Community Cellular Networks

Kurtis Heimerl

Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2013-220
http://www.eecs.berkeley.edu/Pubs/TechRpts/2013/EECS-2013-220.html

December 18, 2013

Community Cellular Networks

Community Cellular Networks

by

Kurtis Heimerl

B.E. (University of Washington) 2007
M.S. (University of California, Berkeley) 2009

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, Co-chair
Assistant Professor Tapan Parikh, Co-chair
Associate Professor Jenna Burrell
Professor Ali Niknejad

Fall 2013

Abstract

Community Cellular Networks

by

Kurtis Heimerl
Doctor of Philosophy in Computer Science

University of California, Berkeley

Professor Eric Brewer, Co-chair
Assistant Professor Tapan Parikh, Co-chair

There are over six billion active cellular subscribers, spending a total of over a trillion dollars a year on communications. Despite this, hundreds of millions of people, primarily in rural areas, are still without network connectivity. There are two primary reasons for this: First, it is too expensive for traditional carriers to install equipment in far flung regions with small markets. Second, traditional carriers are the only organizations allowed to provide coverage; local people cannot.

To bring network connectivity to these remaining hundreds of millions of people, we introduce the concept of *community cellular networks*: small-scale, locally operated networks independent from traditional telecommunication firms. In support of these networks, we also introduce *virtual coverage*, a novel power saving mechanism in GSM networks. Virtual coverage allows us to reduce the total install cost of the network to the point where individual operators can run their own full-scale telecommunication firms.

We deployed the first community cellular network in the community of Desa[1], in rural Papua, Indonesia. This community of approximately 1500 people is far from existing network coverage. We developed numerous services for the community including a credit system, singing competition, and others. The network also made use of virtual coverage.

We evaluated community cellular and virtual coverage through this real-world installation in Desa. Over a six-month period, we saw over 100000 communications, reduced the night power draw by 56.6%, and generated over US$5000 in revenue for the local operator. This network is sustainable, generating over US$368 per month in profit for the local operator, even assuming complete financing of the install. We interviewed users of the system, discovering that most of their communication is with family and friends in other Indonesian areas. Lastly, we show that this network is significantly more locally-focused than a traditional firm in sub-Saharan Africa.

These results validate the community cellular model. We explore future work needed for this model to flourish, including more technological innovations in power saving and alternative models for regulating these small networks.

[1]Name anonymized.

Contents

Chapter 1

Introduction

There are over six billion subscribers of cellular networks [50]. These nearly ubiquitous networks impact the fundamental economics of subscribers' lives [53] and their relationships with friends and family [15]. This impact is particularly large among the disadvantaged; Agüero et al. [3] found that, for users at the "Bottom of the Pyramid", telecommunications displayed the usage patterns of a necessity; as total income falls the share spent on telecommunications *increases*, with some disadvantaged groups spending upwards of 50% of their income on communications. Unfortunately, hundreds of millions of people, primarily in rural areas, still lack cellular coverage. There are two primary reasons for this: economics and structure.

The economic argument is fairly simple; It is too expensive for incumbent providers to enter these remote and rural markets. The primary cost of a rural cellular installation is power; the GSM Association (GSMA) estimates that 95% of the people without cellular coverage in East Africa also lack grid power [36]. Similarly, the ITU estimates that 50% of the operating expense (OPEX) of a rural cellular network is power-related [51]. This is because many rural towers are powered with diesel generators, requiring constant trips to refill reservoirs, roads to support these trips, and fences to protect the valuable fuel and equipment.

For these reasons, powering cellular equipment with renewable energy sources could change the economics of building rural cellular infrastructure and bring coverage to many currently unserved parts of the world. Multiple vendors, including Range Networks [74], Vanu [91], and Altobridge [5] provide cellular equipment for this very purpose. Unfortunately, the equipment sold by these providers draw upwards of 70 watts of power. Though "low power" compared to traditional telecommunications hardware, powering a 70W BTS solely with renewable energy can be difficult, potentially requiring over fifteen 70-pound deep-cycle batteries per BTS site and increasing the total installation cost by up to fifteen thousand dollars [42].

The second problem is more subtle. As Galperin et al. [30] suggest, there is also a structural issue at work; the only organizations currently capable of deploying cellular networks are major telecommunications firms. This model of deployment is essentially *top-down*, with small communities having little power to provide their own solutions. Despite subsidies such as universal service obligations [83], the nation-scale firms' strong profit motives limit their investment in rural connectivity, leaving areas that are outside of existing infrastructure or have too sparse a user base without coverage. As an alternative model, Galperin points to microtelcos: small community, NGO-owned, or otherwise independent telecommunications firms that exist in some countries in Latin America.

In this work we attempt to build and deploy a system capable of bringing cellular connectivity to the remaining hundreds of millions of people without it by resolving these two core concerns: economics and structure. We do this by creating a novel way for reducing power draw in rural cellular networks, *virtual coverage*, and a bottom-up form of cellular networks, the *community cellular network*. We evaluate both of these innovations in the context of a real-world deployment of the technology in rural Papua, Indonesia

Virtual coverage is essentially a smart duty cycling mechanism. Instead of naively turning off equipment, the infrastructure turns itself off ("sleeps") when not in use. It "wakes" when a user takes an action, such as pressing a button on an autonomous radio called the Wake-Up Radio (WUR). It also "wakes" on incoming communications, such as SMS or voice calls directed to users on the network. When "woken", the network stays active for some period of time, enabling communications by subscribers during this period. The network goes back into "sleep" mode, reducing its power draw, after a pre-configured idle time, i.e., period in which no SMS or call is initiated or received.

Community cellular networks (CCNs) are small-scale, community-centric, *bottom-up* cellular networks. Owners and operators of CCNs are able to make all of the decisions about network operations and have complete freedom to design the network, its user interfaces, and services in a way that best serve their users and communities. This includes how calls are routed, what services are available, and billing and pricing policies. This customizability enables heterogeneous network design and local decision making, in contrast to monolithic national-level telco services and policy. CCN operators can re-envision the goals of a cellular network, optimizing their networks for population covered, community service, or even political action. From an economic perspective, researchers have found that small-scale local actors can have lower costs than traditional external actors by leveraging existing infrastructure [30] and local knowledge. Similarly, decentralizing these networks could potentially empower the local community, increasing opportunities and freedom [80].

We evaluate these innovations by implementing both in a small-scale cellular network and deploying it in the village of Desa in rural Papua, Indonesia, in partnership with two local NGOs. This network provides mechanisms for local operation including buying and provisioning SIM cards, buying and selling credits, checking credit levels, communicating with in-network friends via SMS or voice, and out of network contacts via SMS. It is jointly owned and operated by a local education-focused NGO and a for-profit wireless Internet service provider, providing a base of operations and the technical knowledge required for sustainable deployment. This network operates profitably, with our NGO partner charging for SIM cards as well as credits in the network.

We monitored the transactions in our network from February 11th, 2013 (the opening of the network for sales) until August 12th to measure its usage by the community. In that time, the network serviced 187 users; handled over 49,000 outbound (out of network) SMS, 42,000 inbound (into network) SMS, 24,000 local (in-network) SMS, 12,000 local calls and over 20,000 local service requests, including credit transfer, credit check, and others, some generating revenue for community members.

During this evaluation period, the infrastructure also entered the virtual coverage "sleep" mode from 23:00–06:00 every night while hydro power was unavailable. During these periods, we saw 1485 total communications, with 730 (49.2%) being initiated by and 755 (50.1%) received by people in Desa. The BTS itself was "woken" from sleep a total of 428 unique times, with 59 (14%) being caused by local users pressing the button and 369 (86%) by incoming SMS. Night service was used widely with 86 out of 170 (50.6%) of the subscribers making use of the network during the night period. Seventy six of these sent at least one unsolicited SMS, indicating that this participation was not only passive. Last, even though we were able to provide connectivity overnight, the network remained "asleep" for 87% of the night, reducing the night power draw by 56.6%.

We also evaluate the sustainability of the installation. With a model assuming shared power and network costs, complete financing of the rest of the installation at World Bank rates [95], and assuming no subscriber growth, the network would be paid off in five years as well as generate US$368.31 per month in profit for the operator. A model assuming no infrastructure sharing would still be sustainable, generating US$66.43 per month for the operator. We also conducted numerous interviews with users of the network. We find that it meets many of their needs, which revolve around communicating with family and friends abroad. Their primary wish is for outbound voice service and there is apprehension within the community about the network's use by minors.

Lastly, we compare the usage of the Desa network against a traditional firm located in sub-Saharan Africa across a series of graph-theoretic metrics to determine if the community in Desa used it in a more locally-focused way. These metrics include the basic inbound-outbound ratio, transitivity and the average clustering coefficient. The results demonstrate that the Desa network usage is significantly more locally-focused than the traditional telco, supporting our desire for local services and decision-making.

We believe these results demonstrate the value of community cellular networks and virtual coverage. The installation in Desa is sustainable, profitable, highly used by the local population, while still operating at low power and providing for night usage. We believe the core contributions of this thesis are:

- The design and implementation of *community cellular networks*: small-scale, locally operated, cellular networks;
- A taxonomy of the design space for operators of these networks;
- The design and implementation of *virtual coverage*, a system for reducing power draw in rural cellular networks; and
- A real-world evaluation of community cellular, virtual coverage, and the sustainability of these networks in a community in rural Papua, Indonesia
- A comparison of our Desa network against a traditional nation-scale telecommunication firm using graph theoretic community measurement metrics, showing that our network is more community focused.

This work is organized as follows: We begin with a discussion of related work in information and communications technology and development (ICTD), wireless, and power. We then define community cellular networks, including a taxonomy of the design space available to operators of these networks. We follow with a similar definition of virtual coverage including an implementation and lab evaluation showing the technique is technically viable. We then provide the context of our deployment in Desa, Papua, Indonesia and then the implementation of the network installed there. We follow with an evaluation of the installation including usage, power savings, sustainability, and a comparison with traditional telcos. We then discuss the risks to operators of community cellular networks. We follow with a discussion on related aspects of the work, such as local implications, generalizability, and evaluation limitations. We end with future work.

Chapter 2

Related Work

This research touches a broad number of disciplines. Here we describe the existing research related to cellular networks, international development, community and local cellular, power, and the context of our deployment.

2.1 Cellular Networks

2.1.1 Cellular Network Usage

Eagle et al. [26] investigated how calling patterns changed as users migrated among rural and urban areas. Blumenstock [12] used cellular logs to measure the connections between different nations. These analyses informed our expectation of user behavior. Heimerl et al. [46] researched how users in developed, developing, urban, and rural areas viewed and made use of their cellular infrastructure. He found that rural users had a better understanding of network properties, including coverage patterns. Similarly, Donner [24] and Sambasivan et al. [79] found that people in developing regions often expend considerable effort to achieve fundamental wants: communication or media consumption. This informs our design, implying that users would be willing to take actions to get network access.

2.1.2 Cellular Economics

Our work explicitly focuses on the economic models of cellular networks, an area of much research [31, 56]. Johansson et al. [54] proposed a method for estimating the cost of cellular installations, and mechanisms for optimizing coverage decisions based on the results. Our goal is to enable new mechanisms for deploying low-cost cellular equipment, allowing cellular operators to service rural parts of the world profitably.

2.1.3 Rural Cellular

A variety of researchers have explored and deployed custom cellular networks in the developing world [6, 43, 44, 96]. These works, as ours, build off prior research in economics and development advocating for these smaller-scale networks [30]. Rather than evaluating the user experience of the network, as in prior work, this paper focuses on the impact of virtual coverage on the power draw of a rural cellular network.

2.2 Development

2.2.1 Developing Region/Rural Networks

Network connectivity in developing and rural areas is an area of active research [14]. Wireless telecommunications are a common idiom [28, 85, 9], as the cost of deploying these networks is significantly less than traditional wired networks in areas with limited infrastructure. Researchers have investigated long-distance wireless [71, 17] and sustainability in these areas [86, 87]. Commercial companies now provide solutions for long-distance 802.11 networks [90].

Researchers have been looking into using WiFi for voice communications, primarily through mesh networks [11, 1, 29]. Hasan et al. found that these networks are inherently limited [40] and unlikely to provide a stable platform for voice communications. We focus instead on the design and use of cellular systems (and their existing handsets), and community cellular networks do leverage long-distance WiFi for backhaul or to connect different base stations.

2.2.2 Mobile for Development

With the incredible adoption of cellular systems, both industry and academia have investigated the impact of telephony on rural and developing populations. The International Telecommunications Union (ITU) has published aggregate usage statistics [50] as well as lower-level analyses of the cost of running a rural cellular site [51]. The GSM Association has done similar work [35].

Researchers have also analyzed cellular use. Jensen [53] found that the installation of cellular towers in Kerala benefited the users by reducing information asymmetries. Waverman et al. [93] and Deloitte [22] found other benefits in many other developing areas of the world. In general, researchers note that the cell phone has been an important technology in the developing world [15, 48]. These results motivate our desire to enable telecommunications.

Many services have been developed that utilize the mobile phone. Parikh investigated the use of mobile phones in supporting microfinance [69]. ODK [39] uses mobile phones for data gathering. Patel et al. [70] developed a voice-based tool for farmers in rural India. M-PESA [63] is one of many banking solutions and motivated our credit transfer service.

2.3 Community Cellular

2.3.1 Small-scale and Community Telephony

History is scattered with stories of small-scale telephony networks. Early in the history of the United States, there were numerous cooperatively owned networks [27]. Though times have changed, there is a still a robust market for small-scale operators [20]. The dominant model for these carriers is simply to provide roaming service to the major US telcos, and thus there is no opportunity for customizing their networks for specific populations. Similarly, these networks are usually hundreds of nodes, unlike CCNs, which are smaller.

In most countries, a few carriers tend to dominate communications. Researchers have suggested alternative structures. Galperin et al. argues for "microtelcos" for Latin America [30] including the coproduction aspects. Similarly, Zuckerman argues for decentralizing the mobile infrastructure [98] using WiFi networks. We believe community cellular to be the first actual instantiation of these designs that are sustainable and replicable; by leveraging the existing market and install base for GSM phones and providing the basic services people need, these networks can be sustainable [44].

2.3.2 Infrastructure Design

Clark et al. [19] and Clark [18] investigated the design choices made in the Internet and how they can encourage openness and choice. Similarly, Shenker [81] investigated how pricing impacts the use of Internet service. Coproduction [60, 55] is the idea of delivering public services like telephony with community participation and involvement. This idea permeates any community cellular deployment. For example, the Papua network utilizes existing power and network infrastructure to reduce its cost.

2.4 Virtual Coverage

2.4.1 Existing Systems

Seeing the need for rural cellular communications, commercial providers have begun to develop products optimized for these areas. Both Vanu [91] and Range Networks [74] have developed "low power" 2G cellular equipment capable of running on entirely renewable energy sources and using only around 90 watts of power. Altobridge's [5] lite-site product varies its capacity in order to save power. As all of these products provide constant network coverage (rather than virtual coverage), they cannot reduce their power draw below 90 watts.

The OpenBTS project [67] is a full-featured GSM base station using software-defined radios. It bridges GSM handsets to voice-over-IP systems, enabling lower-cost, lower-power cellular equipment. The OsmocomBB project [68] is an open-source GSM handset, capable of interfacing with any 2G GSM station. Our research is built on these two pieces of technology that enable us to modify the GSM standard and implement a low-power GSM telephony solution.

2.4.2 Saving Power

As the world goes "green", reducing power consumption has become a critical goal of system designers [8]. Researchers (and system implementers) have focused on many different mechanisms for saving power, including disabling specific pieces of hardware. Gobriel et al. [33] investigated how to save power by utilizing the "idle time" in data center networks. Zheng et al. [97] used asynchronous wake-up mechanisms to save power in ad-hoc wireless networks. Wake-on-Wireless-LAN [62, 82] enables behavior similar to wake-on-LAN [64] across wireless networks. We similarly save power by utilizing idle time and a wake-up mechanism in cellular networks. However, the potential power savings in rural GSM cellular networks far outweigh smaller wireless deployments, as the range and power consumption of cellular networks are orders of magnitude larger.

Lin et al. investigated multihop cellular networks [61] to broaden the range of GSM cellular towers. In their solution, each handset could potentially act as a router for other handset's messages, directing them to a centralized base station. Others expanded this idea, investigating how one might properly incentivize users to share their cellular connections [52]. Our work utilizes a different mechanism, virtual coverage, to achieve the goal of increased effective network range. These two designs are not mutually exclusive. A solution using both could have a dramatic impact on rural telephony.

A few groups have explicitly investigated reducing power consumption in cellular networks. Bhaumik et al. [10] proposed varying the size of cells; saving power by turning a subset of the cellular towers off during times of low load. Peng et al. [72] took these ideas even farther, demonstrating significant over-provisioning in cellular networks. Unfortunately, their proposed method for saving power is not possible in most rural areas as there is often just one tower providing service.

2.4.3 Sensor Networks

Efficient use of power is of critical importance in wireless sensor networks due to battery-based operation. The core technique we use to save power in cellular networks is very similar to one developed by Gu et al. [38] for wireless sensor networks. Other researchers have explored similar designs [23]. As in our system, the nodes sit idle until communication is needed, and wake each other up using large radio bursts that are distinguishable from noise through some mechanism. Similarly, other researchers also created standalone devices for creating wake-up signals [25]. This is unsurprising, as these technologies inspired our design. The key differences are in scale, intent, and mechanisms.

2.5 Context

Our work in Papua is informed by the wide variety of anthropological work done in the area. O'Brien et al. [66] is a classic work describing the indigenous people near our area of research. Ploeg investigated the highlands in a more recent era [73]. Other research has explored local understanding of HIV/AIDS [16]. The Indonesian government also produces a census every ten years [84], which is also analyzed by other parties [13].

Chapter 3

Community Cellular Networks

Community cellular networks are fundamentally an inversion of the traditional model of rural cellular. Instead of *top-down* nation-scale networks deployed from major population centers into rural areas, CCNs are built *bottom-up* by rural communities themselves with their goals and skills in mind. In this chapter, we describe what defines a community cellular network and what designs operators of CCNs can make.

3.1 Defining Community Cellular Networks

In this section, we define community cellular networks by differentiating them from these more traditional cellular networks across three core metrics: operation, scale, and goals.

3.1.1 Operation

Nearly all cellular networks are run by large (often multinational) corporations (e.g., AT&T) or by the state (e.g., China Telecom). As a result, a handful of nation-wide carriers typically dominate cellular communications. A community cellular network, in contrast, is operated by an entity within a community, such as a local entrepreneur, NGO, or cooperative organization. This allows design to be responsive to community needs and desires.

3.1.2 Scale

Major cellular networks operate across entire countries, providing coverage in large population centers and covering millions of people. One reason for this is economies of scale; networks can amortize the cost of expensive related infrastructure like the mobile switching center. This means that all users across the country get a similar experience; there's little chance to customize the network for individual communities.

Community cellular networks are much smaller scale. They are often just one tower in one community, as are both of our case studies. Though it is possible for an owner/operator to be involved with more communities, we do not expect community networks to encompass more than a dozen nodes. This property enforces the above property of local operation; a large network will invariably be operated by entities further and further from the community. It also allows for more customization, and the network itself will be smaller and need to generalize less.

3.1.3 Goals

Privately owned cellular networks tend to work towards maximizing their individual profit. State-owned firms work towards the goals of the state. Community cellular, by virtue of enabling new models of ownership by different types of organizations, enables new goals for the network.

Though financial sustainability is likely to continue to be a concern for any entrepreneur, a local network is in a prime position to support *local* goals and *community* growth. For example, a health-care NGO may deploy a community cellular network in support of an immunization project and provide tools on the network for that service. A cooperative may limit the network to only local calls and local people, explicitly forbidding out-of-network connections. Any community cellular network must support local goals and community growth as one of its goals.

3.1.4 Conclusion

In conclusion, a community cellular network is:

- operated by a person or organization deeply embedded within the serviced community;
- small-scale, with at most a dozen nodes; and
- trying to achieve goals that benefit the community.

3.2 Designing Community Cellular Networks

There is little literature on the design of core network infrastructure. Usually the decisions made about these technologies are far removed from end users. CCNs change this: suddenly, high-level structural decisions (e.g., licensing) are being made by small-scale operators who live and work in the communities being served. The options available to CCN operators are wide and potentially very confusing. To remedy this, we provide a taxonomy of the many design variables available to *operators* of CCNs.

3.2.1 Licensing

Cellular networks operate in licensed spectrum bands, and buying the necessary licenses is often expensive. For instance, a recent auction of portions of the 3G bands in India netted over 10 billion USD [34]. These high prices stem from the fact that spectrum is considered a valuable public resource and is priced accordingly. Although large firms can pay these high licensing costs, smaller community cellular installations typically can not. A network should operate legally when it can, but operators can decide to operate under a range of licensing regimes, ranging from fully licensed to “pirate” (or “extralegal” [21]) and including the broad grey area, including tacit approval and operation in unregulated areas, in between.

- **Full License** – Total legal operation within the country. This usually includes registration with the spectrum coordinator and cooperation with requests for data from law enforcement (“lawful intercept”). An operator may also benefit from regulations requiring interconnect with existing carriers. The primary benefit of this model is scale; as your network grows it is less likely to run afoul of regulators. Unfortunately obtaining a license can be prohibitively expensive and may invite business interference from competitors. In some cases, it may be possible to use the license of another carrier with permission, either via rental or due to regulatory pressure for “universal access”.

- **Pirate** – Unlicensed operation. Though illegal, this model can work in practice in rural areas as there's little revenue being generated and thus little incentive for regulatory enforcement; rural operators are also unlikely to interfere with the operations of license holders. Operating as a pirate avoids licensing costs and reduces regulatory requirements. An operator can also build a unique network without their vision hampered by regulatory regimes designed for traditional networks. Unfortunately, as the network grows the chance of regulatory enforcement increases.
- **Grey Licensing** – A middle ground between a full license and a pirate station. In this case, regulators may give tacit approval to the network with certain understandings. For example, that it will cooperate with law enforcement and not attempt to operate where other carriers and license holders operate. This is likely where all spectrum has been sold but isn't being used in rural areas.

3.2.2 Access

Access control in GSM cellular networks relies on subscriber identity module (SIM) cards, which uniquely identify a SIM card. A network can decide whether to accept or reject a handset, based on its SIM, when it attempts to attach to a tower. In terms of access, there are two types of networks: open and closed.

- **"Open" networks** – Customers can easily buy/use a SIM card that grants access. This is the common model in most of the world. Users simply buy a SIM from a vendor and are immediately granted network access. Similarly, CCNs can make accepted SIM cards easily available, or they can even be configured to allow *any* SIM to connect.
- **"Closed" networks** – Only specific SIMs, usually defined by a white list, are allowed access to the network. This model is seen in smaller, so-called "private networks" like oil rigs or cruise liners. Access is limited to reduce network contention or drive up prices. It is also possible to have a functionally closed network: Burma at one point had expensive SIMs (around US$3000), effectively creating a private network for the country's elite. In such situations, broadening access would require lowering prices, which may not be in the best interest of the carrier, especially in monopoly situations.

The trade-offs available for operators are primarily about profits and control. The open model brings in more customers and increases total communications at the cost of less control over how people make use of the network. The closed model lets operators closely monitor and control the network resources, reducing network contention. It also lets operators choose their users, perhaps to achieve political, financial, or social goals.

3.2.3 Naming/Numbering

Though people traditionally assume a one-to-one relationship between numbers and phones, or numbers and SIM cards, network operators have wide latitude in how numbers are assigned among their customers. In GSM networks, phone numbers are surprisingly flexible. Though people traditionally associate them with individual SIM cards, this is not actually the case. Instead, each SIM has an IMSI which is *mapped* to a phone number by the network. This flexibility isn't used very often; only for number portability.

