Third-Party Cellular Congestion Detection and Augmentation

Paul Schmitt, Daniel Iland, Mariya Zheleva, and Elizabeth Belding, *Fellow, IEEE*

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, as are common in rural or post-disaster areas, lead to poor network performance and a poor-quality 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 from commercial cellular networks collected 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.

Index Terms—Local small cells, cellular network multi-homing, network measurements, software-defined radio, displaced persons

1 Introduction

Millions of people throughout the world live in areas at the fringes of cellular connectivity, where the cost of increasing local infrastructure exceeds the expected return on investment. In India, for example, roughly 25 percent of the population (nearly 315 M [1]) resides in areas without cellular coverage. Furthermore, coverage that does exist in fringe areas is often spotty and overburdened. As a result, residents typically carry multiple SIM cards in order to obtain connectivity in locations where one provider is present while another is not. Nevertheless, 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.

Recently, local cellular networks have gained traction as a solution for providing cellular connectivity in remote areas that lack coverage [2], [3], [4], [5]. These installations use small-scale base stations running open source software such as OpenBTS [6], which translate 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* [4], [5], [8], [9]. In contrast, we use local cellular technology to create HybridCell[1] (Fig. 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

- *P. Schmitt is with Princeton University, Princeton, NJ 08544. E-mail: pschmitt@cs.princeton.edu.*
- *D. Iland and E. Belding are with the University of California, Santa Barbara, CA 93106. E-mail: {iland, ebelding}@cs.ucsb.edu.*
- *M. Zheleva is with the University at Albany, SUNY, Albany, NY 12222. E-mail: mzheleva@albany.edu.*

Manuscript received 22 Mar. 2017; revised 14 Feb. 2018; accepted 10 Apr. 2018. Date of publication 16 Apr. 2018; date of current version 3 Dec. 2018. (Corresponding author: Paul Schmitt.)
For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below.
Digital Object Identifier no. 10.1109/TMC.2018.2827031

1. This manuscript is an extension of [10].

1536-1233 © 2018 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission.
See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

Authorized licensed use limited to: Univ of Calif Santa Barbara. Downloaded on March 06,2020 at 23:36:16 UTC from IEEE Xplore. Restrictions apply.

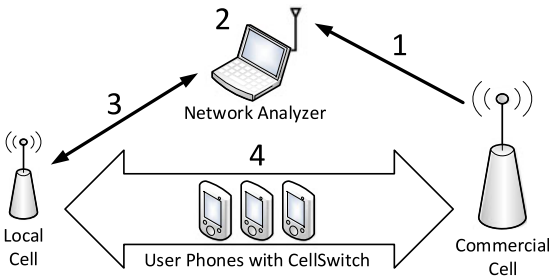

Fig. 1. HybridCell consists of a network analyzer, a local cell and our network switching application CellSwitch. (1) The network analyzer gathers performance data from the commercial cell; (2) the network analyzer algorithm determines commercial cell congestion and failure; (3) the local cell is reconfigured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

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 varying 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 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 ProjectFi [11], as described in Section 4.

To more specifically put the HybridCell system into context, we examine the communication needs of Syrian refugees, and demonstrate how HybridCell could help improve cellular access in such circumstances. The United Nations High Commissioner for Refugees (UNHCR) estimates that the conflict in Syria has caused over 4.9 million Syrians to leave the country as of March 15, 2017 [12]. 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 sprang 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 [12]. 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 3, 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 the cellular network's core 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.

The remainder of this paper is organized as follows. In Section 2 we outline the main components of HybridCell and explore tradeoffs in system design choices. We evaluate the effectiveness of each component of HybridCell in Section 3 using real-world data collected in the Za'atari refugee camp, as well as simulations and experiments. We address related work in Section 4, and conclude with a discussion of the impacts of our solution and potential future work in Section 5.

2 HybridCell System Design

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 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, then 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. Prior work, however, shows that as much as 70 percent of rural networks traffic can be local [13]. Thus HybridCell has the potential for high impact in our target areas.

HybridCell consists of three components as illustrated in Fig. 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

Authorized licensed use limited to: Univ of Calif Santa Barbara. Downloaded on March 06,2020 at 23:36:16 UTC from IEEE Xplore. Restrictions apply.

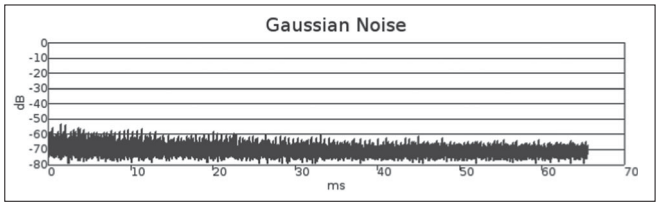

(a) Gaussian noise

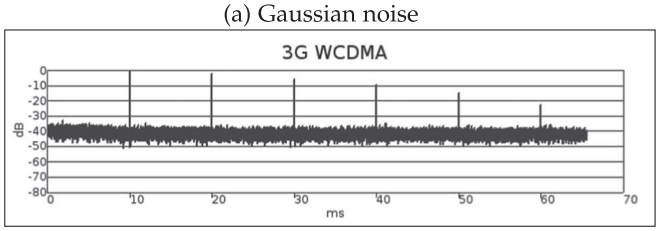

(b) 3G WCDMA

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

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.

2.1 Commercial Network Analyzer

Before taking any action, HybridCell first assesses 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. In what follows, we detail our carrier detection and characterization mechanisms.

2.1.1 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 implement a two-pronged approach: 1) *blind service identification*; and 2) *cellular-aware incumbent detection*. Our techniques provide information about all of the nearby base stations (including technology and ARFCN) that can be augmented by HybridCell should they become overloaded.

Blind Service Identification. Three commonly proposed spectrum sensing methods for blind service identification are energy detection, matched filter detection, and feature detection [14]. 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 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 blind service identification using an Ettus Research USRP2 and GnuRadio to generate the cyclic autocorrelation function (CAF) values for each signal type. Fig. 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 (Fig. 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 Fig. 2b, the CAF of a 3G signal includes clear peaks spaced at 10 ms, corresponding to the frame length for 3G WCDMA. We observe similar periodicity for GSM and 4G LTE signals. These are not included for brevity.

We do not anticipate a high churn rate in terms of neighboring base stations. Therefore, blind service identification can be run with low periodicity, reducing processing burden and allowing the software defined radio to be utilized to observe and characterize commercial cells between scans. Once the analyzer has finalized its blind server identification, it connects to the identified cells in order to collect control channel messages for further characterization as detailed below.

