

GSM Whitespaces: An Opportunity for Rural Cellular Service

*Shaddi Hasan
Kurtis Heimerl
Kate Harrison
Kashif Ali
Sean Roberts
Anant Sahai
Eric Brewer*

Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2013-198

<http://www.eecs.berkeley.edu/Pubs/TechRpts/2013/EECS-2013-198.html>

December 8, 2013



Copyright © 2013, by the author(s).
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

GSM Whitespaces: An Opportunity for Rural Cellular Service

Shaddi Hasan, Kurtis Heimerl, Kate Harrison, Kashif Ali, Sean Roberts, Anant Sahai, Eric Brewer
{shaddi, kheimerl, harriska, kashif, sahai, brewer}@eecs.berkeley.edu, seanroberts@berkeley.edu
Dept. of Electrical Engineering and Computer Science
University of California, Berkeley

Abstract—The GSM network is the largest network on Earth, providing vital communications service to billions of people. Yet hundreds of millions of people live outside coverage of existing cellular providers. Recently, researchers have demonstrated a new model of cellular connectivity, community cellular, that has the potential to bring coverage to extremely rural populations. Although the total capital costs for these networks (<US\$10,000) are affordable for rural communities, the high financial and political cost of access to spectrum makes running a legal community cellular network nearly impossible. This is despite the fact that very rural areas typically have substantial amounts of licensed spectrum that is not actively used.

We define this unused spectrum as *GSM whitespace*. Like TV whitespaces, we argue that GSM whitespaces should be regulated for dynamic spectrum sharing, and that doing so will support the growth of community cellular networks and thereby improve rural access to communications services. We propose a hybrid sensing and database-driven spectrum sharing scheme called *Nomadic GSM* that provides safe coexistence between primary and secondary users without requiring coordination or cooperation from existing license holders. Nomadic GSM also mitigates concerns about “spectrum squatting” by secondaries and provides regulators visibility into and control of the spectrum usage of secondary operators. Lastly, we implement and evaluate a prototype Nomadic GSM system that is fully compatible with existing, unmodified GSM handsets and has been deployed in a real-world community cellular network in Papua, Indonesia.

I. INTRODUCTION

Cellular networks are the largest communication systems on Earth, with billions of people relying on them for basic communication services. These networks have positive economic impacts ([30], [17]), and spending on telecommunication matches that of a necessity among the poor [2]. As such, providing universal access to cellular service is an important policy objective, with universal service obligations (USO) and subsidies being the primary policy mechanisms for accomplishing this. Although these programs are important and have been widely adopted, they have failed to provide truly universal service: hundreds of millions of people worldwide live outside the coverage area of existing cellular carriers. These people largely live in rural areas, where traditional telcos are unable to operate profitably due to high costs and low subscriber density. Even with USO subsidies, carriers are inherently disinclined to invest in rural infrastructure, which is marginally profitable at best, particularly given the opportunity cost compared to investing in lucrative 3G and 4G infrastructure in urban areas.

Fundamentally, USOs assume a *top-down* model of cellular deployment where large-scale organizations bring cellular




Fig. 1. Overview of Nomadic GSM. (1) The secondary initially queries a regulatory database for available channels in its area. (2) The secondary gathers measurement reports from its subscribers’ phones. (3) Secondaries report spectrum usage (both their own and measured) and service usage (e.g., number of calls and SMS) to a database on a regular basis. (4) Secondaries use measurement report data and guidance from the reg. DB to pick future non-interfering channels to use, and regularly change channels. (5) Optionally, primaries update the regulatory database with their own spectrum usage and query it to obtain reports on what spectrum in use by secondary operators.

coverage to remote rural areas for subsidies. Historically this made sense; there were only a handful of organizations capable of building and operating cellular networks in any particular country. Yet recent innovations invalidate this assumption. The cost and complexity of building and maintaining a GSM cellular network has decreased to the point where individuals or rural communities can create their own micro-scale cellular networks [16]. These *bottom-up* “community cellular networks” (CCNs) have demonstrated that local communities can operate their own telecommunications infrastructure and provide essential communication services.

Unfortunately, despite the potential benefits of community cellular networks, regulatory barriers exist. GSM uses licensed spectrum, and gaining access to long-term spectrum licenses is almost impossible for micro-scale rural operators. To solve this, we argue for spectrum sharing in *GSM whitespaces* to provide GSM service in rural areas. GSM whitespace refers to licensed GSM spectrum that is unused in a particular geographical area¹ and thus could be re-used by a secondary operator without interfering with the primary license holder. By allowing CCNs to operate in GSM whitespaces, regulators would empower rural communities to build infrastructure appropriate to their own needs, without waiting for incumbent carriers to begrudgingly allocate resources their way.

To enable this, we propose Nomadic GSM (NGSM), a hybrid sensing and database-driven approach for GSM spectrum sharing (Fig. 1). NGSM takes advantage of the fact that

¹Note we do *not* refer to the space between adjacent GSM channels.

GSM handsets continually measure spectrum occupancy in their neighborhood and report these measurements back to their serving cell. By doing so, we can sense for potential interference at the *client* device, avoiding the hidden node problem. Although certain edge cases necessitate a spectrum occupancy database, NGSM enables secondary operators like CCNs to share licensed spectrum without requiring cooperation or participation from existing licenseholders. NGSM works with existing, unmodified GSM handsets. As such, it is deployable today, and we demonstrate this with a prototype deployment in a CCN in Papua, Indonesia.

The contribution of this paper is as follows. First, we define GSM whitespaces and describe Nomadic GSM, a scheme for dynamic spectrum sharing in GSM whitespaces that enables secondary operators—community cellular networks—to provide service without interfering with each other or with primaries and that does not require explicit cooperation or engagement with primary license holders. Next, we consider the opportunities and risks spectrum sharing presents to major stakeholders and how NGSM addresses these. Finally, we demonstrate the feasibility of our proposal by building, deploying, and evaluating a prototype implementation of NGSM that is compatible with existing, unmodified GSM handsets. We close with a discussion of why the whitespace approach works better than the obvious market-based alternatives and a path forward for regulators.

II. RELATED WORK

A. Policies for Rural Service

One policy mechanism for bringing coverage to rural areas is a universal service obligation (USO) [26]. USOs, originally developed for postal service, refer to a requirement for a baseline level of service to every resident of a country. An example is the US Telecommunications Act of 1996 [32], whose goals were to promote the availability of quality services at reasonable rates, increase access to advanced services, and provide these services to all consumers, including low-income or rural people. Similar regulations exist in many countries, including Indonesia [29], where we deployed our pilot system.

Despite these lofty goals, hundreds of millions of people in the world remain without basic telecommunication services. The reasons for this are fundamentally economic; operators would prefer to work in areas where they are profitable without the headache of dealing with USOs [5]. Researchers have attempted to address some concerns with USO systems through competitive means [8], including USO auctions [31]. This work argues for a fundamentally different model of rural access; one owned and operated by rural entrants and communities themselves. This would free traditional firms from USOs while providing coverage in underserved markets.

B. Locally-owned Infrastructure

Local or community ownership or development of critical infrastructure has a long history. A well-known concept is co-production [18], targeting infrastructure such as irrigation [19]. There is a similar history of small-scale cooperative or locally owned telephony networks [10]. Modern cellular networks have largely ignored these models in most of the world, focusing instead on nation-wide networks and coverage. The

Village Phone [3] was a system where “phone ladies” would buy handsets and sell use, similar to a phone booth; while the network infrastructure was owned by a nationwide carrier (Grameen), local entrepreneurs provided access to the network to their community. Galperin et al. [11] proposed running cellular networks as small-scale cooperatives, using evidence of earlier cooperative telephony networks in Latin America as a motivating example. Elgar made similar arguments for the viability of “bottom up” telecommunications [9]. However, only recently has cellular equipment become inexpensive enough for these models to be economically feasible.

