

# Characterizing the Security Implications of Third-Party Emergency Alert Systems Over Cellular Text Messaging Services

Patrick Traynor  
Georgia Institute of Technology  
traynor@cc.gatech.edu

**Abstract.** Cellular text messaging services are increasingly being relied upon to disseminate critical information during emergencies. Accordingly, a wide range of organizations including colleges, universities and large metropolises now partner with third-party providers that promise to improve physical security by rapidly delivering such messages. Unfortunately, these products do not work as advertised due to limitations of cellular infrastructure and therefore provide a false sense of security to their users. In this paper, we perform the first extensive investigation and characterization of the limitations of an Emergency Alert System (EAS) using text messages as a security incident response and recovery mechanism. Through the use of modeling and simulation based on configuration information from major US carriers, we show emergency alert systems built on text messaging not only can not meet the 10 minute delivery requirement mandated by the WARN Act, but also potentially cause other legitimate voice and SMS traffic to be blocked at rates upwards of 80%. We then show that our results are representative of reality by comparing them to a number of documented but not previously understood failures. Finally, we discuss the causes of the mismatch of expectations and operational ability and suggest a number of techniques to improve the reliability of these systems. We demonstrate that this piece of deployed security infrastructure simply does not achieve its stated requirements.

## 1 Introduction

Text messaging allows individuals to transmit short, alphanumeric communications for a wide variety of applications. Whether to coordinate meetings, catch up on gossip, offer reminders of an event or even vote for a contestant on a television game show, this discreet form of communication is now the dominant service offered by cellular networks. In the United States alone, over five billion text messages are delivered monthly [25]. While many applications of this service can be considered non-critical, the use of text messaging during emergency events has proven to be far more utilitarian.

With millions of people attempting to contact friends and family on September 11th 2001, telecommunications providers witnessed tremendous spikes in cellular voice service usage. Verizon Wireless, for example, reported voice traffic rate increases of up to 100% above typical levels; Cingular Wireless recorded an increase of up to 1000% on calls destined for the Washington D.C. area [28]. While these networks are engineered to handle elevated amounts of traffic, the sheer number of calls was far greater than

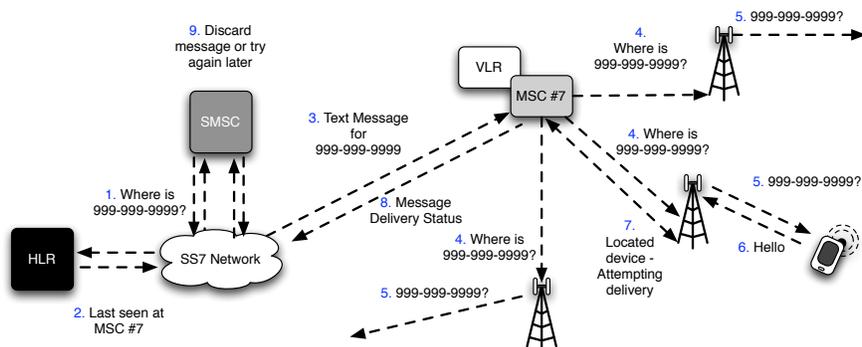
capacity for voice communications in the affected areas. However, with voice-based phone services being almost entirely unavailable, SMS messages were still successfully received in even the most congested regions because the control channels responsible for their delivery remained available. Similar are the stories from the Gulf Coast during Hurricanes Katrina and Rita. With a large number of cellular towers damaged or disabled by the storms, text messaging allowed the lines of communication to remain open for many individuals in need, in spite of their inability to complete voice calls in areas where the equipment was not damaged and power was available.

Accordingly, SMS messaging is now viewed by many as a reliable method of communication when all other means appear unavailable. In response to this perception, a number of companies offer SMS-based emergency messaging services. Touted as able to deliver critical information colleges, universities and even municipalities hoping to coordinate and protect the *physical security* of the general public have spent tens of millions of dollars to install such systems. Unfortunately, these products will not work as advertised and provide a false sense of security to their users.

In this paper, we explore the limitations of third party *Emergency Alert Systems* (EAS). In particular, we show that because of the currently deployed cellular infrastructure, such systems will not be able to deliver a high volume of emergency messages in a short period of time. *This identifies a key failure in a critical security incident response and recovery mechanism (the equivalent of finding weaknesses in techniques such as VM snapshots for rootkits and dynamic packet filtering rules for DDoS attacks) and demonstrates its inability to properly function during the security events for which it was ostensibly designed.* The fundamental misunderstanding of the requirements necessary to successfully deploy this piece of security infrastructure are likely to contribute to real-world, human-scale consequences.

In so doing, we make the following contributions:

- **Emergency Event Characterization:** Through modeling and simulation based on real provider deployments, we provide the first public characterization of the impact of an emergency event on a cellular network. This contribution is novel in that it explores a range of realistic emergency scenarios and provides a better understanding of their failure modes.
- **Measure EAS over SMS for multiple emergency scenarios:** We provide data to debunk the common assertion made by many third-party vendors that large quantities of text messages can be delivered within a short period of time (i.e., seconds to minutes). We evaluate a number of different, realistic emergency scenarios and explain why a number of college campuses have reported “successful” tests of their systems. Finally, we provide a real-world example that very closely mirrors the results of our simulations.
- **Quantify Collateral Damage:** We characterize the presence of the additional traffic generated by third-party EAS over SMS and show that such traffic causes increased blocking of normal calls and text message, potentially preventing those in need of help from receiving it. We also discuss a number of ways in which these networks can cause unexpected failures (e.g., message delay, message reordering, alert spoofing).



**Fig. 1.** Before a message can be delivered, a mobile device must be located. To do so, the MSC requests that towers within a given area all transmit paging requests. If an when a device is found, the MSC forwards the message to the appropriate tower, which attempts to deliver it wirelessly. The status of the delivery attempt is then returned to the SMSC. If delivery failed, the SMSC will attempt delivery at a later time.

## 2 Network Architecture

We specifically examine GSM networks in these discussions as they represent the most widely deployed cellular technology in the world; however, it should be noted that message delivery for other technologies such as CDMA, IDEN and TDMA are very similar and are therefore subject to similar problems.

