Title
Towards Scalable Community Networks

Author
Hasan, Shaddi

Publication Date
2019

Towards Scalable Community Networks

by

Shaddi Husein Hasan

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Computer Science

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Eric Brewer, Chair
Associate Professor Joshua Blumenstock
Professor Scott Shenker
Professor Anant Sahai

Spring 2019

Towards Scalable Community Networks

Abstract

Towards Scalable Community Networks

by

Shaddi Husein Hasan

Doctor of Philosophy in Computer Science

University of California, Berkeley

Professor Eric Brewer, Chair

Over 400 million people live without access to basic communication services, largely in rural areas. Community-based networks, and particularly community cellular networks, can sustainably support services even in these extremely rural areas where traditional commercial network operators cannot. However, community cellular networks face a variety of technical, business, and regulatory challenges that hamper their proliferation.

In this thesis, we aim to develop approaches to enable scale within and across community cellular networks. Through a mixed-methods study of more than 80 rural wireless Internet service providers, a related class of network operators, we identify key scaling challenges rural network operators face. We next present two approaches for addressing these challenges in the context of community cellular networks. The first, GSM Whitespaces, demonstrates that rural community cellular networks can safely share spectrum in bands occupied by incumbent mobile network operators, removing a key barrier to independent operation. The second, CCM, shows how community networks and incumbent mobile network operators can cooperate to share resources to extend service. We will explore each of these approaches through practical systems and longitudinal deployments of community cellular networks in Southeast Asia that provide service to thousands of rural people.

For the people doing the work to make the community networking vision a reality.

Contents

List of Figures

List of Tables

Acknowledgments

I have been told that a dissertation is an individual work, but in putting this one together I have never felt alone: I am lucky to have had a wonderful and supportive community of collaborators, friends, and family accompanying me throughout this journey.

I thank Eric Brewer, my advisor, for his unyielding optimism about my work and enthusiastic mentorship throughout my time at Berkeley. I am also grateful to Scott Shenker for his steadfast support and encouragement throughout the different phases of my time in graduate school; Anant Sahai for his clear insights and fruitful interdisciplinary collaboration; and Joshua Blumenstock for his perspectives and work to connect this work to direct social impact. I am thankful to all of them for serving on my committe and for their thoughtful advice and feedback on this work.

I am also grateful to those who mentored and supported me prior to graduate school and who introduced me to the joys of research and teaching: Jay Aikat, Kevin Jeffay, Henry Fuchs, and Lynne Parker. I would not be here without your mentorship: thank you.

I have been blessed with many excellent collaborators in my time in graduate school. I thank my co-authors, Kurtis Heimerl, Kashif Ali, Kate Harrison, Mary Claire Barela, Matthew Johnson, Sean Roberts, Yahel Ben-David, Max Bittman, and Barath Raghavan for being such wonderful collaborators and friends. Similarly, the TIER folks – Kurtis, Kashif, Yahel, Matt Podolsky, Javier Rosa, Jordan Freitas, Mellisa Densmore, and all the others who have visited our group – have provided a rich and supportive intellectual community for my work. I am grateful for what I've learned from each of you.

I consider myself so lucky to have crossed paths with Kurtis and Kashif in particular, two dear friends who were crazy enough to try to take our community cellular work out of the lab and into the world. To them, as well as the others who shared the Endaga journey with us – especially Lance Condray, Omar Ramadan, and Matt Ball – what a long, strange, and rewarding trip it's been. Tower to the people!

My colleagues at Facebook Connectivity have pushed me to focus on impact while considering the challanges of rural connectivity more deeply. Thank you to Roger Greene, Amar Padmanabhan, Praveen Kumar Ramakrishnan, Emre Tepedelenlioglu, Subbu Subramanian, Ashish Kelkar, Meghana Sahasrabudhe, Neil Bartlett, Philip Liddel, Cait Campos, Nathan South, Alan Norman, Sam Perales, Steve Muir, Evgeniy Makeev, and the many others who I've bene fortunate to work with for your insight and collaboration over the years.

One of the most rewarding aspects of my time in graduate school was the opportunity to work with so many inspiring people building rural community networks. I thank Steve Song, Peter Bloom, Ciaby, Carlos Rey-Moreno, and David Haag for expanding how I think about community networks. Similarly, it was a joy working with all my collaborators who were necessary to make these deployments possible, including the Radio Jënpoj de Tlahuitoltepec crew, the team at Globe Telecom, the University of the Philippines VBTS team, and all the anonymous people throughout Indonesia and the Philippines who did the real work of

building these networks. Similarly, I am grateful for all the WISP operators who took the time to speak with me about their networks. I hope that this work helps advance in some way the hard work that these communities are carrying forward.

This work was supported by a variety of sources, including Endaga, Facebook, the Philippines Commission on Higher Education's PCARI program, USAID, and the NSF GRFP. Thanks!

To the extended 1044 crew (Dave Moore, Allie and Ryan Janoch, Pat and Caitlin and Molly Virtue, Justine Sherry, Paul Pearce, Judy Hoffman, Jonathan Long, Jono Kummerfeld, Ellen Stuart, Greg Durrett, Emily Cogsdill, Michael Cole, and Meghan Kelly), the Chaat crew (Matt Ball, Nora Vivanco, Tharu, Rishi Gupta, Zina and Otavio Good, Trevor Shannon, and Katie Dektar), and all the other Bay Area folks (Maggie Zhou, Katherine Elliott, Buro Mookerji, Gilbert Bernstein, Andy Stock, Christina Felschen, Laura Bloomfield, Ronan Arthur, Fran Moore, Colin Scott, and all the others): thanks for being such awesome friends. As many of you know, getting through grad school is impossible without the support of good friends, and you're all good people.

To my family – thank you for your support. Adam, thank you for being such an awesome brother and travelling buddy (and for helping assemble all those spectrum scanners that time you came to visit for your spring break!). Mom and dad, thank you for instilling in me a love of learning and supporting and encouraging me as I pursued it.

Finally, to Elinor, my wife. You've been my best friend and most steadfast companion on this journey, and I'm grateful for your love, your support, your puns, and for the inspiration you've given me along the way. I'm looking forward to seeing what adventures we get ourselves into next.

Chapter 1

Introduction

This thesis examines how to build networks that provide equitable access to communications.

The last decades have witnessed the proliferation of the Internet to most people and places on earth. The remarkable resilience and flexibility of the Internet's underlying architecture in the face of this growth and the accompanying evolution of usage, more than forty years after the introduction of TCP/IP, is a well-cherished win of the networking community. The construction of the *physical* infrastructure that has enabled this growth–the backbone fiber and microwave links, the wired and wireless access networks–is equally impressive and important.

These physical networks are today a vital public resource, providing access to emergency services, creating positive economic impacts [169, 100], helping organizations and individuals operate more effectively [155], and fulfilling the basic human need to communicate. Spending on communication among the poor reflects that of a necessity [4], underscoring the importance of these networks. Among the different technologies used for this physical infrastructure, cellular networks are the primary way the vast majority of users first get access to the Internet: while the majority of *traffic* on the Internet still originates from WiFi networks, mobile phones are the dominant screen, particularly in the developing world [46]. Indeed, cellular networks are the largest communications systems on Earth, with over 3 billion unique users and a global economic impact in excess of US$10 trillion (3.6% of *global* GDP) [73].

Despite their importance, at least one billion people live outside the reach of any cellular network coverage, largely concentrated in rural areas [70]. In light of this, providing universal access to cellular service has emerged as an important policy objective worldwide and is a target of the United Nations Sustainable Development Goals [120]. The primary policy mechanisms governments use to induce carriers to provide rural service are through conditions on spectrum licenses (i.e., explicit coverage obligations), subsidies (often via universal service funds), and infrastructure-sharing requirements. Although these programs are important and have been widely adopted, they have failed to provide truly universal service. The economics of traditional cellular operators are driven by *subscriber density*: these operators can serve more people and produce more revenue per dollar of infrastructure investment in densely populated urban areas compared to sparsely populated rural ones. Even with subsidies, carriers are inherently disinclined to invest in rural infrastructure, which is marginally profitable at best, particularly given the opportunity cost of investing in lucrative 3G and 4G infrastructure in urban areas. Coverage obligations require reliable monitoring of infrastructure buildout and effective enforcement for failure to comply, both of which may be challenging for under-resourced regulators in poor countries. And while infrastructure sharing can reduce costs of deployment for operators while maintaining a competitive environment for users, some entity must still build physical infrastructure, and with multiple operators splitting a relatively small pool of revenue business cases are often challenging.

The fundamental problem with these historical approaches is that they assume a *top-down* model of cellular deployment. Historically this made sense: only large organizations with deep technical and business expertise were capable of building and operating cellular networks, and only a handful could exist in any particular country. In a similar way, this model has informed the entire regulatory and commercial environment in which cellular networks exist. Obtaining a spectrum license for a cellular network costs millions to billions of dollars and confers upon the license holder an exclusive right to operate a network in that spectrum over a vast geographic area. This leads directly to the structural challenges facing rural communities and preventing equitable and universal access to communications—yet recent innovations invalidate these assumptions.

Community Networking

A team of wireless enthusiasts gathers in a cluttered office in a city for their weekly meeting to flash new software onto a collection of old WiFi routers. An organizer stands in a crowded room explaining that new spectrum regulations have created an opportunity for this remote town to finally have mobile phone service. A student from the regional university returns to her field site to meet with the community she's been working with for months to help diagnose a problem with a new wireless link.

These scenes and others like them are familiar to veterans of the community networks movement. Community networks, defined as "networks built by and for the people which they are intended to serve" [153] have emerged as an alternative to traditional cellular operators as millions worldwide continue to be unserved by traditional communications providers, either due to lack of access to their networks or inability to afford their services [70, 8, 75].

Community networks in various forms have a long history, dating from the earliest days of telephone networks. Rural farming communities in the United States built their own party line phone networks using primitive telephone systems, some even repurposing barbed wire fencing as telephone lines to cut down on deployment costs [61]. Telephone cooperatives in Argentina have provided service to rural areas since the 1960's, with more than 300 active by the 1990's [65]. In more recent years, community wireless networks have leveraged low-cost WiFi devices operating at relatively low power in unlicensed spectrum to provide Internet access [64, 33, 141, 142]. Some particularly sophisticated networks have added fiber to their wireless networks [167].

The choice of technology in each of these networks is less important than the organizational principles that unify each of them. However, the technical choices made by community networks inform the services they can provide as well as the cost structures and regulatory considerations they face.

In the past decade, the cost and complexity of building and maintaining a cellular network has decreased to the point where individuals or rural communities can create their own micro-scale cellular networks [88, 143] that provide effective service while remaining financially sustainable. These *bottom-up* "community cellular networks" (CCNs) demonstrate that local communities can operate their own cellular telecommunications infrastructure and provide essential communication services, and present a promising new approach for providing sustainable access to communication for rural communities. In contrast to other community wireless networks, community cellular networks can cover an area with less infrastructure than WiFi-based networks, while providing service directly to mobile devices that are commonly owned by people worldwide. Today, communities in Indonesia [89], Mexico [143], Canada [126], and elsewhere are running successful, financially sustainable networks in this model. The commercial ecosystem around community cellular is maturing rapidly, with multiple vendors producing hardware and software tailored to small-scale and community network operators [51, 55, 140, 125, 16, 7]. Regulatory authorities themselves have also recognized the important role that community networks can play in meeting universal coverage targets [52, 94].

Despite the enthusiasm, however, community cellular networks in particular face a number of challenges: there are few, if any, examples of community cellular networks with more than a handful of nodes[1]. The largest challenge these networks face is the *lack of a regulatory framework* in which they can exist as independent networks. In most jurisdictions, obtaining a spectrum license to operate such a network legally is either outright impossible or prohibitively expensive. Another challenge is that the *complexity of cellular networks is a barrier for community operators* to design, build, and operate, particularly compared to IP networks. These are the two challenges we address in this thesis.

Part 1: Revisiting the Wireless ISP

As community cellular networks are only just emerging, we can better understand the challenges they will face by considering the challenges faced by other community-based network approaches. Existing literature on connectivity for developing regions focuses on repurposing WiFi equipment to build cost-effective infrastructure networks [134, 157]. In the past decade, the wireless Internet service provider (WISP) ecosystem that emerged around

---

[1] "Node" here generally refers to a cellular base station, or a collection of cellular base stations serving a settlement. Unlike WiFi-based community networks, a single cellular base station can provide service to thousands of people, depending on configuration, capabilities, and demand.

such technologies has rapidly matured. Today, an aspiring WISP operator has a range of commercial software and hardware produced by multiple vendors to choose from and (in some markets) access to WISP associations through which they can learn best practices from other operators. Thus, by considering the challenges of WISPs in these sophisticated markets, we gain insight into the challenges that community cellular networks are likely to face as the basic ecosystem matures.

To do this, we conduct a survey of WISPs in the United States, a highly developed WISP market. Our survey, conducted in 2012, included an online survey with over 75 WISPs operating in both urban and rural areas, as well as follow-up interviews with 12 participants, and sought to understand the operational challenges and barriers to scale that WISPs faced. We find that WISPs today are quite effective at managing and operating their networks, with the primary causes of network faults being due to physical infrastructure issues such as weather and power failure. In contrast to previous work, we find that the challenges these service providers face are largely non-technical, and center on access to adequate and appropriate frequencies. Rural operators in particular face scaling challenges around access to financing to support network expansion, as their networks require specialized equipment and thus high customer acquisition costs.

Part 2: GSM Whitespace

Unlike WiFi equipment, cellular networks use licensed spectrum, and gaining access to spectrum licenses is almost impossible for micro-scale rural operators in most markets. Licensed spectrum is allocated in frequency blocks over a geographic area to provide exclusivity for licenseholders. In practice, the most commercially important cellular spectrum is allocated on a national basis for assignment to operators (often, though not always, via an auction process). As a result, in order to serve profitable urban areas, operators also obtain exclusive access to rural spectrum, with spectrum laying fallow outside those operators' service areas.

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 area[2] and thus could be re-used by a secondary operator without interfering with the primary license holder. Our core insight is that unlike in dense urban areas, spectrum in rural areas is an abundant, not a scarce, resource: without network infrastructure, no spectrum is actually being used, and the needs of low-density rural areas can be met with less spectrum than denser urban areas. By enabling CCNs to operate in GSM whitespaces, regulators would empower rural communities to build infrastructure appropriate to their own needs, without waiting for incumbent operators to allocate resources their way.

Our approach represents a dynamic spectrum sharing approach to removing barriers to spectrum access for community cellular networks. We develop a set of requirements for sharing spectrum between incumbent cellular operators and community cellular network

---

[2] Note we do *not* refer to the guard band or space between adjacent GSM channels.

operators in rural areas, taking into account both the needs of primary users of cellular bands (in this case, incumbent operators) as well as community cellular networks. Unlike Television White Space (TVWS) [59], we require an approach that is backwards compatible with existing cellular standards and devices, while still providing safety and independence for primary users.

We demonstrate such an approach is feasible by implementing Nomadic GSM, and evaluate the system through both lab testing and field deployment in a real community cellular network. Nomadic GSM is able to detect the deployment of a new primary radio within 90 seconds of operation and move to new spectrum without disrupting user traffic, including sensitive traffic like active voice calls. The system adds a shim control layer to off-the-shelf mobile network radio infrastructure and requires no changes to user equipment (such as mobile devices). This allows community cellular network operators to take advantage of the existing mobile hardware ecosystem, reducing costs and deployment complexity. While Nomadic GSM specifically targets 2G GSM networks, the approach is generic to any mobile network that leverages network-assisted handover between cells, which is true of all widely deployed mobile network standards including LTE.

Part 3: Scaling Community Cellular Networks

GSM Whitespace provides access to spectrum for community cellular networks, but it requires significant regulatory changes: few countries have adopted any form of dynamic spectrum sharing, much less sharing in cellular frequency bands. Achieving scale in the short term will likely require additional approaches.

We begin by noting that, fundamentally, community cellular networks and incumbent operators actually have aligned interests in unserved rural areas. Incumbent operators benefit indirectly (in the form of increased usage revenue) when more people are on the phone network and can call their users, but they can benefit directly from improved rural connectivity in the form of reduced regulatory pressure to build out into unprofitable areas. Community cellular operators, in turn, can benefit from resources that incumbent operators already have, such as spectrum licenses, phone number allocations and interconnection, and backhaul infrastructure. By facilitating collaboration between incumbent operators and community cellular networks, we can enable community network operators to gain access to the resources they need to connect rural populations, without major reforms to how spectrum is managed: we reduce the problem from a regulatory one to a commercial transaction.

The key challenge is to ensure *isolation* between incumbent operator infrastructure and community network operators, ensuring safety for the incumbent and flexibility for the community network. Interconnection, in particular, requires integration with an incumbent operator's "core network", vital infrastructure for the stability and security of the mobile network. Sharing resources with community networks cannot require undue burden on the mobile network operator (MNO), so we must minimize transaction costs for new sites.

We present CommunityCellularManager (CCM) as an example of such a system, and evaluate its viability through a multi-year deployment in partnership with a large mobile network operator in the Philippines. CCM consists of a centralized, multi-tenant interconnection and management component, hosted by the MNO, and a distributed core network at each node of the community networks. This allows CCM to expose a collection of community networks through a single, unchanging point of interconnection with the mobile network operator, while providing community networks the flexibility to innovate on services offered to their users. CCM also automates much of the configuration and operation of the community networks themselves, simplifying deployment of these networks. Our system supports the largest community cellular network deployment to date (in terms of number of sites), providing connectivity to thousands of active users across 15 sites, as well as the first community cellular network deployment done using spectrum licensed to a commercial operator.

1.1 Summary of Results

In this dissertation, we consider the challenges that community networks, and the people who hope to build them, face, and present two systems to shape the technical landscape in which community cellular networks, in particular, can operate.

Using the experiences of WISPs in the United States as a model for understanding the challenges community cellular networks face in rural environments, we conduct the largest survey of wireless Internet service providers in the academic literature to understand their operational and technical challenges. We find that while the challenge of building sustainable rural infrastructure is largely a business one, dedicated spectrum and network automation, particularly at the boundary of business processes and network operations, are critical for enabling scale and sustainable operation. We further observe that physical layer failures, rather than network layer ones, are the primary challenge faced in rural operator networks, occurring both more frequently and creating more severe faults when they do occur.

We draw upon insights from the WISP industry to extract lessons for community cellular networks, developing two systems to remove barriers to entry and growth for the nascent community cellular networking ecosystem.

With **GSM Whitespaces**, we propose a new framework for dynamic spectrum sharing of unused spectrum between incumbent MNOs and community cellular networks. We show:

- Sharing spectrum between incumbent licenseholders and community cellular networks is feasible, and may be done while both preserving safety for all parties and requiring no party to act contrary to their incentives.
- A practical implementation of our framework, Nomadic GSM. Our implementation proves that GSM whitespace can be shared without changes existing mobile infrastructure, including user handsets. Because Nomadic GSM is a software control mechanism for existing GSM base stations, new community cellular network infrastructure requires no new hardware.

- Our implementation is able to detect the arrival of a new primary user in less than 90 seconds using handset-based sensing alone, and dynamically change frequencies without dropping user traffic.
- While our mechanism is specific to GSM networks, the underlying methodology is generic for any cellular technology that relies on network-initiated handover (i.e., all modern cellular technologies, including LTE).

Next, we present **CCM**, an end-to-end system for enabling community cellular network deployment at scale. CCM demonstrates:

- Cooperation and resource sharing between community networks and incumbent mobile network operators is both feasible and safe for mobile operators. CCM preserves flexibility for community networks to implement their own services, supports multiple administrative domains among community networks, and a single point of integration (and control) for allocating numbering, spectrum, and interconnect resources from a mobile operator to a collection of community networks.
- Enabling this collaboration can accelerate community cellular network deployments. We illustrate this through a three-year deployment using CCM to support 17 community cellular networks throughout the Philippines, in collaboration with Globe, the largest mobile network operator in the country. This represents the largest community cellular network deployment to date, providing service to 2,800 monthly active users who completed a total of more than 50,000 hours of voice calls and sent or received over three million SMS messages.
- The benefit of supporting autonomy and disconnected operation comes at minimal cost. CCM demonstrates a novel partially-decentralized mobile core network architecture that gracefully supports disconnected operation while preserving enforcement of network policies, including billing. Sites in our deployment were offline 65% of the time due to a combination of intermittent backhaul failure and power outages. By supporting disconnected local services, the networks were able to serve 19% more traffic than without, as well as 16% more local pre-paid credit sales. Our state synchronization mechanism required to support this use case incurs negligible overhead, consuming as much backhaul capacity as an additional 6 second voice call per minute.

Between the two approaches, we provide adoptable *regulatory* and *commercial* paths towards securing spectrum for community cellular networks.

1.2 Dissertation Outline

The remainder of this dissertation proceeds as follows. In Chapter 2, we perform a survey of small-scale WISP networks throughout urban and rural segments of the United States to understand the challenges they face operating and scaling their networks.[3] Chapter 3 introduces the GSM whitespace concept and presents Nomadic GSM, which enables community cellular networks to provide service on allocated and assigned, but unused, licensed spectrum. In Chapter 4, we present Community Cellular Manager, a system for end-to-end management of community cellular networks, along with an evaluation of its usage in a large-scale, multi-tenant deployment supporting community cellular networks in the Philippines. Chapter 5 concludes by discussing the evolution of the community networking and rural connectivity space broadly, with particular attention on the future of community-based network models.

---

[3] While WISP networks are not always community networks, they use many of the same technologies and operate on similar scales as community networks. The lines between community networks and WISP networks are in any case not always clear.

Chapter 2

Understanding the Rural Challenge

Community cellular networks are a recent addition to both community networking and the literature on rural connectivity (the earliest dating from 2012 [88]), and expand on a robust literature on enhancing rural connectivity dating from the creation of the earliest telephone networks. We begin with a brief review of prior work specific to both of these topics to situate our work. To more concretely understand the challenges we expect rural community networks to face, we next review our 2012 study of wireless ISPs (WISPs) in the United States. WISPs represent a particularly useful population to evaluate, as they operate under similar constraints as community cellular networks: they tend to be small scale, rely upon favorable wireless regulations, and largely focus on serving areas unserved by incumbent providers.

2.1 Background

Community Networks. Local or community ownership or development of critical infrastructure has a long history. A well-known concept is coproduction [131, 129, 103], where beneficiaries of services play an active role in producing those services. This concept has found application in endeavors ranging from irrigation infrastructure [108] to urban reform [130]. This is the framework in which community communications networks broadly exist – rather than being simply a service provided for a community, that community plays a role in the provisioning, maintenance, and design of the infrastructure it uses and the set of services it provides.

There is a similar history of small-scale cooperative or locally owned telephony networks [61]. Modern cellular networks have largely ignored these models in most of the world, focusing instead on nation-scale networks. The Village Phone [5] 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. [65] 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 [50]. However, only recently has cellular equipment become inexpensive enough for these models to be economically feasible, and today a variety of researchers and practitioners have built community cellular networks.

Heimerl et al. [87] described building a low-cost cellular network for rural areas, and then built a locally-owned, sustainable, community-run cellular network in Papua, Indonesia [88]. VillageCell [9] and Kwiizya [177] were community cellular networks in Zambia built by researchers with a focus on local communications and services. In addition to the networks we describe in Chapter 4, another large operational community cellular network is the Rhizomatica [143] network in Oaxaca, Mexico, which provides service to thousands of users across ten communities starting in early 2013. These are key examples of the types of networks we seek to support and understand in this thesis.

A number of projects have taken concepts similar to community cellular networks and provided rural connectivity under non-participatory governance structures. The TUCAN3G [113] project built a series of 3G cellular networks in the Peruvian Amazon. While these networks were used to support community services, such as telemedicine, unlike the previous community cellular networks the TUCAN3G sites were built in partnership with an incumbent MNO on a revenue share basis. Similar coordination between rural-focused infrastructure operators and incumbent MNOs to build semi-autonomous rural networks has emerged in multiple markets. In Peru, the experience of TUCAN3G inspired the creation of a new class of service provider, the "Rural Mobile Infrastructure Operator" [40]; this allows an organization to build their own network equipment that uses spectrum held by MNOs in areas where MNOs do not operate. Mayutel [114] is an example of an RMIO in Peru, and Africa Mobile Networks [3] operates under a similar model in a number of West African countries (though solely through commercial arrangements with incumbent operators, rather than via a special regulatory mechanism). In both cases, however, services are still provisioned by the incumbent operator, with these third party companies focused solely on financing, construction, and operation of physical network infrastructure. Thus, while not community networks, these networks still incorporate some elements of coproduction and local cooperation into their business model.

Supporting rural access. Policy efforts to encourage incumbent providers to expand access into rural, underserved areas similarly have a long history. One policy mechanism for bringing coverage to rural areas is a universal service obligation (USO) [154]. 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 [171], designed 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. The United States has also rolled out grant and loan programs focused on rural broadband, including the Connect America Fund [41] and the USDA's ReConnect program [164]. Similar regulations exist in many countries, including Indonesia [163], but many of these are ineffective [74]: hundreds of millions of people in the world remain without basic telecommunication service. As stated before, the reasons for this are economic; operators would prefer to work in areas where they are profitable without the hassle of dealing with USOs [27]. Researchers have attempted to address some concerns with USO systems through competitive means [35], including USO auctions [170]. We acknowledge that these mechanisms can and likely will play a role in expanding rural cellular access; the systems we describe in this work are complementary to this work.

2.2 The Challenges of Scaling WISPs

Given this background, we turn to a specific class of service providers focused on serving rural, underserved populations: wireless Internet service providers (WISPs).[1] Fueled in part by the rapidly decreasing costs and increasing performance of WiFi-based wireless radio equipment in the past decade, WISPs have emerged globally as a means by which rural communities get access to Internet service. Over 2,000 WISPs provider service throughout the United States [173]; worldwide numbers are less clear, but anecdotally thousands more exist.[2] Despite their prevalence in rural areas, these service provider networks are not well studied in the academic literature. During my fieldwork, however, many of the locations in which community cellular networks were being built were served primarily by WISPs; in Papua [88] our local partner was already operating as a WISP, and during exploratory work with Rhizomatica in 2014 and 2015, all of the backhaul to their community cellular network sites was provided by WISPs.

WISPs are often (though not always) examples of community networks, and they provide service to millions of subscribers throughout the US. Those who are unfamiliar with the WISP industry are often surprised by its size: over two million subscribers are served by WISPs in the US alone, and in rural areas, WISPs can be the only source of broadband Internet access besides satellite [14]. WISPs have an even larger impact outside the US. Ubiquiti, a leading hardware vendor for WISPs, reports that "the substantial majority of [their] sales occur" outside the US, and sees emerging markets as a major opportunity for their growth [161]. These facts are unsurprising when one considers the fundamental reduction in capital expenditure required to build a wireless ISP network compared to a traditional wired one. Falling costs and rising performance of commodity wireless equipment, driven by the popularity of WiFi, have allowed the industry to grow as availability of unlicensed spectrum has increased globally. Early research from the academic community demonstrated that the same chipsets used in laptops and phones could be used to build long-distance WiFi

---

[1] This section is based on work that appeared in ACM DEV [83].

[2] Hard numbers on the size of the global WISP industry are hard to pin down since so many operate informally. I personally come across WISPs nearly everywhere I travel doing rural connectivity work; at night, the distinctive pattern of status lights on radios made by Ubiquiti, a major WISP radio vendor, become hard not to notice once you know what to look for. South Africa has an organized WISP trade association, WAPA, and WISPs are very common in all the locations I have worked in Latin America: Brazil, Peru, Argentina, and Mexico.

links [134]; since then, similar technologies have been commercialized and are widely available. Radio equipment for a 50km link providing more than 50Mbps of throughput can be had for under US$200, with each radio consuming under 5W.

Because of the important role WISPs play in providing broadband Internet service to rural areas, as well as their prevalence in areas where early community cellular networks were being built, we argue that the experience of WISPs provides a useful lens through which to understand the broader challenges of providing scalable and sustainable rural connectivity. In the summer of 2012, we conducted a survey of WISPs throughout the US to understand their operations. We sought to answer the following questions:

- What are the demographics of rural WISPs?
- What are the key operational challenges WISPs face?
- What policy support do WISPs require to effectively provide service?

2.3 Methodology

To investigate the operation and characteristics of WISPs, we developed a web-based survey. The survey consisted of 20 questions covering the size of the WISP and its network, budgeting, network failures, and network management; a detailed summary of survey results is available in Appendix 5.2. After completing the survey, participants were invited to participate in a follow-up semi-structured phone interview. We had a total of 75 responses to our survey; 13 of those participated in a follow-up interview. Twelve of those interviewed were active WISP operators, and of those nine operated networks in rural areas. We denote these rural WISPs as R1-R9 below, and the three remaining urban-focused WISPs as U1-U3.

We recruited participants via convenience sampling by distributing announcements on three WISP-focused email lists: the Wireless Internet Service Providers Association's (WISPA) public and members-only mailing lists and the Animal Farm Microwave Users Group (AFMUG) mailing list[3]. The WISPA and AFMUG mailing lists have significant overlap, though the latter focuses primarily on users of a particular manufacturer's equipment (Cambium). We announced the study to the WISPA and AFMUG communities in June 2012.

Participation rate is difficult to calculate since the membership lists of each list are private. Based on public archive records, 434 unique users posted to the WISPA mailing list in the two years preceding our study; of course, this only captures active list participants on a single mailing list (though likely the largest of the three we contacted). The AFMUG mailing list claims that "list membership exceeds 450 members". Nevertheless, our respondent pool represents a wide cross-section of the WISP industry. The vast majority of our survey respondents were involved in the day-to-day operation of a WISP. Other respondents had operated WISPs in the past, but now served primarily in a management role, often as a result of growing their company through acquisition.

---

[3] WISPA is the industry associations for WISPs in the United States.

	1-99	**100-499**	**500-999**	**1000-4999**	**>5000**
Number	4	17	14	32	5
Percentile	5.5%	29.2%	48.6%	93.1%	100%

Table 2.1: Distribution of number of subscribers of US WISPs surveyed.