Cellular-Aware Incumbent Detection. Blind identification suffers from the hidden terminal problem, whereby base stations that are beyond the sensing range of the network analyzer's radio may not be detected. While those farther base stations will not be augmented by HybridCell, we must ensure that HybridCell does not interfere with their frequencies. Cellular-aware incumbent detection is thus used to determine the ARFCNs of those farther base stations. HybridCell observes System Information Types 1 and 2 (Table 1) messages broadcast in the beginning of every multiframe by nearby, detected commercial cells, and extracts *cell channel descriptions* and *neighbor cell descriptions* from those messages. Neighbor cell descriptions include ARFCNs of nearby base stations in order to enable mobility between base stations; phones constantly

TABLE 1
System Messages Used by HybridCell

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 Type 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 Wait Indication	Request Reference
Immediate Assignment Reject	

Authorized licensed use limited to: Univ of Calif Santa Barbara. Downloaded on March 06,2020 at 23:36:16 UTC from IEEE Xplore. Restrictions apply.

monitor the list of frequencies and select the base station with the highest signal strength. By combining the list of ARFCNs generated by our blind service identification scans with ARFCN lists observed in commercial base station system information messages, HybridCell avoids frequencies used by detectable commercial cellular base stations *and* all frequencies used by their neighbors. In this way, HybridCell avoids all frequencies that a carrier *announces* they are using in the local area, rather than just all frequencies HybridCell *observes* a carrier using. This mitigates the hidden terminal problem for HybridCell, making it less likely HybridCell will interfere with base stations it cannot directly detect.

By combining blind service identification and cellular-aware incumbent detection, HybridCell can utilize information from a software defined radio and commercial cellular base station broadcasts to quickly and efficiently identify incumbents. Due to the large amount of channel occupancy information periodically broadcast in GSM System Information messages, HybridCell can operate in licensed spectrum without causing interference, and without relying on a centralized spectrum occupancy database. A feasible path to deploying HybridCell is to use licensed spectrum with the permission of the incumbent(s), who would benefit from reduced congestion due to a HybridCell deployment. As spectrum policies evolve, HybridCell may be permitted to operate as a secondary user of licensed spectrum, or to operate in a 'general authorized access' tier such as the FCC has recently proposed for the 3.55-3.7 GHz range [15].

2.1.2 Commercial Carrier Characterization

We base the characterization of commercial carrier performance on messages exchanged on the Broadcast Control Channel (BCCH) and Common Control Channel (CCCH). In this section we first discuss tradeoffs of different carrier characterization approaches. We then give a brief overview of the BCCH and CCCH messages. Finally, we detail how we use these messages to characterize active carriers' performance. While carrier characterization can also be performed at the end user device, we choose a dedicated device approach (i.e., at the network analyzer) due to the specific system requirements of carrier characterization, as described in our implementation (Section 2.1.3). We note that our analyzer-based approach may lead to inaccuracies if the analyzers' view of the network is different than that of the end users' phones. For the envisioned use cases, however, we anticipate that both the analyzer and the end users will be residing in the same commercial cell, thus the analyzer-based characterization will be accurate from the end user's view point.

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 1. System Information Messages types 1-4 are broadcast on the Broadcast Control Channel and are akin to Wi-Fi beacon frames, 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 characterize the nearby base stations. Immediate assignment messages, sent on the Common Control Channel, 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 1, 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 messages [16]. When HybridCell detects access class restrictions, it learns that the observed cell is overloaded.

Another way to discern an overloaded network is through monitoring *Immediate Assignment Reject* messages. 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 [16]. 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 the exponentially weighted moving average of channel assignment success rate. 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 $\alpha \in [0, 1]$ be a weighting factor for previous measurements where higher α means that past observations will be given more weight in estimating channel availability; let χ be the number of successful Immediate Assignments observed since the last calculation; let ρ be the number of Immediate Assignment Reject messages observed since the last calculation; and let Ψ_m be the estimated availability at time m which is defined as

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

Authorized licensed use limited to: Univ of Calif Santa Barbara. Downloaded on March 06,2020 at 23:36:16 UTC from IEEE Xplore. Restrictions apply.

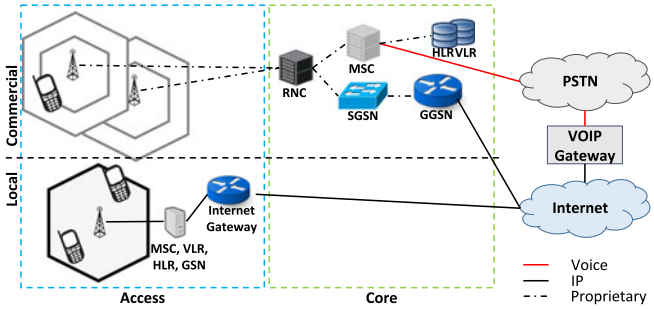

Fig. 3. Comparison of local and commercial cellular network architectures. Key advantages of local cells include their highly-distributed nature, and use of free software, open hardware, and generic backbone for interconnects.

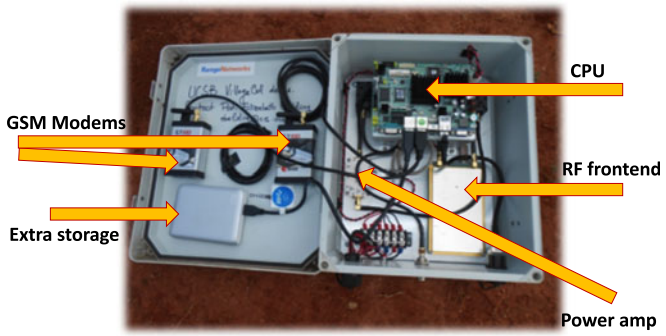

Fig. 4. Example of cellular-network-in-a-box extended with modems for active testing and additional storage, used for the authors deployment in [5].

Intuitively, Ψ_m varies between 0 and 1, where $\Psi_m = 0$ indicates a dysfunctional network, where each attempt to connect to the base station receives an Immediate Assignment Reject, while $\Psi_m = 1$ characterizes a fully-functional network with no occurring Immediate Assignment Rejects.

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 3.

2.1.3 Network Analyzer Implementation

We design the network analyzer to perform blind service identification as well as commercial carrier characterization without the use of carrier SIM cards. We use affordable software-defined radios, including Nuand BladeRFs [17]. We leverage the existing open-source tool `gr-gsm` [18] to tune the SDR to the commercial base station frequency, decode GSM control channel messages, and output the messages using the GSMTAP [19] pseudoheader format which is easily parse-able by network protocol analyzers.