Numbering directly impacts the user experience. CCN operators must balance ease of use, cost, and functionality when choosing their number schemes. Major telcos have optimized for ease of use by packaging all SIMs with numbers. CCNs, having more involvement with their users and less capital, may choose differently. The flexibility afforded by CCNs enables a number of design choices including number scope, sharing, and dynamism.

Scope

Number assignment can be done differently in each network. The network can assign users a globally addressable number, or a number only useful within the community network.

- **Global** – Numbers are addressable by people outside the network. Globally reachable numbers must be assigned by a licensing agency (e.g., a national telecom regulator) and not overlap with any existing allocated numbers.
- **Local** – Numbers are addressable only by people inside the network. Local numbers can be almost anything and can overlap with phone numbers on the global telephone network since they aren't routable from anywhere beyond the local network.

This design decision is about flexibility in the network. Networks that use local numbers have a variety of usability options available to them. They can use shorter, more easily memorable numbers (e.g., 5555), greatly simplifying the use of the network. Similarly, for small networks, numbers could correspond to individual users' names (e.g., 262 for Bob) or other vanity numbering schemes. However, these numbers would not be available globally, potentially a confusing user interface with the wider telephony world.

These problems are nearly identical to issues seen in private branch exchanges (PBXs). PBXs are systems used primarily for internal telephone networks in businesses. Calls are routed locally based on local numbers, avoiding the broader telephone network unless specifically desired. PBXs are a well understood, mature technology with a rich literature describing their design, implementation and use [65].

Sharing

Although most phone networks assign one number per user, it is also possible to *share* numbers, for instance assigning multiple local numbers to share one global number.

- **Individual** – Each user has their own unique phone number. This is the traditional model. It is usually expensive, as each number must be purchased (or provided in some other way) from a telecommunications firm with an existing block of numbers or from a regulatory body. However, it provides a simple experience for users.
- **Shared** – Users share a number. This is conceptually similar to Private Branch Exchanges (PBXs) in the voice over IP (VoIP) world. Individual users have a local number (extension) and share a global number for outbound communications. Incoming communications are multiplexed to the receiver through some mechanism. Two examples would be an interactive voice response (the caller inputting an extension) or network address translation (NAT), directing the call to the user who made the last outbound call. This allows an operator to purchase just one number for a group of users, rather than many, lowering cost.

This design decision is primarily about weighing cost versus usability. Sharing numbers is a great mechanism for reducing a large operational cost: number rental. However, the multiplexing system can be difficult for users to understand.

Dynamism

Numbers are often assumed to be static, tied to individual phones. However, it is actually possible to change the mapping between users and numbers dynamically, based on changes in the network or any other criteria.

- **Static** – Numbers are static and do not change. They always route to the same user. Traditional phone networks use this model.
- **Dynamic** – Numbers change for individual users. This can be in the case of sharing a single number among users; for instance a single number can route to the *current* triage nurse rather than a particular one. Similarly, users themselves may change their numbers if the operator lets them. This could be done to lose old contacts, reinvent themselves, or avoid spam.

This design decision is about enabling new functions at the cost of usability. Allowing users to change their numbers is a powerful tool for customization, but may confuse people familiar with traditional telecommunication idioms.

3.2.4 Pricing

Pricing is a variable networks have long used to influence user behavior [81]. CCN operators have complete freedom to price network services as they see fit. A deeper discussion of these choices is beyond the scope of this thesis, but at a high level pricing models fall into two broad camps:

- **Tiered** – Different prices for different services. For example, lower prices for "local" numbers and higher ones for "long distance" or "roaming". Though this price difference often reflects actual economic realities (such as "in-network" calls being cheaper), sometimes operators want to reduce load at certain times [49] ("peak" airtime).
- **Flat** – One flat price for most or all services. This also includes "all you can eat" plans. This has become more common in large networks, primarily for simplicity. Some countries regulate roaming charges, enforcing flat pricing for cross-country communications.

All local communications and services have near zero marginal cost (as they utilize no outside network resources) but do suffer from scarcity since in-network bandwidth is limited. Outbound communications likely incur costs for the operator. These economic forces must be balanced against the operator's desire for service use.

3.2.5 Payment

If the network requires users to pay for access, they must also design a payment system that fits their community. The primary decision is how to bill users: pre-paid or post-paid?

- **Pre-paid** – Pay for access and each communication before use. Pre-paid systems are common in most of the developing world. Users first buy a SIM from a vendor. Users then buy "credits" for their account, and spend them to make communications. In this case, both access to the network (buying SIMs) and actual use (making communications) are pre-paid.
- **Post-paid** – Monthly billing of communications. Post-paid is common in the developed world, and is often a set monthly fee for unlimited access and use. Users often get a free SIM from the vendor and sign a contract to pay a certain amount a month for access to services.

This design choice is usually influenced by the economics of the targeted population; people with little disposable income benefit from the smaller purchase denominations enabled with pre-paid. A consequence of choosing a pre-paid system is that it requires a mechanism for credit transfer and sales. This is often a specialized "credit sales" person in the community. This person will be a stakeholder of the network, but will require compensation for their time. Networks using a post-paid model require less involved sales staff, as collection happens just once a month and payments are made in bulk. A post-paid network also requires a certain amount of "trust", as legal options for collection may be costly and difficult. There's also a middle ground available; for instance access being pre-paid (buying a SIM without a contract) but use being post-paid (a monthly fee).

3.2.6 Services

Lastly, a network can choose what services to offer their users. Traditional telcos offer services such as SMS, voice, and data, as well as nation-wide services such as American Idol voting or pay-for services like SMS jokes. Researchers have explored ways to use these services to benefit people, building systems like an automated voice mailing list [70], job search systems [75], or a singing competition [58].

CCNs provide the perfect platform for creating network-based services tailored for specific communities. Though the specific design constraints of each service is outside the scope of this work, we note that these applications have access to voice (including local language speech recognition [7]), SMS, data, and a user's location and usage history. A vast number of applications are possible. Any one of disaster warnings, crop prices, entertainment, photo sharing, or community mailing lists can be customized for the individual locales being served by each CCN. This ability to customize network applications to each specific community, rather than nation-wide as with traditional telcos, is an extremely powerful component of community cellular.

Chapter 4

Virtual Coverage

To reduce the power draw of rural cellular installations, we developed *virtual coverage*. Virtual coverage is, at its core, smarter duty cycling. It powers down the most power-hungry portion of individual cellular towers, the power amplifier (see Table 4.1), when the tower is not in use. At the same time, the tower is still operational enough to listen for user broadcasts. When the network is needed, as signaled by a user initiating or receiving a call, the power amplifier is enabled and the tower is available for communications. This design allows us to save a large amount of power in the largest networks on Earth while still providing consistent coverage at all times.

Specifically, individual cellular towers are put into a low power “idle” mode during periods of zero utilization. Although powering down a tower is conceptually simple, waking one is much harder; when a user wishes to make an outbound call, they must wake the network by sending a burst from either a modified GSM phone (our Wake-up Phone) or via a small, low-cost, push-button transmitter (our Wake-up Radio, functionally similar to a garage-door opener). Mobile-terminated calls require no changes to user behavior; the tower simply turns on and holds the call until the requested user connects.

This “idle” mode system comprises virtual coverage. In this chapter, we detail a prototype technical implementation of a base station, handset and autonomous radio supporting virtual coverage. We conduct some micro-benchmarks to demonstrate that this system will function with existing handsets and save power in some example theoretical networks.

4.1 System Implementation

Enabling virtual coverage requires a holistic rethinking of the base station (BTS) itself. First, the BTS must be modified to enable programmatic control of the power amplifier, which consumes upwards of 65%

Component	**Draw (W)**	**% (10W)**	**% (50W)**
Computer	12W	17.4%	7.8%
Radio	12W	17.4%	7.8%
10W Amp	45W	65.2%	
50W Amp	130W		84.4%

Table 4.1: Power draw of the components of our Range Networks 5150 BTS.

of the BTS power (see Table 4.1). This allows us to enter an "idle" mode in which the total power draw is reduced at the cost of limited connectivity while idle.

Second, we implement a mechanism for allowing users to wake the cellular tower remotely and promptly, thus enabling coverage on demand. We implement two models of virtual coverage wake up: 1) a software-only handset modification to send special wake-up bursts, and 2) a custom autonomous low-cost radio that sends the same message, allowing the system to work with existing, unmodified handsets. After detection of this burst, the network exits the idle state and resumes normal operation.

4.1.1 Enabling Low-Power Modes in Cellular Infrastructure

Virtual coverage requires the base station to have a low-power mode when the network is not in use. There are two core changes needed create a low-power mode for a GSM base station. First, the hardware must be modified to provide a mechanism for programmatic control of the power amplifier, the primary power draw. Second, the software must actually cease broadcasting during idle times while still listening to detect wake-up bursts.

Hardware Figure 4.1 shows the internals of our revised Range Networks 5150 cellular base station. The key pieces of equipment are the radio, computer, duplexer, and power amplifier (PA). We added a USB-controlled high-current switch and connected it directly to the power amplifier, allowing us to control the PA's status via serial commands from the computer. When the BTS enters idle mode, the PA is turned off.

Software There are two key software modifications. First, we implement the idle mode and drop all transmissions (including the beacon) while the power amplifier is off. Second, we implement a mechanism for the BTS to receive wake-up signals from user radios.

We implemented idle mode with a service that sends messages to the switch controlling the power amplifier. This daemon, which has access to the GSM and switch state, controls entering and exiting idle mode. Instead of naively returning to idle when all calls have terminated, we use a number of heuristics to improve the user experience. First, we require that the network be active for a minimum of 90 seconds, approximately double what we found to be the worst-case time necessary for a handset to connect to and communicate with the tower (i.e., *camping*) (Table 4.2). This ensures that all handsets waiting to camp will have ample time to do so should a tower return from idle mode. Second, the BTS only transitions to idle if there has been no cellular traffic for 30 seconds. This enables serialized actions, e.g., redialing a dropped call.

Originally, we had hoped to provide a "low coverage mode" (i.e., signal transmission without amplification), where the BTS could still provide coverage for people physically near the radio. In testing, however, we discovered that the radio *must* be disabled when in idle, even to the exclusion of transmitting with the amplifier in pass-through mode. If the BTS broadcasts (even at low power), handsets nearby will attempt to camp (a process near-all handsets perform automatically and periodically). As our burst-detection is a simple power measurement, this legitimate network traffic would be indistinguishable from wake-up bursts.

The BTS provides two primary functions: incoming (mobile-terminated) and outgoing (mobile-originated) calls. Mobile-originated calls are simple. The tower must be in active mode, as only a camped handset can initiate a call. Mobile-terminated calls are more complicated. If the tower is "active" when it receives a call, the call is simply routed to the appropriate handset. If the tower is "idle", the caller either leaves a voicemail (and the tower remains "idle") or they are put on hold and the tower immediately wakes and waits for the handset to camp. According to our measurements (Table 4.2), this can be up to 40 seconds (with most being

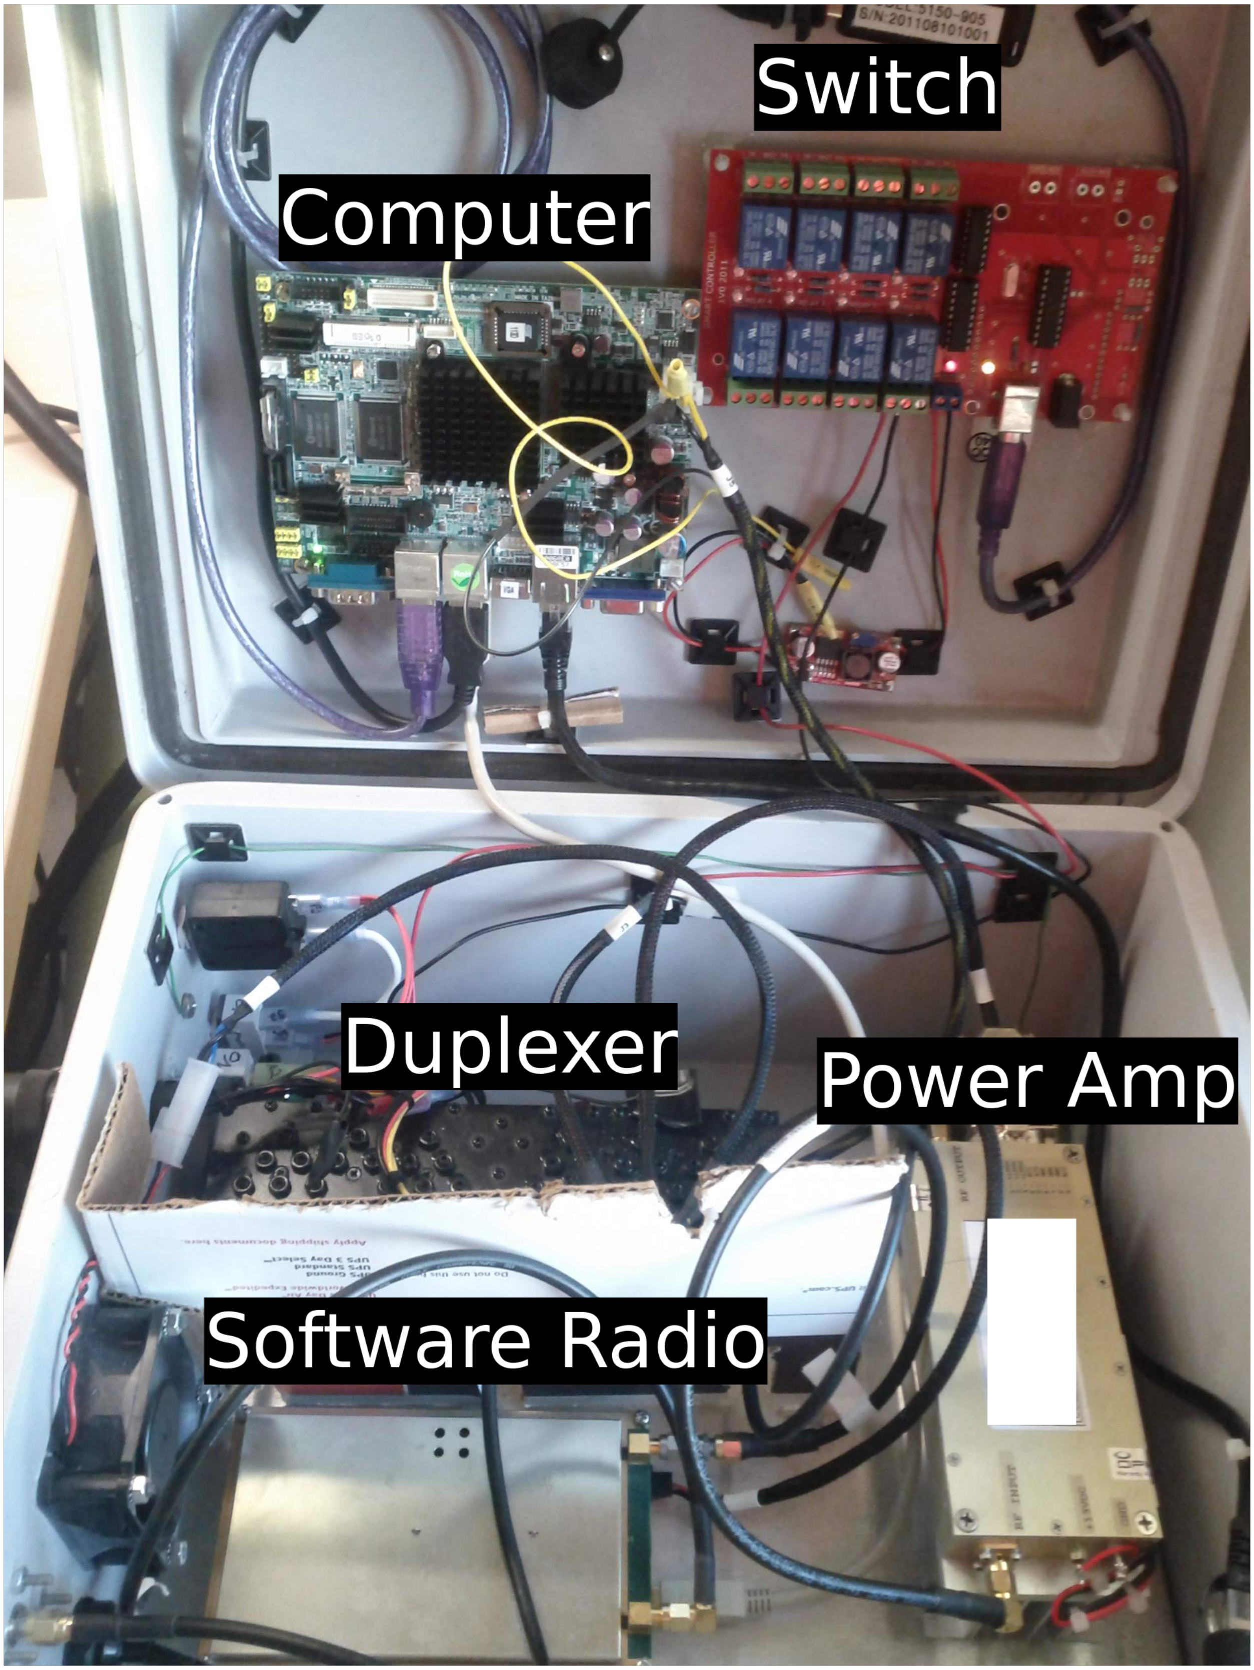

Figure 4.1: Our Range Networks GSM BTS.

under 25 seconds). When the handset eventually camps, the callee is immediately connected to the caller by bridging to the held call or initiating a new call if they hung up.

The basic mechanism for detecting bursts is implemented in the transceiver of the radio. If the radio is in idle mode, any high enough power burst on the tower's Absolute Radio Frequency Channel Number (ARFCN, basically just the frequency the tower listens and transmits on) will cause a message to be sent to the daemon, waking the system. The power required depends on the current noise level as determined by the transceiver. This technique is similar to ones used in sensor networks [38].

As OpenBTS utilizes voice-over-IP (VoIP) as its interconnect, there are no changes required to any other network services. Were we to interconnect using more traditional protocols (i.e., SS7/MAP) the name database (HLR) would have to allow longer registrations from users on virtual coverage enabled BTSs. This is the only change required for inter-operation.

With this system, we are able to provide on-demand voice services for rural networks at the cost of increased call-connection latency. SMS and data traffic are assumed to sync during periods of active voice traffic.

4.1.2 Waking Up in Virtual Coverage

Virtual coverage is not just a change in the cellular tower; it also requires a device capable of sending a "wake up" message. As mentioned in the previous section, we implemented two mechanisms for sending this message: from a handset or via an autonomous radio.

Cellular Handset

We have implemented our base station wake-up mechanism using an osmocomBB compatible mobile phone. We call this the *Wake-up Phone (WUP)*. OsmocomBB [68] is an open-source GSM baseband software implementation, which simplifies changes to the GSM protocol. However, every GSM handset should be able to send a wake-up burst with a software change from the manufacturer. The mechanism for waking up the BTS is sending a burst packet on the BTS's ARFCN. The BTS, though not transmitting, receives this message and exits the idle state, allowing the handset to camp.

Each BTS broadcasts its ARFCN number (as well as the ARFCNs of similar nearby towers) on the beacon channel, which details the exact frequency used to communicate with the BTS. A handset periodically scans the network for towers to camp on, and gathers these numbers. In our system, the handset stores these numbers when the network idles, and then uses them to send "wake up" messages (as above) during periods without network availability.

Mobile-Originated (MO) Call In order to initiate a call, the Wake-up Phone will transmit "wake up" bursts on a selected set of ARFCNs. These ARFCNs are either a list of previously detected base stations or a static configured list. The "wake up" packets are random packets that are transmitted on the selected ARFCN. After transmission, handset scans for a tower broadcasting on the ARFCN just awoken, instead of scanning the whole cellular band (as in normal cell selection). If discovered, the handset camps to this tower and the user is able to communicate.

If a WUP is unsuccessful in camping to the recently awakened base station, the handset will proceed to the next ARFCN in its list, if any, and perform similar operations. This mechanism repeats until the handset is successfully camped or it runs out of available ARFCNs. At this point it will default back to the standard GSM protocol, which scans the entire band looking for available towers.

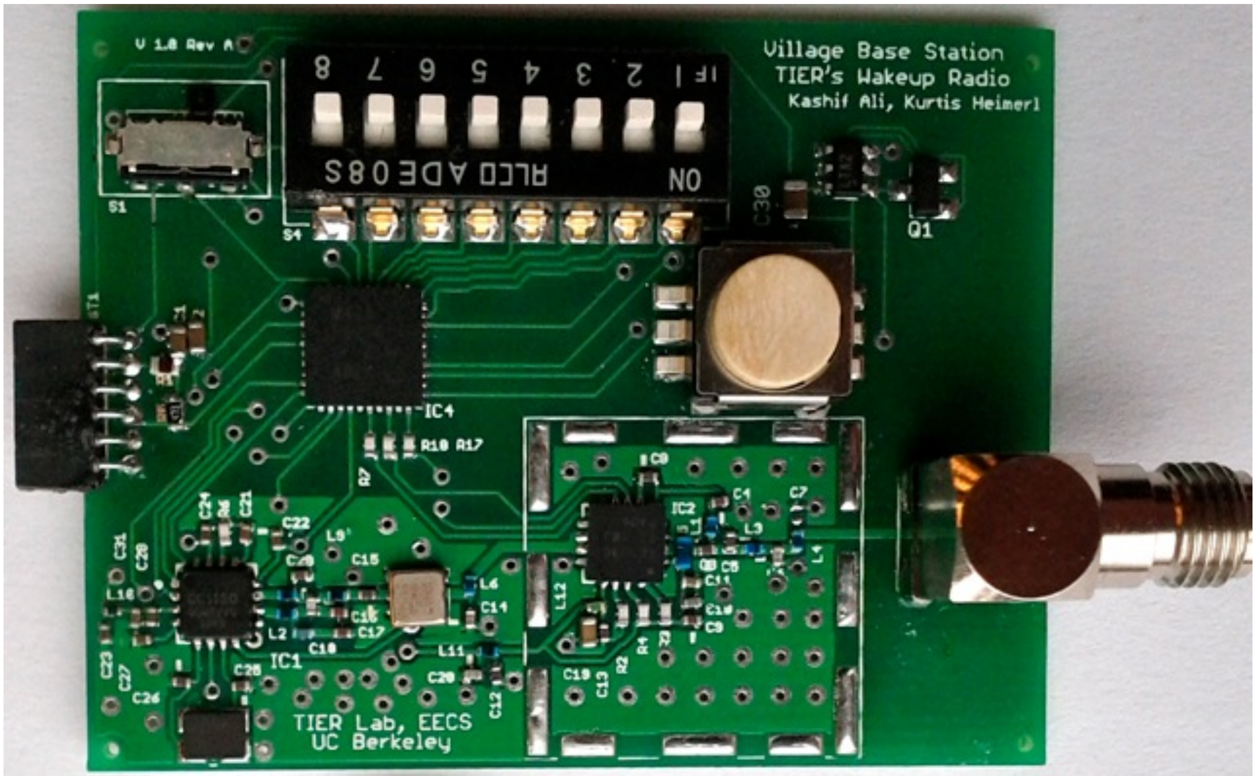

Figure 4.2: Prototype implementation of WUR.

