{\delta t} = 1.99$ seconds with a standard deviation $\sigma_{\delta t} = 3.42$. Finally, in the third case with a sample size of 40 where we randomly sampled the PCCH on a random ARFCN, we observed $\mu_{\delta t} = 0.517$ seconds with a standard deviation $\sigma_{\delta t} = 0.582$. Our findings are summarized in figure 7.

Our goal is to be able to discriminate between the situation where we are on the same ARFCN and therefore on the same tower, and the situation where we are in the same LAC, but on a different tower. The hypothesis is that if we are on the same tower as the victim, we will hear the immediate assignment for that device, which will be issued very close to the paging request in order to provide fast service. Therefore, we set our null hypotheses as follows:

$H_0^1 : \mu_s = \mu_d$, and $H_0^2 : \mu_s = \mu_r$, where

- μ_s is the mean time difference for test condition 1 (same tower)
- μ_d is the mean time difference for test condition 2 (different tower)
- μ_r is the mean time difference for test condition 3 (random).

To compare the time differences obtained from our tests, we use a Welch two sample t-test that is robust with small samples and accounts for populations with different distributions. The results are as follows. Comparing μ_s and μ_d, we obtained a p-value of 0.001199. Comparing μ_s and μ_r, we obtained a p-value of 0.0006942. Both results indicate that finding a test statistic equally or more extreme is very unlikely. We can thus reject both null hypotheses with such low p-values. Therefore, our conclusion is that there is likely to be a meaningful difference that allows us to discriminate between situations where we are listening on the same ARFCN, or on a different one by looking at the paging requests and immediate assignment messages' arrival times.

5.6 Moving objects

Our attack assumes that the victim is relatively stationary for the duration of the test. To reveal the TMSI, we use the entire LAC through which the victim would be expected to be present for a relatively long time period given the large geographic coverage of several square kilometers that we observed. Indeed, we verified that our identity revealing method works on a device moving at an average of 105 km/h through an LAC. The device stayed in the same LAC for 8 minutes, and we only required 2 minutes to reveal the TMSI. Clearly, our "same tower" test would not work for a fast moving object. We observed that mobile stations tend to camp on the same ARFCN until they move further than 1km away. Our test requires about 5 seconds per ARFCN to complete. Depending on the region, there could be 3 to 5 ARFCNs with high enough RSSI values to test. Thus, if the victim lingers within 1km of the tower for a couple of minutes, our attack would succeed. If victim's physical path is known, it would also be possible to predict where the victim would be after a limited number of time steps. Such extensions are left as future work.

6 Carrying out the attack in practice

To carry out a meaningful attack, we need to understand the geographic coverage of the LACs and the distribution of the cell towers. In this evaluation, we focused on the T-Mobile GSM network in a large metropolitan area. No public registry for small cell tower structures are available. In fact, the Code of Federal Regulations (CFR) 47 Part 17.7 (Revision 10/01/1996) specifies that small structures need not be registered. Thus we used a surveying method similar to Chen et al. [16] and applied a wall-following and hill-climbing method to reduce the amount of samples required. We mapped the LAC area by taking sample readings from geographic locations surrounding our LAC of interest, which turned out to cover an area of about 100 km^2. Within that LAC, we mapped a much smaller area to determine the cell tower locations and the coverage area of each ARFCN for those towers. The techniques we used could be applied to other GSM providers as well.

6.1 Mapping an LAC

We used a combination of our patched T-Mobile G1 phone to log the GSM messages from the baseband and a GPS to track our location as we moved through the LAC `747d`. We then followed an approximate wall following exploration, roads permitting, treating the edge of the LAC `747d` as a wall. In doing so, we also mapped the edge of the neighboring LACs. The result of our coarse survey is shown in figure 8. The area corresponding to LAC `747d` is in grey. The red line on the North side of the region is approximately 10km long. The entire region is a little over 100km^2 in area. In the same figure, the small blue area on the East side is a small area we chose to zoom in to map the individual towers as presented in figure 10. The left plot in figure 8 shows the same area with other LACs plotted. The overlapping regions are approximations of the areas where our device could attach to towers in both the target LAC and the neighbor LAC.

