

# SmartCell: Small-Scale Mobile Congestion Awareness

Paul Schmitt, Daniel Iland, and Elizabeth Belding  
 {pschmitt, iland, ebelding}@cs.ucsb.edu  
 University of California, Santa Barbara

**Abstract**—Despite improvements and expansion of cellular coverage in developing regions, a substantial qualitative divide remains. Maps that display the presence or absence of cellular coverage mask critical differences in infrastructure performance and client load. In order to illuminate challenges faced by users of such mobile networks, we collect and analyze GSM network measurements at the local-scale. We discover higher network congestion in developing regions as well as performance differences between available carriers in each location. Based on our findings, we propose an app, called SmartCell, that informs and empowers users in near real-time to seek out improved mobile connectivity.



## 1 INTRODUCTION

A look at the worldwide mobile cellular subscription rate indicates that by the end of 2015, there will be more than 7 billion mobile cellular subscriptions, corresponding to a penetration rate of 97%<sup>1</sup>. While this seems like an immediate cause for celebration, a deeper look is required to more fully understand what this number represents. If we parse the numbers by region, we find that in 2015, 78% of residents of Europe and 77% in the Americas had mobile broadband subscriptions, compared with only 42.3% in Asia and the Pacific and 17.4% in Africa. Breaking coverage into geographic region begins to give us an understanding of access differences; however, stopping there masks critical disparities in the *quality* of available cellular connectivity. In particular, the cellular technology (LTE vs 3G vs 2G), the data rate, the number of users per base station, and the cost of the subscription, are all critical factors that play into the persisting inequities between cellular users in developed and developing regions. Simply looking at who has access does not give an indication of whether the quality of that access is equivalent for different users in different locations.

To dive deeper into these disparities, we conducted a cellular infrastructure measurement campaign in three locations of diverse population characteristics. Specifically, we passively observed, collected and analyzed messages broadcast by cellular base stations in order to assess coverage quality and usability. We collected measurements in San Cristóbal Verapaz, Guatemala, a city of roughly 20,000 in one of the poorest countries in Latin America; in the Za’atari refugee camp in Jordan, the oldest and largest refugee camp in Jordan with a population of roughly 80,000 in only 6 square kilometers; and, for comparison, in Santa Barbara, California, a community of about 90,000 residents in 42 square miles. Through our measurements, we reveal a number of interesting and important anomalies about coverage in these different locations that begin to paint a

clearer picture of cellular quality divides. Most importantly, we discover chronic network congestion in some of the networks, which leads to a consistently poor quality of experience for associated users. Interestingly, we find that in our measurement locations, there are multiple providers available, which have varying levels of traffic load, ranging from minimal load (completely uncongested) to very heavy loads (chronically congested).

Residents in developing countries are acutely aware of coverage disparities between carriers, and often use SIM cards from multiple carriers, switching between carriers depending on location. Our interviews with Za’atari residents indicate that they typically carry SIM cards for each available provider, but do not have any specific algorithm for selecting which to use at a given time. Reduced cost calling and texting between users on the same carrier creates a network effect, and may cause users to prefer a given SIM, simply because the provider is the most popular. This leads to uneven traffic distributions and load between carriers.

Based on our findings, we propose a system called SmartCell, an Android application that gives users awareness of quality of service by detecting congestion on their cellular base station. SmartCell informs the user when congestion is detected, thereby empowering her with information about the local network. With this information, the user has multiple options for obtaining better quality of service. Most simply, the user could decide to use one of her alternate SIM cards, switch to an underloaded network, and obtain higher quality service. If she has a multi-SIM phone, SmartCell can automatically select a SIM from a less congested carrier to use for voice calls, SMS, and mobile data traffic. Without information about the quality of the given cellular network provided by SmartCell, the user is not informed that a better connectivity option is available, and is likely to continue to suffer with unacceptable performance.

1. <http://www.itu.int/en/ITU-D/Statistics/Documents/facts/ICTFactsFigures2015.pdf>

## 2 RELATED WORK

Two common approaches to understanding cellular performance are application-level studies that look at end-to-end characteristics and radio-level studies that focus on the access link. Our measurement of the cellular infrastructure is a radio-level study, that specifically utilizes messages broadcast over the GSM air interface. Recent work has illustrated the potential impact that cellular radio state has on end user experience [1], [2], [3]; and how air interface messages can be used to infer cellular user activity [4], [5]. Our work demonstrates how detailed, small-scale analysis can support research on local-level infrastructure. This approach is in contrast to efforts to measure network performance on a global or nationwide scale, such as the FCC’s Measuring Broadband America<sup>2</sup>, and enables us to separate performance attributable to client-facing infrastructure serving the measurement locations from that which is related to carrier’s core networks.