2.4 WISP Demographics

The WISPs we surveyed were small, by almost any metric one considers. While only 5% of those surveyed had fewer than 100 customers, half had less than 1000, and almost all of those survey had fewer than 5000 customers (Table 2.1). In terms of traffic load, 40% of WISPs surveyed saw a peak traffic demand of under 100Mbps, and 80% had a peak demand of under 500Mbps. Compared to public traffic load data available from PeeringDB as of July 1, 2012 [110], this is small: the median traffic load is between 1-5Gbps, with loads under 100Mbps at approximately the 16th percentile.

One of the characteristic of many of our survey respondents was that they often played multiple roles within their WISP's operations. Most survey respondents reported that they filled a combination of business management, technical management, and marketing roles. This is unsurprising given almost half of US respondents had fewer than 5 employees; 90% had less than 25 employees. For the smallest WISPs we talked to, only one or two people were responsible for the entire operation, though hiring part-time or contract workers for specialized tasks such as tower climbing was common.

The rural WISPs we interviewed ranged in age from just under one year to twelve, with several interview subjects themselves having been involved with telecom or ISP businesses before they started their current WISP. The smallest WISP in our interview population served just over 100 customers [R6], while the largest served more than 4000 [R2]. Our interview participants provided service in the northwest, west, midwest, and southern United States. We additionally interviewed three WISPs focused on urban and peri-urban markets.

2.5 Operational Challenges

The original goal of this study was to develop an understanding of the technical challenges that impacted WISPs, with the hopes of motivating further research on systems to help WISPs operate more efficiently. After collecting data and conducting interviews, we came to realize that some of the assumptions we made were misguided. We initially expected WISPs in the US to face struggles similar to those described in the literature about wireless networks in the developing world [157]: flaky hardware, challenging fault diagnosis, and poor local IT expertise. This turned out not to be the case. In contrast to the "hacked together" wireless systems of the mid-2000's, technology used by WISPs for building and managing their networks has matured sufficiently that operators focus more of their effort on business development than technical issues. One WISP we spoke with in Colorado provided service over a 45,000 sq. mi. area with only 10 employees [R5].

The most common concern among WISPs in our study was spectrum scarcity. Of 43 respondents to the free-response question in our web survey "What is the biggest challenge your organization faces?", 22 (51%) expressed concerns relating to spectrum. The next most common group of concerns was around business development (23%), followed by affordability of upstream bandwidth and backhaul (16%). This was an unexpected result: we had anticipated issues around configuration and manageability of WISP networks to be a major concern, but this was not the case.

Spectrum

We asked about spectrum usage in our interviews. All of the rural WISP operators we interviewed with operated in unlicensed spectrum [R1, R2, R3, R4, R5, R6, R7, R8, R9], but many took advantage of some licensed spectrum as well [R1, R4, R5, R6, R7, R8, R9], most often for important high-capacity backhaul links. In particular, five of the nine rural WISPs [R1, R4, R6, R7, R9] used the 3650MHz "lightly licensed" band, which has been very popular for WISP operators due to the relative quiet of the band compared to unlicensed ones. Multiple interview subjects expressed a desire to have spectrum set aside for WISPs due to overcrowding in the unlicensed bands. According to R1, "basically, we need to use lots of bands because things are so crowded. [The 3.65GHz band] will never have home routers in it. So we can use, especially for backhaul, a relatively obscure chunk of spectrum." Another WISP [R7] expressed frustration with other household devices causing interference with their network equipment: "The biggest problem we have with our network is baby monitors. I hate baby monitors." Higher frequency licensed spectrum (specifically, the 11GHz band) was used in some capacity by five of the respondents [R4, R5, R7, R8, R9], primarily for high-capacity backhaul links.

On the surface, it surprising that spectrum scarcity would be an issue in areas that are largely underserved, but several factors make this an issue for WISPs. The WISPs we spoke with all had a limited number of tower sites for access points to connect customers to their networks. To reduce capital cost of expansion, these WISPs would take advantage of geographic features, re-use existing towers, or re-purpose other tall structures (e.g., grain silos) to avoid building new towers from scratch. Adding more subscribers thus meant co-locating more access points on each tower, leading to interference at the tower site. According to one WISP [U2],

> And then a tower may have anywhere from three to five AP's on it depending on the density of the population and what we need to do in the area. [...] We plan [our network] pretty well, where we try not to self-interfere. The pickiest system about interference, I would say would be the 365's [3.65GHz]. Because you're only dealing with 20 megahertz [available in the band], if you're doing 5

> megahertz channels. And with Cambium [5GHz equipment] you're dealing with 40 megahertz, and it propagates pretty well. We've bumped heads with ourselves on the propagation issue, even 10 miles away from our location. That's definitely an issue.

The second driver was foliage—WISPs serving forested areas reported heavy usage of 900MHz spectrum due to improved foliage penetration. The 900MHz band is the lowest-frequency band commonly available to WISPs, and is only 28MHz wide (compared to over 150MHz for the more commonly-used 5GHz band) yet is shared with a variety of non-WISP users, such as cordless phone systems and smart meters. Several WISPs we interviewed were evaluating the new sub-1GHz spectrum made available by Television White Space (TVWS) to provide service for similar applications to the 900MHz band [R1, R2, R4, R6]. One Ohio WISP [R4] primarily used 2.4GHz and 5GHz radio equipment for their network, with custom installation at each customer site to obtain the best signal possible. This WISP hoped to use the new TVWS spectrum to reach customers that were inside their service footprint, but were unable to be served by their existing equipment:

> We really need the TV white spaces because, out in these rural areas, early indications are we can get six or seven miles, maybe even as much as 10 miles off of an AP, through trees and dales. [...] To our south and to our east, that becomes critical, we have a lot of people we [can't serve] because of trees and where they live. You can only do so much for a $100 install. With a little more money, or a little better radios, we can handle those people. We can get them some service, and that's what the TV white space radios, hopefully, will do for us.

Another WISP in Tennessee [R6] planned to use TVWS for both foliage penetration and the ability to expand into new business lines, including machine-to-machine communication. According to the operator, "I pretty much put all my eggs in the white space basket. I bet the farm on it. If it works we're going to deploy in a big way."

Business Development and Financing

An unexpected theme that emerged from our interviews was difficulties around obtaining financing to expand network growth and meeting demand for service. Particularly for rural WISPs, costs of adding a customer to their network are high, with an installation requiring a site survey, a trained technician, customer premises equipment, and physical installation of the hardware. A humorous but sobering comment from one WISP that responded to our survey was that they switched to using minivans for service calls due to better fuel economy, yet still spend US$2 per user per month on fuel. In an industry with an estimated ARPU of US$30 [14], this represents almost 7% of gross revenue.

We found evidence of “step functions” in WISP growth in our interviews as well. An example from R9 demonstrates a theme that emerged across several of our rural subjects, inability to expand in spite of existing market demand:

> If we had capital, we could light up towers like crazy. The market is there. We get calls from places that were nowhere near close to being serviced but they say, “Hey, we’ve heard good things about your company. Are you in our area?” And we’re like, “well, no.” We could be but capital is the thing, and the other thing is probably spectrum.

Another WISP [R3] went so far as to intentionally slow their growth:

> We’ve actually impeded growth on purpose, because we have to eat what we kill, and we can’t afford to buy CPEs. [... We] tell customers, “look, you’re a month out, [your installation] is scheduled for a month out”, when in reality we just can’t afford to buy the CPEs that fast.

Finding opportunities for stepwise growth, either through acquisition of neighboring ISPs [R2, R4, R5, R9] or other strokes of fortune, was common among the WISPs in our interview set. For example, one WISP operator [R5] stated that their ability to buy used equipment from a bankrupt competitor at a fraction of retail price was instrumental in allowing them to grow their revenue to sustainability. The founder of this ISP reported starting in 2004 with a loan of US$20,000, funding further expansion based on the cash flow from existing subscribers. When the ISP started out, their cost for adding a new subscriber ranged between US$200–$300, driven largely by the cost of customer premises equipment at approximately US$200. According the founder of R5,

> A few years in, I came across a WISP in Minnesota that had 300 Tranzeo radios, and they were switching from Tranzeo to Canopy [another radio vendor]. So, they had all these Tranzeo radios that they were looking to get rid of.
>
> So, I went and I made them a deal. I bought those radios for $50 per radio. I had a come up with $15,000, or whatever it was. That was a deal that completely got us over the hump, because when you can go from paying $200 radio to $50 radio, that’s pretty huge. So, I went and bought these 300 radios and we just went on this install spree. [...] That turned into 300 paying customers, and 300 paying customers that were averaging $40 a month. That’s $12,000 a month, monthly cash flow.

Network Failures

One key area of interest for our study was to understand the frequency and causes of faults in WISP networks. Previous work has identified misconfiguration stemming from management complexity as a leading cause of network outages in enterprise networks [151]; Surana et al. [157] identified poor power quality and harsh deployment environments as leading causes of network downtime, complicated by difficult failure diagnosis due to lack of trained staff.

Do these findings hold in modern WISPs, and do they further apply outside of a developing region?

We asked survey respondents about the frequency of outages due to a variety of different physical- and network-layer causes. Participants were asked to report the frequency of outage due to each cause (“never”, “1–2 times per year”, “monthly”, “weekly”, “daily”), as well as the downtime typically caused by outages of each type (“under 10 minutes”, “10 minutes–1 hour”, “1 hour–6 hours”, “6 hours–1 day”, “Over 1 day”). While outage duration does not take into account scale of outage, we use this as a rough measure of outage severity since it maps closely to service availability. Some failure modes (such as antenna misalignment) by their nature have more localized failure domains than others (such as upstream transit provider failure), while others have high variability (a weather event or a misconfiguration can impact large or small portions of a network).

Figure 2.1 provides a breakdown of failure causes, broken down by outage frequency and duration. We find that outages caused by physical infrastructure faults are both more common and last longer than logical network faults. In line with our findings on limited spectrum availability in Section 2.5, failures due to interference caused the most common outages. Five of our interview subjects also noted this [R1, R6, R7, R8, R9], breaking down between a mix of interference issues at the customer site and self-interference between towers. According to R1, “Our most difficult task is interference mitigation. That’s where I spend most of my time, either upgrading towers or changing channels and whatnot, power levels and so forth, trying to keep the network running smoothly.” In a similar vein, most WISPs we interviewed were proactive about building fault tolerance into their infrastructure to withstand power failures at critical points within their networks [R2, R4, R5, R6, R8] to cope with power outages, often tied to weather events.

Given the harsh nature of the geographies in which these WISPs operate, the fact that physical failures predominate network outages is not surprising and mirrors findings from prior work. In contrast, the relatively limited frequency and duration of outages at the network layer–caused by failures due to upstream transit or transport providers, software updates, misconfiguration, and operator error–stand out.

While we did not find any single conclusive reason for this difference, one possible explanation is that WISPs’ networks tend to be simple. We find evidence for this in both our survey and our interviews. Table 2.2 shows a breakdown of traffic management practices adopted by WISPs in our survey, as well as upstream transit providers. Complex services–that is, ones requiring stateful on-path middleboxes, including those that inspect

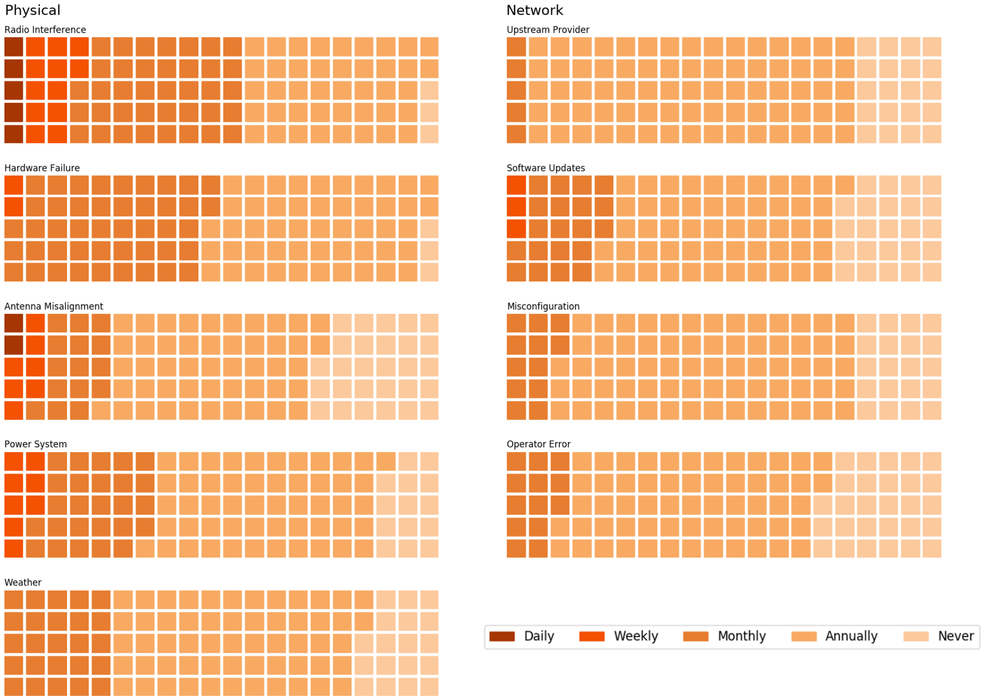

(a) Outage frequency by cause.

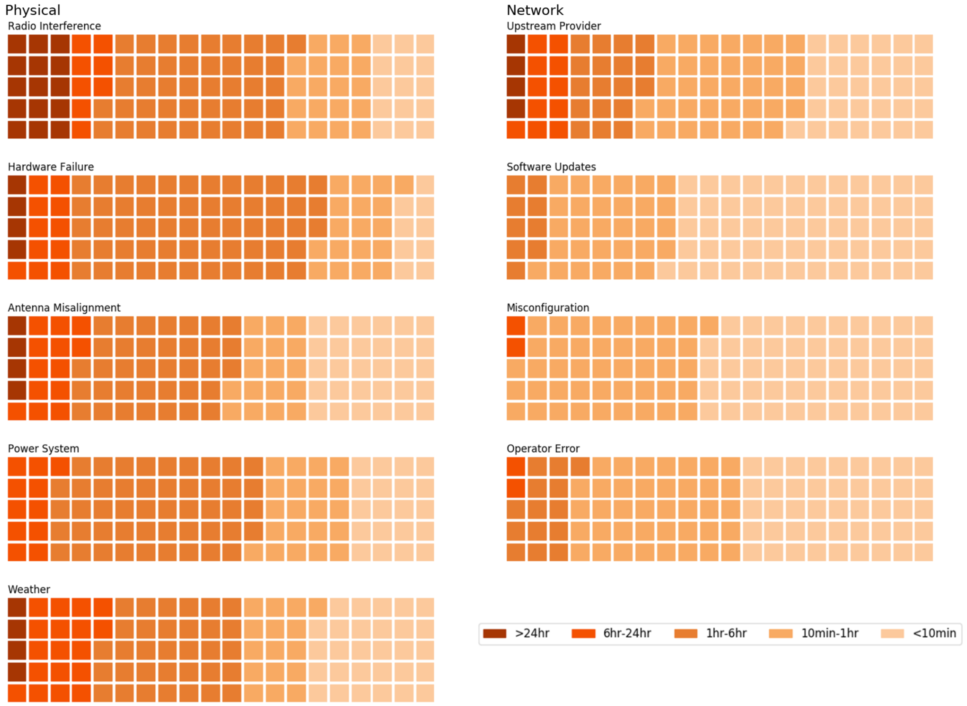

(b) Outage duration by cause.

Figure 2.1: Network outage frequency by failure mode in rural WISPs.

Figure 2.1: Network outage frequency by failure mode in rural WISPs (Previous page). Each square represents 1% of responses. On the left, we see failures due to physical faults, while on the right we see failures due to network management activities. While all of these failure causes are reasonably common, each impacting in some way the majority of respondents, physical faults cause the most frequent outages. Daily outages, for example, are only caused by physical layer faults, while the majority of outages due to physical failures result in outage durations of more than an hour.

application-level traffic–were adopted by a minority of WISPs.[4] Client isolation, QoS, and network segmentation all are features widely implemented on low-end WISP switching and routing hardware, while more sophisticated services, such as caching and throttling, generally require specialized hardware. Thus, they introduce more opportunities for network outages or require more sophisticated failover mechanisms to reliably enforce network policy.

At the other extreme, R7 ran their entire network of 750 clients as a single bridged network, without any form of routing and relying on Rapid Spanning Tree [2] for establishing their logical topology. R7 offered the following rationale for forgoing routers within their network:

> I designed the network that way because routing has its place, but I want to be able to take a vacation and not have to pay somebody 80 or 100 grand a year. It complicates the network that much more. I designed the network to make it as easy as possible to train new employees to be able to operate within our environment. That's why we're a layer-two type network; it's mainly for easier deployment and training.

Billing policies are critical for service provider networks, but the policies we saw in our interview participants were straightforward. The WISPs we interviewed relied on a combination of per-subscriber throttling [R2, R5, R6] or data caps [R1, R4, R7, R8] . Per-user throttling is widely supported in networking equipment used by WISPs, such as Mikrotik [117] and Ubiquiti [162], making this policy simple for WISPs to implement (and, of course, throttling is a stateless policy). Moreover, those that supported data caps did not always technically enforce them. R4 relied on what they called "soft caps": while user plans had data caps, with associated overage charges, the WISP did not actually measure or enforce those caps; R8 took a similar approach. Likewise, R7 included bandwidth caps in their plans, but for enforcement relied upon actually speaking with customers about their usage, focusing on those consuming disproportionate bandwidth during peak usage times:

> I'd give them a call and I'd say, "Hey, I don't really care if you do BitTorrent, but if you're doing it at 7:00 at night we're going to have problems. [...] If you

---

[4] The most popular of these, application filtering, appeared from our interviews to be largely used for filtering peer-to-peer file sharing. However, R7 reported that this traffic has decreased in volume on their network to the point that "it's not even an issue I look for anymore."

Practice	Adoption
Client Isolation	81%
VoIP QoS	80%
Network Segmentation	77%
Throttling	44%
Application filtering	37%
Caching	17%
Content filtering	10%

(a) Network management practices.

# Upstreams	Percentage	Cumulative
1	30.5%	30.5%
2	44.1%	74.6%
3	22.0%	96.6%
4	3.4%	100.0%

(b) Upstream transit providers.

Table 2.2: WISP networks are not complex. On the left, we see network management practices and adoption rates by WISPs serving rural areas. "Client isolation" refers to blocking traffic between subscribers. Management practices that require stateful and/or on-path middleboxes are not adopted as frequently as stateless services implemented at the IP layer or above. On the right, we see the number of upstream transit providers of each WISP; three-quarters have only two or less. In practice, this figure also represents the number of physical interconnections between the WISP network and the broader Internet.

> do it after midnight, and you stop it before 7:00 in the morning, I don't care. Go to town. [...] The reality of it is I'm not making any money when you're doing that. If I can't make money on the account, I don't want the account." They, generally speaking, were very understanding of that.

Our takeaway is that while outages in WISPs do happen due to (logical) network management issues, these are not themselves barriers to WISP scaling today. This is likely due to simple network policies and topologies employed by WISPs.

2.6 The Role of Automation

Most of the challenges we've identified in this chapter are, at least superficially, either non-technical (access to financing and spectrum) or simply the cost of doing business for rural ISPs (weather-related outages, customer support). We have also seen that WISP network outages are most often due to physical failures, likely because WISP networks are simple from a topology perspective. This begs the question–are there opportunities for *technical* solutions to the challenges of scaling WISPs?

In short, yes. A key theme that emerged from our study was the importance of integration and automation into network management systems for wireless ISPs. A number of our study participants had either built their own backend management systems from open source products [R2, R5, R6, R7, R9], were using a commercial product that handled billing and customer support [R1, R6, R7, R9], or relied on vendor-provided tools to manage network devices [R5, R6]. We broadly break the systems and techniques used by the WISPs into two categories: BSS/OSS integration, and proactive monitoring and automation.

BSS/OSS integration

Billing support systems (BSS) and operational support systems (OSS) are integral to any service provider network. "All in one" commercial products that integrated billing and customer management into a WISP's network operations were common in our study. WISPMon [174], Azotel [15], PowerCode [137] were all used in our sample, each of which handled customer billing, report generation, payments, and implementation of service plans. R9, in addition to operating a WISP, also provided consulting services to others starting WISPs. He recommended:

> [Before you] do anything, you should have your backend stuff in place, have your accounting package in place, have the servers in place, that track the customers and authenticate them and if they don't pay their bill, redirect them to a, "Hey, you haven't paid your bill. Here's how to pay it" [web page]. Have all of these systems in place because it will help you do it with less people. The more automation you have from the start, the less people you have to have. That just turns into the less people you have to hire, the more money you can get.
>
> [I advise WISPs that are just starting out] to have as much automation as possible, anything you can automate, billing, subscriber authentication, tracking if they paid their bill or not. Have all that in place and have a good business foundation too.

These systems become integral to WISP operations; no WISP we interviewed described replacing their billing system. Even the smallest WISP we interviewed, R6, had WISPMon configured to automate billing within 9 months of operation.

Related to these customer-facing tools, WISPs we interviewed devised a variety of adhoc or custom solutions for building OSS tools to manage their installation and repair operations. R2 used a detailed shared online calendar to coordinate their staff; salespeople could update the install schedules of technicians, and technicians would add details of their installation (such as measured signal strength) to the notes of each invite for archival. One urban WISP, U3, developed a custom customer relationship management (CRM) tool for tracking interactions with customers; using this tool, every staff member's interactions with each customer were tracked, including changes to network configuration on customer equipment, equipment type, and problem history. R9 expressed a desire for a system quite similar to U3's, augmented by automatic change detection based on polling device configuration in the field.

In all these cases, these tools were used to help organize the technical and business operations of each WISP. Integrating business and operational processes with network management tools helped WISP operators have a clear overview of the *management* plane of their networks and businesses. Standardization and automation also simplified training.

Proactive Monitoring and Automation

Almost all WISPs engaged in some form of proactive network monitoring to identify faults before receiving a user report [R2, R4, R5, R7, R8, R9]. While three WISPs [R2, R4, and R8] configured their monitoring systems to send SMS-based alerts to staff upon detection of outages, proactive monitoring largely took the form of alert systems based on open-source software such as Nagios [119], used to generate dashboards for quickly identifying anomalies in network performance. According to R7,

> I can take a look and see if there's an individual customer having an issue. If they're intermittently having an issue I can probably find it in about– without them calling or anything; me just scanning the network–probably about 15 minutes. If they're down hard I'll know about it in 30 seconds.

Similarly, WISPs employed automation at the site level to both detect and remediate physical network faults, the most common cause of network outage we observed. R5 deployed "ping rebooters" extensively in their network; these devices would simply power cycle a device that stopped responding to ICMP pings; according to R5, "[this fixes] 90% of the issues that come up remotely." Automatic fail-over within the operator's backhaul was deployed in all rural WISPs we spoke with as well where redundant paths were available. R8 deployed automatic backup generators to key tower sites. These simple examples demonstrate the importance of basic level of automation at remote sites, particularly if it can save an in-person maintenance visit and associated cost.

2.7 Conclusions

The initial goal of this study was to develop an understanding of the technical challenges that impacted WISPs, with the hopes of motivating further research on systems to help WISPs operate more efficiently through enhanced network automation. After collecting data and conducting interviews, we came to realize that our motivation for the study was somewhat misguided. The WISP industry has matured to the point where technology and manageability of their networks are not the biggest bottlenecks for growth. We found that the primary operational challenges faced by WISPs are spectrum policy and understanding how to grow their own businesses. Further, physical faults cause more frequent and significant network outages than those caused by logical network faults.

We argue that the state of the modern WISP industry has lessons to provide to the nascent community cellular network community. An obvious takeaway is the vital importance of spectrum availability that meets the needs of these small operators. In the case of WISPs, they seek spectrum for backhaul networks that isn't subject to interference from other uses; while more unlicensed spectrum would certainly be beneficial, it wouldn't directly address this need. Another finding with application to community cellular networks is the importance of utilizing existing client hardware: the fact that community cellular networks can leverage existing client infrastructure [88] simplifies customer acquisition and network expansion, with CCN operators able to focus on building just their core network. Another lesson from the WISP industry's experience is that the management overhead–both business and technical–for WISPs becomes a bottleneck as WISPs scale, and tools to help standardize business processes and automate network management are vital for a WISP's success.

In the next chapters, we consider these problems in further depth, first focusing on a system to support improved access to spectrum for CCNs, taking into consideration their particular needs. We will then consider an integrated system that provides proactive monitoring, automation, and BSS/OSS integration for CCNs.

Chapter 3

Shared Spectrum for Community Cellular

In the previous chapter, we saw that access to spectrum for rural wireless Internet service providers (WISP) is a significant obstacle to scaling. These operators largely rely on unlicensed spectrum in the 900MHz, 2.4GHz, and 5GHz bands, and in our interviews reported needing more spectrum to provide higher service speeds to more users to keep up with increasing demand.

We find even greater spectrum challenges for community cellular networks. Indeed, the primary impediment for a community cellular network operator is access to spectrum – with few exceptions, the frequency bands used by mobile networks (and most importantly, mobile handsets) is licensed spectrum, with high barriers to entry for gaining access. As a result, early deployments of community cellular networks in Papua [88] and Zambia [177] operated without any spectrum license, outside the regulatory framework. The Rhizomatica network [143], while today holding a long-term spectrum concession, initially operated under a short-term experimental license. The lack of enabling regulatory frameworks, including a feasible way for community cellular networks to obtain access to licensed spectrum, creates uncertainty for community cellular networks and their users and thus limits their growth.

Yet prior empirical work has shown that, in general, spectrum is largely underutilized, especially in rural areas. An evaluation of spectrum occupancy across multiple bands in urban and suburban locations in Europe found occupation well under 50% across the 400MHz to 6GHz band, even in commercially important bands used for cellular networks and television [166]. Table 3.1 provides an overview of spectrum occupancy measurement results, focusing on studies that looked at occupancy across a wide band (>500MHz). While these results are best thought of as a lower bound on spectrum utilization – utilization in practice can be higher than observed spectrum occupancy due to frequency hopping and wideband signals that vary in duty cycle [152] – they provide evidence that unused spectrum is common.

This begs the question: why do WISPs consider spectrum to be scarce, particularly in rural areas where we expect relatively low spectrum utilization? In short, because spec-

Type	**Location**	**Year**	**Bandwidth (Band)**	**Occupancy**
Urban	USA (New York City)	2005	2770MHz (30MHz-3GHz)	13% [115]
Urban	New Zealand (Auckland)	2007	1944MHz (806MHz-2.75GHz)	6.2% [34]
Urban	Singapore	2008	5770MHz (80MHz-5.85GHz)	4.54% [93]
Urban	India (Pune)	2013	632MHz (174MHz-806MHz)	<7% [133]
Urban	Romania (Bucharest)	2014	3375MHz (25MHz-3.4GHz)	21% [112]
Urban	Czech Republic (Brno)	2009	2900MHz (100MHz-3GHz)	7% [165]
Rural	USA (Virginia)	2004	2670MHz (30MHz-2.9GHz)	3.4% [116]
Rural	USA (Maine)	2007	2770MHz (30MHz-3GHz)	1.2% [53]
Rural	Romania (Maneciu)	2014	3375MHz (25MHz-3.4GHz)	14.2% [112]

Table 3.1: A brief summary of wide-band spectrum occupancy measurement studies. Each evaluation uses its own methodology for measurement (such as choice of power threshold to consider a band "in use") making direct comparison difficult, but the results overall suggest (a) low spectrum utilization in practice and (b) higher utilization in urban areas than rural areas.

trum has historically been allocated on an application-specific and often exclusive basis to its users, be they radar operators, broadcasters, or paging systems, aircraft, or broadband providers. Thus, in practice, WISPs are limited to approximately 250MHz of "unlicensed" spectrum for their access networks. This spectrum is shared with all other unlicensed spectrum applications, such as in-home WiFi, baby monitors, smart meters, cordless phones, and other widely deployed technologies. On the plus side, this leads to a rich hardware ecosystem around these bands, lowering costs through scale: WISP equipment leverages the same WiFi chipsets used for mobile phones, laptops, and routers. Yet it also yields the paradox that rural broadband providers – WISPs – are unable to (legally) leverage unused spectrum in their service areas, even when their hardware can technically support doing so.[1]

For community cellular networks, we see the same paradox of spectrum abundance faced by WISPS: while spectrum occupancy in key cellular bands (Bands 2, 3, 5, and 8, covering non-contiguous blocks between 824MHz-960MHz and 1710-1990MHz) is among the highest measured in the occupancy studies in Table 3.1, it is still lightly utilized in rural areas. Yet the situation is even more extreme: because cellular spectrum is exclusively licensed, we can even safely say that any location without mobile service has *no* spectrum utilized at all. In these locations, a community cellular network could safely use this licensed frequency

---

[1] Because unlicensed spectrum bands vary somewhat internationally, most WiFi-based equipment supports a fairly wide band. For example, the Ubiquiti Nanobeam M5, a sub-US$100 radio used by WISPs, can support operation over a 705MHz band (5170-5875MHz), but is restricted to just 125MHz (5725-5850MHz) in the United States to comply with FCC regulations.

without any risk of causing interference to licenseholders.

3.1 A Brief History of Spectrum Regulation

To better understand this paradox, as well as to contextualize our proposed solution, a brief overview of how spectrum is regulated is useful. We will focus on the experience of the United States for brevity, but the general approaches to spectrum regulation the United States has employed are instructive for those used worldwide.

The Early Days

Radio spectrum is fundamentally a shared resource: one operator's use of a particular frequency in a particular location at a particular time in general limits other operators from using the same. Since the earliest days of radio, spectrum has represented a "tussle space" [38] with a myriad of competing interests and stakeholders seeking to define the norms and rules for its appropriation. The United States' Radio Act of 1912 was an early and formative attempt at formal radio regulation, which prohibited radio transmission in the US without a license from the Department of Commerce. Yet this seemingly restrictive requirement in practice had no teeth. A series of court decisions restricted the government's ability to deny licenses (*Hoover v. Intercity Radio*, 1923) and to regulate frequency usage and output power (*U.S. v. Zenith Radio Corporation*, 1926), effectively leaving spectrum usage completely unregulated. With the rise of commercial broadcast radio service in the 1920's, and the dramatic increase in the number of commercial users of spectrum in the US, this regulatory regime proved inadequate.