### 2.1 Cellular Network Architecture

**Sending a Message** While most users are only familiar with sending a text message from their phone, known as *Mobile Originated SMS (MO-SMS)*, service providers offer an expanding set of interfaces through which messages can be sent. From the Internet, for instance, it is possible to send text messages to mobile devices through a number of webpages, email and even instant messaging software. Third parties can also access the network using so-called SMS Aggregators. These servers, which can be connected directly to the phone network or communicate via the Internet, are typically used to send “bulk” or large quantities of text messages. Aggregators typically inject messages on behalf of other companies and charge their clients for the service. Finally, most providers have established relationships between each other to allow for messages sent from one network to be delivered in the other.

After entering a provider’s network, messages are sent to the *Short Messaging Service Center (SMSC)*. SMSCs perform operations similar to email handling servers in the Internet, and store and forward messages to their appropriate destinations. Because messages can be injected into the network from so many external sources, SMSCs typically perform aggressive spam filtering on all incoming messages. All messages passing this filtering are then converted and copied into the necessary SMS message format and encoding and then placed into a queue to be forwarded to their final destination.

**Finding a Device** Delivering messages in a cellular network is a much greater challenge than in the traditional Internet. Chief in this difficulty is that users in a cellular

network tend to be mobile, so it is not possible to assume that users will be located where we last found them. Moreover, the information about a user's specific location is typically limited. For instance, if a mobile device is not currently exchanging messages with a base station, the network may only know a client's location at a very coarse level (i.e., the mobile device may be known to be in a specific city, but no finer-grained location information would be known). Accordingly, the SMSC needs to first find the general location for a message's intended client before anything else can be done.

A server known as the *Home Location Register* (HLR) assists in this task. This database acts as the permanent repository for a user's account information (i.e., subscribed services, call forwarding information, etc). When a request to locate a user is received, the HLR determines whether or not that device is currently turned on. If a mobile device is currently powered off, the HLR instructs the SMSC to store the text message and attempt to deliver it at another time. Otherwise, the HLR tells the SMSC the address of the *Mobile Switching Center* (MSC) currently serving the desired device. Having received this location information, the SMSC then forwards the text message on to the appropriate MSC.

**Wireless Delivery** As mentioned earlier, even the MSC may not know more information about a targeted device's location. In order to determine whether or not the current base station serving this device is known, the MSC queries the *Visitor Location Register* (VLR), which temporarily stores information about clients while they are being served by the MSC. In most cases, this information is not known, and so the MSC must begin the extensive and expensive process of locating the mobile device. The MSC completes this task by generating and forwarding paging requests to all of its associated base stations, which may number in the hundreds. This process is identical to locating a mobile device for delivery of a voice call.

Upon receiving a paging request from the MSC, a base station attempts to determine whether or not the targeted device is nearby. To achieve this, the base station attempts to use a series of *Control Channels* to establish a connection with the user. First, the base station broadcasts a paging request over the *Paging Channel* (PCH) and then waits for a response. If the device is nearby and hears this request, it responds to the base station via the *Random Access Channel* (RACH) to alert the network of its readiness to receive information. When this response is received, the network uses the *Access Grant Channel* (AGCH) to tell the device to listen to a specific *Standalone Dedicated Control Channel* (SDCCH) for further exchanges. Using this SDCCH, the network is able to authenticate the client, perform a number of maintenance routines and deliver the text message. By limiting the operations necessary to deliver a text message to the control channels used for call setup, such messages can be delivered when all call circuits, known as *Traffic Channels* (TCHs) are busy.

When the attempt to deliver the message between the targeted device and the base station is complete, the device either confirms the success or failure of delivery. This status information is carried back through the network to the SMSC. If the message was successfully delivered, the SMSC deletes it. Otherwise, the SMSC stores the message until a later period, at which time the network re-attempts delivery. Figure 1 offers an overview of this entire process.

## 2.2 Third-Party Provider Solutions

In the past few years, a significant number of third-parties offering to deliver alert messages (and other information services) via text messaging have appeared. Citing the need for improved delivery targeted to a highly mobile population, many such services advertise text messaging as an instant, targeted disseminator capable of delivering of critical information to tens of thousands of mobile phones when it is most needed. These systems have been extensively deployed on college and university campuses throughout the United States.

The architecture of these systems is relatively simple. Whether activated through a web interface [7, 10, 35, 45, 46], directly from a phone [18], or as software running on a campus administrator's computer [34, 29], these services act as SMS aggregators and inject large numbers of text messages into the network. Colleges and universities subscribing to these services then collect mobile phone numbers from students, faculty and staff. In the event of an alert, all or a subset of the collected numbers can be targeted. While network providers may offer some limited information back to the third party, aggregators are largely unaware of conditions in the network or the geographic location of any specific individual.

## 3 Modeling Emergency Events in Real Environments

To determine whether there exists a mismatch between the current cellular text messaging infrastructure and third party EAS, it is necessary to observe such systems during an emergency. However, because large scale physical security incidents are rare, we apply a number of modeling techniques to help characterize such events.

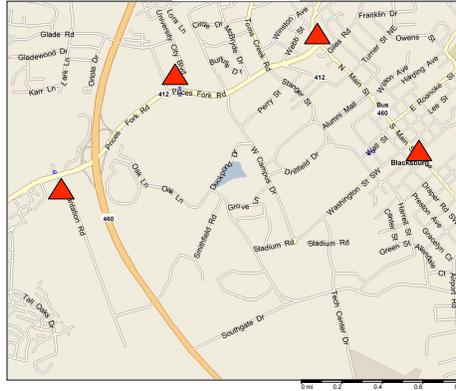
### 3.1 Location Selection & Characterization

The events that unfolded at the Virginia Polytechnic Institute and State University ("Virginia Tech") on April 16, 2007 have become one of the primary motivations behind the calls to use SMS as the basis of an emergency system. Many argue that had such a system been in place during what became the deadliest campus shooting in US history, countless lives could have been saved. However, a thorough examination of such claims has not been conducted. In particular, it is not clear whether or not the messages transmitted by such a system would have reached all students before the Norris Hall shootings. Accordingly, we have selected Virginia Tech as our location to characterize.