Mobile-Terminated (MT) Call As stated above, when the BTS receives a mobile-terminated call it immediately exits idle mode and waits for the handset to camp. The WUP scans the stored ARFCNs much more frequently (10:1), reducing the average time to camp. However, this does not affect the worst case analysis, which is 7 seconds. When found, the phone camps and the call is connected.

Wake-up Radio

We have also designed and implemented a system to wake-up our BTS, the *Wake-up Radio (WUR)*, nicknamed the *garage door opener*. The WUR transmits wake-up bursts, similar to our modified GSM handset, on a specific ARFCN.

The radio is designed to be both cheap and low power. The primary user interface is just a single button. When pressed, this button triggers a burst on the configured ARFCN. The radio produces a signal at approximately 500mW, the minimum power required for a handset in the GSM standard. The WUR uses an on-board battery pack, but provides interfaces for other power sources (e.g., solar) as well. The WUR, with two AAA batteries is capable of 5000+ bursts. The WUR can be configured to produce different ARFCNs with dip switches.

The WUR is only needed for mobile-originated communications. For mobile-terminated communications, the tower simply wakes and waits for the recipient's handset to camp.

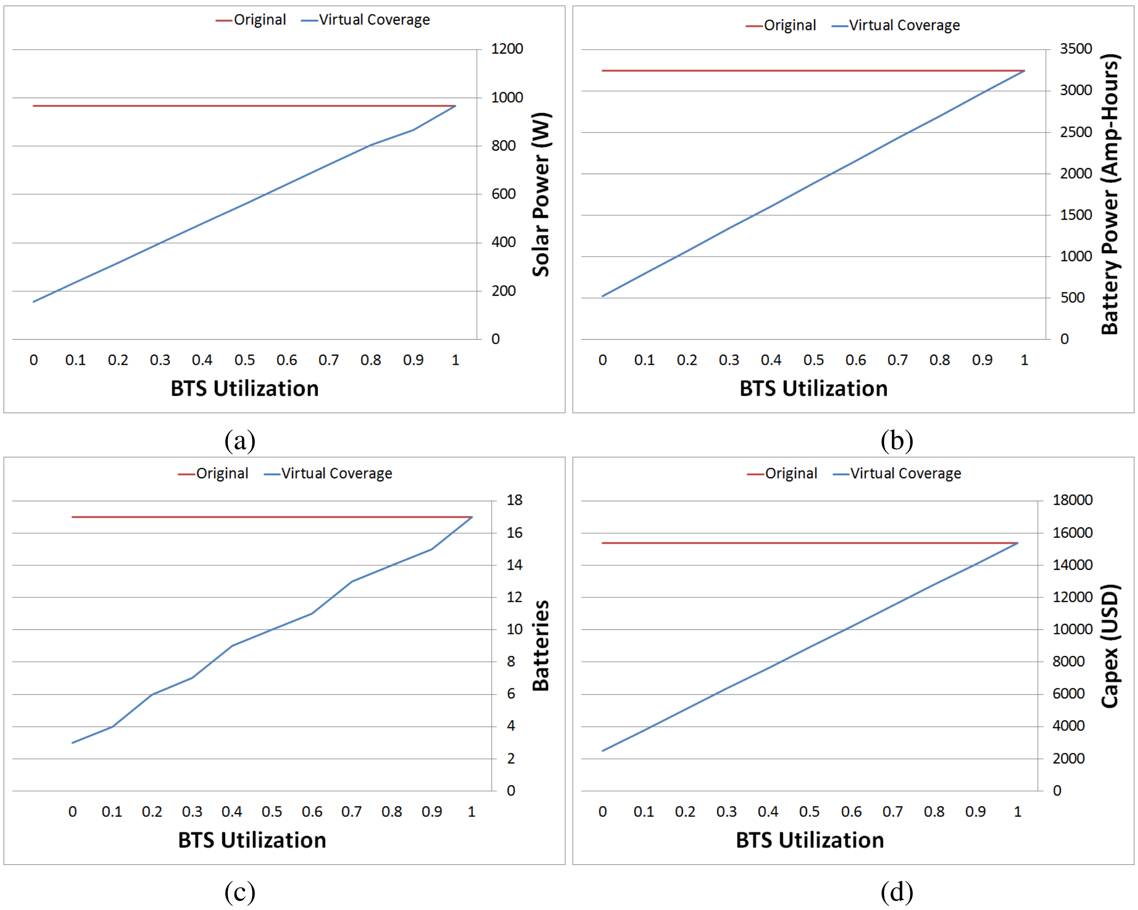

Figure 4.3: Solar Power (a), Battery Power (b), Individual Batteries (c), and total spending (d) required to operate a virtual coverage tower at different levels of utilization for a week in an area with 5 hours of sun, compared to a traditional ("Original") BTS. Note that 100% utilization is equivalent to a traditional tower.

4.2 Technical Evaluation

4.2.1 BTS Power Savings

We begin by evaluating the performance of a single modified Range Networks 5150 BTS. This unit has two power amplifiers available: 10 Watt and 50 Watt. The 10W unit supports two channels and 14 concurrent calls. It is commonly used for low-density areas. The 50W unit is designed for denser areas, produces up to five channels, and is capable of providing 35 concurrent calls. Both cover up to 35 kilometers, depending on geography. Areas with buildings or dense foliage will have worse signal propagation. Table 4.1 shows the relative power draw for each component of the BTS. As expected, the amplifier dominates the overall power draw, consuming 65% of the power for a 10W unit and 84% for a 50W.

In our system, we added a USB-controlled switch to programmatically control the power amplifier. This switch draws negligible power (less than 1W). We also saw no change in the power draw of the computer, as expected with the BTS handling no calls in "idle" mode. As such, we are able to reduce the overall power draw of our BTS by over 65%.

Model	Time (Avg)	Time (Max)
WUP (2G) (MT)	2s	7s
WUP (2G) (MO)	2s	2s
HTC Dream (3G)	12.1s	41.8s
Samsung Nexus S (3G)	23.6s	37.6s
Nokia 1202 (2G)	10.8s	14.6s

Table 4.2: Measurements on how long a handset has to wait, on average, to camp to a specific tower. This is the additional connect time if the network is idle.

4.2.2 Handsets

Wake-up Radio With the wake-up radio (WUR), the user has access to a device tuned to the particular frequency of their local cellular tower. Depending on local cost and logistic constraints, this device can be either widely deployed as an attachment to each local user's individual cellular phone, or singly deployed at some central location as a "phone booth". The user presses the button, sending the wake-up message to the tower, taking the station out of idle mode. When the tower wakes, it broadcasts a beacon signaling its location and ownership.

Traditional GSM handsets periodically scan the airwaves looking for beacons. As such, a user's handset will eventually camp to the newly awoken BTS. We measured the time to camp, after waiting to ensure the initial network search failed, for three different handsets: the Samsung Nexus S (Android), the HTC Dream (Android), and the Nokia 1202 (Symbian S30) over thirteen trials. The results are shown in Table 4.2. The first two phones are quad-band 3G phones, meaning that they scan a wider band than a dual-band 2G-only feature phone (e.g., our Nokia 1202) commonly used in developing regions.

Users must wait for their handsets to camp in order to communicate using the network. Our results mean that using the WUR increases the setup time for all calls by a maximum of approximately 40 seconds. The average wait measured is less than 25 seconds. Both the maximum and average time to camp are highly dependent on the specific phone used. Though this potential wait is a nontrivial amount of time, we believe this is acceptable to rural users who have limited alternatives for communication.

Wake-up Phone In the embedded solution, the WUP is able to camp on the BTS almost immediately after sending the wake-up burst. Since the same hardware device both delivers the burst and camps to the tower, the timing of the two tasks can be optimized for minimal delay. We measured the amount of time needed to register with the BTS [1] and found that, in all cases, it took exactly two seconds from wake-up burst to "camped normally". This is shown in Table 4.2. The practical impact of this result is that, with a virtual coverage network, every mobile-originated communication takes two seconds longer to set up when the BTS starts idle.

For mobile-terminated calls, the handset does not generate the wake-up burst and must instead listen for a cellular beacon. Fortunately, the handset still knows what cellular towers are in the area. Instead of periodically scanning the entire band (as in standard GSM), we can scan just the beacons that are present in the area. This again takes no more than two seconds. Other operating system functions occasionally delay this scan, causing a maximum wait time of seven seconds.

[1] GSM state $A1\_TRYING\_RPLMN$ to $C3\_CAMPED\_NORMALLY$

4.2.3 Deployment Example

To begin to understand the impact of virtual coverage in a real-world situation, we calculate the approximate amount of infrastructure (solar panels and batteries) required to support a 50W (drawing 155W) cellular station year-round using only solar power.

We first frame our “real-world situation”: Providing winter-time network coverage in Pakistan. During the winter the country receives 5 hours of usable sun [32], as solar panels deliver minimal power when the sun is on the horizon. Using this, we are able to calculate the amount of power optimally tiled solar panels generate. We assume an operating temperature of 40 degrees Fahrenheit and 24V batteries. Batteries are priced at 442 USD for 200 Amp-Hours and solar panels are priced at 1.07 USD per Watt. Lastly, the BTS draws 155W at full power (3720 Watt-hours/day) and 25W at idle (600 Watt-hours/day). As there is often inclement weather, we calculate the requirements for powering the station over a week without any power generation. The results are shown in Figure 4.3.

The actual impact of virtual coverage is large; a completely idle tower requires one sixth of the batteries, solar panels, and total infrastructure cost of a traditional tower. As expected, these variables scale linearly with increased utilization. Related work [42] demonstrates that utilization scales sub-linearly with respect to total calls (and thus users), meaning that the price of infrastructure required to support a virtual coverage tower scales sub-linearly with the total number of calls handled. Contrast this with a traditional cellular tower that must install the same amount of solar panels and batteries regardless of the number of calls and users serviced.

We wish to note that these costs are not only monetary. A single traditional BTS wanting week-long backup requires a kilowatt of solar panels and seventeen deep-cycle batteries (each weighing 68 pounds!). This equipment will be hiked into rural areas, an enormous load. Compare that with a virtual coverage station at 20% utilization. There, just 300W of solar panels and 6 batteries must go up the hill. Lastly, virtual coverage also allows for incremental growth; as an area moves from 20% to 30% utilization, new batteries and panels can be installed.

Chapter 5

Deployment

In mid-2012, a wireless Internet service provider in rural Papua, Indonesia contacted us and asked us to visit and install a low-power GSM base station (BTS). This was the beginning of our journey into Papua. Note that all names, including villages, people, and organizations, are anonymized. We first describe the community we were targeting and then the design choices made about the network by ourselves and the community.

5.1 Context

5.1.1 Papua

"Papua" (Figure 5.1) refers to the western half of New Guinea (and nearby islands) in the Southwest Pacific. Claimed by the Dutch since the 1800s, it has been a province of Indonesia since being officially annexed in 1969. The 19th and 20th centuries saw turbulent transitions from a pre-colonial situation of small polities with shifting inter-regional relations, to the Dutch project of colonial state-building, and finally to the current context of incorporation into Indonesia [77]. Often portrayed as "isolated", developments in Papua are historically connected to maneuvers at centers of global power such as the UN and the US [78]. Major international resource companies are active in Papua, notably at the world's largest gold and copper mine located in the western Highlands and owned by US-based Freeport MacMoran. Various forms of unrest have waxed and waned through the colonial and post-colonial periods [78]. In 2003, Papua was divided into two provinces, "West Papua" and "Papua", with the independent state of Papua New Guinea to the east. Our work is in the Province of Papua.

Despite being on the second largest island on earth, there are only 2.8 million people in the province, just 8.9 people per square kilometer. The population, however, is growing rapidly due to settlement of migrants from other Indonesian islands. The province faces many societal ills: high infant mortality (56.6 per 1000 [13]), low life expectancy (63.07 [13]), low Human Development Index score (65.36 [84]), and a high number of HIV/AIDS cases [16]. These issues are often viewed by locals as connected to wider patterns of exclusion from the benefits of official development programs [16].

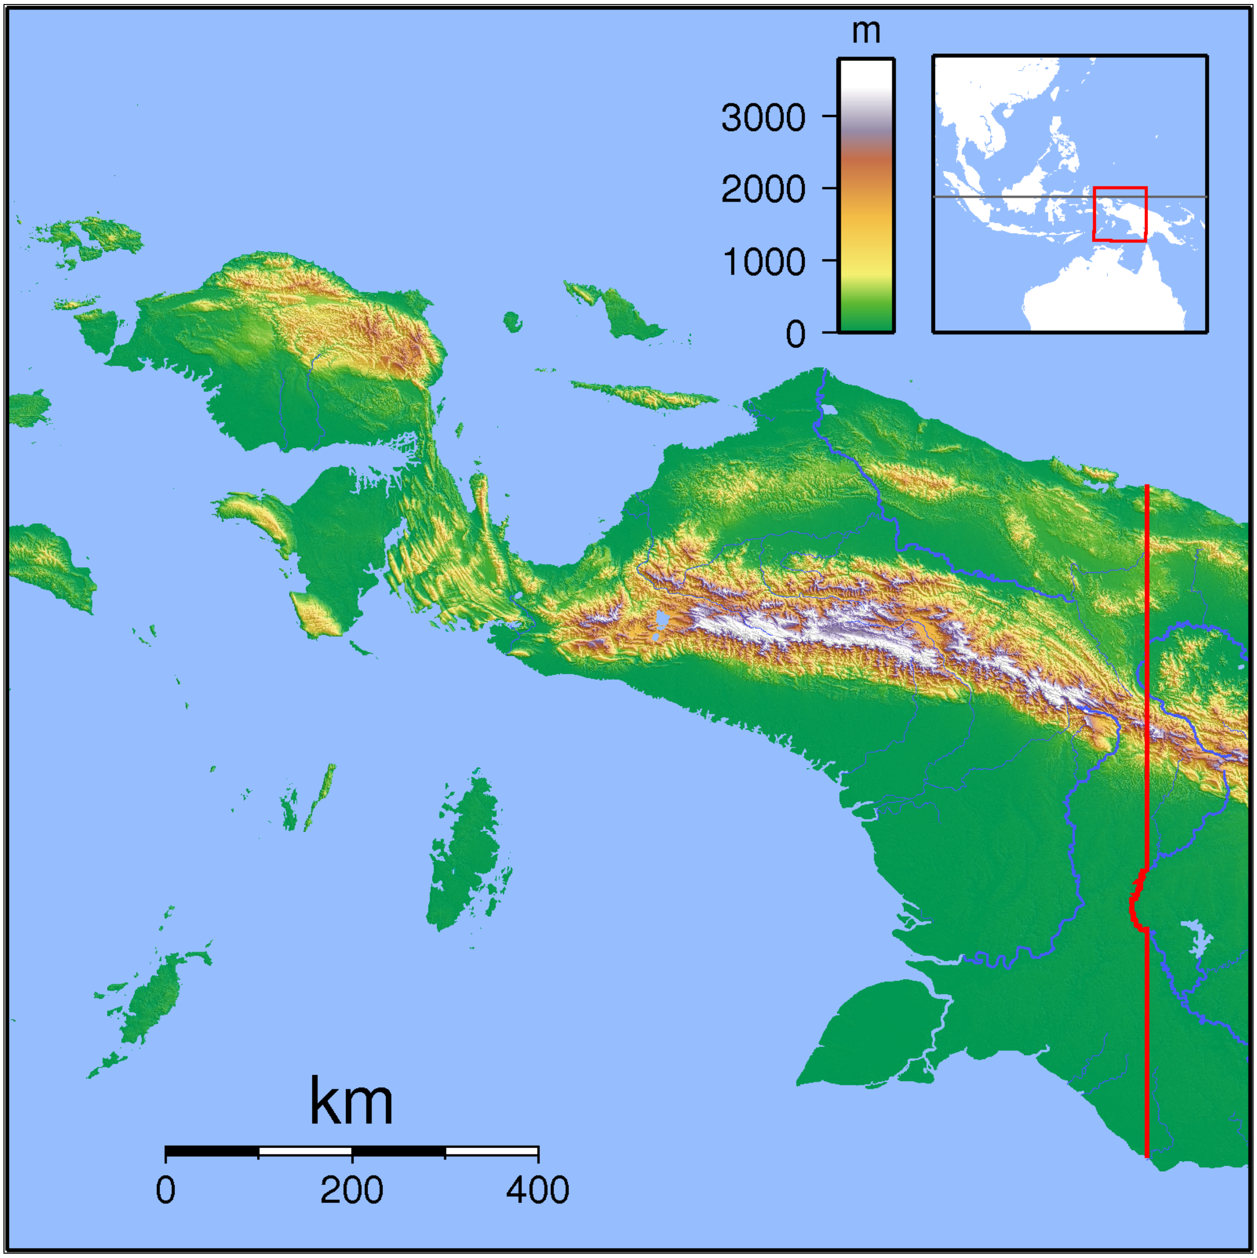

Figure 5.1: Papua, Indonesia.

5.1.2 The Highlands

Our intervention takes place in the Central Highlands of Papua. Despite a long history of regional trade connections [73], the area had no contact with Westerners until the 20th century, when a scientific expedition happened upon the population in 1938. Following this "discovery", the region was host to extensive missionary activities and the establishment of Dutch police posts.

The Central Highlands region has undergone dramatic transformations since the first development programs began in the mid 1900s. Wamena, the region's commercial and administrative center, is now a bustling town with a rapidly growing population and economy, targeted with major state infrastructure investments, and drawing migrants from the surrounding rural area and distant Indonesian islands. These infrastructures are concentrated in the town itself; there are no roads connecting to Papua's other cities, and virtually all goods are transported in by air. A number of companies provide Internet access. A local provider, called WamenaCom, run by an Indonesian and an American, is our technical partner for this installation. Connectivity in the highlands relies on satellites; there are no wired or fiber connections to other Papuan cities, nor off Papua itself.

Infrastructure expansions have recently proliferated outside Wamena, specifically in places selected as new administrative centers. A number of nearby villages have electricity; diesel generators are commonly present and occasionally running in government buildings and homes of wealthier people. However, villages distant from Wamena generally lack village-wide communications or power infrastructure. We installed our network in one such village, *Desa*.

5.1.3 Desa

Desa is located four-hours drive from Wamena and has neither a cellular network nor village-wide electricity. Administratively it is the seat of a district and thus is the site of government buildings, a police station, a military command center and post, a health clinic, two churches and a mosque, a produce market, shops, government schools, as well as a church-owned private primary school. We have been unable to find demographic data for the area; our estimates place Desa's population at 1,500 and the district's at 10,000. Although a majority of the population is indigenous, there is a significant population of migrants as well. The population includes a number of civil servants, church workers, police officers, and soldiers. Agricultural production is the backbone of the regional economy: farmers harvest produce for subsistence and for sale at nearby markets. Desa's economy is shaped by its status as a district center, with cash available from salaries, institutional budgets and produce sales. Still, Desa is economically and administratively subservient to other centers, in particular Wamena and the regency capital, both located several hours drive away.

As an early regional center of Western missionization and Dutch colonial state building, Desa has partially retained its status as a center of administration, commerce, and church organization. The availability of commodities brought by missionaries in the 1950s led Desa to become a hub for regional trade [73]. Patterns of local authority are linked to the history of missionary education, and the church plays a large role in the networks of patronage and obligation that structure local politics. The church-owned school Misionaris Sekolahin (MS) is our partner in Desa. The school is funded through donations from other countries (primarily the US) and is run by an American couple (Regis and Nancy) who have lived in Papua for over a decade.

Infrastructure Prior to our arrival, Misionaris Sekolahin (MS) installed a 5kVA micro hydro generator that powers numerous pieces of equipment: tens of light bulbs, ten laptops, a projector, refrigerators, and

Figure 5.2: The existing infrastructure: Hydroelectric power and VSAT.

other smaller loads. This hydro is on from 6AM to 10PM each day; it must be stopped to allow the water reservoir to recover. The hydro is shown in Figure 5.2.

MS also maintains a VSAT Internet connection. Internet access was primarily used by teachers to keep in contact with their friends and family outside of Desa. It was also used for educational purposes, as teachers download informative YouTube videos or share Wikipedia information. The VSAT is also shown in Figure 5.2. This Internet connection was shared via a switch connecting multiple WiFi routers: one local to the school, one long-distance link to a nearby community, and one large omnidirectional link for village access. Network access was originally controlled by a shared WEP key, but they added a Mikrotik hotspot system after some congestion issues. A small battery bank with 4 deep-cycle batteries was used to power the VSAT (and a printer/copier) when the hydro was down.

5.1.4 Deployment

In mid-October 2012, we flew to Wamena to join WC in bringing cellular access to Desa. WC had been supporting MS for years, providing technical support for the VSAT installed in Desa. This VSAT was intended for school business (e.g., instructional lectures, downloading education games) but half of the traffic is actually Facebook. This is accepted, as Regis believes that the teachers will leave if they can't communicate with their families.

The team (researchers, WC, MS) decided that the network will be owned by WamenaCom, with most profits returning to MS as payment for hosting and protecting the equipment. WC provides maintenance in the event of hardware issues, with the researchers resolving any software problems.

User Base Estimating the potential user base in Desa is challenging. We conducted numerous informal interviews and learned that phones are prevalent, despite the lack of coverage. Entertainment was the primary use case given by interviewees.

However, due to a quirk of GSM's design, we were able to measure the number of unique phones in Desa during the period of our study. Phones will attempt to connect to any BTS at all if signal is unavailable; this is to provide emergency services. We recorded all such attempts to connect for the month of January 2013. During this time period we had 1060 unique handsets attempt to connect to our BTS. 356 of these were present for five or more days, indicating a sustained presence. Our findings understate total handsets, as many users have their handsets off while in Desa. These results indicate that we have a potentially large customer base available.

Preparation and Activation Though we arrived in Papua with a mostly functional system in late October 2012, we spent months resolving technical and social concerns before activating the system. Technical concerns included interconnection failures (with the global telephony system), a drill intersecting a motherboard, and building our billing and management systems. Our social concerns involved community meetings asking for feedback and long-term beta tests among some influential villagers. This allowed us to modify our technology to better suit the community and reduce problems before we had paying customers. Eventually, issues were resolved and the system stabilized. On February 11th 2013, the network opened for business.

5.2 The Papua Network

The Papua network went live on February 24th, 2013 and has operated continuously to this day, December 2013. The network is profitable for both the operators and the community. As of September 2013, the network has served 200 customers and handled over 160,000 communications (voice or SMS).

5.2.1 Community Cellular

Ownership & Governance:

The Papua network is governed by the heads of the two respective operating institutions: Misionaris Sekolahin (MS), a missionary primary education school, and WamenaCom (WC), a for-profit Internet service provider. WC makes all technical decisions: where to install the tower, pricing, numbering, and services available. MS makes social decisions, such as who sells network credits. They also provide management such as reporting network failures to WC, dealing with personal conflicts among buyers and sellers, and selling credits at wholesale to the resellers. Profits are used by MS to provide services within the community, such as dinners for teachers and paying for the school's Internet access. Both groups were advised by the researchers, though final decisions were left to MS and WC.

The Papua network operates as two autocracies; one focused on technical features and another on community operation. This allows the network to listen to and operate within the community but still maintain the level of technical knowledge and profit to operate sustainably. Unfortunately, the key loss in this governance model is that no indigenous people are in a decision-making position for the network. This no doubt limits the level of tailoring for the specific community. The local missionaries, despite their desire to help the population, sometimes make decisions the population would not, such as blocking pornography.

Scale:

The Papua network is currently one node, located in the center of Desa. There are plans to expand the network to two nearby communities through long-distance WiFi [71], though further expansions would be new CCNs.

