Connecting the Last Billion

Yahel Ben David

Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2015-233
http://www.eecs.berkeley.edu/Pubs/TechRpts/2015/EECS-2015-233.html

December 11, 2015

Connecting the Last Billion

by

Yahel Ben David

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Engineering - Electrical Engineering and Computer Sciences

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Eric Brewer, Chair
Professor Scott Shenker
Professor Tapan Parikh

Fall 2015

Connecting the Last Billion

Abstract

Connecting the Last Billion

by

Yahel Ben David

Doctor of Philosophy in Engineering - Electrical Engineering and Computer Sciences

University of California, Berkeley

Professor Eric Brewer, Chair

The last billion people to join the online world, are likely to face at least one of two obstacles:

Part I: Rural Internet Access

Rural, sparsely populated, areas make conventional infrastructure investments unfeasible: Big corporations attempt to address this challenge via the launch of Low-Earth-Orbiting (LEO) satellite constellations, fleets of high-altitude balloons, and giant solar-powered drones; although these grandiose initiatives hold potential, they are costly and risky. At the same time, small local operators, Wireless Internet Service Providers (WISPs), are growing in numbers, in subscribe base and in territory covered. WISPs can play a major role in serving a growing number of rural communities, as well as offer real competition to incumbent operators in urban and semi-urban markets, leading to better service at reduced costs.

The key motivation for this work is to lower the barriers-to-entry for small rural WISPs, and to sustainably grow their operation — this has been my research focus for over 15 years.

The core of this work is based on a case-study of a WISP, **FurtherReach** — which we have built from the ground up. This WISP brings broadband Internet service to hundreds of subscribers at the south coast of Mendocino county in California. Through designing, deploying and operating this venture, we learn about the real challenges faced by WISPs, develop technical solutions as well as business models, operational methodologies and deployment strategies. The FurtherReach case study is presented in chapter 2. Chapter 1 introduces the overall WISP ecosystem. In chapter 3 we discuss the potential of Software Defined Networks (SDN) to aid in WISP operations. Finally, chapter 4 discusses the often overlooked computer-security concerns that are unique to rural communities, especially in developing countries.

Part II: Dissent Networking

Oppressive regimes censor and restrict information flow. Sadly, Internet censorship, in some countries, does not seem to be going away, and presents a growing challenge. The degree and effectiveness of censorship varies greatly, as does the risk of getting caught circumventing it. Similarly, the technologies to aid dissenters vary accordingly. My work in this field predates smart-phones, which I believe could be made to offer safe and effective solutions even in the most dangerous of countries. Should we consider these technologies Internet access? Once again the degree of connectivity, and especially interactivity, from behind a censoring firewall, varies greatly.

In chapter 5, I present our attempt at defining the threats and narrating the exceptionally challenging problem space. I find this chapter quite discouraging as it dictates exceptional restrictions on the design space, yet ignoring these constraints may put users of the technology in greater risks than without it. It makes us question if technology can help at all?

I continue in chapter 6, to present Rangzen, our initial attempt at designing a solution that adheres to the strict constraints presented in 5. While limited in functionality, especially given its delay-tolerant approach, which is incompatible with many Internet applications that expect real-time interactivity, it does successfully follow our design guidelines for dissent technologies as presented in chapter 5. We have built an Android app, Rangzen, based on these design specifications, which undergoes a beta testing program at the time of this writing. The app will be distributed freely on the Android store in January 2016, and the code is open source and available to the public.

To Divya, Nitai, and especially their mother — Yael.

Contents

List of Figures

List of Tables

Acknowledgments

I'm 47 years old — quite a bit older than the average age of doctorate recipients.[1] I had a very rewarding, diversified and rich career before grad-school, and I trust it would only get better moving forward. I have been very fortunate to be guided, directed, assisted, and overall receive the help and support of many wonderful people, both in my professional career as well as in personal life. Of course this is very subjective, but I feel exceptionally fortunate in this regard. That said, of the many supporters and advocates throughout my life, there is one who stands out and tops the list — **Prof. Eric Brewer**. One might say that every doctoral student is highly influenced by their advisor, after all, that is what advisors are for — however, Eric's influence on me, goes much further and deeper. Until Eric suggested it, I never would have envisioned going to grad-school, nor did I ever had such desires or aspirations. Somehow the term *scholar* or *academic* did not fit my vision of myself, still does not, and perhaps never will. Moreover, conventional methods of study were always hard for me — I'm dyslexic and as a kid I suffered from ADHD (I probably still do), and therefore even getting through high school was a scarring experience. With such baggage related to traditional educational systems, the thought of more studies was unpleasant. Eric planted that seed — the idea that I would benefit from, perhaps even enjoy, graduate-school. Eric was able to see in me something I was unable to see in myself.

Throughout the program Eric's hands-off approach to advising worked wonders — I was given absolute freedom to explore and pave my own path, yet with the occasional, well timed, sound advices. Often these would be a single-line email, carefully worded, that penetrated deep and helped refocus my efforts. Overall, Eric's ability to say more with less, in person as well as in writing, is a skill I'll be pursuing well into the future. The idiom *Still waters runs deep* well describes this soft spoken individual who was always accessible and available, with a remarkably short response times. Thank you, Eric!

In fall 2010 I took CS-268 (Computer Networks) with **Prof. Scott Shenker**. From the first class something clicked and it was like we've made an unspoken agreement to disagree — about everything — mostly we disagreed about how researchers design networks vs. how it is done outside academia. Sounds like a scary and unhealthy relationship for a student to have with a professor, but it was a kind of *love-hate relationship* that was lots of fun. Not only I was driven by my unofficial role of disagreeing with the professor, which deepened my study, but I believe the whole class benefited from our heated discussions. A year later, as a TA for Scott's undergraduate class, he asked that I'll revise the curriculum and deliver the wireless lecture to an auditorium with over 500 students — a scary proposition. Scott never heard me lecture before, how could he knew that I can handle such a task? This was a uniquely enriching experience which I have happily repeated in following years, with growing confidence. Scott then became my co-advisor and help inspire and guide my passion for the SDN philosophy in the design of computer networks.

[1] Median age of doctorate recipients in engineering in the USA is 30. http://www.nsf.gov/statistics/2016/nsf16300/data/tab27.pdf

Prof. Vern Paxson was a name I knew from my carreer as a computer-security professional. Naturally, when we moved to Berkeley, and before I applied to the graduate program, I asked to audit Vern's class. Vern agreed, on the condition that I will do all the assignments as well as the term paper. To make a long story short — this was the best course I ever took, and if I still had any doubts about graduate-school, Vern squashed these concerns and replaced them with excitement and enthusiasm. Twice a week, Vern would write pages-long reviews for all of my assignments — I felt privileged. Towards the end of the semester I realized that Vern commented with that degree of detail to all the students in his class. I'm not sure how is that humanly possible, but no doubt this drove students to study deeper (myself included) and overall made this class so much better than most. Suspicions about Vern not being human, or being super-human, remain disproved till this day — to me, he is a super-hero.

The wonderful **De Novo Group** team and the **FurtherReach** project team didnt only made my research possible, but made it fun and enriching in so many ways. One could not have hoped to be surrounded by a finer group of people. Each with their individual qualities made our work together enriching, rewarding and fun. I look forward to continued collaboration with these wonderful souls: **Barath Raghavan, Ron Steinherz, Paul Lamb, Krispin Scanlon-Hill, Zean Moore, Jason Coleman, Michael Hanson, Aurelien Personnaz, Max Bittman, Jenny Ryan, Alex Papazoglou, Leanne Fan, Charlie Fisher, Jasper McMurtry, Jesus Garcia, Adam Lerner, Monte Meyers, julie Lee and Jane Cavlina.**

The EECS PhD program at Berkeley is designed to foster collaboration and teamwork, and it does this well. I was fortunate to collaborate and interact with many remarkable individuals, too many to list here. From random class-projects co-researchers, co-authors of publications (and papers which never made it into a publication), through members of the **TIER (Technology and Infrastructure For Emerging Regions)** research group (both the old gang and the younger batch), folks at the **NetSys Lab**, and other groups that I was more loosely affiliated with. Most on this list are good friends and I hope this friendship would be long lasting: **Shaddi Hasan, Matthias Vallentin, Giulia Fanti, Collin Scott, Joyojeet Pal, Saurabh Panjwani, Philipp Gutheim, Jay Chen, Sebastian Benthall, Kristin Stephens-Martinez, Seth Fowler, Paul Pearce, Albert Kim, Sangjin Han, Aurojit Panda, Justine Sherry, Matt Podolsky, Jordan Kellerstrass, Achintya Madduri, Javier Rosa, Kurtis Heimerl, Kashif Ali, Sonesh Surana, Rabin Patra, Sergiu Nedevschi, Manuel Ramos and Lakshminarayanan Subramanian.**

Special thanks to my quals and dissertation committee members not already mentioned earlier - **Prof. Tapan Parikh** and **Prof. Sylvia Ratnasamy.**
To **Prof. Jenna Burrell** for an exciting co-teaching experience and support/guidance in many of the less-technical aspects of the FurtherReach research (and there are many).
Finally, to **Dean AnnaLee Saxenian**, for her willingness to always support a good cause, and for inspiring and empowering — always with a smile.

Part I

Rural Internet Access – the premise of WISPs

Chapter 1

Introduction – What are WISPs and why should we care?

1.1 Why WISPs

As populated areas and urban centers enjoy rapidly growing broadband speeds with advent of new technologies rural, sparsely populated communities remain unserved. This is a global situation not an isolated phenomenon, and it has to do with a mix of technological, business and regulatory concerns. Essentially, using the conventional technologies of broadband delivery, it is often unfeasible to deploy such infrastructure in rural settings, or financially too risky. This fundamental technological ill-fitness is augmented by business monopolies, and practices such as *mutual forbearance*[1] by operators to stifle competition. In many countries, the large operators also drive the regulatory domain, and deter regulators from supporting new entrants. Overall, there has been very little, and very slow progress in broadband access for rural communities, worldwide. Apart from grandiose and highly risky, thereby questionable, plans by giants such as Facebook, Space-X and Google to use solar-powered drones, low-orbiting satellite constellations, and high-altitude balloons, there seem to be no viable solutions in sight.

At the same time, a relatively new industry has emerged, that although minuscule in size in comparison to the grandiose plans mentioned above, and certainly minuscule in face of the billions that are not yet online, has been steadily growing and bringing connectivity to the world's most rural, and most poor communities these are Wireless Internet Service Providers (WISPs), and especially small WISPs.[2]

Small and local operators are well focused and not only harder to deter, they typically fly under the radar of the bigger players and go unnoticed. The regulatory domain that is at times so heavily

[1] Mutual forbearance is the reduction of the intensity of the competition through familiarity and deterrence. Trying to avoid a price war, competitors, particularly if they have asymmetric positions, can have interests to slow down their commercial behavior in each other market.

[2] Often refereed to as *Fixed-Wireless Operators*, and sometimes also called *Terrestrial Fixed-Wireless Operators* to distinguish the technology from satellite-based service in which subscribers' antennas are aimed at a satellite in the sky.

skewed in favor of giant incumbents, may often overlook and exclude the operational methodologies of very small operators, especially if they require no licenses nor public funding for their roll out. At the same time, solutions have emerged to enable exceptionally low-cost deployments, based on fixed-wireless technologies, and a whole industry is growing around these hardware vendors who build low-cost, low-power and outdoors-ready networking devices, software vendors who design tools for this industry, and a vast network of distributors, world-wide, to make these solutions accessible and affordable. This industry, unlike with incumbent telcos, enjoy constructive competition and strive with exceptionally low profit margins, driven by large quantities.