In response to increasing interference between radio users due to lack of coordination, and in collaboration with the radio and broadcasting ecosystem, Commerce Secretary Herbert Hoover organized a series of "Radio Conferences" in the mid-1920's, culminating in the passage of the Radio Act of 1927 [24]. This Act established the Federal Radio Commission (FRC), the precursor to the modern Federal Communications Commission (FCC), and provided the general outline of spectrum regulation in the United States (and by example, for much of the rest of the world) for decades to come. The Act required all radio transmission to operate under a license issued by the Commission, and empowered the Commission to establish technical regulations (such as power output and frequency usage) and to allocate licenses to serve "the public interest, convenience, or necessity." This standard would be the basis for so-called "beauty pageant" administrative assignment of spectrum licenses: licenses granted at no cost when the Commission found the application would serve the "public interest". In practice, the Commission often held that the "public interest" was well served by large-scale (and sometimes monopoly) actors such as AT&T and RCA, with the early years of the FRC and FCC marked by regulatory capture [175].

Market based regulation

In 1951, Leo Herzel suggested the “much more controversial” approach of treating access to spectrum as a market problem [90]. Eight years later, Ronald Coase published his foundational article advocating for assigning spectrum via auction, rather than beauty contest [39]. DeVany et al. built on this work, proposing a framework for property rights for spectrum utilization, defining a so-called “TAS package” – the combination of time, geographic area, and spectrum utilization (both frequency and power level) – that would capture the rights of a licenseholder in a market-based system [44]. The critical insight across this work was that contention for access to spectrum was an inherently economic question and that a regulatory regime that granted a notion of property rights to spectrum could achieve a more efficient and flexible allocation of spectrum resources compared to administrative assignment.

The first spectrum auctions took place in New Zealand in 1989 [43], forty years after Coase’s initial proposal. India followed with auctions starting in 1991 [97], and by 1994 the FCC conducted the United States’ first auction [42]. Today, allocation by auction is the predominant means for allocating cellular spectrum. The FCC has conducted over one hundred auctions as of 2019, and regulators worldwide have raised billions of dollars via auctions [72, 6]. The details of actually implementing a spectrum auction are non-trivial – what type of auction should be conducted? What reserve price, if any, should be established? What TAS package should the license cover? Nevertheless, it is fair to say that nearly 70 years since Coase’s proposal auctions have proven to be an effective and efficient means for assigning spectrum [85].

From the perspective of users outside of mobile coverage, however, spectrum auctions leave much to be desired. First, auctions are a source of government revenue; selling a small number of high-value spectrum licenses is most efficient for resource-constrained regulators, even if high license costs potentially lead to diminished consumer surplus due to network operators passing the license costs to users via higher service costs [72]. Second, and more importantly, in most countries cellular spectrum is allocated on a nationwide basis. This means that organizations who wish to obtain spectrum in profitable dense urban areas must also pay for the right to use that spectrum in rural areas, and vice versa. Broad geographic licenses enable policy tools like coverage requirements (that is, conditions of a spectrum license requiring the licenseholder to provide service to a particular threshold of the area or population covered by their license), but providing exclusivity over wide geographic areas is neither desirable nor necessary for small-scale community cellular networks. A spectrum licensing regime designed to provide access to spectrum exclusively to large-scale organizations limits the actors capable of building mobile networks and, more importantly, constrains the growth of those networks to the investment profile of these large-scale organizations. Even if a mobile operator could profitably serve low return-on-investment rural areas, the opportunity cost of doing so compared to investing in network expansion elsewhere within their license area are likely prohibitive.

The other practical challenge with a market-based approach to spectrum licensing is that regulators must contend with nearly fifty years of non-market based spectrum assignment [85]. In the United States, most spectrum below 6GHz had been administratively assigned to a mix of private and public sector users, including military users, on an exclusive basis. This "legacy assignment problem" is perhaps the crux of our challenge as we consider opportunities to make spectrum available for rural community cellular networks.

Dynamic spectrum sharing

One of the underlying assumptions behind exclusive licensing is that spectrum can only be used by a single user at a time; more specifically, allowing multiple users to transmit at the same time will result in undesirable interference. This may have been strictly true when analog broadcast services were the predominant commercial use of radio, but advances in wireless technology and the emergence of wireless data services have invalidated this assumption.

The ALOHANet system [26] pioneered random access packet radio as early as 1971. Proposals for wider area packet radio networks using related techniques following in subsequent years (e.g., RADIONET in 1976 [104]), as well as performance improvements to the original ALOHA protocols such as Slotted ALOHA [144] and MACAW [25] directly influenced today's widely deployed wireless technologies such as 802.11 WiFi and GSM (Global System for Mobile Communications, the predominant 2G cellular standard worldwide). Not only did these systems demonstrate the viability of packet radio, they also showed that uncoordinated wireless systems could effectively share spectrum resources at fine granularity. The near ubiquitous deployment of WiFi worldwide further proved that unlicensed spectrum could support commercially-meaningful data services without the need for any licenses, let alone exclusive use. Indeed, today more than 50% of data usage worldwide comes from users on WiFi networks in unlicensed spectrum [36] – far exceeding that over mobile networks in exclusive-use licensed spectrum.[2]

These systems all rely on MAC layer protocols designed for cooperative sharing of spectrum resources. Dynamic spectrum sharing (or dynamic spectrum access) refers to techniques that generalize this concept to enable spectrum sharing across services, independent of service and protocols used. This is particularly relevant for improving utilization of brownfield spectrum – that is, spectrum that has already been allocated and assigned. Dynamic sharing techniques can be broadly categorized into those that use sensing [68] and those that use databases for coordination [118, 78]; hybrid approaches are also feasible.

For example, a portion of the 5GHz band adjacent to the unlicensed portion of the band

---

[2] This figure is somewhat misleading because it does not consider the impact of pricing. WiFi tends to be deployed indoors as a "last mile" technology connected to fixed network infrastructure. These networks often have higher data caps, higher speeds, and lower relative prices compared to mobile networks, and are deployed in locations where people spend more time, such as homes, schools, and workplaces. Nevertheless, the broader points that (a) WiFi networks are commercially important and (b) unlicensed spectrum supports a massive portion of modern Internet usage are undeniably true.

is allocated for Terminal Doppler Weather Radar, a type of short-range radar system used at some airports to detect wind shear. This spectrum is unused when and where these radar systems are not in use and could otherwise easily be used by WiFi equipment. In 2003, the FCC established rules for unlicensed equipment to operate in this U-NII band [58], establishing the first shared spectrum regulation in the United States. Under these rules, unlicensed transmitters must both observe strict transmission limits and implement a sensing capability to detect whether the primary users of the band, radar systems, are in operation. More recent policy making has focused on spectrum sharing in television whitespaces (TVWS) [59, 60], spectrum made available by the transition to digital television from analog, and the resulting decrease in channel bandwidth used for broadcast television. TVWS spectrum sharing proposals leverage databases as the primary mechanism of spectrum coordination, with sensing taking a backseat [81, 80].

Partly in response to this, in 2012, the President's Council of Advisors on Science and Technology (PCAST) issued a report outlining the scope of this challenge and identifying opportunities to make lightly-utilized federal spectrum bands available to more users [91]. Recognizing the difficulty of reallocating federal spectrum, the principal finding of the PCAST report was that "the norm for spectrum use should be sharing, not exclusivity", and goes on to identify over 1,000MHz of federal spectrum that could be shared with commercial users. While rulemaking for both U-NII and TVWS services was well underway when this report was released, it informed the next major shared spectrum band, the Citizens Broadband Radio Service (CBRS). Similar to U-NII, CBRS facilitates sharing of spectrum with a limited number of existing users including both satellite ground stations and some ship-based naval radar.

3.2 Dynamic Spectrum for Community Cellular

Building on previous dynamic spectrum sharing proposals, we argue for spectrum sharing in *GSM whitespaces* to make spectrum available to community cellular networks providing GSM service in rural areas.[3] Here, GSM whitespace refers to licensed GSM spectrum that is unused in a particular geographical area[4] 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 (Figure 3.1). NGSM takes advantage of the fact that GSM handsets continually measure spectrum occupancy in their area and report these measurements back to their serving cell. By doing so, it can sense for

---

[3] This remainder of this chapter is based of work that appeared at DySPAN 2014 [84]; a similar chapter appears in my co-author Kate Harrison's dissertation [79], as she was instrumental in developing key high-level ideas presented here, such as nomadic operation.

[4] Note we do *not* refer to the space between adjacent GSM channels.

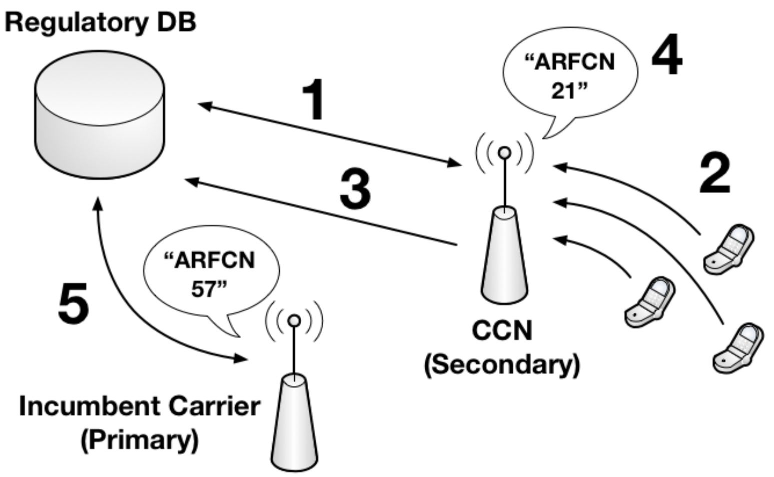

Figure 3.1: Overview of Nomadic GSM. (1) The secondary initially queries a regulatory database for available channels in its area. (2) The secondary gathers measurement reports from its subscribers' phones. (3) Secondaries report spectrum usage (both their own and measured) and service usage (e.g., number 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.

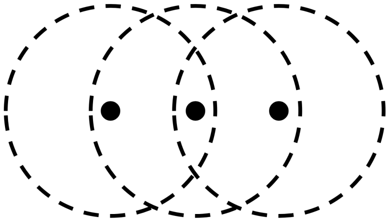

Figure 3.2: The hidden node problem. The middle node is within the transmission and reception range of either of the two exterior nodes. If the left node transmits at the same time as the right node, the center node will experience interference; however, the left node has no way of detecting this on its own since it is outside the transmission range of the right node. Note this example assumes all connections between nodes are balanced links, i.e. the transmission and reception range is identical.

potential interference at the *client* device, mitigating the practical impact of the "hidden node problem" (Figure 3.2). Although certain edge cases necessitate a spectrum occupancy database, NGSM enables secondary operators like CCNs to share licensed spectrum without requiring cooperation or even 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 motivation for using GSM whitespaces to regulate coverage for community cellular networks stems from two key observations. First, utilization of spectrum allocated to GSM networks is much lower in rural areas than urban ones, but spectrum is often allocated to operators uniformly across rural and urban areas. Second, the GSM protocol itself is well suited for whitespace; every one of the tens billions of 2G GSM phones in the world are capable of providing measurements for spectrum policing and enforcement.

The contribution of this chapter 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.

3.3 Community Cellular Networks

Historically, cellular networks have been expensive to build and complicated to operate; this is particularly the case for rural cellular networks [86]. 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 [71]. 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 [127] – have challenged this status quo. A rural community can build and operate their own cellular network for under $10,000 in capital expenditure [88]. 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.

As we saw in Chapter 1, these advancements have enabled *community cellular networks* as a new model for advancing rural connectivity. Community cellular networks are locally owned and operated, and tend to consist of at most a handful of base transceiver station (BTS) sites sites. Such networks exist in Papua, Indonesia [88] and Oaxaca, Mexico [143]. 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.[5]

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 [88], service was *coproduced* [103, 108] 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 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 do this.

---

[5] Indeed, 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 3.2: Bandplan for the GSM900 band in Indonesia [148]. The entirety of the band has been granted to these three carriers under nationwide licenses.

3.4 GSM Whitespaces

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. Figure 3.3 shows the national cellular coverage map for Indonesia.[6] Although the entire GSM900 and GSM1800 bands have been licensed to carriers in Indonesia (Table 3.2), vast swaths of the nation remain without any coverage. The largest provider, Telkomsel, claims to cover "over 95%" of the population as of 2013 [160], 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.

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 all others from offering service. Licenses have traditionally been offered in this manner 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.

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

---

[6] Obtaining 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 [159] and crowdsourced coverage maps [147].

Figure 3.3: Indonesian cellular coverage as of 2013 across all carriers. 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.

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.[7]

---

[7] This idea was advanced in a public conversation by John Chapin during the 2012 ISART workshop in

5. *Backwards compatibility.* Existing, unmodified GSM phones should work with secondaries' networks.

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.*

3.5 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 (Figure 3.1). 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.

Interference Detection

A key issue for dynamic spectrum sharing schemes that rely on sensing is the hidden node problem [25]. 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.

Boulder, CO.

For example, in the Papua network roughly 66% of traffic is outbound [88]. 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 [135].

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 ARFCNs 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 the 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 channel 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. Conceptually similar to sharing spectrum sensing results as proposed in CORVUS [32], 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 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.

This example demonstrates that it is possible for a secondary to coexist with a primary – even in the same geographical area. The inverse scenario, where a secondary wants to enter a coverage area already served by a primary, is essentially the same. However, we have a new problem – the cooperative BTS must be careful not to choose a channel occupied by a primary user. Ideally, the secondary can check the database to find available frequencies in its area; if this is possible, we can simply use those channels and enter interference avoidance mode.

If the database doesn't have such information, however, other options are available. One such approach would be for the secondary site to operate at very low power so that only nearby handsets associate with the tower. Once a handset is camped to the secondary, we can begin using handsets to survey the surrounding area for other users. Assuming the handsets find no such interference, the secondary can continue increasing its power, allowing more distant subscribers to connect to the base station and thus allowing measurement reports from a wider area. Once the secondary is at full power, it can enter normal interference avoidance mode. Depending on the capabilities of the secondary's hardware, it could also sense what channels are in use in its area before transmitting; this introduces new problems related to the hidden node problem, however. We leave a full investigation of the challenges of enabling secondary usage in areas already served by primaries to future work, noting that areas with primaries providing service already have access to communication services.

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 3.6). 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 as in practice incumbent users often turn their transmitters on for testing well in advance of commercial operation, 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 [86]). 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 [86]. By cycling frequencies only when the BTS is not in use, we can avoid interrupting ongoing calls while minimizing perceived downtime to users. However, this doesn'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 a 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 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. Our technique for determining when to perform handover is novel, but the mechanism by which we move clients from one frequency to another is completely standard. NGSM performs handover between channels on the *same* BTS for the purposes of frequency agility, but does not prevent handover for mobility between different BTS units.

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 3.6).

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).

3.6 Opportunities, Risks, and Incentives

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 [88] 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.

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, reducing risk. Similarly, wholesale revenue opportunities for incumbents may also be possible; we explore this in more detail in Chapter 4.

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.

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 perceived "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.

Other Stakeholders

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

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 initially operated under a two-year experimental license, but while this enabled legal operation in the short term, they had no guarantee their license will be extended in the future. Such experimental licenses could also restrict CCNs from making profit with their networks, which hinders efforts to provide sustainable and reliable service to 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 3.6, 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 shutdown 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.

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, are fraught with political baggage, and have had mixed results in practice [74].

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.

3.7 Evaluation and Deployment

In Section 3.5, we claimed that NGSM achieves all five goals for spectrum sharing in GSM whitespaces (Section 3.4). 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.

Implementation

We implemented NGSM as a software control layer based on OpenBTS [127], which uses a flexible software-defined radio and a commodity PC to implement a GSM base station. We support dual-ARFCN operation, with one software radio per ARFCN. We also support single-ARFCN operation, with the 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 RSSI[8] 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 $K = 1$ and $R = 0$ 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

---

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

Variable	**Value**	**Purpose**
R	0	For ARFCNs *not* in use by the secondary, the RXLEV threshold for determining whether an ARFCN is in use.
K	1	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 3.3: Parameters for NGSM.

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.

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. 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`.

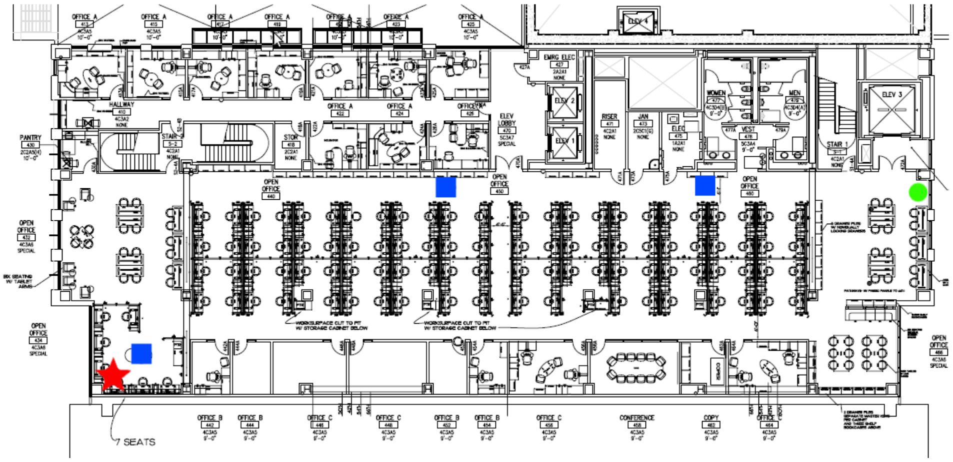

Figure 3.4: Floor layout showing the primary (green circle, right), CCN (red star, left), and three cellphones (blue squares). The monitoring site was co-located with the middle phone.

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 3.5, 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. Figure 3.4 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.

Figure 3.5 shows the results of this test in the uplink band. This figure shows the usage of the phone on a call while the BTS alternates between ARFCNs. As expected, the phone completes handover successfully and the call continues without interruption. In Figure 3.6, we see the spectrum usage on the downlink band during a simulated appearance of a primary user. Initially, the secondary BTS is alternating between ARFCNs 20 and 30. The primary

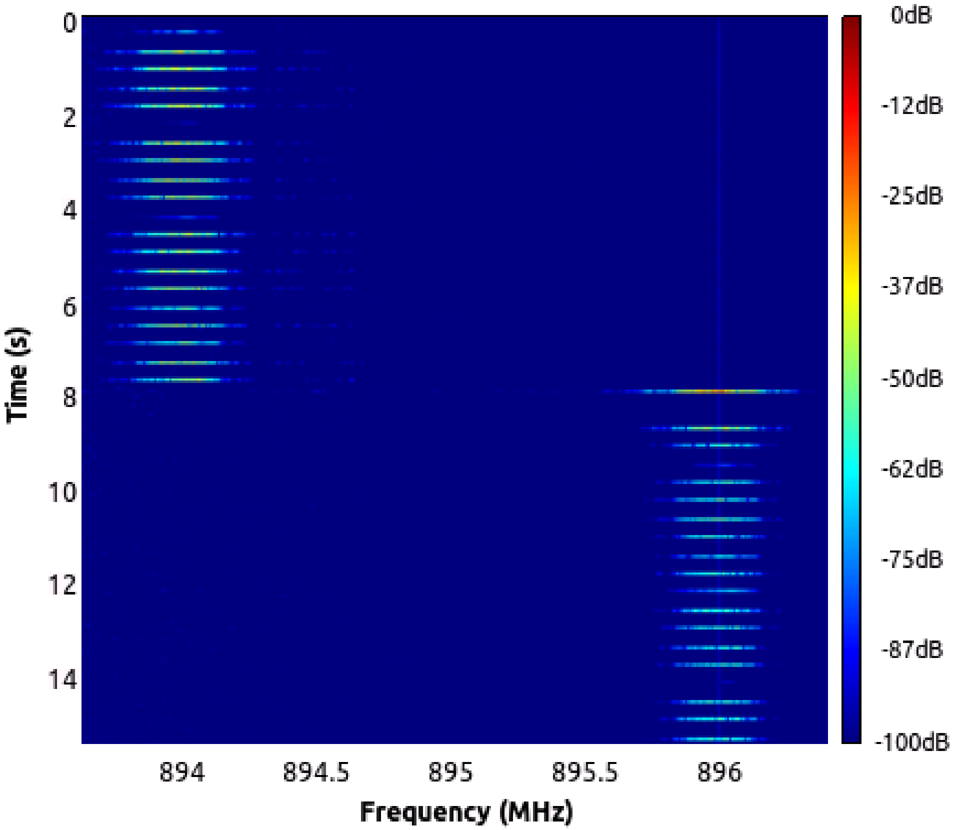

Figure 3.5: Spectrum usage of a handset during a call in the uplink band (i.e., from the phone to the BTS). The phone switches ARFCNs without interrupting the ongoing call.

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.

In all cases, the secondary BTS was able to detect the primary BTS and cease use of the interfering ARFCN within six minutes, and in one case less than two minutes. Moreover, the secondary experienced minimal downtime, generally less than thirty seconds between detecting the primary and restarting again on a new, safe ARFCN. Because the secondary used two ARFCNs, any phones that were camped to the ARFCN that did not change experienced no downtime at all.

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

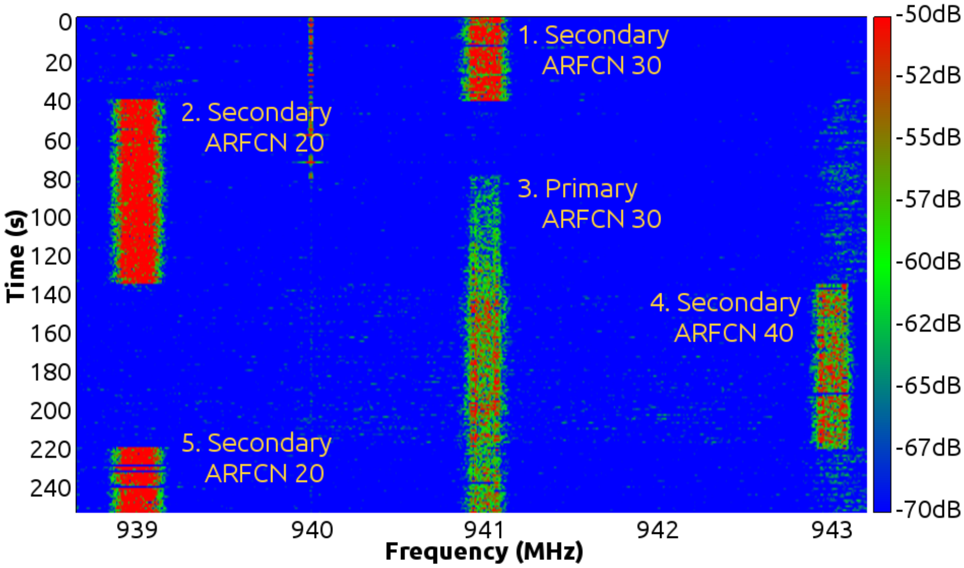

Figure 3.6: Spectrum usage during simulated arrival of a primary into the secondary's service area. Initially, the secondary is alternating between using ARFCNs 20 and 30 (939MHz and 941MHz, respectively). When a primary BTS appears on ARFCN 30, the secondary detects its presence and switches to using ARFCN 40 instead (943MHz). Thereafter, the secondary alternates between ARFCNs 20 and 40.

usage accordingly within a matter of minutes. Finally, it shows that secondaries can rapidly change channels upon detecting interference with minimal service disruption.

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), 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

Trial	Detection (s)	Switchover (s)
1	187	23
2	59	33
3	353	11
4	241	23
5	124	12

Table 3.4: Results of coexistence test. Detection is the time between when the "primary" BTS was activated to when the secondary BTS recognized its presence; switchover time refers to how long the secondary BTS took to change frequencies and become fully operation after detection. The secondary was able to detect the primary and change frequencies in less than six minutes.

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, Figure 3.7 shows the distribution of interarrival 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. Figure 3.8 shows the number of measurement reports received during the operational hours of the CCN, collected in 10 minute bins. The minimum number of 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.

Deployment

NGSM operated as expected when we deployed it onto the Papua CCN. Figure 3.9 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

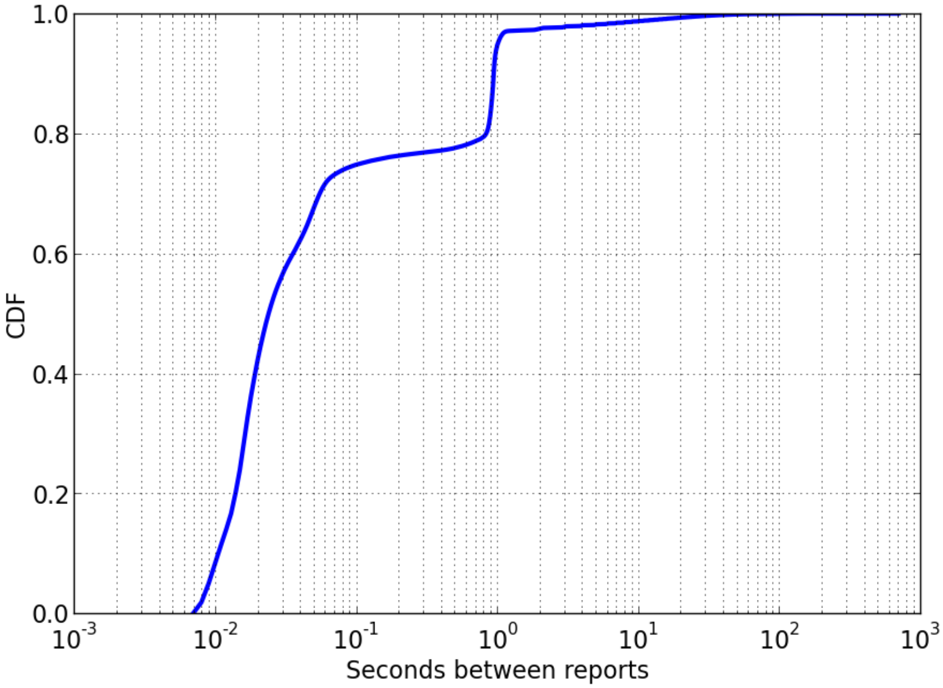

Figure 3.7: CDF of measurement report interarrival times received during daytime hours (7AM–12AM). Night time hours are not included since the CCN was typically powered off at night. $\mu = 0.64$, $\sigma = 4.02$. Note logarithmic scale.

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 ($K = 1, R = 0$).

Despite these frequent channel changes, we observed no negative impact on network usage after deploying our system. Table 3.5 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, 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).

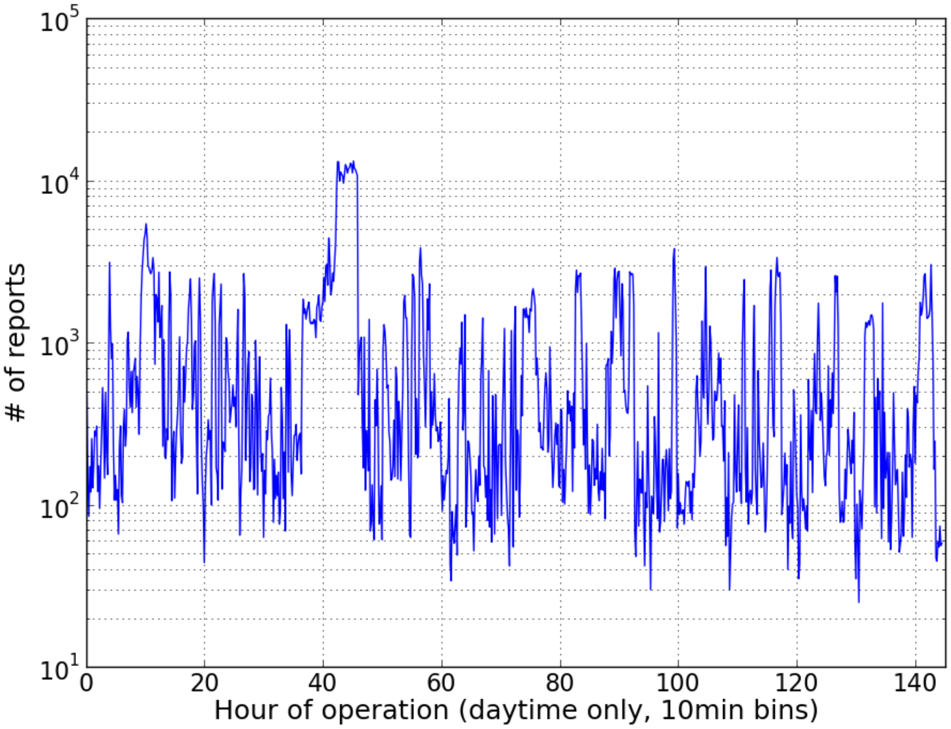

Figure 3.8: Time series of number of measurement report received per hour during daytime hours (7AM–12AM). Night time hours are not included since the CCN was typically powered off at night. Note logarithmic scale.

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 3.5: 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.