2.2 Local Cells

Local cellular networks, based on open source software and open hardware, have gained traction as affordable means to provide cellular connectivity in infrastructure-challenged environments [5]. Recent deployments [8], [9] operate in areas without, or with limited existing commercial coverage in order to satisfy residents' communication needs. Key benefits of these local cells, as depicted in Fig. 3, are their highly-distributed nature, and use of free software, open hardware and generic IP backbone for interconnects. As shown in the figure, local cells push mobile core functionality to the network edge by running local software processes that perform the functions of corresponding GSM architecture entities (e.g., Mobile services Switching Center (MSC), Home Location Register (HLR), etc.), whereas commercial networks' architecture is highly-centralized and relies on proprietary backhaul links to connect the edge with such GSM entities residing in the carrier core. The commercial network backhaul requirements, along with RF frontend investments, often render commercial deployments in infrastructure-challenged environments economically-infeasible [20]. Local cellular networks, in turn, are often realized as a cellular-network-in-a-box (Fig. 4). As such, both the radio-frontend and the core components reside in the access network areas, making the design highly-distributed and more feasible for remote deployment.

HybridCell is intended for scenarios where global connectivity is unreliable due to user demand or lack of sufficient commercial infrastructure. Therefore, we leverage local cellular networks as they do not rely on always-on global connectivity. In areas with mixed local and commercial networks, HybridCell augments cellular coverage by moving users between the commercial network and the local network when it detects congestion or failure of the commercial network. This behavior eases congestion by shifting some of the commercial network load to the local network. In the case of complete failure, the local network provides an 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 Effective Isotropic Radiated Power (EIRP) can provide coverage with a radius of a few kilometers. In densely populated areas, a number of local cells with reduced coverage can 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.

2.3 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, as detailed below.

2.3.1 CellSwitch Network Switching Protocol

CellSwitch employs duty cycling to define the amount of time users spend connected to the commercial and local networks. As illustrated in Fig. 5, our application divides the

Authorized licensed use limited to: Univ of Calif Santa Barbara. Downloaded on March 06,2020 at 23:36:16 UTC from IEEE Xplore. Restrictions apply.

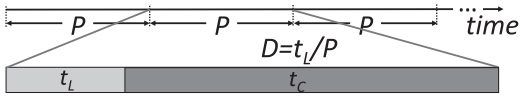

Fig. 5. Frame structure of CellSwitch.

time into equal periods P; within each period it switches between the commercial and the local cell according to its duty cycle D. The period P is equal to the sum of the *time spent on local network (t_L)* and the *time spent on commercial network (t_C)*, and defines how often the cycle repeats. The duty cycle D, in turn, is defined as the ratio of time spent on the local network over the entire period: $D = t_L/P$.

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 2.1.2. The duty cycle percentage is set using the formula: $D = 1 - \Psi$. For instance, a Ψ value of 0.2 will result in users spending 20 percent of P connected to the commercial network and 80 percent of P connected to the local network. The analyzer communicates the congestion level to the CellSwitch users via the local network. Thus, CellSwitch requires each phone to spend a minimum amount of time t_{Lmin} on the local network within each cycle for the congestion updates to be received. As the congestion updates are a single floating point value, the minimum time for transfer is a fraction of a second, even assuming relatively slow GPRS throughputs of roughly 35 Kbps. We test the time to authenticate and join a local network in Section 2.3.2. 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.

In the beginning of each period, the phone connects to the local network and later switches to the commercial network, as defined by its duty cycle. This operation requires two network switches per cycle, which incur overhead in call establishment delay and power consumption. Thus in selecting our duty cycle, we need to balance the tradeoff between network availability on one hand, and user reachability and power consumption on another. In the remainder of this section, we provide measurement-driven quantification of these two overheads and present evaluation to justify our duty cycle selection.

2.3.2 Key Tradeoffs

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 networks. 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

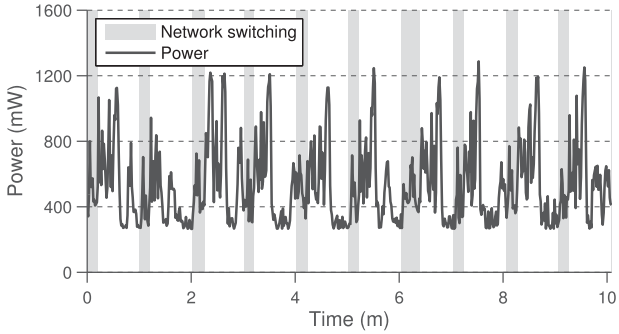

Fig. 6. Energy consumption increases during transitions, but peaks immediately following movement between cellular networks.

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.

Impact of Network Switching on Power Consumption. We first explore how switching networks affects phone power consumption. For this experiment, we configured a phone to repeatedly switch between two OpenBTS base stations running in our lab. 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 choose this experimental setup as opposed to switching across commercial networks as such behavior may appear to commercial carriers as a denial of service or resource exhaustion attack. We note that while our power consumption results are from controlled 2G OpenBTS networks, the results generalize to commercial networks and to 3G technology as all technologies utilize similar attach / detach procedures as specified by 3GPP.

As shown in Fig. 6, 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 [21], [22] 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.

Impact of Network Switching on 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. Fig. 7 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

Authorized licensed use limited to: Univ of Calif Santa Barbara. Downloaded on March 06,2020 at 23:36:16 UTC from IEEE Xplore. Restrictions apply.

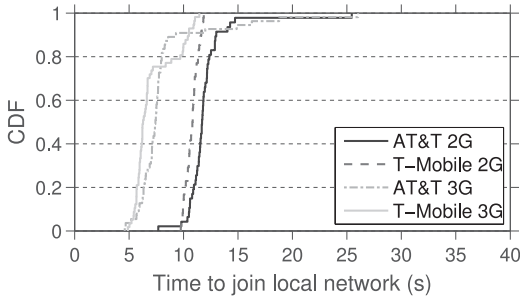

Fig. 7. Time spent without service when moving to the local network from commercial networks.

explore times for phone transitions in the opposite direction, from the local network to the commercial network. Fig. 8 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 Fig. 7. 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.

Impact of Duty Cycle Period on User Experience. We now evaluate the impact of duty cycle duration on user experience, as it relates to incurred power consumption and disconnected time. 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 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 (i.e., avoiding the use of large P values) while balancing the constraints imposed by limited battery life offered by a phone (i.e., avoiding very small P values).

