HybridCell: Cellular connectivity on the fringes with demand-driven local cells

Paul Schmitt*, Daniel Iland*, Mariya Zheleva† and Elizabeth Belding*
*University of California, Santa Barbara
Email: {pschmitt, iland, ebelding}@cs.ucsb.edu
†University at Albany, SUNY
Email: mzheleva@albany.edu

***Abstract*—While cellular networks connect over 3.7 billion people worldwide, their availability and quality is not uniform across regions. Under-provisioned and overloaded networks lead to poor network performance and an aggravated user experience. To address this problem we propose HybridCell: a system that leverages locally-owned small-scale cellular networks to augment the operation of overloaded commercial networks. HybridCell is the first system to allow a user with their existing SIM card and mobile phone to seamlessly switch between commercial and local networks in order to maintain continuous connectivity. HybridCell accomplishes this by identifying poorly-performing networks and taking action to provide seamless cellular connectivity to end users. Using traces collected from observing the cellular infrastructure during our visit to the Za'atari refugee camp in Jordan, we demonstrate HybridCell's capability to detect and act upon commercial network overload, offering an alternate communication channel during times of congestion. We show that even in scenarios where provider networks deny calls due to overload, HybridCell is able to accommodate users and facilitate local calling.**

I. Introduction

Millions of people throughout the world live in areas at the fringes of cellular connectivity, as the cost of increasing local infrastructure exceeds the expected return on investment. In India, for example, roughly 25% of the population (nearly 315M [1]) resides in areas without cellular coverage. Coverage that does exist in fringe areas is spotty and often overburdened. Residents typically carry multiple SIM cards in order to obtain connectivity in locations where one provider is present while another is not, and the inability to obtain any service during busy times of day is common. Cellular networks in these areas are simply unable to service the demand placed on them, yet upgrading the infrastructure is infeasible due to cost. Further, despite continued expansion of commercial networks, hundreds of millions of people reside in areas without coverage. Clearly, ubiquitous coverage will not be achieved without exploring the use of more cost effective technologies.

Rural areas are just one example where infrastructure capacity is unable to meet the demand. When disasters strike, people often move to makeshift camps located in areas with available space, where existing cellular infrastructure is not provisioned for the increase in user load. Likewise, political conflicts lead to displacement of people to refugee camps. Such camps are often located in rural areas, on the fringes of infrastructure, in order to reduce disruption of the residents of the host territory. Even infrastructure in urban areas can experience sudden spikes in demand (e.g. 100,000 people gathered for a protest), where the expense of building the infrastructure necessary to meet peak demand makes little financial sense.

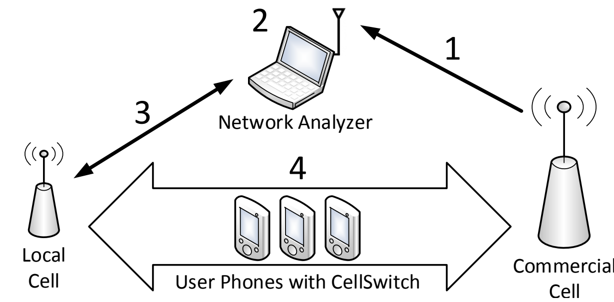

Fig. 1: HybridCell consists of a Network Analyzer, a Local Cell and our network switching application CellSwitch. (1) The Network Analyzer gathers performance data from the commercial cell; (2) the Network Analyzer algorithm determines commercial cell congestion and failure; (3) the Local Cell is reconfigured based on the Analyzer's measurements; and (4) CellSwitch shifts users between Local and Commercial cell for seamless connectivity.

Recently, local cellular networks have gained traction as a solution for providing cellular connectivity in remote areas lacking coverage [2], [3], [4], [5]. These installations use small-scale base stations running open source software such as OpenBTS [6], which transitions GSM to voice over IP (VoIP), and offer free or low-cost cellular services. Unlike femtocell technology [7], local cellular networks can exist autonomously and do not require a reliable connection to a commercial provider in order to operate. When a reliable connection to a provider is unavailable, a local cellular network can still provide voice and SMS capability between users connected within that local network. The question arises, can local cellular networks be leveraged to alleviate and bolster poor commercial connectivity in areas where providers are either unable or unwilling to augment their infrastructure? Prior work on local cellular networks has focused on areas with *no existing coverage*. In contrast, we use local cellular technology to create HybridCell (Figure 1), a self-contained, community-owned and operated cellular network that intelligently *augments* *and coexists* with existing commercial coverage. HybridCell provides supplemental capacity when commercial networks become overloaded or non-operational by shifting a portion of local calls and SMS messages between HybridCell users to the local network. We design HybridCell to improve connectivity in areas with existing, yet under-performing cellular coverage rather than areas devoid of coverage.

The design of HybridCell raises several research challenges. For instance, how can we determine when a cellular network is overloaded without having access to that network, and without adding any traffic load? How can we define congestion given that base stations can have various capacities? Can we create a solution that respects spectrum occupancy and does not cause interference?

HybridCell solves these challenges through three key components: (i) a commercial network analyzer, (ii) small local cells and (iii) an Android application dubbed CellSwitch. HybridCell's analyzer performs non-intrusive, passive measurements to quantify the level of congestion in the commercial network and thus determine the performance and health of the commercial network. When the commercial network is fully functional, HybridCell operates quietly in the background, simply monitoring the commercial network load. As congestion and failures increase, HybridCell adaptively shifts clients from the commercial to the local network in order to decrease the load on the commercial network. The CellSwitch application controls each user's network association programmatically, removing the need for users to explicitly choose networks via physically changing SIM cards or reconfiguring network association settings. HybridCell works alongside any commercial carrier, in contrast to the connectivity model used by Google Fi [8], as described in section IV.