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

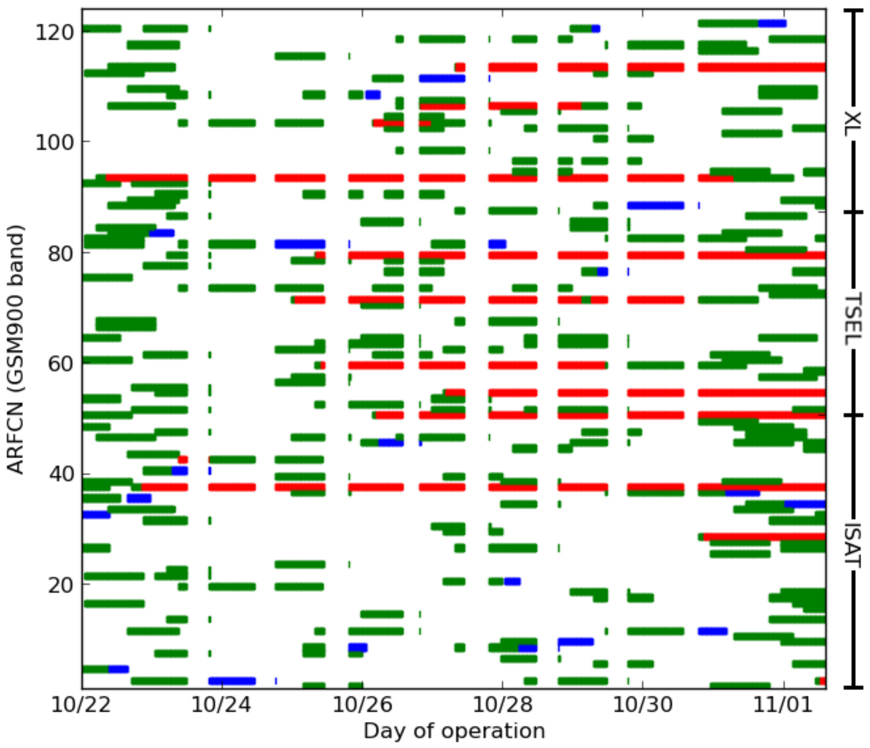

Figure 3.9: Spectrum usage and reports for the Papua CCN from Oct 22 through Nov 1. Blue represents the ARFCN in use by the CCN. Green represents an ARFCN that was scanned and considered “safe” to use, and red represents an ARFCN which was scanned but appeared to be in use, and was thus “unsafe”. Unsafe 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.

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.

3.8 Discussion

Market alternatives to GSM whitespaces

The obvious market-based alternative 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 the secondary spectrum market to obtain local usage rights in that area. Conceptually tempting and simple, this approach is problematic. The empirical evidence that such transactions do not actually happen in the real world suggests that something must indeed be wrong with this approach. As we will see in Chapter 4, even when this approach is possible, agreeing on commercial terms can be expensive and time consuming.

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.

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 would be traded? There are two possibilities: spectrum (where CCNs pay spectrum license holders to be able to deploy their systems) or some form of credit for USOs (where CCNs bid to accept USOs in some area).

Second, once there are market transactions in an asset or liability, the asset or 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. 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 transactions. The NGSM approach to GSM whitespaces avoids having any dollar transactions by directly giving primaries USO credit from the regulator for the actions of a third party. This avoids the risk of setting a price on service obligations. Further, by leveraging a database system to facilitate access, NGSM's approach enables such exchange without direct coordination between primaries and secondaries, minimizing transaction costs.

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.

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.

Relationship to 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. [146] propose reusing timeslots in a GSM cell for adhoc networks during periods when the GSM cell is lightly utilized. Buddhikot et al. [31] describe a system for indoor femtocells to dynamically share spectrum with incumbent carriers by operating over ultra wide bands. Yin et al. [176] proposes a similar system and provides measurement results indicating 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 whitespaces 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.

3.9 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. Community cellular networks appear to offer substantial advantages for providing sustainable rural service without subsidies or external support, but 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.

Chapter 4

Scaling Community Cellular Networks with CCM

Using GSM whitespaces, systems like Nomadic GSM could enable a “use it or share it” policy regime within the cellular bands, giving full flexibility to community cellular networks while preserving the protections needed by incumbent licenseholders. While we argue that our proposals are feasible, the approach we outlined in the previous chapter still requires significant regulatory change, which can take many years. Are there other approaches we can take that are feasible within the *existing* regulatory framework?

One approach we could take is to somehow encourage MNOs to willingly allow community cellular networks to operate within their spectrum assignments under commercial terms. To do this, we need to facilitate cooperative interaction between community networks and incumbent MNOs. Thankfully, such opportunities readily present themselves.

We have already seen that exclusive spectrum licensing of cellular spectrum effectively allows only a small set of nation-scale actors to provide commercial services in most countries. Community cellular networks also face other practical barriers to entry such as securing necessary operating licenses to obtain access to telephone numbers and interconnection agreements with the global phone network. And, in spite of their size, incumbent mobile network operators (MNOs) are capital constrained and struggle to justify investment in rural areas with marginal business cases when compared to more profitable and lower-risk urban markets. At the same time, these same MNOs benefit from improved rural connectivity both directly in the form of increased traffic between rural users and their existing subscribers, as well as indirectly through decreased pressure from regulators to expand their infrastructure.

As we have seen, community networks built “by and for” the people they serve in a community-centric, often cooperative, fashion [23, 29, 167], allow for creative localized schemes for sustainable operation that work even when covering low-income, remote, or sparse communities. The strength of the community network model is its ability to flexibly adapt to local needs using local resources and local insights. A network serving one community can provide services suited to that community, and can institute their own policies that would not make sense in a nation-scale network (for example, free local SMS mailing lists). They can also tailor their network deployments to leverage local capacity, such as existing towers and power systems, lowering deployment costs.

We are left with an unfortunate market failure: incumbent MNOs hold resources and capabilities, in the form of licenses and interconnect, necessary to provide rural coverage, but can't justify investment in underserved areas. Concurrently, community networks could provide sustainable service, but lack a combination of the commercial, legal, and technological capacity to operate in a telecommunications regime designed for large carriers. We argue that rather than standing in opposition to each other, incumbent operators and community networks have an opportunity for cooperative interaction, leveraging the strengths of each to fulfill mutually beneficial business and social goals. Fundamentally, such a partnership requires a technology stack enabling collaborative *resource management* and *provisioning of common services* between MNOs and community networks.

In this chapter, we present CommunityCellularManager (CCM), a solution for operating community cellular networks at scale within the existing telecom ecosystem.[1] CCM enables multiple community networks to operate under the control of a single, multi-tenant controller, preserving flexibility for each community network to implement their own unique services while allowing them to share MNO resources, such as spectrum and phone numbers. For MNOs, CCM resolves the administrative problems they face granting many small third parties access to their core network by aggregating their traffic behind one logical point of interconnection, while providing a point of control to mitigate the risk they take on by sharing their resources and licenses. We evaluate CCM in a deployment to expand the service of an MNO into unserved areas via community networks using the MNO's spectrum and interconnect (Section 4.4).

The contributions of this chapter are as follows. First, we describe the design and implementation of CommunityCellularManager, a system for deploying rural community cellular networks at scale by facilitating cooperation and resource sharing with traditional mobile network operators. Second, we present a multi-year, large-scale community network deployment in the Philippines that provides basic communications service to 17 communities and over 2,800 monthly active users, powered by CCM and in partnership with Globe, the largest MNO in the country. We further demonstrate CCM's ability to support key use cases necessary for rural community cellular networks, such as offline operation, as well as its ability to support independent local services on a subset of the community networks in our deployment.

4.1 Related Work

Rural Access. Access to connectivity supports numerous services, such as health [107], education [105], and finance [28]. However, access remains unequal, with connectivity

---

[1] This chapter is based on work that appeared at NSDI [82].

significantly better in dense, urban, and developed environments than in sparse, rural, and developing ones. In attempts to resolve this disparity, researchers have proposed a variety of novel technologies such as long-distance WiFi [134], delay tolerant networking [56], and “sneaker” nets [69, 136, 77].

Although these technologies were all successes from a research perspective, they have had limited impact on the global problem of access. These technologies exist alongside an array of other important infrastructure, and innovations at the network layer alone will be insufficient. Surana et al. [157] find that looking at the system holistically is critical for its success [156]. This includes both physical infrastructure such as power *and* the social structures and people surrounding the technology. We take these lessons to heart in CCM, bringing issues like local ownership and customization to the forefront of the system design.

Community Networks. Researchers have explored wireless community networks [98] for years with contributions such as specific wireless technologies [62], topologies [168, 167], and mesh protocols [20] more appropriate for their unique constraints. The above systems all use WiFi for access, primarily because of low cost and use of unlicensed spectrum. These networks have shown to have the potential to empower local communities [48] and create more resilient networks [23], but the use of WiFi as an access technology limits these networks to higher-end devices and small coverage areas without mobility.

More recently, researchers and practitioners have built community *cellular* networks [9, 88, 143, 45]. Ideologically aligned with the goals of community networking, these networks use cellular protocols to provide wide-area coverage to basic phones at lower cost. Unfortunately, they remain limited by the technical affordances of cellular, including the need for licensed spectrum, phone numbers, and interconnect with the global phone network. CCM provides a platform for community cellular that resolves these concerns.

Mobile cores. Cellular core network standards are developed by the 3GPP [132]; recent work focused on improving scalability and reliability of core networks. PEPC [139] implements a high-performance LTE core network (an Evolved Packet Core, or EPC) by consolidating per-user state typically distributed across distinct network elements to reduce synchronization between core network and facilitate scale-out. SCALE [21] decomposes the standard MME (Mobility Management Entity, responsible for mobility signalling within an LTE network) into two components, load balancers and backend workers, to facilitate virtualization and thus scaling. KLEIN [138] considers the broader design space afforded by virtualization of mobile core network elements and presents a backwards-compatible mobile core architecture designed to leverage network function virtualization. Similar to SoftCell [101], CCM is a clean-slate architecture, and was designed for public cloud environments like ECHO [123]. While each of these systems considers a refactoring of the traditional mobile core architecture, CCM differs in its focus on shifting the *administrative boundaries* of mobile core network elements to enable cooperation between MNO and community networks.

4.2 Design Goals

CCM is targeted to a specific but important use case: extending mobile network service provided by an incumbent MNO in a cooperative partnership with community cellular networks. As we have discussed earlier, our interest in expanding community cellular networks stems from their demonstrated ability to provide sustainable service in areas unserved by MNOs. Such cooperative arrangements [40, 113] promise to unblock regulatory and commercial barriers faced by community cellular networks, such as access to spectrum. The architectural rigidity of mobile networks, however, makes achieving our federated use case challenging, since it provides no affordance for having multiple administrative domains providing service in a loosely coupled fashion.

The 3GPP specifications [132] define the architecture of traditional mobile networks. Although these specifications continue to evolve, the network architecture prioritizes (1) mobility, (2) operator control of subscribers and network policy, and (3) efficient utilization of radio resources as first-order priorities. To do this, 3GPP networks tunnel traffic between end-user devices and a (logically) centralized "core network", which provides services such as voice calling, messaging, and data service. These logical services are offered over the physical radio access network (RAN), which is responsible for providing users with physical connectivity to the core network and ensuring efficient use of the limited radio resource. The two components of the network – the core and the RAN – are tightly coupled: without a core network, the RAN is not capable of providing network services to end users. This approach stands in stark contrast to that of the Internet [37], with its focus on supporting datagram services among independent networks.

Both architectural approaches – 3GPP's centralized control and the Internet's decentralized flexibility – have clearly enabled successful networks, but the 3GPP architecture has several significant drawbacks for community networks in rural areas. First, the strength of community cellular networks is their adaptability to their local context, providing customized services that are relevant to users in the local community. Centralized core networks make a single entity responsible for all service provisioning, preventing innovation at the local level. Second, community cellular networks are heterogeneous and administratively decentralized, with diverse IP transport and transit configurations and different sets of actors responsible for network operation and maintenance at each site. Tightly coupling the RAN equipment deployed in a community to a core network requires coordination between the community network and the MNO's core network. Fundamentally, this requires the MNO to expose their most critical network assets – their core network – to a third party, so they need ways to mitigate the risk of doing so. As a result, this interaction incurs high administrative overhead due to the configuration required for each new site and the extensive security reviews required when connecting third-party infrastructure to MNO core networks.

These challenges suggest the following design goals:

1. Enable community networks to *deploy local services* within their footprint, while pro-

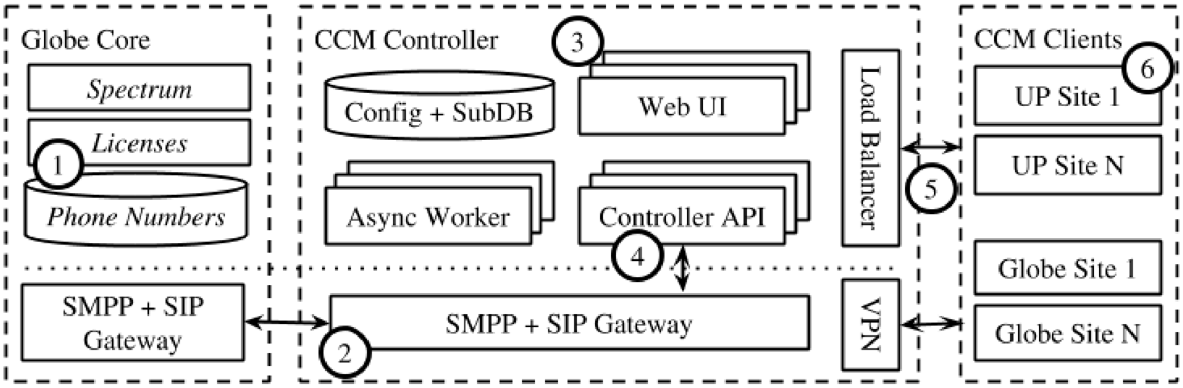

Figure 4.1: The CCM high level architecture, consisting of Clients, the Controller, and the interconnect with Globe. Phone numbers, licenses, and spectrum [1] are provided by Globe, with a number block manually allocated into the Controller's subscriber database (SubDB). The CCM Controller [3] manages network configuration and status, provides a point of control for Globe to enforce global network policies and perform administrative tasks, and a web interface for network management. The CCM Controller also interconnects with Globe's core for routing traffic between many individual community networks and Globe's network [2]. Clients communicate with the Controller via an HTTPS API [5], which is also used to determine call and SMS routing [4]. Note that CCM utilizes local breakout at each component, so traffic may be handled as the SMPP/SIP gateways closest to the two users. CCM is multitenant, so Clients in many communities [6] can have unique local services and operate asynchronously if the backhaul is unavailable.

viding a basic set of services from the mobile network operator.

2. Minimize the cost and risk of *adding new cell sites and community networks* for both the MNO and the community network operator.

3. Provide service at *low incremental cost* beyond that required for RAN deployment, even in small networks.

4.3 System Design

To achieve these goals, CCM builds upon our experience building and operating an early community cellular network [88, 87]. We implemented the first version of the software that would evolve into CCM on location during our fieldwork in Papua, Indonesia, and the constraints of that environment – reliably unreliable satellite Internet and off-grid power – shaped our design.

In CCM, *each network node can provide service autonomously*, even without a connection to the Internet. The ability to operate without a connection to the Internet (that is, autonomously) for extended periods of time was directly informed by our experience in Papua;

by keeping autonomous operation as a top-level design objective, we ensured that our system could fail gracefully in the face of unreliable supporting infrastructure. This decision would define the overall system architecture.

CCM is broadly divided into two components (Figure 4.1). The first is the CCM Client, which is co-located with a community network's RAN equipment and fills the role typically played by a 3GPP core network, terminating the logical connection from a user's phone and providing services to the end user. To support voice and SMS services, we translate all 3GPP voice and SMS to SIP and SMPP at the Client, which allows us to perform call and message routing at the Client (also knon as *local breakout*). By making each network node a logically independent cellular network, CCM elegantly supports autonomous service deployment in each community network and allows core network capacity to scale up and down with the deployed RAN network.

The second component, the CCM Controller, provides a set of common services across community networks, such as RAN configuration and network management tools, including a web application that serves as the primary user interface for both MNO staff and community network operators. Importantly, the Controller also manages state distribution across the different networks: rather than centralizing both policy and network state in on-path devices, we synchronize application state among network nodes. The Controller also provides interconnection between community networks and the MNO network, regulating access to one of the most important resources the MNO owns: globally routable phone numbers for voice calls and SMS.

This decomposition bridges the centralized world of traditional mobile networks with the decentralized world of community networks and IP. Network management functions – subscriber provisioning, device management, network policy, and necessarily integration with Globe (Section 4.3) – are all highly *centralized*. In contrast, user traffic and signalling (the data and control planes, from a 3GPP perspective) and service provision are *decentralized* where traffic doesn't interact with Globe. This hierarchy enables scale by allowing individual community networks to take advantage of common infrastructure while preserving their ability to deploy new services. It also reflects a pattern we have observed in other scaled (non-cellular) community networks, with "second level organizations" solving expensive or complicated problems once while more local groups operate network infrastructure [17].

CCM is designed for GSM (2G) networks and provides only voice and SMS service, not data (though we do rely on IP data for management, we do not provide data service to users' phones), as the plurality of devices in our deployment only supports 2G [150]. Similarly, we emphasize that our particular decomposition of responsibilities represents just one point in a wider design space with its own political and technical tradeoffs. Section 4.6 discusses extensions toward LTE and further exploration of this design space.

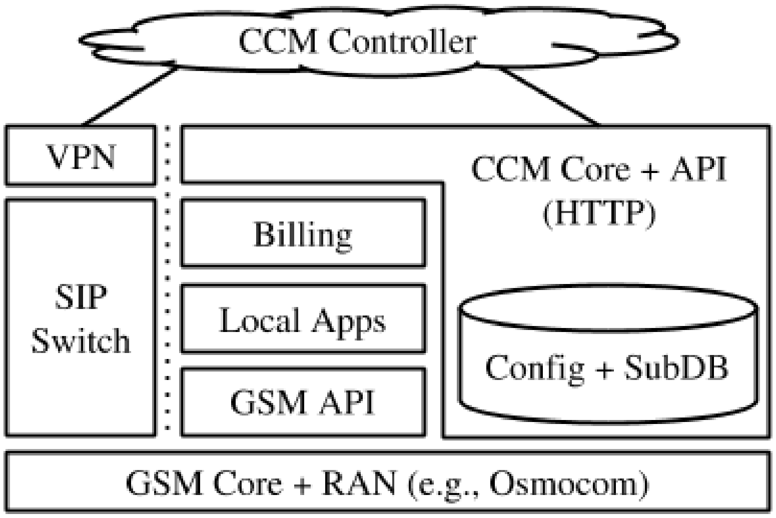

Figure 4.2: The CCM Client is deployed at each site and provides service even when the connection to the Internet is not available. The Client also provides an opportunity for customization of local web services and local applications.

CCM Client

The CCM Client is responsible for all operation at the community network site. It has two key responsibilities: provide connectivity services to users at the site, and manage the site itself. Physically, the CCM Client is a suite of software tools that runs on a off-the-shelf x86 PC co-located with the radio equipment of a rural cellular site. In our deployments, this PC is either located physically within the RAN enclosure or located nearby in a separate enclosure. Figure 4.3 shows the components of one example install.

The Client is the primary element in our system responsible for delivering communication service (in our case, voice and SMS) to end users. We designed CCM to work with the two most popular open-source GSM network implementations, OpenBTS [127] and Osmocom [128], each of which implements a complete GSM core network and, importantly, translates cellular traffic and signalling to VOIP (SIP, RTP, and SMPP). CCM extends this translation and implements the business logic to properly route calls and SMS, implement local services, and provide basic business support services such as charging and user provisioning.

User management and billing. In GSM networks, the SIM card holds a unique IMSI number, the user's identity on the network. Each IMSI corresponds to a particular subscriber's profile in CCM, which associates both a phone number and an account balance with the subscriber. Similar to previous work [88], we provide basic SMS-based applications for provisioning a new SIM card (i.e., associating it with a subscriber profile and assigning a phone number) and transferring network credit among users. The network operator may add or remove credit directly to a subscriber's account via the CCM Controller as well.

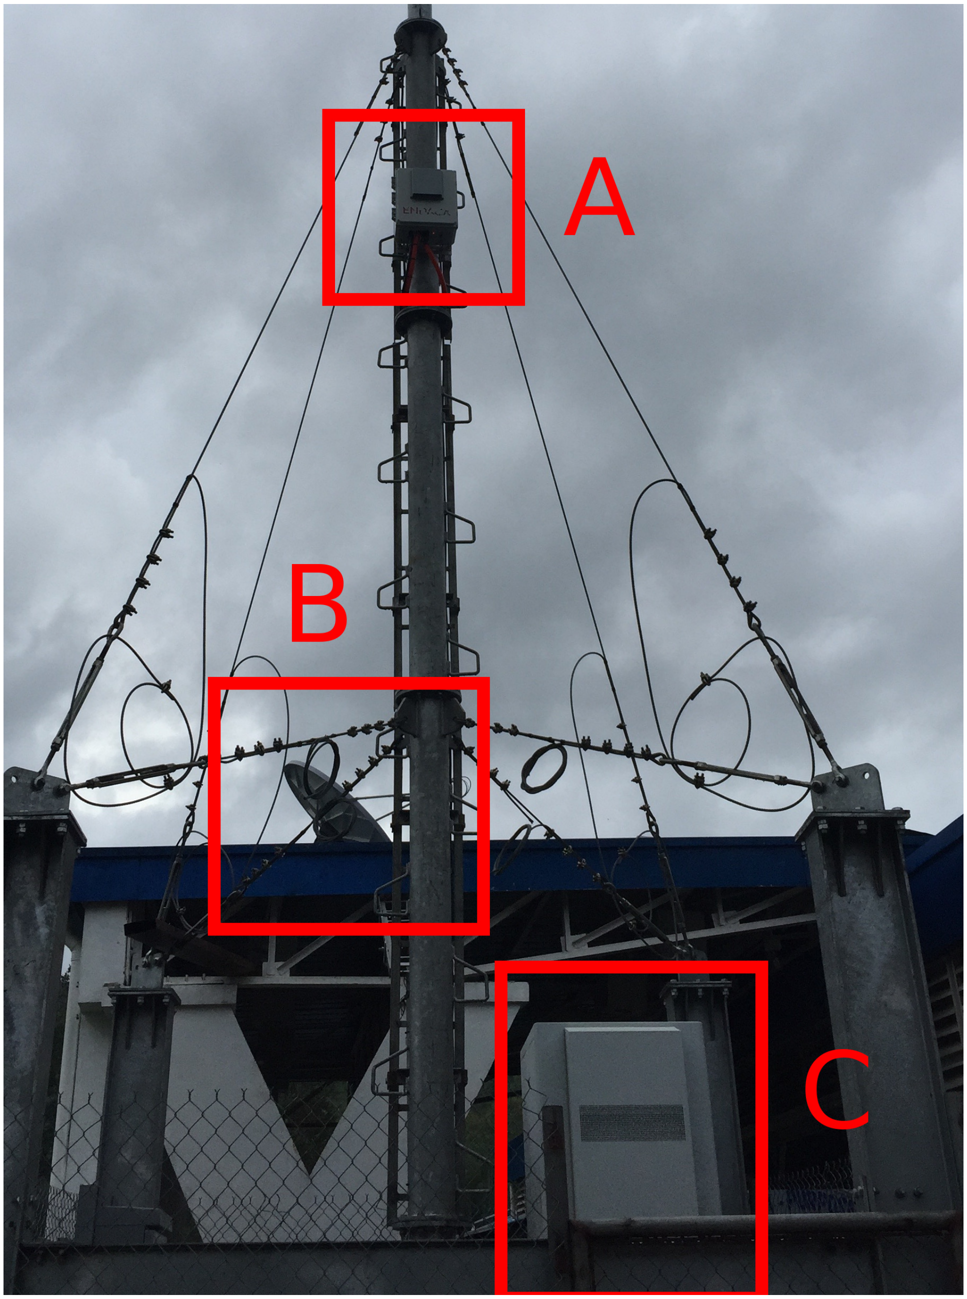

Figure 4.3: The CCM Client hardware components. (A) PC and GSM radio (in combined enclosure). (B) The satellite antenna used for backhaul and solar panels (not visible). (C) Closet for power system and satellite modem.

CCM implements a simple prepaid billing system in which operators can set the price of different network actions (e.g., price per minute for a voice call to a certain number prefix); time-based account balances are also supported. This user profile is one of the key elements of network state in our system, which we discuss further in Section 4.3.

Local breakout. We run a SIP softswitch [63] on each CCM Client to route calls and SMS. Because CCM leverages local breakout, this softswitch can fully route and connect local calls and messages without having any traffic (signalling or media) leaving the site. This gives CCM the ability to support arbitrary local services, including custom interactive voice response applications or SMS shortcodes. Any outbound traffic addressed to a user not currently active on a site is routed to the CCM Controller, which allows the MNO to set policy for these communications and potentially route them between different sites as well as to the global telephony network (Section 4.3).

Site provisioning. Community networks are built and operated by ordinary people, not specialized technical staff. As a result, we focused on automating as much of site provisioning and setup as possible: our goal was to make setup at least as easy as installing a home WiFi router. To provision a site, each CCM Client generates and assigns itself a unique identifier during a "manufacturing" phase; we expect that either the vendor or experts assembling the unit will perform this step. The unique identifier is printed on the outside of the hardware to be deployed at the site.

Devices start in an unprovisioned state; CCM prevents the radio equipment from transmitting while unprovisioned. To provision the device, installation technicians or community members register the device by inputting the unique identifier into the CCM Controller's web interface. Once the device is powered on and connected to the Internet, it connects to the CCM Controller to provision itself. If the device has been registered with the CCM Controller, the device will receive VPN configuration information and a secret API token for future requests to the CCM Controller.[2] Using the API token, the site can begin a process of generating secondary credentials required for operation, such as a VPN keypair. Once the Client establishes secure connectivity to the Controller, configuration and network state information is synchronized and the device is ready to begin operation.

Under nominal conditions, the only user interaction required to provision a site is entering the unique ID into the web UI (see Figure 4.4). No prior coordination is required between equipment providers, the CCM Controller operator, and the community network team actually deploying the site. This is in contrast to traditional 2G networks, which typ-

---

[2] The gap between when a device is registered with the CCM Controller and when it actually connects could allow an attacker with knowledge of the unique ID to register a fake device with the CCM Controller. To mitigate this, we instruct users to register the device only after they've verified it is powered on and connected to the Internet. In practice, this is done in a "staging" lab environment prior to field installation, thus minimizing the window of opportunity for an attacker. This attack could be completely mitigated if the device-specific secret API tokens or a keypair were pre-provisioned with a CCM Controller prior to associating a device with a user's account; however, doing this would require coordination between device manufacturers and the entity operating the CCM Controller, which was not feasible in our deployment.

Tower: VBTS [redacted]

from 2018-09-13 at 5:29PM to 2018-09-20 at 5:29PM

from 2018-09-13 at 5:29PM to 2018-09-20 at 5:29PM

Figure 4.4: A screenshot of the CCM web user interface.

ically require RAN, core, and network management software from the same vendor to be manually provisioned by one entity.

CCM Controller

The CCM Controller is responsible for managing traffic among the collection of community networks it manages, as well as with the outside world. Moreover, while in practice CCM supports geographically distinct cell sites without overlapping coverage, users can move between networks and CCM must be able to coordinate state across networks; the Controller also handles this task.

Voice and SMS Interconnection

Internally, CCM routes all calls and SMS between SIP switches at each CCM Client site. Within the administrative domain of a single CCM Controller, we can assign phone numbers to subscribers much like an enterprise can assign extensions to internal users or RFC1918 addresses for private networks. However, for users of CCM networks to make and receive communication from the global phone network, we need to interconnect CCM sites with an entity that has been assigned globally-routeable phone numbers and has the ability to route voice and SMS communication. CCM is capable of interconnecting with wholesale VOIP providers (such as Nexmo [121] or Bandwidth [19]) as well as an MNO's VOIP infrastructure; we will focus on the latter in this work, though the mechanism is similar for both.

CCM interconnects with an MNO through its VOIP gateway (typically used for supporting enterprise customers) as well as the MNO's SMPP gateway, an industry standard for SMS exchange. In this arrangement, the MNO allocates the CCM Controller blocks of phone numbers to be used exclusively by users of CCM sites. On the CCM side, we run a corresponding SIP switch and SMPP gateway, as well as any VPN infrastructure required for connecting to the MNO's systems. The CCM SIP and SMPP infrastructure is stateless, and determines inbound call routing by querying the CCM Controller API to determine subscriber location; all billing and charging is handled on the CCM Client.

The integration architecture we use has two consequences. First, since the CCM Controller handles multi-tenancy, there is a single point of interconnection between the MNO and potentially many community networks. This is crucial: setting up our initial integration with Globe took just over four months from start to finish, and required the team we were working with at Globe to obtain approvals and request configuration changes with a number of different teams within their organization. Further, the technical work required to set up this interconnection, while simple in theory, is complicated by the fact that MNOs rely on custom and legacy systems that must be carefully managed to prevent downtime. Going through this process for dozens of community networks would be impossible, so while our integration approach creates a single point of failure, the benefits of a one-time integration outweigh the risks. These risks are also straightforward to mitigate: the network functions involved (SIP, RTP, and SMPP gateways for signaling and traffic, as well as IPSEC VPNs for transport) are commonly deployed in "high-availability" configurations in production telecom networks. Further, the protocols we use have some robustness to failure: the loss of a SIP gateway would not impact ongoing call traffic since the user traffic is handled separately by the RTP media gateway, and SMS messages transferred via the SMPP gateway can be queued and resent once connectivity is established.

We also note that such a single point of integration provides a useful isolation between the IP world that CCM operates in and the legacy technologies used by the MNO. This extends beyond merely the technical choices used within each domain to the overall design and operational philosophy embraced by each. For example, CCM was designed to receive regular automated software updates to both the Controller and the Client, leaving us free to innovate within the CCM domain as long as we maintained only this single interconnection abstraction. This is distinctly different from the approach most MNOs take to evolving their network infrastructure, where the status quo is best described as "if it ain't broke, don't fix it". Having a single logical point for integration also acts as a sort of firewall, allowing the entity operating the CCM Controller to filter (potentially malicious) traffic from Clients before it reaches the sensitive core network infrastructure of the MNO.

Second, this tightly ties CCM networks to the MNO who assigns phone numbers: the MNO can at any point shut off service to these users; users would need new phone numbers if the partner MNO changed. This is advantageous from the perspective of the MNO, since they retain ultimate control over their users, which was important for gaining approval from Globe to allow community networks to operate under their spectrum and operating licenses.

While other approaches [84, 57, 143] enable community cellular without hard dependencies on third parties, our MNO partnership approach required no regulatory change; the project would not have been feasible in the Philippines in the timeframe we've taken otherwise.

State Management

System state in CCM consists of per-site configuration, network policy, and subscriber authentication, billing and location data. In order to support disconnected operation, CCM needs to ensure this state is distributed across all sites administered by a CCM Controller.