We explore the relationships between the duty cycle period P and power consumption, as well as between P and

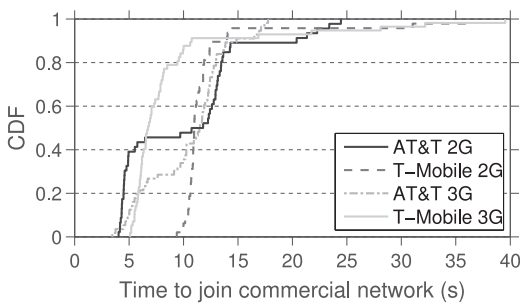

Fig. 8. Time spent without service when moving to commercial networks from a local cell.

TABLE 2
Average Measured Transition Times and Power Consumption

Meaning	Notation	Empirical value
Time to join commercial, 2G	τ^{2G}_{L-C}	11s
Time to join local, 2G	τ^{2G}_{C-L}	11s
Time to join commercial, 3G	τ^{3G}_{L-C}	7s
Time to join local, 3G	τ^{3G}_{C-L}	6s
Power to join commercial	p_{L-C}	4.507 mWh
Power to join commercial	p_{C-L}	8.33 mWh

disconnected time, in order to understand the effects of duty cycle duration on user experience. For this analysis, we use the empirical results for power consumption and disconnected time presented in Figs. 6, 7 and 8. Specifically, we use the average measured transition times and power consumption, as indicated in Table 2. Recall that in each duty cycle a CellSwitch phone has to transition across networks twice: once from local to commercial and one more time from commercial to local. Thus, the average power draw $\overline{p}$ per hour for CellSwitch's operation can be calculated as follows:

$$\overline{p} = n * (p_{L-C} + p_{C-L}), \tag{2}$$

where n is the number of network transitions per hour, calculated as $n = \lfloor \frac{3600}{P} \rfloor$.

Similarly, the percentage of time in an hour a user spends disconnected can be calculated as

$$\Delta^T = 100 * \frac{n * (\tau^T_{L-C} + \tau^T_{C-L})}{3600}), \tag{3}$$

where Δ^T is the percentage of disconnected time, while T is the used technology (2G or 3G).

We study the power consumption and disconnected time within an hour for a phone whose duty cycle P grows from 60 to 600 seconds in increments of 60 seconds. Our results are presented in Fig. 9, and show that shorter P values incur more power consumption and disconnected time within an hour, compared to longer P. Thus, a short P value (less than 200 seconds) may reduce the overall battery life of a device beyond what is likely to be 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 the aggregate energy draw and, in turn, on the device battery life. Our disconnect time results

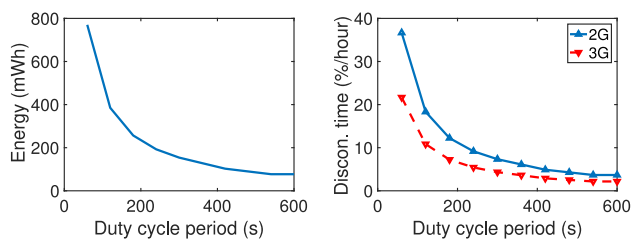

(a) Power consumption increases exponentially as P nears zero.

(b) Disconnected time increases exponentially as P nears zero.

Fig. 9. As the duty cycle period P is increased, the energy burden and the time a phone spends disconnected from both networks is decreased.

Authorized licensed use limited to: Univ of Calif Santa Barbara. Downloaded on March 06,2020 at 23:36:16 UTC from IEEE Xplore. Restrictions apply.

(Fig. 9b) indicate similar trends for power consumption. Thus, these combined results suggest that cycle lengths of 200-300 seconds provide satisfactory balance of power consumption, call initialization delay, and disconnected time.

2.4 Detailed HybridCell Operation

HybridCell's components work in concert to dynamically adapt usage of the local cellular network depending on the health of the nearby commercial network. Our current prototype implements both the network analyzer and the local cell components on a single computer. However, this is not required, as the network analyzer and the local cell machines must simply have connectivity between each other. Recall, the network analyzer produces two outputs: (1) cellular frequencies that are available for use; and (2) the Channel Availability Ψ_m of the nearby commercial cell. The list of cellular frequencies are used by the local cell to select frequencies that will not interefere with neighboring networks and are stored in a text file. The value for Ψ_m is a floating point value between 0 and 1, and is recorded to a file that is readible by the local cell processes. When CellSwitch phones join the local network during each duty cycle, the app connects to the local cell base station and retrieves the latest Ψ_m value, which is then used to determine the time spent on the local and commercial networks (t_L and t_C, respectively).

2.5 Rendezvous

HybridCell reduces the load on nearby overburdened commercial networks by adaptively duty cycling users onto a local network that in some cases may be disconnected to the global telephony network. This dynamic leads to a number of potential use cases depending on the current home networks of the caller and the callee. 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:

1) 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;
2) 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;
3) 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; the caller will be sent to voicemail. An SMS will be sent and queued by the commercial carrier. The recipient will receive the queued message, or voicemail notification, when they connect to the commercial carrier during their duty cycle; and
4) 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 and C-L 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. Thus, though there may be a delay, HybridCell can provide successful communication during times when the commercial network cannot. 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. This operation will result in minimal perceived delay, and in turn, seamless user interaction with the commercial network.

3 Evaluation

In this section, we provide an overview of the data collected during a site visit to Za'atari, and leverage this data to evaluate critical components of HybridCell. 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.

3.1 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" [23]. 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. In addition, our team conducted interviews and administered a survey on mobile phone and Internet use to camp residents.

3.1.1 Cellular Network Measurements

During the three days we spent in Za'atari, we used BladeRF [17] software defined radios to scan the 900 and 1,800 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` [24], 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.

3.1.2 Surveys and Interviews

Our team surveyed 228 residents of the refugee camp.[2] Of the respondents, 174 were youth aged 15 to 25, while 54 were

2. Data collection methodology and analysis recieved IRB approval.

Authorized licensed use limited to: Univ of Calif Santa Barbara. Downloaded on March 06,2020 at 23:36:16 UTC from IEEE Xplore. Restrictions apply.

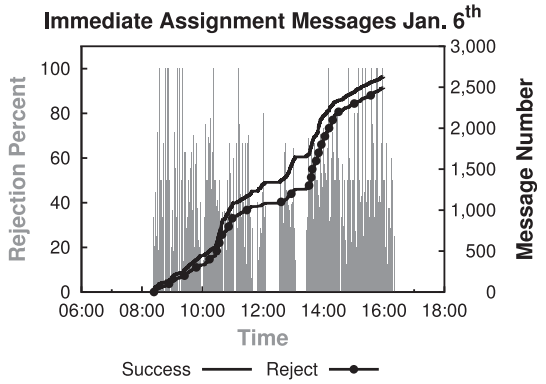

Fig. 10. Immediate assignment rejection percentage per minute, and successful and failed immediate assignment messages collected from Carrier A over 1 day in Za'atari.