To more specifically put the HybridCell system into context, we examine the plight of Syrian refugees and how HybridCell could help improve cellular access. The UNHCR estimates that the conflict in Syria has caused over 4 million Syrians to leave the country as of July 15, 2015 [9]. The influx of refugees has led to the establishment of roughly 30 refugee camps in neighboring countries. The Za'atari refugee camp, located in Jordan, has become one of the largest refugee camps in the world. We visited Za'atari in January 2015 to evaluate current cellular network coverage and usage within the camp. Like many refugee camps, Za'atari quickly sprung into existence, forming virtually overnight. Within 9 months of its July 28, 2012 establishment, the population of the camp had increased to over 200,000 people [9]. As a result, infrastructure struggled to keep up with the rapidly growing population. We captured measurements of the existing cellular infrastructure serving the camp and conducted surveys and interviews of camp residents as well as administrative staff. In section III, we use our measurement data to both convey the dire, present need for a system such as HybridCell, as well as to evaluate the potential for HybridCell in such an environment.

This paper makes several key contributions:

- We design a first of its kind method for passively quantifying cellular congestion as third party observers without requiring access to cellular network core network traffic.
- We use our congestion detection mechanism to assess network load and performance in a real-world multi-carrier environment in Jordan's Za'atari refugee camp.
- We implement an Android application, CellSwitch, that programmatically shifts phones between cellular networks, allowing phones to use both local and commercial cellular networks without manual intervention.
- We integrate our analyzer and CellSwitch into HybridCell; the first system to combine commercial and local cellular networks for seamless user connectivity.

II. HybridCell System Design

HybridCell consists of three components as illustrated in Figure 1: (i) a network analyzer that detects active commercial cellular networks and characterizes their performance; (ii) a local cell that augments cellular services in the face of a failing commercial network; and (iii) the CellSwitch application that resides on users' phones and transparently migrates users between the commercial and the local network in order to assure seamless connectivity. In this section we provide a detailed description of HybridCell's components and operation.

A. Commercial Network Analyzer

Before taking any action, HybridCell first characterizes the performance of the nearby commercial network using a network analyzer, which operates in two phases. The first *detects* available carriers and the second *characterizes* the performance of these carriers. The detection phase identifies all the Absolute Radio Frequency Channel Numbers (ARFCNs), i.e. the operating frequencies of individual cells, along with the technology they use (i.e. GSM, 3G or LTE). The characterization phase then taps into the control channels of each carrier and identifies the health status of the network based on the control messages.

1) Commercial carrier detection: HybridCell begins by detecting all active carriers in its vicinity. This detection is necessary for two reasons. First, it is needed in order to identify available ARFCNs that can be used by the local cell without interfering with the commercial carriers. Second, the detection determines the technologies (GSM, 3G or LTE) used by the commercial carriers, which in turn informs the network characterization. For the purpose of detection we make use of a *blind service identification*, which pinpoints all the base stations (including technology and ARFCN) nearby that can be augmented by HybridCell should they become overloaded.

Three commonly proposed spectrum sensing methods for blind service identification are energy detection, matched filter detection, and feature detection [10]. Energy detection cannot discern the underlying technology behind a signal and matched filter requires a priori knowledge of the signal for which we are searching. Hence they are not good choices for our solution. On the other hand, feature detection is able to exploit the known periodicity (i.e. cyclostationary characteristics) of the target signals. The pilot signals present in cellular standards result in our ability to use a single detection scheme for

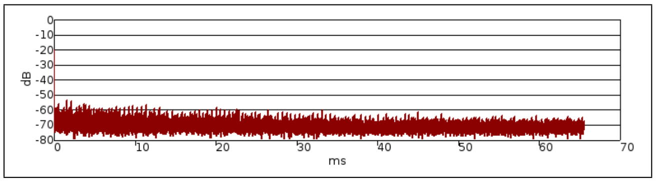

(a) Gaussian noise

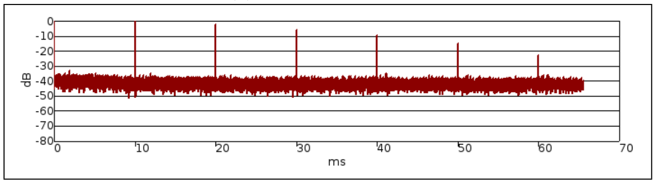

(b) 3G WCDMA

Fig. 2: Peaks with fixed periodicity are observed in the cyclic autocorrelation function, corresponding to timeslots used by cellular synchronization channels.

GSM, 3G, and 4G LTE. Furthermore, cyclostationary detection is more robust in noisy environments than simple energy detection, making it ideally suited for use in HybridCell.

We implement our detector using an Ettus Research USRP2 and GnuRadio to generate the cyclic autocorrelation function (CAF) values for each signal type. Figure 2 shows the CAF for Gaussian noise and a 3G WCDMA signal. Peaks in the figures mean that the observed signal exhibits periodicity at the corresponding x-axis (time) value. The Gaussian noise plot (2a) demonstrates the case of a channel with no service provided. As expected, we do not see any clear peaks in the figure as there is no periodicity. In contrast, as shown in Figure (2b), the CAF of a 3G signal includes clear peaks spaced at 10ms, corresponding to the frame length for 3G WCDMA. We observe similar periodicity for GSM and 4G LTE signals. These are not included for brevity.

2) Commercial carrier characterization: We base the characterization of commercial carrier performance on messages exchanged on the Broadcast Control Channel and Common Control Channel. In this section we first give a brief overview of the messages. We then detail how we use these messages to characterize active carriers' performance.