Checkin. The fundamental mechanism for state distribution in CCM is the "checkin". At least once per minute, each CCM Client sends a HTTPS request to an API endpoint on the Controller. The content of this request includes the site ID, usage logs (also known as call data records, or CDRs), diagnostic information (such as CPU utilization), and subscriber data to the Controller. Based on the computed configuration for the site, the Controller's response to this checkin request is the desired state of the system's configuration and the set of subscribers a site should be able to serve.

Both the request and response are JSON dictionaries divided into independent sections: for example, billing records are stored in the `usage` section. Only changed state is transferred in the server response after the initial checkin. The Client and Controller each maintain a shared context of previously received configuration by tracking a hash of the contents of each checkin section; the Client includes the last hash of each section it has received in its subsequent requests. If the Client hash matches the last recorded Controller hash, it only sends the difference between the last state and the current one. Otherwise, it sends the full contents of the section. Because most configuration changes rarely, this minimizes received checkin response sizes.

Subscriber data. Subscriber data consists of authentication information, location (i.e., which Client a user is currently attached to), and billing information. The first two of these are straightforward to handle. Authentication information consists of SIM card keys as well as whether a subscriber is allowed or not allowed on the network; the network operator can add users via the Controller or through an SMS-based short code application in the field. In both cases, this information is directly written to the Controller and replicated to all sites in the network. Each CCM Client reports the list of subscribers currently attached to the site during checkin; this allows the Controller to have a global view of subscriber location for routing inbound traffic and traffic between sites.

Billing state is more complicated, as it can be mutated both at any particular site (e.g., decremented after a subscriber makes a call) as well as at the CCM Controller (e.g., when an operator adds credit to a subscriber's account directly). A site may also be disconnected from the CCM Controller for arbitrary lengths of time, while still providing communications services within its coverage area. Nevertheless, we need to be able to apply network policies around charging to local calls while disconnected from the tower.

To achieve this, we restrict modifications to subscriber balance to commutative operations (add or subtract), and then represent subscriber balances as a CRDT [109]. Subscriber balances are synchronized across the network during the checkin process: as each site checks in, the Controller merges its subscriber balance state with what the site reports, and provides the merged state back to the site in the response if it differs. This enables each node of the system to both read and write subscriber balance independently – when a site is disconnected, users can continue to communicate locally and transfer balance between each other, and balance will converge once the site comes back online. Users can also move between multiple offline sites and communicate locally on each.

This raises the possibility of double spending and presents the challenge of maintaining the desired invariant (users must have sufficient balance in their account before using the network) during partitioned operation. Much like the Brewer's ATM machine that favors availability over consistency to maximize revenue [30], our partners made the business decision to allow the possibility of overspending to facilitate usage. Resolving the invariant violation after consistency is restored is straightforward: users who do somehow overspend will hold a negative account balance, and must buy enough prepaid credit to make their account balance positive before they are able to continue using the network. In any case, the risk of this particular attack is low in practice because of the substantial distances between different sites, and could be further mitigated by setting a threshold for spending during disconnected operation at each site.

Site Management All other configuration state is similarly managed by the Controller. We provide a web-based UI for viewing network-wide state (including user activity) and defining network policy and configuration. Some information must be manually defined by the network operator (such as pricing plans), and other information is automatically generated to simplify operation (such as the radio channel that each site uses). Outside this controller-based interface, we do not support any other means of configuring devices.[3]

The UI provides fine-grained permissions and control, allowing the operator of the Controller (Globe) to determine which types of users may access different aspects of the UI. In our deployment, spectrum usage is controlled by Globe, and access to Globe sites is restricted to their staff. Globe has allocated a separate "network" for a set of sites run directly by communities, but has restricts what administrators of these networks can do. The superuser has view and edit access to change configuration across sites, networks, and subscribers, as well as the ability to adjust credit and download activity reports. The superuser can also create additional users based on predefined roles: business analyst (view-only), partner (view-only + manage subscriber) and loader (view-only + manage subscriber + adjust credit). The system allows Globe to adjust specific permissions per user.

---

[3] Each Client site does have a command line interface, but this is only used for debugging or emergency recovery purposes.

Figure 4.5: Components of a typical UP community cellular installation: solar panels (top left), battery bank and network hardware (middle left), VSAT antenna (bottom left), outdoor GSM radio system, tower and shelter (right).

4.4 Deployments

We deployed CCM as part of a long-term partnership between the researchers, the University of the Philippines and Globe Telecom [66], the largest MNO in the Philippines with over sixty million subscribers. The researcher-driven portion of the project, conducted in the Aurora province, is funded by the Philippine Government and focuses on bringing the benefits of community cellular to remote parts of the Philippines; we refer to these as the “UP sites”. The MNO created two different administrative domains for their subnetworks, one owned entirely by Globe and another for the UP sites. For both of these networks, users are using phone numbers from Globe, though the service is branded separately from their main service bundles to make clear to users that the services and quality expectations on these extremely rural sites are different from Globe’s main network. In this section we describe each of these subnetworks including their pricing, interconnect, and context. For site deployment dates and locations, please see Table 4.1.

UP Community Networks

UP selected the province of Aurora for their CCM installations. Aurora is a coastal province of Luzon, the main island of the Philippines, on the Pacific Ocean and the Philippine Sea. The capital, Baler, is a relatively affluent town with readily available connectivity and robust land routes to urban areas such as Metro Manila. Further south down the coast, dozens of fishing communities lie outside of any existing connectivity. In these areas, people travel by boat to Baler for any needs that cannot be met in their home communities, including device repair and connectivity. Despite the distance from coverage, initial site surveys done in 2016 found that mobile phones were prevalent and demonstrated demand by uncovering a variety of existing (though complex) connectivity solutions [22]. These communities were selected by the UP team for (1) their relative proximity to Metro Manila, (2) lack of connectivity, and (3) connection to a local university for assistance in deployments. These deployments began in late 2017, with all seven planned sites launched as of February 2019. The community GSM sites provide voice and SMS access to over 1,500 subscribers.

Structure and Context. The Aurora networks are organized in line with most definitions of community networks [167]. The operation and maintenance of the community network is handled by local cooperatives, in partnership with the local government units (LGUs), the local state college and the UP research team. These partners were identified during the initial site surveys. The LGU facilitated the legal appropriation of land for towers and helped expedite the permits and clearances required for civil works at each site. In addition, the LGUs extended assistance to mobilize local labor for installation and deployment activities.

The cooperatives conduct the day-to-day operations, business management, and maintenance of the network. They are in charge of pre-paid credit distribution from Globe to the community retailers, who are mostly existing *sari-sari* (general merchandise) store owners. The cooperative orders credit from Globe at least once a month. After payment has been sent through a bank deposit and verified, Globe tops up the cooperative's accredited mobile number via the CCM Controller web interface. From here, the cooperative distributes it to their authorized retailers via a locally-hosted SMS-based credit transfer application. The cooperative receives a wholesale discount from Globe, part of which is passed on to the retailers. The retailers earn additional income by charging a small "convenience" fee per sale. Finally, the gross revenue from all charged calls and texts on the network is split based on a revenue sharing scheme between Globe and the cooperative, where the cooperative gets 80% and 20% goes to Globe. Earnings are used by the cooperative to pay personnel and as savings.

In terms of maintenance, the researchers employ a three-tier support system. The first tier, L1, is composed of local maintenance personnel hired by the cooperative and residing in the community. They are assigned to do daily upkeep and basic troubleshooting of the cell site. Any issues that are not resolved at the L1 tier are escalated to the local state college and/or the LGU, the L2 tier, which can provide intermediate-level technical assistance.

Figure 4.6: Conclusion of the UP network launch in Dibut (left) and deployed site in Diotorin (right).

Finally all other issues that are not resolved at the L1 or L2 tiers are escalated to the research team, which provides L3 support.

Prior to launch, the researchers facilitated social enterprise training sessions with cooperatives that had no prior experience in conducting business-related activities pertaining to operations of a community network, and technical training with community maintenance personnel. During the network launch, the researchers held a forum with the broader community where the unique properties of their networks were explained (e.g., lack of roaming to Globe's main networks) to the subscribers. The event was also a venue for the community to raise questions and concerns regarding the network and services, and distribute SIM cards. The distribution of SIMs in the UP networks is currently tightly controlled as part of an ongoing randomized control trial [49] on the impact of cellular networks on rural communities.Requiring dedicated SIM cards for these networks makes them opt-in and allows Globe to emphasize to users that these networks are not the standard Globe service offering, protecting their brand.

Globe's Community Networks

Concurrent to the installation of the UP sites, Globe also installed CCM Client access points in eleven other rural communities throughout the Philippines. Of these, the first two were "proof of concept" (PoC) sites installed in Tanay, Rizal province (60km from Manila) and another in Talisay, Quezon province (130km from Manila). These two sites can be seen in Figure 4.8. Following the successful trials of CCM in these two PoC sites, a further eight networks were deployed in Eastern Visayas. All of the sites are rural and lack any existing network coverage, with populations ranging from one to five thousand people. Though census information for rural areas in the Philippines is spotty, the demographics and economics of Tanay have been described in detail in other work [150]. Another twenty networks are

Figure 4.7: Drone shot of UP site in Dilasag, Aurora.

Figure 4.8: Globe CCM towers in two PoC sites. Tanay (left) and Talisay (right).

planned with rollouts expected throughout 2019.

Structure and Context. The MNO's sites consist of two groups: the PoC sites in Luzon and partner sites in Visayas. Both are organized in a more traditional fashion with Globe handling marketing, credit distribution, and installation. For the PoC sites, Globe agents conducted selection interacting with the LGUs to procure locations and timings amenable to the local community. For the Visayas sites, Globe instead partnered with a non-governmental organization (NGO) to find partner communities and negotiate installations. This NGO has a long history of projects in Visayas and was able to find suitable communities in rural areas as well as assist in the day-to-day operations of the sites. The NGO also handles the SIM card and pre-paid credit distribution from Globe to the community retailers.

The installations were done by Globe's tower deployment team using their standard site equipment, aside from the custom RAN equipment itself. The sites use a two-tiered maintenance system with one level of lightweight local support and the main support provided by Globe staff in Manila. One local community member is selected by the deployment team and tasked with using the network to send messages back to Globe's technical team in case of ongoing issues. If any failures disable the network, Globe's engineers use the CCM Controller's web interface to observe the network failures and send engineering staff out to resolve the issue. Credit sellers buy their credits by depositing money into Globe's account at a nearby bank outside of the community. They then take a small premium when distributing these credits throughout the community.

Project Evolution

Rural connectivity projects do not lend themselves to neat stories, and this deployment was no exception. Our deployment of CCM originated from discussions between Globe and the researchers' company, Endaga, in early 2015, and implementation work began in the summer of 2015. Originally, Endaga operated both the CCM Controller and directly supported all field equipment running the CCM Client software. After Endaga joined Facebook in late 2015, CCM was released by Facebook as an open source project to enable continued development by Globe and others. At the same time, we interconnected the CCM Controller with Globe's core and then transferred administrative control of the CCM Controller to Globe in mid-2017 after a successful early proof-of-concept. The researchers continued to be involved in the development and aspects of the deployment of CCM, but after this point day-to-day operation of the service passed to Globe; this continues today. In practice, this means that Globe controlled access to the CCM Controller (including software updates), and the researchers could only deploy local services on the sites we directly had access to: the UP sites. This deployment configuration was not what we originally designed CCM for, but CCM was able to continue providing an interconnection abstraction towards different groups of community networks, even while the administrative boundaries between system elements shifted. In this partnership, the researchers set up and operated CCM on AWS and interconnected it with Globe's core network as described in section 4.3, and then transferred administrative control of CCM to Globe in mid-2017, after an early proof-of-concept deployment.

4.5 Evaluation

Our team completed interconnection between CCM and Globe's network in mid-2016, with two trial sites launched in early 2017. After a year of evaluation, Globe launched nine additional sites and granted permission to the UP to launch seven community networks that year (Table 4.1), of which six are live as of December 2018. While a total of 17 sites have been launched throughout the project, only 9 were in operation as of December 2018 due to hardware failures; the failed sites are being replaced with new hardware that is currently undergoing testing. The UP sites used a mix of hardware (including some self-assembled), and their deployments have not been impacted as severely as Globe's sites, which used a single vendor for their deployments. Supporting heterogeneous hardware was not an explicit goal of CCM, but its ability to do so proved useful in our challenging rural context.

Site launch dates. Sites were launched throughout the three year duration of the project. Initial testing at "proof of concept" sites began as early as 2016, and the bulk of sites were deployed in late 2017 and throughout 2018. We note that while the UP sites were deployed at a rate of less than once per month, Globe sites were deployed in bursts, with as many as four sites being deployed within the same week. Table 4.1 depicts launch dates and locations for each site, and Figure 4.9 shows the locations of deployed sites.

Site Costs

Tables 4.2, 4.3, and 4.4 provide an example of site fixed costs and operating expenses for a typical site in our deployment. Note that prices are in USD, and are approximate. These prices are estimates that reflect cost of equipment once it has already cleared customs and is in country. Similarly, civil and telecom works costs can vary depending on the difficulty of access to a particular site.

Globe preferred to have a single vendor for each component, with the goal of standardizing their deployments and reducing costs. In contrast, the UP installations used different setups and vendors due to a combination of funding limitations, procurement difficulties and delays. For example, we experienced challenges in the procurement and importation of GSM radio hardware. As such, some sites used equipment assembled from spare components we already had in stock.

The monthly backhaul subscription constitutes the bulk of the OpEx costs. Since the sites are very remote, the only feasible option is VSAT (a satellite Internet technology) which is also expensive. Backhaul prices vary significantly depending on the provider; Globe has existing bulk contracts for capacity as well as their own VSAT hub and network infrastructure, lowering costs compared to end-to-end VSAT providers. Other OpEx components

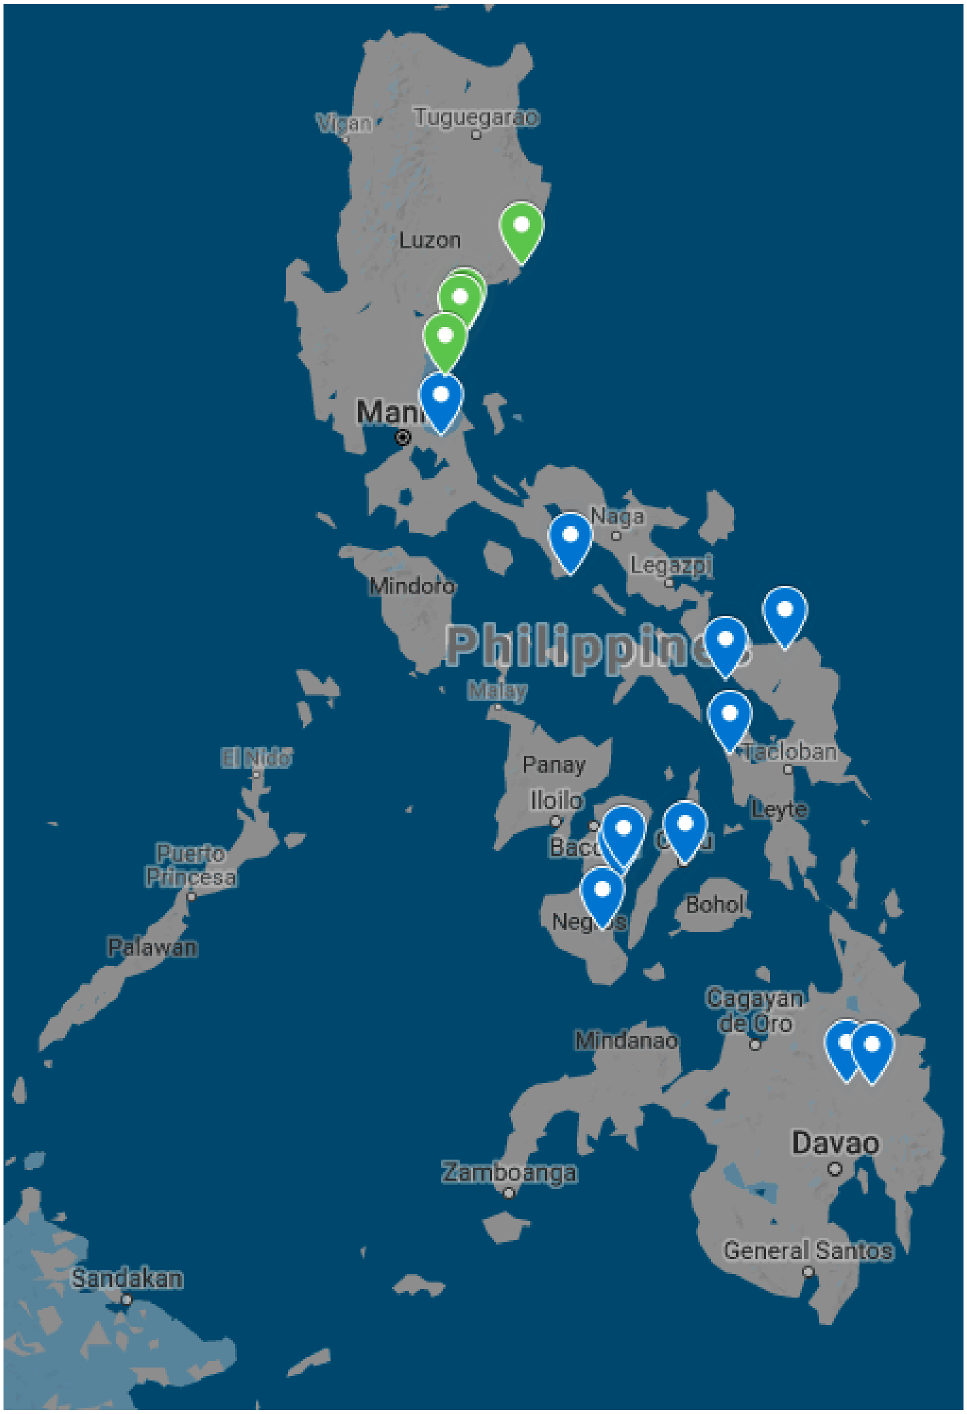

Figure 4.9: Map of deployed sites. UP sites are green, Globe sites are blue.

Site Name	**Type**	**Commercial Launch**
Sabang-Limbok	UP	Sept 13, 2017
Dikapinisan	UP	Oct 25, 2017
Dibut	UP	Feb 1, 2018
Diotorin	UP	May 30, 2018
Bacong-Market	UP	Aug 29, 2018
Dianao	UP	Oct 17, 2018
Tanay	Globe	Jan 29, 2016
Talisay	Globe	Jan 21, 2017
Binobohan	Globe	Feb 2, 2018
Ginulagan	Globe	Apr 3, 2018
Balogo	Globe	Apr 3, 2018
Casalaan	Globe	Apr 5, 2018
Banat-i	Globe	Apr 5, 2018
Mayaposi	Globe	Jun 30, 2018
Golden Valley	Globe	Aug 10, 2018
San Mariano	Globe	Aug 11, 2018
Binucayan	Globe	Aug 11, 2018

Table 4.1: Site launch dates as of December 2018. Tanay and Talisay were proof of concept sites deployed by our MNO partner to demonstrate viability of expanding coverage through community networks and evaluate experimental hardware and software, while the bulk of deployments started after a year later. UP sites are community network sites facilitated by the University of the Philippines. Note that these dates are commercial launch dates; both lab and on-site testing typically took place for several weeks before launch.

include transportation costs by the local cooperative for credit distribution and collection, and monthly honoraria for maintenance staff.

Infrastructure re-use – for example, using an existing tower or building to mount equipment – can reduce costs significantly when possible. For the UP sites, the lack of existing towers or other high structures required constructing new towers from scratch. We identified a local metalworker to fabricate of towers, which we expect will reduce transportation costs compared to shipping tower components from Manila. Moreover, while grid power is provided in Aurora by a local electric cooperative, the grid infrastructure was deemed unreliable by the locals, who recommended that we use an off-grid solar power system instead. The local cooperative also favored this to avoid paying for the site's electrical consumption.

In our deployments, we used equipment that ran on several different voltages: 24VDC (common for low-power wireless equipment), -48VDC (common for telecom equipment), and

Item	Cost	Notes
GSM Radio	US$ 5,200	Combined GSM Radio + CPU for CCM Client, 2x10W
GSM Installation Accessories	US$ 480	Includes cables, mounting brackets, etc.
GSM Antennas	US$ 570	Two high-gain omnidirectional GSM900 antennas.
GSM Radio spare	US$ 125	Budget per-site for spare radios.
VSAT System	US$ 1,790	Includes installation (subcontracted).
Tower	US$ 1,550	10m pole.
Lightning Protection	US$ 200	
Telecom Works	US$ 1,100	Installation for networking and power.
Civil Works	US$ 3,100	Transport and construction of site infrastructure.
Site Survey and Testing	US$ 800	Pre- and post-installation evaluation.
Power System (Solar)	US$ 3,800	Two day backup power for off-grid sites.
Power System (Grid)	US$ 3,100	Inverter + batteries for locations with grid power.
Total CapEx (lower bound)	US$ 18,015	
Total CapEx (upper bound)	US$ 18,715	

Table 4.2: An example breakdown of a deployed site cost for a Globe site. Two different power systems are considered, one for fully off-grid sites and another for sites with grid power. The radio vendor is anonymous due to a non-disclosure agreement with Globe.

220VAC (grid voltage in the Philippines). This not only required additional equipment to perform the necessary conversions, but also led to decreased power efficiency for the entire site, driving up the cost of power, the second-largest component of site CapEx after the radio itself.

These costs also do not include fees for permits, as they were waived by the partner local government. As the project's main intention is research and not profitability, the local government units recognized the project's potential to help their constituents. We do not have direct knowledge of what these fees cost, though anecdotally we understand that they can both vary significantly by municipality and can constitute a significant portion of site costs; the permitting process is a point of leverage for local governments. Although fees were waived in our case, we still have to submit requisite documents such as construction plans and electrical plans.

Item	Cost	Notes
GSM Radio (Endaga CCN1)	US$ 4,700	Locally-assembled, 2x10W.
GSM Radio (NuRAN LiteCell 1.5)	US$ 6,000	2x10W.
x86 Computer	US$ 300	
GSM Antenna and Accessories	US$ 340	9dBi omni antenna, RF cables, connectors, grounding.
Networking Equipment	US$ 200	Off-the-shelf switch and cables.
Power System (Solar)	US$ 4,220	Three day standby power. Includes 800W panels, batteries, controller, inverter and other accessories.
VSAT System	US$ 2,000	VSAT modem and antenna; includes installation.
Civil Works	US$ 3,370	12m tower, equipment shelter, foundation, fencing.
Installation	US$ 2,000	Includes personnel, transport and community training.
Total CapEx (lower bound)	US$ 17,160	
Total CapEx (upper bound)	US$ 18,760	

Table 4.3: CapEx cost breakdown for a UP site in Aurora.

Usage

As of December 2018, CCM supports a total of about 2,800 monthly active users across 17 launched sites (Figure 4.10). Spikes in the number of active users are due to site launches, which are typically accompanied by marketing campaigns to raise awareness for the launch of service in a new town.

Table 4.5 provides an overview of traffic volumes across the network. Across all sites, we observed that inbound call traffic is much more common than outbound call traffic, and is in fact the predominant form of usage on both networks. This is indicative of “call me” or “flashing” [47] behavior as subscribers are aware that they can save money by letting their contacts call them instead; on all networks in our deployment, subscribers are only charged for user-initated calls and SMS. Narrowing our analysis to only user-initiated traffic, we find that SMS is the predominant form of communication, with roughly 7x more messages sent than minutes spent on calls, and over 13x the number of calls made. This is in line with the fact that SMS rates are significantly lower than per-minute rates and extremely popular in the Philippines (the “texting capital of the world” [76]).

We observed more local traffic among the UP sites than the Globe sites. Interviews conducted by the team suggest that one contributing factor is the fact that communities

Item	**Cost (monthly)**	**Notes**
VSAT service	US$ 100-400	512x512kbps, Price varies across vendors.
Transportation	US$ 40	For credit distribution, remittance or technical visits.
Local maintenance	US$ 40	Two maintenance personnel at $20 each. May be subsidized by LGU.
Total OpEx (lower bound)	US$ 180	
Total OpEx (upper bound)	US$ 480	

Table 4.4: OpEx cost breakdown for a UP site in Aurora.

Service Type	**Volume**		
	UP	**Globe**	**Total**
Calls Out (min)	55,459	114,294	169,754
Calls In (min)	1,128,849	1,886,642	3,015,492
SMS Out	440,767	713,017	1,153,784
SMS In	367,212	701,538	1,068,750

Table 4.5: Volume of usage by service type. Call volumes are reported as duration in minutes, while SMS volume is number of messages sent. Inbound communication is free.

served by the UP sites can be clustered into two groups where the sites are located relatively near each other, and because the community networks support existing locally relevant services. In the UP sites in San Luis, Aurora, locals frequently conduct trade activities and have personal connections with residents from the other sites, relationships which existed before the arrival of the community cellular network. In another example, the only high school in the area is located in Dikapinisan and students from the nearby sites need to relocate temporarily for their studies. Parents used the community cellular network to call their children and get updated on other current events. We note though that for these "inter-cove" (the sites are located in a series of coves) transactions, although they are classified as local traffic for billing purposes, these communications are routed through the Controller's SIP switch. Community members also told our team that they used the network for local events, such as a community beauty contest that used SMS for voting.

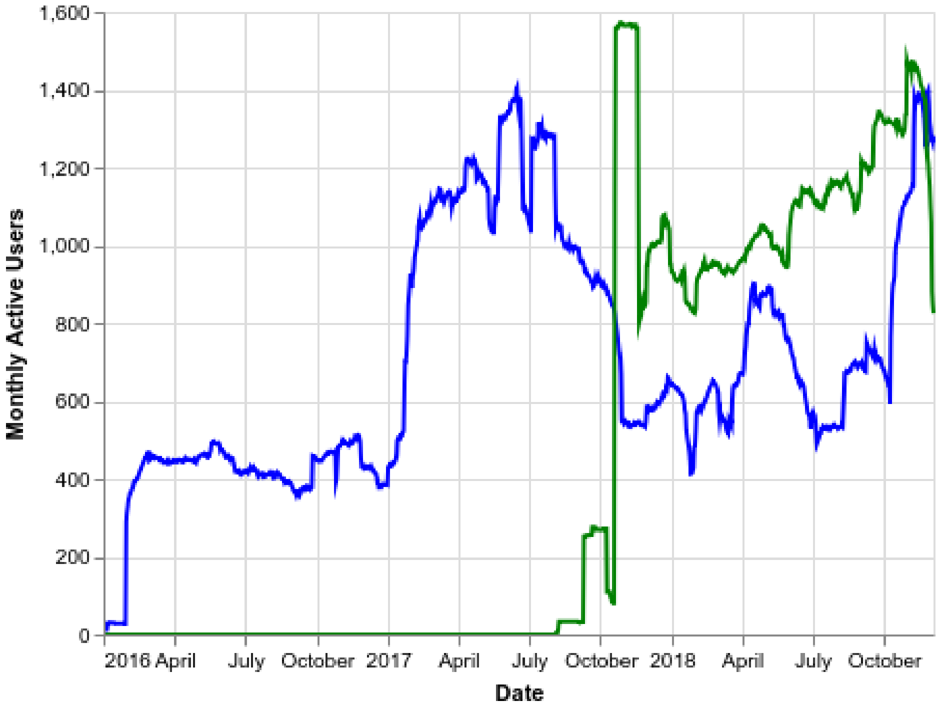

Figure 4.10: Monthly active users for Globe (blue) and UP (green) networks.

Disconnected Operation

Figure 4.11 depicts the launch date and uptime for each site. Downtime was a common case in our extremely rural sites. All sites use satellite backhaul, and even where grid power is available, its poor reliability necessitates battery backups or fully off-grid solar systems; some sites were also turned off on a nightly basis to conserve power. Overall, the mean site uptime is only 35% across all sites, with a median of 27% for Globe sites and 40% for UP sites.[4] This reality motivated our need to prioritize disconnection and downtime as common cases to be handled by CCM.

To understand whether users benefited from CCM's disconnected operation support, we examined CDRs from time periods while sites were offline. We only considered user-initiated "local" communications and credit transfers, since communication out of a site is not possible while it does not have a connection to the CCM Controller. This is still a conservative assumption, since communication between two sites in the same network (such as between two different UP sites) is also considered "local". We include credit transfer as well since it is a crucial utility used by all users of the network to buy service: if credit transfer breaks, sales are directly impacted, hurting network sustainability.

We find that overall, offline local traffic accounts for 16% of the total local traffic, com-

---

[4] Not all sites are unreliable. Three UP and two Globe sites achieve >60% uptime, even despite planned nightly downtime to conserve power.

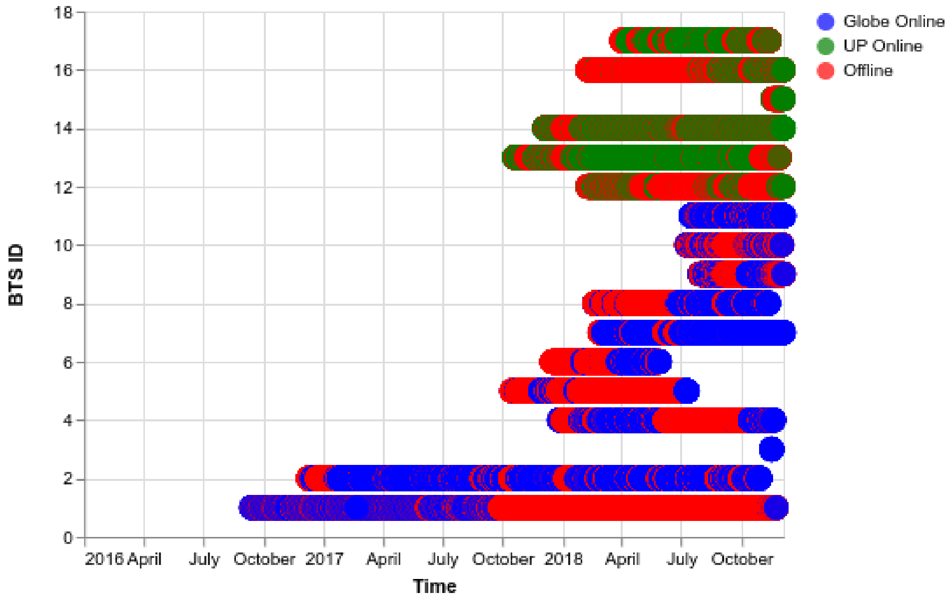

Figure 4.11: Site uptime. Sites 1-11 are Globe operated, while 12-17 are UP operated. Each point is one hour.

prising 7% of local traffic in the Globe sites and 23% of local traffic in the UP sites. This accounts for 4,678 minutes of calls, 81,434 SMS messages, and 11,970 credit transfers representing approximately US$4,600 worth of transfer activity. These credit transfers are largely top-up sales to end users and themselves represent 14% of the transfer volume across all sites, benefiting Globe, retailers, and the community networks at large. Without effective support for offline operation, these sites would have been completely down during backhaul outages, further complicating the already difficult business and sustainability cases for these rural network sites.