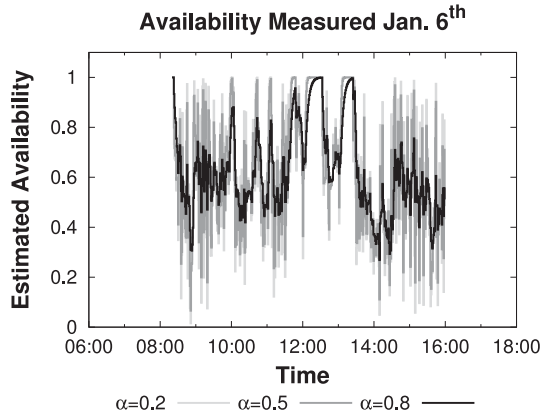

Fig. 11. Estimated availability for Carrier A with different α values.

adults 26 and over; 114 respondents were female and 108 were male, with 6 failing to indicate. Based on our survey, mobile phone ownership is nearly ubiquitous in the camp: 86 percent of youth and 98.1 percent of adults surveyed own a mobile phone. Android phones are very popular in Za'atari: 64 percent of the respondents own an Android device, 22.4 percent own a Nokia device, while about 4 percent own Apple iPhones. We believe the universal availability of Android handsets (e.g., 85 percent global market share [25]), and 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.

3.2 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.

Using one minute bins, Fig. 10 shows observed Immediate Assignment messages and Immediate Assignment Rejects collected over the course of one day for Carrier A, as well as the percent of immediate assignment messages that were rejections during our measurement window. We see different rates and percentages of rejection messages throughout the day. 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. Fig. 11 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 in the data presented in Fig. 10; we see low availability values during time windows of high rejection percentage, and a return to high values during windows of few rejections (e.g., 12:00). 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 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 $\alpha = 0.8$ is satisfactory in balancing responsiveness with congestion detection.

Using $\alpha = 0.8$ 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. Fig. 10 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 Fig. 12, 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. Fig. 12 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

Authorized licensed use limited to: Univ of Calif Santa Barbara. Downloaded on March 06,2020 at 23:36:16 UTC from IEEE Xplore. Restrictions apply.

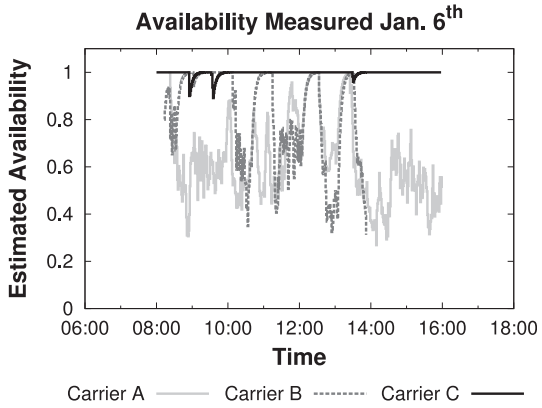

Fig. 12. Estimated availability for all carriers ($\alpha = 0.8$).

Assignment Reject messages in our traces, which may because it is the least popular carrier according to our survey. Over our entire collection period in Za'atari, Immediate Assignment Reject messages totaled 5.9 percent of channel assignment messages for Carrier C, compared to 33.7 percent for Carrier A and 15.2 percent for Carrier B. As shown in Fig. 12, 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.

3.3 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 that represent a phone requesting a resource (e.g., to place a call or SMS). For simplicity, in our simulation we treat all events, both success and reject, as requests for voice calls. We run two simulations in order to understand: 1) the anticipated effect of different values calculated by the availability metric on system usage; and 2) the behavior of HyridCell in a real-world environment.

3.3.1 Effect of Availability Metric

We first seek to learn the impact of the availability metric on system usage and end-user experience. We are particularly interested in the locations of caller and callee at the time of a call as well as any call initialization latency introduced by HybridCell in L-C or C-L events. Our first 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 immediate assignment 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 2.4.

Fig. 13a shows the breakdown for simulated events with varying levels of estimated availability and a duty cycle period set to 300 seconds. As expected, decreased availability leads to increased local network usage. We also see that mixed events (L-C or C-L) are most prevalent when clients spend 50 percent of their time on both networks. 'Misses' account for roughly 25 percent of calls in the worst case; however, this case would be caused by an estimation of 50 percent availability on a commercial network. In other words, we would observe a rejection rate of 50 percent for all requests if our system did not exist. HybridCell's presence in such a situation, on the other hand, would result in shifting 50 percent of call origination to the local cell (25 percent L-L plus 25 percent L-C calls), leaving 25 percent solely on the commercial network (C-C). Overall, we expect 75 percent of calls will succeed without a miss in this scenario, an overall increase of 25 percent compared with 50 percent failure.

Fig. 13b shows a CDF of latencies (i.e., call initialization delays) for all 'non-miss' events (e.g., L-L, L-C, C-C). 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 Fig. 13c and observe that as estimated availability decreases, users spend more time on the local cell, leading to shorter initialization latencies. Lastly, we run the simulation with various duty cycle periods with a constant 0.2 estimated availability. Fig. 13d shows the resulting CDF and the effect of cycle time on call setup delays. These results can be used to inform duty cycle configuration as each situation in which HybridCell can be deployed is

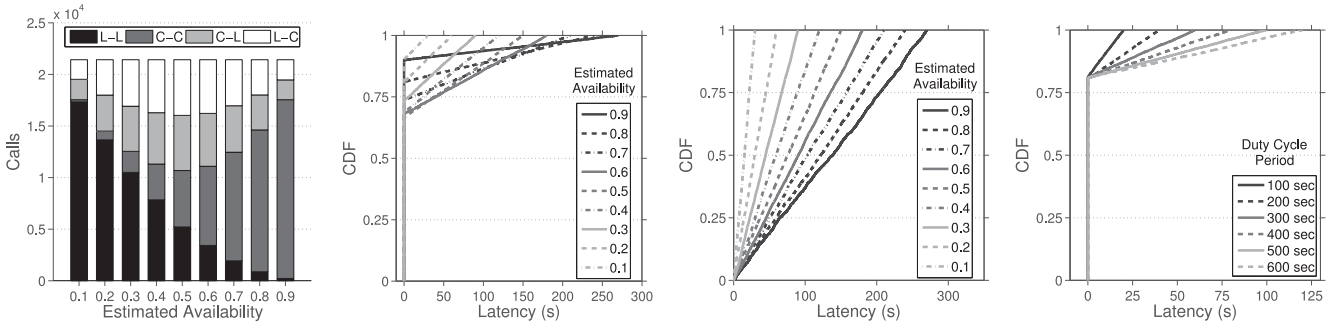

(a) Event designations with various estimated availability. As the commercial network is less available, more traffic is shifted to the local network.