Located in southwestern Virginia, this land grant university is home to over 32,000 students, faculty and staff [48]. For the purposes of this work, we assume that just under half (15,000) of these individuals subscribe to a GSM network. As is shown by the red triangles in Figure 2, the major GSM provider in this area provides service to the campus of Virginia Tech from four base stations.<sup>1</sup> Given that each base station has three sectors (each covering a 120 degree range), we assume that the campus itself is covered by 8 of the 12 total sectors in the area.

---

<sup>1</sup> This is the actual configuration of the major GSM carrier in this area, as confirmed through conversations with this provider.



**Fig. 2.** The placement of base stations (red triangles) for a major GSM provider near Virginia Tech. Given that each base station has three sectors, the campus itself receives service from approximately eight total sectors.

### 3.2 Mathematical Characterization of Emergencies

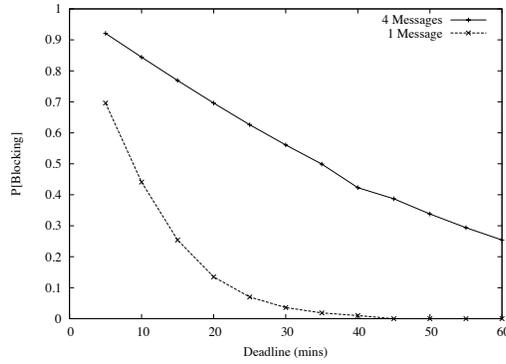
The first step in characterizing a cellular network during an emergency is determining capacity. In particular, we are interested in understanding the minimum time required to deliver emergency messages. If this time is less than the goal of 10 minutes set forth in by the current public EAS policies and the WARN Act [40], then such a system may indeed be possible. However, if this goal can not be met, current networks can not be considered as good candidates for EAS message delivery.

Given that most sectors have a total of 8 SDCCHs, that it takes approximately four seconds to deliver a text message in a GSM network [9, 28] and the information above, the capacity of the GSM network serving the campus of Virginia Tech would require the following amount of time to deliver a single message to 15,000 recipients:

$$\begin{aligned}
 C &= 15,000 \text{ msgs} \times \frac{1 \text{ campus}}{8 \text{ sectors}} \times \frac{1 \text{ sector}}{8 \text{ SDCCHs}} \\
 &\quad \times \frac{4 \text{ secs}}{1 \text{ message}} \\
 &\approx 938 \text{ sec} \\
 &\approx 15.6 \text{ mins}
 \end{aligned}$$

Because the contents of emergency messages are likely to exceed the 160 character limit of a single text message, providers and emergency management officials have estimated the number of messages is likely to increase by at least four times:

$$\begin{aligned}
 C &= 15,000 \text{ msgs} \times \frac{4 \text{ msgs}}{\text{user}} \times \frac{1 \text{ campus}}{8 \text{ sectors}} \\
 &\quad \times \frac{1 \text{ sector}}{8 \text{ SDCCHs}} \times \frac{4 \text{ secs}}{1 \text{ msg}} \\
 &\approx 3752 \text{ secs} \\
 &\approx 62.5 \text{ mins}
 \end{aligned}$$



**Fig. 3.** Calculated blocking probabilities versus delivery windows for emergency SMS traffic.

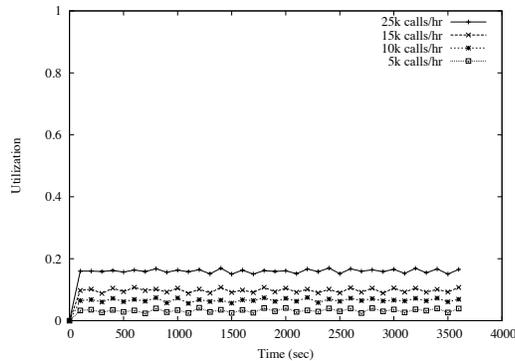
**Table 1.** Simulation parameters

$\mu_{TCH}^{-1}$	120 sec [32]
$\mu_{SDCCH,call}^{-1}$	1.5 sec [32]
$\mu_{SDCCH,SMS}^{-1}$	4 sec [28, 9]
$\lambda_{call,regular}$	10,000 calls/campus/hr .347 calls/sector/sec
$\lambda_{SMS,regular}$	21K msgs/campus/hr 0.75 msgs/sector/sec

The above calculations represent an optimistic minimum time for the delivery of all messages. For instance, it is highly unlikely that all eight SDCCHs will be available for delivering text messages as these channels are also used to establish voice calls and assist with device mobility. Moreover, contention between emergency messages for SDCCHs will also be a significant factor given that the SMSC is unaware of traffic conditions in individual sectors. Finally, depending on conditions within the network, each message is likely to experience different delays. To better characterize these factors, we apply a simple Erlang-B queuing analysis of the system. In a system with  $n$  servers and an offered load of  $A = \frac{\lambda}{\mu^{-1}}$ , where  $\lambda$  is the intensity of incoming messages and signaling traffic and  $\mu$  is the rate at which a single server can service incoming requests, the probability that an incoming emergency message is blocked is:

$$P_B = \frac{\frac{A^n}{n!}}{\sum_{l=0}^{n-1} \frac{A^l}{l!}} \quad (1)$$

Figure 3 compares an imposed deadline for delivering all SMS-based emergency messages against the expected blocking. We note that while Poisson arrival is not appropriate for modeling traffic on the Internet, it is regularly used in telecommunications. Like the capacity equations, *this calculation shows that such large volumes of messages can not be delivered in a short period of time, even without the presence of traffic from normal operations.*



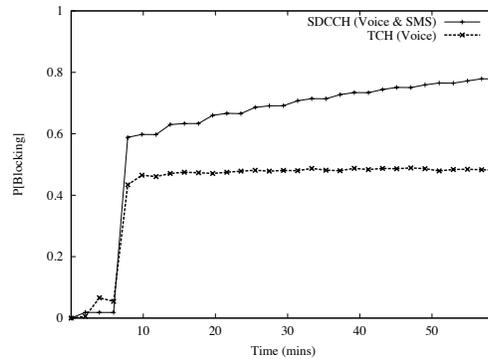
**Fig. 4.** The average utilization of control channels (SDCCHs) for a variety of traffic intensities.