The primary cost CCM incurs for supporting offline operation is the overhead of state synchronization, carried by the checkin protocol. Figure 4.12 shows the distribution of checkin request and response sizes before and after our optimization. The median checkin request and response is 2.6kB and 283B, respectively, corresponding to median “unoptimized” request and response sizes of 14.4kB and 283B. Our optimization is much more effective at the tails: 99th percentile request and responses sizes are reduced by 86.3% and 91.2% respectively. The optimized checkins place minimal burden on our networks: the median checkin consumes less bandwidth than a single 6 second call (using a 64kbps codec), and even at the 99th percentile is equivalent to a 49 second call; this is not a significant increase in load on our network. We note that these sizes are prior to `gzip` compression over the wire, further reducing overhead.

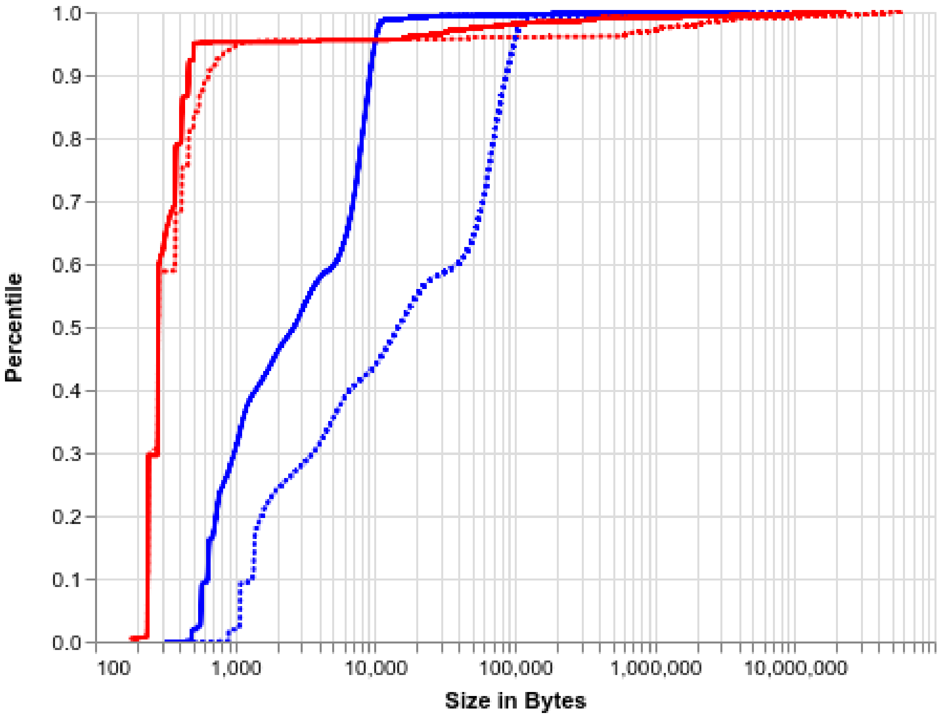

Figure 4.12: CDF of checkin sizes. Blue are requests from the Client to the Controller, and red are responses from the Controller to the Client. Dashed lines are raw data sizes and solid lines are size after optimization for transport.

Local Services

A consequence of CCM's decentralized architecture, and one that differentiates it from traditional centralized cellular networks, is that it allows customization of the individual networks to local needs, requirements, and desires. During our deployment, we and our partners took advantage of this capability to implement a number of unique services in the UP sites. Specifically, we implemented (1) a local repair support tool aiming to empower lay actors from within the community to conduct routine maintenance and repair, (2) a custom local billing solution to allow our team to explore the demand curves for rural communication access without requiring costly changes to Globe's billing systems, and (3) a local "outage hotline" connecting users directly to the UP team. These applications demonstrate CCM's ability to support flexible and dynamic community networks.

Repair. Repair is central to sustainable rural networking interventions [157, 96]. Further, in rural areas, we have the unique advantage that rural users are natural repairers [95]. While Globe installations use traditional maintenance practices, the UP community networks sought to address network repair and maintenance through a local-only repair support service that leveraged latent skills and abilities present in the community [99]. The researchers implemented a set of services leveraging the fact that community networks *interact directly*

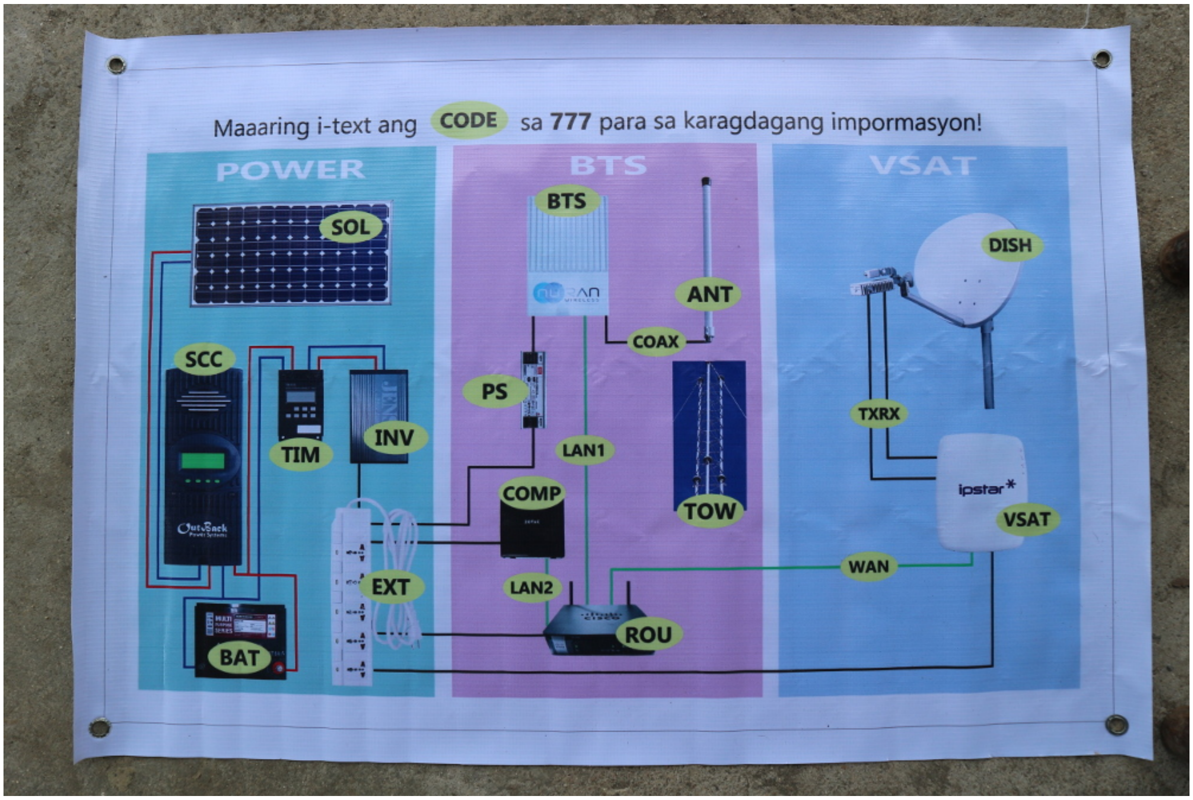

Figure 4.13: The infographic for the local repair service.

with users and can help guide them in conducting repair. The service consisted of digital "repair manuals" embedded into the community cellular infrastructure. Figure 4.13 shows an infographic of the system. Network components are labeled with small codes (e.g., "ANT" for antenna) and an SMS shortcode (e.g., 777) that, when texted, provides background and debugging information about that particular system element. Labels are printed on a large poster in a community building near the tower. The labels are provided as a large plastic poster placed in the community building nearest the cellular installation. For instance, a user texting "hot" to the shortcode will receive a message instructing them to turn off the system to allow it to cool. If the problem reoccurs, the message also instructs to text another keyword to the system (MAINITPARIN) which would trigger a different guide. This system was implemented on the Client's local softswitch (see Figure 4.2), and provides a unique, local mechanism for repair that is particularly appropriate for the UP sites, which don't have dedicated commercial technicians.

Billing. The UP sites also have unique requirements for their billing system. Though CCM handles tariffs between Globe and the community network, the community sets their prices within their domain for their users. For the UP systems, the billing system was extended to enable *promos*, a well-known local pricing scheme among Filipino MNOs. Unlike the per-minute and per-SMS fees charged by default, users pay up-front for a set of network functions, usually a discrete volume of SMS and voice minutes. These promo bundles give more value for money to the subscribers by providing them more texts, call minutes or a combination of both for a fixed price, but expire after a specified duration. These services were implemented in CCM as Client-only databases storing current promo offerings and promo counts for each subscriber. The implementation also allows the administrator to grant promos to subscribers, similar to a rewards system.

Promos exist at a per-site level and are not synchronized across sites; quotas or discounted tariffs are stored locally. These are defined when the network administrator creates the promos via a web form or a CSV file upload. The system supports multiple promos per user. It also offers a SMS interface for users to check their promo status and usage.

Outage Hotline. The UP sites also offer an SMS-based outage hotline. This free service accepts questions, comments, reports and other service-related inquiries from the communities. The hotline logs received messages to a file on the CCM Client's storage, which are synchronized to a remote server (distinct from the Controller) at regular intervals. The hotline also helps the researchers in the detection of technical issues in the field as community members can easily report any problems that they may experience on the ground. This is also the mechanism that local maintenance personnel utilize to send network status updates.

Reliability

Network downtime is common in our deployment. According to one user at the start of our deployment, "I'm OK even if [the network] is turned off at night. At least now, we have something that we can use to communicate. Unlike before, we totally have none." This attitude changed during the course of the deployment; despite the utility of disconnected operation, users still came to expect continuous, reliable operation from these networks: "Why don't we have signal during the night? They shouldn't turn it off during the night, because it is still important in case something bad happens here." Downtime is part of normal operation in these networks. In our deployments, outages occurred at the boundaries of the network's physical infrastructure: power systems, backhaul networks, and site hardware. We were able to mitigate these to an extent through planning and careful site engineering, but we regard reliability in community cellular networks as an open challenge.

All sites use satellite backhaul, and unreliable or non-existent grid power requires use of battery backups or off-grid solar systems. We do not have ground truth for all site outages, but outage records from the UP sites provide an approximate distribution of outage causes (Table 4.6). At these sites, the most common cause of outage was backhaul failure, followed by power outages, collectively accounting for about 66% of the downtime, with hardware or RF issues accounting for the remainder. Backhaul failures had minor impact since they tended to be transient and since CCM gracefully handles offline operation, but power failures were more serious, often resulting in extended downtime. Improving resilience to power failures is hard, as increasing battery capacity adds to site cost; approaches like "virtual coverage" [86] address this challenge to an extent, but require non-standard hardware. To

Cause	**Percentage**	**Example**
Backhaul	42%	VSAT offline due to weather
Power	24%	Discharged batteries
Site Hardware	21%	Overheating
Site RF	12%	Broken antenna cable

Table 4.6: UP site outages by duration. This data comes from manual record keeping by site caretakers and is approximate; we do not have ground truth for all site outages.

reduce long-term outages, sites are shut down at night to reduce power utilization during off-peak hours, either manually or via an automatic switch.

All significant offline periods at Globe sites (Figure 4.11) were due to hardware failure; this impacted 9 of 11 launched sites. When Globe took on operation of the CCM Controller, they selected a new hardware vendor to provide equipment for Globe sites. This vendor included the CCM Client into their existing mature rural base station product, which promised to allow Globe sites to use proven hardware already being produced at scale, reducing cost. This hardware used an SD card as a disk, which was used in read-only mode in the vendor's standard product offering; this is a best practice for rural networking hardware [157]. To support CCM, which requires mutable persistent storage to support offline operation and local applications, the vendor continued to use the SD card, but in a writeable mode. This led to a pernicious failure mode where devices would fail after several months' successful operation in the field due to disk corruption driven by power failures and excessive writes. Community networks benefit from applications like CCM and those outlined in Section 4.5 that require mutable persistent storage, but supporting these applications in rural environments remains a challenge. With the commercial availability of various forms of persistent memory, utilizing NVRAM may be one approach to support such workloads in environments prone to unexpected power failures [18, 111].

4.6 Discussion

Towards LTE. The rapidly maturing ecosystem of open source telecom software and newly-affordable RAN equipment designed for rural cellular networks makes community cellular networks technically and financially feasible. As noted, CCM only supports GSM (2G) service because basic GSM-only phones are still the plurality in the rural areas we target [150]. However, CCM's architecture represents a framework for MNOs and community networks to partner and share resources. We expect the shift towards LTE (4G) to make community LTE networks viable in the coming years and we are actively pursuing this agenda through efforts such as CoLTE [149] and Magma [54].

The shift to LTE creates new opportunities for community networks. While GSM networks are essentially deployed in four spectrum bands globally, LTE devices support dozens of band combinations, some of which support spectrum sharing (such as Band 48 [57]) and others that fall within "digital dividend" bands (e.g., Band 71); the latter is ideal for rural networks due to improved propagation of low-band spectrum. Even without shared spectrum, demand for LTE spectrum for capacity in urban areas is likely to result in substantial underutilization in rural geographies, opening up opportunities for partnership and spectrum re-use as we've done in our deployment. Driven by use cases like private LTE networks for enterprises and industrial "Internet of Things" deployments, as well as applications like mobile edge computing, the LTE base station ecosystem is already larger than that for GSM, lowering costs for community networks. Radio equipment of varying degrees of openness is under development [158, 7], and multiple open source software core networks, analogous to OpenBTS and Osmocom, are available [122, 67, 124].

Given this progress, community LTE networks could provide broadband service independently of MNOs, connecting to the Internet just as a small ISP would. Recognizing this, we recently proposed a *decentralized* LTE architecture [102] that does not require a telecom's (or any centralized organization's) participation. We expect a fully distributed architecture should ease deployment and empower community networks at the cost of increased difficulty scaling. Exploration of the political and technical tradeoffs across the centralized/decentralized and telecom partnership/independent network design spaces remain open research questions.

Fault diagnosis. Community cellular networks rely on a collection of systems to provide service to end users; just to send an SMS, a user's traffic interacts with a radio implementation (proprietary hardware, or a software radio), a GSM stack (OpenBTS or Osmocom), a SIP engine (Freeswitch [63]), an SMPP gateway (Kannel [106]), and of course CCM. Even for experts, diagnosing faults in the mobile stack is challenging due to the the need to manage state across layers and components. We relied on regular automated end-to-end health checks of the CCM Client to identify and rectify faults; automated failure diagnosis will be essential for community networks.

For example, we faced variety of bugs in both OpenBTS and (to a lesser extent) Osmocom that would often manifest as diminished performance, rather than outright outages. For issues with the latter, CCM performs periodic system-level tests of the underlying GSM stack, such as overall channel load, to detect faults and simply restart failed components. We also developed a end-to-end VOIP tester, `fake_phone`, which replicated the behavior of a user's device for debugging. However, we had no similar visibility into the performance of Globe's side of the interconnect. A drop in call volume at the interconnect point could signal a variety of problems, but since Globe was a black box beyond the interconnection, we had no way of diagnosing faults further. There is no `ping` or `traceroute` equivalent for telecom networks, making fault diagnosis challenging at best.

Sustainability. As of today, only a few of the UP sites (and none of the Globe sites) are financially sustainable, taking into account the ongoing operating costs for the sites. This was expected for the unsustainable UP sites, as they were too small to provide enough revenue. We are hopeful that future research endeavors into novel business models and cost structures will resolve this issue. The Globe sites were designed to all be sustainable but poor system reliability hurt usage. We remain hopeful that stability improvement will increase revenue to sustainable levels.

Repeatability. One of our broader objectives is to develop a scalable and repeatable model for community cellular networks. We were fortunate that our project did not require any significant regulatory changes and that spectrum sublicensing was permissible under existing Filipino regulations. This is not always the case, and where MNOs are unable to allow third parties to use their spectrum, they may not be able to engage in this particular business model. Nevertheless, CCM reduces the challenge of starting many community networks to a *commercial* negotiation rather than a *regulatory* discussion (or even *legislative* action), often a much lower bar.

Security A full treatment of security challenges in mobile networks is outside the scope of this paper, CCM presents a few specifically challenging use cases. Because each CCM Client holds a copy of users' SIM card keys on disk and in memory to authenticate subscribers, an attacker with physical access to a CCM Client device has the potential to gain access to critical subscriber information typically safeguarded in an MNO's datacenter, allowing SIM cloning attacks and snooping on user traffic. Unfortunately, due to flaws in the GSM specification, similar attacks are generally trivial even with traditional GSM networks. We assume that the community provides physical security protection to the CCM Client devices to mitigate this risk.

If physical security can't be assumed, encrypting keys at rest on disk would be a useful first step, but this would still leave keys exposed to an attacker with physical access since they could be read in memory. We have a few options for solving this challenge. First, we could simply require key operations be conducted by the CCM Controller on behalf of the Client, with the obvious downside being that authentication operations could not occur when the site was offline. Second, we could rely on hardware enclave support (such as ARM TrustZone [10]) to perform key operations in a secure portion of memory. Doing this would require remote attestation of the enclave, which could be done by the controller. Unlike our first option, this approach allows user authentication operations to proceed when the Client is disconnected from the Controller once attestation is completed. We leave these extensions to future work.

4.7 Conclusion

For the millions who live outside of basic mobile coverage, community cellular networks present a way to sustainable mobile coverage. Finding ways to remove the regulatory and commercial barriers to their growth is key to realizing this potential. Working with MNOs provides a straightforward, if not necessarily easy, path to doing this that requires neither major regulatory changes in many jurisdictions nor any of the actors involved to act contrary to their first order interests. The challenge that remains is building the platform to connect these different entities together.

Our work shows that these challenges are surmountable, and provides an example of ways to deploy community networks at scale. We identified critical design goals for such a system – autonomous services, minimal transaction costs and risk, and minimal absolute deployment costs – and implemented CCM to realize these goals. CCM is designed to scale down well for small-scale community networks and tolerates intermittent disconnected operation while continuing to provide service to end users and enforce network policy. Through a large-scale deployment, we demonstrate CCM's ability to effectively support these community network – MNO partnerships, connecting 17 communities and thousands of users. Our system is open source,[5] and we hope others will find it useful for replicating this model.

---

[5] `https://github.com/co-cell/ccm`

Chapter 5

Conclusion

In this thesis, we have examined the challenges network operators face in providing connectivity to unserved rural areas and developed mechanisms for supporting community-based cellular networks. We have found that modern, mature small-scale rural wireless networks, exemplified by rural WISPs in the United States, face bottlenecks that are largely non-technical in nature. Access to appropriate spectrum and financing are key needs, and physical faults such as equipment failures and lightning strikes are the leading causes of outages and maintenance on these networks. Purpose-built hardware, extensive use of proactive monitoring, and integration between network operations and business support services have enabled WISPs to build networks of significant scale across geographically large service areas with limited personnel.

Informed by the experience of WISPs in the United States, we then sought to identify and remove scaling barriers to growth for the nascent community cellular networks ecosystem. We introduced the concept of **GSM whitespaces** and devised a system for sharing spectrum amenable to the context of rural community cellular networks. We demonstrated that GSM whitespaces enables sharing of unused cellular spectrum between incumbent mobile network operators and community cellular networks without disruption to incumbent services, with minimal overhead on community cellular networks, and without any changes to existing cellular devices. We implement Nomadic GSM, a practical system to take advantage of GSM whitespaces, which leverages the network-driven handover mechanism used in all modern cellular standards; Nomadic GSM enables community cellular networks to dynamically change frequency without impacting user traffic while detecting the presence of an incumbent operator within 90 seconds. We validate our work through a deployment of NGSM in a production community cellular network in Indonesia.

Next, with **CCM**, we built an end-to-end system for operating collections of community cellular networks at scale and in partnership with existing mobile network operators. Our system draws upon experience from operating rural community cellular networks to solve a number of practical challenges facing these types of networks. Today, CCM supports the largest community cellular deployment to date in the Philippines, serving over

2,500 monthly active people and handling more than three million text messages and 50,000 hours worth of calls during its three years of operation. To achieve this, CCM facilitates sharing spectrum, network identity, and interconnection resources between mobile network operators and community cellular networks while integrating business support and network management and automation tools for community network operators. We introduce a novel partially-decentralized federated mobile core network architecture that supports intermittently disconnected operation as well as per-site (in addition to standard per-user) network policies. We show that this architecture supports user traffic in the face of the most common cause of network failure in our deployment, intermittent backhaul outages, and does so with negligible overhead.

5.1 Open Challenges

The work we have presented here highlights a number of open challenges; we consider three of these (as well as possible avenues of attack) here.

Rural challenges outside the United States. In Chapter 2, our survey of rural ISPs focused on those in the United States. While many of our results align with prior work outside the United States [157], how well our results carry over into other markets is an open question. Our key findings speak to the impact of the regulatory, financial, and vendor environment of the United States market; as these all vary significantly from country to country, we would expect to see variation in these results. Further, while our study participants come from across the United States, we would expect that revenue per user for US-based WISPs could be higher than in markets with lower per-capita incomes.

Ideally, we would replicate this survey with WISPs outside the United States. A key challenge in replicating this study outside the United States is the lack of organized industry associations for WISPs, which makes recruiting WISP operators difficult; we were able to broadcast a call for participation in our survey thanks largely to the Wireless Internet Service Provider Association (WISPA) [173] mailing lists. A number of ISP associations exist globally, but few focus specifically on the interests of WISPs. The Wireless Access Providers Association (WAPA) in South Africa [172] appears to be similar in terms of membership to WISPA, and ABRINT (Associação Brasileira de Provedores de Internet e Telecomunicações) in Brazil [11], while open to all types of ISPs, counts many WISPs among its membership. Working with vendors or distributors of equipment used by WISPs may be one avenue to identify WISPs in markets with a less organized WISP ecosystem.

In addition, mutual assistance between WISPs in online forums is common – this is both a primary use case for the WISPA and AFMUG mailing lists we used in our work as well as numerous Facebook groups catering to WISPs.[1] These groups could also provide an avenue

---

[1] One of the largest today is the “Wisp Talk” Facebook group, with more than 6,000 members at this writing.

for replicating our study and more broadly gaining a rich view into the practical challenges faced by these operators.

Cellular spectrum sharing in dense environments. In Chapter 3, our Nomadic GSM (NGSM) system considers the case of a single isolated community network cell. However, overlapping community cellular networks or dense cellular networks acting as a secondary users (for example, a urban community cellular network) are common, and the NGSM approach would need to be adapted in order to be effectively used in those contexts.

Specifically, NGSM uses the handover mechanism to facilitate client-based power sensing of adjacent channels. Doing so requires we add additional entries to the neighbor list of the serving cell in order to receive power level measurement reports from clients on these channels, creating a tradeoff between the throughput of channels we can sense occupancy on and the number of actual adjacent cells for handover purposes. More importantly, N overlapping NGSM cells would be continually alternating between $2N$ channels, incurring control plane overhead to manage handover between cells; this effect would pronounced for client devices at the boundary between cells. Coordination among NGSM sites within a single administrative domain could mitigate this by synchronizing channel switching periods (note that independent secondaries in this scheme would *not* need to coordinate, since they could treat each other as interference to avoid).

Another approach would be to cycle the unused "sensing" channel among a set of cells. Consider a community cellular network using a cell reuse factor of $\frac{1}{7}$, requiring seven minimum channels for communication. Assuming sites coordinate their spectrum usage, each cell could determine whether a primary is using its current channel by switching to an eighth channel one at a time. This minimizes the amount of excess spectrum necessary for detecting interference on a cell's usual channel at the cost of increasing the time to detect interference.

We note that while this is a wide design space, but the underlying premise of our GSM whitespace framework holds throughout: cellular spectrum can be safely shared without modifications to existing incumbent mobile networks or user devices.

Achieving the next level of scale. Our deployment in Chapter 4 represents the largest single deployment of community cellular networks to date; the collection of network deployments coordinated by Rhizomatica is of similar size and scale. Both our deployments and those of Rhizomatica required significant effort to "unlock the market" for community cellular. Rhizomatica spent years working with the Mexican telecommunications regulator IFT to obtain first an experimental spectrum concession and later a permanent spectrum concession for community networks; they were fortunate that a small portion of the GSM band in Mexico had never been assigned and was thus available for their use. Similarly, work to start the deployments CCM supported in the Philippines began in 2013 and relied upon a fortuitous series of personal connections between our research team at UC Berkeley and the University of the Philippines and senior executives at Globe. The fortuitous and long-term nature of these efforts limit their replicability.

Thus, while our goal is to identify ways to scale community cellular networks, we must admit that today there is no obvious path to scale for community cellular networks worldwide: both our efforts and those of Rhizomatica required years of dedicated effort to create the enabling environment for community cellular networks before any networks could be built. Indeed, today, the number of people served by community networks globally is almost certainly under 100,000, largely concentrated in Latin America and Southeast Asia. While the systems and approaches we have discussed – spectrum sharing enabled by NGSM and cooperation with incumbent mobile operators enabled by CCM – provide ammunition for those who seek to replicate our work, the path ahead for them is long and fraught with profoundly non-technical challenges.

Organizing and funding support is essential for the ongoing development of the community networks ecosystem. Groups like the Association for Progressive Communications [12] and the Internet Society [92] have ongoing efforts to organize community networks, facilitate information sharing among groups building community networks, performing advocacy on behalf of community networks, and generally raising the profile of community network efforts worldwide. We hope that systems like those proposed in this thesis support the work that these organizations are doing and help to further stimulate the growth of the community networks ecosystem.

5.2 Lessons Learned and Closing Thoughts

Rural networks face physical failures. Across our interviews with WISPs and in our experiences with our deployments in the Philippines, we consistently found that network outages were primarily due to physical failures: lightning strikes, power outages, damaged cables, typhoons. This makes sense given rural networks tend to operate in challenging environments. The lesson for designers building systems for rural networks, however, is that physical component failure should be assumed to be a common case. This goes not only for network elements, but for the entire stack of supporting infrastructure necessary for operating a rural site, such as power and backhaul infrastructure. Recalling the end-to-end principle [145] is instructive: fate sharing is beneficial in such environments. We saw this in CCM, where by moving the functions of a cellular core network to be co-located with the radio, we were able to support some level of network services as long as the equipment in the community was powered on.

Spectrum can be abundant for rural cellular networks. Past empirical results have shown that spectrum occupancy is generally quite low in both urban and rural environments, and we can deduce that mobile spectrum in uncovered rural areas is not utilized at all due to exclusive-use spectrum regulations being the global default. Our work shows that enabling new actors, like community networks, to use this underutilized resource can be done safely, both through empirical results from our deployment of NGSM in Chapter 3 as well as our large-scale CCM deployment in the Philippines. The key lesson we can take from this work is that, unlike in urban areas, spectrum in rural environments need not always be considered a scarce resource, and that by doing so we prevent rural people from having access to communications services. Moving from the assumption of scarcity to an assumption of abundance for rural spectrum presents new regulatory and technical opportunities for expanding rural access.

Expanding rural connectivity requires more than technical innovation. Our work shows that the challenge of bringing connectivity are a mix of technical and non-technical challenges. From our interviews with WISPs, we saw that financing was a significant barrier to scale, even in the face of reliable, low-cost hardware. In our exploration of GSM Whitespaces, the infrastructure we sought to deploy "simply" needed an enabling policy framework to support its operation. Technical innovation can unlock certain use cases – consider CCM's ability to handle intermittent backhaul failures gracefully improving utility of rural GSM networks. Yet the critical challenges rural network infrastructure operators face, be they community networks or otherwise, are often not due to the technical limitations of the systems they deploy. Proposals to provide connectivity from balloons or planes do not inherently resolve the power, weather, or financial constraints faced by rural network operators and the other actors involved in delivering service to end users, such as retail agents selling top-ups or service technicians helping diagnose faults caused by user error. Designing interventions to expand access to rural connectivity requires a holistic view of both the technical systems operators deploy as well as the regulatory, commercial, operational, and maintenance aspects of their networks: there is no silver bullet for rural connectivity.

Appendix A

This appendix presents raw survey responses for the WISP survey conducted in Chapter 2. As discussed in that chapter, this survey include 75 responses obtained from announcements posted on the Animal Farm Microwave Users Group (AFMUG) and Wireless Internet Service Providers Association (WISPA) mailing lists in June 2012.

For the analysis in Chapter 2, respondents were further categorized by their response to Question A.2; any respondents that selected “Rural, inside the United States” was classified as a rural WISP, along with the three respondents who provided the free form responses “Rural per FCC”, “Semi-rural”, and “Small towns”. This includes respondents who selected the “Urban, inside the United States” category in addition to these. Here, we present the overall (that is, uncategorized) results pacross all respondents. We do not remove outliers, such as one respondent who claimed to represent a large mobile operator outside the United States, though these are not considered in our analysis.

Demographics

Table A.1: Q1. What are your roles in your organization?