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

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

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

Fig. 13. Simulation results with statically set estimated availability.

Authorized licensed use limited to: Univ of Calif Santa Barbara. Downloaded on March 06,2020 at 23:36:16 UTC from IEEE Xplore. Restrictions apply.

TABLE 3
Trace-Driven Simulation Call Totals

		Call Classification				
		L-L	C-C	C-L	L-C	Total
Caller-Callee	**A-A**	1,586	9,980	121	120	11,807
	A-B	46	1,166	8	133	1,353
	A-C	97	2,303	7	273	2,680
	B-A	33	1,219	159	16	1,427
	B-B	2	156	1	1	160
	B-C	6	311	4	1	322
	C-A	88	2,316	278	14	2,696
	C-B	9	332	1	3	345
	C-C	25	607	3	1	636

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.

3.3.2 HybridCell Augmentation of Real-World Networks

Next, we simulate the presence of HybridCell in Za'atari in order to understand system usage and impact on end-user experience. We split a pool of 83,500 users (the estimated population of Za'atari) among the three commercial carriers based on the results of our user surveys. 74.1 percent (61,873) are assigned to carrier A, 17.1 percent (14,279) to carrier B, and 8.8 percent (7,348) carrier C. We use 300 second duty cycles for all users. We use the immediate assignment events from our trace to calculate the availability for all three networks and adjust duty cycle percentages based on our findings. For this simulation we again assume that all immediate assignments, successes and rejections, are voice calls. We randomly select caller and callee pairs from the user pool and record the call designation (L-L, C-C, C-C, L-C) as well as carriers involved and any latencies introduced by L-C or C-L calls.

The resulting breakdown of calls, along with caller and callee designations is shown in Table 3. We observe that, as expected, the majority of calls involve carrier A due to its popularity. We see that the majority (85.83 percent) of all calls are still placed solely on commercial networks. This is expected given the estimated availability shown in Fig. 12. 8.83 percent of calls are L-L, meaning the burden of those calls has been entirely removed from the commercial networks. We also see a relatively small number of L-C and C-L events (2.62 and 2.72 percent, respectively).

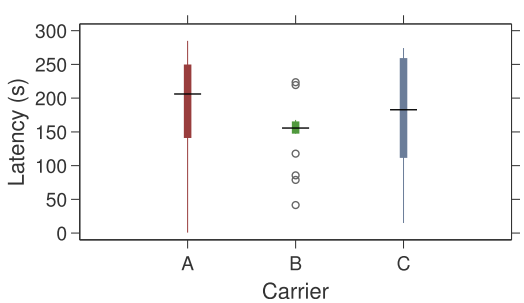

Fig. 14. L-C initialization latency for callers assigned to different carriers.

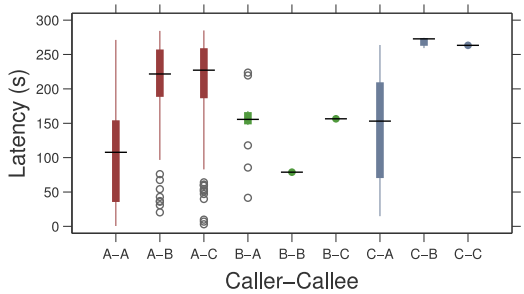

Fig. 15. L-C latency with caller-callee designations.

Overall L-C Call Initialization Latency. We plot L-C latency distributions for the three carriers in Fig. 14 with the tops and bottoms of the boxes representing the 25th and 75th percentiles, respectively and the horizontal line the median for each distribution. Calls originated by carrier A users have the highest median latency at roughly 205 seconds, while carrier B callers have the lowest. We see that carriers A and C have wide distributions of latency, whereas carrier B experienced a much smaller range. This can be partly attributed to small number of L-C calls in Table 3, as well as the availabilities shown for the carriers in Fig. 12. Carrier B experiences brief, severe periods of roughly 50 percent availability at the same time as carrier A, which should result in latencies of roughly half the duty cycle period P, or 150 seconds. Availability for carrier C, on the other hand, remains rather high throughout the observation period and does not appear synchronous with the other carriers, leading to a broad range of potential latency values.

L-C Latency with Caller-Callee Pairs Identified. The previous results, with carrier A callers experiencing the highest median latency, are less than ideal. In order to understand the underlying cause, we plot L-C latency distributions for all caller-callee pairs in Fig. 15. We observe that L-C latencies between carrier A users are roughly 110 seconds, less than half of the duty-cycle period, while A-B and A-C calls have much higher median latencies. This plot illustrates a challenge for HybridCell and an opportunity for improved design. When caller and callee are from the same carrier (e.g., A-A), the duty-cycles of that carrier will be the same for all users, resulting in low L-C latencies when the commercial network experiences low availability as predicted in Fig. 13c. Our selection of random caller-callee pairs results in a *worst-case scenario* for HybridCell, as the availability for the three carriers is independent (as shown in Fig. 12). While one carrier may be suffering from low availability at a given point in time, resulting in a higher percentage of calls from that carrier's users placed originated on the local network, callees from the other carriers, which could have high availability, may be using duty cycles where they spend the majority of time on the commercial network. Intuitively, we anticipate that callers and callees are *not* random, families and social groups may tend to use a single carrier, thus avoiding the potential synchrony variance. This worst-case scenario is illustrated by L-C and C-L events with caller-callee pairs involving A and the other two carriers. For instance, the highest percentage of C-L calls are between B-A and C-A. Because carrier A experienced sustained congestion throughout our observation period, carrier A users are likely to spend a higher proportion of their duty cycle connected to the local network. Conversely, users of the other

Authorized licensed use limited to: Univ of Calif Santa Barbara. Downloaded on March 06,2020 at 23:36:16 UTC from IEEE Xplore. Restrictions apply.

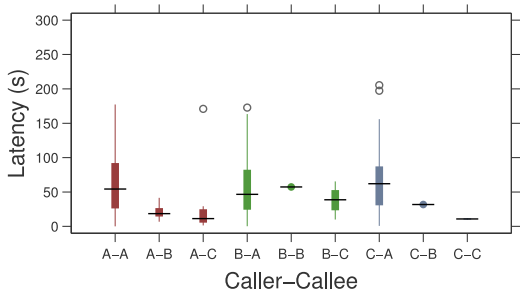

Fig. 16. C-L latency with caller-callee designations.

carriers, which have generally higher estimated availabilities, are more likely to be connected to the commercial network.