## 4 Simulating Emergency Events

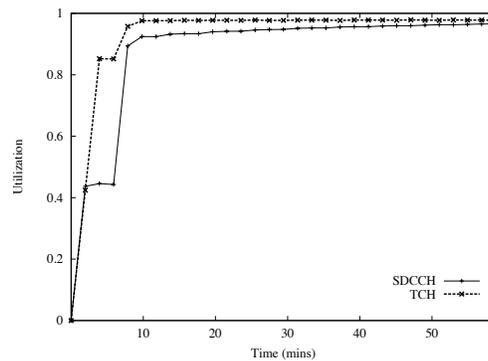
EAS over SMS traffic may still improve the physical security of its intended recipients even though it can not be delivered to the entire population within a 10 minute time period. If such information can be sent *without interfering with other traffic*, it could be argued that it would remain beneficial to at least some portion of the receiving population.

To better understand the impact of this security incident response and recovery mechanism on other traffic, we further characterize a number of emergency scenarios. While the calculations provided in the previous section and a post-9/11 government study on national text messaging capacity[28] are a good start, neither of these approximations help us understand the complex dynamics of the range of emergency scenarios. We therefore use a GSM simulator developed in previous work [41, 42, 44] and extend it for our needs. This tool focuses on the wireless portion of the network and allows the interaction between various resources to be characterized. This simulator was designed according to 3GPP standards documents, input from commercial providers and given optimal settings where applicable [22] so that our results are as conservative as possible.<sup>2</sup> Table 1 provides a summary of additional parameters representing busy hour load conditions (i.e., rush hour) and channel holding/service times. All experiments represent the average of 500 runs, the inputs for which were generated according to an exponential interarrival time using the Mersenne Twister Pseudo Random Number Generator [16]. Confidence intervals of 95% for all runs were less than two orders of magnitude from the mean, and are therefore too small to be shown. Given this system, we are able to explore the details of an emergency without having to wait for such an event occur or requesting log data from cellular providers. In the following subsections, we offer views of normal operations, surges of messages and a full emergency situation with EAS over SMS deployed.

<sup>2</sup> We note that some providers configure their network such that incoming text messages use four of the eight SDCCHs to decrease delivery time. However, this configuration results in higher blocking during busy periods, so we do not consider it further.



**Fig. 5.** The average blocking experienced during a large-scale emergency. Note that blocking on TCHs remains steady in spite of increasing call loads due to increased blocking on the SDCCH.



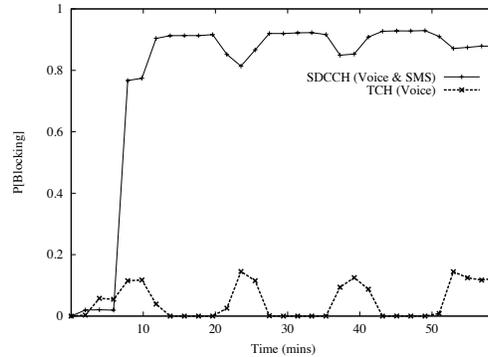
**Fig. 6.** Channel utilization observed during a large-scale emergency. The network becomes saturated almost immediately after the emergency event is realized.

#### 4.1 Normal Traffic

Our first set of experiments represent normal network behavior. Figure 4 shows the robustness of these networks to high traffic, illustrating very low SDCCH utilization rates for all of the offered loads. This graph reinforces the case for using SDCCHs for SMS delivery. Even in the 25,000 calls per hour case, during which nearly more than 55% of incoming calls can not be completed, SDCCHs are utilized at approximately 18%.

#### 4.2 Emergency Scenarios

Users having received notification of an emergency are unlikely to maintain normal usage patterns. In particular, users are likely to attempt to contact their friends and/or family soon after learning about such conditions. Whether by text message or phone call, however, such instinctual communication leads to significant congestion in cellular networks. This phenomenon lead to a spike in the number of attempted calls to the Washington D.C. are by over 1000% percent on September 11th [28]. Accordingly,



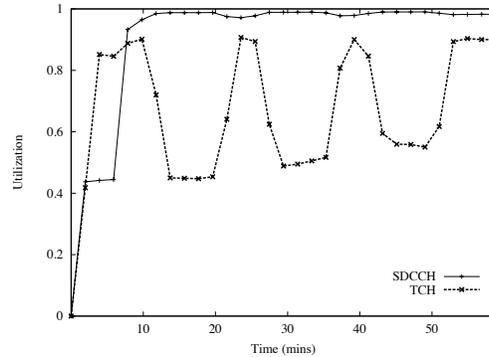
**Fig. 7.** Average blocking during a large-scale emergency in the presence of EAS over SMS. The network experiences blocking rates of approximately 90% when EAS messages are transmitted.

increases of varying intensities and characteristics representing reactionary usage must be considered when designing text messaging-based EAS. We explore two such scenarios, which assume that the third-party EAS over SMS provider has configured their system to deliver all messages within the WARN Act’s 10 minute requirement [40], that SMSCs retransmit previously undeliverable messages once every 15 minutes and assume that 4 messages per user are transmitted by the EAS over SMS system when an emergency occurs.

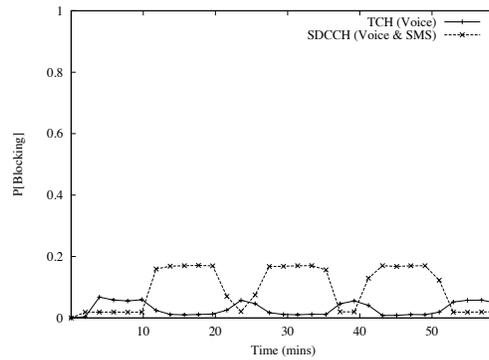
**Large-Scale Emergencies** Whereas small events may have a gradual increase in the volume of traffic, large-scale emergencies are often characterized by substantial and rapid spikes in usage, followed by continued gradual growth. We explore this worst case to understand the full extent of the problems such third party solutions may create. We therefore model a Virginia Tech-like event in which normal traffic increases by 1000% [28], with a 500% increase occurring over the course over a few minutes and the outstanding 500% being distributed across the remaining hour. Like the previous scenario, we conduct these experiments with and without the presence of EAS over SMS.

As expected, the sudden surge of traffic during the emergency almost immediately makes communications difficult. Figure 5 shows blocking rates of approximately 47% for TCHs and between 59% and 79% for SDCCHs. With both SDCCHs and TCHs experiencing near total utilization as shown in Figure 6, the network is already significantly overloaded and unable to deliver additional traffic.