Goals:

There are three goals for the Papua network. First, bring communication to an area without it. Second, do so in an economically sustainable way, for both the owners and the community. Third, do these in a way that encourages local communication, rather than only being a globalizing force.

5.2.2 Design Choices

Licensing:

The Papua network effectively runs with a “grey” license: the regulators know of the network and have tacitly agreed to its existence. This was accepted by the regulators because they know that there is no coverage in the community and no plans to bring it in the future. Though no formal paperwork has (or will) be filed, this means the network can continue to operate in the near future without the fear of a crack down. It also means that no major fees were paid. The operators have made it clear to both the regulators and law enforcement that, if they ever have a request for access, it will be filled. Similarly, if a national carrier were ever to serve the area, the network would be shut down. The primary negative effect of this on the network is that the operators have been hassled by local law enforcement officials trying to hustle a free connection. This was not given, and the officers realized that the removal of the network would not benefit them.

Access:

Papua is an open network. This was done for two reasons: First, the community is relatively small and unlikely to saturate the network resources, reducing the need for admission control. Second, the operators demand profitable operation as sustainability is a primary goal. Opening the network seemed likely to maximize the chance of that. The access implementation for the Papua network is described in Section 6.3.

Naming/Numbering:

In order to simplify the understanding of the network, Papua uses the traditional model of naming. Each user is provisioned a globally addressable number when the SIM is sold and this number remains associated with the SIM indefinitely. These numbers happen to be Swedish, despite the network being in Indonesia. This was done as these numbers were the cheapest available for purchase.

Some local people have short-codes assigned to them. The doctor and chief of police both have 3-digit numbers associated with their SIMs (100,110). This number is advertised on the documentation subscribers receive when buying their SIM. More generally, 3-digit numbers are reserved for system services, 4-digit for broadcast services, and 9-digit for users. The implementation of our naming/number system is given in Section 6.3.

Pricing:

In the Papua network, there are three tiers of prices: Out-of-community, local, and free. Out-of-community prices are the most expensive and this is the primary profit generating activity in the network. Local calls and SMS are priced more cheaply as these calls have no marginal cost for the operator, but generate less profit than out-of-community calls. This was chosen so that the network would encourage local connections, rather than just be a link to the outside world. Despite the fact that supporting these communications is effectively free, the operators wanted to charge for them to limit use and avoid network congestion. Last, free communications are for services and individuals the network explicitly wishes to support, such as emergency medical and police communications. The actual prices set by our partner are given in Section 6.5.

Payment:

The Papuan users demanded a pre-paid system. Apparently hurt by post-paid systems in banks and other institutions, any post-paid system (including monthly fees) would be looked at as swindling, as the operator would still get paid despite not providing any active service. As such, a user on the Papuan network need only buy a SIM card from the credit seller for 100000 rupiah (US$10.30) and then credits for use from the same person. No monthly fees were involved. The payments implementation is given in Section 6.2.

Services:

The Papua network utilizes SMS as its primary mode of outbound communication. This was done to reduce network bandwidth; the network uses a shared VSAT that is often congested. All outbound communications from users occur using SMS. Locally, the network supports both voice and SMS. The network technically supports outbound phones [1] but this service was not made available to subscribers.

The operators in Papua wanted to build a network that encouraged local communications and community. They deployed two key services to support that: Village Idol and SMS mailing lists. Village Idol was similar to Bali et al. [58]. Users call in to record a song, and after a week users could call in to vote on the best song, with the winner receiving network credits. The SMS mailing list allowed users to send "broadcast" SMS to a small group of people and was primarily used by the school to organize social events. Another service, Find-a-Friend, was implemented but not used. This service would randomly connect a caller to another subscriber. Users who talked for five or more minutes would receive credits; the goal was to break clique lines in the community. Unfortunately, during initial testing the operators were informed that there was concern that gender lines may be crossed, and a married person speaking with an unmarried person of the other gender could lead to violence. As such, the system was shelved. More functional services were also deployed. Credit transfer allowed any user to send network credits to any other user. Originally built to enable credit sales, community members soon began using the service to send small payments. Due to community feedback expressing concerns about the reliability of SMS, the operators added an SMS delivery confirmation service. The implementation of these (and other) services are described in more depth in Section 6.4.

[1] The researchers made numerous calls back home, in fact.

Design Choice	Papua	Traditional
Licensing	Grey	Legal
Access	Open	Open
Numbering		
Scope	Global	Global
Sharing	Individual	Individual
Dynamism	Static	Static
Pricing	Tiered	Flat
Payment		
Access	Pre-paid	Both
Use	Pre-paid	Both
Services		
Local	SMS, Voice	SMS, Voice
Global	SMS	SMS, Voice
Community	Many	None

Table 5.1: Design choices made in the Papua network compared to traditional networks.

Chapter 6

System

In support of the Papuan network, we built an entire telecommunications company worth of infrastructure: hardware, billing, and services for our partners and users. The implementation of our system was guided by three key principles:

- Simplicity: Minimize potential failure points;
- Familiarity: Mimic existing networks when appropriate; and
- Locality: Involve the community in design decisions.

With these principles in mind we detail our system; the hardware installed, the software used, and the services built.

6.1 Hardware

Traditional cellular hardware is complicated; instead we wanted a smaller, lighter, more flexible system. OpenBTS [67] allows us to do this; it eschews the need for a variety of related infrastructure and instead allows operators to build a complete GSM base station using only a commodity PC and appropriate radio equipment. Here we detail all of the hardware required for running our rural telco, including the base station and tower.

6.1.1 Base Station

We purchased a 10W Range Networks 5150 BTS and mounted it in a weatherproof box acquired in Papua. We added a 24V→12V voltage regulator to insulate the machine against power failures and fluctuations [87]. The complete GSM base station consists of the Range Networks radio equipment and a commodity PC: no other supporting infrastructure is needed, unlike traditional GSM networks. We tower-mounted the base station to reduce cabling and RF loss and protect the hardware from theft. The completed installation is shown in Figure 6.1. The BTS was plugged into WC's switch for Internet access and attached to their battery bank for overnight power. Two batteries and a larger battery charger were installed to support the extra power draw from the BTS.

Figure 6.1: Our Base Station and its installation.

6.1.2 Tower

Building a tower out in the village is a difficult proposition; concrete and other building materials are heavy and expensive. Instead, we mounted the BTS on a pole, which we then mounted in a large tree with few branches. This is also shown in Figure 6.1. Other unrelated wireless equipment, which shares the local VSAT with other locations, is mounted on the same pole.

6.1.3 Phones

As Desa is three hour's drive from coverage, we were not certain that there were phones available to the community. Though this would no doubt eventually be resolved; we were concerned. To resolve this, we recorded every phone's attempt to attach to our tower before the network opened. This would provide a lower bound on the number of phones available in the village, as many phones would be powered down (as there was no signal until we opened the network to subscribers).

From December to February, we observed over two thousand unique SIM cards in Desa. This is supported by informal interviews indicating that phones are not rare; they were commonly owned and used for playing music and games even without network coverage. As such, we expected users to provide their own phones for use on the network.

6.2 Network Credits

Unlike prior work [29, 76, 96], our network is for-profit; our users actively pay our NGO partner and affiliates to use it. We believe that local ownership (and local profit) incentivizes local actors to support the network; if local people profit from the network, they are more likely to care when it fails.

Most users in Indonesia use prepaid cellular plans. Local entrepreneurs sell network credits to users for a small premium, giving profit to both the network operator and these entrepreneurs. We built a similar system for the network in Desa. Users pay our NGO partner money for credits which they then use for communications.

Our NGO partners don't sell credits directly to potential users. Instead, they sell credits to a primary reseller, who then resells them to local merchants who resell to the users. These resellers (recruited by the primary reseller) set their own prices for credit sales (see Table 6.1). Each seller in that system makes a profit. The primary reseller only sells in bulk to the other sellers, who sell smaller amounts to local buyers. The sellers are, at this moment, only non-indigenous people, for a variety of reasons. First, they had the capital to invest (buy credits from the primary). Second, the two secondary sellers both own local stores that have power, allowing them to charge their phones. Third, a few of the Papuans with the capital and infrastructure to potentially sell credits have personal or political issues with MS. Lastly, the fact that anyone can do a credit transfer has not yet been publicly announced, keeping people from building smaller agencies. It is an explicit goal of MS and the researchers to involve local Papuans in this process, and we hope to do so in the future.

Interestingly, this was not the original system. We originally had an indigenous man (whose wife was a teacher at the school) as our only seller. This seller, like all teachers, lived on church property. The church took umbrage at us operating a commercial service on church grounds and forced him to stop. This happened when the researchers were away, and the primary reseller and MS decided to set up the current system without researcher participation.

Credits	Primary (Rev)	Secondary 1	Secondary 2
1000000	1050000 (5%)	-	-
500000	535000 (7%)	-	-
100000	-	110000 (10%)	105000(5%)
50000	-	57000 (14%)	55000 (10%)
20000	-	25000 (25%)	24000 (20%)
10000	-	13000 (30%)	12000 (20%)

Table 6.1: The price to purchase credits in our system as set by each reseller (and their revenue on each sale).

6.3 Software Components

Our system is a modified version of the Village Base Station [43] (VBTS). VBTS is a set of extensions to OpenBTS [67] we developed allowing for multimodal applications utilizing voice, text, and other mediums. This toolkit, in combination with FreeSWITCH, was used to build all of the services shown below. Our system also utilized some custom web components to speak with our SMS provider Nexmo.

6.4 Services

A full-scale cellular network provides a variety of functions including voice, SMS, and billing. In order to provide coverage in Desa, we implemented these and other services: provisioning new numbers, a pre-paid credit system (including credit check, transfer, and purchase), number checking, and delivery receipts.

6.4.1 User Provisioning

The first step in our system is the selling and provisioning of SIM cards to users. In GSM, SIM cards and phone numbers are different entities; they must be associated together by the network operator.

In our network, SIM cards were purchased from a vendor online and phone numbers are provided by Nexmo. The cheapest numbers provided by Nexmo are Swedish, which is what we sold to our users. We then use Nexmo to route communications (just SMS today) to the outside world. This means that all users in our network have Swedish phone numbers, not local Indonesian ones. The reasoning for this decision is described in more depth in Section 8.2.

SIM Cards

Our network provides standard 2G GSM coverage. This means that any existing GSM handset can connect to the network, assuming it has a valid SIM card. What SIMs can connect is a configuration setting in the BTS; we could require specific SIMs or accept any cards. We chose to manufacture our own SIM cards instead of utilizing the ones from existing carriers (Figure 6.2). Only our SIM cards can connect to the network in Desa. We did this to simplify the user's model of the network; when in Desa our SIM card works. When in the larger town, our SIM does not. SIM switching is common in the developing world [88] and Indonesia.

Figure 6.2: Our SIM cards.

With this design choice made, we fabricated 1000 SIM cards (at US$0.65 per card) and programmed them with the Indonesian Mobile Country Code (510) and our unique Mobile Network Code (55). This made certain that no other networks would accept our SIMs. Potential customers buy a SIM from our primary reseller at cost; the reseller makes no profit on that sale. This card is accompanied by documentation detailing the network, duration of service, our communication prices, and services in the network. The card is inserted into a user's phone and the seller provisions a new phone number that is then associated with that SIM.

Our SIM cards are sold with one year of coverage for 100000 Rupiah, approximately US$10. This price covers one year of number rental from Nexmo (US$0.65 per month) as well as the manufacturing of the SIM itself (US$0.65); our NGO partner makes approximately US$1.80 per SIM sold per year.

Number Provisioning

As mentioned above, we use Nexmo for outbound SMS service. We implemented a mechanism to automatically buy new numbers from Nexmo when provisioning a new SIM card. The system works as follows. When a new SIM card is purchased, it does not come with a set phone number. Instead, the seller places it in the customer's phone and immediately sends an SMS to 101 signaling their desire to join the network. The system queries our database of users and our stock of purchased Nexmo numbers. If there is an available number (i.e., purchased from Nexmo but not assigned to a user) we assign that to the user. If there is no such number available, we instead purchase a new number (at random) from Nexmo and assign that to the user. The user receives an SMS from 101 with their new phone number. The SIM seller is trained in this particular short code; it was not advertised to the community.

6.4.2 Credit System

We utilize a pre-paid credit system similar to that used by other Indonesian cellular providers. To support this, we had to build services for users to buy, check, and transfer their credits.

Credit Creation

At the highest level, our NGO partner *creates* credits in the system to be sold to users for basic communications; they do this via a password protected web interface. This service was not originally written by us, but modified from Range Networks software. After credits are purchased by the primary reseller and added to their account, he or she uses our *credit transfer* service to resell to secondary sellers. The credit purchase website is only available to the owner of the network: our NGO partner.

Credit Transfer

After credits are inserted into the primary reseller's account, they sell their newly acquired credits to secondary sellers via *credit transfer*. Likewise, these secondary sellers use the same system to transfer credits to actual users.

Credit transfer is initiated by sending a specially formatted SMS to the number 887. This format is TARGETNUMBER*AMOUNT. If the SMS is malformed, the user is sent a set of instructions. If the user has insufficient credit, a message informing them of this is sent. If correctly formatted, the user is given a short message detailing the transfer and asking for confirmation, with a four-digit confirmation code included. This was originally 7 digits, but seller feedback indicated that 7 digits were hard to input via SMS.

If they respond (again to 887) with the confirmation code, their transfer is completed and both the buyer and seller are sent a message indicating that the transfer was finished. All SMS are received by the users from the number 887. The researchers or primary reseller instructed each secondary reseller on how to use the credit system. The service was initially kept from the public, but was advertised via broadcast SMS on May 7, 2013.

887 was chosen because of it was close to 888 (our credit check service) and we wanted users to understand that these numbers were all used for BTS services. This service was never advertised. Instead, we demonstrated it to our primary credit seller. He, in turn, demonstrated it to any secondary sellers purchasing credits from him.

Credit Check

To check their current credit level, a user of our system sends an SMS or call to 888. An SMS response is sent (even if they called) from 888 indicating their current credit level. The number 888 was selected because the national telecom provider uses this as their credit check number. This service was advertised to users on the paperwork provided with their SIM card.

6.4.3 Basic Communications

Our network provided both voice and SMS services locally, but only SMS service for outbound communication. The reasons for this are complicated. As mentioned in Section 8, interconnecting with other networks is very difficult. Interconnecting with *both* voice and SMS is even more difficult. As such, we quickly decided that we would only be able to provide one such service for bidirectional communication; users would either be able to SMS or call back and forth with their out-of-network friends.

We conducted a series of informal focus groups with people in Desa to decide if they would prefer outbound SMS or voice service. The results were surprising; though most users communicated primarily through SMS, they would prefer to have only voice service if they had to choose. The reasoning was that some contacts were voice-only (such as older family members) and they wanted the ability to communicate with those contacts.

At the same time, the limitations of our shared infrastructure were beginning to become apparent. SMS is asynchronous and can be delayed; voice requires significantly higher bandwidth and quality of service. We were unable to provide quality voice service without severely impacting other VSAT users or changing providers. As such, we decided to support only SMS for two-way out-of-network communications for now.

6.4.4 Other Services

We opened our network for beta testing in January of 2013. Ten key Desa community members, including missionaries, Papuans, and non-Papuan Indonesians, were invited to participate and given free SIM cards. On February 11th 2013, the network opened for general use. During this time we monitored usage and conducted informal interviews. We discovered a few issues, and we implemented services to remedy these concerns.

Number Check

First, users did not remember their phone numbers. This is likely a consequence of the fact that they had Swedish numbers; these were strange and did not look like traditional numbers. To resolve this, we implemented a service for checking your phone number.

A user sends an SMS or call to 889 requesting their current number in the system. An SMS response is sent (even if they called) from 889 giving their current phone number. 889 was chosen as it is close to 888, and signals that it is a network service. The number check service was advertised on the paperwork provided when a user purchased a SIM card.

Delivery Receipts

The second issue we noticed was one of trust; users were often uncertain about the status of their messages. Indonesia is rife with SMS-based fraud, including spam and phishing. This fact, coupled with our strange Swedish numbers and the expectation that Desa lacked coverage, meant that many SMSs were initially ignored by recipients. Users did not know if the lack of response meant a network failure, a rejection by the receiver, or something else.

To remedy this, we implemented delivery receipts. Delivery receipts inform a user when an SMS has been delivered to the target. We chose to implement the service in an opt-in manner; users signal their participation and then all SMS are given receipts. This choice was made to again mimic the national carrier (or certain handsets), who also use an opt-in model.

To enable delivery receipts, a user sends an SMS to 300. They then receive an SMS from 300 signaling their participation. With delivery receipts enabled, a user would receive receipts again from 300. To disable delivery receipts, a user sends an SMS to 301. As this issue was identified after service activation, delivery receipts were advertised via a broadcast SMS.

6.5 Pricing

In collaboration with our NGO partners, we set a pricing scheme for the basic services our system provides. We had three key goals. First, do not take a loss on any part of the system. This was because our technical partner, WC, was risk averse, having had numerous experiences where operations at a loss were abused by users. Second, we wanted to make a reasonable profit for the operator, MS. They were doing a lot of work and powering the equipment, so they wanted a return on that investment. Lastly (and in that order), we wanted to provide as much coverage as possible to the community.

Table 6.2 shows the result. Note that non-local calls are not yet implemented; this is due to limited VSAT capacity degrading call quality and an uncooperative VSAT provider. These prices may seem high but are justified by the actual costs covered in the next section. They are also not far from those of early networks in Africa [45], which also had limited infrastructure and low population density. We hope to explore other pricing schemes and levels in future work.

Service	**Shortcode**	**Price (US$)**	**Usage (%)**
Number Provisioning	101	free	100%
Delivery Receipt	300/301	free	17.1%
Credit Transfer	887	free	24.6%
Credit Check	888	free	91.4%
Number Check	889	free	23.0%
Local SMS	–	.02/sms	82.4%
Global SMS	–	.09/sms	93.0%
Local Calls	–	.02/min	84.0%

Table 6.2: Services supported, their prices, and usage rates.

Chapter 7

Evaluation

We evaluate our deployment of a community cellular network in rural Papua, Indonesia, via a variety of metrics. We begin by calculating the basic usage statistics of the installation and how the subscribers make use of the network services. We follow this with an analysis of virtual coverage, investigating the night usage of the base station and the Wake-up Radio and the total night power saved in the network. We then explore the sustainability of the installation by calculating network revenue and costs including financing the complete installation. We follow with a set of interviews with BTS users, inquiring into their usage of, perceptions of, desires with and concerns over the cellular network. Lastly, we analyze how our network's call graphs differ from a traditional network in sub-Saharan Africa.

7.1 Data Collection

All of these evaluations utilize one primary data source: the "call detail records" CDR. Every communication attempt in the system is recorded into a CDR. These are recorded even if the connection fails, as would occur, for instance, for calls or SMS to invalid numbers. These records include the caller and callee's unique identifiers in the system, the time and duration of the call, as well as the starting end ending towers for both sides of the call. In our network, some data (e.g., receiver tower ID) is not available. Similar data is recorded for SMS. These records are used to determine the total number of communications. CDRs were recorded for the operation of the BTS, starting February 11th to the date of authoring, Dec 1st 2013. Each evaluation uses a different ending date for their own analysis, as each evaluation was conducted at different time in the project's life.

7.2 Usage

The system opened for customers on February 11th, connecting 187 subscribers by August 12th (Figure 7.1).

7.2.1 Services

Figure 7.2 shows usage of our most important service types over time. This figure excludes calls or SMS made in error (typically caused by attempting to call out-of-network numbers, which we do not

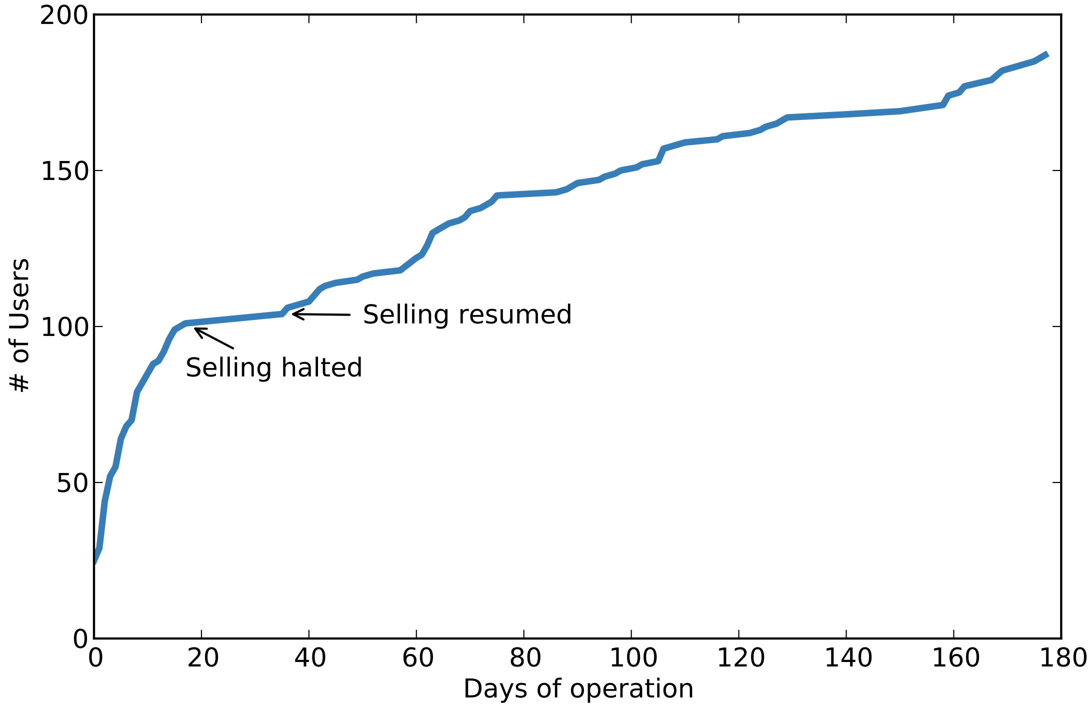

Figure 7.1: Number of users over time. Note the period of inactivity where selling was halted due to a shortage of numbers.

support, or attempts made when the user's account was depleted) and use of free services like credit checks. Out-of-network SMS was by far the most popular service in our network, with over 1500 messages sent per week; SMS from out-of-network to our users ("incoming SMS") was tightly correlated with outgoing SMS frequency and comprised our second most popular service type. Local, in-network usage was typically lower, though in week 8 we saw a large spike in local usage, primarily driven by a sudden increase among a group of 5–10 users.

The preference for out-of-network communication seen in our network is not driven by a handful of heavy users; on the contrary, most of our users communicated more with out-of-network contacts than local ones. This distribution is shown in Figure 7.4. Notably, we had more than 10 users who *only* sent out-of-network SMS, and only 20 users who sent more local SMS than out-of-network SMS. This is despite the fact that out-of-network SMS costs over four times as much as local SMS.

Table 6.2 shows the proportion of users who used each network service at least once. 91% of users checked their account credit, the most popular network service. Though all users received a credit transfer (that was the only way to purchase credits), 24% of users transferred credit to another user, despite it not being advertised until halfway through the study.