Network management	60	80.0%
Business Management	56	74.7%
Customer support	40	53.3%
Service or Installation	37	49.3%
Sales or Marketing	29	38.7%
Not involved with a WISP	0	0.0%
Other (free form)		
Owner	3	3.9%
Founder	1	1.3%
Provides funding to WISPs	1	1.3%
RF Engineering	1	1.3%
No response	1	1.3%

75 responses. Total response exceeds 100% as multiple options were possible.

Table A.2: Q2. Which of the following describe where your organization provides service?

Rural, inside the United States	64	86.5%
Urban, inside the United States	23	31.1%
Other (free form)		
Small towns	1	1.4%
India	1	1.4%
Semi-rural	1	1.4%
US Virgin Islands	1	1.4%
Rural per FCC	1	1.4%
Nigeria	1	1.4%

74 responses. Total response exceeds 100% as multiple options were possible.

Table A.3: Q3. How many people work for your organization?

1–5	31	41.3%
6–25	37	49.3%
26–50	3	4.0%
51–100	2	2.7%
100+	1	1.3%
Prefer not to answer	1	1.3%

75 responses.

Table A.4: Q4. How many subscribers does your network have?

1–99	4	5.3%
100–499	17	22.7%
500–999	14	18.7%
1000–4999	32	42.7%
5000+	6	8.0%
Prefer not to answer	2	2.7%

75 responses.

Table A.5: Q5. How many network devices are used in the backhaul portion of your wireless network?

1–25	10	13.9%
26–50	13	18.1%
51–75	9	12.5%
76–100	6	8.3%
101–125	6	8.3%
126–150	4	5.6%
151–175	0	0.0%
176–200	3	4.2%
201–225	3	4.2%
226–250	2	2.8%
251–300	5	6.9%
301–400	4	5.6%
401–500	4	5.6%
501–1000	1	1.4%
1000+	2	2.8%

72 responses. Free response.

Table A.6: Q6. How many network devices are used in the client-facing portion of your network?

1–100	15	21.4%
101–200	11	15.7%
201–300	6	8.6%
301–400	4	5.7%
401–500	1	1.4%
501–600	2	2.9%
601–700	2	2.9%
701–800	2	2.9%
801–900	2	2.9%
901–1000	0	0.0%
1001–1100	1	1.4%
1101–1200	1	1.4%
1201–1300	2	2.9%
1301–1400	0	0.0%
1401–1500	2	2.9%
1501–1600	0	0.0%
1601–1700	0	0.0%
1701–1800	0	0.0%
1801–1900	1	1.4%
1901–2000	1	1.4%
2001–2500	3	4.3%
2501–3000	6	8.6%
3001–3500	2	2.9%
3501–5000	3	4.3%
5001–10000	2	2.9%
10000+	1	1.4%

70 responses. Free response.

Table A.7: Q7. If your organization offers other services besides fixed wireless broadband, please list them here.

Email hosting	49	79.0%
Web Hosting	44	71.0%
Colocation	34	54.8%
Dedicated Technical Support*	24	38.7%
DSL Internet service	19	30.6%
Other (free form)		
Voice	13	21.0%
IT Services	10	16.1%
Dedicated Connectivity	6	9.7%
Marketing and Design Services	4	6.5%
Fiber (residential)	3	4.8%
Cellular Service	3	4.8%
Television	2	3.2%
Repair Services	2	3.2%

62 responses. Dedicated technical support is independent of support provided to subscribers. Free form responses are coded into the listed categories. Total exceeds 100% as multiple options are possible per respondent.

Network Management

Table A.8: Q8. How many upstream providers do you have?

1	20	28.2%
2	31	43.7%
3	13	18.3%
4	2	2.8%
5	2	2.8%
6	0	0.0%
7	1	1.4%
8	0	0.0%
9	2	2.8%

71 responses. “Upstream” refers to transit providers to the ISP.

Table A.9: Q9. What is the total bandwidth commitment to your upstream providers?

1–100 Mbps	24	32.0%
101–500 Mbps	28	37.3%
501 Mbps–1 Gbps	10	13.3%
1–10 Gbps	10	13.3%
10 Gbps+	1	1.3%
Prefer not to answer	2	2.6%

75 responses.

Table A.10: Q10. What is the total bandwidth commitment from your subscribers to you?

1–100 Mbps	26	34.7%
101–500 Mbps	13	17.3%
501 Mbps–1 Gbps	13	17.3%
1–10 Gbps	15	20.0%
10 Gbps+	3	4.0%
Prefer not to answer	5	6.7%

75 responses.

Table A.11: Q11. What is the typical peak usage you see on your network?

1–100 Mbps	34	45.3%
101–500 Mbps	29	38.7%
501 Mbps–1 Gbps	8	10.7%
1–10 Gbps	2	2.7%
10 Gbps+	0	0.0%
Prefer not to answer	2	2.7%

75 responses.

Table A.12: Q12. Which of the following do you provide to your customers?

Static public IPv4 addresses	63	84.0%
Static private IPv4 addresses	30	40.0%
Dynamic public IPv4 addresses	27	36.0%
Dynamic private IPv4 addresses	25	33.3%
NATed private IPv4 addresses	33	44.0%
IPv6 addresses (any type)	8	10.7%
Other (free form)		
DIDs[1] for VoIP	1	1.3%
IPv6 planned	2	2.7%

75 responses.

[1] "Direct inbound dialing" refers to phone numbers used for voice over IP.

Table A.13: Q13. How large is your total public IP address allocation?

1–32 (< /27)	3	4.1%
33–256 (/27–/24)	11	14.9%
257–1024 (/24–/22)	10	13.5%
1025–4096 (/22–/20)	24	32.4%
4097–16384 (/20–/18)	20	27.0%
16385–131072 (/18–/15)	4	5.4%
131073 (> /15)	0	0.0%
Prefer not to answer	2	2.7%

74 responses.

Table A.14: Q14. How much does your organization spend per month on:

	$0	$0–499	$500–999	$1000–9999	$10,000+
Upstream bandwidth?	0 (0.0%)	2 (2.9%)	5 (7.1%)	51 (72.9%)	12 (17.1%)
New equipment purchases?	0 (0.0%)	7 (10.4%)	13 (19.4%)	31 (46.3%)	16 (23.9%)
Maintenance of existing equipment?	2 (3.0%)	27 (40.3%)	14 (20.9%)	20 (29.9%)	4 (6.0%)
Software/appliance licenses?	17 (25.8%)	31 (47.0%)	8 (12.1%)	9 (13.6%)	1 (15.1%)
Network operations staff?	6 (8.8%)	4 (5.9%)	6 (8.8%)	38 (55.9%)	14 (20.6%)
Customer support?	7 (10.3%)	10 (14.7%)	10 (14.7%)	29 (42.6%)	10 (14.7%)
Non-network operations?	14 (20.9%)	27 (40.3%)	16 (23.9%)	8 (10.9%)	2 (3.0%)
Other network expenses?	4 (6.3%)	23 (36.5%)	13 (20.6%)	19 (30.2%)	4 (6.3%)

Spending on network operations. Total responses varied per category.

Table A.15: Q15. Which best describes how often your network experience faults due to each of the following:

Never	Yearly	Monthly	Weekly	Daily
Software updates?				
17 (23.9%)	39 (54.9%)	12 (16.9%)	3 (4.2%)	0 (0.0%)
Misconfiguration?				
14 (20.0%)	46 (65.7%)	9 (12.9%)	0 (0.0%)	1 (1.4%)
Hardware failure?				
0 (0.0%)	3 (4.3%)	34 (49.3%)	32 (46.4%)	0 (0.0%)
Power failure?				
10 (13.9%)	35 (48.6%)	20 (27.8%)	7 (9.7%)	0 (0.0%)
Antenna misalignment?				
18 (25.4%)	34 (47.9%)	13 (18.3%)	4 (5.6%)	2 (2.8%)
Radio interference?				
3 (4.2%)	28 (39.4%)	27 (38.0%)	9 (12.7%)	4 (5.6%)
Upstream failure?				
14 (19.7%)	51 (71.8%)	4 (5.6%)	2 (2.8%)	0 (0.0%)
Weather?				
14 (19.7%)	36 (50.7%)	20 (28.2%)	0 (0.0%)	1 (1.4%)
Network operator error?				
19 (26.4%)	40 (55.6%)	11 (15.3%)	1 (1.4%)	1 (1.4%)
User error?				
3 (4.3%)	9 (12.9%)	21 (30.0%)	19 (27.1%)	18 (25.7%)

Outage frequency by cause. Total responses varied per category.

Table A.16: Q16. How much downtime per fault event does your network experience as a result of:

	<10 min	10 min–1 hr	1–6 hr	6 hr–1 day	>1 day
Software updates?	40 (58.8%)	19 (27.9%)	8 (11.8%)	0 (0.0%)	1 (1.5%)
Misconfiguration?	34 (50.7%)	30 (44.8%)	1 (1.5%)	2 (3.0%)	0 (0.0%)
Hardware failure?	8 (11.9%)	12 (17.9%)	38 (56.7%)	7 (10.4%)	2 (3.0%)
Power failure?	13 (19.4%)	15 (22.4%)	30 (44.8%)	8 (11.9%)	1 (1.5%)
Antenna misalignment?	17 (26.6%)	11 (17.2%)	24 (37.5%)	9 (14.1%)	3 (4.7%)
Radio interference?	9 (14.1%)	14 (21.9%)	26 (40.6%)	7 (10.9%)	8 (12.5%)
Upstream failure?	21 (31.8%)	21 (31.8%)	14 (21.2%)	8 (12.1%)	2 (3.0%)
Weather?	19 (28.8%)	12 (18.2%)	21 (31.8%)	11 (16.7%)	3 (4.5%)
Network operator error?	30 (44.1%)	25 (36.8%)	10 (14.7%)	2 (2.9%)	1 (1.5%)
User error?	8 (12.3%)	17 (26.2%)	22 (33.8%)	9 (13.8%)	9 (13.8%)

Outage frequency by cause. Total responses varied per category.

Table A.17: Q17. Does your organization perform any of the following traffic engineering practices?

	Yes	No	Prefer not to answer
Per-customer bandwidth caps	36 (49.3%)	37 (50.7%)	—
Traffic filtering by content	8 (11.3%)	62 (87.3%)	1 (1.4%)
Traffic filtering by application	25 (34.2%)	47 (64.4%)	1 (1.4%)
Caching proxies	13 (18.6%)	57 (81.4%)	—
QoS for VoIP	55 (76.4%)	16 (22.2%)	1 (1.4%)
Network segmentation	56 (76.7%)	16 (21.9%)	1 (1.4%)
Client isolation	57 (78.1%)	15 (20.5%)	1 (1.4%)

Traffic engineering policies. Total responses varied per category.

Free form responses

The following questions solicited free form responses. Responses have been edited to remove identifying information.

Table A.18: Q18. Describe any other traffic engineering your organization performs.

- “Absolute priority for committed data rate customers over shared access customers.”
- “QoS [quality of service]”
- “We provide network engineering and management for a neighboring WISP.”
- “Dynamic user bandwidth allocation based on recent user activity.”
- “MPLS/TE [multiprotocol label switching/traffic engineering] for backbone links, PCQ [per-connection queues] for congestable links.”
- “We use a hub system for delivering bandwidth to the outer parts of our network.”
- “OpenDNS is default with an opt out capability”
- “CPE [customer premises equipment] – AP [access point] restrictions. Bandwidth is limited before it reaches the AP”
- “QoS per SM [subscriber module, CPE] software and overall PRTG [monitoring software] monitoring”
- “We prioritize all traffic using 802.1p. We provide layer two transport over MPLS, so we use MPLS-TE on our backbone.”
- “QoS between tower locations at our main internet connection to prevent fair bandwidth to each tower based on load.”
- “Extensive use of routed segments and VLAN [virtual local area network] switching. No over subscription of bandwidth. We only service B to B customers. No Residential. Static IP’s assigned to all customer interfaces for monitoring. Most large customers are touched twice. A note about customer counts, we do very large companies: [list of companies].”
- “I was unsure how to answer your first question above about caps. When WISPs think of caps we usually think a limits of amount of bandwidth total available for the month which is how cellular companies generally setup caps. WISPs generally use maximum allowable speed rules per customer and we do not limit how much bandwidth they consume per month. That way someone gets fairly consistent speed at all times and we can adequately serve the needs of all served from a sector on a tower. Also users are allowed to watch Netflix, Youtube, Hulu, etc. without fear of being ‘capped out’ or getting a huge bill at the end of the month.”
- “P2P Shaping”
- “Network maintenance and LAN management. APR anti-virus file sharing.”
- “Site Survey’s, USGS Map Studies, Path Profiles, Microwave Path Calculations, Testing w/ the radios we will use at a given installation - distribution node or Client site. System Upgrade Planning using all the above.”

Table A.18 – *Continued from previous page*

- “Fair access policy. No one user is allowed to use all capacity on any given link. The link knows it’s max speed and all users are given access to that.”
- “Each user gets more if the network is not busy. If it is busy then everyone shares what’s available.”
- “Dynamic speed control based on daily volume.”
- “Susbscription Rate”

Table A.19: Q19. What is the biggest challenge your organization faces that you want a solution to?

- “More spectrum”
- “We need more unlicensed spectrum.”
- “Spectrum that is dedicated to the WISP industry that is not up for bids by big telcos or others. A solution for non-LOS for customers in rural areas.”
- “We utilize mostly unlicensed spectrum in 2.4, 3.65, and 5 GHz.(some 900) and it is getting VERY crowded. We would utilize licensed frequencies, but the hardware cost go up dramatically, especially for CPE equip. We would like to see either more unlicensed spectrum so it isn’t so crowded, or a better way for those using them to share this spectrum without killing the other providers signals. or more and easier to obtain licensed spectrum with cheaper CPE hardware.”
- “1. Foilage penetration with unlicensed frequency equipment. 2. Support from local county governements. 3. Availability of funds (low interest, long term)”
- “Spectrum and commercial grade reliable equipment.”
- “More spectrum”
- “Spectrum availability (unlicensed mostly), PtMP [point to multipoint] performance/quality with affordable gear.”
- “Throughput due to non-licensed spectrum interference. Terrain (NLOS [non-line of sight]). Trees (NLOS).”
- “Usable spectrum in the lower band to penetrate trees and hills.”
- “Capital, spectrum”
- “Capital to expand and keep up with customer demand at the same time”
- “Expensive upstream backhaul, limited licensed frequencies for our use, more spectrum.”
- “We need more spectrum that isn’t shared with everything.”
- “RF Interference”
- “Additional spectrum and more flexibility in FCC rules.”
- “We need cheaper tower rent.”
- “Last-mile unlicensed bandwidth constraints. Most last-mile systems are built for residential needs, but we’re a business-only provider. We need carrier-grade equipment/features at reasonable prices to compete with Comcast.”

Table A.19 – *Continued from previous page*

- “Getting out of the last mile delivery business.”
- “Reliable, affordable upstream providers”
- “Price to fiber between sites.”
- “Typically it’s the intergration of all OSS [operational support system] platforms (CRM [customer relationship management], Billing, Inventory, Monitoring, field services, trouble tickets, etc...”
- “Higher capacity back hauls 300-500 Mbs at distance is not enough. More licensed spectrum. Better financing for the industry. More granular Speed control. Truly automatic fault tolerance. Better NOS products.”
- “Serving people behind massive amounts of foliage. TV Whitespaces spectrum will likely correct this but until equipment is made available it is not uncommon for me to tell 50% of my potential clients asking for service that we cannot serve them due to obstructions, often trees.”
- “Redundant backhaul mesh implementation, and lack of interference free frequencies.”
- “Competition from ILEC [incumbent local exchange carrier]”
- “county planning codes that we (and every other local wisp) pretend don’t exist.”
- “Interference analysis and mitigation.”
- “Need for more spectrum is the biggest problem. Also, some vendors have supply issues and it gets frustrating to have to wait for equipment.”
- “Platform agnostic monitoring and management software.”
- “Access to cheap bandwidth”
- “More spectrum in sub 5GHZ bands allocated specific to fixed wireless”
- “customer support, techs are good at fixing problems, not supporting end users, and no one who has any useful skill set wants to do customer support, so finding good people for end user support is next to impossible.”
- “The spread of the virus quickly.”
- “Under capitalized”
- “Spectrum sharing. Ownership of spectrum. Customer acceptance of required billing rates. They can’t cancel HBO and expect to get movies for free via the internet. All data traffic has a cost associated to it and video content is, by a huge margin, the highest cost data stream.”
- “Availability of back haul in rural areas.”
- “Capital for expansion.”
- “More unlicensed frequency from the FCC”
- “The cost of upstream bandwidth.”
- “There is no easy way to reroute traffic across our backbone when a section fails, must be done manually due to different equipment manufacturers and wireless vs fiber connectivity.”
- “Clear clean spectrum”
- “High cost of upstream bandwidth”

Table A.20: Q20. Anything else you want to share about your network and organization?

- “Still a startup: [url]”
- “Keep the government out of my business!!!!”
- “We are 8 months old and have about 115 subscribers. We have lost only 3 total since OCT 2011. We are very excited about the value and robustness the TV White Space Technology will bring to our WISP.”
- “We basically serve the rural area unserved by the telcos, and cable companies. We spend on average over $2 per customer per month for gas. We do must of our service calls with minivans to get better economy. We average less than 1 customer per square mile and many WISP average fewer customer per sqaure mile then we do. Most WISP get no government money, or other subsidies to get this done.”
- “We grew from 0 to 1400 customers in 18 months.”
- “We have been building this network for 15 years, there are many different generations of access points and CPE types.”
- “We are small but continue to grow and invest as the capital comes in. Company does not owe any money other than monthly operating expenses, e.g. payroll, tower rent, bandwidth purchase from upstream...”
- “Our network is very reliable because of many levels of redundancy. The outages/downtimes listed above are the effects on one customer or circuit. We almost never have network-wide or tower-wide issues anymore because of the redundancy and high-quality switching equipment we have put in place.”
- “[WISP] is one of the largest and oldest Fixed Wireless Providers in [service area]. The network currently reaches 80% of the [metro area] with spot coverage in other areas. In 2008 Coverage was extended to Customers in [second city in different state]. In 2014 projects are scheduled for [first city] and [third and fourth cities in third state]. The company also specializes in convention & hospitality Internet services and hosting & co-Location services. We also own [an Internet exchange point] a large peering and IP transit provider.”
- “I believe the US federal government is making a huge mistake in not making USF [Universal Service Fund] / CAF [Connect America Fund] funding more accessible to WISPs. It is as if they want us to fail as an industry when WISPs have been the champions of serving broadband to rural America for over a decade. It is deplorable.”
- “Transitioning from PPPoE/bridged to MAC/routed”
- “We cover 40,000 square miles of very rural areas in [multiple Great Plains states], and maintain over 2500 miles of our own microwave backhaul systems - completely avoiding telco connections.”
- “google facebook”
- “I am short of time”

Table A.20 – *Continued from previous page*

- “This question: ‘What is the total bandwidth commitment from your subscribers to you?’ (i.e., how much total capacity have you sold to your customers?) mks: All customers are given a best effort service. They get up to the max that the network will support. But each individual customer’s hardware is capable of very high speeds far in excess of what the upstream network can support. I wasn’t sure how to answer it. Feel free to call me [phone number].”
  Also on outages. Our network is larger than the state of Connecticut. There are over 50 ‘tower sites’ out there. Something is always breaking. Most outages are fixed within minutes of a customer call. Many segments will auto reboot and fix themselves. No outages take down the entire network so I wasn’t exactly sure how to answer your questions. They were not granular enough give accurate answers.
- “the Broadband Stimulus program is a nightmare. creating publicly funded bloated wasteful unsustainable ‘corporations’ that compete with WISP’s”
- “Prices will have to go up, because ‘viewers’ are shifting away from legacy sat/cable delivery methods in favor of the ”random access” nature of streaming. After over 10 years of fairly steady growth in network traffic, our load has literally doubled since the beginning of this year.”
- “We are a CLEC [competitive local exchange carrier] and provide internet and voice service to some customers over fixed wireless”
- “We are a small [New England state] cooperative (non-profit), run by volunteers. This has been in existence and provided bandwidth to the area for 7-8 years. Originally a single T1, then 2; bandwidth now furnished via 5 mile wireless link from a major Level 3 ISP (20 mbps). Fiber Optic will come through town in a few months. We will purchase approx 40-50 additional mbps from this provider. This information is submitted only for comparison to the usual huge for-profit operation.”

Bibliography

[1] 3GPP. *Radio subsystem link control*. TS 05.08. 3rd Generation Partnership Project (3GPP), Nov. 2005. URL: http://www.3gpp.org/ftp/Specs/html-info/0508.htm.

[2] *Rapid Spanning Tree, IEEE 802.1w Standard*. IEEE. 2001.

[3] Africa Mobile Networks. Retrieved 1/2019. URL: http://www.africamobilenetworks.com.

[4] Aileen Agüero, Harsha de Silva, and Juhee Kang. "Bottom of the Pyramid Expenditure Patterns on Mobile Services in Selected Emerging Asian Countries". In: *Information Technologies & International Development* 7.3 (2011). ISSN: 1544-7529. URL: http://itidjournal.org/itid/article/view/758.

[5] Quamrul Alam, Mohammad Abu Yusuf, and Ken Coghill. "Village Phone Program, Commodification Of Mobile Phone Set And Empowerment Of Women". In: *J. Community Informatics* 5.3-4 (2009).

[6] John Alden. "Exploring The Value And Economic Valuation Of Spectrum". In: *ITU Reports on Broadband* (2012).

[7] Kashif Ali and Kurtis Heimerl. "Designing Sustainable Rural Infrastructure Through the Lens of OpenCellular". In: *Communications of the ACM* 61.8 (2018), pp. 22–25.

[8] Alliance for Affordable Internet. *2018 Affordability Report*. Oct. 2018. URL: https://a4ai.org/affordability-report/report/2018.

[9] Abhinav Anand, Veljko Pejovic, Elizabeth M. Belding, and David L. Johnson. "VillageCell: Cost Effective Cellular Connectivity in Rural Areas". In: *Proceedings of the Fifth International Conference on Information and Communication Technologies and Development*. ICTD '12. Atlanta, Georgia: ACM, 2012, pp. 180–189. ISBN: 978-1-4503-1045-1. DOI: 10.1145/2160673.2160698. URL: http://doi.acm.org/10.1145/2160673.2160698.

[10] Arm. *TrustZone*. URL: https://developer.arm.com/ip-products/security-ip/trustzone.

[11] Associação Brasileira de Provedores de Internet e Telecomunicações. Retreived 5/2019. URL: http://www.abrint.com.br/.

[12] Association for Progressive Communications. *Connecting The Unconnected: Supporting Community Networks And Other Community-based Connectivity Initiatives.* URL: `https://www.apc.org/en/project/connecting-unconnected-supporting-community-networks-and-other-community-based-connectivity`.

[13] GSM Association. *Mobile Development Intelligence, Network Coverage Population.* URL: `https://mobiledevelopmentintelligence.com`.

[14] Robert C Atkinson and Ivy E Schultz. *Broadband In America: Where It Is And Where It Is Going.* 2009.

[15] Azotel Technologies. Retrieved 5/2019. URL: `http://www.azotel.com/`.

[16] BaiCells. Retrieved 12/2018. URL: `http://na.baicells.com/`.

[17] Roger Baig, Ramon Roca, Felix Freitag, and Leandro Navarro. “Guifi.net, a Crowdsourced Network Infrastructure Held in Common”. In: *Computer Networks* 90 (2015), pp. 150–165.

[18] Katelin Bailey, Luis Ceze, Steven D Gribble, and Henry M Levy. “Operating System Implications of Fast, Cheap, Non-Volatile Memory”. In: *HotOS*. Vol. 13. 2011, pp. 2–2.

[19] Bandwidth. Retrieved 9/2018. URL: `https://www.bandwidth.com`.

[20] Somprakash Bandyopadhyay, Kazuo Hasuike, S Horisawa, and S Tawara. “An Adaptive MAC Protocol For Wireless Ad Hoc Community Network (WACNet) Using Electronically Steerable Passive Array Radiator Antenna”. In: *Global Telecommunications Conference, 2001. GLOBECOM’01. IEEE*. Vol. 5. IEEE. 2001, pp. 2896–2900.

[21] Arijit Banerjee, Rajesh Mahindra, Karthik Sundaresan, Sneha Kasera, Kobus Van der Merwe, and Sampath Rangarajan. “Scaling the LTE Control-Plane for Future Mobile Access”. In: *Proceedings of the 11th ACM Conference on Emerging Networking Experiments and Technologies*. ACM. 2015, p. 19.

[22] Mary Claire Barela, Mae Sincere Blanco, Philip Martinez, Miguel Carlo Purisima, Kurtis Heimerl, Matthew Podolsky, Eric Brewer, and Cedric Angelo Festin. “Towards Building a Community Cellular Network in the Philippines: Initial Site Survey Observations”. In: *Proceedings of the Eighth International Conference on Information and Communication Technologies and Development*. ACM. 2016, p. 55.

[23] Sarbani Banerjee Belur, Meghna Khaturia, and Nanditha P Rao. “Community-led Networks for Sustainable Rural Broadband in India: the Case of Gram Marg”. In: *Community Networks: the Internet by the People, for the People* (), p. 193.

[24] Louise Benjamin. “Working It Out Together: Radio Policy From Hoover To The Radio Act Of 1927”. In: *Journal of Broadcasting & Electronic Media* 42.2 (1998), pp. 221–236.

[25] Vaduvur Bharghavan, Alan Demers, Scott Shenker, and Lixia Zhang. "MACAW: a Media Access Protocol for Wireless LAN's". In: *Proceedings of the Conference on Communications Architectures, Protocols and Applications*. SIGCOMM '94. London, United Kingdom: ACM, 1994, pp. 212–225. ISBN: 0-89791-682-4. DOI: 10.1145/190314.190334. URL: http://doi.acm.org/10.1145/190314.190334.

[26] Richard Binder, N Abramson, Franklin Kuo, A Okinaka, and D Wax. "ALOHA Packet Broadcasting: A Retrospect". In: *Proceedings of the May 19-22, 1975, National Computer Conference and Exposition*. ACM. 1975, pp. 203–215.

[27] Colin R Blackman. "Universal Service: Obligation or Opportunity?" In: *Telecommunications Policy* 19.3 (1995), pp. 171–176. URL: http://EconPapers.repec.org/RePEc:eee:telpol:v:19:y:1995:i:3:p:171-176.

[28] J. E. Blumenstock, M. Callen, T. Ghani, and L. Koepke. "Promises and Pitfalls of Mobile Money in Afghanistan: Evidence from a Randomized Control Trial". In: *Conference Proceedings of the Seventh International Conference on Information and Communication Technologies and Development*. NY, USA: New York, 2015.

[29] Bart Braem, Arjuna Sathiaseelan, Andres Arcia-Moret, Marco Zennaro, Ermanno Pietrosemoli, and Jose Saldana. "Alternative Network Deployments: Taxonomy, Characterization, Technologies, and Architectures". In: (2016).

[30] Eric Brewer. "CAP Twelve Years Later: How The "Rules" Have Changed". In: *Computer* 2 (2012), pp. 23–29.

[31] Milind M Buddhikot, Irwin Kennedy, Frank Mullany, and Harish Viswanathan. "Ultra-Broadband Femtocells via Opportunistic Reuse of Multi-Operator and Multi-Service Spectrum". In: *Bell Labs Technical Journal* 13.4 (2009), pp. 129–143.

[32] Danijela Cabric, Shridhar Mubaraq Mishra, Daniel Willkomm, Robert Brodersen, and Adam Wolisz. "A Cognitive Radio Approach for Usage of Virtual Unlicensed Spectrum". In: *In Proc. of 14th IST Mobile Wireless Communications Summit 2005*. 2005.

[33] Champaign-Urbana Community Wireless Network. Retrieved 12/2018. URL: http://www.govtech.com/dc/articles/CUWiN-Wirelessing-the-Revolution-with-Open.html.

[34] Robin IC Chiang, Gerard B Rowe, and Kevin W Sowerby. "A Quantitative Analysis Of Spectral Occupancy Measurements For Cognitive Radio". In: *2007 IEEE 65th Vehicular Technology Conference-VTC2007-Spring*. IEEE. 2007, pp. 3016–3020.

[35] Philippe Chone, Laurent Flochel, and Anne Perrot. "Allocating and Funding Universal Service Obligations in a Competitive Market". In: *International Journal of Industrial Organization* 20.9 (2002), pp. 1247–1276. ISSN: 0167-7187. DOI: http://dx.doi.org/10.1016/S0167-7187(01)00077-7. URL: http://www.sciencedirect.com/science/article/pii/S0167718701000777.

[36] Cisco. *Global Mobile Data Traffic Forecast Update, 2017–2022*. Feb. 2019.

[37] David Clark. “The Design Philosophy of the DARPA Internet Protocols”. In: *ACM SIGCOMM Computer Communication Review*. Vol. 18. 4. ACM. 1988, pp. 106–114.

[38] David D Clark, John Wroclawski, Karen R Sollins, and Robert Braden. “Tussle In Cyberspace: Defining Tomorrow’s Internet”. In: *ACM SIGCOMM Computer Communication Review*. Vol. 32. 4. ACM. 2002, pp. 347–356.

[39] Ronald H Coase. “The Federal Communications Commission”. In: *Journal of Law and Economics* 2 (1959), pp. 1–40.

[40] Congreso de la Republica del Peru. *Ley No. 30083*. 2013. URL: http://www.leyes.congreso.gob.pe/Documentos/Leyes/30083.pdf.

[41] Connect America Fund. Retreived 2/2019. URL: https://www.fcc.gov/general/connect-america-fund-caf.

[42] Peter C Cramton. “Money Out Of Thin Air: The Nationwide Narrowband Pcs Auction”. In: *Journal of Economics & Management Strategy* 4.2 (1995), pp. 267–343.

[43] Robert W Crandall. “New Zealand Spectrum Policy: A Model For The United States?” In: *The Journal of Law and Economics* 41.S2 (1998), pp. 821–840.

[44] Arthur S DeVany, Ross D Eckert, Charles J Meyers, and Donald J O’Hara. “A Property System For Market Allocation Of The Electromagnetic Spectrum: A Legal-economic-engineering Study”. In: *Stan. L. Rev.* 21 (1968), p. 1499.