The presence of traffic generated by an EAS over SMS system makes this scenario considerably worse. As shown in Figure 7, call and SMS blocking on SDCCHs almost immediately reaches between 80 and 85%. Like the previous scenario, call blocking on TCHs actually decreases. Such a decrease can again be attributed to the elevated blocking on the SDCCHs, as Figure 8 demonstrates that TCHs remain idle in spite of an increased call volume.



**Fig. 8.** Channel utilization during a large-scale emergency with EAS over SMS. TCH utilization falls significantly when EAS messages are sent, meaning fewer voice calls are delivered.

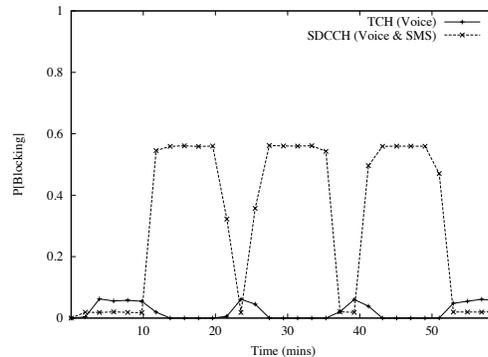


**Fig. 9.** The average blocking observed during a test (one message) of a third-party EAS over SMS system with only 25% of students registered.

### 4.3 Testing Campus Alert Systems

The discrepancy between the scenarios presented thus far and the reports of successful tests of deployed systems is a result of a number of factors. As previously mentioned, the 160 character limit per text message often requires the transmission of multiple text messages during an emergency. Most system tests, however, typically involve sending a single message. Traffic in these tests is therefore sent at one-fourth the volume of more realistic emergency scenarios. The second difference is the size of the affected population. While many universities offer these systems as an optional service to their students, an increasing number are beginning to make enrollment mandatory. Accordingly, *current tests attempt to contact only a subset of students with a smaller volume of traffic than would be used in a real emergency.*

We use reports of successful tests as input for our final set of experiments. In particular, we attempt to recreate the environment in which these tests are occurring. We site information from officials at the University of Texas Austin [20] and Purdue University [31], each of which have reported transmitting messages to approximately 10,000 participants. Note that this represents roughly 25% of the undergraduate student body



**Fig. 10.** The average blocking observed when four messages are transmitted and all other traffic remains constant.

at these institutions. We therefore reduce the receiving population at Virginia Tech to 7,500, of which only half are subscribers to the GSM provider.

Figure 9 shows the probability of blocking for this scenario. With approximately 18% blocking, such a system would appear to replicate current deployments - over 80% of recipients are reached within the first 10-minute long transmission. However, as is shown in Figure 10, by increasing the number of messages sent to this small group by a factor of four to allow for a longer emergency message, the probability of blocking increases to 58%. Because the transmission of multiple messages is more likely, campus emergency coordinators should test their systems based on this setting to gain a realistic view of its performance and behavior.

These two cases provide a more complete picture of the issues facing these systems. Whereas a third-party security response and recovery system may be able to deliver a small number of messages to one quarter of the students on campus, attempts to send more messages and therefore more meaningful communications quickly result in high blocking. Such systems are simply unable to support the rapid delivery of emergency messages to the entire population of the campus.

As corroboration of this final assertion and to further ground our results in reality, we note the results of a campus alert system deployed on the campus of Simon Fraser University in Burnaby, British Columbia, Canada. In April 2008, the University attempted to send test alert messages to 29,374 people; however, only 8600 were able to receive these messages [37]. Only 6500 of those having received the message were able to do so within five hours of it being sent, representing nearly an 80% rate of blocking. Worse still, many students reported an elevated rate of busy signals for many hours. These results are very similar to those shown in Figure 7, which while showing a slightly higher load, shows extremely close levels of blocking (approximately 85%). The analysis in this paper, in concert with this real-life test, clearly explains the failure of this security response and recovery mechanism to meet its requirement.

## 5 Discussion

### 5.1 3G Networks

We profiled the use of GSM networks in this work because they represent the most widely used cellular technology in the world. However, much faster third generation (3G) cellular systems are beginning to be deployed. With high speed data service available in many metropolitan areas, it would appear as if the analysis made in this paper will not remain relevant.

The migration to these new systems will not address these issues for a number of reasons. First, all cellular networks expend significant effort when establishing a connection. As demonstrated in Section 2, these operations include locating a targeted mobile device and performing significant negotiations before a single packet can be delivered. While the delivery rates of cellular data services have been steadily improving over the past decade, this setup and delivery of the first bit of information remains a significant bottleneck in the process. This means that while it is possible to download large files relatively quickly using such networks, beginning the download remains expensive. Second, many providers currently configure their 3G networks for the circuit switched delivery of text messages. Accordingly, such messages will continue to compete with voice calls for resources, leading to the same kinds of blocking conditions.

### 5.2 Message Delivery Order

Implicit in the misunderstanding of text messaging as a real-time service are misconceptions about the order in which messages will be delivered to targeted devices. Specifically, it is often assumed that messages will be delivered in the order in which they were injected by the sender. Message delivery order is in fact not always predictable.

The order in which messages are delivered can be affected by a number of factors. For instance, Traynor et al [41] showed that the SMSCs of different providers implement a variety of service algorithms, including FIFO and LIFO service disciplines. Accordingly, it is possible for two providers to deliver the same stream of messages in opposite order. Even if all carriers implemented the same delivery algorithm, congestion in the network can cause further disordering of packets. If an incoming text message is unable to be delivered due to a lack of resources on the air interface, the SMSC will store the message for a later attempt. However, if subsequent messages have been sent before this message fails and manage to gain the required resources, they will be delivered out of the sender's intended order. In an emergency such as a tornado, which may frequently change directions, such out of order delivery may actually send subscribers directly into the storm as opposed to away from it.

There are a number of emergency scenarios in which the above has occurred. During a wildfire evacuation at Pepperdine University in 2007, multi-part messages were transmitted to students and faculty to provide relocation instructions. However, some reported that the messages were not useful. One student later noted that "Each notification that was sent came through in six to eight text messages... And they were jumbled, not even coming in in order" [4]. More serious conflicts in message delivery order were noted at the Georgia Institute of Technology [6]. After a chemical spill in 2007,

a message alerting students and faculty to evacuate campus was transmitted. Later, instructions to ignore the evacuation notification were also sent. However, a number of students noted receiving the messages out of order [36], adding greater confusion to the situation. Similar problems have been reported at a number of other universities [8, 14].