One example of cheap cellular equipment is OpenBTS [20]. OpenBTS is an open-source GSM base transceiver station (BTS) implementation which has enabled a wide range of projects aimed towards building small-scale “community cellular” networks [15]. Heimerl et al. demonstrated the viability of independently run, locally operated cellular networks [16]. Similarly, Rhizomatica has deployed several community-run cellular networks in Oaxaca, Mexico [22]. Zheleva et al. [34] deployed a similar system for purely local communications in Zambia. Of these networks, only the Oaxaca network has a short-term experimental spectrum license; the rest operate without licenses. This reality motivates our desire to develop a mechanism for effectively licensing and regulating spectrum access for community cellular networks.

C. Cognitive Radio

The literature on cognitive radio, whitespaces, and dynamic spectrum sharing is vast; while most work in the space focuses on TV whitespaces (TVWS), our work is more closely related to work on re-use of cellular spectrum. Sankaranarayanan et al. [23] propose reusing timeslots in a GSM cell for adhoc networks during periods when the GSM cell is lightly utilized. Buddhikot et al. [6] describe a system for indoor femtocells to dynamically share spectrum with incumbent carriers by operating over ultra wide bands. Yin et al. [33] proposes a similar system and provides measurement results which indicate that unused spectrum (i.e., whitespace) exists even in a dense, urban environment (Beijing). The assumption in the community, however, seems to be that cellular spectrum is efficiently used and that finding GSM whitespace is challenging.

In contrast to these, we focus on reusing GSM whitespaces to provide GSM service by means of macrocells in rural areas. Moreover, rather than relying on fine-grained spectrum sharing, we rely on spatial separation to provide coarse-grained sharing at the level of full GSM channels. This high margin for error—due to the large distance between primary and secondary networks—along with our novel sensing strategy is likely to be more appealing to incumbents.

III. COMMUNITY CELLULAR NETWORKS

Historically, cellular networks have been expensive to build and complicated to operate; this is particularly the case for rural cellular networks [14]. A single rural GSM macrocell can cost upwards of US\$500,000 to build, not including the supporting network core infrastructure that the network operator must already possess. Macrocells have high power consumption, and in areas without reliable grid power must rely on diesel generators; the fuel for these generators is a

major ongoing operational expense and target for theft [12]. These factors have created a situation where only a handful of entities, primarily large corporations or governments, are able to operate cellular networks. Spectrum licensing compounds this: not only must an organization who wants to obtain a license pay large amounts of money, they also must understand how spectrum is regulated, how and when auctions take place, and how to participate in those auctions, all factors which raise the barrier to entry for small organizations.

Recent technological innovations—notably, low-cost software defined radios and open-source software such as OpenBTS [20]—have challenged this status quo. A rural community can build and operate their own cellular network for under \$10,000 in capital expenditure [16]. Low-power equipment can be operated using solar panels, dramatically reducing operational expenses. These networks rely on voice over IP (VoIP) technology and can thus use any available Internet backhaul to connect to the global telephony network), including satellite or fixed wireless broadband.

These advancements have enabled a new model, the *community cellular network* [15]. Community cellular networks are locally owned and operated, and they consist of at most a handful of BTS sites. Such networks exist in Papua, Indonesia [16] and Oaxaca, Mexico [22]. Not only are these networks effectively serving rural communities where incumbent carriers have failed (or even refused) to do so, they are financially sustainable for the local operators. The Papua network, for example, generates a revenue of around US\$1,000 per month, which while minuscule by traditional telco standards represents a good business opportunity for a local entrepreneur. Moreover, both of these networks were built and are operated without any involvement or coordination with existing operators.²

Compared to traditional cellular networks, the core advantage of CCNs is that they enable local independent entrepreneurs to solve their own communication problems. There’s no reason existing telcos cannot take advantage of low-cost equipment targeted towards CCNs to build out rural infrastructure, but access to low-cost equipment isn’t enough to ensure sustainable operation in rural areas. A key finding from prior work on community cellular networks is that locally operated microtelcos have the flexibility to make decisions that traditional telcos cannot. In the example of the Papuan CCN [16], service was *coproduced* [18], [19] with the local community: pricing decisions were made locally, and electricity and backhaul were sourced from a school in the community. The microtelco in Papua was also able to set prices that were appropriate for their own community and costs, thus ensuring sustainability. A large-scale telco does not have this flexibility—the overhead of managing small, potentially informal, relationships with many widely distributed partners is prohibitively expensive and time consuming. Yet these relationships and the understanding of local community structure and norms are the key advantages of local entrepreneurs.

Beyond simply being more affordable, CCNs also have inherent advantages for providing rural service. Although other technologies and spectrum bands (e.g., WiFi) could provide rural communications services, using operating GSM base

²Indeed, the network in Papua is operating without a license, though it has not received any complaints.

Uplink (MHz)	Downlink (MHz)	Licensee
890.0 - 900.0	935.0 - 945.0	Indosat
900.0 - 907.5	945.0 - 952.5	Telkomsel
907.5 - 915.0	952.5 - 960.0	XL

TABLE I. BANDPLAN FOR THE GSM900 BAND IN INDONESIA [25]. THE ENTIRETY OF THE BAND HAS BEEN GRANTED TO THESE THREE CARRIERS UNDER NATIONWIDE LICENSES.

stations in spectrum traditionally used for GSM networks leverages the wide installed base of billions of existing handsets with existing charging, repair, and distribution infrastructure. Inexpensive and ubiquitous, existing GSM phones ease adoption by providing a familiar experience for end users. People want to be able to use their existing phones, and it’s unlikely any manufacturer will produce a cheap, durable phone just for rural areas using a novel protocol.

CCNs put operating cellular network infrastructure within reach of individuals. It is technically and economically feasible for individuals to deploy this infrastructure for their communities on their own initiative, as many already do with WiFi infrastructure. The primary obstacle is access to spectrum: unlike WiFi, devices for cellular networks operate in licensed bands. Removing this barrier is vital to widespread deployment of community cellular networks, and their unique strengths argue for policy mechanisms to support their growth. GSM whitespace presents an opportunity to resolve this tension.

IV. GSM WHITESPACES

A. Defining GSM Whitespace

GSM whitespace refers to spectrum that has been licensed to carriers for GSM networks but is unused in a particular geographic area. As defined, GSM whitespaces are incredibly common worldwide: due to exclusive licensing of 2G GSM spectrum, any areas that are unserved by telcos are guaranteed to have unused spectrum in the 2G GSM bands.

Consider the case of Indonesia. Fig. 2 shows the national cellular coverage map for Indonesia.³ Although the entire GSM900 and GSM1800 bands have been licensed to carriers in Indonesia (Table I), vast swaths of the nation remain without any coverage. The largest provider, Telkomsel, claims to cover “over 95%” of the population as of 2013 [28], meaning close to 10 million people live outside of coverage in Indonesia alone. In contrast, the GSMA suggests [13] this number could be as high as 90 million. The number of people living outside of coverage (and hence in areas with ample GSM whitespace) could exceed a billion in developing countries alone.

Exclusive licensing of GSM spectrum creates significant amounts of unused spectrum. Regulating spectrum in rural areas in the same way as urban areas inflicts a significant social cost: although low potential revenue makes it difficult for incumbent carriers to justify providing service in remote areas, exclusive license agreements prevent any others from offering service. Licenses have traditionally been offered in this manner

³Obtaining accurate data on what areas are actually served is very difficult. The data from this figure comes from the map of international roaming coverage published by AT&T. It generally matches self-reported tower locations found in annual reports [27] and crowdsourced coverage maps [24].



Fig. 2. Indonesian cellular coverage. Wide swaths of sparsely populated parts of the country lack any cellular coverage, which includes at least 10 million people. The red star on the right marks the location of the Papua CCN.

because there was no local competition for the rural spectrum and it was easier for carriers to plan their networks assuming an exclusive license. We recognize the latter reason as valid, but the rise of CCNs puts the former out of date.