Traditional communication infrastructure, whether based on old-age copper pairs, coaxial cables, or fiber optic cables take years to Return On Investment (ROI), sometimes decades, while as for most WISPs using unlicensed fixed-wireless technologies, the ROI is measured in months, at times as low as three months! Not only the ROI is faster, but the initial Capital Expenditures (CAPEX) needed for deployments is much smaller. While small fiber-based projects are in the hundreds of millions of dollars,[3] a WISP operation to serve hundreds of subscribers in highly-rural area is likely to cost in the sub-million dollar range, and may ROI within a couple of years. Moreover, given the low cost of equipment used by WISPs many vendors offer attractive financing to reduce WISPs' cash-flow constraints.

Without side tracking into development theories, I like how this bottom-up industry had emerged to address the needs of communities neglected, or completely overlooked, by conventional top-down approaches. Without public funding, without intervention by academics and development theory researchers, the WISP industry is booming. It prospers in almost complete isolation from adjacent ecosystems, nearly unaffected by the traditional telecommunication industry, unhindered by regulatory capture and unfavorable policies. I envision fascinating opportunities for future multidisciplinary research around the growth of WISPs and the social impact they bring about.

WISPs use stationary wireless devices, typically using directional antennas,[4] to establish radio links between subscriber's rooftops and a nearby base station site. By and large, the radio technology used by WISPs requires a clear line of sight (LoS) between the subscriber's device[5] and nearby base station. Typically, multiple subscribers would be fed from a single base-station antenna using a point to multi-point (PTMP) topology. This part of the service has been traditionally called the ”last mile”, although it's common for WISPs to use these links and topology over much longer distances. Some WISPs feed their basestations using some form of wired connection from an upstream Internet provider, but many use radio technology as a back-haul link between their base-station sites and an upstream operator. These links are commonly known as ”middle mile” and are typically over a longer distance and higher capacity then those feeding a single subscriber.

[3] Based on recent years of both the state of California (CPUC) or federally funded grants for rural broadband projects.

[4] Typically one side of the link, the subscriber's side, uses a high-gain directional antenna aimed at a base station that uses a wider-angle *sector* antenna.

[5] In Telco jargon, subscriber's devices are typically referred to as CPEs – Customer Premises Equipment.

Unlike cellular operators, WISPs do not offer service to mobile devices, and their notion of coverage-area means that with high likelihood, often thanks to the use of a sizable antenna mounted on a tall mast, and at times pending some tree trimming, coverage is likely to be achieved.

These technical properties are common across most WISPs, while many other properties, both technical and operational, may vary substantially and will be discussed later.

The WISP industry is already nothing short of booming, nevertheless, this research focuses on how to drive it further and allow WISPs to overcome their current and future challenges.

1.2 Many small WISPs vs. giant Telcos

For most WISPs, being acquired by a larger operator is the most common route to scale. Such acquisition typically stops the organic growth, and limits expansion into further rural and challenging areas. It's common that with an acquisition, the service quality degrades, and the personal touch, and local care of the original founder goes away with him (the use of masculine pronoun here is not by mistake – I have yet to come across a female WISP founder). But there is an even bigger problem – most WISPs don't grow enough, and don't do well enough to be bought by a larger operator (not even their remains when they go under), and certainly not well enough to support themselves. Many WISPs simply fade-away, and at times cause much harm to the community they served with their (often sudden) demise. Obviously, there are exceptions, some WISPs are doing very well, others operate a sustainable business, often referred to as a lifestyle business, that provides for their families, and that of their employees, while providing excellent service to their subscribers. It's this second type of entrepreneur that I believe will play a key role in reaching the very rural, very sparse, and very poor communities. These operators would happily operate on low profit margins, and would work hard to build-out, and maintain their network, while providing high quality of service to their local communities. I anticipate such operators would have between 100 and 1000 subscribers, would employ two to ten employees, and would cover an area with a radius of a few hours drive (without the need for overnight stay to service the most remote corners of the network). These are some of the natural boundaries for WISPs, although there are other containing factors discussed later[1]. In other words, we're looking at millions of such small operators to serve the last billion – we need to lower their barriers-to-entry, and we need to equip them with tools for sustainable operation and growth.

1.3 Wireless meets fiber

Regardless if a WISP operates a single base-station or many, the wireless portion of the network must be fed from a wired Internet gateway, or multiple such gateways. This often becomes a barrier-to-entry for WISPs operating in rural areas – where would they get their upstream bandwidth? Most WISPs, especially the smaller ones, begin their operation by purchasing a residential-grade Internet service from the local telco or cable operator. In the USA these would be ADSL-based (over copper pair) offerings from the likes of AT&T and Verizon, or the likes of Comcast cable using DOCSIS technology over coaxial cable. Although such operators would not cover rural and sparsely populated areas, they would probably offer service in a nearby town which the WISP could use and wirelessly extend to the more rural areas. There's nothing critically wrong with this approach and this type of upstream bandwidth; indeed many WISPs use that to feed their entire operation. That being said, there are some notable challenges and limitations to the "residential-grade" Internet gateway as used by WISPs:

- Legal concerns — many ISPs clearly forbid resale of their bandwidth. While often hard to enforce, at times when usage pasterns are particularly high or otherwise indicative of resale, the ISP would either employ automated measures to slowdown and harm the connection (a practice often refereed to as data-cap-based-throttling), or would initiate an investigation leading to legal proceedings.

- Reliability — the above mentioned Internet plans that target residential service are not known for their reliability. Essentially, not only there is no resiliency built into the service, it relies on copper cables running to the home from a nearby central-office (CO) or more commonly from a street-side cabinet. These copper cables are typically aerial, strung over shared utility poles. Overall, these sort of systems introduce many potential points for failure, subject to harm by weather, hit by vehicles, typically dependent on grid-power without backup, and could even get harmed by wildlife. Essentially, the focus of these offerings is on low-cost and reuse of existing copper infrastructure, as old as it may be. The shared coaxial cable infrastructure is especially fragile, and a malfunctioning cable modem (or simply a dead-short) at one home could harm the service to much of the neighborhood. Service levels to these sort of packages are also limited by design — it could take quite a few days for a repairman to arrive. Imagine a WISP that goes dark at the beginning of long weekend and cannot provide service for a few days all because one bird too many decided to rest on a Comcast cable — obviously, the subscribers of that WISP would quickly seek alternatives for their business (if there are any).

- Availability — one obvious way to improve the reliability of the upstream service is to purchase multiple residential packages and load-balance the WISPs traffic across all of them. This approach is not without challenges and limitations either. First, buying multiple services from the same operator is not going to improve reliability by much - the service is still dependent on the same poles, the same street-side cabinet, and the same main cable (even if in the case of ADSL it would be over different copper pairs). Moreover, the operators are unlikely to allow multiple services to the same address as this would raise suspicion that the bandwidth is used for resale. (WISPs are resourceful and known to circumvent this by ordering service to multiple suites in an office building, or multiple residents sharing a house — but these tricks don't scale too well, and are essentially subject to the limitations of a single cable to the home, both in terms of reliability and capacity). Purchasing service from two or more operators is a wiser approach. However, it's common that in rural towns there would only be a single operator, if any. Moreover, the diversity in technologies introduces another problem — while Cable-TV systems offer speeds surpassing 100mbps, ADSL only scale to about 45mbps, if in good proximity to the CO.[6] Load-balancing over multiple links with varying capacities adds challenges and if the faster of the two goes down the remaining link might get congested.

- Capacity and over subscription — At the time of this writing, the fastest offering from Comcast, targeting small businesses, is for ”up to” 150mbps download and 20mbps up. The price is $250 per month (with two years commitment). AT&T’s fastest offer, over two copper pairs (pending short distance from CO), is for ”up to” 75mbps down, and 8mbps up, and goes for $110 per month (with two years commitment). These download capacities are sufficient to support a decent WISP operation, depending on what is an acceptable over-subscription ratio for the target market. I came across rural markets in the USA where a total Internet gateway of 100mbps or less, was serving over a thousand (mostly happy) subscribers. It’s common for rural WISPs to offer plans of 1mbps (cost is in the range of $25–50 per month) therefore yielding an over-subscription rate of 10:1 which is considered acceptable for rural access networks. Even if most subscribers splurge and go for the 2mbps plans, it’s probably still okay to split a 100mbps across a thousand subscribers. Such ratios are still common in rural parts of North America, and over-subscription is much higher in rural parts of developing countries. However, the nature of over-subscription is mostly misunderstood, even (or especially) by small WISPs. To better discuss over-subscription ratios, one must understand that such capacity planning decisions are based on complex probabilistic models and cannot be compared with ease. For example, let us examine three scenarios:
  A. An ISP has a 100mbps gateway and a thousand 1mbps subscriptions.
  B. An ISP has a 10mbps gateway and a hundred 1mbps subscriptions
  C. An ISP has a 100mbps gateway and a hundred 10mbps subscriptions.
  In all scenarios the over-subscription rate is the same - 10:1. But will the service be the same? In other words, would the perceived users’ satisfaction be the same? The answer depends on the probability for a small number of users to congest the gateway. Therefore scenario B and C are the worst, as only 10 simultaneous users are capable of congesting the gateway, while as in scenario A it would require 100 users — a lower probability. Theoretically speaking, we could have derived a relatively simple equation for such capacity planning, as studies have shown that packet arrivals could be well modeled as Poisson distribution. However, realities differ widely based usage patterns and these differ widely by market. For example, recent uptake of Netflix video streaming been shown to dramatically affect traffic patterns across the whole Internet[2], and this trend is driving substantial lowering of over-subscription rates by ISPs, in-turn hurting ISPs profit margins[3] and heated debates around net-neutrality[4]. We therefore should not compare over-subscription ratios as fixed numbers as such comparison is mostly pointless without context. At the very least

[6] Operators now offer faster speeds over ADSL by bundling two copper pairs - at the time of writing, the highest speed offered by AT&T to small businesses, using copper-pairs was ”up to” 75mbps, with an ”up to” 8mbps upload capacity.

we should add the total capacity of the Internet gateway, if we're to derive any value from such comparisons.

- Asymmetry — the inherently asymmetric nature of the above mentioned technologies for bringing Internet to the home, is also detrimental to capacity planning and its overall unsuitability for feeding a WISP operation. If we recall the fastest offering from Comcast, of 150mbps down, it only had an upload speed of 20mbps, which is the typical maximum in most of Comcast deployments (although the current prevailing DOCSIS3.0, or 3.1 specifications allow for faster upload speeds, if the operator is willing to compromise TV channels in favor of more upload capacity).[7] With asymmetric bandwidth, our probabilistic capacity planning needs to be based on the much lower upstream capacity. If we take for example the favorable scenario A above, which could have allowed for reasonable service to thousand subscribers of 2mbps each, and we'll limit the upload speed of these subscribers to an acceptable minimum of 512kbps, we learn that our 20mbps might get congested by 40 subscriber who are uploading at full capacity. With an install base of a thousand the odds for that to happen are quite high, and such instances would bring down the network. It is therefore why asymmetric services are less than ideal as a main feed for a WISP. Additionally, the wireless technologies used by WISPs (half-duplex radios) are symmetric, so its a shame to restrict the service and offer asymmetric speeds to subscribers (through rate-limiting).