During our survey, we found that there was a significant overlap between areas covered by different LACs. By

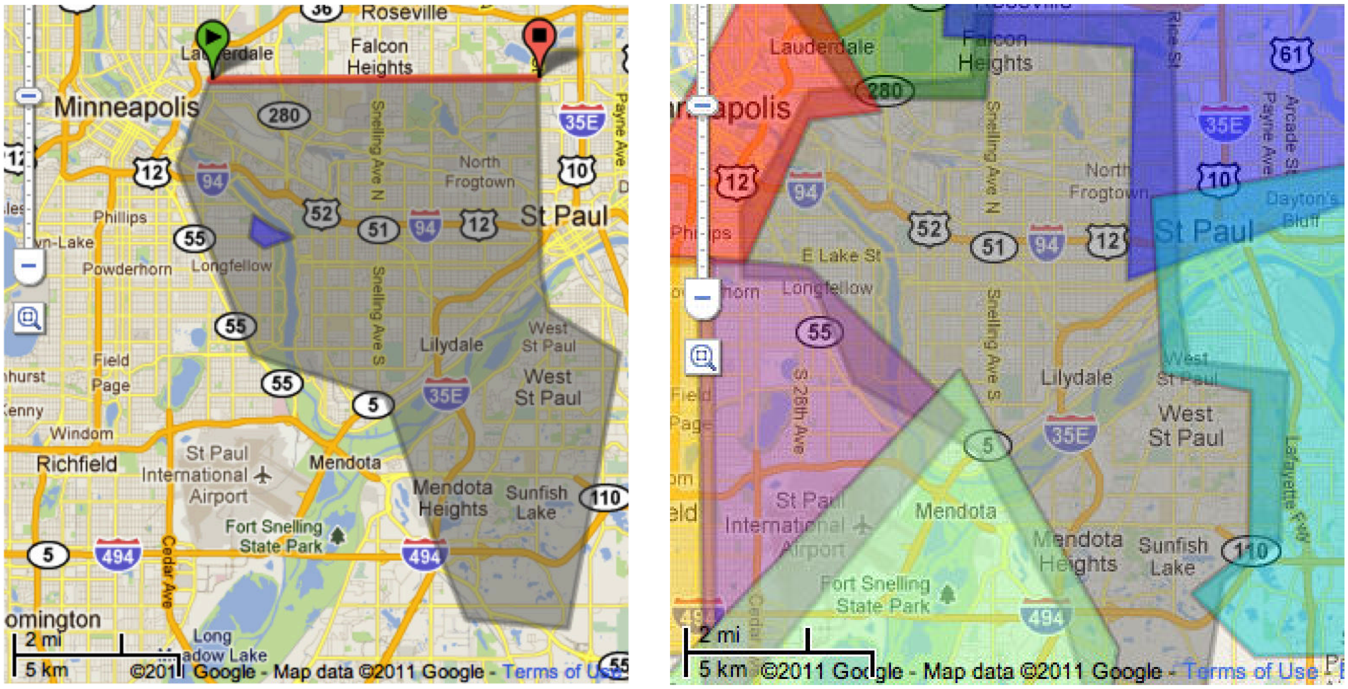

Figure 8. The LAC 747d (grey) and surrounding LACs (right) on T-Mobile.

choosing an area carefully, an attacker can maximize the areas monitored without having to move. In our experiment, we estimated that we would need about 6 systems to cover our metropolitan area with a population of around 2 million residents [9].

6.2 Cell tower location

Since the highest granularity of tests that we can perform is the determination that we are on the same tower, it is important to know the location of that tower precisely. We used the hill climbing method with the objective of maximizing the RSSI in the RF field of the target tower. We used our modified G1 to make point measurements of the field strength and we moved in the field following a variant of the classic hill climbing algorithm where we overshot the maximum point by 50m or more, then backtracked before taking a perpendicular direction to ensure that we were not stuck in a local minimum due to non-uniform RF signal attenuation. Figure 9 shows an example of our hill climbing and the result with the detected cell phone tower.

6.3 Mapping towers and ARFCNs

Next, we wanted to determine the distance between a tower and a mobile station before it would switch to a different tower. That will give us a region where the victim is likely to be. We used the same system from our LAC mapping study to conduct our individual cell tower mapping study. This time, we moved along every city block in the target region within the LAC of interest. We recorded the 5 strongest RSSIs and their corresponding ARFCNs. The results of the survey for 3 ARFCNs with the strongest RSSI values along with inferred cell towers using the previously described hill climbing method is shown in figure 10.

A mobile station will attempt to camp on the ARFCN with the highest RSSI. We therefore determined that the victim is likely to be within the 1km^2 area around the tower it is attached to, depending on the signal attenuation within the area we surveyed. We also note that there are several regions where we could camp on multiple ARFCNs. Thus, an attacker can survey multiple towers at once, while the victim has to be close to the tower. In this particular experiment, one testbed was able to monitor all 3 towers, for a total area of around 2.5×2.5km from a single location.