[45] Aditya Dhananjay, Matt Tierney, Jinyang Li, and Lakshminarayanan Subramanian. “WiRE: A New Rural Connectivity Paradigm”. In: *SIGCOMM*. 2011, pp. 462–463.

[46] Jonathan Donner. *After Access: Inclusion, Development, And A More Mobile Internet*. MIT Press, 2015.

[47] Jonathan Donner. “The Rules of Beeping: Exchanging Messages Via Intentional “Missed Calls” on Mobile Phones”. In: *Journal of Computer-Mediated Communication* 13.1 (2007).

[48] Marisa Elena Duarte. *Network Sovereignty: Building the Internet Across Indian Country*. University of Washington Press, 2017.

[49] Sarah Earl, Chaitali Sinha, and Matthew L Smith. “Innovations in Evaluating ICT4D Research”. In: *Connecting ICTs to Development* (2013), p. 241.

[50] T.M. Egyedi and D.C. Mehos. *Inverse Infrastructures: Disrupting Networks from Below*. Edward Elgar, 2012. ISBN: 9781781952290. URL: http://books.google.com/books?id=wG4c7s8uKjUC.

[51] Endaga. Retrieved 11/2014. URL: http://www.endaga.com/.

[52] Ente Nacional de Comunicaciones. *Resolución 4958/2018*. Aug. 2018. URL: http://servicios.infoleg.gob.ar/infolegInternet/anexos/310000-314999/313590/norma.htm.

[53] Tugba Erpek, Mark Lofquist, and Ken Patton. *Spectrum Occupancy Measurements, Loring Commerce Centre, Limestone, Maine, September 18-20, 2007*. Tech. rep. Shared Spectrum Company, Sept. 2007. URL: http://www.sharedspectrum.com/wp-content/uploads/Loring%5C_Spectrum%5C_Occupancy%5C_Measurements%5C_v2%5C_3.pdf.

[54] Facebook. *Magma*. URL: https://code.fb.com/open-source/magma/.

[55] Fairwaves. Retrieved 11/2014. URL: http://www.fairwaves.co/.

[56] Kevin Fall. "A Delay-tolerant Network Architecture For Challenged Internets". In: *SIGCOMM '03: Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications*. Karlsruhe, Germany: ACM, 2003, pp. 27–34. ISBN: 1-58113-735-4. DOI: http://doi.acm.org/10.1145/863955.863960.

[57] Federal Communications Commission. *Citizens Broadband Radio Service (FCC 15-47)*. 2015. URL: https://docs.fcc.gov/public/attachments/FCC-15-47A1.pdf.

[58] Federal Communications Commission. "Unlicensed NII Devices in the 5GHz Range (62 FR 4649)". In: *Federal Register* 62.21 (Jan. 1997), pp. 4649–4657.

[59] Federal Communications Commission. "Unlicensed Operation in the TV Broadcast Bands (74 FR 7314)". In: *Federal Register* 74.30 (Feb. 2009), pp. 7314–7332. URL: https://federalregister.gov/a/E9-3279.

[60] Federal Communications Commission. "Unlicensed Use of TV Band and 600 MHz Band Spectrum". In: *Federal Register* 79.225 (Nov. 2014), pp. 69709–69746.

[61] Claude Fischer. *America Calling: A Social History of the Telephone to 1940*. Berkeley, CA: University of California Press, 1992.

[62] Pantelis A Frangoudis, George C Polyzos, and Vasileios P Kemerlis. "Wireless Community Networks: An Alternative Approach For Nomadic Broadband Network Access". In: *IEEE Communications Magazine* 49.5 (2011).

[63] FreeSWITCH. Retrieved 4/2013. URL: http://www.freeswitch.org/.

[64] Freifunk Wireless Network. Retrieved 12/2018. URL: https://freifunk.net/.

[65] Hernan Galperin and François Bar. "The Microtelco Opportunity: Evidence from Latin America". In: *Information Technologies and International Development*. Vol. 3. 2. 2006.

[66] Globe Telecom, Inc. Retrieved 9/2018. URL: https://www.globe.com.ph/.

[67] Ismael Gomez-Miguelez, Andres Garcia-Saavedra, Paul D Sutton, Pablo Serrano, Cristina Cano, and Doug J Leith. "srsLTE: an Open-Source Platform for LTE Evolution and Experimentation". In: *Proceedings of the Tenth ACM International Workshop on Wireless Network Testbeds, Experimental Evaluation, and Characterization*. ACM. 2016, pp. 25–32.

[68] Vânia Gonçalves and Sofie Pollin. "The Value Of Sensing For TV White Spaces". In: *New Frontiers in Dynamic Spectrum Access Networks (DySPAN), 2011 IEEE Symposium on*. IEEE. 2011, pp. 231–241.

[69] Jim Gray, Wyman Chong, Tom Barclay, Alex Szalay, and Jan Vandenberg. "TeraScale SneakerNet: Using Inexpensive Disks For Backup, Archiving, And Data Exchange". In: *arXiv preprint cs/0208011* (2002).

[70] GSM Association. *Mobile Access – The Last Mile*. July 2014. URL: https://gsmaintelligence.com/files/analysis/?file=140725-mobile-access.pdf.

[71] GSM Association. *Powering Telecoms: West Africa Market Analysis Sizing the Potential for Green Telecoms in Nigeria and Ghana*. Retrieved 10/2013. URL: http://www.gsma.com/mobilefordevelopment/wp-content/uploads/2013/03/GPM-Market-Analysis-West-Africa-.pdf.

[72] GSM Association. *Spectrum Pricing In Developing Countries: Evidence To Support Better And More Affordable Mobile Services*. URL: https://www.gsma.com/latinamerica/wp-content/uploads/2018/07/Spectrum-pricing-developing-countries.pdf.

[73] GSM Association. *The Mobile Economy 2014*. URL: http://www.gsmamobileeconomy.com/GSMA%5C_ME%5C_Report%5C_2014%5C_R2%5C_WEB.pdf.

[74] GSM Association. *Universal Service Fund Study*. Apr. 2013. URL: http://www.gsma.com/publicpolicy/wp-content/uploads/2013/04/GSMA-USF-Main-report-final.pdf.

[75] GSM Association. *Unlocking Rural Coverage: Enablers for commercially sustainable mobile network expansion*. July 2016. URL: https://www.gsma.com/mobilefordevelopment/wp-content/uploads/2016/07/Unlocking-Rural-Coverage-enablers-for-commercially-sustainable-mobile-network-expansion%5C_English.pdf.

[76] GSMA Intelligence. *Country Overview: Philippines Growth Through Innovation*. Retrieved 9/2018. URL: https://www.gsmaintelligence.com/research/?file=141201-philippines.pdf.

[77] S. Guo, M. H. Falaki, E. A. Oliver, E. A. Oliver, S. Ur Rahman, S. Ur Rahman, A. Seth, M. A. Zaharia, and S. Keshav. "Very Low-cost Internet Access Using Kiosknet". In: *SIGCOMM Comput. Commun. Rev.* 37.5 (2007), pp. 95–100. ISSN: 0146-4833. DOI: http://doi.acm.org/10.1145/1290168.1290181.

[78] David Gurney, Greg Buchwald, Larry Ecklund, SL Kuffner, and John Grosspietsch. "Geo-location Database Techniques for Incumbent Protection in the TV White Space". In: *New Frontiers in Dynamic Spectrum Access Networks, 2008. DySPAN 2008. 3rd IEEE Symposium on*. IEEE. 2008, pp. 1–9.

[79] Kate Harrison. "A Quantitative Approach To Wireless Spectrum Regulation". PhD thesis. UC Berkeley, 2015.

[80] Kate Harrison. "Cognitive Radios in the TV Whitespaces: Challenges and Opportunities". MA thesis. UC Berkeley, 2011.

[81] Kate Harrison, Shridhar Mubaraq Mishra, and Anant Sahai. "How Much White-space Capacity Is There?" In: *New Frontiers in Dynamic Spectrum, 2010 IEEE Symposium on.* IEEE. 2010, pp. 1–10.

[82] Shaddi Hasan, Mary Claire Barela, Matthew Johnson, Eric Brewer, and Kurtis Heimerl. "Scaling Community Cellular Networks with CommunityCellularManager". In: *16th USENIX Symposium on Networked Systems Design and Implementation.* NSDI'19. 2019, pp. 735–750.

[83] Shaddi Hasan, Yahel Ben-David, Max Bittman, and Barath Raghavan. "The Challenges of Scaling WISPs". In: *Proceedings of the 2015 Annual Symposium on Computing for Development.* ACM. 2015, pp. 3–11.

[84] Shaddi Hasan, Kurtis Heimerl, Kate Harrison, Kashif Ali, Sean Roberts, Anant Sahai, and Eric Brewer. "GSM Whitespaces: An Opportunity For Rural Cellular Service". In: *2014 IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN).* IEEE. 2014, pp. 271–282.

[85] Thomas W Hazlett, David Porter, and Vernon Smith. "Radio Spectrum And The Disruptive Clarity Of Ronald Coase". In: *The Journal of Law and Economics* 54.S4 (2011), S125–S165.

[86] Kurtis Heimerl, Kashif Ali, Joshua Blumenstock, Brian Gawalt, and Eric 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: USENIX Association, 2013.

[87] Kurtis Heimerl and Eric Brewer. "The Village Base Station". In: *Proceedings of the 4th ACM Workshop on Networked Systems for Developing Regions.* NSDR '10. San Francisco, California: ACM, 2010, 14:1–14:2. ISBN: 978-1-4503-0193-0.

[88] Kurtis Heimerl, Shaddi Hasan, Kashif Ali, Eric Brewer, and Tapan 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: ACM, 2013, pp. 2–12.

[89] Kurtis Heimerl, Shaddi Hasan, Kashif Ali, Tapan Parikh, and Eric Brewer. "A Longitudinal Study Of Local, Sustainable, Small-scale Cellular Networks". In: *Information Technologies & International Development* 11.1 (2015), pp–1.

[90] Leo Herzel. ""Public Interest" And The Market In Color Television Regulation". In: *University of Chicago Law Review* 18.4 (1951), pp. 802–816.

[91] J Holdren and E Lander. "Realizing The Full Potential Of Government-held Spectrum To Spur Economic Growth". In: *President's Council of Advisors on Science and Technology, Washington, DC, USA, Tech. Rep* (2012).

[92] Internet Society. *Community Networks*. Retrieved 4/2019. URL: https://www.internetsociety.org/issues/community-networks/.

[93] Md Habibul Islam, Choo Leng Koh, Ser Wah Oh, Xianming Qing, Yoke Yong Lai, Cavin Wang, Ying-Chang Liang, Bee Eng Toh, Francois Chin, Geok Leng Tan, et al. "Spectrum Survey In Singapore: Occupancy Measurements And Analyses". In: *2008 3rd International Conference on Cognitive Radio Oriented Wireless Networks and Communications (CrownCom 2008)*. IEEE. 2008, pp. 1–7.

[94] ITU. *Telecommunication For Rural And Remote Areas*. Recommendation D.19. ITU Telecommunication Development Sector, 2014. URL: https://www.itu.int/en/ITU-D/Conferences/WTDC/Documents/D-TDC-WTDC-2014-PDF-E.pdf.

[95] Steven J. Jackson, Alex Pompe, and Gabriel Krieshok. "Repair Worlds: Maintenance, Repair, and ICT for Development in Rural Namibia". In: *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work*. CSCW '12. Seattle, Washington, USA: ACM, 2012, pp. 107–116. ISBN: 978-1-4503-1086-4. DOI: 10.1145/2145204.2145224. URL: http://doi.acm.org/10.1145/2145204.2145224.

[96] Steven J. Jackson, Alex Pompe, and Gabriel Krieshok. "Things Fall Apart: Maintenance, Repair, and Technology for Education Initiatives in Rural Namibia". In: *iConference*. 2011.

[97] RS Jain. "Spectrum Auctions In India: Lessons From Experience". In: *Telecommunications Policy* 25.10-11 (2001), pp. 671–688.

[98] Saurabh Jain and Dharma P Agrawal. "Wireless Community Networks". In: *Computer* 36.8 (2003), pp. 90–92.

[99] E. Jang, M. C. Barela, M. Johnson, P. Martinez, C. Festin, M. Lynn, J. Dionisio, and K. Heimerl. "Crowdsourcing Rural Network Maintenance and Repair via Network Messaging". In: *Conference on Human Factors in Computing Systems (CHI)*. 2018.

[100] Robert Jensen. "The Digital Provide: Information (Technology), Market Performance, and Welfare in the South Indian Fisheries Sector". In: *The Quarterly Journal of Economics* 122.3 (Aug. 2007), pp. 879–924.

[101] Xin Jin, Li Erran Li, Laurent Vanbever, and Jennifer Rexford. "SoftCell: Scalable and Flexible Cellular Core Network Architecture". In: *Proceedings of the Ninth ACM Conference on Emerging Networking Experiments and Technologies*. ACM. 2013, pp. 163–174.

[102] Matthew Johnson, Spencer Sevilla, Esther Jang, and Kurtis Heimerl. "dLTE: Building a more WiFi-like Cellular Network: (Instead of the Other Way Around)". In: *Proceedings of the 17th ACM Workshop on Hot Topics in Networks, HotNets 2018, Redmond, WA, USA, November 15-16, 2018*. 2018, pp. 8–14. DOI: 10.1145/3286062.3286064. URL: https://doi.org/10.1145/3286062.3286064.

[103] Anuradha Joshi and Mick Moore. "Institutionalised Co-production: Unorthodox Public Service Delivery in Challenging Environments". In: *The Journal of Development Studies* 40.4 (2004), pp. 31–49. URL: http://ideas.repec.org/a/taf/jdevst/v40y2004i4p31-49.html.

[104] Robert Kahn. "The Organization Of Computer Resources Into A Packet Radio Network". In: *IEEE Transactions on communications* 25.1 (1977), pp. 169–178.

[105] M. Kam, A. Mathur, A. Kumar, and J. Canny. "Designing Digital Games for Rural Children: A Study of Traditional Village Games in India". In: *Conference on Human Factors in Computing Systems (CHI)*. 2009.

[106] Kannel. Retrieved 2/2019. URL: https://www.kannel.org/.

[107] N. Kumar, T. Perrier, M. Desmond, K. Israel-Ballard, V. Kumar, S. Mahapatra, A. Mishra, S. Agarwal, R. Gandhi, P. Lal, and R. Anderson. "Projecting Health: Community-led Video Education For Maternal Health". In: *Seventh International Conference on Information and Communication Technologies and Development*. 2015.

[108] Wai Fung Lam. "Institutional Design of Public Agencies and Coproduction: A Study of Irrigation Associations in Taiwan". In: *World Development* 24.6 (June 1996), pp. 1039–1054. URL: http://ideas.repec.org/a/eee/wdevel/v24y1996i6p1039-1054.html.

[109] Mihai Letia, Nuno Preguiça, and Marc Shapiro. "CRDTs: Consistency Without Concurrency Control". In: *arXiv preprint arXiv:0907.0929* (2009).

[110] Aemen Lodhi, Natalie Larson, Amogh Dhamdhere, Constantine Dovrolis, et al. "Using PeeringDB to Understand the Peering Ecosystem". In: *ACM SIGCOMM Computer Communication Review* 44.2 (2014), pp. 20–27.

[111] Brandon Lucia and Benjamin Ransford. "A Simpler, Safer Programming And Execution Model For Intermittent Systems". In: *ACM SIGPLAN Notices* 50.6 (2015), pp. 575–585.

[112] Alexandru Martian. "Evaluation Of Spectrum Occupancy In Urban And Rural Environments Of Romania". In: *Revue Roumaine des Sciences Techniques-Electrotechn. et Energ* 1 (2014), pp. 87–96.

[113] Andrés Martínez Fernández, José Vidal Manzano, Javier Simó Reigadas, Ignacio Prieto Egido, Adrián Agustín de Dios, Juan Paco, and álvaro Rendón. "The TUCAN3G Project: Wireless Technologies For Isolated Rural Communities In Developing Countries Based On 3G Small-cell Deployments". In: *IEEE Communications Magazine* 54.7 (2016), pp. 36–43.

[114] Mayu Telecommunications SAC. Retrieved 1/2019. URL: http://www.mayutel.com.

[115] MA McHenry, D McCloskey, and G Lane. *Spectrum Occupancy Measurements, Location 4 Of 6: Republican National Convention, New York City, New York, August 30, 2004-september 3, 2004*. Tech. rep. Shared Spectrum Company, Aug. 2005. URL: `http://www.sharedspectrum.com/wp-content/uploads/4%5C_NSF%5C_NYC%5C_Report.pdf`.

[116] Mark A McHenry and Karl Steadman. "Spectrum Occupancy Measurements, Location 1 Of 6: Riverbend Park, Great Falls, Virginia". In: *Shared Spectrum Company Report* (2005), pp. 1–55.

[117] MikroTik. Retrieved 5/2019. URL: `https://www.mikrotik.com/`.

[118] Rohan Murty, Ranveer Chandra, Thomas Moscibroda, and Paramvir Bahl. "Senseless: A Database-driven White Spaces Network". In: *Mobile Computing, IEEE Transactions on* 11.2 (2012), pp. 189–203.

[119] Nagios. Retrieved 5/2019. URL: `https://www.nagios.org/`.

[120] United Nations. "Transforming Our World: The 2030 Agenda For Sustainable Development". In: *Resolution adopted by the General Assembly* (2015).

[121] Nexmo. Retrieved 4/2013. URL: `http://www.nexmo.com/`.

[122] NextEPC. Retrieved 5/2019. URL: `https://nextepc.org/`.

[123] Binh Nguyen, Tian Zhang, Bozidar Radunovic, Ryan Stutsman, Thomas Karagiannis, Jakub Kocur, and Jacobus Van der Merwe. "ECHO: A Reliable Distributed Cellular Core Network for Hyper-Scale Public Clouds". In: *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking*. ACM. 2018, pp. 163–178.

[124] Navid Nikaein, Raymond Knopp, Florian Kaltenberger, Lionel Gauthier, Christian Bonnet, Dominique Nussbaum, and Riadh Ghaddab. "OpenAirInterface: an Open LTE Network in a PC". In: *Proceedings of the 20th Annual International Conference on Mobile Computing and Networking*. ACM. 2014, pp. 305–308.

[125] Nuran Wireless. Retrieved 11/2014. URL: `http://nuranwireless.com/`.

[126] Susan O'Donnell, George Kakekaspan, Brian Beaton, Brian Walmark, Raymond Mason, and Michael Mak. "A New Remote Community-owned Wireless Communication Service: Fort Severn First Nation Builds Their Local Cellular System With Keewaytinook Mobile". In: *Canadian Journal of Communication* 36.4 (2012).

[127] OpenBTS. Retrieved 4/2013. URL: `http://openbts.org`.

[128] Osmocom. Retrieved 9/2018. URL: `https://osmocom.org/projects/cellular-infrastructure`.

[129] Elinor Ostrom. "Crossing the Great Divide: Coproduction, Synergy, and Development". In: *World Development* 24.6 (1996), pp. 1073–1087.

[130] Elinor Ostrom. “Metropolitan Reform: Propositions Derived From Two Traditions”. In: *Social Science Quarterly* 53.3 (1972), pp. 474–493.

[131] Roger B Parks, Paula C Baker, Larry Kiser, Ronald Oakerson, Elinor Ostrom, Vincent Ostrom, Stephen L Percy, Martha B Vandivort, Gordon P Whitaker, and Rick Wilson. “Consumers as Coproducers of Public Services: Some Economic and Institutional Considerations”. In: *Policy Studies Journal* 9.7 (1981), pp. 1001–1011.

[132] Girish Patel and Steven Dennett. “The 3GPP and 3GPP2 Movements Toward An All-IP Mobile Network”. In: *IEEE Personal Communications* 7.4 (2000), pp. 62–64.

[133] Kishor Patil, Knud Erik Skouby, and Ramjee Prasad. “Cognitive Access To TVWS In India: Tv Spectrum Occupancy And Wireless Broadband For Rural Areas”. In: *2013 16th International Symposium on Wireless Personal Multimedia Communications (WPMC)*. IEEE. 2013, pp. 1–5.

[134] Rabin Patra, Sergiu Nedevschi, Sonesh Surana, Anmol Sheth, Lakshminarayanan Subramanian, and Eric Brewer. “WiLDNet: Design and Implementation of High Performance WiFi Based Long Distance Networks”. In: *4th USENIX Symposium on Networked Systems Design and Implementation*. 2007.

[135] Jon M Peha. “Sharing Spectrum through Spectrum Policy Reform and Cognitive Radio”. In: *Proceedings of the IEEE* 97.4 (2009), pp. 708–719.

[136] Alex (Sandy) Pentland, Richard Fletcher, and Amir Hasson. “DakNet: Rethinking Connectivity in Developing Nations”. In: *Computer* 37.1 (2004), pp. 78–83. ISSN: 0018-9162. DOI: `http://dx.doi.org/10.1109/MC.2004.1260729`.

[137] Powercode. Retrieved 5/2019. URL: `https://powercode.com/`.

[138] Zafar Ayyub Qazi, Phani Krishna Penumarthi, Vyas Sekar, Vijay Gopalakrishnan, Kaustubh Joshi, and Samir R Das. “KLEIN: A Minimally Disruptive Design for an Elastic Cellular Core”. In: *Proceedings of the Symposium on SDN Research*. ACM. 2016, p. 2.

[139] Zafar Ayyub Qazi, Melvin Walls, Aurojit Panda, Vyas Sekar, Sylvia Ratnasamy, and Scott Shenker. “A High Performance Packet Core for Next Generation Cellular Networks”. In: *Proceedings of the Conference of the ACM Special Interest Group on Data Communication*. ACM. 2017, pp. 348–361.

[140] Range Networks. Retrieved 4/2013. URL: `http://www.rangenetworks.com/`.

[141] Red Hook WiFi. Retrieved 12/2018. URL: `https://redhookwifi.org/`.

[142] Carlos Rey-Moreno, William D Tucker, Domonic Cull, and R Blom. “Making A Community Network Legal Within The South African Regulatory Framework”. In: *Proceedings of the Seventh International Conference on Information and Communication Technologies and Development*. ACM. 2015, p. 57.

[143] Rhizomatica. Retrieved 4/2013. URL: `http://rhizomatica.org/`.

[144] Lawrence G Roberts. "ALOHA Packet System With And Without Slots And Capture". In: *ACM SIGCOMM Computer Communication Review* 5.2 (1975), pp. 28–42.

[145] Jerome H Saltzer, David P Reed, and David D Clark. "End-to-end Arguments In System Design". In: *Technology* 100 (1984), p. 0661.

[146] Srivatsan Sankaranarayanan, Panagiotis Papadimitratos, Amitabh Mishra, and Steven 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*. IEEE. 2005, pp. 279–288.

[147] Sensorly. Retrieved 11/2013. URL: `http://sensorly.com`.

[148] Denny Setiawan. *Overview of Spectrum Issues in Indonesia*. London, June 2010. URL: `http://denysetia.files.wordpress.com/2010/09/290610-broadband-wireless-gsma-london-denny-ins.ppt`.

[149] Spencer Sevilla, Pathirat Kosakanchit, Matthew Johnson, and Kurtis Heimerl. "Building Community LTE Networks with CoLTE". In: *The Community Network Manual: How to Build the Internet Yourself*. Ed. by Luca Belli. 2018. Chap. 3, pp. 75–102.

[150] K. Shah, P. Martinez, G. M. D. Cruz, J. D. Joshua Blumenstock, and K. Heimerl. "An Investigation of Phone Upgrades in Remote Community Cellular Networks". In: *International Conference on Information and Communication Technologies and Development*. Ed. by Pakistan Lahore. 2017.

[151] Justine Sherry, Shaddi Hasan, Colin Scott, Arvind Krishnamurthy, Sylvia Ratnasamy, and Vyas Sekar. "Making Middleboxes Someone Else's Problem: Network Processing As A Cloud Service". In: *ACM SIGCOMM Computer Communication Review* 42.4 (2012), pp. 13–24.

[152] Lixin Shi, Paramvir Bahl, and Dina Katabi. "Beyond Sensing: Multi-GHz Realtime Spectrum Analytics". In: *12th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 15)*. 2015, pp. 159–172.

[153] Steve Song, Carlos Rey-Moreno, Anriette Esterhuysen, Mike Jensen, and Leandro Navarro. "The Rise And Fall And Rise Of Community Networks". In: *Global Information Society Watch* 19 (2018), pp. 7–10.

[154] Menahem Spiegel. "Universal Service Obligation". English. In: *Expanding Competition in Regulated Industries*. Ed. by Michael A. Crew. Vol. 37. Topics in Regulatory Economics and Policy Series. Springer US, 2000, pp. 83–96. ISBN: 978-1-4419-5006-2. DOI: `10.1007/978-1-4757-3192-7_5`. URL: `http://dx.doi.org/10.1007/978-1-4757-3192-7%5C_5`.

[155] Janaki Srinivasan and Jenna Burrell. "Revisiting the fishers of Kerala, India". In: *Proceedings of the Sixth International Conference on Information and Communication Technologies and Development: Full Papers-Volume 1*. ACM. 2013, pp. 56–66.

[156] Sonesh Surana, Melissa Ho, Rabin Patra, and Lakshminarayanan Subramanian. "Designing Healthcare Systems In The Developing World: The Role Of Computer Science Systems Research". In: New York University: Technical Report. URL: `http://www.cs.nyu.edu/~lakshmi/systemsforhealth.pdf`.

[157] Sonesh Surana, Sergui Nedevschi Rabin Patra, Manuel Ramos, Lakshminarayanan Subramanian, Yahel Ben-David, and Eric Brewer. "Beyond Pilots: Keeping Rural Wireless Networks Alive". In: *5th USENIX Symposium on Networked Systems Design and Implementation*. 2008.

[158] Telecom Infra Project. *OpenRAN Project Group*. URL: `https://telecominfraproject.com/openran/`.

[159] Telkomsel. *Telkomsel, 2006 Annual Report*. 2007. URL: `http://www.telkomsel.com/media/upload/pdf/AR2006.pdf`.

[160] Telkomsel. *Telkomsel, 2012 Annual Report*. Feb. 2013. URL: `http://www.telkomsel.com/media/upload/annualreport/(FINAL)%20AR%20TELKOMSEL%20CONTENT.pdf`.

[161] Ubiquiti Networks. *Ubiquiti Networks Q1 Fiscal Year 2015 Earnings*. Nov. 2014.

[162] Ubiquiti Networks. Retrieved 5/2019. URL: `https://www.ubnt.com/`.

[163] International Telecommunications Union. *Implementation of USO Program in Indonesia*. Jakarta, Indonesia, Mar. 2011. URL: `http://www.itu.int/ITU-D/.../Indonesia/Session2-BP3TI.pdf`.

[164] United States Department of Agriculture ReConnect Fund. Retreived 2/2019. URL: `https://www.usda.gov/reconnect`.

[165] Václav Valenta, Zbynek Fedra, Roman Marsalek, Geneviève Baudoin, and Martine Villegas. "Towards Cognitive Radio Networks: Spectrum Utilization Measurements In Suburb Environment". In: *2009 IEEE Radio and Wireless Symposium*. IEEE. 2009, pp. 352–355.

[166] Václav Valenta, Roman Maršálek, Geneviève Baudoin, Martine Villegas, Martha Suarez, and Fabien Robert. "Survey On Spectrum Utilization In Europe: Measurements, Analyses And Observations". In: *2010 Proceedings of the Fifth International Conference on Cognitive Radio Oriented Wireless Networks and Communications*. IEEE. 2010, pp. 1–5.

[167] Davide Vega, Roger Baig, Llorenç Cerdà-Alabern, Esunly Medina, Roc Meseguer, and Leandro Navarro. "A Technological Overview Of The Guifi.Net Community Network". In: *Computer Networks* 93 (2015), pp. 260–278.

[168] Davide Vega, Llorenc Cerda-Alabern, Leandro Navarro, and Roc Meseguer. "Topology Patterns Of A Community Network: guifi.net". In: *Wireless and Mobile Computing, Networking and Communications (WiMob), 2012 IEEE 8th International Conference on*. IEEE. 2012, pp. 612–619.

[169] Leonard Waverman, Meloria Meschi, and Melvyn Fuss. "The Impact of Telecoms on Economic Growth in Developing Countries". In: *Africa: the Impact of Mobile Phones* 2.1980 (2005), pp. 1–23. URL: `https://www.businessgrowthinitiative.org/KeyPracticeAreas/Documents/Forms/DispForm.aspx?ID=148`.

[170] Dennis Weller. "Auctions for Universal Service Obligations". In: *Telecommunications Policy* 23.9 (1999), pp. 645–674. ISSN: 0308-5961. DOI: `http://dx.doi.org/10.1016/S0308-5961(99)00048-8`. URL: `http://www.sciencedirect.com/science/article/pii/S0308596199000488`.

[171] Michael E. Whitman. "A Look at the Telecommunications ACT of 1996". In: *IS Management* 14.3 (1997), pp. 82–86. URL: `http://dblp.uni-trier.de/db/journals/ism/ism14.html#Whitman97`.

[172] Wireless Access Providers Association. Retrieved 5/2019. URL: `http://wapa.org.za/`.

[173] Wireless Internet Service Providers Association. Retrieved 2/2019. URL: `http://www.wispa.org/`.

[174] WISPMon. Retrieved 5/2019. URL: `https://www.wispmon.com/`.

[175] Tim Wu. *The Master Switch: The Rise And Fall Of Information Empires*. Vintage, 2010.

[176] Liang Yin, Kaishun Wu, Sixing Yin, Shufang Li, and Lionel M Ni. "Reuse of GSM White Space Spectrum for Cognitive Femtocell Access". In: *Parallel and Distributed Systems (ICPADS), 2012 IEEE 18th International Conference on*. IEEE. 2012, pp. 1–7.

[177] Mariya Zheleva, Arghyadip Paul, David L. Johnson, and Elizabeth Belding. "Kwiizya: Local Cellular Network Services in Remote Areas". In: *Proceedings of the 11th International Conference on Mobile Systems, Applications, and Services*. MobiSys '13. Taipai, Taiwan: ACM, 2013.