### **5.3 Message Delay**

Examples of the delay that can be experienced during times of high volume are most easily observed during New Years Eve celebrations or the most recent US Presidential Inauguration. As hundreds of millions of users around the globe send celebratory greetings via SMS, service providers often become inundated with a flood of messages. Accordingly, the delivery of such messages has been noted to exceed more than six hours [11]. Even though providers often plan and temporarily deploy additional resources to minimize the number of blocked calls, the sheer volume of messages during such an event demonstrates the practical limitations of current systems. In spite of temporarily deploying additional towers, such delays are experienced even when cellular providers are aware that a high volume event will take place.

Why then has SMS been a successful means of communication during other national emergencies such as September 11th and Hurricanes Katrina and Rita? Numerous sources cite SMS as an invaluable service when both man-made and natural disasters strike [15, 26]. The difference between these events and other emergencies is the magnitude of messages sent. For instance, at the time of the attacks of September 11th, text messaging was still largely a fringe service in the United States. Had most users across the country attempted to communicate using SMS as their primary mode of communication, however, a report by the National Communications System estimates that current network capacities would need to be expanded by 100-fold [28] in order to support such a volume. The reliability of text messaging during Hurricane Katrina is due to similar reasons. Because only a very small number of people were communicating via text messaging, the towers undamaged by the storm were able to deliver such messages without any significant competition from other traffic. Moreover, because the network automatically attempted retransmission, users were more likely to receive text messages than calls. If SMS use during these events approached emergency levels, it would have experienced delays similar to those regularly observed on New Years Eve.

## **6 Improving Incident Response Communications**

From the discussions, mathematical characterizations and simulations in the previous sections, the mismatch between the current cellular infrastructure and current response mechanisms is clear. Accordingly, such systems can not currently form the basis of a reliable alert system in the timescales required by the WARN Act, regardless of promises made by third party systems. However, the ubiquity of cellular phones gives them a potential role in the delivery of critical information during an emergency. This role would be complementary to the other platforms of the Emergency Broadcasting System (Television, radio, etc.).

There are a number of solutions currently under consideration that may help in this space. The most well known is *cell broadcast*. Unlike the point to point operations required for the delivery of messages in current networks, cell broadcast would allow the rapid dissemination of emergency information through point to multipoint communications. Such a system could ideally reach the majority of cellular users in an area and would not require knowledge of each particular user's location. This option is backed by the Commercial Mobile Service Alert Advisory Committee, which is currently working on developing standards documents. However, while cell broadcast will significantly improve communications over current mechanisms, a number of critical problems remain. First, like traditional text messaging, information delivered via cell broadcast will not be authenticated. Second, the channels used for cell broadcast are relatively bandwidth limited. The rate with which complex messages can be delivered during highly dynamic situations (e.g., an on-campus gunman) may therefore be lower than desired. Third, cell broadcast does little to address issues of coverage, which may become exacerbated during an emergency. For instance, in the event of a natural disaster or attack, it is highly likely that some cellular towers will be damaged or unpowered. Third, cell broadcast does not currently provide special options for hearing or visually impaired users. Finally, while cell broadcast is designed to deal with the overload of many simultaneous point to point connections, this technology is still relatively immature (i.e., standards are pending) and has not been deployed and measured in any large scale, public, scientific fashion. When this mechanism is deployed it may indeed improve communications during such scenarios; however, it is critical to recognize that our widely deployed current infrastructure is a deeply flawed information system when used for emergency communications.

The increasing capabilities of mobile devices could potentially be leveraged to improve the reach of communications during an emergency. For instance, cell broadcast could potentially be used to signal mobile phones equipped with 802.11 wireless cards to connect to a specific website containing regularly updated information. This information could potentially include sound or video clips similar to traditional EAS broadcasts to assist hearing and visually impaired users. The presence of 802.11 cards could also assist in improving coverage. Borrowing techniques from delay tolerant networking [13, 19, 50] may allow phones passing through areas with poor or no cellular reception to inform other devices of the current alert. The recent addition of AM/FM radio tuners in a variety of phones [30, 17] may further assist in this process. Specifically, mobile devices could be used to immediately tune into the more traditional Emergency Broadcast system, ensuring consistent dispersal of information. The presence of AM/FM radios would also significantly improve the robustness of communications in a large scale emergency as cellular or 802.11 outages in a user's immediate vicinity would not prevent information from continuing to be delivered.

These suggestions also face a number of research questions. Like the cell broadcast case, a strong method of authenticating incoming notifications would be necessary. This issue may potentially be addressed by directing phones to an SSL-based webpage run by the university. Moreover, studies focused on latency and data provenance for delay tolerant networks in densely populated urban areas and campuses would also need to be conducted. Until such systems are realized, however, legislators and the general

public should not rely upon text messaging or third party EAS providers for delivering emergency information.

## 7 Related Work

Following the events of September 11th, 2001, curiosity about the ability to use text messaging as the basis of a reliable communications system during times of crisis arose. In response, the National Communications System (NCS) conducted an investigating the use of text messaging during a nation-wide emergency, which through simple calculations concluded that current systems would require “100 times more capacity to meet [the] load” created by widespread use of text messaging [28]. A related study by the European Telecommunications Standard Institute (ETSI) identified the increasing prevalence of spam as a significant threat to the operation of cellular networks during an emergency [12]. However, both studies were limited to high-level calculations of a single emergency scenario and neither considered the use of third party EAS over SMS systems. Our study conducted the first characterization and simulation of multiple scenarios for EAS over cellular services and compared them directly to real-world, on-campus testing.