Our system explores advising users to manually switch between cellular providers based on observed quality of service. In April 2015 Google announced Project Fi, which shifts users between the T-Mobile and Sprint networks based on expected data speeds<sup>3</sup>. The goals of Project Fi and SmartCell are thematically related, but the implementations are distinct due to Google’s integration with two existing commercial providers and active user switching. Additionally, while SmartCell infers performance of BTSs through passive and independent observation, we expect Google has access to back-end carrier metrics for determining performance. The always-best-connected concept [6] also touches on the use of multiple networks, however this work assumes business relationships between providers. We do not require such relationships and simply inform the user of current conditions, leaving the choice of using an alternate carrier to users themselves.

## 3 BACKGROUND

Our measurement studies are focused on GSM networks as that is the most prevalent cellular technology in our selected communities. In this section we provide an overview of the GSM system and control messages we use for infrastructure characterization. We then briefly describe the locations we use for our case study.

### 3.1 GSM technology

The GSM core network is a complex hierarchy composed of many logical entities. The pertinent objects for our study are mobile stations (MS) and base station transceivers (BTS). The MS is the user device, commonly a mobile phone equipped with a SIM card. BTSs are the components that communicate with the user device over the air interface. MSs communicate with only one BTS at a time.

GSM control messages enable the inference of BTS control channel congestion serving the measurement locations. Figure 1 displays the message sequences that take place when an MS requests a private communication channel from the BTS, which is necessary for voice/SMS or data.

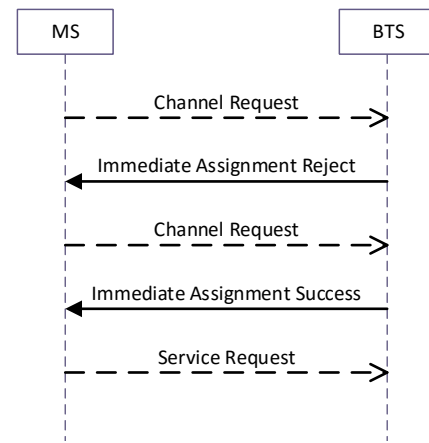


Fig. 1: GSM network immediate assignment procedure. We capture messages broadcast from the BTS to MS over the GSM air interface.

We capture messages broadcast from the BTS to the MS (solid arrows in Figure 1) on the common control channel. When an MS needs to use the network for a voice/SMS or data session, it issues a channel request to the base station. If a channel is available, the base station responds with an *immediate assignment success* message that provides information about the available channel. A BTS operating at full capacity such that it is unable to allocate a channel will issue an *immediate assignment reject* message, in accordance to the 3GPP 04.08 specification<sup>4</sup>, indicating no channel is available. Because these control messages are broadcast to all MSs connected to the BTS, we can leverage the immediate assignment success rate to approximate stand-alone dedicated control channel (SDCCH) blocking [7], [8], a GSM industry-standard key performance indicator (KPI) used to measure BTS congestion and to inform capacity expansion. We calculate the immediate assignment success rate by dividing the number of observed successful immediate assignment messages by the total number of immediate assignment (success and reject) messages.

#### 3.1.1 Message capture

We capture broadcast messages using a total of eight cellular phones, which consist of Samsung Galaxy Nexus, S2 and Galaxy S4 handsets with radio debug mode enabled, which logs all cellular communications to a computer via USB. We use `xgoldmon`, an open source tool that converts debug logs into packet capture files using the GSMTAP pseudo-header. Each phone records all of its own uplink traffic as well as all broadcast traffic sent by BTSs. Critically, our message capture is **non-invasive** and **non-intrusive**. Messages broadcast over common control channels are received in plaintext and intended for all MSs connected to a BTS. The GSMTAP pseudo-header does not include private user data. In all locations we capture using two phones per carrier, one set to prefer 2G and another set to prefer 3G. We do not capture 4G LTE as LTE is only available in the U.S. location.