6.4 Silent watermarking to expose the TMSI

In order to avoid detection, we need a method that will cause a paging request with the target TMSI, without causing any user-observable side effects on the mobile station. Previous methods have used malformed SMS messages to achieve this goal [24], but those techniques are not universal across carriers and devices. Our method is to initiate a regular phone call from the PSTN, and abort before the first ring occurs on the mobile station. We know from section 4 that

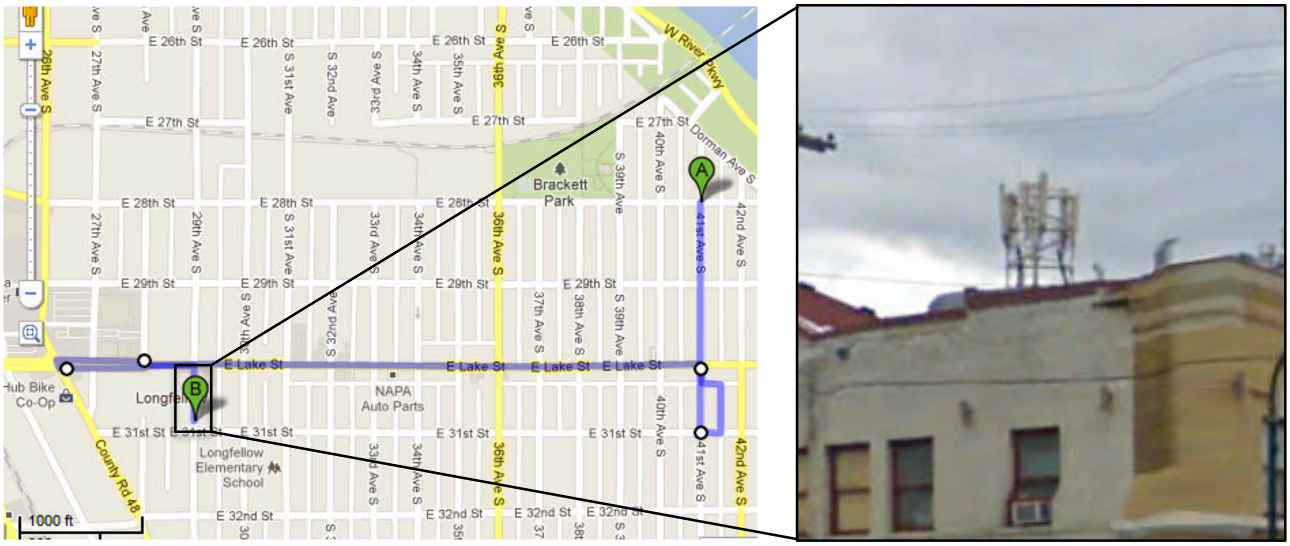

Figure 9. Hill climbing while moving in the RF field of a target cell tower.

the time delay between a call initiation on the PSTN and the appearance of the paging request on the PCCH is about 2.94 seconds on average and the actual ring on the device takes about 7.0 seconds in the median case. We also know that hanging up a call within 5 seconds avoids a ring on the device. We therefore conclude that we can safely hang up at the 4 second mark to silently cause the paging requests and immediate assignment messages to be issued without alerting the user.

6.5 Evaluation

We now know how to reveal the TMSI of a target device, how to determine if we are listening on the same ARFCN and how to find the cell tower responsible for that ARFCN. We applied those methods to test for the absence and proximity of a set of 8 phones on the T-Mobile and AT&T GSM networks. The models of phones tested include the T-Mobile G1, the iPhone 3G (iOS 3.1.3), iPhone 4 (iOS 4.3.3) and iPhone 4S (iOS 5) all set with the 3G network off, and an unmodified Motorola C118. We did the call initiations both from a land line and another cell phone. We used $t_{min} = 2$ and $t_{max} = 5$. We chose a rural location at a relatively quiet time of the day (8pm local time) and a metropolitan area during a busy time of the day (noon local time).