B. Spectrum Sharing in GSM Whitespaces

Our proposal to resolve this disconnect between spectrum licensing and rural service is simple: allow CCNs to utilize spectrum available in GSM whitespaces. Although we can draw some lessons from work on TVWS, the opportunities presented by GSM whitespaces have fundamental differences. Most importantly, our usage scenario is far simpler than those envisioned for TVWS. Our proposal aims to broaden access to basic communications services, not to maximize spectrum utilization. We are only interested in enabling a single type of service in the whitespace, GSM cellular service, and this service has well-defined and minimal spectrum requirements (each channel is 200kHz wide). We are also primarily concerned with operation in rural areas with ample available spectrum. Finally, the economics of CCNs suggest that few secondary operators will coexist in the same area at the same time; we stress again that the localities CCNs are designed to serve are *unprofitable* for traditional telcos. This constrained design space simplifies our task.

Our goals for GSM whitespace spectrum sharing are:

- 1) *Safety*. Secondary operators should be able to provide cellular service in unused spectrum in standard GSM bands without interfering with primaries or other secondary operators.
- 2) *Independence*. Primary operators should have no new burdens restricting their usage, and should not need to cooperate with (or be aware of) secondary operators. Similarly, secondaries should not require special permission from or coordination with a primary.
- 3) *Verifiability*. Regulators and primaries should have visibility into what spectrum secondaries are using, and they should be able to verify that secondaries are actually providing service.
- 4) *Spectrum flexibility*. Secondary users should not be able to claim that use of any particular channel is necessary for their operation.⁴
- 5) *Backwards compatibility*. Existing, unmodified GSM phones should work with secondaries' networks.

⁴This idea was advanced in a public conversation by John Chapin during the 2012 ISART workshop in Boulder, CO.

We can achieve *safety* and *independence* by demonstrating a robust and reliable mechanism for detecting spectrum usage of other nearby operators, both primary and secondary. By reporting spectrum utilization measurements and usage by subscribers to a regulatory database, secondaries can provide *verifiable* rural coverage. *Spectrum flexibility* comes from ensuring secondaries have an actively and often exercised mechanism for frequently changing their broadcast channel without compromising their ability to provide service. By only leveraging existing mechanisms in the GSM specification, we can do all of this while maintaining *backwards compatibility*.

V. NOMADIC GSM

The linchpin of our proposal is the feasibility of implementing a GSM base station that can achieve our goals for sharing spectrum in GSM whitespaces; this is Nomadic GSM (Fig. 1). NGSM is able to:

- quickly detect when it may be causing interference to a primary or another secondary operator (*safety, independence*);
- rapidly and frequently adjust its frequency usage to avoid causing interference (*spectrum flexibility*);
- accurately report its own frequency usage, as well as the frequency usage of other users in its area, to a regulatory database (*safety, verifiability*);
- and achieve the above without requiring modifications to existing client devices or significant interaction with existing license holders (*backwards compatibility*).

In this section, we describe the mechanisms by which NGSM meets these goals. We discuss the first three points in turn while continuously addressing the fourth.

A. Interference Detection

A key issue for dynamic spectrum sharing schemes that rely on sensing is the hidden node problem [4]. By definition, interference occurs at a receiver, so two transmitters may be interfering with each other even if they are unable to detect each other's transmissions by sensing the medium.

One solution to this problem that has been proposed for TVWS is a regulatory database of frequency usage. A similar database-driven approach to spectrum sharing also fits GSM whitespace. By their nature, GSM base stations will be connected to the Internet in order to provide service to their users; a local-only GSM network is only useful in limited cases. For example, in the Papua network roughly 66% of traffic is outbound [16]. We can report frequency usage and information on unused channels in the BTS's area to a database using this Internet connection. We assume secondary GSM operators will be willing to accept new regulatory requirements, such as registering their spectrum usage with a regulatory database. However, it is impractical (and contrary to our goals) to assume incumbent operators will accurately register their systems to a database; in effect, they will not be cooperating with secondary operators. We need a system to enable non-cooperative base stations to coexist with cooperative ones; this is a form of coexistence-based spectrum sharing [21].

NGSM leverages part of the GSM standard to overcome this challenge [1]. Every GSM BTS operates on one or

more channels, known as ARFCNs (Absolute Radio Frequency Channel Number); because GSM employs frequency-division duplexing, an ARFCN specifies a particular pair of frequencies used for downlink (from the BTS to phones) and uplink (from phones to the BTS). In order to support handover of a phone between cells, base stations provide a list of frequencies for up to six “neighbor” cells (the “neighbor list”) to phones that are camped to (i.e., associated with) the base station. Since BTSs initiate handover, phones regularly scan each of these frequencies and report back the received signal strength (RSSI) for each, along with one for the current base station. The report also contains network and base station identification codes for each active ARFCN discovered.

By intelligently selecting the neighbor list at the BTS, NGSM can induce phones to report usage on frequencies of our choosing, without any modifications to the phones. Suppose we wanted to monitor whether ARFCN 20 is in use. NGSM would add this ARFCN to its neighbor list and then wait for measurement reports from handsets. If ARFCN 20 were not in use, handsets would report back as such. However, if another provider was actively using that band handsets would detect the other signal and inform our base station of its use. Importantly, this approach solves the hidden node problem by measuring interference at handsets, rather than at the BTS. However, all new logic required by NGSM is implemented at the BTS, ensuring backwards-compatibility with existing handsets. While conceptually similar to sharing spectrum sensing results as proposed in CORVUS [7], backwards compatibility with unmodified devices sets NGSM apart.

Monitoring the BTS’s current ARFCN is slightly more complicated. Measurement reports are ambiguous in this case: if a handset reports a high RSSI for our ARFCN, it’s impossible to know if that reading is due to the handset being near our tower or because we are interfering with another tower. Fortunately, there is a simple solution: configure our base station to use two or more ARFCNs simultaneously, rather than one. This is a common and well-supported configuration for GSM base stations, since a cell’s capacity is directly related to the number of ARFCNs it supports.

NGSM handles this case as follows. First, we ensure that the neighbor list transmitted by the BTS contains on each of its ARFCNs contains both of the BTS’s ARFCNs. Next, we *alternate* between each ARFCN, turning one completely off. Because the phones continue to receive both ARFCNs in their neighbor list, however, the BTS continues to receive measurement reports for both ARFCNs. If a primary user operates on the same ARFCN as one of our two ARFCNs, phones will continue to report the ARFCN is in use, even during periods when we have turned that ARFCN off, allowing us to detect which of our ARFCNs are no longer safe for use. The faster the rate at which we switch ARFCNs, the sooner we are able to detect potential interference.

Finally, we note that we can set the threshold for considering a channel occupied quite low since (1) switching to another frequency is easy to do and (2) there are likely many GSM channels available. Note that this technique can work with any number of ARFCNs per BTS, not just two, by always leaving one ARFCN off.

B. Changing Frequencies

The secondary’s BTS changes its frequency use in three cases. First, to avoid causing interference: once a BTS detects that it may be causing interference, whether via measurement reports from handsets or the regulatory database, it needs to be able to quickly modify its frequency usage. Second, a secondary needs to cycle through different frequencies on a regular basis. Doing so prevents secondary operators from claiming a particular frequency is essential for their operation, thus protecting primaries from spectrum squatting (Section VI-A). Finally, the BTS must switch between two channels during regular operation in order to detect interference on its own channels. The final two cases differ in timescale: while changing frequencies once per day may be sufficient for the former, in the latter case we want to be able to switch between channels quickly, on the order of minutes or even seconds.

What mechanism should we use to change channels? A naive solution would be to simply change the ARFCN on which the secondary’s BTS operates. From the perspective of a phone, this is equivalent to shutting off the BTS on the old ARFCN and bringing up a new BTS on a different ARFCN. However, this approach has a serious downside: phones will have to re-associate with the BTS after each channel switch, causing downtime for users (phones take up to two minutes to reassociate [14]). Active calls would also be disrupted during an ARFCN switch. Given one of our primary design criteria is compatibility with existing, unmodified handsets, there is a tension between frequency agility and system usability.