- Burst-ability vs. hard limit — Burstable bandwidth, in layman terms, is a way to allow much faster speeds for short bursts, as long as the main pipe is not congested. This technology work wonders for increasing users' satisfaction when using applications that are of "bursty" nature, like web-surfing. The users gets their web pages much faster and with good probability there's no harm to other subscribers who share the bandwidth. The nature of most residential services, however, is that they're hard limited and do not allow bursts. In ADSL deployments the operator restricts the circuit speed based on the purchased plan, and even if the proximity of the subscriber to the CO allows for faster modulations (which could have offered bursting) these are not enabled in most scenarios.[8] Cable-TV systems, using DOCSIS, often do allow "burstable" bandwidth in the download direction, but rarely upstream. We later discuss wholesale bandwidth options for WISPs, in contrast with residential-grade bandwidth, where the potential of "burtsable" capacity is brought to the fore.[9]

- Price — Price is the main reason for WISPs to consider the purchase of residential-grade bandwidth in the first place. Contrary to common belief, it is not an availability concern - the Telcos will be happy to deliver as much capacity as one is willing to purchase to wherever it is needed, for the right price — that price is "by design" not attractive for resale.

[7] In other words, a single business-grade subscriber of Comcast, could bring the whole neighborhood down if using their maximum uplink capacity - such as when uploading videos, etc.

[8] More commonly, the purchased speed cannot be achieved due to length or poor copper quality, hence the speeds are capped by the limitations of this aging infrastructure.

[9] with wholesale bandwidth the circuit capacities are typically very high, while the client (the WISP) commits to purchase a monthly bulk capacity metered over the faster circuit.

> The telcos set a high price tag to stifle competition, it does not matter if they already have sufficient unused capacity at the target location (which is the case for most rural Central Offices in North-America, as their majority is fiber-fed already) it is simply a business decision and has little to do with actual infrastructure or capacity. The exceptions would be in locations where there's high competition, such as in telco-hotels[10] or Internet exchanges - but it's unlikely to find these in rural areas where WISPs would benefit from such low bandwidth prices.[11] As for rural areas, fast business circuit prices will indeed drop somewhat if bandwidth is available from more than a single carrier at that location, but due to the practice of ”mutual forbearance” by telcos, these prices would still be much higher than that of wholesale bandwidth at a Telco hotel.[12] To summarize, while residential-grade bandwidth has many limitations, it's price as well as availability in semi-rural areas is what makes it an attractive choice for Internet gateway by small WISPs. It's important to note, however, that prices of wholesale bandwidth are yet lower, at least the price per Mbps, and do not share the above mentioned limitations, but are harder to come-by for WISPs, as we discuss later.

Residential-graded Internet gateway, although it's many limitations, is often the only option for small WISPs. Moreover, for many subscribers of WISPs, there are no other viable options, thereby subscribers' acceptance and tolerance to downtime, congestion, and overall slow service is quite high.

[10] Telco-hotels are data-centers where multiple carriers have presence and often interconnect - bandwidth prices in these locations is therefore cheapest, with little variation between the competing carrier.

[11] It's not completely impossible, there are quite a few Telco hotels and large Internet exchanges that are within reasonable proximity to unserved areas (such as Sacramento, CA, or Navato, CA, for example).

[12] While the price of wholesale bandwidth, as well as of fast business circuits may vary much across location, the price of residential service is typically fixed across the service area of the large telcos - which is often nation-wide.

I'm often asked if there's a way to tell, for an area served by multiple WISPs, which one is better? Obviously, there's no easy way to know, but I believe an important differentiating factor is the WISP's choice of Internet gateway; Use of residential-grade upstream bandwidth for the WISPs gateway is typically a bad sign — it indicates acceptance by the WISP for the above mentioned limitations, and essentially means the service must be of lower quality than that of a basic residential service on which it is built. It also typically means small-scale WISP, as the limited Internet gateway is likely to restrict growth. Many WISPs grow by employing multiple residential-grade gateways purchased at different locations, essentially operating multiple isolated networks each fed by a separate gateway. Overall, not essentially a bad thing, but perhaps indicative of the WISP's goals, capabilities and business practices.

The question therefore, is now to know what sort of Internet gateway is used by which WISP? Obviously, skillful tracing of traffic from that WISP could provide good hints, but how can one chose before purchasing the service needed for such tracing?

I believe the tell-tale sign is the plans offered by the WISP and published on their website; If the plans suggests asymmetric speeds, it is probably an indication that the WISP is using a residential-grade service as their Internet gateway — why else would a WISP offer asymmetric plans?

Its intriguing to study this hypothesis further, but few centralized databases of WISPs exists, and these do not include the nature of plans offered. Moreover, few WISPs would voluntarily confirm they are using residential-grade gateways.

So what are alternatives to residential-grade bandwidth that are available to WISPs? Essentially, the solution is to buy bandwidth from carriers that are not telcos, or in other words, that do not offer service to end-users, and thereby do not enforce an artificially high price tag on bandwidth to stifle competition. In the USA that would preclude companies like AT&T, Verizon, Comcast and similar operators. The challenge therefore is how to find other carriers that actually reach rural areas, and how likely is that? Realities, at least for North America, suggest that many such options do exist! Carriers of long-haul fiber employ *In-Line Amplification* (ILA) stations (often called ”regeneration huts”) every 10-20 kilometers along the cable's path, to regenerate the signal (there are cases of longer legs without regeneration, but these are less common). Companies like Level-3-communications, for example, own a dense network of long-haul fiber across the nation (and world-wide) and there are many such companies, and thus a staggering number of ILA huts, with their vast majority in rural areas. That said, although owned by many, often competing companies, their fiber routes would mostly overlap as they are traditionally laid along the same paths, typically railroad tracks. There are, however, many exceptions of other fiber that reaches rural areas, such as power-grid operators who run fiber along their high-voltage transmission lines. As an example, see figures 1.1 and 1.2 for a partial map of ILA huts and other Level-3 facilities in California, where bandwidth at wholesale prices is likely to be available.

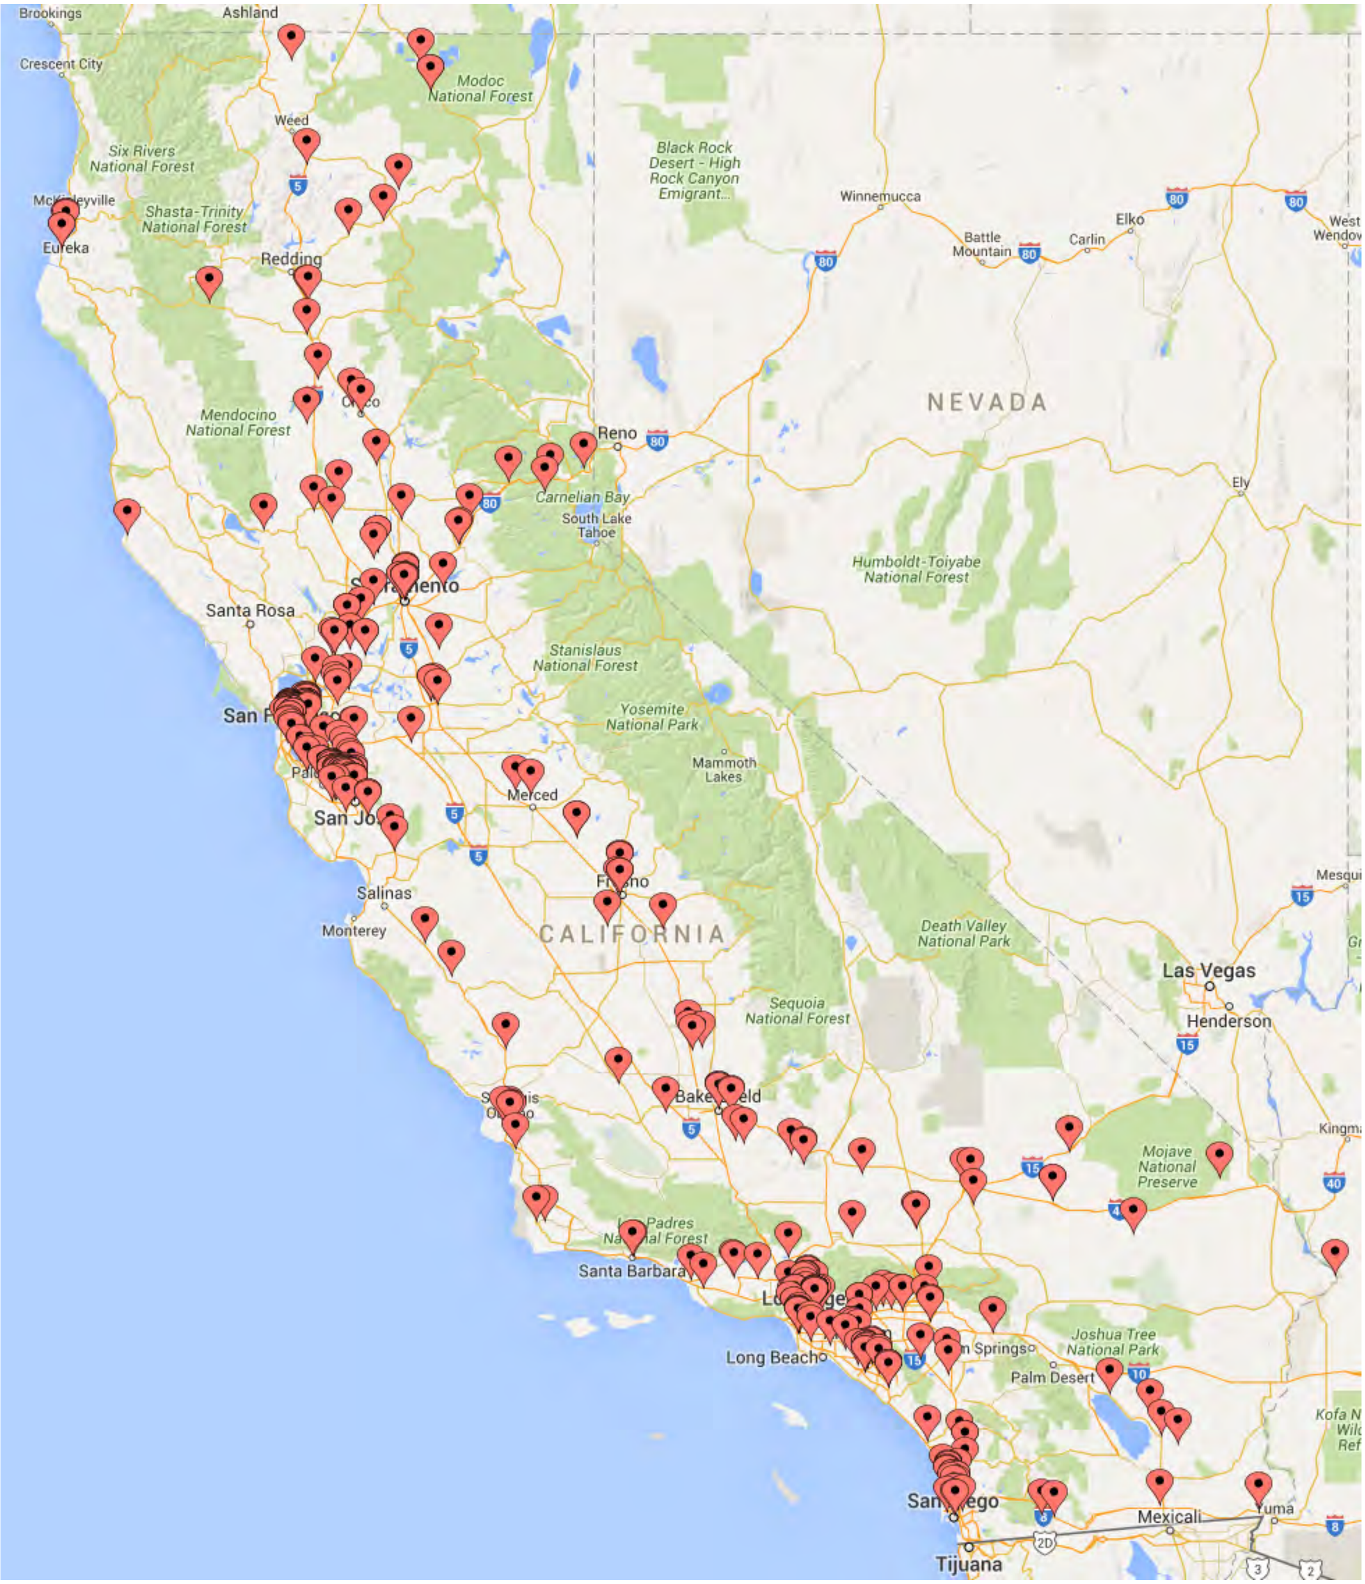

Figure 1.1: Level-3 communications facilities in California without IP service - partial list.