2. <https://www.fcc.gov/measuring-broadband-america>

3. <https://fi.google.com/about/faq/#network-and-coverage-1>

4. <http://www.3gpp.org/DynaReport/0408.htm>

	Capture Duration (s)	Mean Immediate Assignments per second	Total Immediate Assignment Success Messages	Total Immediate Assignment Rejection Messages	Immediate Assignment Reject Percentage
AT&T (US)	840,114	0.04	34,539	9	0.03%
T-Mobile (US)	379,897	0.05	19,112	2	0.01%
Claro (GT)	625,051	0.01	7,661	0	0.00%
Movistar (GT)	625,073	0.21	127,996	1,055	0.82%
Tigo (GT)	622,340	0.21	129,837	846	0.65%
Orange (JO)	243,678	0.09	21,675	1,355	5.88%
Umniah (JO)	243,678	0.12	27,385	1,550	5.36%
Zain (JO)	243,678	0.12	19,804	8,800	30.76%

TABLE 1: Overall mobile network statistics.

## 3.2 Capture locations

We analyze GSMTAP captures from three locations, each of which represents a different level of development, in order to discover manifestations of cellular infrastructure digital divides. In this section we provide background on each of the locations.

### 3.2.1 San Cristóbal Verapaz, Guatemala

Guatemala is one of the poorest countries in Latin America with a GDP per capita of around \$3,512. Despite this, mobile carriers within Guatemala have invested heavily in cellular infrastructure, providing connectivity to a large portion of the country. We collected measurements in the rural city of San Cristóbal Verapaz over eight days in January 2015.

Three GSM carriers offer service in Guatemala: Tigo, Claro, and Movistar. All three are available in San Cristóbal Verapaz. Tigo is Guatemala’s most popular carrier and has built the most extensive cellular infrastructure network. Claro is the least popular carrier in San Cristóbal Verapaz and coverage is more focused on 3G in urban areas. Movistar is the second most popular and its network consists of a higher percentage of 2G base stations compared to the other carriers. It is difficult to compare prices for the carriers as there is no equivalent package offered on all three. Generally, Tigo is slightly more expensive than the others (e.g. 30GTQ for 500MB versus 25GTQ for 500MB from Movistar). Movistar is the most affordable carrier and prepaid scratch cards used for credit are widely available.

### 3.2.2 Za’atari refugee camp, Jordan

Za’atari is the largest of four refugee camps in Jordan. Located in a previously rural desert region near the border with Syria, the camp was opened in July 2012 and its population quickly rose to 120,000. It currently hosts about 80,000 residents in roughly 6 square kilometers. We collected cellular measurements in Za’atari over three days in January 2015.

Mobile service in the camp is offered by three carriers: Zain, Umniah, and Orange. Much of their infrastructure is located outside the camp and was initially were built to serve the surrounding rural community that existed prior to the camp’s development. We discussed carriers with camp residents and were told that Orange is popular in the more urban areas of Amman, while Zain is popular for rural customers. Umniah was generally not used by people with whom we spoke. Zain is the most popular carrier in the camp as SIM cards for the network are given to residents by camp administration as they first arrive and register at the camp.

### 3.2.3 Santa Barbara, California

For comparison purposes, we capture messages on two cellular carriers in and around Santa Barbara, CA (~90,000 residents) including a mix of urban and suburban base stations. We focus on the two major GSM carriers in the U.S.: AT&T and T-Mobile. We collected traces over roughly ten days for AT&T and five days for T-Mobile in October 2015.

## 4 CELLULAR MEASUREMENT ANALYSIS

We investigate network performance and congestion via immediate assignment success and rejection messages. As described in Section 3, an MS issues a request to the BTS for a communication channel at the time that the MS needs to use the network for voice/SMS or data. In response, the BTS broadcasts an immediate assignment message to inform the MS whether or not its request is granted.

### 4.1 BTS messaging load

We examine ‘busyness’ for our observed BTSs on each of the carrier networks by calculating the number of immediate assignment messages per second. Table 1 summarizes our results across the measurement locations. It is important to note that immediate assignment messages are BTS-specific and they do not indicate the service footprint or population served by a BTS. As shown in the table, AT&T and T-Mobile in the U.S. have relatively lightly-used BTSs, with 0.04 and 0.05 immediate assignments per second, respectively. These results could indicate networks comprised of a larger number of BTSs, where each BTS serves a smaller area and hence a smaller number of users. Claro also exhibited very few messages. We believe this may be caused by Claro’s lack of popularity in our measurement location. Tigo and Movistar exhibit similar message loads and have the highest rates of immediate assignment messages, roughly four times that of the U.S. networks. The three carriers in Jordan have similar messaging rates compared with one another and fall in between the observed load on the U.S. networks and those in Guatemala.