The specific impacts on the reliability and security of such networks under torrents of text messages have also been explored. Traynor et al. [41, 43] noted that an attacker could exploit connections between the Internet and cellular networks to cause significant outages. With the bandwidth available to a cable modem, an attacker could send a small but targeted stream of text messages to a specific geographic region and prevent legitimate voice and text messages from being delivered. While subsequent research was able to better characterize and provide mitigations against such attacks [42], it was ultimately discovered that a more basic problem was responsible. Instead of simply being a matter of using a low-bandwidth channel to deliver data, the real cause of such attacks was a result of fundamental tension between cellular networks and the Internet. Specifically, because cellular networks can not amortize the significant cost of connection establishment when delivering data, they are fundamentally vulnerable to such attacks [44]. Accordingly, as long as text messages are delivered in the point to point fashion as is done now, the expense of establishing connections with each and every phone in an area will remain prohibitively expensive.

Whether as an unintended consequence or deliberate act, the flooding behavior exhibited in this above work closely resembles Denial of Service (DoS) attacks on the Internet. The research community has responded with attempts to classify [27] and mitigate [1–3, 5, 21, 23, 24, 33, 39, 38, 47, 49] such attacks. However, such attacks are only beginning to be understood in the context of cellular networks, making the direct application of these solutions unsuitable.

## 8 Conclusion

Cellular networks are increasingly becoming the primary means of communication during emergencies. Riding the widely-held perception that text messaging is a reliable method of rapidly distributing messages, a large number of colleges, universities and municipalities have spent tens of millions of dollars to deploy third-party EAS over cellular systems. However, this security incident response and recovery mechanism simply

does not work as advertised. Through modeling, a series of experiments and corroborating evidence from real-world tests, we have shown that these networks can not meet the 10 minute alert goal mandated by the public EAS charter and the WARN Act. Moreover, we have demonstrated that the extra text messaging traffic generated by third party EAS will cause congestion in the network and may potentially block upwards of 80% of normal requests, potentially including calls between emergency responders or the public to 9-1-1 services. Accordingly, it is critical that legislators, technologists and the general public understand the fundamental limitations of this mechanism to safeguard physical security and public safety and that future solutions are thoroughly evaluated before they are deployed.

## Acknowledgements

This work was supported in part by 3G Americas. Any opinions, findings, conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of 3G Americas. We would also like to thank the cellular providers that helped us more accurately model this issue. This work was also supported in part by the US National Science Foundation (CNS-0916047). Any opinions, findings, conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the National Science Foundation.

## References

1. D. Andersen. Mayday: Distributed Filtering for Internet Services. In *Proceedings of the USENIX Symposium on Internet Technologies and Systems (USITS)*, 2003.
2. T. Anderson, T. Roscoe, and D. Wetherall. Preventing Internet Denial of Service with Capabilities. In *Proceedings of ACM HotNets*, 2003.
3. K. Argyraki and D. R. Cheriton. Scalable Network-layer Defense Against Internet Bandwidth-Flooding Attacks. *ACM/IEEE Transactions on Networking (TON)*, 2009.
4. S. Blons. Emergency team aids efforts. <http://graphic.pepperdine.edu/special/2007-10-24-emergencyteam.htm>, 2007.
5. M. Casado, P. Cao, A. Akella, and N. Provos. Flow Cookies: Using Bandwidth Amplification to Defend Against DDoS Flooding Attacks. In *Proceedings of the International Workshop on Quality of Service (IWQoS)*, 2006.
6. T. Christensen. Ga. Tech Building Cleared After Blast. [http://www.11alive.com/news/article\\_news.aspx?storyid=106112](http://www.11alive.com/news/article_news.aspx?storyid=106112), 2007.
7. CollegeSafetyNet. Campus Alert, Campus Security, Emergency Warning, college safety Crisis notification, Reverse 911, Mass emergency notification, Emergency Alert System, Cell phone alerts, Email alerts, Text Message Alerts, Student warning system, Student notification, campus notification, and Mass notification at CollegeSafetyNet.com. <http://www.collegesafetynet.com/>, 2008.
8. Courant.com. University Emergency SMS service doesn't deliver. <http://www.courant.com>, November 13, 2007.
9. B. K. Daly. WIRELESS ALERT & WARNING WORKSHOP. [http://www.oes.ca.gov/WebPage/oeswebsite.nsf/ClientOESFileLibrary/Wireless%20Alert%20and%20Warning/\\$file/ATT-OES-2.ppt](http://www.oes.ca.gov/WebPage/oeswebsite.nsf/ClientOESFileLibrary/Wireless%20Alert%20and%20Warning/$file/ATT-OES-2.ppt), 2007.
10. e2Campus. Mass Notification Systems for College, University & Higher Education Schools by e2Campus: Info On The Go! <http://www.e2campus.com/>, 2008.
11. A.-M. Elliott. Texters to experience 6 hour delays on New Year's Eve. <http://www.pocket-lint.co.uk/news/news.phtml/11895/12919/palm-new-years-text-delay.phtml>, 2007.
12. European Telecommunications Standards Institute. Analysis of the Short Message Service (SMS) and Cell Broadcast Service (CBS) for Emergency Messaging applications; Emergency Messaging; SMS and CBS. Technical Report ETSI TR 102 444 V1.1.1.
13. K. Fall. A Delay-Tolerant Network Architecture for Challenged Internets. In *Proceedings of the Conference on Applications, Technologies, Architectures and Protocols for Computer Communications (COMM)*, 2003.
14. L. Ganosellis. UF to test texting alerts after LSU glitch. [http://www.alligator.org/articles/2008/01/08/news/uf\\_administration/lsu.txt](http://www.alligator.org/articles/2008/01/08/news/uf_administration/lsu.txt), 2008.
15. D. Geer. Wireless victories, September 11th, 2001. *Wireless Business & Technology*, 2005.
16. J. Hedden. Math::Random::MT::Auto - Auto-seeded Mersenne Twister PRNGs. <http://search.cpan.org/~jhedden/Math-Random-MT-Auto-5.01/lib/Math/Random/MT/Auto.pm>. Version 5.01.