We can address this concern in part by only cycling frequencies while the BTS is not being used (i.e., no active calls, SMS or data transfers). Rural cellular networks tend to be lightly utilized, especially during off-peak hours [14]. By cycling frequencies only when the BTS is not in use, we can avoid interrupting ongoing calls while minimizing perceived downtime to users. However, we don’t actually reduce the amount of time that the BTS remains out of service.

We can take this one step further by leveraging the GSM handover mechanism and the fact that our BTS operates as two cells (i.e., already operates on two ARFCNs for the purposes of detecting interference on our own ARFCNs, as described earlier). Handover is designed to move a handset between cells of GSM network during a call and is instigated by the network infrastructure (the base station controller in a traditional GSM network) when call quality degrades. In this model, once the BTS decides to change one of its two ARFCNs, it first initiates handover for all phones camped to that ARFCN and moves them to the other ARFCN. Once all phones have camped to the new ARFCN, the BTS can safely turn off that channel or tune it to a new ARFCN. Phones will experience no downtime; even in-progress calls are not interrupted.

Importantly, GSM handover is universally adopted and widely used functionality for all GSM networks globally and as a result is widely implemented and tested in client devices. For example, handover allows users to make uninterrupted calls while in a moving vehicle. While our technique for determining when to perform handover is novel, the mechanism by which we would move clients from one frequency to another is completely standard.

C. Policing and Reporting Usage

We've described mechanisms for detecting interference by leveraging reports from phones and changing channels frequently without significantly impacting users. We now turn to reporting usage and policing spectrum use. As discussed earlier, we assume that all secondaries' BTSs will have Internet access; this is reasonable given such access is essential to provide service to the public telephone network. Given this, these systems have two unique capabilities. First, these BTS units measure actual spectrum usage in their service area. Measurement reports gathered by phones can be used to determine ground-truth regarding spectrum usage in an area. This applies to both the spectrum the CCN gathering the reports is using as well as others in the area, enabling secondary users to "police" their area and report the existence of nonconforming operators. Secondly, CCNs know actual aggregate usage statistics about their users, such as number of calls or SMS served per day. Reporting both of these measurements to a database would give regulators insight into the scale and nature of rural service, and provide an effective mechanism for policing compliance with regulations on usage of GSM whitespaces. Incumbent licenseholders can also benefit from this data by using it to plan their network expansion into rural markets or for determining what portions of their spectrum are being used where by CCNs to obtain credit towards fulfilling USOs (Section VI-A).

Unlike the TV whitespaces (and others), we believe that the core role of a GSM whitespace database is to enable reporting, rather than to guarantee noninterference by appropriately herding devices (in frequency). This is important because it means that the GSM database does not require any action on the part of the incumbent before systems can *safely* begin using the GSM whitespaces. However, regulators can also respond to actual interference events by using the database to rapidly direct CCNs away from the frequencies on which interference is being perceived (a "frequency kill switch" of sorts).

VI. OPPORTUNITIES, RISKS, AND INCENTIVES

A. Already Licensed Carriers

Although GSM frequency bands are heavily utilized in urban areas with high subscriber densities, spectrum is plentiful in unserved rural areas. We argue that sharing this rural spectrum imposes little if any cost to incumbents. For example, the Papuan CCN [16] operates in a frequency that has been licensed to Telkomsel, the largest Indonesian carrier. The network is serving a village four hours by car from the nearest place with cellular coverage. Although not legal, strictly speaking, the Papua CCN is isolated and does not impact the licenseholder's operations. It's even plausible that Telkomsel could provide service concurrently with the CCN if it decided to serve the same village; due to low subscriber density, the CCN is able to effectively serve its community with two ARFCNs (0.5MHz), under 7% of Telkomsel's GSM900 license allocation. That leaves most spectrum available to Telkomsel if it ever decides to serve the same area, even though their spectrum needs would be similar to that of the CCN.

Carriers need more than just assurance that sharing their spectrum won't impose any direct costs. We consider how to mitigate some potential risks and outline benefits spectrum sharing provides incumbent carriers.

1) *Benefits and Incentives:* GSM whitespaces offer several potential benefits for incumbent operators.

Fulfilling universal service obligations. Sharing spectrum with CCNs could serve to fulfill a carrier's universal service obligations. Whether service is provided by a carrier or a CCN is functionally similar: rural customers receive access to communications service in both cases. Allowing carriers to take "credit" for CCNs operating in their spectrum for the purposes of demonstrating providing rural service would provide a strong incentive for carriers to support policies that enable spectrum sharing in GSM whitespaces. Requiring secondary users to report their spectrum usage and subscriber activity to a regulatory database (per our goal of verifiability) enables a simple way for carriers and regulators to determine such credit in a trustworthy way. Carriers would be able to automatically generate reports from the regulatory database to learn what areas and how many people CCNs that use the carrier's spectrum are serving. Since these reports come from CCNs, which are unlikely to have formal arrangements with carriers, "faking" this data will be difficult for carriers without directly supporting CCNs or providing service themselves. We envision the database being public, allowing civil society to call out untruthful providers.

Opening up new rural markets. By their nature, CCNs open up new markets for cellular service in rural areas. These markets start small, but grow as the community and infrastructure expand. The CCN's presence encourages local investment in cellular phones and businesses to adopt the technology to improve their processes. Eventually, these markets may become economically viable for incumbents, and the CCN's presence has prepared the community for their arrival.

Incumbent carriers could take advantage of this progression in more immediate ways as well. One example could be entering into a partnership with the local CCN where the carrier captures the CCN's customers as its own when these users travel into the carrier's network coverage. This approach preserves the autonomy of the independent CCN operator and has low overhead for the carrier while providing a channel for the carrier to acquire new, otherwise hard-to-reach, customers. When the incumbent eventually enters the rural market, those customers are immediately available.

2) *Mitigating Risks:* At the same time, spectrum sharing carries significant potential risks, the most significant of which is "spectrum squatting". NGSM's spectrum flexibility mitigates this to an extent, but GSM whitespaces offer inherent protections for primary users as well.

The "Grandparent Problem". The "grandparent problem" is a potential risk that carriers face when they allow another entity to provide service in spectrum they own. If the carrier ever wants to reclaim that spectrum from the other entity (e.g., once the agreement expires), customers of the secondary entity may lose service and be upset at the carrier for being the "cause" of their service disruption. If those customers are a politically important constituency, such as a grandparent who is no longer able to communicate with grandchildren, the carrier may find itself in the crosshairs of negative public opinion and under pressure from policymakers and regulators to continue allowing the secondary entity to provide service in their frequency band. This isn't a concern with

GSM whitespaces. Spectrum flexibility ensures secondaries can easily switch frequencies and continue providing service to their users even if a few GSM primaries decide to put a portion of their spectrum to use. Low population density means rural markets have minimal spectrum requirements, providing plenty of room for secondaries to coexist with primaries. Moreover, since all users of the band would be using GSM technology, the customers of secondaries could easily switch to the new provider by changing a SIM card.

Avoiding enabling new competitors. Another significant concern carriers are likely to have is that by sharing spectrum with CCNs they are enabling new competitors. At a high level, CCNs are not competing with incumbent carriers in the most significant markets, urban areas. There could be situations where a CCN and an incumbent carrier try to serve the same area. Competitively, this could resolve in two ways. First, the CCN could have intrinsically lower costs of operation than an incumbent carrier. If this is the case, the incumbent should simply use the CCN as a roaming partner, an established model in the cellular business, that would allow the carrier to receive the network connectivity they need at a lower price. On the other hand, if the CCN has higher intrinsic costs than the incumbent, the CCN will be unable to compete when the incumbent begins providing service, and competition will force them out of the market. Traditional telcos will almost certainly be able to undercut CCNs in the most lucrative markets, urban and suburban areas, protecting their own business.