Control channel overview. Cellular phones or software defined radios can be used to collect System Information Messages broadcast by nearby cellular networks. This message collection is non-intrusive and non-invasive: all captured messages are broadcast by base stations in plain text on control channels and are intended to be received and processed by all phones associated with a given base station.

HybridCell makes use of the messages listed in Table I. System Information Messages types 1-4 are broadcast on the Broadcast Control Channel (BCCH) and are akin to Wi-Fi beacons, providing basic information about the base station configuration and the services it offers. The messages include frequencies occupied by the base stations as well as those of neighboring cells, access classes to notify the handset of the level of service available, and the mobile network and country codes. HybridCell uses these messages to identify and

Broadcast Message	**Information used by HybridCell**
System Information Type 1	Cell Channel Description: List of ARFCNs Band Indicator Access Control Classes (ACC)
System Information Type 2	Neighbor cell descriptions: List of ARFCNs Access Control Class (ACC)
System Information Type 3 and 4	Mobile Country Code (MCC) Mobile Network Code (MNC) Location Area Code (LAC) Cell Identity (CI) Access Control Class (ACC)
Immediate Assignment	Timeslot Single channel ARFCN
Immediate Assignment Reject	Wait Indication Request Reference

TABLE I: System messages used by HybridCell.

characterize the nearby base stations. Immediate assignment messages, sent on the Common Control Channel (CCCH), indicate whether the base station is able to reserve resources for a mobile device (e.g. assignment of a voice traffic channel for a call). We leverage these messages to estimate the health of the observed base station by recording the number of successful immediate assignments as well as rejection messages.

Congested and failed cell detection. A critical function of the system is the ability to characterize service availability and quality of service of nearby commercial base stations through passive observation. To this end, HybridCell uses observations of channel assignment broadcasts to determine the quality and reliability of provided cellular service. We estimate the operational state of commercial networks and configure local cells accordingly, shifting traffic to local cells when commercial networks are congested.

Several messages are indicative of network overload. For example, overloaded base stations may attempt to reduce load on the system by preventing users from using the network, by barring one or more *access classes*. As shown in Table I, access class settings are broadcast in system information messages. Each of the bits in the Access Control Class (ACC) field of a System Information message represents a class of users. Base stations can block one or more Access Classes from connecting by modifying the ACC value broadcast in System Information Type 2 and Type 3 messages [11]. When HybridCell detects access class restrictions, it learns that the observed cell is overloaded.

A cellular base station operating at full capacity that cannot allocate radio resources to serve a user will issue Immediate Assignment Reject messages. In the GSM 04.08 specification, Immediate Assignment Reject messages are defined to indicate that no channel is available for assignment [11]. The link between available channels and Immediate Assignment Reject messages makes this message an excellent indication that a base station is overloaded.

We use observed radio resource management messages to infer congestion. Our system estimates congestion every minute, based on a weighted moving average of channel assignment success rates. Our goal is to estimate the likelihood of a user being allocated a channel when they request one.

Specifically, we define a *Channel Availability* metric based on the ratio of observed Immediate Assignment messages and Immediate Assignment Reject messages. Let α be a weighting factor for previous measurements; let ρ be the number of Immediate Assignment Reject messages observed since the last calculation; let χ be the number of successful Immediate Assignments observed since the last calculation; and let Ψ be the estimated availability at time m which is defined as:

$$\Psi_m = 1 - ((\alpha \times \Psi_{m-1}) + \left((1-\alpha) \times \frac{\rho}{\rho + \chi}\right))$$

We analyze the Channel Availability of all three Jordanian cellular network operators by using our Za'atari traces to compute and evaluate this metric in section III.

B. Local Cells

HybridCell augments cellular coverage by moving users between the commercial network and the local network when it detects congestion or failure for users of the commercial network. This behavior eases congestion by shifting some of the load on the commercial network to the local network. In the case of commercial failure, the local network provides a functional alternate means of connectivity. We leverage existing local cellular network technologies in order to provide local cellular connectivity without requiring a reliable Internet connection, a key limitation of femtocells.

Depending on antenna location and height, local base stations with as little as 1 watt of amplification can provide coverage with a radius of a few kilometers. In densely populated areas, a number of local cells with reduced coverage would be deployed to provide additional capacity. The local network can easily be extended by interconnecting base stations using point-to-point Wi-Fi infrastructure, as voice and SMS traffic is encapsulated by standard UDP packets.

C. CellSwitch Android Application

The final component of HybridCell is CellSwitch, an Android application that allows a user's mobile phone to switch to specific nearby cellular networks without user intervention. We target Android as affordable Android-based smartphones are widely available and popular throughout the world, including, as we observed, in the Za'atari camp. Our application uses an Android system function in the telephony framework's GSMPhone class to instruct the baseband processor to register on a specific network. This function allows the application to programmatically *duty cycle* between cellular networks without requiring user involvement. Tradeoffs inherent in configuring the behavior of CellSwitch are based on balancing user reachability with power consumption.

Reachability. The main goal of HybridCell is to improve user connectivity when a nearby commercial network is unable to provide adequate service. We design HybridCell with the intention that phones should remain on commercial networks as much as possible since HybridCell is meant to be a secondary, rather than primary means of connectivity. When users are connected via their commercial network, they are

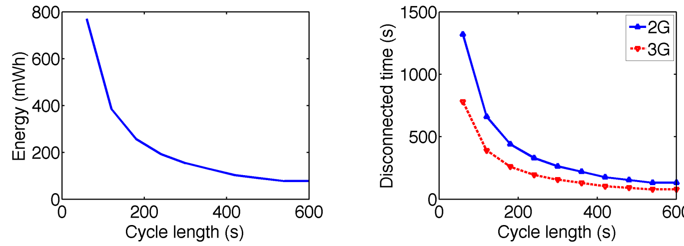

Fig. 3: Power consumption and disconnected time increase exponentially as P drops close to 15 seconds. As P increases the energy burden of CellSwitch and the time a phone spends disconnected from both networks decreases.

globally reachable and can receive incoming calls and text messages from users outside the local network. Reachability of users on the local cellular network is deployment-dependent: if the local cell does not permit incoming calls through a backhaul link using VoIP, the only way to call someone on the local cell is for the caller and callee to both be on the local cell network. In such a case, the traffic HybridCell is able to offload is limited to that between local users. Note that local cellular coverage can be extended in an area using point-to-point WiFi links between cells. HybridCell takes advantage of cellular locality of interest, as previous work has shown a large percentage of cellular traffic in rural areas is between users connected to the same base station [12].