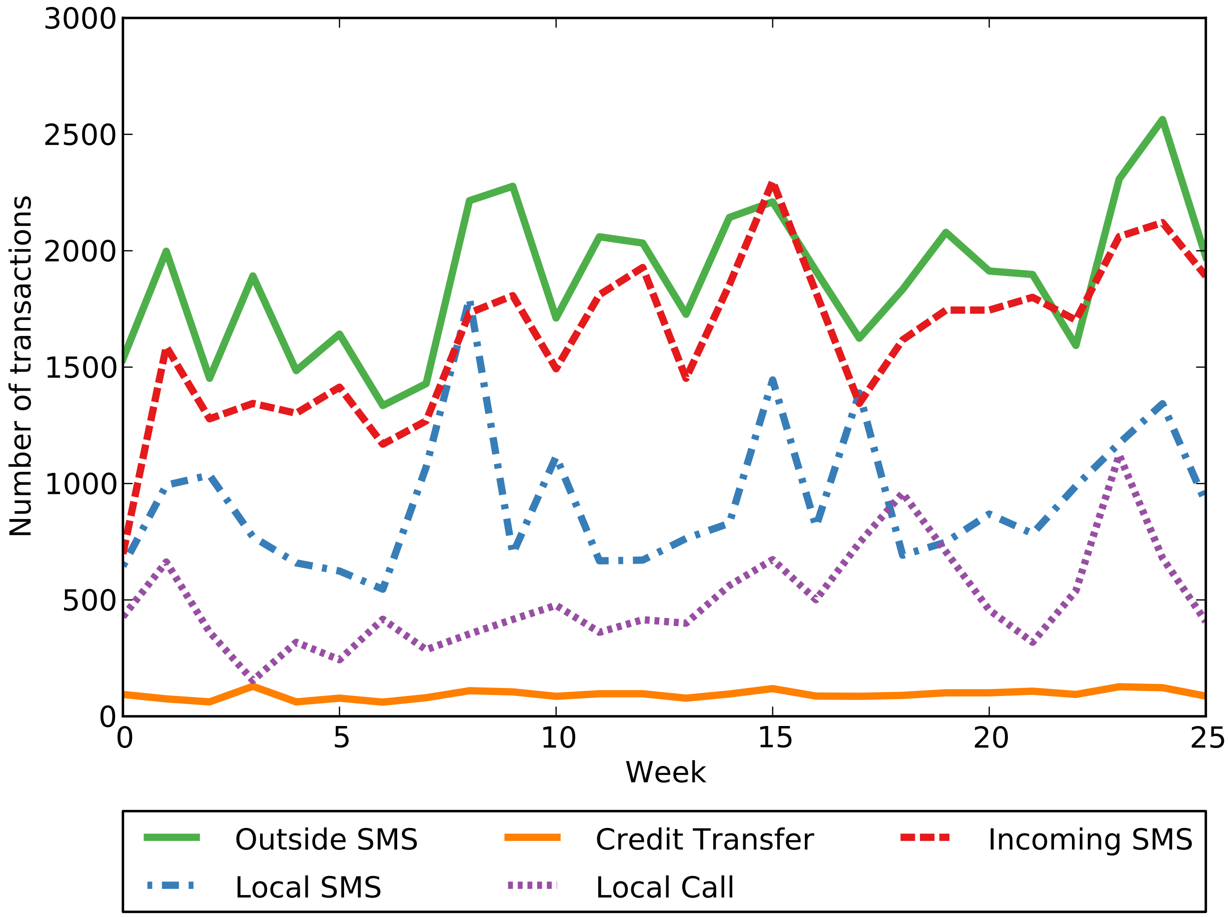

Figure 7.2: Per week overall usage of key service types. Out-of-network messaging was most popular.

7.2.2 Demographics

Table 7.1 shows user demographics for the first 100 SIM card sales. This data reflects individuals who bought a SIM card from the primary reseller; we have no way of knowing who actually used the SIM card once purchased. Although there is no census data to compare against, we believe non-Papuan Indonesians and Westerners are modestly overrepresented in the network, presumably because these groups tend to be wealthier. We observe a gender gap between male and female users exceeding the 17% gender gap seen in the rest of South East Asia [37]. At least part of this gap may be explained by men purchasing SIM cards on behalf of women.

7.3 Virtual Coverage

We evaluate virtual coverage in three core ways. First, we investigate the use of the system and whether subscribers were able to understand and utilize virtual coverage. This is done though an analysis of call records through the night periods, showing the volume of, nature of, and participation in night

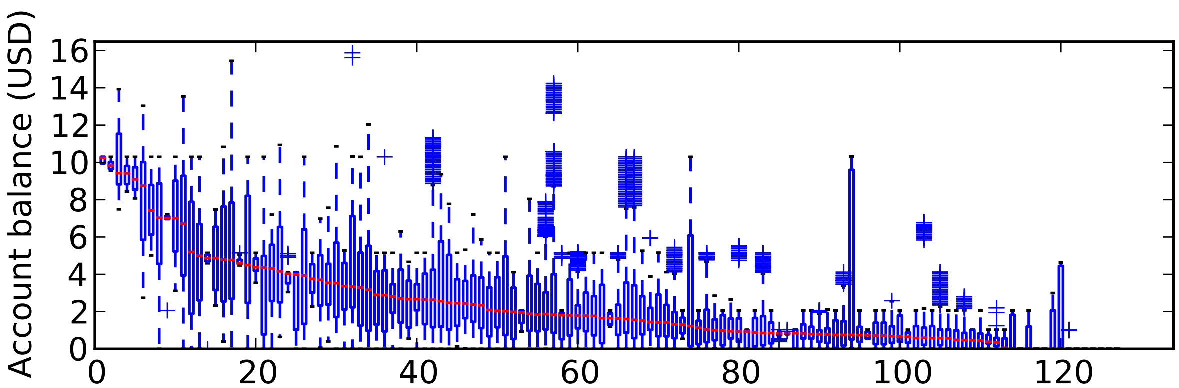

Figure 7.3: Distribution of account balances for each user. This omits 7 users who maintained extremely large credit balances, either due to being involved in a credit counterfeiting scheme or official credit resellers.

Origin	
Non-Papua Indo.	34%
Highland Papuan	54%
Coastal Papuan	6%
Western	6%

Gender	
Male	82%
Female	18%

Table 7.1: Demographics of first 100 SIM card purchasers.

communications. We then evaluate if our network was able to meet the core goal of virtual coverage: reducing night power consumption. This is done through BTS logs of "sleep" and "wake" events, showing the amount of power saved through the use of virtual coverage. We also consider the actual design of the WUR through observations of its usage. To begin, we first discuss the relevant background system statistics and the methods of data collection.

We emphasize that most communications handled by our BTS happen during the day; the night service provided with virtual coverage is, by its very nature, limited. Desa has no power at night and the majority of the users go to sleep after the sun goes down. The point of virtual coverage is to provide service for instances where users have a pressing need for communications, such as emergencies or other time-sensitive events, while still avoiding the wasted energy of broadcasting during times of low utilization. It is our view that even light night use validates this model.

7.3.1 Additional Data Collection

As a consequence of the network design, there can be multiple "wake up" events for each recorded CDR. For instance, a user may use the wake-up radio multiple times before actually sending an SMS. Similarly, a user may send multiple SMS after waking the BTS just once. To evaluate power usage of the network, we also record a series of events in the BTS log itself and use them in this analysis. Each instance of "sleeping" and "waking" is recorded in this log. Unfortunately, misconfigured log rotation caused us to lose the earliest BTS logs. As such, we only have power measurements from March 24th to July 16th.

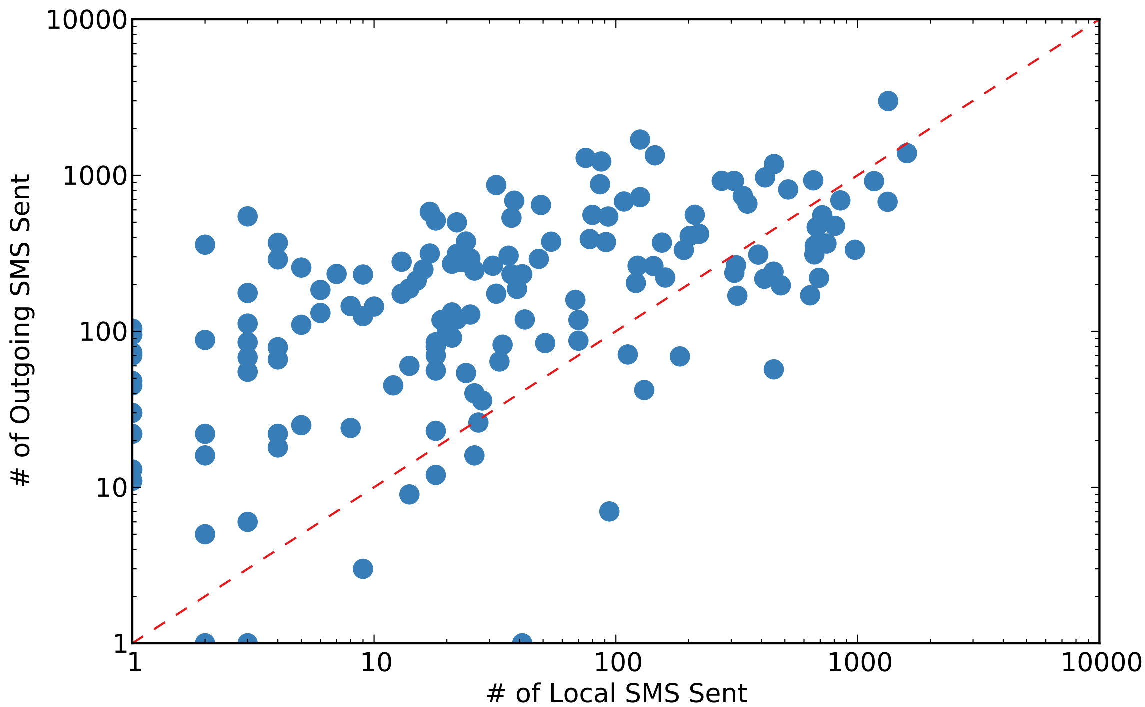

Figure 7.4: Comparison of in-network versus out-of-network SMS volume sent per user. Users tend to send more out-of-network (i.e, outside Desa) messages than local ones; more than 10 users sent no local SMS messages.

7.3.2 Night Usage

Using the both the CDRs and BTS logs, we address the core questions about the usability of the virtual coverage. We first analyze if users were able to make use of the "wake-up" mechanisms, demonstrating their understanding of the basic idea of virtual coverage. We then examine how many users communicated during the "night" periods and in what volume, demonstrating that the technology was actually used by a large portion of the population and therefore presumably useful to the network subscribers in general.

"Wake-up" Usage As recorded in the BTS log, the Desa base station was woken from low power mode a total of 566 times in four months, an average of 4.9 wake-ups per night. This does not include the daily wake-up due to duty cycling or extraneous wake-ups due to reboots. The majority (65%) of wake-ups were triggered by incoming SMS. The remaining 35% of wake-ups were the result of using the WUR button. Most of the WUR events happened concurrently; the user was either holding the button down or pressing it repeatedly. 138 of the wake-ups were repeated, meaning they immediately followed the BTS going to sleeping after a previous WUR burst. Removing these from the analysis leaves 59 independent WUR events. Using these numbers instead, we see 14% of our unique wake-ups were caused by the WUR and the remaining 86% from incoming SMS, as shown in Figure 7.5. Generally, the BTS woke significantly more from incoming SMS than from the use of the WUR. This is likely because the main island of Indonesia

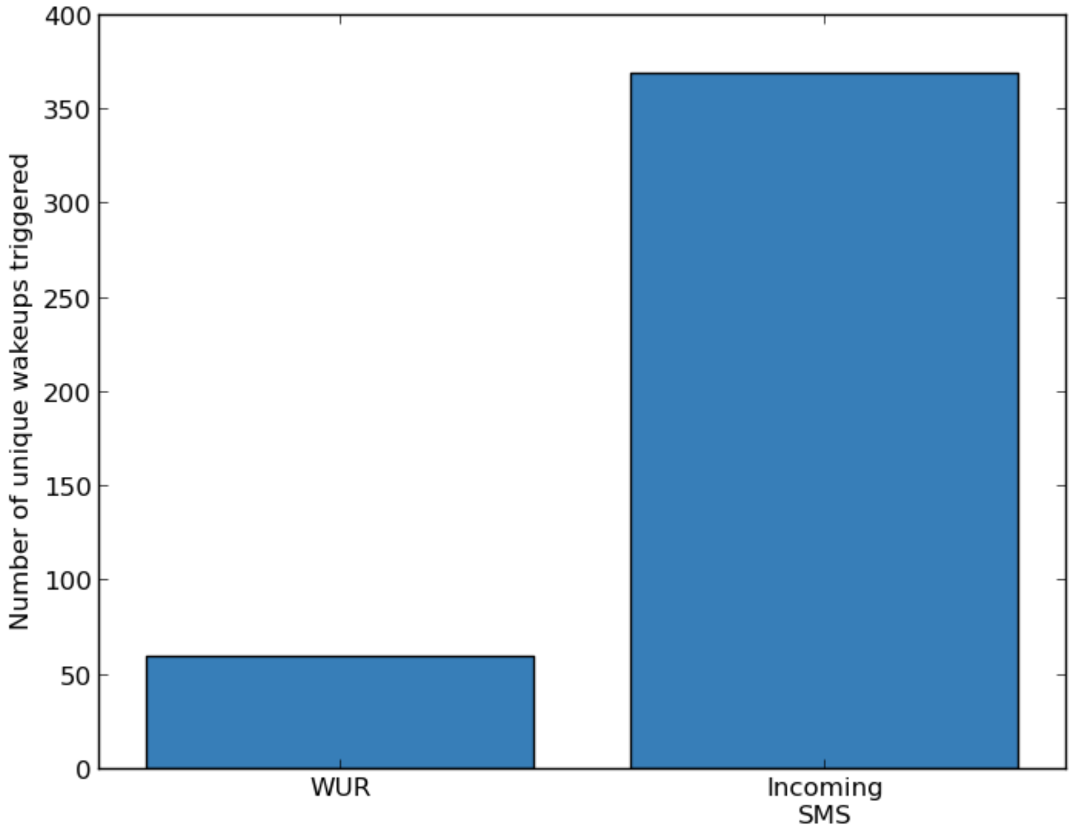

Figure 7.5: Unique causes of wake-up during night time hours. The majority of wake ups were caused by incoming SMS rather than use of the WUR.

(Java) is 2 time zones behind Papua; our 11pm timeout is only 9pm in Java. It is also possible that actively using the WUR is just less convenient than passively receiving a message.

Night Network Usage Table 7.2 shows the frequency of each type of communication on the base station, for both day and night, taken from the CDRs. We note that the "night" usage statistics only include communications that took place after the base station entered low-power mode, i.e., actually slept for the first time each day. On days of heavy usage the base station did not enter low power mode until after midnight. In general, there was a large amount of night usage by the community. A total of 1485 messages were sent or received by people in Desa over the evaluation period of six months. 86 individuals used the network during this time, just over half of the total subscriber population. Moreover, 76 of these users sent at least one "unsolicited" SMS or call, initiating a communication without having received any for at least an hour.

The bulk of communications during the night were with contacts outside of Desa (77.1%). This is slightly higher than the same ratio for daytime communications (58.1%). In general, night usage was more focused on external communications than day usage. This isn't surprising; most people in Desa are asleep

Service	Night Use (%)	Day Use (%)
Inbound SMS	755 (50.8%)	31549 (27.0%)
Outbound SMS	390 (26.3%)	36560 (31.2%)
Free Call/SMS	182 (12.3%)	16788 (14.3%)
Local SMS	77 (5.2%)	17972 (15.3%)
Local Call	34 (2.3%)	9104 (7.8%)
Other	47 (3.2%)	3237 (2.7%)
Total	1485 (100%)	117208 (100%)

Table 7.2: Popularity of services on the BTS for both night and day. “Night” refers to the times when the BTS has been awoken from sleep. “Other” refers to misdialed numbers and attempts to call or send SMS by users with no credits in their account.

at night, limiting communication options. Everyone lives within a few kilometers of each other; if you’re going to wake someone up, you may actually visit in person rather than send a text, especially if you need to walk to the WUR anyway.

We observed activity at all hours of the night. The bulk of night time usage occurred “early” in the night, shown in Figure 7.6. Unsurprisingly, this time was close to our peak daily usage period of 19:00 - 21:00. Due to a bug, the BTS was forced to briefly reset itself at 03:00 every morning, “waking” the BTS, and we see a small bump in communications at this time. We note that there are nearly 80 times more communications during the day (117208) than at night (1485). Though the day period is longer than night (17 hours versus 7 hours), this is insufficient to account for the difference. Instead, we believe the night is simply a time when user communication needs are reduced. This supports our design; our goal is to save power during these periods of network idleness. The need to go to a WUR to turn on service also disincentives usage.

These results validate the findings in earlier work [45] that naive duty cycling does not meet user needs; people in rural areas will occasionally need to communicate with outsiders during periods of relative inactivity. Virtual coverage successfully meets this need.

7.3.3 Power Savings

One key goal in implementing virtual coverage was to reduce night time power usage of our cellular infrastructure. To evaluate if our system met this goal, we measured power consumption in two ways. First, using the BTS logs we are able to precisely measure the actual power consumption of the base station. Secondly, using the call and SMS records in the CDRs, we calculated an “ideal” power consumption profile. The second method ignores periods where the base station was turned on but not actually used; this situation turned out to be far more common than we expected.

Figure 7.7 shows the cumulative distribution of hours the base station was active each night (i.e., between 23:00 and 06:00) under each of these models. The BTS remained in low-power mode for the bulk of most nights, with a median “awake” time of 48 minutes per night. In general, below the 90th percentile the actual and ideal power consumption were closely matched, suggesting that power consumption roughly matched what was required for the desired usage. However, we observed a large increase in actual usage over the ideal for nights in which the BTS was turned on the longest.

As stated above, we found six instances where the base station received a sustained burst from one of

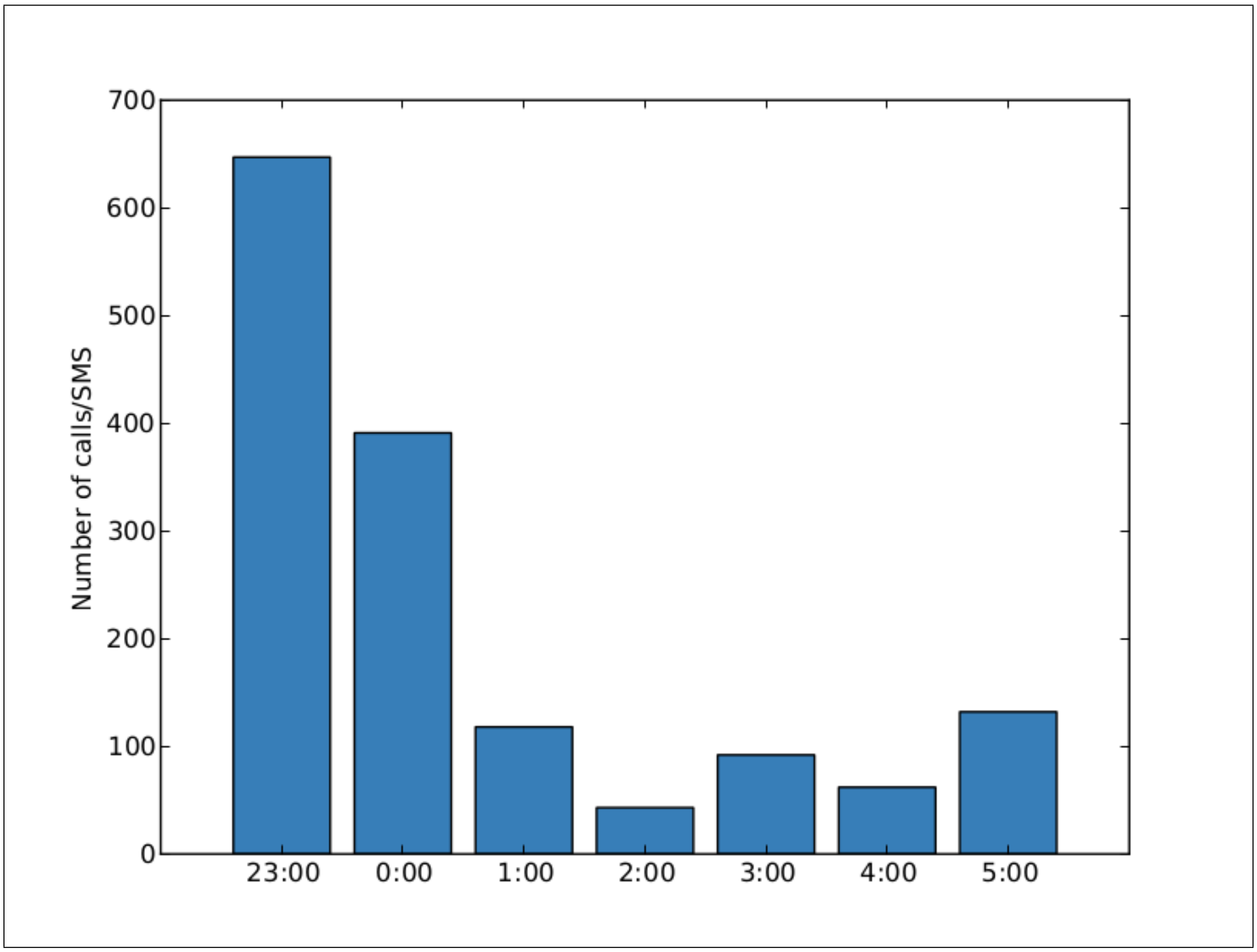

Figure 7.6: Night hourly usage. Most usage happens in the 23:00 hour, but we observed usage at all hours of the night.

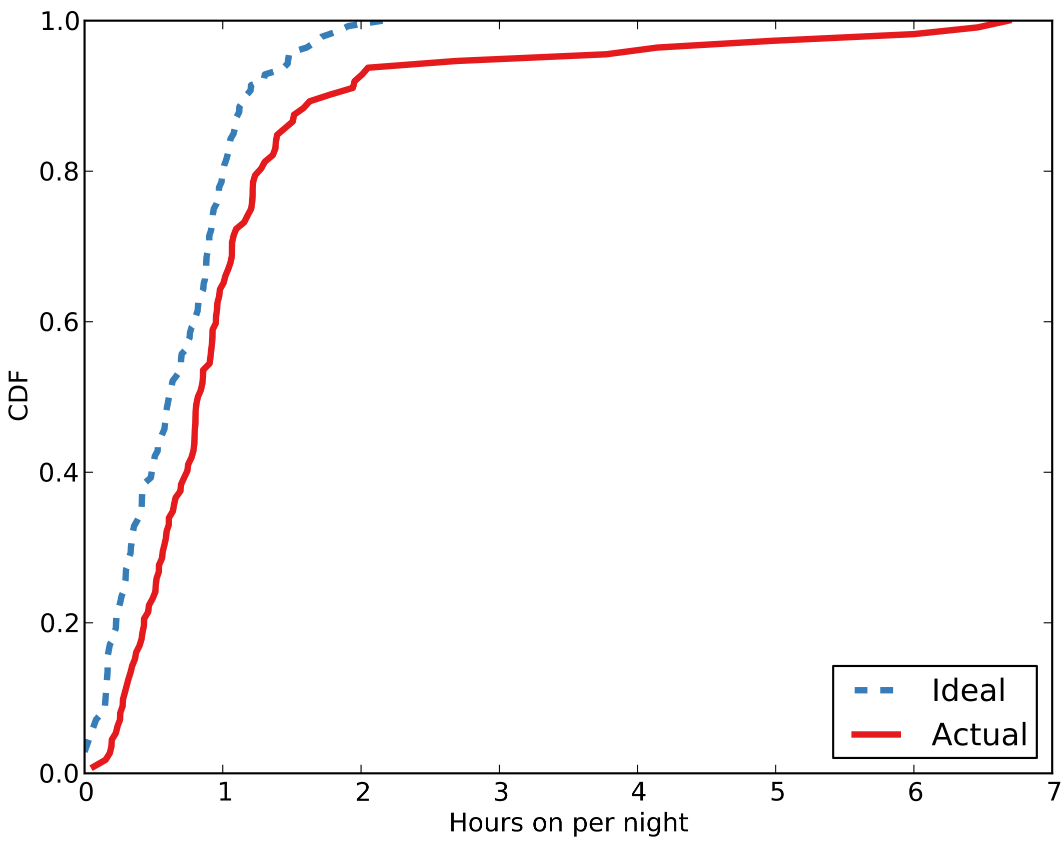

Figure 7.7: Cumulative distribution of time on per night. The actual and ideal power consumption diverge at the high end due to several instances of the WUR being kept active for several hours, preventing the BTS from going to sleep.

the WUR for more than 30 minutes, the longest of which lasted over three hours. These events took place throughout the course of our study and were not concentrated only at the beginning. Though we cannot be sure, we believe a user may have taped down the switch of a WUR to force the base station to remain on throughout the night. This may have stemmed from a misunderstanding of how virtual coverage works; when the base station is off, handsets will show "no service" to their users. Someone expecting an incoming message could be understandably worried that they may miss their message during this period, even though the base station would have woken up upon its receipt. We hope to further investigate this behavior in the future.