Finally, because NGSM checks with a database, regulators have control over where secondary users can operate. Spectrum sharing could be explicitly disallowed in “non rural” areas. This level of control and oversight should further mitigate competitive concerns from primary users.

B. Other Stakeholders

The two other major stakeholders with an interest in the regulation of GSM whitespaces are community cellular network operators and regulators.

1) Community Cellular Operators: Community cellular networks have a range of options for how they deal with spectrum licensing besides GSM whitespaces. Small, isolated networks may choose to operate as pirates under the expectation that they are beyond the reach (or interest) of regulators; the Papuan CCN took this route. The Oaxacan CCN operates under a two-year experimental license, but while this enables legal operation in the short term, they have no guarantee their license will be extended in the future. Their license also restricts them from making profit with their network, which hinders their efforts to provide sustainable and reliable service to their users. While we are unaware of any examples, a CCN could also partner with a carrier and operate under their license or simply buy a commercial spectrum license outright. As a small entity, obtaining a carrier partnership or a standard commercial license is out of the reach for most CCNs. Without proper incentives and risk mitigation factors like those outlined in Section VI-A, carriers have little reason to cooperate with CCNs; even finding an audience with a carrier to discuss a partnership is challenging for small entrepreneurs in rural areas. Commercial licenses can cost millions of dollars, well beyond the budget of a CCN.

Our proposal for spectrum sharing in GSM whitespaces represents a middle ground for CCNs with a range of attractive properties. Unlike pirate operation, it would allow CCNs to operate “above the board”, reducing risk to both the operator and the long-term sustainability of a CCN while at the same time maintaining the flexibility and independence of pirate operation. Experimental licenses have similar drawbacks to pirate operation—while legal, CCNs still face the risk of shut-down should their temporary license not be renewed. Explicitly supporting commercial operation indefinitely is necessary for incentivizing local entrepreneurs to operate CCNs so they can confidently plan to recoup their initial investments.

2) Regulators: First and foremost, by enabling the operation of CCNs regulators can expand access to service in rural areas. This fits the social mission of regulatory bodies to ensure that telecommunications access is available to their nation’s citizens and that spectrum, a vital public resource, is used equitably and efficiently.

GSM whitespaces give regulators a fundamentally new tool by which to achieve this mission. Today, decisions to build out rural infrastructure rest solely with incumbent license holders, and regulators are only able to indirectly influence these decisions through mechanisms like universal service obligations. Spectrum allocated for GSM networks is poorly utilized in rural areas: beyond the coverage area of existing cellular carriers, exclusively licensed spectrum is simply not used. Telcos are inherently disinclined from serving rural areas due to high costs of service and low revenue potential (due to low subscriber density), so this spectrum lies fallow. Existing mechanisms for incentivizing carriers to serve rural areas, such as universal service obligations, have high overhead and are fraught with political baggage.

In contrast, light regulation of GSM whitespaces, as we propose, allows local entrepreneurs to operate small-scale community cellular networks without requiring regulators to engage in expensive oversight of these operations. Moreover, since pirate operation is currently among the most attractive “licensing” options for CCN operators, providing a low-touch mechanism for these operators to register and regularly report on their spectrum usage gives regulators control over an emerging trend: rural, isolated communities have a strong demand for cellular communications, and it’s foolish to hope that existing penalties will prevent them from building their own infrastructure if they are able.

VII. EVALUATION

In Section V, we claimed that NGSM achieves all five goals for spectrum sharing in GSM whitespaces (Section IV). In this section, we justify that claim by implementing and evaluating NGSM. In addition to testing in a controlled environment, we also deployed NGSM into a real-world, operational CCN in Papua, Indonesia.

A. Implementation

We implemented NGSM as a software control layer based on OpenBTS [20], which uses a flexible software-defined radio and a commodity PC to implement a GSM base station. We support both dual-ARFCN, with one software radio per ARFCN. We also support single-ARFCN operation, with the

Variable	Value	Purpose
R	0	For ARFCNs <i>not</i> in use by the secondary, the RXLEV threshold for determining whether an ARFCN is in use.
K	0	Number of reports exceeding R needed to declare an ARFCN unsafe.
T	90	Number of seconds between ARFCN switches for the two ARFCNs in use by the BTS.

TABLE II. PARAMETERS FOR NGSM.

limitation that a single-ARFCN BTS is unable to detect other users of its own channel. Our implementation of NGSM monitors all control traffic between OpenBTS and phones, including measurement reports, and configures the ARFCN and neighbor list used by OpenBTS as appropriate.

In particular, we randomly select 5 ARFCNs for phones to scan every N hours. All ARFCNs are initially considered “unsafe”, save the ones initially in use. In our implementation, we randomly pick initial ARFCNs, though we expect a wider deployment might also be able to use a database query to pick the initial ARFCN. Once we receive measurement reports indicating an ARFCN is not in use, we consider the ARFCN “safe”. An ARFCN remains safe as long as we receive no more than K reports indicating an RSSI⁵ on that ARFCN exceeding R . Once these thresholds are exceeded, the ARFCN is demoted to being unsafe; in our implementation, once an ARFCN was demoted to being unsafe it remained so for at least $4N$ hours, at which point it could be scanned again and marked safe. ARFCNs that had been used by the BTS were similarly marked as unsafe once they had been used to ensure the BTS would use a different ARFCN each time it switched channels.

In our implementation, we chose a cycle length N of four hours to allow the BTS to scan a quarter of the GSM900 band every day, though this was itself chosen arbitrarily. We set $K = R = 0$ to be as conservative as possible in detecting other users of the band. These values are essentially the sensing threshold for the system; we leave full consideration of how to set these values appropriately to future work.

Given our scanning results, we select a pair of safe ARFCNs for the BTS to use. We alternate use of each ARFCN every $T = 90$ seconds by adjusting the TX attenuation on each software radio between 0 and 100 dB. One ARFCN operates without attenuation, while other operates at high attenuation. To change which ARFCN is active, we gradually increase the attenuation of the current ARFCN while reducing the attenuation on the inactive ARFCN to 0 dB. As the attenuation on the former ARFCN increases, phones automatically handover to the latter; from the phone’s perspective, we’ve simulated moving away from one cell and towards another.

If any of our in-use ARFCNs become unsafe, we immediately cease use of that ARFCN and switch to different safe one. This doesn’t result in any service disruption since we can only detect use of one of our own ARFCNs while that ARFCN is fully attenuated. To remain spectrum flexible, we select a new pair of ARFCNs to use every night. Additionally, if the BTS is restarted for any extrinsic reason (such as a power failure), we also pick a new pair upon restarting.

We deployed our implementation of NGSM in both a

controlled environment in our lab for testing, as well as the Papua CCN referred to earlier. We began running NGSM in Papua on October 14, 2013. We ran the system for testing for one week before we began collecting data, which we did from October 22 through November 1, 2013.⁶

As stated before, we used two independent software radios for dual-ARFCN operation. This was solely due to time constraints, and we stress that there is no fundamental reason our approach requires two radios: although implementing support for multiple channels on a single radio would require more engineering effort, it’s within the hardware capabilities of existing software radios. We are currently in the process of adding single-radio operation to our prototype and see no major engineering obstacles ahead of us. An unfortunate consequence of this is that due to hardware limitations in the Papua CCN, we were not able to deploy the full dual-ARFCN version of NGSM and were constrained to a single ARFCN. This also means we were unable to detect potential interference on the ARFCN currently in use by this BTS. We were, however, still able to collect measurement reports and scan the band for other users on that network, and we still changed the ARFCN the BTS used at least once per day.

Finally, although all of our experiments were conducted in the GSM900 band, our implementation could be easily modified to operate on any other band used for GSM. Our implementation of NGSM is available at <http://github.com/shaddi/gsmws>.