### 4.2 BTS congestion

In terms of congestion, we find qualitative performance divides between locations. Predictably, we capture almost no immediate assignment rejection messages on the U.S. provider networks. This indicates adequately provisioned BTSs for the given client load. Apart from Claro, which has a very low load, the Guatemalan networks experience higher rejection rates (minimum 21 times higher) than AT&T, the

U.S. carrier with the highest rejection rate. Turning to the refugee camp network we see an even greater divide. The Jordanian carriers all have vastly higher rejection percentages than in the other locations, with Orange and Umniah both experiencing more than 150 times higher rejection rates than AT&T. Zain, the most popular carrier in the camp, has an extremely high rejection message rate (more than 1,000 times higher than AT&T). For perspective, prior work considers a 5% SDCCH blocking rate to be ‘high congestion’ [9].

While the rejection percentages may seem low, their impact can be drastic. Each rejection includes a value for a backoff timer that forces the MS to wait up to 255 seconds before resubmitting a resource request. In terms of user quality of experience, such backoffs result in very poor performance. For instance, the GSM specification classifies call setup times, the time between a user pressing the call button and the call connection, into three categories<sup>5</sup>. ‘Fast’ call setup is defined as 1 - 2 seconds; ‘normal’ 2 - 5; and ‘slow’ 5 - 10. We analyze rejection messages in the traces and find the Jordanian networks frequently send backoffs of up to 128 seconds during congestion events. Clearly, the additional time attributable to backoffs we observe extends call setup times far beyond the classification of ‘slow.’

### 4.3 Temporal congestion characteristics

Because the carriers serving Za’atari experienced the greatest congestion, we take a deeper look into carrier-specific performance in the traces to investigate any cross-network traffic load relationships. Figure 2 shows the percent of immediate assignment messages that were rejections in one minute bins over the course of a single day during our measurement window.

We see that the carriers’ rejection patterns do not resemble one another. The Orange network was congested in short, severe bursts. It also appears as though congestion on Orange was higher at particular times of the day that correspond with workday schedules (i.e. before and after lunch). Umniah, on the other hand, exhibited almost no evidence of congestion throughout the day. The Zain network experienced sustained congestion, occurring throughout the day and frequently reaching rejection percentages above 50%.

The asynchronous nature of congestion across providers implies that users can potentially achieve improved connectivity and performance by switching points of attachment when the BTS or carrier on which they are connected experiences congestion.

## 5 SMARTCELL

Informed by our analysis, we propose SmartCell, an Android application that helps users avoid congested cellular base stations. SmartCell observes GSM control channels, allowing it to passively detect when the BTS serving a user is congested. Essentially, SmartCell exposes the SDCCH blocking KPI in near real-time to end-users; such metrics are typically only known to the cellular carriers themselves.

5. [http://www.etsi.org/deliver/etsi\\_gts/02/0267/05.00.01\\_60/gsmnts\\_0267v050001p.pdf](http://www.etsi.org/deliver/etsi_gts/02/0267/05.00.01_60/gsmnts_0267v050001p.pdf)