This finding exposes an opportunity for design improvement in HybridCell. There are several potential avenues to pursue in order to solve this problem. Carrier-specific actions, such as CellSwitch maintaining availability state of all area carriers and triggering specific phone behaviors based on the callee carrier, may be warranted. We leave this for future exploration.

C-L Call Initialization Latency with Caller-Callee Pairs Identified. Recall from Section 2.4 that C-L events will result in a queued SMS or a voicemail message as the callee is not globally available at the time of the event. We plot the latencies for the callees to receive an SMS or voicemail notification in Fig. 16. Across all three carriers, C-L latencies are lower than L-C latencies. This is expected, as the estimated availability of the commercial carriers tends to be above 0.5, meaning users generally spend more time on their commercial carriers than the local network. It is difficult to give general observations for calls that do *not* involve carrier A, given the low count as shown in Table 3. For calls that do involve carrier A, we see two trends. First, when the callee is a carrier A user (e.g., A-A, B-A, C-A), the distributions are similar and have median values of roughly 60 seconds. This is because carrier A callees, no matter the caller, have the same likelihood to be connected to the local network when a call is initiated. Second, for C-L where the caller is a carrier A user and the callees are from B or C, the latency medians are quite low at roughly 25 seconds. We believe this is because carriers B and C were generally more available that carrier A, so callees of those networks are likely to spend the majority of their time connected to the commercial network.

Overall, our simulations show that HybridCell successfully offloads a significant percentage of traffic from the nearby commercial networks, while not introducing an unacceptable amount of latency in mixed local and commercial call events. We anticipate that a real-world deployment would result in tuning some of the assumed settings used in our simulation (e.g., duty-cycle period, α used for estimated availability, and frequency of availability calculation) for environment-specific characteristics.

4 Related Work

HybridCell is related to two bodies of work: (i) wireless networks for remote areas and (ii) heterogeneous networks.

Wireless Networks for Remote Areas. 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 [3], [4], [5]. 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 [11]. 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 [26] 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 [27]. 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.

Heterogeneous Networks. HybridCell falls under the general umbrella of heterogeneous networks (HetNets), whose goal is to complement *incumbent* cellular networks with *augmenting* technologies such as Wi-Fi, LTE-U, and small cells [28], [29], [30], [31], [32]. HetNets can be compared across several key criteria including (i) whether the incumbent and the augmenting technology are cooperating, (ii) what is the target coverage area of the augmenting technology, (iii) what are the backhaul connectivity requirements of the augmenting technology, (iv) permanence of the augmenting technology and (v) what user handset capabilities does the HetNet require. HybridCell is fundamentally different than the state-of-the-art across all criteria. First, while HetNet technologies are presumed as desired by the incumbent, HybridCell is designed to provide connectivity in areas such as rural refugee camps or urban protests, where infrastructure augmentation can often be discouraged by local authorities [33]. As such, HybridCell

Authorized licensed use limited to: Univ of Calif Santa Barbara. Downloaded on March 06,2020 at 23:36:16 UTC from IEEE Xplore. Restrictions apply.

assumes no cooperation from the incumbent, whereas current HetNet solutions require cooperation. Second, while HetNets aim at improved spectrum reuse by reducing the coverage area of each cell to a nano- or pico-scale, HybridCell assumes wide-area cell footprint to maximize user coverage while minimizing the required hardware. Third, while all of the current HetNet technologies require robust backhaul connectivity in order to operate, HybridCell can scale its operation depending on the degree to which backhaul is available. When no backhaul is available, HybridCell can still augment for local calls. Where backhaul is available, HybridCell can bridge calls globally. Fourth, while HetNets assume permanence of deployment, HybridCell is designed with an outlook towards infrastructure mobility; one can take their augmenting network from one location to another and HybridCell will be able to adapt to its new environment. Last, but not least, HybridCell can operate on off-the-shelf phones, whereas some of the emerging HetNet modalities such as cellular-LTE-U require additional chipsets, that might not be readily available in our target user populations.

5 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 initialization 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.

Acknowledgments

This work was funded through US National Science Foundation Network Science and Engineering (NetSE) Award CNS-1064821 and US National Science Foundation Catalyzing New International Collaborations (CNIC) Award IIA-1427873. The authors thank Dr. Nijad al-Najdawi and UNHCR staff for facilitating access to Za'atari. P. Schmitt was with the University of California, Santa Barbara, while this work was done.

References

[1] Closing the network 'coverage gaps' in Asia. (2017). [Online]. Available: https://gsmaintelligence.com/research/2015/06/closing-the-network-coverage-gaps-in-asia/, Accessed on: May 01, 2015.
[2] So much going on!. (2017). [Online]. Available: http://rhizomatica.org/2015/01/14/so-much-going-on/, Accessed on: Apr. 18, 2015.
[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 *Proc. IEEE Int. Symp. Dyn. Spectr. Access Netw.*, Apr. 2014, pp. 271–282.
[4] K. Heimerl and E. Brewer, "The village base station," in *Proc. 4th ACM Workshop Netw. Syst. Developing Regions*, Dec. 2010, Art. no. 14.
[5] M. Zheleva, A. Paul, D. L. Johnson, and E. Belding, "Kwiizya: Local cellular network services in remote areas," in *Proc. 11th Annu. Int. Conf. Mobile Syst. Appl. Services*, Jun. 2013, pp. 417–430.
[6] (2017). [Online]. Available: http://www.openbts.org
[7] V. Chandrasekhar, J. Andrews, and A. Gatherer, "Femtocell networks: A survey," *IEEE Commun. Mag.*, vol. 46, no. 9, pp. 59–67, Sep. 2008.
[8] K. Heimerl, S. Hasan, K. Ali, E. Brewer, and T. Parikh, "Local, sustainable, small-scale cellular networks," in *Proc. 6th Int. Conf. Inf. Commun. Technol. Develop.*, 2013, pp. 2–12.
[9] (2017). [Online]. Available: https://www.endaga.com/
[10] P. Schmitt, D. Iland, M. Zheleva, and E. Belding, "HybridCell: Cellular connectivity on the fringes with demand-driven local cells," in *Proc. IEEE INFOCOM*, Apr. 2016, pp. 1–9.
[11] FAQ - Project Fi. (2017). [Online]. Available: https://fi.google.com/about/faq/#network-and-coverage-1, Accessed on: May 01, 2015.
[12] Syria Regional Refugee Response. (2017). [Online]. Available: http://data.unhcr.org/syrianrefugees/regional.php, Accessed on: Mar. 16, 2017.
[13] M. Zheleva, P. Schmitt, M. Vigil, and E. Belding, "Bringing visibility to rural users in Cote d'Ivoire," in *Proc. 6th Int. Conf. Inf. Commun. Technol. Develop.*, Dec. 2013, pp. 179–182.
[14] A. Ghasemi and E. Sousa, "Spectrum sensing in cognitive radio networks: Requirements, challenges and design trade-offs," *IEEE Commun. Mag.*, vol. 46, no. 4, pp. 32–39, Apr. 2008.
[15] FCC Makes 150 Megaherts of contriguous spectrum availabile for mobile broadband and other users through innovative sharing policies. (2017). [Online]. Available: https://www.fcc.gov/document/fcc-makes-150-megahertz-spectrum-available-mobile-broadband, Accessed on: Apr. 18, 2015.
[16] 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
[17] (2017). [Online]. Available: http://www.nuand.com/blog/product/bladerf-x40/, Accessed on: May 01, 2015.
[18] (2017). [Online]. Available: https://github.com/ptrkrysik/gr-gsm
[19] (2017). [Online]. Available: https://osmocom.org/projects/baseband/wiki/GSMTAP
[20] Maximizing Mobile. (2017). [Online]. Available: http://www.worldbank.org/en/topic/ict/publication/ic4d-2012, Accessed on: Feb. 02, 2017.