Without a priori knowledge of the ARFCN where the phone is camping, we obtained a list of audible frequencies in the vicinity and we carried out the high resolution test for the 3 frequencies with the best RSSI in the area. We were able to determine the appropriate ARFCN with a sample of 8 silent pages, concluding a proximity to the tower using the detected ARFCN. We proceeded to use the hill climbing technique to find the tower and verified that the phones were within the 1km^2 area surrounding their respective towers. At the conclusion of the test, we verified that we had the proper TMSI and ARFCN. We note that the field test application on the T-Mobile G1 was the only one to report the TMSI allowing a direct verification of our method. For the other phones, we used multiple calls at random intervals to confirm the TMSI used. The time difference between the PSTN call initiation and the paging request were consistent with our preliminary observations above.

7 Mitigations

7.1 Paging multiple LACs

The service provider will try to limit the consumption of limited radio resources. Therefore, when looking for a device, it will only page the LAC where the device is most likely to be located. However, this enables the absence test as we described. It is clear that paging all the LACs for all devices is very wasteful. We therefore propose to only page a set of LACs where the device is present most of the time. From a study by González, Hidalgo and Barabási [20] on human mobility based on tracking their mobile devices, and later confirmed by other works including De Mulder [19], it is known that human mobility patterns are very regular. By repeating the pages at locations where the device is frequently located, we would obfuscate the real location of the device. An implementation of this scheme could automat-

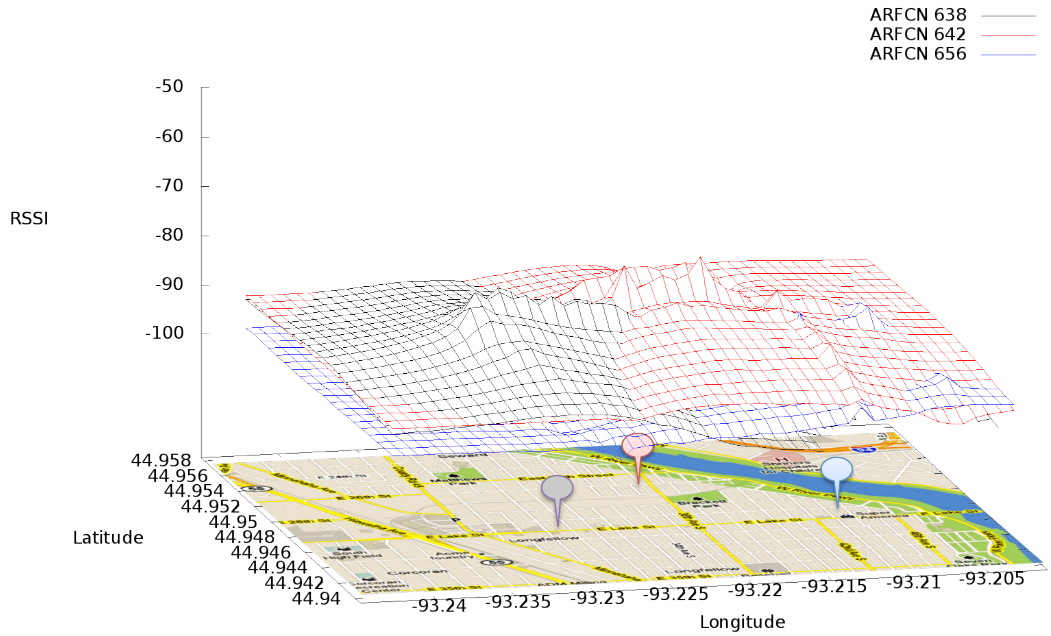

Figure 10. Signal strengths for 3 ARFCNs in a 1.5 $\times$ 2.5km area. The peaks of the meshes correspond directly to the geographic position of the observed BTS.

ically learn the common locations of a user and page the appropriate LACs. From the readings in table 1, we observe that most pages are of type 1, (that can carry at most 2 identities) and only one of type 3 (that would allow 4 simultaneous identities). We interpret this as an indication that the paging channel is operating well below its maximum capacity and suggest that doubling the paging channel traffic should be possible without significantly impacting the service delivery.

7.2 Frequent TMSI change and TMSI allocation

Part of the goal of the TMSI is to hide the correspondence between messages from the PSTN and the PCCH. In this respect, it is parallel to those introduced by Chaum [15]. The first goal is to make the input and output bitwise unlinkable. The second goal is to defend against traffic analysis for which we apply a method introduced by Danezis [18].