Despite these instances, the system was successful in reducing power consumption and remained in low-power mode for the bulk of each night. In total, virtual coverage allowed the power amplifier to stay off for 87.0% of the night (when hydro power is unavailable) providing 56.6% total power savings at night. In the ideal model where users only wake the BTS when they actually need to communicate, the BTS would be asleep 90.3% of the night. This results indicate that our BTS, by virtue of using virtual coverage, can utilize half the batteries and power generation of a traditional low-power cellular base station.

7.3.4 WUR Design Findings

In our deployment of the WURs, we found that the basic design of the radio was lacking in two primary ways: the use of the WUR button, and the design of the case.

WUR Button The WUR was designed to be as simple as possible; a user walks up, presses the big red button on the case, and the BTS turns on, providing coverage. However, we did not anticipate how users would understand and make use of this simple interface.

When the WUR button is pressed, it broadcasts a radio burst on a specific frequency that is picked up by the BTS. Users understood this, as demonstrated by the numerous wake-up events detected by the network. However, users did not understand the effect on the *battery life* of the WUR. As mentioned above, on six occasions users seemingly **held down** the button for long periods of time, causing it to continuously broadcast and wasting the WUR's battery. We assume this was done because they did not understand that the BTS did not need to be awake for incoming SMS to be delivered. A quick change to the WUR design, causing each *press* of the button to cause a broadcast (instead of broadcasting continually when pressed) would reduce the incentive to do this. Similarly, adding a light when transmitting or when the BTS is on may reduce this behavior. We propose other solutions in the discussion.

WUR Case The WUR case was designed to be portable, as one potential model of deployment was attaching the radio to individual users or phones [42]. This was a mistake given the way we deployed WURs in Desa; there was no great way to tether the WUR to a specific location. Instead, we only duct taped the radios to covered posts. Despite our best social efforts, one WUR has already been stolen, though this happened near the end of the study and is unlikely to have meaningfully affected any results. A case that had hooks for locks, or zip-ties, or simply just screw holes would be much easier to mount and secure.

7.4 Sustainability

One of our key goals was to demonstrate that small local cellular networks can be financially sustainable. We investigated this by performing an analysis of expenses, revenues, and profitability over the period from February 11th to August 12th.

Item	**Shared**	**Non-shared**	**Frequency**
BTS RF Equip.	US$5,400	US$5,400	One time
BTS CPU	US$100	US$100	One time
VSAT Equipment	US$500	US$5,500	One time
Power	US$2,750	US$4,000	One time
Cabling	US$150	US$150	One time
Enclosures	US$100	US$100	One time
VSAT Service	US$15	US$130	Per month
Maintenance	US$100	US$100	Per month
Phone Number	US$0.65	US$0.65	Per number-month
Outgoing SMS	US$0.01	US$0.01	Per message
SIM Card	US$0.65	US$0.65	Per card

Table 7.3: System costs under shared-infrastructure and non-shared infrastructure models. The largest capital expenditure is the radio equipment, and the largest operational expenditure is the monthly cost of VSAT service for backhaul. Note that users paid US$10.30 up-front for one year of both their phone number and SIM card.

Table 7.3 describes both the capital and operational costs of the system. We consider two infrastructure models — one in which the network shares power and backhaul infrastructure with the community ("shared") and one in which there is no shared infrastructure ("non shared"). We compute the cost of shared infrastructure as 5% of the capital costs of our partner's VSAT and micro-hydro generator; non-shared infrastructure assumes independently purchasing a small VSAT and a solar system with sufficient battery capacity to power both the BTS and VSAT for two days without sunlight. The total capital cost associated with the shared model is US$9,000, and the non-shared model US$15,250. Maintenance covers wages for a part-time technician to maintain the system. Excluding VSAT service, other recurring costs are the same for both models.

The core components of the base station's cost are the radio and its associated amplifier and duplexer. However, as more companies enter the market for low-cost BTS radio equipment we expect the costs to fall. On the operational side, backhaul VSAT service was the largest expense. Although connecting a rural community cellular system to the outside world is challenging — typically expensive, low-performance satellite connections are the only option — our usage results show that out-of-network communications are the most valuable to users in Desa. We also faced particularly high recurring monthly costs for phone numbers due to buying numbers at retail prices rather than on the wholesale market. Because of this, WamenaCom required users to prepay one year's worth of charges for their number. This ensures that we do not operate at a loss. Users will need to pay again each year to maintain service.

On the revenue side (Table 7.4), the highest grossing service was out-of-network SMS, accounting for almost 90% of total revenue. Calls accounted for under 2% and the remaining revenue came from local SMS. Over the first 25 weeks of operation, the system grossed US$5001.70, or approximately US$830.65 per month. Our distribution of revenue per user is shown in Figure 7.8. Our median user spent US$2.01 per month, and 93% of users spent under US$15.00 per month. Together, the top 10% of users accounted for 43% of revenue. Our average revenue per user per month (ARPU) was US$4.44, compared to the ARPU

Usage Type	Amount	Revenue
Outside SMS	49,105	US$4,438.24
Incoming SMS	42,525	–
Local SMS	24,177	US$490.83
Free SMS	17,733	–
Local Call	12,965	US$72.63
Free Call	5,271	–
Credit Transfer	2,452	–
Total	**154,228**	**US$5,001.70**

Table 7.4: Number of usages of each service type for the period February 11th through August 12th. Revenue shown is for the network operator, and does not include reseller revenue.

Growth Rate	Shared	Non-shared
0 users / month	US$368.31	US$66.43
5 users / month	US$400.71	US$98.83
10 users / month	US$433.11	US$131.23
17 users / month	US$478.47	US$176.59

Table 7.5: Projected mean monthly profit over 5 years. Note that current growth rate is 17 users per month.

for prepaid users of the largest Indonesian carrier Telkomsel of US$3.60 [89].

Table 7.5 presents our projected monthly profit for the network under a variety of growth rates. While the network currently gains 17 users per month (excluding the first month of operation, which experienced rapid growth), we do not have sufficient data to project a long-term growth rate. Our model uses actual network revenues and costs for the first six months. For subsequent months, the projection assumes that current average monthly cash flow remains constant (i.e., usage patterns do not change) and that each new user will generate the current median monthly cash flow per user of US$2.01. We further assume that capital costs are financed by a 5 year loan at 12.4% APR [95], paid monthly. Thus, the analysis allows for complete system replacement (including BTS, solar equipment, batteries, and VSAT) every 5 years. The system's current monthly profit using these assumptions is US$368.31 and US$66.43 under the shared and non-shared model, respectively.

Our profit estimates, combined with the system's ARPU, are key results. Local cellular networks have latitude to set prices appropriately to the costs of serving their area. In contrast, incumbent national providers must set prices uniformly across all their users and compete on a large scale. As an example, Telkomsel charges 150 rupiah per SMS to Indonesia and our partners charge 900.

This may seem high, but other goods in Desa are similarly expensive: 1kg of sugar is 25000 rupiah (vs. 8000 in Jayapura), 1kg of rice is 20000 rupiah (vs. 5000), and 1kg chicken is 65000 rupiah (vs. 14000). These prices put the system just above the break-even point under the non-shared infrastructure model. Meanwhile, the ability to share common infrastructure (as we shared power and backhaul connectivity with our local partner) substantially reduces costs, particularly capital expenditure. As a result, shared

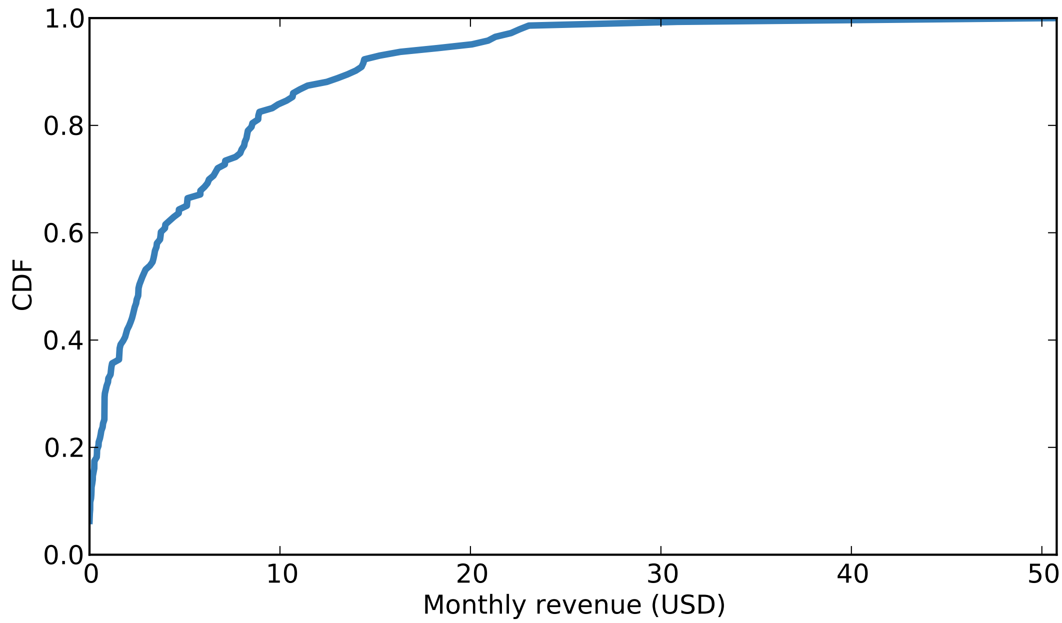

Figure 7.8: CDF of revenue per user per month. While 50% of users spent less than US$2.01/month, these users account for 7.7% of monthly revenue; the top 10% of users account for over 40% of monthly revenue.

infrastructure networks are able to take advantage of the lower cost to operate more profitably in rural environments like Desa.

Last, we note that the network is also profitable for the local resellers (Table 7.6). It is hard to gauge the relative impact of these funds on the sellers. However, as each seller conducts numerous transactions every day, it is likely a valuable income source.

	Primary	**Secondary 1**	**Secondary 2**
Amount Sold	US$3,955.93	US$987.07	US$2,852.12
Cost	US$3,955.93	US$1,091.95	US$2,932.45
Revenue	US$4,177.24	US$1,177.45	US$3,303.93
Avg. Markup	5.6%	19.3%	15.8%
Profit per Month	US$36.48	US$18.19	US$85.07
Sales per Month	9.7	71.1	290.4

Table 7.6: Income for credit resellers in dollars. Secondary 1 began sales on March 25 and Secondary 2 on April 3.

7.5 Subscriber Interviews

Throughout the deployment, we conducted numerous design discussions with users. We also conducted formal interviews with 6 users in Desa, and a group of three people in Jayapura, the capital of the Province of Papua, who were in contact with users in Desa. One of the people in Jayapura is a resident of Desa. The formal interview targets were entirely indigenous Papuans, with five males and four females. Ages ranged widely, with two interviewees being students at the local government school (not MS), two being older church organizers, and the remaining five being adults. The interviews were done in Bahasa Indonesia and conducted and translated by a fluent speaker on our team. We first note the uptake of the service:

Interviewer*: Among your friends in Desa, are many people using cards?*
Garet*: Oohhhh! Lots! You can see for yourself, right? Later, you can ask Matias, he is the one who sells [credits], there are already so many customers. The other day I saw the list, there are almost... 200 there. From the villages they also want to come buy.*

We also sought to answer key questions about our intervention:

- How did the intervention benefit you and/or the community?
- How was the network used?
- What problems have you seen, or expect to see?
- What's your perspective on who owns the network? and
- What would you change about the network?

7.5.1 Benefits

One message that permeated most of the interviews (8/9 of the formal) was the value of the network to the community. Users made sure to mention their appreciation of what was done and how important it is:

Benjamin*: We are here as customers of WamenaCom, here in Distrik Desa. We are very grateful for the installation of the WamenaCom network. This is a very big change, an extraordinary change, it can build communication for these two districts. So I now say thank you just for this.*

Garet*: People have been coming and saying "Desa has a special network." So we are proud. This is one of the things that can bring us to be... advanced. Not like others, others use Telkomsel, we have a different network. This is what I was thinking. I felt very happy. So I was thinking, 'I don't have to go to Wamena anymore to send SMS, to phone in Wamena, I can relax here in Desa, sitting in the forest, roasting cassava.*

Paulus*:So generally people are happiest because now they know that Desa already has a network [that works]. So they can SMS to family in Desa, from there they can SMS to Jakarta, to anywhere, because Desa kids, many of them are taking studies in Jakarta, in Sentani here, in Manado, in Jogja, so parents are happy, because, right away, if one needs anything one simply contacts one's parents by SMS. So they are happy...*

7.5.2 Use

Communication with family was incredibly important; every interviewee in Desa listed it as a primary use of the network. Leaving your home village for opportunities elsewhere is routine in Papua. Similarly, many non-Papuan Indonesians from nearby islands travel to Papua for work, including military service. These migrations, combined with strong familial bonds, led to this desire among our interviewees:

Garet*: At the time of the [large social gathering], I already got my card. I sent SMSs, but I didn't send here, because I miss my family over there, I try sending SMS over there, and... to Wamena. So, I try like this, and - eh! - smooth!*

Selina*: With the younger sibling in Jayapura, as it happens she is now studying. So, she is still studying, and she is also in her final semester... We checked the developments, how is it, after long in Wamena, maybe is she sick, or what, tired. We asked like this, and are there any shortages or not, we checked so we can complete her needs. That's why we checked with the family that is over there.*

Surprisingly, there was no discussion of business use of the network. This is likely a consequence of only interviewing Papuans, who own very little of the commercial infrastructure in Desa. Instead, two interviewees used it for local politics. As an example:

Selina*: [Important politician] is my older male relative, so I follow developments: "how are things going?" and so on. If there are obstacles, challenges, we can support in prayer, so it is smooth.*

Lastly, in-network communications were surprisingly common; two interviewees used them to organize activities in the village. This was unexpected, given the small size of Desa. As one of the youths noted:

Interviewer*: Why do you send SMS within Desa, even if it is close, you could speak directly, but you send SMS. What is the purpose?*
Letty*: Well, maybe, for example I'm at home, and I SMS a friend, or contact them, say "hey! please come pick me up, ok?". So we can go walking around.*

7.5.3 Problems

Although local communications no doubt benefited the youth; others in the community were less optimistic about these messages. Two of the interviewees in Desa and one in Jayapura expressed concern about the youth dating, and the network's role in that:

Benjamin*: People say having a network is a big change — some older people say it is not so good. This difference between good and that bad or negative, the issue is that someone who already has a [cell phone] and has purchased this WamenaCom card, maybe they can send SMS like "OK how about we meet up" and so on with girlfriends or boyfriends.*

and

Selina*: But on the negative side, meaning, but even that, it all depends and returns to us. We parents, especially we parents who have sons and daughters who are teenagers, who all this time maybe have been under tight supervision by their parents, but now with WamenaCom being here, yes, maybe parents are not watching closely and children pull...*

We note that this concern isn't about WamenaCom in particular, but rather an apprehension about mobile phones in general. Nor is it paranoia; some of our other interview targets spoke about contacting potential or existing romantic partners.

Lastly, we note that a few of the users did mention technical issues. One of the students was unable to communicate with anyone in Wamena, for some unknown reason, and this user was the only one to not speak highly of the network. There was also an issue one interviewee brought up with how users outside of the network understood our numbers, and sometimes confused them with spam:

Garet*: I call my brother. So for instance I needed money, so I had to go through, communicate with them. But they didn't believe, because only SMS, even though I already gave them the number clearly and my identity clearly, so I went straight there, so then they believed.*

7.5.4 Ownership

As our network is fairly untraditional, we were curious how the network ownership was viewed by the users. As the researchers spent months in the field and interacted with the population in Desa, we were well known. As one person noted:

Garet*: Well [local person] said, the one who owns this network, is from... over there, they, this, team for - what - this network, this - there is this guy who left already, right?*
Interviewer*: Kurtis?*
Garet*: Right.*

Two other interviewees attributed ownership to the NGO that housed the equipment: Misionaris Sekolahin and its leader Regis.

Interviewer*: [Who owns the network?]*
Selina*: All we know is, it is through uncle Regis. Through Regis, so... we are sure, because Regis is here, he is here together with us always, so we are sure, this is through Regis, so we believe, and we buy that card.*

and

Paulus*: Indeed, that day, I asked Garet, and he said, "Mr. Regis made this", he said that, Mr. Regis made it, not Telkomsel.*

Lastly, one of the interviewees seemingly believed that the network was literally owned by the community:

Herman: *Well according to me, this network, indeed, those who installed it, I don't really know, but it helps, so it is public.*
Interviewer: *Public, owned by the public?*
Herman: *Yes, it belongs to the public. Because at this time it is in Desa, so I think... people there use it, so it is publicly owned.*

7.5.5 Wants

We also asked the users what future features were most important to them. The answer universally was voice service. When asked to justify this desire, respondents tended to answer in rather vague terms which suggested an intangible, affective value of communicating by voice as opposed to text. As an example, a prominent local organizer for a new political party went so far as to say that Desa needed voice service so the "voice of this rural area can be heard." We hope to deliver this service shortly.

7.6 Comparison with Existing Carriers

We compare call logs from the Papua network against similar logs from a traditional telecommunication firm located in Sub-Saharan Africa. The Papua network had the goal of supporting local communications and community, so we evaluate how the this network differs from an existing carrier across three metrics of the strength of community structure among their users: local versus outbound ratio, transitivity, and the network average clustering coefficient.

7.6.1 Methodology

Data

As mentioend above, we collected complete call and SMS CDRs from the Papua network for a period of five months. These detail the source, target, time, and duration of every communication. Failed communications, such as busy or rejected calls, were recorded but ignored. We compare these with one month of similar call logs from 276 towers owned by a telco in Sub-Saharan Africa (SSA). These logs do *not* contain failed calls. All communications in the network are recorded, as well as the specific towers where the communications originated. SMS were not available. There were no price differences between in-tower and out-of-tower calls for SSA.

The CDRs were first split into one month periods and then converted to a set of per-tower directed graphs of communications. Each individual user on a tower is converted into a node in the graph. The weight of each edge between nodes is the number of individual calls and SMS sent between them. For example, if A calls B 3 times and SMSs B 5 times, there will be an edge from node A to node B of weight 8. These graphs were converted to undirected, unweighted graphs as needed in the following analysis.

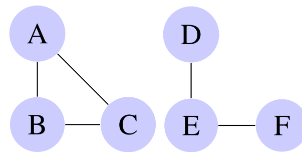

Figure 7.9: A triangle (left) and a triad (right). A triad is effectively a potential triangle.

Metrics

Our primary goal was to evaluate the "localness" of the Papua network in relation to a traditional telco. We begin by evaluating ratio of local to outbound communications on each tower. We then attempt to measure the clustering of local users to show that they were more tightly connected to each other. We do this through two clustering metrics, transitivity and average network clustering coefficient.

Local vs Outbound: The simplest measurement we have for "localness" is the ratio of local to outbound communications. This includes both calls and SMS, but excludes calls to local services. A more "local" network could be assumed to make a higher proportion of local communications. This metric is known to be naive, as a single user communicating primarily externally (or internally) would skew the results.

Transitivity: The transitivity of a graph is the ratio of the number of triangles versus the number of triads (see Figure 7.9). This method is well understood in social graph analysis [92], and roughly measures the degree to which the graph clusters together. We converted our graph to an undirected, unweighted graph for this analysis. Our graphs are per-tower; we do not see links between off-tower users. As such, we also remove all non-local nodes from the graph and only compute the transitivity of the local graph. High transitivity on the local graph implies the local users cluster together strongly, a good metric of the "localness" of the tower.

Network Average Clustering Coefficient: The clustering coefficient (CCE) [47] is a measure of the degree to which nodes in a graph tend to cluster together. The local clustering coefficient is measured per node and is the ratio of the number of its neighbors who are connected to each other (triangles) versus the number of possible connections among its neighbors (triads where the selected node is the center). It is similar to transitivity, but is analyzed on a per-node, rather than per-graph basis. In order to make this a tower-wide metric, we compute the network average clustering coefficient [57], which is just the average CCE among all the nodes in the graph. As with transitivity, we first convert the graph to an unweighted, undirected graph and we only measure the CCE of local (on-tower) nodes. NCCE, like transitivity, roughly measures the degree to which the nodes cluster together, but is more influenced by low-degree nodes.

7.6.2 Results

Local vs Outbound

Figure 7.10 shows the local vs. outbound ratio of all communications for each tower in SSA and each of the five months of the Papua network. The Papua network is not particularly different from the SSA network; each month of the Papua tower's usage is close to the average for the SSA towers. Welsh's two-sample t-test [94] confirms this, as we are unable to show a significant difference in the local versus outbound ratio for the SSA and Papua towers ($p < 0.074$).

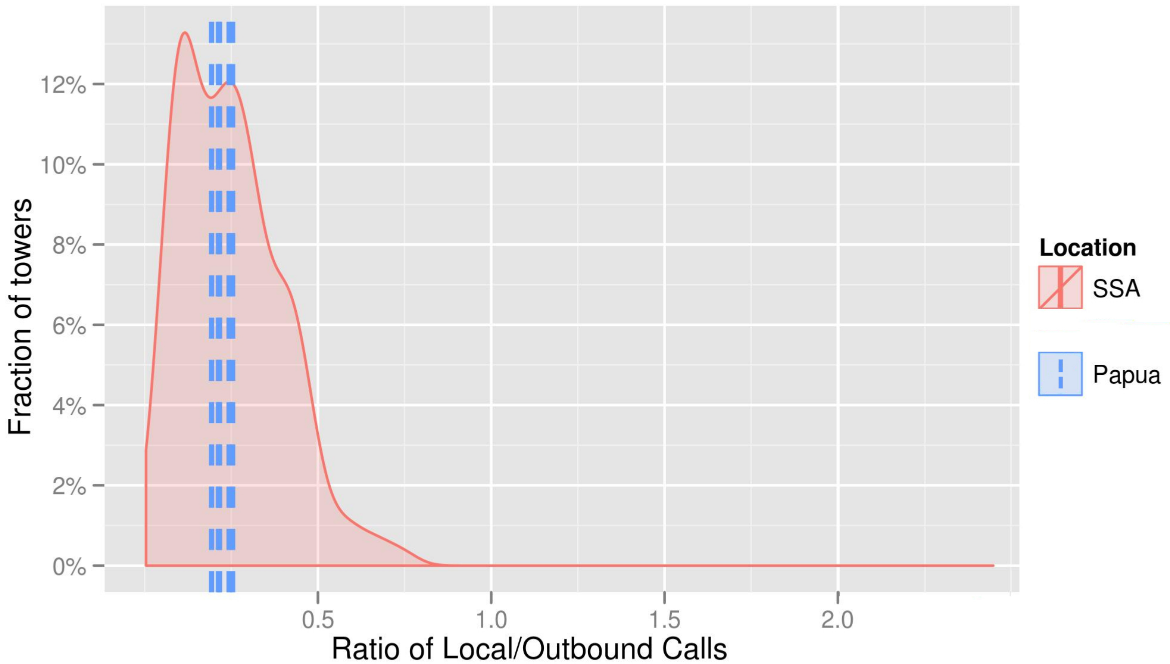

Figure 7.10: The distribution of the local versus outbound ratio for each month of the Papua network and the SSA towers. SSA towers comprise the distribution with each month of the Papua tower being represented by an overlaid line.