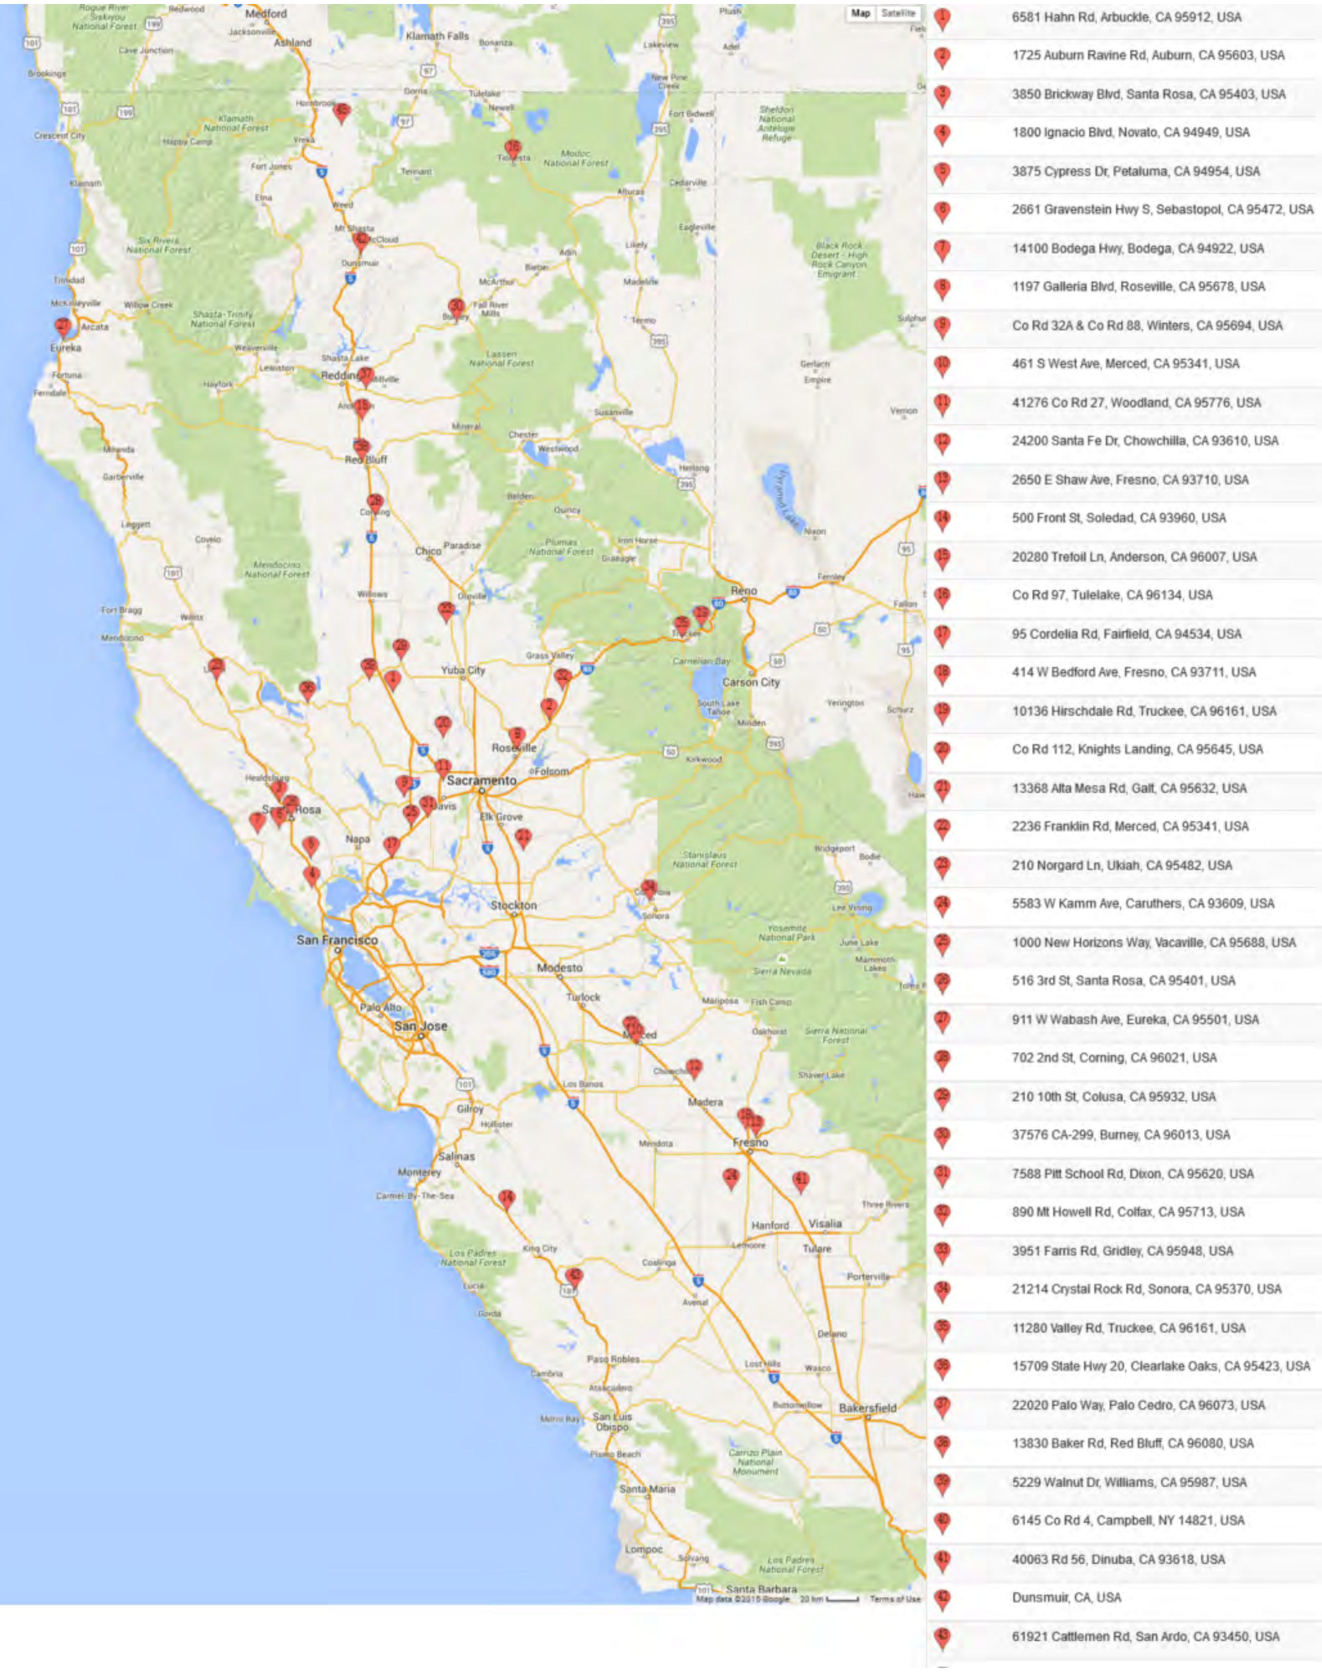

1 6581 Hahn Rd, Arbuckle, CA 95912, USA

2 1725 Auburn Ravine Rd, Auburn, CA 95603, USA

3 3850 Brickway Blvd, Santa Rosa, CA 95403, USA

4 1800 Ignacio Blvd, Novato, CA 94949, USA

5 3875 Cypress Dr, Petaluma, CA 94954, USA

6 2661 Gravenstein Hwy S, Sebastopol, CA 95472, USA

7 14100 Bodega Hwy, Bodega, CA 94922, USA

8 1197 Galleria Blvd, Roseville, CA 95678, USA

9 Co Rd 32A & Co Rd 88, Winters, CA 95694, USA

10 461 S West Ave, Merced, CA 95341, USA

11 41276 Co Rd 27, Woodland, CA 95776, USA

12 24200 Santa Fe Dr, Chowchilla, CA 93610, USA

13 2650 E Shaw Ave, Fresno, CA 93710, USA

14 500 Front St, Soledad, CA 93960, USA

15 20280 Trefoil Ln, Anderson, CA 96007, USA

16 Co Rd 97, Tulelake, CA 96134, USA

17 95 Cordelia Rd, Fairfield, CA 94534, USA

18 414 W Bedford Ave, Fresno, CA 93711, USA

19 10136 Hirschdale Rd, Truckee, CA 96161, USA

20 Co Rd 112, Knights Landing, CA 95645, USA

21 13368 Alta Mesa Rd, Galt, CA 95632, USA

22 2236 Franklin Rd, Merced, CA 95341, USA

23 210 Norgard Ln, Ukiah, CA 95482, USA

24 5583 W Kamm Ave, Caruthers, CA 93609, USA

25 1000 New Horizons Way, Vacaville, CA 95688, USA

26 516 3rd St, Santa Rosa, CA 95401, USA

27 911 W Wabash Ave, Eureka, CA 95501, USA

28 702 2nd St, Corning, CA 96021, USA

29 210 10th St, Colusa, CA 95932, USA

30 37576 CA-299, Burney, CA 96013, USA

31 7588 Pitt School Rd, Dixon, CA 95620, USA

32 890 Mt Howell Rd, Colfax, CA 95713, USA

33 3951 Farris Rd, Gridley, CA 95948, USA

34 21214 Crystal Rock Rd, Sonora, CA 95370, USA

35 11280 Valley Rd, Truckee, CA 96161, USA

36 15709 State Hwy 20, Clearlake Oaks, CA 95423, USA

37 22020 Palo Way, Palo Cedro, CA 96073, USA

38 13830 Baker Rd, Red Bluff, CA 96080, USA

39 5229 Walnut Dr, Williams, CA 95987, USA

40 6145 Co Rd 4, Campbell, NY 14821, USA

41 40063 Rd 56, Dinuba, CA 93618, USA

42 Dunsmuir, CA, USA

43 61921 Cattlemen Rd, San Ardo, CA 93450, USA

Figure 1.2: Level-3 communications facilities in Northern California where IP services are available - partial list.