The GSM specifications in clause 2.4 of the *Numbering, addressing and identification* specification [5] does not mandate a specific structure for the TMSI other than preventing the use of `0xFFFFFFFF`. Thus, the operators are free to choose the value of the TMSI since it has relevance only to the VLR. From a 51 hour observation on the AT&T network, we observe that for 211,094 paging requests using TMSIs there were only 59,329 unique TMSIs, so a TMSI was reused on average 3.56 times. It is clear that with $2^{32}-1$ possibilities for the 4-octet TMSI, there is still plenty of room for TMSI allocation. For that AT&T network, most the TMSIs had a value above `0x90000000` with the bulk being between `0xBE000000` and `0xFB000000`.

Our attack against TMSIs works because of the long lasting TMSI allocation compared to the short delays between calls. By making the time that a TMSI is allocated shorter than the time delay between two calls, the TMSI now becomes unrecognizable and we achieve our first goal. We note that the TMSI doesn't need to change based on time (although recommended in the GSM specification), but can be reallocated after successfully receiving a page. Thus, the TMSI allocation time t_{TMSI} will be at least as small as the calling delay ($t_{\mathrm{TMSI}} \leq t_{\mathrm{call}}$), since the system can also choose to change the TMSI between calls.

7.3 Continuous time mixes applied to the PCCH

Even with bitwise unlinkability between phone numbers and TMSIs, it is still possible to apply traffic analysis techniques to determine probabilistically if we are seeing pages caused by our calls, especially during quiet periods. With a low traffic rate, the candidate set of TMSI satisfying the inequality $t_{\min} \leq t_{\mathrm{TMSI}} \leq t_{\max}$ is very small.

To prevent timing-based traffic analysis as discussed, we propose to delay the outgoing paging request messages at the BSS. From a previous work on continuous time

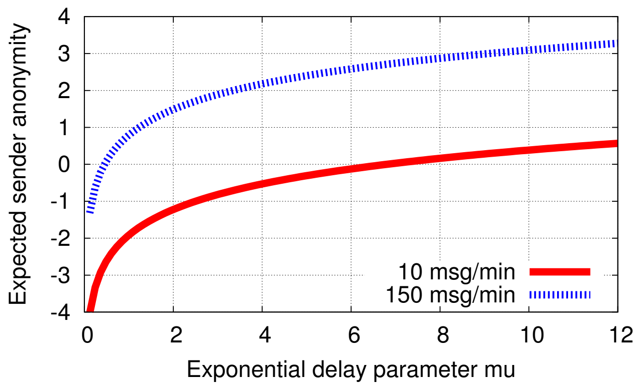

Figure 11. Expected sender anonymity for varying delay parameters with the lowest observed traffic rate of 10 messages/minute and the highest observed traffic rate of 150 messages/minute.

mixes [18], we chose an exponential delay. Specifically, we know that $\mathcal{A} = -log\frac{\lambda_\alpha e}{\mu}$, where applied to our case $\mathcal{A}$ is the expected sender anonymity, λ_α is the PSTN call arrival rate following a Poisson distribution, and μ is the delay parameter following an exponential distribution. We show in figure 11 the predicted expected anonymity for varying exponential delay parameters μ under the observed traffic conditions of $\lambda_\alpha = \frac{1}{6}$ and $\lambda_\alpha = 2.5$. Note that we are bounded by the limits $\mu < \lambda_\alpha e$ for which this estimate is considered accurate. The intuition behind this inequality is that the departure rate should be lower than the arrival rate. We also note that we obtain better anonymity with a higher message arrival rate. Given that the TMSI allocation space is underutilized, with low traffic conditions it would be possible to increase the apparent arrival rate by introducing cover traffic composed of decoy pages containing unassigned TMSIs. By keeping the total traffic over 150 pages/min, it would be possible to keep the expected anonymity high ($\mathcal{A} > 1$), even for a small delay parameter $\mu \geq 2$.

8 Conclusion

We have shown that there is enough information leaking from the lower layers of the GSM communication stack to enable an attacker to perform location tests on a victim's device. We have shown that those tests can be performed silently without a user being aware by aborting PSTN calls before they complete. We demonstrated our attacks using cheap hardware and open source projects and showed mapping techniques to supplement cell tower databases to a granularity acceptable for our attacks. We finally proposed some solutions by applying low cost techniques with good anonymity properties to the GSM stack that could be implemented without hardware retrofit.

Due to the possibility of physical harm that could result from the location leaks in the GSM broadcast messages, we are in the process of drafting responsible disclosures for cellular service providers and the technical standards body of the 3^{rd} Generation Partnership Project (3GPP).