Transitivity

Figure 7.11 shows the transitivity of each measured tower. In this case, the transitivity of every SSA tower is less than the lowest transitivity of any month of the Papua network's operation. The transitivity of our highest month (month 1: 0.291) is more than double that of the highest SSA tower (0.125). In this case, the transitivity of the Papua network is statistically significantly different from that of the SSA towers ($p < 0.00002$).

Network Average Clustering Coefficient

Figure 7.12 shows the network average clustering coefficient (NACC) for each measured tower. Again, every SSA tower's NACC is lower than the lowest month of the Papua network's NACC (0.20). Also, the highest Papua NACC (Month 1: 0.27) is more than double the highest SSA (0.12). Lastly, the tower NACCs for Papua and SSA are statistically significantly different ($p < 0.0002$).

These results suggest that usage of the Papua network is significantly different than that of traditional cellular networks. Though it is not possible to attribute the difference to a particular design decision, the Papua network was designed to encourage stronger local communities, and this goal was achieved.

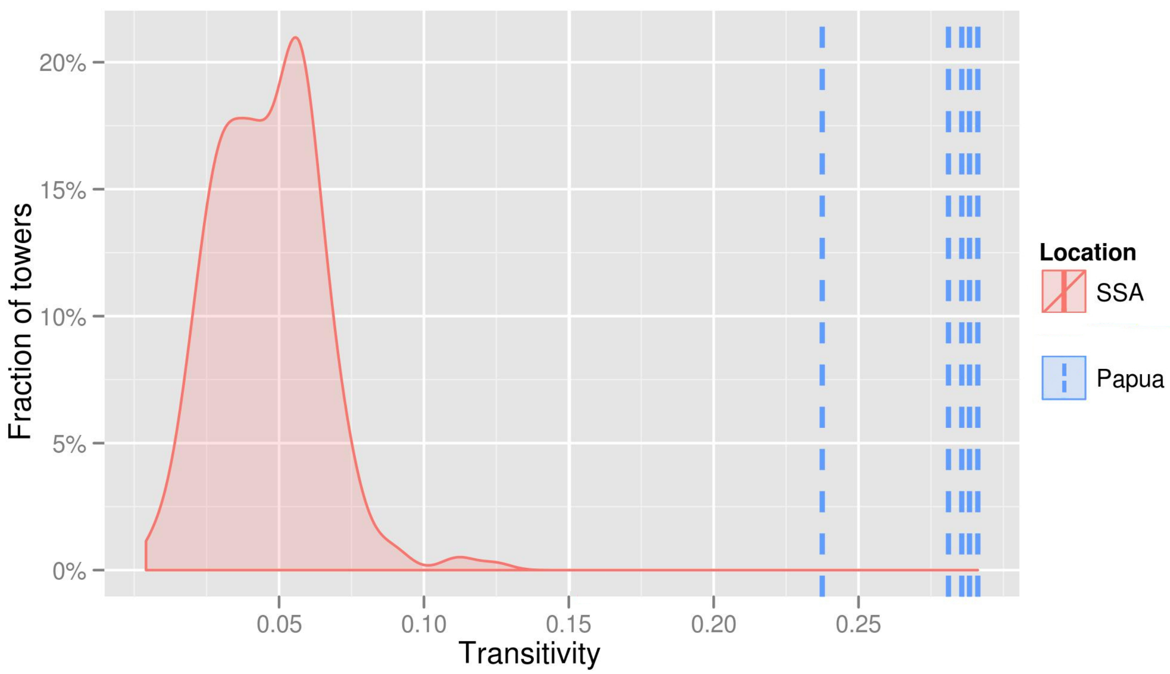

Figure 7.11: The distribution of the transitivity for each month of the Papua network and the SSA towers.

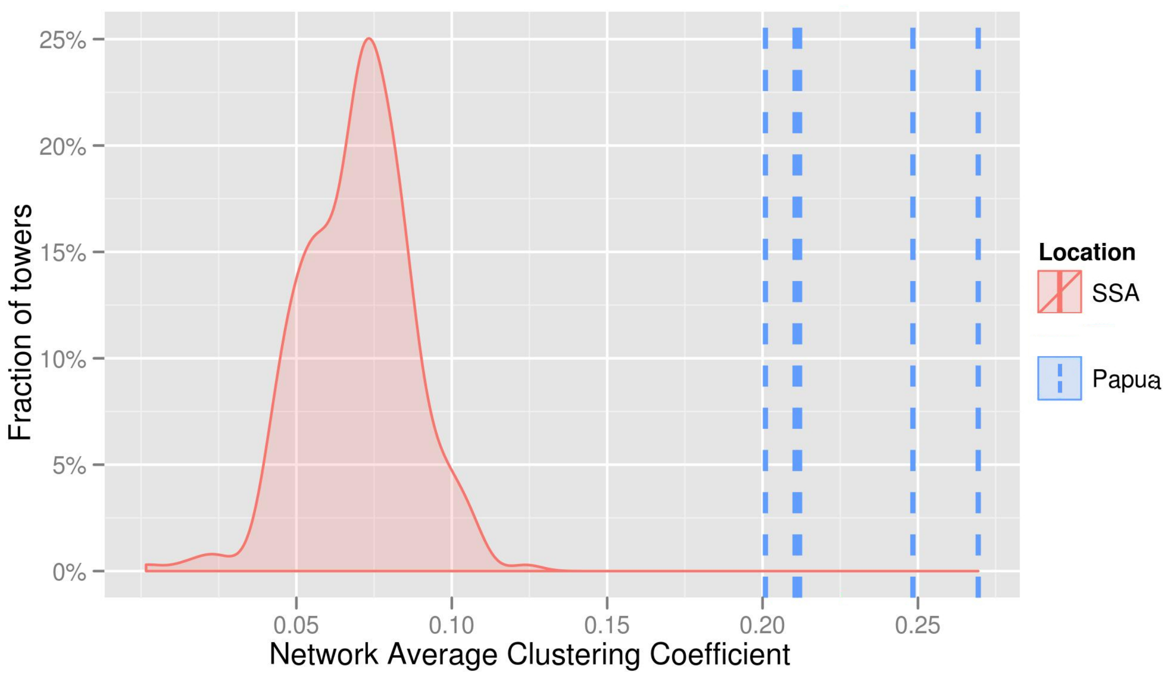

Figure 7.12: The distribution of the network average clustering coefficient for each month of the Papua network and the SSA towers.

Chapter 8

Risks

Chapter 7 demonstrated that our community cellular network in Papua is sustainable, highly used, and beneficial to the community, but this does not mean that this would always be the case. Our experience exposed numerous challenges and conflicts that will likely emerge in other local cellular installation.

8.1 Licensing

GSM is licensed spectrum, meaning that to operate a GSM tower legally you must have government oversight and approval. Typical spectrum licenses are country-wide and cost billions of dollars. Though we have tacit approval from high-level members of the Indonesian government to operate the BTS, we do not have an actual license to operate in Indonesia. This means our installation is not, technically speaking, legal. Although the assurances we've been given reduce our risk of government interference, it is still present. Operating in a rural area also affects this risk. Papua, like many rural areas, has lax enforcement of wireless spectrum laws. In Desa, key authority figures from the police and military have expressed their acceptance of the installation, and local members of the security forces are among its most prolific users. As a result, our installation exists in an administrative grey zone, having achieved a degree of legitimacy even though the legal framework does not yet encourage decentralized mobile telephony networks. In general, spectrum licensing laws and enforcement vary greatly across nations, and any operator attempting to run a GSM network should be aware of the regulations in their area.

8.2 Interconnect

Though earlier work [29, 76, 96] explored local-only networks, our usage indicates that out-of-network connectivity is a critical need for any local telco. Unfortunately, this is often not an easy task. In the cellular world, interconnection is a wild and woolly business, and who you know is as important as the services you bring. Being a small, independent rural operator is an extreme limitation in such an environment.

Originally, we had hoped to connect directly to a local Indonesian carrier. Meeting with these operators brought a lot of hope but few results. We tried tens of other SMS-routing companies, but none seemed to serve Indonesia well. It turns out that each company maintains multiple potential routes to Telkomsel, each with differing properties. Some could send to Telkomsel customers but not receive from them; others receive but not send. Still others could do both but the return phone number would be mangled or randomized to avoid spam filters.

We eventually ended up with Nexmo, the only company we found who could consistently route to Indonesia. Each provider connects different countries, so it is likely that an exhaustive search for the correct partner will be required whenever deploying in a new country. We hope that, given enough time and a large enough market, service providers will eventually value their *breadth* of countries enough to provide solid service across the board.

8.3 Theft

Local ownership is supposed to reduce theft; stealing from people in your community who you know (and who know you) is much more dangerous than stealing from faceless corporations. We found this to be generally the case; no equipment was broken or stolen, nor was there any evidence of tampering. However, we had one instance of theft: a few users discovered how to generate counterfeit credits. The bug was quickly discovered and shut down. A similar situation happened again later, when a few subscribers discovered that you could send a message so long as you had *any* credit in your account; you didn't need the full amount. Users would send each other small amounts of credit (e.g., 10 rupiah) and then make an outbound communication, which would zero their account. This bug was also quickly resolved.

Though we have yet to investigate these incidents in depth, our feeling is that stealing *digital* items has less social stigma in our community. Our NGO partner has had to deal with similar issues in the past, such as tech-savvy users who were stealing Internet access. As such, theft of network service is likely to be a long-term concern for operators of networks like ours.

Chapter 9

Discussion

9.1 Community Cellular

9.1.1 Local Implications

Although we describe our system as local, we wish to emphasize that it is not run by indigenous people. Both operating organizations, WC and MS, are run by non-Papuans; WC is run by an American technologist in Wamena. MS is run by an American couple who have lived in Papua for over a decade. We describe our project as "local" based on the fact that the tower is located within the community and operated by community members. Similarly, the operators are "locals" based on their long-term presence in, involvement in, and knowledge of Desa.

Being local does not mean that the technology or intervention will always benefit the community. Far from it – even "local" agents may be as destructive as any other. As an example, this system was deployed in partnership with an existing Papuan-run Christian mission. The long colonial history of Christian missions in Papua should not be ignored. Part of our work reinforces existing colonial power structures, potentially at the expense of indigenous ones. Though we believe very strongly in the good intentions of our missionary partners, this system could cause harm. Our goal in this work is to make cellular systems more accessible from a cost and technology perspective. Our belief is that this will enable more local, community-based actors to provide connectivity. Eventually, any group with sufficient will and a small amount of start-up capital will be able to run their own networks; including groups with less resources and organizational capacity then WC and MS.

9.1.2 Policy Support

Substantive policy changes could reduce the risk for rural entrants and encourage innovation in small-scale telephony networks. We hope that our work, showing meaningful gains for rural users from local cellular networks, will positively benefit ongoing spectrum licensing reform movements. One approach is that of The Netherlands, which has set aside a portion of the DCS1800 band for unlicensed low-power GSM networks [2]. We have proposed an alternative mechanism for regulating community cellular networks called Nomadic GSM [41]. This technique utilizes advances in cognitive radio design to allow for a second class of cellular operators that can peacefully coexist with existing operators.

The relatively low margins of small-scale cellular networks make them particularly sensitive to high interconnection fees that incumbent providers use to price out new competitors. Another (perhaps unintentional) barrier to entry for small operators is zealous blocking of SMS by incumbent providers to reduce spam; opaque blocking criteria and the difficulty of communicating with incumbent providers compounds this issue. Regulation to ensure fair interconnection with incumbent providers would be a major benefit for rural operators. These regulations exist in the US and are a key reason for the high number of rural operators.

Our results suggest that some portion of rural connectivity goals can be met by local entrepreneurs with small or even no subsidies; this has interesting implications for universal service funds [83], a common mechanism for subsidizing rural access. Utilizing these funds to encourage local entrepreneurs and reduce universal service obligations could make the other policy changes necessary to support small-scale cellular networks more palatable to incumbent providers.

9.1.3 Generalizability

Papua is a unique place in the world; there are few places with so many geographical obstacles to infrastructure development, so much natural wealth, and such wide cultural diversity. Because of this, we cannot know if our results will generalize to other parts of the world. The specific technologies (e.g., VSAT, wifi, hydro, solar) and specific social and technological structures (e.g., prepaid, prices, credit selling) used will no doubt depend on factors related to the target location. However, our design is based on a fundamental principle that *does* generalize: empowering local agents to solve their own communication problems with the materials and knowledge available in their community. For this reason, we believe our designs and technology will work in other areas, and plan to demonstrate that in future work.

9.1.4 Alternative Models of Local Ownership

Our deployment follows a simple model of local ownership; our partner organizations are led by individual immigrants to Papua. These leaders control the organizations completely; other members are employees or customers. Although this is a limitation of our study, we believe our system could work with other organizational structures as well, such as cooperatives or franchisees. Grameen Telecom [4] used a franchise model, but the phone ladies only had to buy phones and not base stations. Our franchisees would need more capital and technical capacity, but would also extend coverage (unlike Grameen).

Cellular networks, like many other kinds of infrastructure, are vital to the whole community; 187 people in Desa are customers of our service. We envision an alternative model: a cooperative network [30] in which each member both owns and pays into the system. Rhizomatica [76] is an example: a community-owned OpenBTS-based network in Mexico. We envision services in the network that would support this vision of participatory governance; for example SMS voting and broadcasting community meetups. Our work is just a first step demonstrating the basic feasibility of small-scale local cellular networks, we plan to continue to explore supporting other models of ownership as we move forward.

9.1.5 Evaluation Limitations

Community cellular networks, by design, bring coverage to new areas. This makes it difficult to compare these networks, as there are simply too many variables. For example, the location, culture, and population size all differ between our CCNs and the SSA networks. As such, it is impossible for us to say that our designs are the factor that caused the CCN tower usage to differ so much. Instead, we can only say that these networks *are different*.

9.2 Virtual Coverage

9.2.1 Data Services

An obvious critique of virtual coverage is that we do not support data services on the base stations. This is not fundamental to the design, but rather a use case we have not yet needed to implement as the Papua network has no data access. There are numerous ways to provide such a service. For example, if a user requests data service while the tower is idle, the BTS can handle this in a similar fashion to placing a phone call: the tower is sent a wake-up burst (via either phone handset or wake-up radio), wakes up and connects to the Internet. If idle time is significantly reduced by data requests (e.g., by data-channel apps seeking updates from the web), users can be incentivized to turn off these features when negotiating service rates with the provider.

9.2.2 Mobility

In the design of virtual coverage, we explicitly avoid the issue of mobility. Our equipment is designed to create "islands" of coverage, simplifying the architecture dramatically. It is assumed that a rural virtual coverage cell will not intersect any other covered region. However, as this work moves forward, we recognize that the issue of mobility should be addressed.

The GSM tower broadcasts not only its own ID, but also those of nearby towers. This helps in two ways. First, a handset can try waking up towers in succession to increase its chances of successfully waking a BTS, at the cost of extra delay. Second, during a call the handset could try to wake-up nearby towers proactively, either due to low signal from the current tower or just in case. Once awake, hand-off to neighboring towers works as usual. Finally, on higher-end phones, a GPS-indexed database could inform the handset of exactly what tower(s) to wake in a specific location.

9.2.3 Inter-operation with existing infrastructure

The GSM specification assumes constant coverage; a by-product of a system designed for developed, urban areas with strong power infrastructure. Virtual coverage changes this, turning cellular towers into dynamic agents. The interaction of these networks is complex.

Our system handles this already: Modified handsets *always* connect and call through existing static systems if possible: we are only capable of waking a tower if we are not attached to any existing tower. This is done primarily to save power; the static tower is on (whether or not it fields an additional call), but we'd prefer to keep our local tower idle. Similarly, if the other tower is dynamic but active, we'd prefer to send two calls through the powered tower, rather than waking an idle tower.

9.2.4 Security

The system, as designed, has no mechanism for authenticating users before they wake the BTS. This means that one dedicated attacker could launch a power-based denial of service (DOS) attack by constantly sending the wake-up bursts to the tower. Though this would obviously not be to the benefit of most users, one could easily imagine situations where this would happen. For example, a dedicated user may DOS the tower to leverage an information asymmetry they have developed.

The two mechanisms for waking up the BTS have different trade-offs for preventing such DOS attacks. For the WUR it is impossible to identify the user sending the wake-up burst, as the device is totally separate from the phone. Instead, we could add a cost to the use of the physical radio; perhaps by placing it in a phone-booth-like structure and charging a fee. This way, DOSing the tower could be made prohibitively expensive.

For the modified handset, we can identify the users by changing the protocol slightly. Currently, users send the burst, camp, and then wait. They must choose separately to make a call. We can instead enforce that the user must make a call (or any communication) immediately after camping to the station. This would allow us to know who did the waking. Unfortunately, this is still susceptible to another, similar attack; a user could send the wake-up burst and immediately pull the battery. In this case, we could modify the BTS to note there has been no traffic and immediately re-enter the idle state, reducing the impact of the attack.

Lastly, it is possible to send identifying information, such as the IMSI (unique SIM ID), in the burst message. This would allow us to charge users for waking the tower and prevent DOS attacks. We leave this to future work.

9.2.5 User Behavior

Heimerl et al. [45] found that cellular subscribers in rural areas often make dramatic changes to their own behavior in order to make efficient use of their networks. They would commonly climb trees, memorize coverage patterns, and plan their travel in order to communicate at critical points.

Our system is ripe for such optimization. For instance, if we were able to charge users for waking the BTS, it is likely they would begin to batch their calls in order to make the best use of the network availability. An alternative mechanism for this would be to divide the cost of the power over all users concurrently communicating.

Similarly, if we only enabled virtual coverage at night (as planned), users would potentially shift their calling to daytime periods. These optimizations are hard to predict before deploying a system, and as such we hope to investigate how to support them in future work.

9.2.6 Evaluation Limitations

The primary limitation of our evaluation of virtual coverage is the lack of qualitative data about the use of the network during the "night" period. We have no data concerning the relative value of these communications to the users or the community, only quantitative data on the actual use of the network. This was an unfortunate side effect of the potential nature of talks late at night; they traditionally (at least to us) are more private and sensitive than similar conversations during the day. In short, we felt uncomfortable asking users about their discussions that took place at 03:00. In the future, we hope that this will be resolved through better interpersonal relationships with users so that they might be willing to trust us with this information. As it stands, we were not there yet.

A second limitation concerns the generalizability of the result. The night communication patterns and needs of the people of Desa are likely to differ from those of other communities. For example, in Desa there is no grid power at night; this means that relatively few people are awake and potentially limits the amount of local communication at night. Similarly, Indonesia spans multiple time zones. This likely increases the chances of incoming SMS earlier in the night at Desa (GMT +9), when people are still out and about in Jakarta (GMT +7). Each potential deployment of virtual coverage should assess their own specific communication needs before deploying and configuring the system.

9.2.7 Holistic Power Savings

We note that, despite the BTS power savings shown in this work, the eventual goal is to reduce the power draw of the entire system. Throughout our earlier virtual coverage work [42] and this work, there is little discussion of the other pieces of infrastructure required for operation. In our case, that would be the VSAT system. Altobridge [5] has developed modifications to the core network protocols to minimize VSAT usage. We believe that it is possible to implement similar power saving techniques in the VSAT, continuing to lower the total power draw of the system.

Chapter 10

Future Work

10.1 Community Cellular

10.1.1 Policy

As mentioned previously, regulation is a core problem in community cellular network. Running your own cellular network is currently illegal in most countries and the only alternative is an experimental license that often requires nonprofit operation. We've been investigating alternative mechanisms for licensing small-scale rural cellular operators and have devised a system called Nomadic GSM [41]. While this is still under active development, we hope to pilot this technology soon and demonstrate its feasibility.

10.1.2 Deployments

The deployment of a community cellular network in Desa is just the beginning. We have already begun installation of a second site on a nearby mountain. The site should cover multiple communities from the mountaintop and is powered using a solar array installed in parallel.

Outside of Indonesia, we are planning to deploy multiple sites in the Philippines in partnership with local telecoms (e.g., Smart) and the University of the Philippines. These deployments will cover the multitude of islands without coverage.

Lastly, other organizations are similarly piloting community cellular installation. Rhizomatica [76] has had a network in the community of Talea in Oaxaca, Mexico for nearly a year and has installed different equipment and software in two other nearby communities with a long-term experimental license from the Mexican government. We hope to continue to work with them to create a shared platform for community operated networks.

10.1.3 Services

Our installation in Papua provides a number of basic services: local voice/SMS and outbound SMS. Traditional telcos also provide global voice and data access. We hope to enable these services on the Desa network. Voice is the most requested feature addition, but a number of subscribers also have smartphones and have asked for data. We've very interested in the impact of generic Internet access on the community and how it might affect the population.

We also hope to develop a number of value-added services for the Desa community specifically. One target is a personal audio/SMS "website" where users can post songs or other information they value. This is conceptually similar to the Spoken Web [59] or Avaaj Otalo [70]. We believe the fact that local calls are very cheap (having zero marginal cost for operators) can enable many of these services that are fundamentally too expensive when operating on a traditional telco.

10.2 Virtual Coverage

10.2.1 Radio Design

As mentioned earlier, the project would have benefited from better industrial design: a WUR that only broadcast the wake-up burst for a limited duration for each button press (rather than broadcasting continuously) and better discouraged theft. The team had little experience with user-facing hardware design, and the lesson is well learned. Our next deployment will use better designs.

10.2.2 Smarter Wake-Up Protocol

As mentioned in related work [42], it is possible to use a smarter "wake up" burst when moving the BTS out of idle mode. We could program a unique ID into the actual burst. This could be used to limit the number of transmissions from any device by ignoring repeat bursts, potentially disincentivizing the above "holding the button down" issue.

10.2.3 Smarter Power Amplifier

Virtual coverage uses WURs, deployed throughout the community, to enable network access during times of relative inactivity. The users walk to the site of a WUR to "wake" the BTS. This has worked successfully in the Desa deployment site. However, it may be possible to eliminate the need for these devices entirely while still supporting legacy phones and reducing BTS power draw.

We are currently investigating the design of a smarter power amplifier that does duty cycling on a millisecond basis. This would allow us to put it to "sleep" on a per-slot basis, instead of for minutes at a time as we do now. GSM is time-division multiplexed, meaning that each time slot is a different logical channel. When the BTS is operating at low utilization, it may be possible to put the amplifier to "sleep" for just unused channels, while keeping it "awake" for important channels like the beacon channel. This would allow us to build a power-proportional BTS that scales its power consumption linearly based on load. Handsets would see the expected channels (Beacon, Synchronization, and Frequency Correction) and still be able to camp to the BTS and make or receive calls even if most other channels are inactive.