A problem with ILA huts, is that they typically don't have any IP routers, so IP-transit is not offered, and the only product available is 10GE transport to a remote datacenter. Naturally, that may make the costs unattainable for many WISPs, who would never need anywhere that much capacity in their wildest dreams. Moreover, in some of the smaller regen huts, even 10GE is unavailable, and the only possibility would be to lease a dark-fiber pair, often a prohibitively costly endeavor in most markets, unless you are a giant telco at at least a CLEC.[13]

Figure 1.1 is a partial map of Level-3 facilities in California. The majority of these facilities are ILA huts where IP services are not offered. As seen, apart from the dense clusters around urban centers, there are quite a few facilities in rural parts of the state, although there are clearly vast areas that do not have Level-3 fiber anywhere nearby.

Figure 1.2 shows facilities owned by Level-3 communications, where IP services are offered. Some of these locations are strikingly rural.

The bottom line is that high-capacity bandwidth is available, in many rural corners of the world, and often at affordable, wholesale prices. That being said, the ability to do business with these large operators is often challenging and inaccessible to anyone who is not a large operator themselves - small WISPs might struggle with these business dealings.

1.4 Topography, Antennas, Frequencies and how it all fits together

Line of sight and Fresnel zone

As discussed earlier, the radio technologies used by WISPs are predominantly dependent on Line of Sight (LoS). Seen in figure 1.3 clear LOS is typically not enough for a good signal. The microwave radio signal propagates in an ellipse pattern known as a Fresnel zone and obstructions within this pattern may attenuate the signal. This is the main reason for use of communication towers, to overcome obstacles on the ground and the curvature of the earth, which becomes an issue for long-distance links. The radius of the Fresnel clearance zone is not dependent on distance alone, but is also affected by frequency, the lower the frequency (or the longer the wave-length) the larger is the radius of the Fresnel clearance zone.

$$Fz = \sqrt{\lambda \frac{d^2}{2d}}$$

- Fz is the Fresnel zone radius (in meters)
- λ is the signal's wavelength (in meters)
- d is the distance between receiver and transmitter (in meters)

[13] CLEC - In the United States, a CLEC (competitive local exchange carrier) is a telephone company that competes with the already established local telephone business by providing its own network. The term distinguishes new or potential competitors from established local exchange carriers (LECs) and arises from the Telecommunications Act of 1996, which was intended to promote competition among both long-distance and local phone service providers.

In 1.3, the tree marked ”borderline” will be outside the Fresnel clearance zone and would pose no obstacle if we use higher frequency, while it may hinder the link unusable at lower frequencies. Base on this intuition alone, we might feel inclined to use higher frequencies whenever faced by a Fresnel clearance challenges — this is true, but must be weighted in relation to another radio-propagation phenomenon, that of free-space path loss (FSPL).

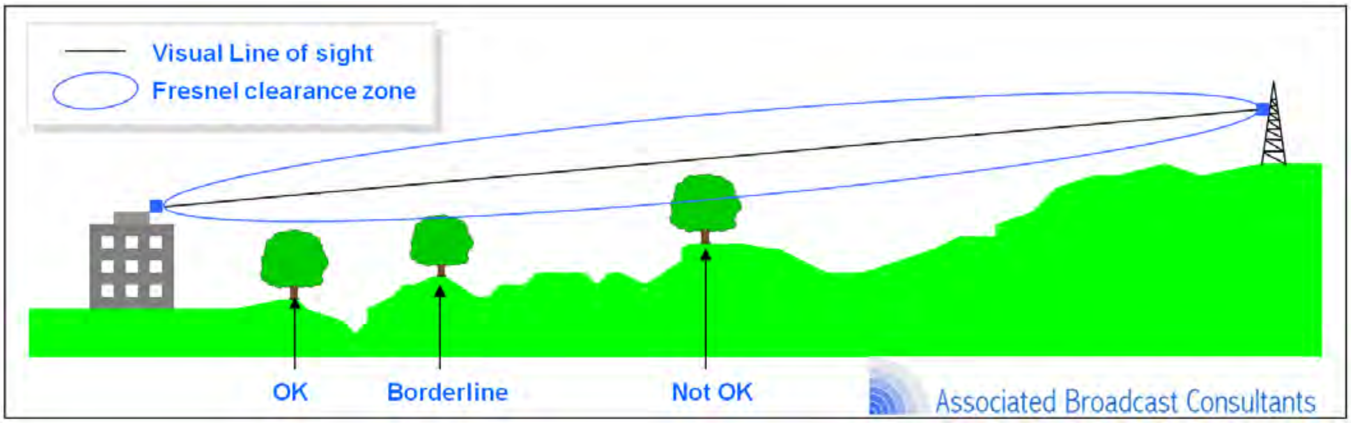

Figure 1.3: Line-of-sight and Fresnel clearance zone.
Image use with permission - Glyn Roylance, Associated Broadcast Consultants

FSPL is the loss in signal strength of an electromagnetic wave that would result from a line-of-sight path through free space (usually air), with no obstacles nearby. It’s important to note that the formula for FSPL is non-linear, similarly to the formula for Fresenl zone clearance, they are proportional to the square of the distance between the transmitter and receiver, and also proportional to the square of the frequency of the radio signal.

$$FSPL = \left(\frac{4\pi d}{\lambda}\right)^2 = \left(\frac{4\pi df}{c}\right)^2$$

- λ is the signal’s wavelength (in meters)
- d is the distance from transmitter to receiver (in meters)
- f is the signal’s frequency
- c is the speed of light

Use of lower frequencies is commonly thought of as a way to improve signal propagation; as seen by the FSPL equation above, lowering the frequency indeed appears promising, and foliage penetration is also considered better with lower frequencies. However, wider Fresnel clearance is not the only deterrent from using lower frequencies:

- We may not have enough spectrum at lower frequencies that we can use legally, and indeed much regulatory progress is being made, word-wide, to free more spectrum for unlicensed use. Most notable are the new regulations for use of TV-White-spaces (TVWS), which are indeed at lower frequencies than commonly used unlicensed bands, like WiFi at 5Ghz, 2.4Ghz and even 900Mhz.[14]

- Antenna sizes grow with wavelength, or in other words, the lower the frequency the larger the antenna. This could be detrimental to the viability of the overall solution. Larger antennas are not only more costly, they are also heavier and result in higher wind loads. In effect, larger antennas call for stronger towers, stronger mounts, and less usable tower space overall. Let us not forget the wider Fresnel clearance radius, which should drive towers to be higher than if used with higher frequencies.

- Cost is of-course driven by size and quantity of material, thereby larger antennas are more expensive, not to mention the overall cost of the beefier towers and mounts. However, economies of scale are to be considered as well; the large quantities of WISPs-targeted gear in the 5Ghz band, resulted in huge cost reductions, and the reuse of the same silicon as in consumer devices (all use WiFi radios in the same band) is also critical for driving prices down. However such economies of scale do not favor TVWS gear or other lower frequencies solutions. It seems unlikely that TVWS technology will become nearly as ubiquitous as the 5Ghz band, certainly not for fixed-wireless broadband delivery, and we should therefore not expect a meaningful cost reduction of this technology.

Antennas

Skillful choice of antennas and their placement is crucial for WISPs' success. We desire wireless links to have a high Signal to Noise Ratio (SNR), to enable a stable connection with high capacity as well as resistance to interference. An obvious way of achieving high SNR is to increase the radiated power at the transmitters, however, that approach is typically not only illegal, it is also harmful as it would interfere with other devices in the area, and in turn they would be forced to increase their power and interfere with you. A better approach for improved SNR is the use of directional antennas, often called high gain antennas. Antennas are passive devices and therefore high-gain means more focused and more directional. With high gain antennas, more energy goes to where it needs to go with less interference in other directions.

Two high-gain antennas pointed at each other is indeed the solution for middle-mile links used by WISPs to cover long distances, the WISP's back-haul, a point-to-point link. However, for connecting subscribers' rooftops, the last-mile, the desired topology is that of a point-to-multi-point, where a single basestation serves multiple subscribers. Not only is it unfeasible to setup a dedicated directional antenna at the base-station site aimed at each subscriber individually, there would not be enough non-overlapping frequencies to allow such system. The typical WISP deployment therefore uses higher gain antennas at subscribers' rooftops (CPEs) and lower gain antennas, with

[14] In the USA, the TVWS spectrum that is used for communications is in 470-698 MHz range.

wider angle at the base-stations (BTS). The wider angle antenna at the BTS allows it to serve multiple CPEs that share the channel.

Small WISPs that serve extra-rural areas often choose to use omni-directional antennas. Omni antennas directs the signal in a circle around it, 360 degrees, and would serve CPEs in all directions. Omni-directional antennas, however, are by definition of lower-gain than directional antennas and the resulting SNR is lower. The low SNR, even if resulting in low-capacity link, might not be a problem for some small WISPs, if the WISP's bottleneck is their Internet gateway, as often is the case with small WISPs who use residential-grade service for their Internet uplink, they may not care about the capacity of their last-mile as much. Moreover, if the plans offered by the WISP are for low-speeds, the last-miles links might not matter. A 7dBi omni directional antenna, at two miles away from a 19dBi CPE could provide 1-4mbps capacity if there are no interfering signals on the same channel and with the overall low-noise floor as common in rural settings. Such speeds are enough for many WISPs, but omni-based service would not scale to higher bandwidths, nor to a large number of subscribers. The typical number of subscribers that is supported by access-points used by WISPs, such as the Ubiquiti[5] Rocket-M5, or NanoStation-M5, is of about 30 subscribers.[15] This is not a hard limit and depends on the usage patterns of the subscribers, and especially on the signal quality (flapping links are tolling to the whole AP), but realities have shown the above mentioned limits, which are commonly accepted across the WISP industry. What this limit means for a WISP that uses omni-directional antennas, is that the whole base-station site is not to expect more than 30 subscribers in all directions. While quite a small number, that may jeopardize the viability of the whole site, it is not uncommon in rural areas.

Omni antennas are also a pretty bad interference source, as they emit their energy in all directions, but if there's nothing to interfere with, as might be the case is some extra-rural areas, this might be fine. It is interesting to note that contrary to common belief, omni-directional antennas are costlier than many of their directional counterparts. A common panel antenna with an embedded router that is used by many WISPs at base-stations, is the remarkable NanoStation-M5 by Ubiquiti[5], which costs $89 (MSRP). This device will serve 20-30 subscribers without a problem, and its specifications suggest its gain is 14db at an angle of 55 degrees. 55 degrees might seem narrow for a base-station in the center of town, but it would provide decent coverage across 120 degrees, with only some signal attenuation. The omni-directional alternative would be a dual-polarized 13dBi antenna from Ubiquiti, AMO-5G13 at $165, that would be coupled to a Ubiquiti Rocket-M5 router at $89.[16] To get a complete 360 degrees coverage using the directional NanoStations, we would need 3 such devices, so the omni directional approach is still slightly cheaper,[17] but how often do we really need 360 degrees coverage?

How often do we really have a base-station on a tower smack in the middle of town? It is

[15] Newer access points, such as the recent 802.11AC devices support larger numbers of subscribers.

[16] The relatively high cost of the omni antenna is due to the dual-polarization requirement - the whip shape of omni antennas does not land itself well to horizontally polarize emitters. Vertical-only omni antennas are very cheap.