Acknowledgments

This work was supported in part by the National Science Foundation award CPS-1035715 and a grant from the Korean Advanced Institute of Science and Technology. We would like to thank N. Asokan and Valtteri Niemi of Nokia for their insight and support. We would also like to thank Lisa Lendway for guidance in the statistical methods used and Alison Sample for logistical assistance during the geographic mapping of the Local Area Codes.

References

[1] 3GPP TS 01.02 V6.0.1 – General Description of a GSM Public Land Mobile Network (PLMN). http://www.3gpp.org/ftp/Specs/html-info/0102.htm, November 1998.

[2] 3GPP TS 031.60 V7.9.0 – General Packet Service, Service Description. http://www.3gpp.org/ftp/Specs/html-info/0102.htm, November 1998.

[3] 3GPP TS 03.02 V7.1.0 – Network architecture. http://www.3gpp.org/ftp/Specs/html-info/0302.htm, January 2000.

[4] 3GPP TS 04.01 V8.0.0 – Mobile Station - Base Station System (MS - BSS) interface; General aspects and principles. http://www.3gpp.org/ftp/Specs/html-info/0401.htm, March 2000.

[5] 3GPP TS 03.03 v7.8.0 – Numbering, addressing and identification (release 1998). http://www.3gpp.org/ftp/Specs/html-info/0303.htm, January 2003.

[6] 3GPP TS 04.08 v7.21.0 – Mobile radio interface layer 3 specification. http://www.3gpp.org/ftp/Specs/html-info/0408.htm, January 2004.

[7] Part 16: Air interface for broadband wireless access systems. http://standards.ieee.org/getieee802/download/802.16-2009.pdf, May 2009.

[8] 3GPP TS 36.201 V10.0.0 – LTE physical layer; General description (Release 10). http://www.3gpp.org/ftp/Specs/html-info/36201.htm, December 2010.

[9] Census 2010. `http://2010.census.gov/2010census/data/`, 2010.

[10] United nations international telecommunication union sees 5 billion mobile subscriptions globally in 2010. `http://www.itu.int/net/pressoffice/press{\textunderscore}releases/2010/06.aspx`, 2010.

[11] The android mobile operating system. http://www.android.com/, 2011.

[12] The apple ios mobile operating system. http://www.apple.com/ios/, 2011.

[13] The osmocombb project – open source gsm baseband software implementation. http://bb.osmocom.org/, 2011.

[14] J. Caffery and G. Stuber. Overview of radiolocation in cdma cellular systems. *Communications Magazine, IEEE*, 36(4):38–45, 1998.

[15] D. Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, 24(2):84–90, 1981.

[16] M. Chen, T. Sohn, D. Chmelev, D. Haehnel, J. Hightower, J. Hughes, A. LaMarca, F. Potter, I. Smith, and A. Varshavsky. Practical metropolitan-scale positioning for gsm phones. *UbiComp 2006: Ubiquitous Computing*, pages 225–242, 2006.

[17] R. Cheng, Y. Zhang, E. Bertino, and S. Prabhakar. Preserving user location privacy in mobile data management infrastructures. In *Privacy Enhancing Technologies*, pages 393–412. Springer, 2006.

[18] G. Danezis. The traffic analysis of continuous-time mixes. In *Privacy Enhancing Technologies*, pages 742–746. Springer, 2005.

[19] Y. De Mulder, G. Danezis, L. Batina, and B. Preneel. Identification via location-profiling in gsm networks. In *Proceedings of the 7th ACM workshop on Privacy in the electronic society*, pages 23–32. ACM, 2008.

[20] M. Gonzalez, C. Hidalgo, and A. Barabási. Understanding individual human mobility patterns. *Nature*, 453(7196):779–782, 2008.

[21] N. Husted and S. Myers. Mobile location tracking in metro areas: malnets and others. In *Proceedings of the 17th ACM conference on Computer and communications security*, pages 85–96. ACM, 2010.

[22] P. Kalnis, G. Ghinita, K. Mouratidis, and D. Papadias. Preventing location-based identity inference in anonymous spatial queries. *IEEE Transactions on Knowledge and Data Engineering*, pages 1719–1733, 2007.

[23] G. Krishnamurthi, H. Chaskar, and R. Siren. Providing end-to-end location privacy in ip-based mobile communication. In *Wireless Communications and Networking Conference, 2004. WCNC. 2004 IEEE*, volume 2, pages 1264–1269. IEEE.

[24] K. Nohl. Wideband gsm sniffing. http://events.ccc.de/congress/2010/, 2010.