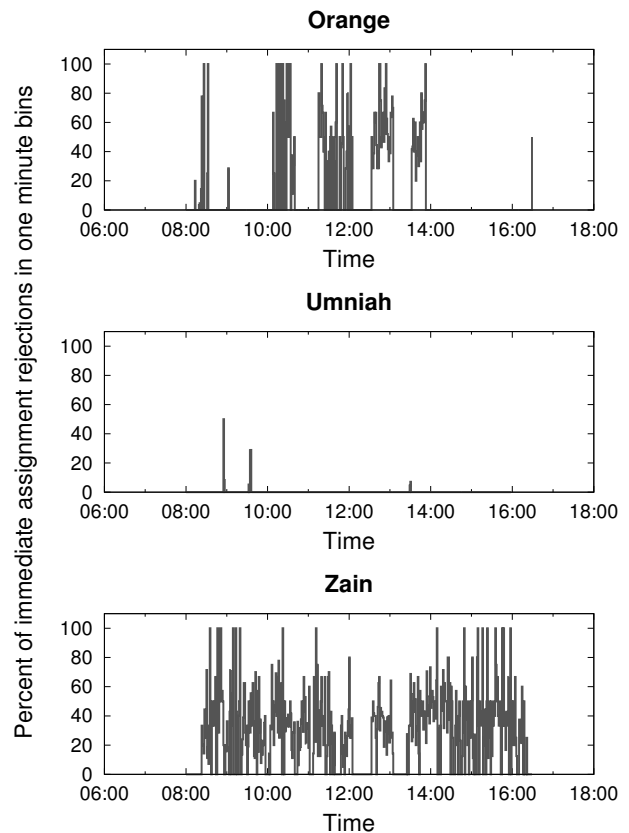


Fig. 2: Percent of immediate assignment messages that were rejections in one minute bins on January 6<sup>th</sup> 2015 in Za’atari.

The operation of SmartCell depends on a user’s preferences and the features of their mobile device. SmartCell periodically logs a timestamp, mobile network code, cellular base station ID, approximate location, and a BTS Availability (BTSA) value based on observed cellular network congestion. If permitted by the user, SmartCell may share these measurements with other SmartCell users, by periodically uploading this information to SmartCell’s servers. By combining measurements from multiple nearby users, SmartCell enables users to make an informed choice from among their possible providers, even if their phones do not support the diagnostic interface that permits SmartCell to directly observe the GSM air interface.

For users with multi-SIM phones, SmartCell may operate automatically, relocating users away from congested base stations using the phone’s observations and observations collected by other users nearby. SmartCell can select which SIM card should be the phone’s ‘primary’ SIM, which is used for outbound calling and mobile data. If a user’s phone only supports a single SIM card, SmartCell will notify the user if it determines that they might benefit by switching mobile networks or cells. This determination begins with SmartCell observing congestion on the user’s current network above a threshold.

As congestion approaches this threshold, SmartCell will provide context to the measurements by querying its servers to determine if a user is observing worse congestion than

nearby users on other cells. If congestion on the user’s current cell is above average for their area, SmartCell alerts the user which carriers offer superior service in their area, and of better performing cells of their own carrier. This may prompt users to switch SIMs to use a less congested carrier, or to relocate in order to switch to a less congested base station of the same carrier.

### 5.1 Hardware platform

During our visit, our team surveyed 228 residents of the Za’atari refugee camp. We found that Android phones are very popular in Za’atari: 64% reported owning an Android device, 22.4% owned a Nokia device, while 4% owned an iPhone. We believe the universal availability of Android handsets (e.g. 85% global market share), powerful features and APIs, and wide range of models at varying prices, make Android an ideal platform for SmartCell.

SmartCell relies on a diagnostic interface of the phone’s cellular baseband to collect the cellular broadcast messages used to detect congestion. This data collection technique has been used to detect baseband attacks, and is supported on a wide range of Android handsets with popular baseband chipsets [1].

### 5.2 Use of multiple networks

SmartCell relies on the use of multiple SIM cards, which is common in developing regions. Unlike areas where ‘good’ cellular connectivity is ubiquitous, carriers in rural and developing regions often have vastly disparate coverage and quality of service. In contrast to the multi-year cellular contracts common in countries like the US and Japan, the availability of low cost and contract-free prepaid SIMs enables users in developing regions to use multiple cellular networks. This leads to users carrying SIM cards from multiple carriers and switching as needed to obtain acceptable connectivity. For example, the number of cellular subscriptions per 100 residents in Guatemala is 106.6, and in Jordan it is 147.8, while in the United States it is 98.4<sup>6</sup>. Respondents to the survey administered in Za’atari use three SIM cards on average, switching SIMs to take advantage of less congested networks, ‘same network’ discounts, and cheaper data-only plans. This level of comfort in switching carriers is promising for a system such as SmartCell.

Switching SIMs traditionally requires powering off the phone. However, phones that support multiple SIM cards are increasingly popular. OpenSignal reports 57% of Android users in Guatemala have multi-SIM phones, and more than 50% of Android users in several other developing countries own dual-SIM phones<sup>7</sup>. When SmartCell is used on a multi-SIM phone, physically swapping SIM cards is unnecessary. Instead, SmartCell can take the user directly to the SIM settings activity, allowing the user to select which SIM to use for voice calls, SMS messaging, and data traffic. Alternatively, SmartCell can reconfigure the phone automatically, attempting to select the least congested mobile network in the area.