B. Coexistence

The ability to detect and respond to potential interference is a crucial requirement for NGSM. To demonstrate our ability to do this, we set up two BTS units. The first was a standard, unmodified GSM BTS, configured to simulate a “primary” user’s BTS broadcasting on a single ARFCN. The second ran NGSM with two ARFCNs as outlined in Section V, simulating a BTS run by a “secondary” user, a CCN. We also configured three phones as customers of the secondary BTS. The primary BTS used the same ARFCN as the secondary, but its other parameters (such as network ID) were distinct from those of the secondary: to phones, the secondary and primary BTS units appear to belong to two completely separate network operators. Each BTS was configured to transmit at 100mW per ARFCN. Fig. 3 shows the layout of the two BTS units and the 3 phones on in our testing environment, a single floor of an office building. Additionally, we placed two spectrum analyzers next to the middle phone, tuned to both the downlink and uplink bands used by the three BTS in this experiment.

We started NGSM on the secondary BTS and began alternating between its two ARFCNs. We waited for the three phones to camp to the secondary BTS and begin transmitting measurement reports, simulating a CCN operating in a steady state (i.e., with phones camped to its tower, but not necessarily in use). One phone, the middle one, had an ongoing call to the BTS. We then turned on the primary BTS to simulate the appearance of a primary in the vicinity of the secondary.

Fig. 4 shows the results of this test in the uplink band. This figure shows the usage of the phone on a call while

⁵RSSI is defined in GSM 04.08; specifically, we use the RXLEV-NCELL value, defined from 0 (−110 dBm) through 63 (> −47 dBm).

⁶NGSM is still running today; we would plan to include a longer dataset in a final camera-ready submission.




Fig. 3. Floor layout showing the primary (green circle, right), CCN (red star, left), and three cellphones (blue squares). The monitoring site was co-located with the middle phone.




Fig. 4. Spectrum usage of a handset during a call in the uplink band (i.e., from the phone to the BTS). The phone switches ARFCNs without interrupting the ongoing call.

the BTS alternates between ARFCNs. As expected, the phone completes handover successfully and the call continues without interruption. In Fig. 5, we see the spectrum usage on the down-link band during a simulated appearance of a primary user. Initially, the secondary BTS is alternating between ARFCNs 20 and 30. The primary appears on ARFCN 30 halfway through the experiment. Detecting this, the secondary BTS picks a new, unused ARFCN to use instead of ARFCN 30 (in this case, ARFCN 40). The secondary then begins alternating between ARFCNs 20 and 40, while the primary continues operation on ARFCN 30 without interference.

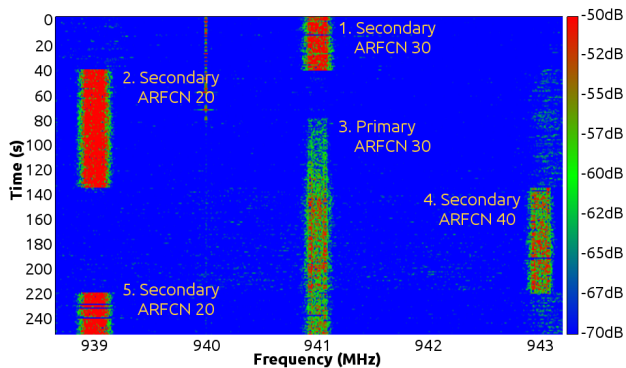


Fig. 5. Spectrum usage during simulated arrival of a primary into the secondary's service area. Initially, the secondary is alternating between using ARFCNs 20 and 30 (939MHz and 941MHz, respectively). When a primary BTS appears on ARFCN 30, the secondary detects its presence and switches to using ARFCN 40 instead (943MHz). Thereafter, the secondary alternates between ARFCNs 20 and 40.




Fig. 6. CDF of measurement report interarrival times received during daytime hours (7AM–12AM). Night time hours are not included since the CCN was typically powered off at night. $\mu = 0.64$, $\sigma = 4.02$. Note logarithmic scale.

Although this test may seem simple, it demonstrates a few important points. First, using handsets to detect interference is possible, without making any modifications to existing GSM infrastructure like phones or network infrastructure of existing carriers. Second, even when serving a handful of phones, the secondary BTS is able to detect interference and adapt its usage accordingly within a matter of minutes. Finally, it shows that secondaries can rapidly change channels upon detecting interference with minimal service disruption.

C. Measurement Reports

The time a secondary takes to detect a primary is inversely proportional to the frequency of measurement reports. Although phones constantly send measurement reports when in active use (e.g., during a call or when receiving an SMS), they only do so once every six minutes otherwise. Thus, measurement report frequency is directly related to the number of users a CCN has and how active those users are. In other words, we have measurements exactly when we need them: the potential for harmful interference is higher from an active network with many handsets, which in turn will have more frequent measurements.

We evaluate this empirically with our deployment in the Papua CCN, which has over 200 subscribers, more than 70 of which are active each day. Although we could not directly evaluate speed of detection—the operators of the CCN don't have the equipment necessary to replicate the previous experiment, namely a second BTS unit with which to simulate the appearance of a primary—we can evaluate the frequency of measurement reports the Papua CCN BTS receives. To do this, we logged every measurement report received by the Papua CCN from October 22 through November 1. A small, consistent interarrival time is valuable because it allows us to put bounds on how quickly a secondary can detect the primary.

We received approximately 846,000 measurement reports during our 10 days of operation. Of these, we only consider those received during “daylight” hours between 7AM through 12AM; because the CCN typically turns off at night, outside these hours we do not receive measurement reports. This removes about 12,000 reports (1.5%) from our analysis. With this in mind, Fig. 6 shows the distribution of interarrival

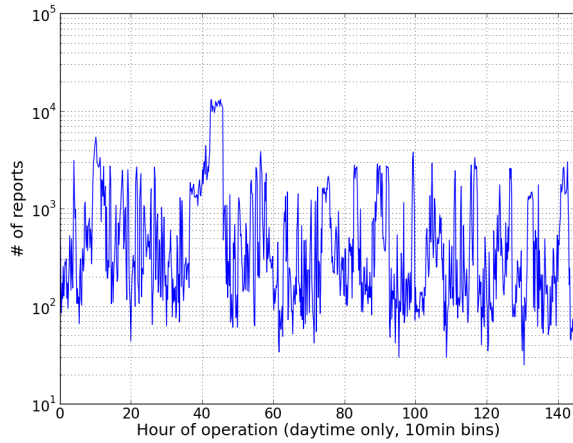


Fig. 7. Time series of number of measurement report received per hour during daytime hours (7AM–12AM). Night time hours are not included since the CCN was typically powered off at night. Note logarithmic scale.

times between measurement reports received over the course of our two week deployment. The maximum spacing between measurement reports seen during daytime hours was 11.7 minutes; the 99.9th percentile was 56 seconds. This result suggests real-world CCNs will enjoy faster detection times than we observed in our lab.

Finally, measurement reports arrive consistently while the BTS is operational. Fig. 7 shows the number of measurement reports received during the operational hours of the CCN, collected in 10 minute bins. The minimum number of reports received in any 10 minute window was 25; the median was 300 (obviously, when the BTS was off or not in use at night there were periods in which no reports were received). Combined with the previous distribution of report interarrivals, this demonstrates that we can rely on receiving regular measurement reports, placing an upper bound on the time to detect a primary’s BTS on the order of minutes during normal usage.

D. Deployment

NGSM operated as expected when we deployed it onto the Papua CCN. Fig. 8 shows the measurement results from the deployment. In this figure, the in-use ARFCN is blue, while ARFCNs considered “safe” or “unsafe” are colored green and red, respectively. During the experimental period, the operator’s primary source of electric power failed, causing several prolonged outages. Nevertheless, the CCN switched ARFCNs frequently, as designed. We were also able to verify through measurement reports that many ARFCNs were available for use around the Papua CCN, even when using the most sensitive detection thresholds ($K = R = 0$).