Power consumption. HybridCell employs duty cycling to define the amount of time users spend connected to the commercial and local networks. The two defining variables for such a system are the *period*, and the *cycle*. The period defines how often the cycle repeats. The duty cycle (D) is defined as the percentage of time spent connected to the local network divided by the period (P). We consider both facets of the system with regards to system responsiveness as well as power consumption.

Duty cycle period. The length of the duty cycle period P dictates the maximum length of time between leaving a network and returning to it. That is, a handset will join both the local and commercial networks at least once in each period. User phones must connect to the local network at least briefly each period to check for queued SMS messages and to learn the most recent estimated availability in order to adjust their duty cycle percentage accordingly.

Therefore, the P value is an upper bound on call initialization and SMS delivery latency, as calls can begin when both users are on the same network and SMS messages are queued until the recipient connects to the network. An effective cycle period can be thought of as one in which the system minimizes call initialization or SMS delivery delays while balancing the constraints imposed by limited battery life offered by a phone.

We explore the relationships between P and power consumption as well as 'disconnected time' in Figure 3. We estimate the power consumption and disconnected time for a phone for one hour for different P values. Our estimates for power consumption and disconnected time are based on empirical findings on commercial carriers detailed further in section III-C. Disconnected time is the amount of time a phone spends not connected to any network and is a product of the cellular network registration process. We see in the figure that shorter P values lead to higher power consumption as well as more disconnected time.

Duty cycle percentage. HybridCell adaptively adjusts the duty cycle percentage based on estimated channel availability (Ψ) observations made by the network analyzer detailed in section II-A. For instance, a Ψ value of 0.2 will result in users spending 20% of P connected to the commercial network and 80% of P connected to the local network. As congestion on the commercial network increases, Ψ decreases. More users connected to the local network will lead to an increased proportion of calls and SMS messages between pairs of users that happen on the local network, offloading a portion of the load away from the commercial network. It is possible to overload a local network given enough users. This can be avoided by controlling the number of HybridCell users or supplementing local network capacity by increasing the number of local cells.

D. Detailed HybridCell functionality

HybridCell reduces the load on nearby overburdened commercial networks by adaptively duty cycling users onto a local network. This dynamic leads to a number of potential use cases based on to which network a caller and callee are connected. A participating phone's SMS or voice call is attempted on the network that the phone is associated with at the time of the attempt. As a result, there are four possibilities at the moment a voice / SMS event is triggered: (i) a local-local (L-L) call, where both users are connected to the local network, in which a call or SMS proceeds immediately with no delay; (ii) a commercial-commercial (C-C) call, where both users are connected to the commercial networks, and a call or SMS proceeds on the commercial network with no delay; (iii) a commercial-local (C-L) call, where the caller is on the commercial network while the callee is on the local network. The call will be placed immediately but will result in a 'miss' due to the callee's local association. An SMS will be sent and queued by the commercial carrier. The recipient will receive the queued message when they connect to the commercial carrier during their duty cycle; and (iv) a local-commercial (L-C) call, where the caller is on the local network and the callee is on the commercial network. When the call is placed the local network will place the caller on 'hold' until the callee associates with the local network. Alternatively, an SMS message sent by the local user will be queued on the local cell and will be delivered when the callee associates with the local cell in their duty cycle.

HybridCell therefore introduces call initialization and SMS latencies as well as the possibility of missed calls for its users. We are particularly interested in L-C events as they result in latencies not present in traditional cellular systems. However, the system is designed for deployment in areas where calls and SMS messages are failing due to overburdened networks. When the commercial network is operating well, the duty cycle can be set such that users spend only a few seconds on the local cell each period to gather Ψ updates and queued SMS messages, until the analyzer detects a user specified threshold of commercial congestion.

III. Evaluation

In this section, we provide an overview of the data collected during a site visit to Za'atari as well as evaluate critical components of HybridCell including: (i) network analyzer congestion detection; and (ii) CellSwitch duty cycle evaluation. Lastly, we combine in-situ measurements from Za'atari with a simulation environment to examine the HybridCell user experience. The simulation is required as an actual deployment in Za'atari is not feasible given regulatory and pragmatic constraints.

A. Data set: Za'atari refugee camp

We traveled to the Za'atari refugee camp in Jordan for several days in January 2015 as part of a multidisciplinary team studying Internet and cellular phone use within the camp. Our goal was to gather measurements to objectively quantify the cellular coverage in the camp, as prior reports have characterized the camp cellular networks as "unusably slow for most of the day" [13]. We used software-defined radios and mobile phones to record raw spectrum measurements as well as cellular broadcast messages for 2G and 3G service on all three cellular providers offering service in the camp. Our team conducted interviews and administered a survey on mobile phone and Internet use to camp residents.