6. <http://www.itu.int/en/ITU-D/Statistics/Documents/publications/misr2015/MISR2015-w5.pdf>

7. <http://opensignal.com/reports/2015/01/android-devices-dual-sim/>




Fig. 3: BTS Availability (BTSA) for carriers in Za’atari on the morning of January 6<sup>th</sup> 2015.

### 5.3 SmartCell operation

SmartCell observes cellular control messages, specifically immediate assignment and immediate assignment rejection messages, to understand network availability and load on a user’s current BTS. Recall that these messages are BTS-specific and allow an attached MS to infer SDCCH blocking affecting not only itself, but also other MSs connected to the BTS. SmartCell detects these messages and uses them to compute a *BTS Availability* (BTSA) metric. When this metric surpasses a threshold, SmartCell notifies the user that the network is congested or automatically switches the priority of a user’s SIM cards to an alternate network. We use 0.9 as the threshold based on doubling the classification of 5% blocking as high congestion [9].

#### 5.3.1 Availability algorithm

SmartCell estimates BTSA periodically based on an exponential weighted moving average of immediate assignment success rate, defined in Section 3.1, as shown in Equation 1. Recall that we infer SDCCH blocking from the immediate assignment success rate.

$$BTSA_t = A_t \times \alpha + (1 - \alpha) \times BTSA_{t-1} \quad (1)$$

Using our algorithm, we plot the BTSA for each of the carriers serving Za’atari on the morning of January 6<sup>th</sup> 2015 in Figure 3. We use a period of one minute and an  $\alpha$  of 0.25. Note that  $\alpha$ , the periodicity, and the threshold are tunable values as system responsiveness requirements may differ depending on the environment. We see that the carriers have distinct availabilities that roughly correspond to the immediate assignment rejection percentage plotted in Figure 2. Importantly, congestion on one carrier does not necessarily correlate with congestion on another. If the value crosses below the threshold (indicated by the red line in Figure 3), SmartCell informs the user that their current BTS is congested, prompting the user to switch BTSs or networks, or, in the case of a multi-SIM phone, switches to a different network automatically.

An example scenario of SmartCell’s operation can be observed in Figure 3. A Zain customer running our application would cross the availability threshold at around 8:30 (shaded region) and be advised they may wish to switch SIMs or move to a different location. At that time, Orange also experienced congestion, though to a lesser degree, while Umniah remained congestion-free. The Zain user could switch to either network and obtain a BTS with higher availability, and hence receive better quality of service.

### 5.3.2 Reachability

The use of multiple SIM cards negatively impacts user reachability, which is the ability for a correspondent to connect with a user via one of the user’s phone numbers. Issues with reachability manifest as delays in correspondence, as calls and SMSs sent to an inactive SIM will not be received until the user inserts or re-activates the inactive SIM. SmartCell may exacerbate or improve this problem, depending on usage and the cellular networks in the area. In cases like Za’atari, where one carrier has significantly lower availability than others, users of SmartCell may visit the worst performing network less often, decreasing their reachability via their phone number from that network. However, as widespread use of SmartCell would likely result in loose load balancing across networks, this decrease in reachability is likely to be short-lived.

By contrast, in more balanced network environments, SmartCell may encourage users to switch between networks more frequently than they would have otherwise, reducing the average time their SIMs remain inactive and therefore reducing the delay before a user receives a missed call notification or SMS. We do not believe that attempting to contact a user of SmartCell will be drastically different than contacting non-users, as users commonly switch between cellular networks currently. In contrast, SmartCell should *improve* reachability when a user is receiving a call on an active SIM, as SmartCell users prefer base stations with higher availability, which are more likely to have the capacity to provide service for the incoming call.

### 5.3.3 Global congestion

In the case of global congestion in an area, where all carriers are severely congested, SmartCell may cause users to switch between carriers when the switch will not actually improve connectivity. However, assuming a large population of users, SmartCell should lead to homeostasis where all carriers in an area have roughly the same availability, rather than drastically different levels as witnessed in Za’atari.

In cases of severe congestion on a single carrier, shifting many users onto a different carrier network may lead to simply migrating congestion from the egress network to the ingress network. As such, SmartCell can be tuned to avoid rapid, drastic changes in estimated availability by increasing the weight assigned to past measurements and/or setting the threshold at a lower value. By increasing the weight assigned to past measurements the availability metric is smoothed.