Authorized licensed use limited to: Univ of Calif Santa Barbara. Downloaded on March 06,2020 at 23:36:16 UTC from IEEE Xplore. Restrictions apply.

[21] N. Balasubramanian, A. Balasubramanian, and A. Venkataramani, "Energy consumption in mobile phones: A measurement study and implications for network applications," in *Proc. 9th ACM SIGCOMM Conf. Internet Meas.*, Nov. 2009, pp. 280–293.
[22] F. Qian, Z. Wang, A. Gerber, Z. M. Mao, S. Sen, and O. Spatscheck, "Characterizing radio resource allocation for 3G networks," in *Proc. 10th ACM SIGCOMM Conf. Internet Meas.*, Nov. 2010, pp. 137–150.
[23] M. Pizzi, "Logging on in Zaatari: Part I." (2017). [Online]. Available: http://www.smex.org/logging-on-in-zaatari-part-i/, Accessed on: Apr. 25, 2015.
[24] (2017). [Online]. Available: https://github.com/2b-as/xgoldmon, Accessed on: May 01, 2015.
[25] Android and iOS Squeeze the Competition. (2017). [Online]. Available: http://www.idc.com/getdoc.jsp?containerId=prUS25450615, Accessed on: Jul. 13, 2015.
[26] E. Gustafsson and A. Jonsson, "Always best connected," *IEEE Wireless Commun.*, vol. 10, no. 1, pp. 49–55, Feb. 2003.
[27] D. Iland and E. Belding, "EmergeNet: Robust, rapidly deployable cellular networks," *IEEE Commun. Mag.*, vol. 52, no. 12, pp. 74–80, Dec. 2014.
[28] J. G. Andrews, S. Buzzi, W. Choi, S. V. Hanly, A. Lozano, A. C. K. Soong, and J. C. Zhang, "What will 5G be?" *IEEE J. Sel. Areas Commun.*, vol. 32, no. 6, pp. 1065–1082, Jun. 2014.
[29] S. Sagari, S. Baysting, D. Saha, I. Seskar, W. Trappe, and D. Raychaudhuri, "Coordinated dynamic spectrum management of LTE-U and Wi-Fi networks," in *Proc. IEEE Int. Symp. Dyn. Spectr. Access Netw.*, Sep. 2015, pp. 209–220.
[30] A. Balasubramanian, R. Mahajan, and A. Venkataramani, "Augmenting mobile 3G using WiFi," in *Proc. 8th Int. Conf. Mobile Syst. Appl. Services*, 2010, pp. 209–222.
[31] M. Bennis, M. Simsek, A. Czylwik, W. Saad, S. Valentin, and M. Debbah, "When cellular meets WiFi in wireless small cell networks," *IEEE Commun. Mag.*, vol. 51, no. 6, pp. 44–50, Jun. 2013.
[32] I. Hwang, B. Song, and S. S. Soliman, "A holistic view on hyper-dense heterogeneous and small cell networks," *IEEE Commun. Mag.*, vol. 51, no. 6, pp. 20–27, Jun. 2013.
[33] P. Schmitt, D. Iland, E. Belding, B. Tomaszewski, Y. Xu, and C. Maitland, "Community-level access divides: A refugee camp case study," in *Proc. 8th Int. Conf. Inf. Commun. Technol. Develop.*, Jun. 2016, Art. no. 25.

Paul Schmitt is a postdoctoral research associate with the Center for Information Technology Policy, Princeton University. His research focus is on network systems design, and network measurement and performance analysis. His work spans a wide range of topics including local cellular networks, commercial cellular data core characterization, wireless spectrum sensing, and network connectivity in resource-limited environments.

Daniel Iland is working toward the PhD degree in the Department of Computer Science, University of California, Santa Barbara (UCSB). He is a senior software engineer on sensing, inference, and research with Uber. His work primarily focuses on using sensor fusion to improve mobile device location accuracy. At Uber, he helped form a cross organizational team to coordinate Uber's technology response to natural disasters. His research focused on enabling robust and reliable communication in emergency and disaster scenarios, wireless localization, and ICTD.

Mariya Zheleva received the PhD degree in computer science from the University of California, Santa Barbara, in 2014. She is an assistant professor with the Department of Computer Science, University at Albany SUNY. Her research interests include the intersection of wireless networks and information and communication technology for development. She has done work on small local cellular networks, dynamic spectrum access, spectrum management, and sensing and network performance and characterization. She is the founder and director of the UbiNet Lab, University at Albany.

Elizabeth Belding is a professor with the Department of Computer Science, University of California, Santa Barbara (UCSB). Her research focuses on mobile networking, including network performance analysis and information and communication technology for development (ICTD). She is the author of more than 140 technical papers and has served on more than 70 program committees for networking conferences. She is currently an editor-at-large of the *IEEE Transactions on Networking*. She is a fellow of the IEEE and an ACM Distinguished Scientist.

▷ **For more information on this or any other computing topic, please visit our Digital Library at** www.computer.org/publications/dlib.

Authorized licensed use limited to: Univ of Calif Santa Barbara. Downloaded on March 06,2020 at 23:36:16 UTC from IEEE Xplore. Restrictions apply.