1) Cellular network measurements: During the three days we spent in Za'atari, we used BladeRF [14] software defined radios to scan the 900 MHz and 1800 MHz cellular spectrum for active GSM base stations a total of 14 times. To capture commercial cellular network broadcast messages, we used Samsung Galaxy S2, Galaxy Nexus, and Galaxy S4 handsets with radio debug mode enabled. Radio debug mode logs all cellular communications to a computer via USB. We used `xgoldmon` [15], an open source tool that converts debug logs into easily parsable and human readable packet capture (pcap) files. Each phone recorded all of its own uplink traffic as well as all broadcast traffic sent by cellular base stations. Using eight Android handsets, we were able to log more than 95,000 cellular radio messages.

2) Surveys and interviews: Our team surveyed 228 residents of the refugee camp. Of the respondents, 174 were youth aged 15 to 25, while 54 were adults 26 and over; 104 respondents were female and 98 were male, with 6 failing to indicate[1]. Based on our survey, mobile phone ownership is nearly ubiquitous in the camp: 86% of youth and 98.1% of adults surveyed own a mobile phone. Android phones are very popular in Za'atari: 64% of the respondents own an Android device, 22.4% own a Nokia device, while about 4% own Apple iPhones. We believe the universal availability of Android handsets (e.g. 85% global market share [16]), and

[1] Data collection methodology and analysis recieved IRB approval.

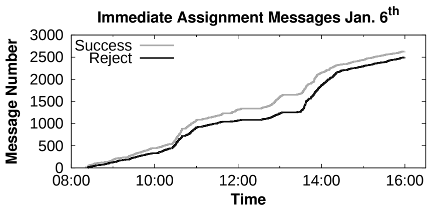

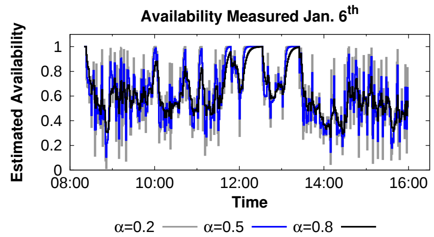

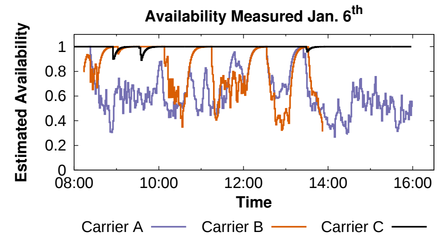

Fig. 4: Successful and failed channel assignment messages collected from Carrier A over 1 day in Za'atari.

Fig. 5: Estimated availability for Carrier A with different α values.

Fig. 6: Estimated availability for all carriers (α = 0.8).

wide range of models at varying price levels, makes Android an ideal platform for our system. One key insight gained from our survey is the utilization of multiple cellular networks and SIM cards. Respondents use 3 SIM cards on average, switching SIMs to take advantage of less congested networks, 'same network' discounts, and cheaper data-only plans.

Interviews with camp administrators and NGO staff identified the use of multiple SIMs as a key problem impeding communication with refugees. When residents switch between SIMs, no single phone number is guaranteed to reach them at any given time. Lists of contacts collected during registration are quickly out of date, and there is no guarantee the corresponding SIM will be in use when a call is placed or SMS is sent. Because of this, NGOs and UNHCR currently supplement SMS broadcasts with megaphone announcements and community outreach. HybridCell can rectify this problem by providing a unified platform for rapid and reliable information dispersal as all users connect to the local network automatically, where queued messages can then be delivered.

B. Network analyzer congestion detection

We evaluate HybridCell's network analyzer algorithm by computing the estimated availability (Ψ) for each of the three Jordanian cellular carriers using the data we collected in Za'atari. First, we explore how the selection of α values impacts our congestion metric using data from Carrier A, the most congested and most popular carrier. Then, we compare the results of our congestion detection algorithm for all three carriers using the same α value for each carrier. This demonstrates the efficacy of our metric in differentiating between heavily congested, sometimes congested, and rarely congested networks.

Figure 4 shows observed Immediate Assignment messages and Immediate Assignment Rejects collected over the course of one day for Carrier A. We see different rates of rejection messages throughout the day, loosely corresponding to a workday schedule (e.g. roughly 10:30, shortly before 14:00). This result points out the need for flexible, adaptive offloading as congestion is not constant. We calculate Ψ over time for the carrier using a range of α values, from 0.2-0.8. Figure 5 shows the computed Ψ over the course of one day for Carrier A with each α value. As expected, smaller α values result in higher estimated availability variance as less weight is assigned to history. Our metric succeeds in detecting congestion corresponding to the timing for messages seen in Figure 4 as we see sustained low values around 10:00 and 14:00. Higher α values, such as 0.8, do not impact the detectability of congestion events, but reduce the variance of Ψ. A potential drawback of using a small α values, thus resulting in rapid fluctuations, is that Ψ could increase rapidly after a congestion event ends, causing HybridCell to quickly shift many users to the recently congested network, potentially causing further congestion. This would result in unnecessary and frequent migration between networks, with recent estimations no longer accurately detecting congestion. Smoothing out large fluctuations in the availability metric is desirable for HybridCell, as channel availability is used to configure duty cycling. From these measurements, we determine that our metric for Estimated Availability with α = 0.8 is satisfactory in balancing responsiveness with congestion detection.