## 6 DISCUSSION AND CONCLUSION

SmartCell opens multiple possibilities for further research. A current limitation of SmartCell is that each phone has

limited a-priori information about whether selecting alternate carriers or BTSs will result in improved service. Each phone with SmartCell relies on its own measurements of the user’s current cell, and recent measurements from nearby SmartCell users. This presents a challenge in areas with a low density of SmartCell users. Unless other nearby users have recently reported congestion information for multiple nearby carriers and cells, the selected network could be as congested or worse post-switch. An algorithm to *predict* availability for carriers and BTSs based on past data submitted to the SmartCell repository is one area for future improvement. The repository could serve tiles to clients in a similar manner to map applications, where each tile contains long-term congestion history of BTSs in an area. This would reveal long-term trends, give users a better intuition of which carriers and cells are likely to be congested at any given time, and enable a more informed choice of networks and cells in the absence of real-time measurements.

The use of a central repository in congested environments can also be a limitation. If poor connectivity precludes the use of a central repository, cooperative peer measurements could be shared using direct channels such as local WiFi networks, WiFi direct, or Bluetooth. As each measurement is only a few bytes of information, a measurement can be encoded into short strings. For example, the string “1458049642:410:54012:34.4312:-119.7598:.8” indicates that at epoch time 1458049642, a cell on ATT (MNC 410) with cell ID 54012 was observed near 34.4312:-119.7598 with availability .8. These strings can be shared in a connectionless fashion, such as in Wi-Fi SSIDs or in the UUID field of Bluetooth Low Energy beacons.

Although cellular coverage is rapidly increasing throughout developing regions, it is clear that disparities in the quality of the coverage exist, and will likely continue to exist for the foreseeable future due to a variety of economic and geographic factors. Our work scratches the surface of understanding these inequities in two specific, very different regions. Clearly, more work remains to be done to fully characterize the problem.

## 7 ACKNOWLEDGEMENTS

This work was funded through NSF Network Science and Engineering (NetSE) Award CNS-1064821 and NSF Catalyzing New International Collaborations (CNIC) Award IIA-1427873.

## REFERENCES

- [1] N. Vallina-Rodriguez, A. Auçinas, M. Almeida, Y. Grunenberger, K. Papagiannaki, and J. Crowcroft, “RILAnalyzer: A comprehensive 3G monitor on your phone,” in *IMC ’13*, (Barcelona, Spain), 2013.
- [2] S. Rosen, H. Luo, Q. A. Chen, Z. M. Mao, J. Hui, A. Drake, and K. Lau, “Discovering fine-grained RRC state dynamics and performance impacts in cellular networks,” in *MobiCom ’14*, (Maui, Hawaii, USA), 2014.
- [3] F. Lu, H. Du, A. Jain, G. M. Voelker, A. C. Snoeren, and A. Terzis, “CQIC: Revisiting cross-layer congestion control for cellular networks,” in *HotMobile ’15*, (Santa Fe, NM), 2015.
- [4] D. F. Kune, J. Köndorfer, N. Hopper, and Y. Kim, “Location leaks over the GSM air interface,” in *NDSS*, (San Diego, CA), 2012.
- [5] A. Achtzehn, J. Riihijärvi, I. A. Barriáa Castillo, M. Petrova, and P. Mähönen, “CrowdREM: Harnessing the power of the mobile crowd for flexible wireless network monitoring,” in *HotMobile ’15*, (Santa Fe, NM), 2015.

- [6] E. Gustafsson and A. Jonsson, "Always best connected," *IEEE Wireless Communications*, vol. 10, pp. 49–55, Feb 2003.
- [7] S. Kyriazakos, G. Karetsos, E. Gkroustiotis, C. Kechagias, and P. Fournogerakis, "Congestion study and resource management in cellular networks of present and future generation," in *10th IST Mobile Summit*, (Barcelona, Spain), 2001.
- [8] B. Haider, M. Zafrullah, and M. Islam, "Radio frequency optimization & QoS evaluation in operational GSM network," in *World Congress on Engineering and Computer Science*, (San Francisco, CA), October 2009.
- [9] S. Kyriazakos, N. Papaoulakis, D. Nikitopoulos, E. Gkroustiotis, C. Kechagias, C. Karambalis, and G. Karetsos, "A comprehensive study and performance evaluation of operational GSM and GPRS systems under varying traffic conditions," in *11th IST Mobile Summit*, (Thessaloniki, Greece), June 2002.