Despite these frequent channel changes, we observed no negative impact on network usage after deploying our system. Table III shows the distribution of network usage metrics per day before and after the deployment. We only consider calls and SMS initiated by users of the CCN; incoming communication is not included in these statistics. Active users refers to the number of subscribers who initiated either a call or an SMS that day. A number of factors—the aforementioned power failures, natural variation in usage (e.g., people travelling), etc.—preclude statistical testing, but we observe that usage remains roughly the same in terms of active users and SMS,

Metrics	Pre-NGSM		With NGSM	
	Mean	Deviation	Mean	Deviation
Calls	95.1	60.5	138.4	65.4
SMS	656.5	113.5	633.7	147.6
Active users	62.8	4.5	62.4	7.4

TABLE III. USAGE PER DAY IN THE PAPUA CCN, BEFORE IMPLEMENTING NGSM (09-09-2013 TO 10-09-2013) AND AFTER IMPLEMENTING NGSM (10-10-2013 TO 11-08-2013). DEPLOYING NGSM DID NOT SIGNIFICANTLY IMPACT USAGE.




Fig. 8. Spectrum usage and reports for the Papua CCN from Oct 22 through Nov 1. Blue represents the ARFCN in use by the CCN. Green represents an ARFCN that was scanned and considered “safe” to use, and red represents an ARFCN which was scanned but appeared to be in use, and was thus “unsafe”. Unsafe ARFCNs could become safe ARFCNs once they had been scanned again and found to not be in use. Empty columns represent times when the CCN was out of service; the CCN operator’s primary source of electric power failed during the evaluation period, causing prolonged nightly outages. Carrier allocations are shown on the right axis; the system jumps between all three.

and actually increases for number of calls. This is not a surprising result—we designed our system to only change ARFCN during periods of little or no activity to avoid impact on usage. Nevertheless, it shows that a CCN can operate effectively even when it changes its ARFCN relatively frequently.

We expected the GSM900 band to be completely unused around the Papua CCN, as the network operators informed us that the nearest cellular service was almost 30 kilometers away, beyond several mountainous ridges (the Papua CCN itself is located in a small valley). In general, we found this to be the case, but there were a few interesting exceptions.

Somewhat surprisingly, we detected a usage of several ARFCNs during our deployment, many of which were licensed to carriers who do not provide any service in the Papuan highlands. For example, on October 26 the BTS received 19 reports over a 2 hour period indicating ARFCN 50 was in use. The Papua CCN’s BTS performed as designed and did not use those ARFCNs going forward. Unfortunately we have no way of knowing what may have caused these reports, nor can we necessarily discount the possibility they were simply spurious reports. However, this highlights a crucial point: spectrum regulations might already be flouted in rural areas, and regulators (and licenseholders) have no way to detect these violations until they actually interfere with operations.

VIII. DISCUSSION

A. Market alternatives to GSM whitespaces

The obvious market-based alternatives to GSM whitespaces is to have the CCN operator enter into a contractual relationship with the license-holder of the spectrum (as a franchisee of sorts or an alternative local brand) or to engage in a transaction in secondary spectrum market to obtain local usage rights in that area. While conceptually tempting and simple, there are problems. The empirical evidence of such transactions not actually happening in the real world suggests that something must indeed be wrong with this approach.

1) *Individually-negotiated contracts:* The theory of Coasian bargaining says that the problem must be in transaction costs. Indeed, it is hard to imagine a carrier engaging a lawyer and engineers to travel to a remote area to negotiate and implement/verify a contract to split some profits that amount to a mere \$1000 per month, even without factoring in the uncertain future squatting risk that they might feel they face from having the CCN operate solely in their spectrum. The NGSM approach to GSM whitespaces here eliminates that transaction cost and the uncertainty.

2) *Standardized markets:* In principle, a clean online secondary market could also eliminate some of these transaction costs. However, there are subtle issues here. First, what is being traded? There are two possibilities: spectrum (where CCNs pay spectrum license holders to be able to deploy their systems) or USOs (where CCNs bid to accept USOs in some area).

Second, once there are market transactions in an asset or liability, the asset/liability can be quantified in dollar terms. Hence by accounting principles, the entire asset or liability must be quantified in dollar terms on the firm's books. For spectrum, this could challenge the (speculative) high valuations that firms carry on spectrum to serve as collateral for loans, etc. For USOs, it would suddenly cause USO obligations to show up as dollar-valued liabilities as opposed to vague risks.

The net effect on the books of having such secondary market transactions is likely to be negative and so the managers of firms are disinclined to explore such markets/transactions. The NGSM approach to GSM whitespaces avoids having any dollar transactions by directly getting USO credit from the regulator for the actions of a third party.

B. Application to urban and marginal coverage areas.

Our GSM whitespace scheme is designed for rural areas. It assumes large amounts of unused spectrum and relies upon spatial separation to avoid interference. We do not believe it is a good fit for situations where spectrum is scarce and highly utilized. This is not a problem in our minds: existing spectrum allocation policy has proven to be adequate for ensuring widespread cellular coverage in urban areas.

Compared to areas completely beyond coverage, an even larger portion of the planet's population likely lives in areas with marginal cellular coverage: while they may be able to access a cellular network, coverage may be sporadic or otherwise spotty. These areas, like the completely unserved areas beyond them, are likely to have GSM whitespaces available, but the potential interactions between primary and secondary license

holders will be more complex. We leave a full consideration of these issues for future work. In the meantime, the simpler case of spectrum sharing in completely unserved areas is tractable and should be considered by policymakers in the short term.

C. Distribution of clients

Our proposed interference detection mechanism relies on reports from handsets and thus is sensitive to their geographic distribution. In the degenerate case, all phones could be clustered in a small area and unable to detect other networks within range of our BTS site. This concern is mitigated partially by the fact that both primary and secondary users will be trying to serve the same people in a rural area, so the physical distribution of a secondary's user base is likely to be correlated with that of the primary. Additionally, the GSM standard provides a mechanism for obtaining geolocation information from phones, the radio resource location services protocol (RRLP). This information could be used in conjunction with measurement reports to identify potential "blind spots" where the secondary is unable to detect interference. Beyond simply shutting down or requesting guidance from the regulator database, a CCN can take a number of creative actions upon detecting such a blind spot. For example, the CCN could automatically send an SMS to a user requesting them to wander over to the blind spot area, perhaps incentivizing them with free network credit.

IX. CONCLUSION

Rural areas are fundamentally hard for traditional telcos to serve profitably, leaving hundreds of millions of people beyond the reach of existing cellular phone networks for structural reasons. While community cellular networks appear to offer substantial advantages for providing sustainable rural service without subsidies or external support, their growth is stymied by a lack of rights to spectrum. However, exclusive spectrum licensing has created large areas of GSM whitespaces, areas in which GSM spectrum is allocated to a carrier but not actually used, as is the case in many rural areas worldwide.

This spectrum need not be wasted: we believe it represents an opportunity to enable community cellular networks to provide service in rural areas. In this paper, we've proposed Nomadic GSM, a spectrum sharing technique for GSM whitespaces to leverage this opportunity. NGSM uses a combination of a spectrum database and a novel distributed spectrum scanning technique, leveraging the reporting capability of mobile phones, to ensure rapid detection of potential interference. Our proposal allows CCNs to safely share spectrum with incumbent carriers without their explicit cooperation, while mitigating key concerns that licenseholders might have with sharing their spectrum. By reporting spectrum measurements to a database, it enables regulators to verify what spectrum is actually in use in an area so that carriers can receive USO credit.

NGSM is compatible with existing, unmodified GSM phones: we've demonstrated its feasibility with both a prototype implementation in our lab as well as a real-world deployment on an existing community cellular network in Papua, Indonesia. We've demonstrated that with 70 daily active users, we are able to receive a measurement report at worst every 11.7 minutes while the BTS was on, with a 99.9th percentile interarrival of 56 seconds.