The overall power savings of this design are likely to be comparable to virtual coverage; though we would still waste power broadcasting control information at times of zero utilization, we'd save more power during times of light utilization where virtual coverage requires the entire power amplifier be awake to handle just one call. This is a promising research agenda and we hope to have a prototype by our next field deployment.

10.2.4 Price Changes

We are curious how changes in service prices might affect the usage of the WUR. There's two lines of thought on the topic. First, communications should be *more* expensive at night as they require the use of more power infrastructure. Second, communications should be *less* expensive at night as the backhaul network is less congested. We are unsure which is correct, and so we kept prices the same as during daytime operation in this study. It also simplified the analysis. We plan to evaluate the data and eventually attempt to vary the prices to more efficiently use the network and power resources, perhaps changing the prices based on the amount of power in the batteries or the level of network congestion.

Chapter 11

Conclusion

The goal of this research is to bring telecommunications to the remaining 700 million people on Earth currently without coverage. We attempt to do this by addressing the two core reasons for this lack of coverage, expensive installations and "top-down" networks, with two innovations. The second, virtual coverage, is a smarter duty cycling system for cellular networks. It reduces the total power draw by over 50% in many networks, reducing the largest capital expenditure in rural cellular networks: power. The second, community cellular networks, is small-scale, locally operated, "bottom-up" cellular networks. These networks can be operated by individuals, NGOs, or other members of the community and customized to meet community needs and fulfill community goals.

We demonstrated the value of these two innovations through a real-world deployment in rural Papua, Indonesia that has handled over 100000 communications and generated over US$5000 in revenue in six months of active operation. The network was deployed in partnership with a local education NGO and a for-profit wireless Internet service provider, using the NGOs existing micro-hydro and satellite installations. In the Papua cellular network, Virtual coverage reduced the night power draw by 56.6%. The network is also sustainable, with the sizable profits staying within the community. We interviewed subscribers, finding that most communications were with family and friends in other parts of Indonesia. Lastly, we compared the usage of our installation with similar installations from a nation-scale firm in sub-Saharan Africa, finding that our network was more community focused.

These results validate the community cellular model: bringing telecommunications to rural areas through empowering local people and organizations to do it themselves. We plan to continue this work in a variety of ways. We have initial designs for a new power amplifier that can power cycle on a per-slot basis, potentially enabling virtual coverage without the necessity of a wake-up radio. We're looking into mechanisms for regulating community cellular networks so they interoperate peacefully with existing firms. Lastly, we have numerous deployments planned in the near future, with contacts in the Philippines, Pakistan, and more installations in Papua. As this work progresses and matures, it has the potential to bring communications to the millions of people currently without coverage.

Acknowledgments

This research was funded by the Blum Center for Developing Economies and the Development Impact Lab (USAID Cooperative Agreement AID-OAA-A-12-00011), part of the USAID Higher Education Solutions Network.

I first want to thank my thesis committee (Eric Brewer, Tapan Parikh, Jenna Burrell, and Ali Niknejad) for taking the time out to provide advice during my qualification exam, attend my dissertation talk, and read this document. Thanks again!

I had a lot of technical and social support for this work including from my colleagues Shaddi Hasan, Kashif Ali, Jacob Nerenberg, Brian Gawalt, Joshua Blumenstock, Peter Bloom, Stephen Dawson-Haggerty and Andrew Krioukov; my advisors Eric Brewer and Tapan Parikh; Range Networks, particularly David Burgess and Harvind Samra; Ciaby; Onno Purbo; David Haag; anonymized people of Desa; Jerry Ng; Nathan Eagle; and Amy Wesolowski. Thanks to all of you for all of the help.

Numerous individuals also personally supported me in this research including my advisors, Eric Brewer and Tapan Parikh; my parents, Rose and Randy Heimerl; my sisters Julie and Michelle (and Ben too); Kelly Buchanan; Orianna DeMasi; Yoe-One Ariestya Niovitta; Kuang and Alice Chen; Heather Underwood; Matei Zaheria; Peder Reiland; the Super Awesome League; the developers of Starcraft 2, and anyone I've ever played a board game with. Thanks to all of you as well.

Bibliography

[1] M. Adeyeye and P. Gardner-Stephen. The Village Telco Project: A Reliable And Practical Wireless Mesh Telephony Infrastructure. *EURASIP J. Wireless Comm. and Networking*, 2011:78, 2011.

[2] Agentschap Telecom. Regeling gebruik van frequentieruimte zonder vergunning 2008. `http://wetten.overheid.nl/BWBR0023553/volledig/geldigheidsdatum_23-04-2013`, April 2013.

[3] A. Agüero, H. de Silva, and J. Kang. Bottom of the Pyramid Expenditure Patterns on Mobile Services in Selected Emerging Asian Countries. *Information Technologies & International Development*, 7(3), 2011.

[4] Q. Alam, M. A. Yusuf, and K. Coghill. Village Phone Program, Commodification Of Mobile Phone Set And Empowerment Of Women. *J. Community Informatics*, 5(3-4), 2009.

[5] Altobridge. `http://www.altobridge.com/`.

[6] A. Anand, V. Pejovic, E. M. Belding, and D. 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, pages 180–189, Atlanta, Georgia, 2012. ACM.

[7] K. Bali, S. Sitaram, S. Cuendet, and I. Medhi. A Hindi Speech Recognizer for an Agricultural Video Search Application. In *Proceedings of the 3rd ACM Symposium on Computing for Development*, ACM DEV '13, pages 5:1–5:8, New York, NY, USA, 2013. ACM.

[8] L. A. Barroso and U. Hölzle. The case for energy-proportional computing. *Computer*, 40(12):33–37, Dec. 2007.

[9] G. Bernardi, P. Buneman, and M. K. Marina. Tegola tiered mesh network testbed in rural scotland. In *Proceedings of the 2008 ACM Workshop on Wireless Networks and Systems for Developing Regions*, WiNS-DR '08, pages 9–16, New York, NY, USA, 2008. ACM.

[10] S. Bhaumik, G. Narlikar, S. Chattopadhyay, and S. Kanugovi. Breathe To Stay Cool: Adjusting Cell Sizes To Reduce Energy Consumption. In *Proceedings of the First ACM SIGCOMM Workshop on Green Networking*, Green Networking '10, pages 41–46, New Delhi, India, 2010. ACM.

[11] J. Bicket, D. Aguayo, S. Biswas, and R. Morris. Architecture and evaluation of an unplanned 802.11b mesh network. In *Proceedings of the 11th Annual International Conference on Mobile Computing and Networking*, MobiCom '05, pages 31–42, New York, NY, USA, 2005. ACM.

[12] J. E. Blumenstock. Using mobile phone data to measure the ties between nations. In *Proceedings of the 2011 iConference*, iConference '11, pages 195–202, New York, NY, USA, 2011. ACM.

[13] BPS-Statistics Indonesia. `http://www.datastatistik-indonesia.com/`. Retrieved 4/2013.

[14] E. Brewer, M. Demmer, B. Du, M. Ho, M. Kam, S. Nedevschi, J. Pal, R. Patra, S. Surana, and K. Fall. The case for technology in developing regions. *Computer*, 38(6):25–38, 2005.

[15] J. Burrell. Livelihoods And The Mobile Phone In Rural Uganda, January 2008. Report for the Grameen Foundation.

[16] L. Butt. "Lipstick Girls" and "Fallen Women": AIDS and Conspiratorial Thinking in Papua, Indonesia. *Cultural Anthropology*, 20(3):412–442, Aug. 2005.

[17] K. Chebrolu, B. Raman, and S. Sen. Long-distance 802.11b links: Performance measurements and experience. In *Proceedings of the 12th Annual International Conference on Mobile Computing and Networking*, MobiCom '06, pages 74–85, New York, NY, USA, 2006. ACM.

[18] D. Clark. The Design Philosophy of the DARPA Internet Protocols. In *ACM SIGCOMM Computer Communication Review*, volume 18, pages 106–114. ACM, 1988.

[19] D. D. Clark, J. Wroclawski, K. R. Sollins, and R. Braden. Tussle in Cyberspace: Defining Tomorrow's Internet. In *Proceedings of the 2002 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, SIGCOMM '02, pages 347–356, New York, NY, USA, 2002. ACM.

[20] Competitive Carriers Association. `http://competitivecarriers.org/`.

[21] H. De Soto et al. *The other path*. Harper & Row New York, 1989.

[22] Deloitte. Economic Impact of Mobile in Bangladesh, Malaysia, Thailand, Pakistan, Serbia, and Ukraine, January 2008. A report prepared for Telenor ASA.

[23] I. Demirkol, C. Ersoy, and E. Onur. Wake-up receivers for wireless sensor networks: benefits and challenges. *Wireless Commun.*, 16(4):88–96, Aug. 2009.

[24] J. Donner. The Rules of Beeping: Exchanging Messages Via Intentional "Missed Calls" on Mobile Phones. *Journal of Computer-Mediated Communication*, 13(1), 2007.

[25] B. V. d. Doorn, W. Kavelaars, and K. Langendoen. A prototype low-cost wakeup radio for the 868 mhz band. *Int. J. Sen. Netw.*, 5(1):22–32, Feb. 2009.

[26] N. Eagle, Y.-A. de Montjoye, and L. M. A. Bettencourt. Community computing: Comparisons between rural and urban societies using mobile phone data. In *Proceedings of the 2009 International Conference on Computational Science and Engineering - Volume 04*, CSE '09, pages 144–150, Washington, DC, USA, 2009. IEEE Computer Society.

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

[28] R. Flickenger, C. E. Aichele, C. Fonda, J. Forster, I. Howard, T. Krag, and M. Zennaro. Wireless networking in the developing world. 2006.

[29] V. Gabale, R. Mehta, J. Patani, K. Ramakrishnan, and B. Raman. Deployments Made Easy: Essentials of Managing a (Rural) Wireless Mesh Network. In *Proceedings of the 3rd ACM Symposium on Computing for Development*, ACM DEV '13, pages 10:1–10:10, New York, NY, USA, 2013. ACM.

[30] H. Galperin and F. Bar. The Microtelco Opportunity: Evidence from Latin America. In *Information Technologies and International Development*, volume 3, Winter 2006.

[31] B. Gavish and S. Sridhar. Economic aspects of configuring cellular networks. *Wirel. Netw.*, 1(1):115–128, Feb. 1995.

[32] Global Energy Network Institute. `https://www.geni.org/`. Retrieved 8/2012.

[33] S. Gobriel, C. Maciocco, and T.-Y. Tai. Long idle: Making idle networks quiet for platform energy-efficiency. In *Systems and Networks Communications (ICSNC), 2010 Fifth International Conference on*, pages 340 –347, aug. 2010.

[34] Government of India. Official Results for 3G Spectrum Auctions: Department of Telecommunications. May 2010.

[35] GSM Association. 3 Billion GSM Connections on the Mobile Planet. http://www.gsmworld.com/newsroom/press-releases/2008/1108.htm.

[36] GSM Association. Powering Telecoms: East Africa Market Analysis Sizing the Potential for Green Telecoms in Kenya, Tanzania and Uganda. http://www.gsma.com/mobilefordevelopment/powering-telecoms-east-africa-market-analysis. Retrieved 2/2013.

[37] GSM Association. Women and Mobile: A Global Opportunity. http://www.gsma.com/mobilefordevelopment/wp-content/uploads/2013/01/GSMA_Women_and_Mobile-A_Global_Opportunity.pdf.

[38] L. Gu and J. A. Stankovic. Radio-triggered wake-up for wireless sensor networks. *Real-Time Syst.*, 29(2-3):157–182, Mar. 2005.

[39] C. Hartung, Y. Anokwa, W. Brunette, A. Lerer, C. Tseng, and G. Borriello. Open Data Kit: Tools to Build Information Services for Developing Regions. In *Conference on Information and Communication Technologies for Development (ICTD)*, 2010.

[40] S. Hasan, Y. Ben-David, G. Fanti, E. Brewer, and S. Shenker. Building Dissent Networks: Towards Effective Countermeasures against Large-Scale Communications Blackouts. In *Proceedings of the 2nd USENIX Workshop on Free and Open Communications on the Internet*, FOCI'13, Washington, D.C., 2013. USENIX Association.

[41] S. Hasan, K. Heimerl, K. Harrison, S. Roberts, A. Sahai, and E. Brewer. Gsm whitespaces: An opportunity for rural cellular service. In Submission.

[42] K. Heimerl, K. Ali, J. Blumenstock, B. Gawalt, and E. Brewer. Expanding Rural Cellular Networks with Virtual Coverage. In *Proceedings of the 10th USENIX Symposium on Networked Systems Design and Implementation*, NSDI'13, Lombard, Illinois, 2013. USENIX Association.

[43] K. Heimerl and E. Brewer. The Village Base Station. In *Proceedings of the 4th ACM Workshop on Networked Systems for Developing Regions*, NSDR '10, pages 14:1–14:2, San Francisco, California, 2010. ACM.

[44] K. Heimerl, S. Hasan, K. Ali, E. Brewer, and T. Parikh. Local, Sustainable, Small-Scale Cellular Networks. In *Proceedings of the Sixth International Conference on Information and Communication Technologies and Development*, ICTD '13, Cape Town, South Africa, 2013. ACM.

[45] K. Heimerl, R. Honicky, E. Brewer, and T. Parikh. Message Phone: A User Study and Analysis of Asynchronous Messaging in Rural Uganda. In *SOSP Workshop on Networked Systems for Developing Regions (NSDR)*, 2009.

[46] K. Heimerl and T. S. Parikh. How users understand cellular infrastructure. Technical Report UCB/EECS-2012-42, EECS Department, University of California, Berkeley, Apr 2012.

[47] P. W. Holland and S. Leinhardt. Transitivity in Structural Models of Small Groups. *Small Group Research*, 2:107–124, 1971.

[48] H. Horst and D. Miller. *The Cell Phone: An Anthropology of Communication*. Berg Publishers, October 2006.

[49] J. Hou, J. Y. 0008, and S. Papavassiliou. Integration of pricing with call admission control to meet qos requirements in cellular networks. *IEEE Trans. Parallel Distrib. Syst.*, 13(9):898–910, 2002.

[50] International Telecommunication Union. Measuring the information society 2012. 2012.

[51] International Telecommunications Union. Green Solutions to Power Problems (Solar & Solar-Wind Hybrid Systems) for Telecom Infrastructure. http://www.itu-apt.org/gtas11/green-solutions.pdf. Retrieved 1/2013.

[52] M. Jakobsson, J.-P. Hubaux, and L. Buttyn. A micro-payment scheme encouraging collaboration in multi-hop cellular networks. In W. Hunt and F. Somenzi, editors, *Computer Aided Verification*, volume 2742 of *Lecture Notes in Computer Science*, pages 15–33. Springer Berlin / Heidelberg, 2003.

[53] R. Jensen. The Digital Provide: Information (Technology), Market Performance, and Welfare in the South Indian Fisheries Sector. *The Quarterly Journal of Economics*, 122(3):879–924, 08 2007.

[54] K. Johansson, A. Furusk, P. Karlsson, and J. Zander. Relation between base station characteristics and cost structure in cellular systems. In *15th IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, volume 4, pages 2627 – 2631, 2004.

[55] A. Joshi and M. Moore. Institutionalised Co-production: Unorthodox Public Service Delivery in Challenging Environments. *The Journal of Development Studies*, 40(4):31–49, 2004.

[56] D. Katsianis, I. Welling, M. Ylnen, M. Y. (nokia, T. Sphicopoulos, N. K. Elnegaard, B. T. Olsen, D. Varoutas, and L. B. (swisscom. The economic perspective of the mobile networks in europe, 2001.

[57] A. Kemper. *Valuation of Network Effects in Software Markets: A Complex Networks Approach*. Contributions to Management Science. Physica-Verlag, 2009.

[58] Z. Koradia, P. Aggarwal, A. Seth, and G. Luthra. Gurgaon Idol: a Singing Competition over Community Radio and IVRS. In *Proceedings of the 3rd ACM Symposium on Computing for Development*, ACM DEV '13, pages 6:1–6:10, New York, NY, USA, 2013. ACM.

[59] A. Kumar, N. Rajput, D. Chakraborty, S. K. Agarwal, and A. A. Nanavati. Wwtw: The world wide telecom web. In *Proceedings of the 2007 Workshop on Networked Systems for Developing Regions*, NSDR '07, pages 7:1–7:6, New York, NY, USA, 2007. ACM.

[60] W. F. Lam. Institutional Design of Public Agencies and Coproduction: A Study of Irrigation Associations in Taiwan. *World Development*, 24(6):1039–1054, June 1996.

[61] Y.-D. Lin and Y.-C. Hsu. Multihop cellular: a new architecture for wireless communications. In *INFOCOM 2000. Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies.*, volume 3, 2000.

[62] N. Mishra, K. Chebrolu, B. Raman, and A. Pathak. Wake-on-wlan. In *Proceedings of the 15th international conference on World Wide Web*, WWW '06, pages 761–769, New York, NY, USA, 2006. ACM.

[63] O. Morawczynski. Examining the Usage and Impact of Transformational M-Banking in Kenya. In *Proceedings of the 3rd International Conference on Internationalization, Design and Global Development: Held as Part of HCI International 2009*, IDGD '09, pages 495–504, Berlin, Heidelberg, 2009. Springer-Verlag.

[64] S. Nedevschi, J. Chandrashekar, J. Liu, B. Nordman, S. Ratnasamy, and N. Taft. Skilled in the art of being idle: reducing energy waste in networked systems. In *Proceedings of the 6th USENIX symposium on Networked systems design and implementation*, NSDI'09, pages 381–394, Berkeley, CA, USA, 2009. USENIX Association.

[65] J. Nellist. *Understanding Telecommunications and Lightwave Systems: An Entry-Level Guide*. IEEE Press Understanding Science & Technology Series. Wiley, 2004.

[66] D. O'Brien and A. Ploeg. Acculturation Movements among the Western Dani. *American Anthropologist*, 66(4):281–292, 1964.

[67] OpenBTS. `http://openbts.org`.

[68] OsmocomBB. `http://bb.osmocom.org/`.

[69] T. Parikh. Designing an Architecture for Delivering Mobile Information Services to the Rural Developing World, 2007.

[70] N. Patel, D. Chittamuru, A. Jain, P. Dave, and T. S. Parikh. Avaaj Otalo - A Field Study of an Interactive Voice Forum for Small Farmers in Rural India. In *Proceedings of ACM Conference on Human Factors in Computing Systems (CHI)*, 2010.

[71] R. Patra, S. Nedevschi, S. Surana, A. Sheth, L. Subramanian, and E. Brewer. WiLDNet: Design and Implementation of High Performance WiFi Based Long Distance Networks. In *4th USENIX Symposium on Networked Systems Design and Implementation*, 2007.

[72] C. Peng, S.-B. Lee, S. Lu, H. Luo, and H. Li. Traffic-Driven Power Saving in Operational 3G Cellular Networks. In *Proceedings of the 17th Annual International Conference on Mobile Computing and Networking*, MobiCom '11, pages 121–132, Las Vegas, Nevada, USA, 2011. ACM.

[73] A. Ploeg. Wealth items in the Western Highlands of West Papua. *Ethnology*, 43(4):291–313, 2004.

[74] Range Networks. `http://www.rangenetworks.com/`.

[75] A. A. Raza, F. Ul Haq, Z. Tariq, M. Pervaiz, S. Razaq, U. Saif, and R. Rosenfeld. Job opportunities through entertainment: Virally spread speech-based services for low-literate users. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '13, pages 2803–2812, New York, NY, USA, 2013. ACM.

[76] Rhizomatica. `http://rhizomatica.org/`.

[77] M. Ricklefs. *A History of Modern Indonesia Since C. 1300*. Stanford University Press, 1993.

[78] D. Rutherford. Sympathy, State Building, and the Experience of Empire. *Cultural Anthropology*, 24(1):1–32, 2009.

[79] N. Sambasivan and T. Smyth. The Human Infrastructure of ICTD. In *Proceedings of the 4th ACM/IEEE International Conference on Information and Communication Technologies and Development*, ICTD '10, pages 40:1–40:9, London, United Kingdom, 2010. ACM.

[80] A. Sen. *Development As Freedom*. Alfred A Knopf, Oct. 1999.

[81] S. Shenker, D. Clark, D. Estrin, and S. Herzog. Pricing in Computer Networks: Reshaping the Research Agenda. *Telecommunications Policy*, 20(3):183 – 201, 1996.

[82] E. Shih, P. Bahl, and M. J. Sinclair. Wake on wireless: an event driven energy saving strategy for battery operated devices. In *Proceedings of the 8th annual international conference on Mobile computing and networking*, MobiCom '02, pages 160–171, New York, NY, USA, 2002. ACM.

[83] M. Spiegel. Universal service obligation. In M. A. Crew, editor, *Expanding Competition in Regulated Industries*, volume 37 of *Topics in Regulatory Economics and Policy Series*, pages 83–96. Springer US, 2000.

[84] Statistics Indonesia. `http://www.bps.go.id/eng`. Retrieved 4/2013.

[85] L. Subramanian, S. Surana, R. Patra, S. Nedevschi, M. Ho, E. Brewer, and A. Sheth. Rethinking wireless in the developing world. In *In Proc. of the 5th Workshop on Hot Topics in Networks (HotNets*, 2006.

[86] S. Surana, R. Patra, S. Nedevschi, and E. Brewer. Deploying a Rural Wireless Telemedicine System: Experiences in Sustainability. *Computer*, 41(6):48–56, 2008.

[87] S. Surana, S. N. Rabin Patra, M. Ramos, L. Subramanian, Y. Ben-David, and E. Brewer. Beyond Pilots: Keeping Rural Wireless Networks Alive. In *5th USENIX Symposium on Networked Systems Design and Implementation*, 2008.

[88] E. Sutherland. Counting Mobile Phones, Sim Cards & Customers. http://ssrn.com/abstract=1468847, 2009. Available at SSRN.

[89] Telkomsel. Telkomsel, 2011 annual report. `http://www.telkomsel.com/media/upload/annualreport/AR_Telkomsel_2011_ENG_rev.%2021Sept2012_Web%20Version.`

pdf, September 2012.

[90] Ubiquity. http://www.ubnt.com/.

[91] Vanu. http://www.vanu.com/.

[92] S. Wasserman, K. Faust, and D. Iacobucci. *Social Network Analysis : Methods and Applications (Structural Analysis in the Social Sciences)*. Cambridge University Press, Nov. 1994.

[93] L. Waverman, M. Meschi, and M. Fuss. The Impact of Telecoms on Economic Growth in Developing Countries. *Africa the impact of mobile phones*, 2(1980):1–23, 2005.

[94] B. L. Welch. The Generalization of 'Student's' Problem when Several Different Population Variances are Involved. *Biometrika*, 34(1/2):28–35, 1947.

[95] World Bank. World Development Indicators 2012. https://openknowledge.worldbank.org/handle/10986/6014, 2012.

[96] M. Zheleva, A. Paul, D. L. Johnson, and E. Belding. Kwiizya: Local Cellular Network Services in Remote Areas. In *Proceedings of the 11th International Conference on Mobile Systems, Applications, and Services*, MobiSys '13, Taipai, Taiwan, 2013. ACM.

[97] R. Zheng, J. C. Hou, and L. Sha. Asynchronous wakeup for ad hoc networks. In *Proceedings of the 4th ACM international symposium on Mobile ad hoc networking & computing*, MobiHoc '03, pages 35–45, New York, NY, USA, 2003. ACM.

[98] E. Zuckerman. Decentralizing the Mobile Phone: A Second ICT4D Revolution? *Information Technologies and International Development*, 6(Special Edition 2010: Harvard Forum II Essays):99–103, 2010.