Using α = 0.8 our algorithms detected several periods of congestion for each carrier. Carrier A suffered from the most congestion. This follows logically from our survey data, which indicated that Carrier A is by far the most popular carrier in the camp. Figure 4 shows that of five thousand channel assignment messages observed on Carrier A, successful channel assignments barely outnumbered rejected channel assignments collected during the same time period. As shown in Figure 6, our congestion metric identified continuous congestion throughout the course of one day in the camp for Carrier A. Carrier B suffered from both ephemeral and prolonged congestion events, but generally recovered between incidents. Figure 6 shows that Carrier B's Ψ fluctuates throughout the observation, again corresponding roughly to busy hours impacted by the workday, dropping as low as 0.4, but typically returning to 1.0 within an hour. Carrier C is the least congested network based on the frequency of Immediate Assignment Reject messages in our traces, which may be due to being the least popular carrier according to our survey. Over our entire collection period in Za'atari, Immediate Assignment Reject messages totaled 5.9% of channel assignment messages for Carrier C, compared to 33.7% for Carrier A and 15.2% for Carrier B. As shown in Figure 6, Channel Availability is near 1 for Carrier C throughout the test period, and never drops below 0.8. A HybridCell user of Carrier C therefore would spend the least time on the local cell, often only the minimum time required to receive updates and queued SMS messages.

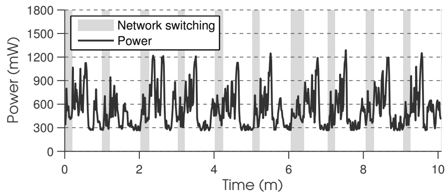

Fig. 7: Energy consumption increases during transitions, but peaks immediately following movement between cellular networks.

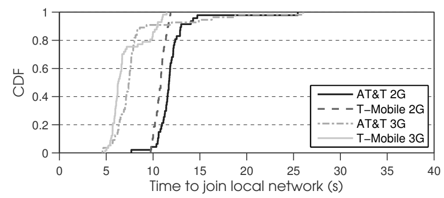

Fig. 8: Time spent without service when moving to the local network from commercial networks.

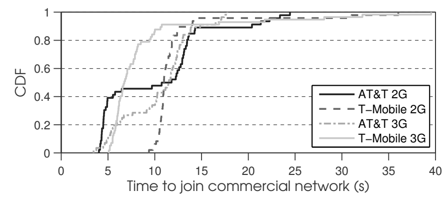

Fig. 9: Time spent without service when moving to commercial networks from a local cell.

C. CellSwitch duty cycle evaluation

Power consumption and disconnected time are both important factors when configuring the duty cycle functionality of CellSwitch. To characterize the additional power burden imposed by switching networks, and to determine the time a phone takes to move between networks, we perform several experiments. We configure three different models of Samsung Android phones to programmatically switch between a commercial network and a local network once per minute. We measure the power consumed by the phone each second during the test using an in-line USB current and voltage monitor based on the Texas Instruments INA219 DC current sensor. Our test devices were fully charged before testing to minimize the impact of charging the battery on power consumption. We create an Android test driver application that triggers transitions between networks. Our application registers a PhoneStateListener, allowing it to receive and log notifications when the phone goes in and out of service. Our test driver logs the current time in milliseconds when transitions are triggered and completed. The phones are synced to the same Network Time Protocol server as the computer recording power measurements to allow for accurate alignment of power logs with user transition logs. We force phones to connect to 2G and 3G commercial networks, allowing us to examine the impact of commercial cell technology on transition time.

1) Impact of network switching on power consumption: We first explore how switching networks affects phone power consumption. As shown in Figure 7, our experiments measured an increase in power consumption during the transition between networks. The power consumed during a migration is roughly double the idle power consumption, but, interestingly, consumption remains high for up to 30 seconds after joining a network, with peaks at almost four times idle power consumption. This result echoes prior findings that mobile radios remain in high power states for some time after usage [17], [18] in an effort to avoid the latency penalty incurred when forced to transition from idle to active. Such power behavior must be taken into account when configuring client duty cycling in HybridCell.

Duty cycle period length directly impacts power consumption of mobile handsets. As we previously saw in Figure 3, power consumption decreases as cycle length increases. A short P value (less than 200 seconds) would reduce the overall battery life of a device beyond what is likely to be considered acceptable by many users. As the period increases above 200 seconds, the steepness of the power consumption curve decreases, and variations in P length have less impact on device battery life. Based on our observations, we believe periods greater than 200 seconds are required to mitigate power consumption concerns.

2) Disconnected time: During the transition between networks, a phone is disconnected from both networks; hence the frequency, determined by the duty cycle period, of shifting between networks must be kept relatively low.

To understand how the cellular technology used by commercial carriers and 'direction' of transition impacts disconnected time, we perform migration experiments in both ingress and egress directions, with phones configured to prefer 2G or 3G networks. Figure 8 shows transition times for phones moving from the two major U.S. commercial GSM networks to a local GSM network. We observe distinct performance differences between 2G and 3G devices, with 3G transition times roughly half of 2G transition times. We also explore times for phones transitioning in the opposite direction, from the local network to the commercial network. Figure 9 shows a cumulative distribution function (CDF) of time required to join the commercial networks from the local network. We observe slightly improved best case performance compared to the ingress case in Figure 8. Interestingly, for AT&T 2G we observe two distinct groupings at 4-5 and 11-12 seconds. We posit that this could indicate the connection attempt was delayed or experienced retransmissions in the longer cases. Overall, we see transition times are typically between 4 and 7 seconds for 3G users, while 2G users most often spend approximately 10 to 12 seconds disconnected during a transition.