The implications of our system are important. Rural communities will build their own community cellular networks in increasing numbers, many of which will be operating in GSM whitespace. The situation is akin to that of WiFi in countries that had not yet adopted policies allowing unlicensed spectrum use—strong demand compels community cellular network operators to flout regulations and operate illegally, outside the control of regulators and at risk of interfering with the operation of licenseholders and other CCNs. Unlike WiFi, however, CCNs are still in their infancy, and an enlightened regulatory approach towards them can allow countries to maximize their benefits for providing rural service while mitigating impact on other users of the GSM bands.

Our proposal has attractive properties for achieving this, and, importantly, is deployable today, requiring no changes to existing mobile phones, network infrastructure, or operational practices of incumbent network operators. As such we feel it represents a strong first step towards a comprehensive policy for enabling legal coexistence of community cellular networks. We suggest that regulators take the following steps:

- Legalize use of GSM whitespaces with requirements that CCNs using them regularly (a) move between unused frequencies and (b) use NGSM (or similar) to monitor local GSM frequency and avoid causing interference.
- Facilitate creation of a GSM whitespace reporting database.
- Give carriers USO credit for CCNs operating in their spectrum allocations using the database reports as evidence for such claims.

REFERENCES

- [1] 3GPP. Radio subsystem link control. TS 05.08, 3rd Generation Partnership Project (3GPP), Nov. 2005.
- [2] A. Agüero, H. de Silva, and J. Kang. Bottom of the Pyramid Expenditure Patterns on Mobile Services in Selected Emerging Asian Countries. *Information Technologies & International Development*, 7(3), 2011.
- [3] Q. Alam, M. A. Yusuf, and K. Coghill. Village Phone Program, Commodification Of Mobile Phone Set And Empowerment Of Women. *J. Community Informatics*, 5(3-4), 2009.
- [4] V. Bharghavan, A. Demers, S. Shenker, and L. Zhang. MACAW: a Media Access Protocol for Wireless LAN's. In *Proceedings of the Conference on Communications Architectures, Protocols and Applications, SIGCOMM '94*, pages 212–225, New York, NY, USA, 1994. ACM.
- [5] C. R. Blackman. Universal Service: Obligation or Opportunity? *Telecommunications Policy*, 19(3):171–176, 1995.
- [6] M. M. Buddhikot, I. Kennedy, F. Mullany, and H. Viswanathan. Ultra-Broadband Femtocells via Opportunistic Reuse of Multi-Operator and Multi-Service Spectrum. *Bell Labs Technical Journal*, 13(4):129–143, 2009.
- [7] D. Cabric, S. M. Mishra, D. Willkomm, R. Brodersen, and A. Wolisz. A Cognitive Radio Approach for Usage of Virtual Unlicensed Spectrum. In *In Proc. of 14th IST Mobile Wireless Communications Summit 2005*, 2005.
- [8] P. Chon, L. Flochel, and A. Perrot. Allocating and Funding Universal Service Obligations in a Competitive Market. *International Journal of Industrial Organization*, 20(9):1247 – 1276, 2002.
- [9] T. Egyedi and D. Mehos. *Inverse Infrastructures: Disrupting Networks from Below*. Edward Elgar, 2012.
- [10] C. Fischer. *America Calling: A Social History of the Telephone to 1940*. University of California Press, Berkeley, CA, 1992.
- [11] H. Galperin and F. Bar. The Microtelco Opportunity: Evidence from Latin America. In *Information Technologies and International Development*, volume 3, Winter 2006.
- [12] GSM Association. Powering Telecoms: West Africa Market Analysis Sizing the Potential for Green Telecoms in Nigeria and Ghana. <http://www.gsma.com/mobilefordevelopment/wp-content/uploads/2013/03/GPM-Market-Analysis-West-Africa-.pdf>. Retrieved 10/2013.
- [13] GSMA. Mobile Development Intelligence, Network Coverage Population. <https://mobiledevelopmentintelligence.com>.
- [14] K. Heimerl, K. Ali, J. Blumenstock, B. Gawalt, and E. Brewer. Expanding Rural Cellular Networks with Virtual Coverage. In *Proceedings of the 10th USENIX Symposium on Networked Systems Design and Implementation, NSDI'13*, Lombard, Illinois, 2013. USENIX Association.
- [15] K. Heimerl, S. Hasan, K. Ali, P. Bloom, J. Blumenstock, E. Brewer, and T. Parikh. Designing Community Cellular Networks. In *In Submission*.
- [16] K. Heimerl, S. Hasan, K. Ali, E. Brewer, and T. Parikh. Local, Sustainable, Small-Scale Cellular Networks. In *Proceedings of the Sixth International Conference on Information and Communication Technologies and Development, ICTD '13*, Cape Town, South Africa, 2013. ACM.
- [17] R. Jensen. The Digital Divide: Information (Technology), Market Performance, and Welfare in the South Indian Fisheries Sector. *The Quarterly Journal of Economics*, 122(3):879–924, 08 2007.
- [18] A. Joshi and M. Moore. Institutionalised Co-production: Unorthodox Public Service Delivery in Challenging Environments. *The Journal of Development Studies*, 40(4):31–49, 2004.
- [19] W. F. Lam. Institutional Design of Public Agencies and Coproduction: A Study of Irrigation Associations in Taiwan. *World Development*, 24(6):1039–1054, June 1996.
- [20] OpenBTS. <http://openbts.org>. Retrieved 7/2013.
- [21] J. M. Peha. Sharing Spectrum through Spectrum Policy Reform and Cognitive Radio. *Proceedings of the IEEE*, 97(4):708–719, 2009.
- [22] Rhizomatica. <http://rhizomatica.org/>. Retrieved 4/2013.
- [23] S. Sankaranarayanan, P. Papadimitratos, A. Mishra, and S. Hershey. A Bandwidth Sharing Approach to Improve Licensed Spectrum Utilization. In *New Frontiers in Dynamic Spectrum Access Networks, 2005. DySPAN 2005. 2005 First IEEE International Symposium on*, pages 279–288. IEEE, 2005.
- [24] Sensorly. <http://sensorly.com>. Retrieved 11/2013.
- [25] D. Setiawan. Overview of Spectrum Issues in Indonesia. <http://denysetia.files.wordpress.com/2010/09/290610-broadband-wireless-gsma-london-denny-ins.ppt>, June 2010.
- [26] M. Spiegel. Universal service obligation. In M. A. Crew, editor, *Expanding Competition in Regulated Industries*, volume 37 of *Topics in Regulatory Economics and Policy Series*, pages 83–96. Springer US, 2000.
- [27] Telkomsel. Telkomsel, 2006 Annual Report. <http://www.telkomsel.com/media/upload/pdf/AR2006.pdf>, 2007.
- [28] Telkomsel. Telkomsel, 2012 Annual Report. <http://www.telkomsel.com/media/upload/annualreport/%28FINAL%29%20AR%20TELKOMSEL%20CONTENT.pdf>, February 2013.
- [29] I. T. Union. Implementation of USO Program in Indonesia. <http://www.itu.int/ITU-D/.../Indonesia/Session2-BP3TI.pdf>, March 2011.
- [30] L. Waverman, M. Meschi, and M. Fuss. The Impact of Telecoms on Economic Growth in Developing Countries. *Africa: the Impact of Mobile Phones*, 2(1980):1–23, 2005.
- [31] D. Weller. Auctions for Universal Service Obligations. *Telecommunications Policy*, 23(9):645 – 674, 1999.
- [32] M. E. Whitman. A Look at the Telecommunications ACT of 1996. *IS Management*, 14(3):82–86, 1997.
- [33] L. Yin, K. Wu, S. Yin, S. Li, and L. M. Ni. Reuse of GSM White Space Spectrum for Cognitive Femtocell Access. In *Parallel and Distributed Systems (ICPADS), 2012 IEEE 18th International Conference on*, pages 1–7. IEEE, 2012.
- [34] M. Zheleva, A. Paul, D. L. Johnson, and E. Belding. Kwiizya: Local Cellular Network Services in Remote Areas. In *Proceedings of the 11th International Conference on Mobile Systems, Applications, and Services, MobiSys '13*, Taipei, Taiwan, 2013. ACM.