[17] In fairness, to connect multiple directional devices we would also need a network switch, while as the omni antenna is a single device. However, Nanostations allow their daisy-chaining on the tower to alleviate the need for a switch and multiple cable runs. Nonetheless, we've found it best not to daisy-chain more then 3 such devices due to power limitations of the PoE feed.

more common for a WISP to have a base-station on a hill overlooking town, or a tall building. The tall building, if in the center of town would need a high mast for the omni-antenna on top of the roof, to ensure the building itself is not blocking the signal from subscribers. While as with the Nanostation-like devices, one could be mounted on each corner of the roof, without additional height and would have clear LoS to subscribers. Moreover, the omni antenna will be able to serve a maximum of about 30 subscribers, while with the Nanostation approach there would be 30 subscribers per each. Obviously the SNR for each link would be higher with the directional approach than with omni, as well as less interference (it's contained in a given direction) and better resistance to interference. There is another important aspect to consider in this comparison, that of the vertical angle, in other words can we serve subscribers right under the tower as well as those on a hill above it? The omni-directional antennas must have a very narrow vertical angle to achieve any decent gain. This not only limits their vertical diversity, it also dictates a very solid mount as if it would swing, the very narrow vertical beam may overshoot or undershoot subscribers as the wind blows, leading to packet loss and flapping links.

Omni directional antennas must be mounted at the very top of a mast, or else the mast itself would shadow their emitted energy to a particular direction. Top mast mounting is problematic as the antenna might get hit by lightning instead of the lightening arrester that should typically be at the top of the tower. Use of stand-offs, to mount antennas away from the tower in order to minimize the shadow effect is a commonly used solution for omni antennas, but it's not without cost and complications (wind-load and other strength concerns, not to mention difficult to reach for service). Directional antennas are ideally mounted at the side of a mast or tower, below it's top and are therefore not subject to the shadow effect nor these complicated mounting concerns. Finally, the omni directional antennas would be quite a bit heavier and larger than 3 or 4 of the Nanostations. It would require a solid mast and strong mounts to support the higher wind loads.

Overall, with the exception of the ultra-rural settings, and for a WISP that only provide low speeds, it would be unwise to use omni-directional antennas.

Another tell-tale sign for a small and unaspiring WISP is the use of omni-directional antennas. These antennas are easy to recognize in rural settings, and while indicative of an area served by a WISP - expectations regarding service quality and speed should be contained.

As shown, the choice of antennas directly depends on the deployment topology, and wise choices for the locations of base-station sites are probably the most important aspect for any WISP's deployment. In 2009 I wrote about the use of further-away mountains or hills for base-station placement as a preferred approach over a tower in the middle of the village[6]. As seen in figure 1.4 similar or greater footprint could often be achieved from a remote mountain using a directional antenna. An omni antenna in the center of the village is likely to offer inferior SNR to that of the remote directional site. Such approach may also eliminate the need to erect a costly tower at the populated area, if indeed a mountain or a hill exists at relevant distance.

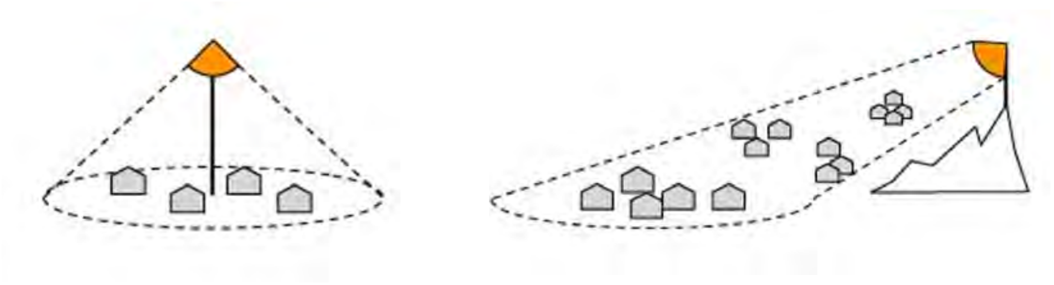

(a) Nearby omni-directional antenna.

(b) Far directional antenna.

Figure 1.4: Similar coverage with different choice of antennas

1.5 Who are WISPs entrepreneurs and some of their challenges

I came across many types of entrepreneurs who start a WISP, and the conditions driving them to attempt this operation are quite diversified. There's obviously a need for a wide-scale study of these attributes, as we must better understand who our target partners might be. In the meanwhile, we present discussion based on anecdotal evidence and my own personal experience.

There's always a real, first-hand, need: it's unlikely that a business-minded entrepreneur would wake up one day and decide to start a WISP for a rural area. Although it appears the market potential is quite high, a graduating MBA would typically envision bigger fish to fry, and certainly lower-hanging fruit to pick. Most WISP's founders simply experience the need themselves; they are from an unserved area, without Internet access, or cannot afford it. My experience suggest that many WISP's founders are strikingly non business-minded, or more accurately, not business-savvy nor business-trained or educated. If I had to suggest a single primary reason for small WISPs' failures, that would be it. That being said, small WISPs are not doomed to fail, their common lack of business experience could be assisted and perhaps remedied by making available sane and well-tested methodologies and models for operating a WISP — a key goal for this study. Ideally, these methodologies should be embedded within the tools and the technology to manage such operations, as the novice WISP entrepreneurs tend to emphasize the technical solutions and often overlook the business and financial aspects until it's too late.

Obviously, WISP founders possess strong entrepreneurial skills. Even if they don't always know to define themselves as such — they essentially are leaders who strive to make things better, and believe there is a solution to a worthy problem. Finally, all WISPs have a good degree of technical inclination, or at least are not afraid to experiment and explore new technologies. I came across completely non-technical individuals who experienced the problem, were entrepreneurial in nature, and understood the market potential, but were unable to go beyond an initial experiment or pilot phase. As soon as basic technical challenges presented themselves (and things get complicated very quickly), they were unable to surmount these themselves, and were unable to recruit technical assistance in their locality, or to afford it. Although the above are typically the main driving forces for entrepreneurs to attempt an ISP business, these are by no means the only required skill-sets for success. Quite the contrary, the WISP business is unique in the demand for versatility and diversity of skills required from its leaders. A WISP operation is characterized by exceptional dynamism, and difficulty in long-term planning; you may plan to serve an area using relay sites at particular locations, but the likelihood of securing these specific locations when the time comes is low. The deployment is so dependent on topography and location, that even the slightest shift may render the whole plan nonviable. Capricious (or simply human) land owners might inadvertently sabotage an enormous effort, by merely asking to move an antenna to the other side of the yard or house, as he/she does not like the color — see figure 1.5 and 1.6 and their captions. WISP leaders who cannot tolerate a high level of uncertainty, and who are not fast on their feet to rapidly adapt to such changing conditions, are not likely to do well.

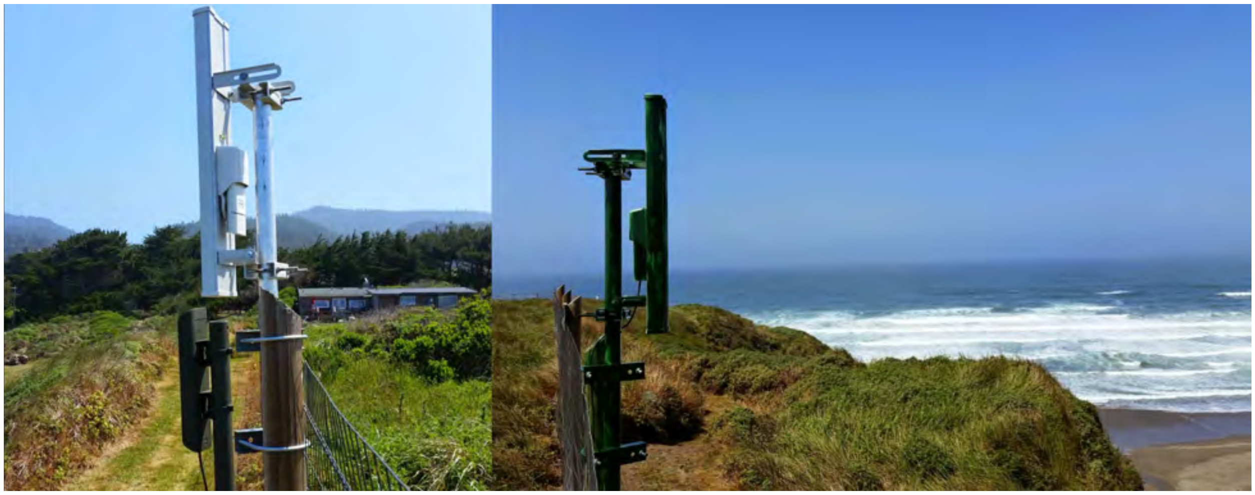

Figure 1.5: Small distributed relay site with low visual impact.

In the left photo is a base-station antenna during an upgrade process. The smaller *NanoStation-M5* device, painted green, below it's replacement sector antenna that is not yet painted. In the background is the landowner's residence with the larger uplink dish on the roof. This is a distributed site, in which antennas are not co-located on the same mast/tower, but rather spread over a larger area. Aesthetics was a prime concern for this site as the base-station antenna is in the view of the land-owner - between his residence and the ocean. Initially, they only allowed an antenna lower then the existing fence post, painted green, and generally not visible from their home with the naked eye. Later on, as they learned to appreciate the service and perhaps to value it more in relation to the minor impact on their view, the landowners allowed the upgrade shown in this photo. The photo on the right is the upgraded antenna, freshly painted, with the older one removed. The cliff over the ocean in the background, suggesting why this antenna is valuable to serve the village across the valley to the right (not shown). Anecdotally, this upgrade did not result in major improvement as was expected. Although a bit higher, a wider angle, and higher-gain sector antenna, it did not improve the signal over the original *NanoStation-M5* unit.

Human nature aside, it's difficult to predict the growth of a rural network. Often, a site that was initially considered an end node, perhaps even considered the end of the world, may become a major core node as the network expands in unforeseen directions. Not only market demand dictates expansion in unexpected directions, but also unforeseen topography features might present themselves, such as the ability to make a long-distance link to a new area from an existing rooftop.

It is common that WISP operators discuss and compare their work to a war-zone situation – they often use military jargon (building a relay is a ”mission”, surveys are ”reconnaissances runs”, etc.). I've met quite a few such entrepreneurs and WISP employees with a military background, and it does not seem merely coincidental. WISP operators should also posses exceptional ”people skills”. The key to a successful WISP deployment is diversity of relay sites. Not only does a denser topology allow for shorter distances, and hence higher capacity and higher quality links, it also decreases the dependability upon any single major site. Securing the cooperation of multiple landowners is probably the most challenging task for most WISPs. Indeed, the majority of WISPs do not diversify much, and put all their eggs in a few baskets: these would typically be preexisting communication towers on which they rent space for their antennas.