These results must inform the chosen duty cycle period P, as shorter lengths will result in a larger percentage of time spent disconnected. An extremely short P (e.g. 60 seconds) will cause a user to be unreachable for up to 120 4-12 second intervals over the course of an hour, or up to 33% of the time. As P increases, disconnected time decreases. For a 3G device that averages a 5.5 second transition time, P=120 seconds results in 9.16% disconnected time, while P=200 seconds results in 5.5% disconnected time and P=300s decreases disconnected time to 3.66%. Combined with our

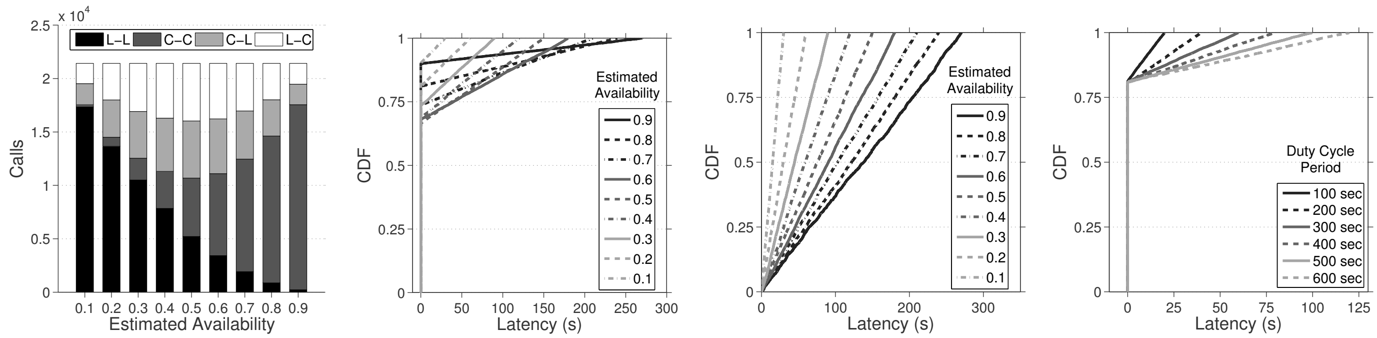

(a) Event designations with various estimated availability. As the commercial network is less available, more traffic is shifted to the local network.

(b) Latency values for non-miss events. Higher availability results in longer latencies.

(c) Latency values for L-C events. Higher availability results in longer latencies.

(d) Duty cycle period comparison for networks with 0.2 estimated availability.

Fig. 10: Simulation results.

power consumption results, these data suggest that cycle lengths of 200-300 seconds provide satisfactory balance of power consumption, call initialization delay, and disconnected time.

D. Simulation

Lastly, we evaluate the offloading behavior and expected user experience for HybridCell via simulation. We parse the cellular traces gathered in Za'atari to record timestamps for 21,426 Immediate Assignment success and reject events in response to a phone requesting a resource (e.g. to place a call or SMS). For simplicity, in our simulation we treat all events as requests for calls. Our simulation includes two pools of 500 users, callers and callees, with user duty cycle start and end times uniformly distributed across the duty cycle period P as CellSwitch timers are kept by each phone running CellSwitch and are not centrally synchronized. For simplicity, we assume that all users are HybridCell users, meaning they have access to the local network. For each request event in the trace, we randomly select a caller-callee pair and categorize the pair based on the users' associated networks at the time of the event as one of L-L, C-C, C-L, or L-C as described in section II-D.

Figure 10a shows the breakdown for simulated events with varying levels of estimated availability and a duty cycle period set to 300 seconds. As expected, increased congestion leads to increased local network usage. We also see that mixed events (L-C or C-L) are most prevalent when clients spend 50% of their time on both networks. 'Misses' account for roughly 25% of calls in the worst case; however, this case would be caused by an estimation of 50% availability on a commercial network. In other words, we would expect 50% of requests to be rejected with no intervention. HybridCell, on the other hand, will shift 50% of the burden to the local cell (L-L and L-C calls) and leave 25% on the commercial network, an overall increase of 25% for expected call success.

Figure 10b shows a CDF of latencies for all 'non-miss' events. We see that a large percentage of events experience no latency, corresponding to L-L or C-C events. L-C events cause latencies greater than zero; different congestion estimates clearly affect user latency. We single out L-C events in Figure 10c and observe that as estimated availability decreases users spend more time on the local cell, leading to shorter latencies. Lastly, we run the simulation with various duty cycle periods with a constant 0.2 estimated availability. Figure 10d shows the resulting CDF and the effect of cycle time on latencies. These results can be used to inform duty cycle configuration as each situation in which HybridCell can be deployed is unique. As noted previously, short cycle times will result in excessive power consumption. In areas where low latencies are most critical and device battery life may not be a concern, short P values can be selected. On the other hand, where latency is less important and battery life must be maximized, long P values work best.

IV. Related Work

Our work utilizes recent research on local cellular networks. OpenBTS [6] is an open-source GSM base station that has been used to provide community-scale cellular coverage in rural and underdeveloped areas. Prior works use OpenBTS to provide coverage where no commercial carriers exist [4], [5], [3]. In contrast, our focus is on using local cellular networks to *characterize* and *augment* the coverage of incumbent wireless carriers in areas where commercial coverage *does* exist, but does not provide acceptable quality of service.

HybridCell explores moving users between independent cellular providers based on observed quality of service. In April 2015 Google announced Project Fi, which moves users between T-Mobile and Sprint base stations based on expected mobile data speeds [8]. The goals of Project Fi and HybridCell are related, but the implementations are distinct due to Google's integration with two existing commercial providers. While HybridCell learns about nearby cellular networks through passive and independent observation, we expect Google has access to carrier metrics for base station performance. Additionally, HybridCell is designed to be backwards compatible with existing SIM cards and Android devices, while Project Fi requires a particular phone model with a special SIM card. The always best connected concept [19] also touches on the use of multiple networks, however it assumes business relationships exist between providers, whereas HybridCell includes a completely independent local network.

Nomadic GSM also addresses non-interfering frequency selection for base stations [3]. However, this work relies on user handsets to scan the GSM frequency range, requiring active local cell users to discover incumbents. In contrast, our system monitors incumbent control channels to determine frequencies used by incumbent carriers without relying on local user handsets. Each System Information message may reveal up to 16 frequencies in use by commercial carriers, and each System Information message we use is broadcast multiple times per second. This allows our system to more quickly identify incumbents, and to identify them before transmitting.