17. HTC Corporation. HTC Tattoo Specifications. <http://www.htc.com/europe/product/tattoo/specification.html>, 2009.
18. Inspiron Logistics. Inspiron Logistics Corporation WENS - Wireless Emergency Notification System for Emergency Mobile Alerts. <http://www.inspironlogistics.com/>, 2008.
19. S. Jain, K. Fall, and R. Patra. Routing in a Delay Tolerant Network. In *Proceedings of the Conference on Applications, Technologies, Architectures and Protocols for Computer Communications (COMM)*, 2004.
20. E. Jaramillo. UT director: Text alerts effective. <http://www.dailytexanonline.com/1.752094>, 2008.
21. A. Keromytis, V. Misra, and D. Rubenstein. SOS: Secure Overlay Services. In *Proceedings of ACM SIGCOMM*, 2002.
22. C. Luders and R. Haferbeck. The Performance of the GSM Random Access Procedure. In *Vehicular Technology Conference (VTC)*, pages 1165–1169, June 1994.
23. R. Mahajan, S. M. Bellovin, S. Floyd, J. Ioannidis, V. Paxson, and S. Shenker. Controlling High Bandwidth Aggregates in the Network. *Computer Communications Review*, 32(3):62–73, July 2002.
24. A. Mahimkar, J. Dange, V. Shmatikov, H. Vin, and Y. Zhang. dFence: Transparent Network-based Denial of Service Mitigation. In *Proceedings of USENIX Networked Systems Design and Implementation (NSDI)*, 2007.
25. K. Maney. Surge in text messaging makes cell operators :-). [http://www.usatoday.com/money/2005-07-27-text-messaging\\_x.htm](http://www.usatoday.com/money/2005-07-27-text-messaging_x.htm), July 27 2005.
26. J. McAdams. SMS does SOS. [http://www.fcw.com/print/12\\_11/news/92790-1.html](http://www.fcw.com/print/12_11/news/92790-1.html), 2006.
27. J. Mirkovic and P. Reiher. A Taxonomy of DDoS Attacks and DDoS Defense Mechanisms. *ACM SIGCOMM Computer Communication Review*, 34(2):39–53, 2004.
28. National Communications System. SMS over SS7. Technical Report Technical Information Bulletin 03-2 (NCS TIB 03-2), December 2003.
29. National Notification Network (3n). 3n InstaCom Campus Alert - Mass Notification for Colleges and Universities. <http://www.3nonline.com/campus-alert>, 2008.
30. C. Nettles. iPhone 3 to have Broadcom BCM4329, 802.11N/5GHz Wireless, FM transmitter/receiver. <http://www.9to5mac.com/broadcom-BCM4329-iphone-802.11n-FM>, 2009.
31. M. Nizza. This is only a (text messaging) test. <http://thelede.blogs.nytimes.com/2007/09/25/this-is-only-a-text-messaging-test/?scp=5&sq=Emergency%20Text%20Messaging&st=cse>, 2007.
32. Nyquetek, Inc. Wireless Priority Service for National Security. <http://wireless.fcc.gov/releases/da051650PublicUse.pdf>, 2002.
33. B. Parno, D. Wendlandt, E. Shi, A. Perrig, and B. Maggs. Portcullis: Protecting Connection Setup from Denial of Capability Attacks. In *Proceedings of ACM SIGCOMM*, 2007.
34. Reverse 911. Reverse 911 - The only COMPLETE notification system for public safety. <http://www.reverse911.com/index.php>, 2008.
35. Roam Secure. Roam Secure. <http://www.roamsecure.net/>, 2008.
36. shelbinator.com. Evacuate! Or Not. <http://shelbinator.com/2007/11/08/evacuate-or-not/>, 2007.
37. Simon Fraser University. Special Report on the April 9th Test of SFU Alerts. [http://www.sfu.ca/sfualerts/april08\\_report.html](http://www.sfu.ca/sfualerts/april08_report.html), 2008.
38. A. Stavrou, D. L. Cook, W. G. Morein, A. D. Keromytis, V. Misra, and D. Rubenstein. WebSOS: An Overlay-based System For Protecting Web Servers From Denial of Service Attacks. *Journal of Computer Networks, special issue on Web and Network Security*, 48(5):781 – 807, 2005.
39. A. Stavrou and A. Keromytis. Countering DOS Attacks With Stateless Multipath Overlays. In *Proceedings of ACM Conference on Computer and Communications Security (CCS)*, 2005.
40. The 109th Senate of the United States of America. Warning, Alert, and Response Network Act. [http://thomas.loc.gov/cgi-bin/query/z?c109:H.R.1753](http://thomas.loc.gov/cgi-bin/query/z?c109:H.R.1753;), 2005.
41. P. Traynor, W. Enck, P. McDaniel, and T. La Porta. Exploiting Open Functionality in SMS-Capable Cellular Networks. *Journal of Computer Security (JCS)*, 2008.
42. P. Traynor, W. Enck, P. McDaniel, and T. La Porta. Mitigating Attacks on Open Functionality in SMS-Capable Cellular Networks. In *IEEE/ACM Transactions on Networking (TON)*, volume 17, 2009.
43. P. Traynor, M. Lin, M. Ongtang, V. Rao, T. Jaeger, T. La Porta, and P. McDaniel. On Cellular Botnets: Measuring the Impact of Malicious Devices on a Cellular Network Core. In *Proceedings of the ACM Conference on Computer and Communications Security (CCS)*, 2009.
44. P. Traynor, P. McDaniel, and T. La Porta. On Attack Causality in Internet-Connected Cellular Networks. In *Proceedings of the USENIX Security Symposium*, 2007.
45. TXTLaunchPad. TXTLaunchPad provides Bulk SMS text message alerts to businesses, schools, and advertisers. <http://www.txtlaunchpad.com/>, 2007.
46. Voice Shot. automated emergency alert notification call - VoiceShot. <http://www.voiceshot.com/public/urgentalert.asp?ref=uaemergencyalert>, 2008.
47. M. Walfish, M. Vutukuru, H. Balakrishnan, D. Karger, and S. Shenkar. DDoS Offense by Offense. In *Proceedings of ACM SIGCOMM*, 2006.
48. Wikipedia. Virginia Polytechnic Institute and State University. [http://en.wikipedia.org/wiki/Virginia\\_Tech](http://en.wikipedia.org/wiki/Virginia_Tech), 2008.
49. X. Yang, D. Wetherall, and T. Anderson. TVA: A DoS-limiting Network Architecture. *IEEE/ACM Transactions on Networking (TON)*, 2009.
50. W. Zho, M. Ammar, and E. Zegura. A message ferrying approach for data delivery in sparse mobile ad hoc networks. In *Proceedings of the International Symposium on Mobile Ad Hoc Networking & Computing (MOBIHOC)*, 2004.