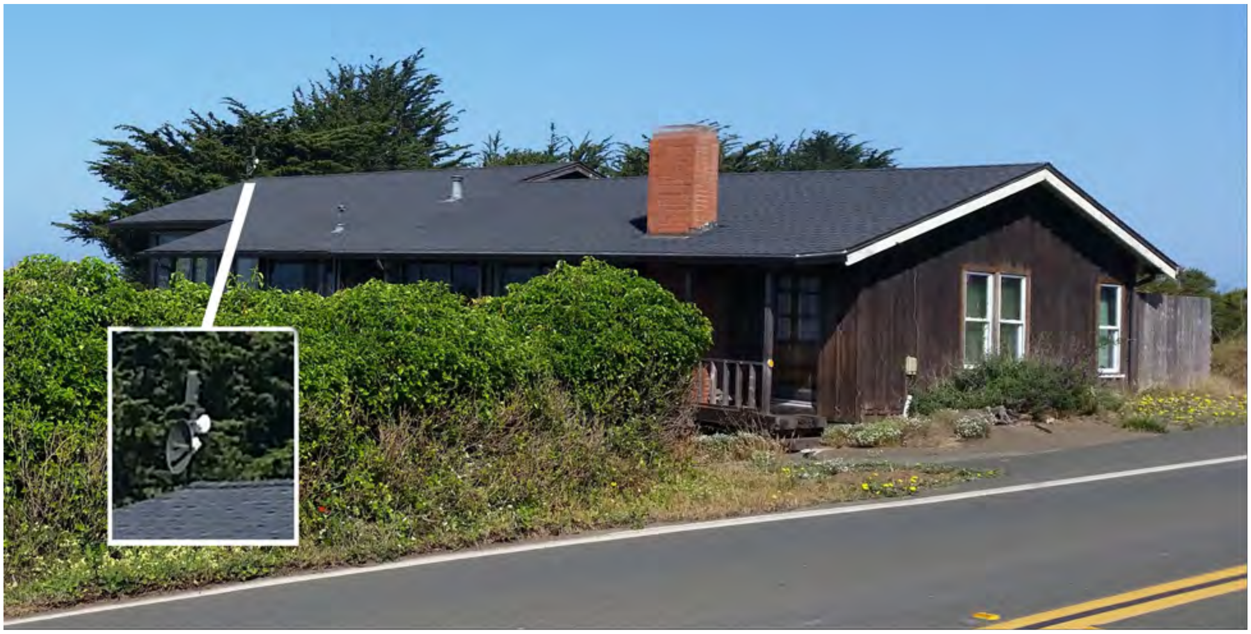

Figure 1.6: A small neighborhood relay site.

From the road or from the neighbor's view, this small relay site is barely visible. As can be seen in the enlarged section, the only element that may attract attention is the newly added, and therefore yet unpainted, small disc antenna. The large (400mm) dish serving as uplink as well as the small rectangular *NanoStation-M5* BTS above it, are painted to merge with the tree in the background and are hardly noticeable.

Use of existing towers presents a key limitation — it's unlikely that WISPs who based their expansion on such dependency would reach far into the rural and sparsely populated areas, which is perhaps where the WISPs potential for social impact is highest. Moreover, if an existing tower is available it would typically mean that competition is also present, at the very least cellular coverage, which could jeopardize the WISP's sustainability.

Intuitively, one may argue that operating multiple sites would make the operation so much more expensive, yet upon deeper consideration the economics generally tilt towards the multiple, smaller, and vastly distributed approach. Rent is costly on communication towers, and "vertical real estate", i.e. tower space and mounts, are scarce, not to mention the high levels of radio interference. Erecting a large communication tower is probably an insurmountable barrier for most small WISPs, given the costs of the tower and its erection, not to mention the lengthy (often never ending) permitting process. Multiple small relays on rooftops of subscribers, or in their yards, are likely to be cheaper on all fronts. Such a distributed approach would not only offer more capacity and shorter links (also smaller and cheaper antennas at subscriber premises), but would result in a resilient network that is less dependent on a single main site. Maintenance of gear on high towers is also costly, as it requires trained tower climbers and special gear, while the back-yard-relay approach may not need more than a simple household ladder.

Identifying good relay locations, a good understanding of topography, and mastery of GIS tools, are another set of required skills for WISP operators. However, successfully negotiating a reasonable land-use deal (once the land owner is even found) for multiple sites, is a skill not often easy for the average WISP operator and typically does not always reside in technically strong founders. When securing these deals, a good degree of legal understanding and document preparation is often required. Even if the WISP cares less about a binding legal contracts (more on that later), the landowner might insist to have one in place. The inability to draft and adopt such legal documents in a timely manner while the iron is hot, could be detrimental to WISP's success in securing critical relay sites.

Generally speaking, where it comes to land-use contracts, especially when done in exchange for free-Internet service and nothing else, the contractual grounds gets murky. The landowner rightfully feels that they're not getting much in return for a commitment on their end, while the WISP would be worried to make an investment without guarantees. However, with the exceptionally short ROI for WISPs there should not be a major concern. Essentially, for a small roof-top relay, its very likely the WISP would recoup the costs for such a setup in 3-6 months. It's unlikely the realities behind the landowner's willingness would change in such a short time, so the risk to the WISP is low.

Moreover, the WISP would probably only setup such a relay when subscribers are lined-up to get service from it, thereby further shortening the ROI time. There are typically more incentives to keep this arrangement going, neighbors being served by the relay will become unhappy if the landowner terminates it, and thereby further lowering the risk for the WISP. I'm therefore under the assumption that WIPS should not bother with contracts when building a small relay site, when the exchange is free or discounted service to the land owner. In fact, there could probably be more that the WISP could lose from such a contract if things go south. At the same time, the landowner might benefit from a contract, and especially from inclusion of the site under the WISPs general liability insurance (assuming the WISP has one). Naturally, as with other things, there's no one-size-fits-all approach for binding legal contracts, but essentially these collaborations are based on short-term mutual interest. As soon as a new WISP comes around with cheaper or better offering (or willingness to pay real money for the site), the mutual interests are likely to end, and binding long-term contracts are unlikely to be agreeable to the land-owner in the first place; The best the WISP could hope for is ample warning to allow the move of the subscribers without down time.

Chapter 2

FurtherReach

2.1 Introducing FurtherReach

FurtherReach is project by the non-profit *De Novo Group* in which we deploy, operate and maintain a large-scale wireless network that brings broadband Internet access to rural communities at the coast of Mendocino county in California. The project was lunched through a generous $2 million grant from Google.Org with the emphasis on networking technology research that is well grounded and tested through the FurtherReach network deployment. The research elements of this project are through a partnership of De Novo Group with *UC Berkeley* and *Stanford*.

The non-profit company De Novo Group bridges the gap between research and impact by turning novel university research into applicable technological solutions to address the needs of deserving communities.

In summer 2013, upon receipt of the grant we announced[1] our intent to lunch a network deployment in northern California and called out to communities who are unserved. We received hundreds of letters from such communities asking us to chose them for the project. Although the term “community” in this context is quite vague, these letters always came from a group of people who, at the very least, considered themselves a community by their own definition. The sheer number of requests was surprising to us — we did not expect such a large unserved population in northern California, and only a three-hour drive from San Francisco, as was specified in the announcement. The many letters were both encouraging — as it demonstrated the need and the potential for impact from our work, so close to home — but also overwhelming and made the choice quite difficult. The key driver for choosing Manchester, California as our starting point was due to the sizable and yet well organized effort from that community. They were also very vocal regarding their need of broadband service and overall louder than most other communities who

[1] The main announcement and call for communities was through an SF-Chronicle article: http://blog.sfgate.com/techchron/2013/10/22/google-grants-2m-to-researchers-bringing-broadband-to-rural-california/

approached us. The availability of high-quality, affordable, bandwidth from *Level-3 Communications*, thanks to the trans-pacific fiber-optic cable landing station at Manchester certainly helped tilt the scale towards that community.

2.2 Business model and finances:

It's not by chance that I begin this chapter with a discussion of FurtherReach's financial model; albeit we are a nonprofit driven to make social impact, we understand well that nothing of lasting impact can be achieved without a long-term stable operation based on sane financial planning. We've seen too many well-intended initiatives that survive only as long as grant money is being pumped into the operation, but crash miserably when external support seizes. In turn, such a sudden demise often results in more harm to the community than not starting in the first place. The South Mendocino coast area where FurtherReach chose to begin, had seen two cases of sudden demise of previous WISP operators, first Esplanade and recently CVC, leaving a painful vacuum that was very harmful to the community. The account of Esplande's demise and the resulting harm to the community is narrated in Greg Jirak's report to the Mendocino board of supervisors[7]. CVC's pull out of Point Arena area was less harmful thanks to FurtherReach's ability to quickly feed most of CVC's relay sites and provide service to their subscribers without interruption (and free of cost). CVC also gave early notice of their plans to terminate the operation, about two weeks ahead of time.

At FurtherReach we understood that financial sustainability for a long-term operation is the key goal, and maintaining this emphasis superseded all other objectives and decisions. At the time, what seemed to be an endless internal discussion, and a heated one, was about the prices of the plans we'll offer, especially for the basic plan. While these decisions were based on detailed planning, like every business plan, these are essentially "guesstimates" and resulted in a plausible range not a fixed number. We decided to choose the higher end of each range, knowing that choosing too low might lead to disastrous results and harm the community in the long term. In hindsight, I am very happy for that choice, and even these seemingly costly plans may still turn out to be too low should we come across some unforeseen expenses like a natural disaster or a lawsuit, for example.

How FurtherReach seamlessly kept CVC's subscribers alive
Central Valley Cable (CVC), is a cable-TV operator serving Gualala and The Sea Ranch area. Following the sudden demise of Esplanade[7], in February 2011 and until August 15th, 2014 — CVC provided limited, wireless Internet access to subscribers in and around Point Arena, Manchester, Irish-Beach and even north as far as Elk. Towards the end of July 2014, CVC announced to their wireless subscribers, that service would seize on August 15th, due to their inability to renew their lease for their main tower. FurtherReach was just starting, and we had our share of challenges, but decided it's our duty to feed CVC's secondary relay sites as to ensure their subscribers don't go dark. Getting high capacity bandwidth to feed CVC's sites was demanding and the team worked around the clock to achieve this goal. On August 15th, the day they shutdown their main tower, FurtherReach began feeding CVC's routers at their secondary relay locations. Not only the service to subscribers was not interrupted, the bandwidth FurtherReach provided to each of these towers was far greater than CVC ever had at all of their sites combined. Moreover, the links between CVC's relay sites and their management[2] server was severed, resulting in the individual bandwidth restrictions for each subscriber being lifted. Needless to say subscribers were happy, not to mention they no longer had to pay for their service, nor would they know who is running the service. Unless we had kept CVC's subscribers alive that way, not only would the community be harmed, but FurtherReach would get swamped with requests for new installs, far beyond our ability. By continuing service to CVC's customers, we could focus on our planned growth, and solve the problem for the handful of businesses who were served directly from the main CVC tower which was no longer accessible. We had hoped that CVC's subscribers would approach us to become our subscribers and pay for their service, but most, if not all, did not — why would they volunteer to pay for what they get for free? We did not had the contact information for these subscribers, although we were responsible for their service. All we had are the short device names which CVC used for their CPE antennas - most of which were meaningless to us. We therefore started to disconnect these subscribers, one by one, until they contacted us when they lost the service and asked to be re-installed as FurtherReach subscribers. A few months later we have completely shut down all of CVC's former sites, as no subscribers were served off of them any longer.

In rolling out the FurtherReach deployment, we tried to simulate, as much as possible, the constraints and environments faced by a novice WISP founder. Not that we would deliberately make unwise business decisions, just for the sake of simulating an error that a novice entrepreneur might make, but we did ensure as nimble as possible operation from the beginning.

That being said, we have made many unwise business decisions, and some mistakes we probably do not even know about yet. I don't consider myself a seasoned businessman, but I am probably a notch more experienced than most novice rural WISP operators, and I had much help on all fronts.