To provide *global reachability* while a user is on the local network, prior work detailed integration of a local cellular network with Skype [20]. This enables users of the local network to make and receive Skype audio calls from any GSM handset, and to send and receive chat messages via SMS. However, this work was focused entirely on the reliability, rapid deployability, and VoIP gateway aspects of the system and did not address incumbent detection, user migration, or utilization of multiple cellular networks.

V. Conclusion and Future Work

In this paper we design HybridCell, a system that makes use of independent local cells to augment cellular connectivity where commercial networks are overloaded or failing. HybridCell's mixture of local networks with commercial networks and leveraging of cellular communication locality creates a new connectivity model, offering service to the millions of people currently on the fringes of cellular connectivity.

Through our interviews, we know that users in these regions routinely use multiple SIM cards to enable connection to whichever network provides coverage and has capacity in the user's current location. The use of multiple SIM cards clearly can lead to missed calls and queued SMS messages. An added benefit of HybridCell is that in these scenarios, HybridCell automates a process that now occurs manually.

While HybridCell builds fundamental mechanisms for network characterization and switching, there are several questions that remain open related to the operation of the system. The current design incurs call completion latency that is manageable yet larger than that of a single-network communication. This aspect could be improved by the design of a smart duty-cycling mechanism that is also informed by social graph analysis to schedule duty cycles such that time on the local network coincides with the schedules of the user's frequent contacts. While the system currently supports SMS and voice, we are working towards adding data offloading to local cells. This will require advanced per-service network characterization and an improved suite of network switching protocols that cater to data offload. Additionally, while the current prototype includes occupied channel avoidance mechanisms to ensure that the system does not interfere with existing commercial networks, we acknowledge that HybridCell operates in licensed frequencies. We believe that with licensed shared access regulations progressing in both Europe and the U.S., a system such as HybridCell will be feasible in the near future. Given a real-world deployment, we would also be able to incorporate call data record (CDR) analysis to illuminate areas for further optimization such as intelligently scheduled duty cycles and user synchronization.

VI. Acknowledgments

This work was funded through NSF Network Science and Engineering (NetSE) Award CNS-1064821 and NSF Catalyzing New International Collaborations (CNIC) Award IIA-1427873. We thank Dr. Nijad al-Najdawi and UNHCR staff for facilitating access to Za'atari.

References

[1] "Closing the network 'coverage gaps' in Asia," https://gsmaintelligence.com/research/2015/06/closing-the-network-coverage-gaps-in-asia/, accessed: 2015-05-01.
[2] "So much going on!" http://rhizomatica.org/2015/01/14/so-much-going-on/, accessed: 2015-04-18.
[3] S. Hasan, K. Heimerl, K. Harrison, K. Ali, S. Roberts, A. Sahai, and E. Brewer, "GSM whitespaces: An opportunity for rural cellular service," in *DYSPAN 2014*, Mclean, Virginia, Apr 2014.
[4] K. Heimerl and E. Brewer, "The village base station," in *ACM DEV*, London, England, Dec 2010.
[5] M. Zheleva, A. Paul, D. L. Johnson, and E. Belding, "Kwiizya: Local Cellular Network Services in Remote Areas," in *MobiSys'13*, Taipei, Taiwan, Jun 2013.
[6] http://www.openbts.org.
[7] V. Chandrasekhar, J. Andrews, and A. Gatherer, "Femtocell networks: a survey," *IEEE Communications Magazine*, vol. 46, no. 9, Sep 2008.
[8] "FAQ - Project Fi," https://fi.google.com/about/faq/#network-and-coverage-1, accessed: 2015-05-01.
[9] "Syria Regional Refugee Response," http://data.unhcr.org/syrianrefugees/regional.php, accessed: 2015-07-01.
[10] A. Ghasemi and E. Sousa, "Spectrum sensing in cognitive radio networks: requirements, challenges and design trade-offs," *IEEE Communications Magazine*, vol. 46, no. 4, Apr 2008.
[11] 3GPP, "Mobile radio interface layer 3 specification; Radio Resource Control (RRC) protocol," TS 44.018, Sep. 2014. [Online]. Available: http://.3gpp.org/ftp/Specs/html-info/45001.htm
[12] M. Zheleva, P. Schmitt, M. Vigil, and E. Belding, "Bringing visibility to rural users in Cote D'Ivoire," in *ICTD '13*, Cape Town, South Africa, Dec 2013.
[13] M. Pizzi, "Logging on in Zaatari: Part I," http://www.smex.org/logging-on-in-zaatari-part-i/, accessed: 2015-04-25.
[14] http://www.nuand.com/blog/product/bladerf-x40/.
[15] https://github.com/2b-as/xgoldmon.
[16] "Android and iOS Squeeze the Competition," http://www.idc.com/getdoc.jsp?containerId=prUS25450615, accessed: 2015-07-13.
[17] N. Balasubramanian, A. Balasubramanian, and A. Venkataramani, "Energy consumption in mobile phones: A measurement study and implications for network applications," in *IMC*, Chicago, Illinois, Nov 2009.
[18] F. Qian, Z. Wang, A. Gerber, Z. M. Mao, S. Sen, and O. Spatscheck, "Characterizing radio resource allocation for 3g networks," in *IMC*, Melbourne, Australia, Nov 2010.
[19] E. Gustafsson and A. Jonsson, "Always best connected," *IEEE Wireless Communications*, vol. 10, no. 1, Feb 2003.
[20] D. Iland and E. Belding, "EmergeNet: Robust, Rapidly Deployable Cellular Networks," *IEEE Communications Magazine*, vol. 52, no. 12, 2014.