It is also an unfair comparison, or skewed simulation in many other ways — being well funded (a $2 million Google grant) certainly eased our decision making process. Unlike small entrepreneurs who are risking their own dime, we had little personal risk, and the lower anxiety levels are known to critically help lead to better decisions.

Nevertheless the total direct investment in the FurtherReach deployment was less than $254,000 as shown in table 2.2, thanks to revenue generated from subscription fees. About $146,000 went to compensation of local staff, approximately $110,0000 went to hardware installed at relay sites, $38,600 went to equipment at subscriber's premises and about $41,000 went to tools. In addition $69,000 went to miscellaneous expenses, such as bandwidth, insurances, rent, gas and travel, cloud computing, etc.

These are not unattainable figures for even the smallest entrepreneurs (by North American standards). Admittedly, my own salary was paid from a different accounting bucket, which is unfair as I spent more than 50% of my time, during the first two years, solely on FurtherReach.

More important is the monthly breakdown in spending for the FurtherReach deployment, versus the growth in number of subscribers — this bears directly on risk-taking and critically affects investors' anxiety levels. As the install base begins to grow, one can better predict further growth and the time to recoup the investment. It is no longer a shot-in-the-dark guess as subscribers are happy and paying (or willing to pay as in the case of FurtherReachs first 3-4 months of beta), and financial estimates get refined and anchored in reality, leading to safer spending on further growth.

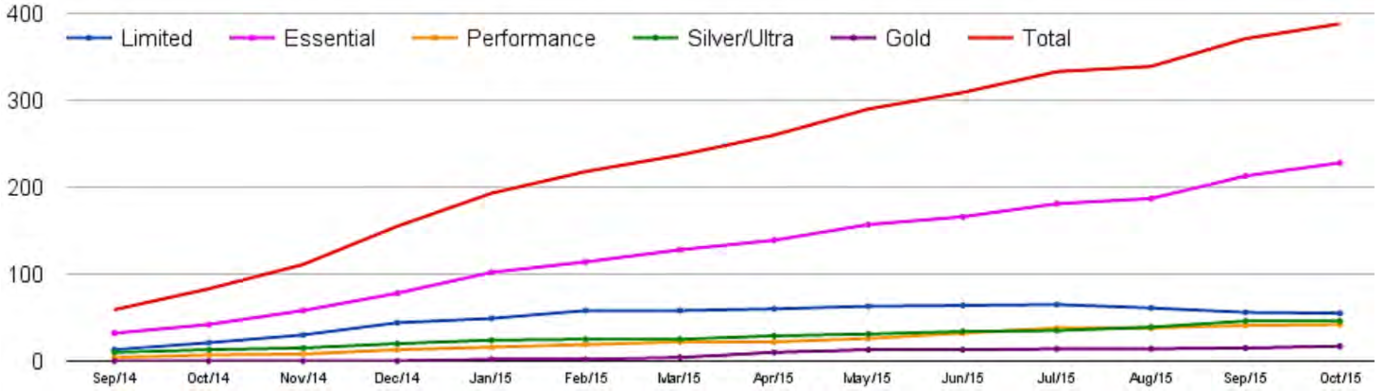

Figure 2.1: FurtherReach subscribers, per plan, over time.

Figure 2.1 depicts Furtherreach's growth in number of subscribers from September 2014 until October 2015 (the time of this writing). We began to connect subscribers in July 2014 but did not yet have our *Celerate* controller (discussed later in this chapter) to track these installs — by the end of September we already had 59 subscribers connected, and I estimate about a third got connected in July and August. Our beta program ended on September 30th. Subscribers connected during the beta program did not pay for the service nor the installation, which helped speed-up new installs and resulted in excellent public relations that helped drive adoption rates for months later. Overall, we could think of the beta program as a major promotion campaign. The cost of that campaign, if we take into account about 55 free installs (4-5 or five are at relay sites who continue to be free in exchange for land use), would be $8250 and approximately another $7000 in unearned revenue from subscription fees during the beta program (obviously there are other costs we incurred during the beta, but most we would have been forced to incur without it as well). Arguably, $15,000 is quite a costly promotional campaign for most new entrant WISPs, yet I believe it has been a wise approach for FurtherReach, as would it be for many young WISPs — I would therefore encourage starting WISPs to consider such approach and overall not underestimate marketing campaigns, as they typically do, especially during the early months.

	Limited	Essential	Performance	Ultra	Silver	Gold
Monthly cost	$30	$70	$100	$130	$130	$200
Minimum speed	*	4Mbps	8Mbps	15Mbps	15Mbps	40Mbps
Typical speed	*	8Mbps	15Mbps	30Mbps	30Mbps	60Mbps

Table 2.1: FurtherReach's plans.

* As of summer 2015, we no longer offer the "limited" plan. This plan, which we referred to as the *subsidized plan*, was designed to provide reasonable web-surfing experience, VoIP telephony and limited video streaming, such as low-quality YouTube videos. In practice, this plan was limited to 450Kbps (in both directions) with a handsome buffer for bursts of 2500KB. We had hoped that only low-income families would sign-up for this plan, given its stated limitations, and we failed to enforce any restrictions on who is allowed to sign up for it. In reality, many subscribers signed up for the *limited* plan, and it was especially popular among vacation rental properties and part-time residents. Sure enough there are quite a few low-income families among them, but not the majority. This was harmful to FurtherReach in two ways — not only we loose money by offering this plan, as our perceived[3] average cost for providing service to a single subscriber hovers around $65 per month, but the slow speeds resulted in negative impact on our public relations and contradicted our ongoing campaign advocating high speeds. It is especially harmful at the vacation rentals, where multiple guests complain about the slow speeds, thinking it's our network and not the poor choice of plan by their landlord. After no longer offering the *limited* plan, we asked subscribers to voluntarily upgrade to *essential* — less than a handful did. We still have many subscribers on *limited* who should not be. In the near future, we'll need to find a solution for faster service to low-income families, at the same or similar price point, and also to terminate this plan for non-qualifying families - even at the cost of losing these subscribers.

As seen in figure 2.2, we only started to charge subscription fees from October first, at the end of our beta program, and since we bill backwards, for services rendered, unlike the industry standard of billing before service is provided, our first payments from subscribers only came in during November 2014. With 59 subscribers in September 2015, who did not need to pay installation fees but only subscription fees based on the ARPU at the time of $60, and another 24 subscribers connected in October, that should have also paid the one-time installation fees of $150, the revenue for November was expected to be about $8500, yet the actual was short of $7000. Revenue in December, with 111 subscribers at the end of November (28 new installs in November) was supposed to be almost $11,000 yet plummeted still to just a notch over $6000.

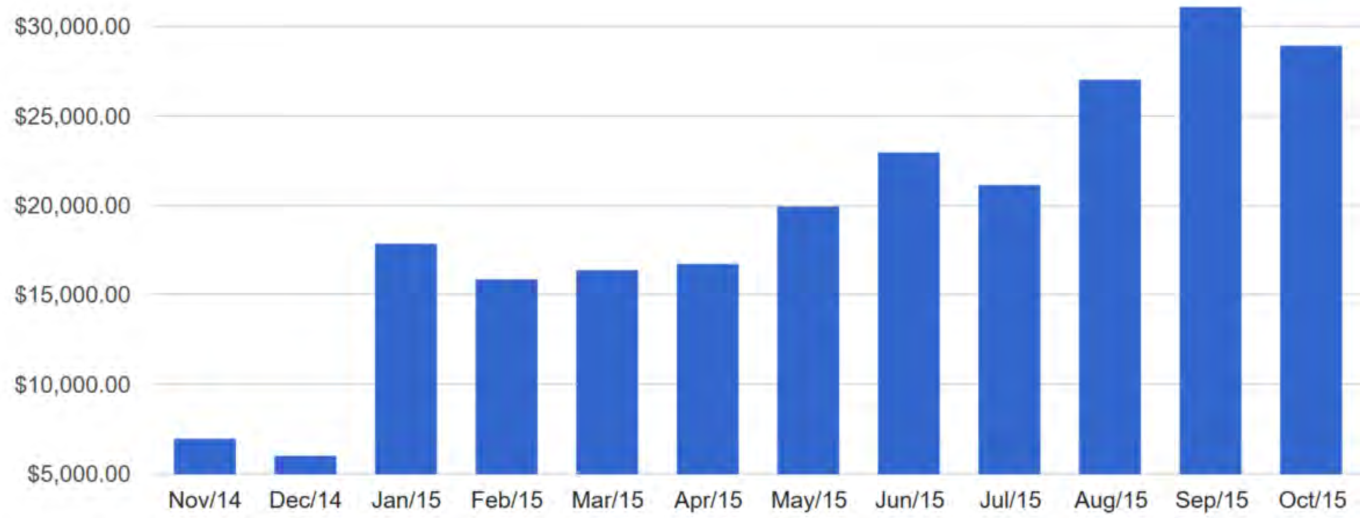

Figure 2.2: FurtherReach revenues from subscription and installation fees.

There are a number of reasons for the phenomenon; first, we had a number of non-paying subscribers from the beginning — these are non-profit organizations, landowners of core relay sites, and landowners in who's back yards we've buried our fiber-optic cables. Then, there are at time difficulties with collection of payment — subscribers who forget to pay, subject to temporary financial crises or just don't pay for unknown reasons. We have underestimated this problem and the resources we now must devote for collection — naturally the problem worsen as the number of subscribers grow. Finally, we believe the main reason for the delayed payments in November and December of 2014 (the two first months in which subscribers paid), was due to subscribers' misunderstanding our policy of paying for service rendered in previous month. Although this is clearly noted in the terms and conditions, people don't actually read these, which is not surprising. This policy, that we believe is in favor of our subscribers as it gives them freedom and shows our trust in them as a community project, seems to be hurting us — contrary to our initial expectation —- people don't read and don't understand that policy, not even after paying for quite a few months, but are getting confused by it when asked to pay past-due fees or whenever any sort of billing issue arises. In a few cases, it appears this policy contributed to subscribers accumulating an unpaid debt which became difficult for them to pay off — we believe these unfortunate unpleasantness could have been avoided if we used the traditional "pay before" policy. Overall, countless billing-related support tickets may have been avoided if we had not chose this unconventional policy, not to mention the resources we have to put into collection — to me, this was quite counter intuitive and surprising. We plan to change this policy and adopt the industry's standard of collecting payment before service is provided, yet this is quite challenging to do for existing subscribers. It's not until

January 2015 that we've finally caught up with the delayed payments from previous months, as seen in figure 2.2.

> As FurtherReach grew beyond its initial service areas of Manchester, Point Arena and Elk, we begun to see a rise in the ARPU. It appears our subscribers in the Albion area, especially Pacific Reefs, as well as the newly connected oceanfront properties south of Point Arena, are more affluent and opt for costlier plans. At the same time, we no longer offer the *limited* plan, and also got some *limited* subscribers to upgrade, which obviously contributed to the higher ARPU.

As with any infrastructure businesses, the beginning is hard and requires an upfront investment, but once rolled-out, maintenance and support costs are relatively low (and lowering these is a core aspect of this study) and thereby enables both ROI and continued growth. The beauty of a WISP operation versus other infrastructure projects, is in its much lower initial investment as well as the potential for exceptionally fast return on investment. The balancing act, however, is how to ensure continuous growth without periods of slow-down to recoup investment in which some staff needs to be let go and later re-hired (if possible). This model lends itself well to a tiny operation of a one or two partners that could afford to tighten their belts initially, knowing revenues are being generated. For FurtherReach, which employs more people, this is the key challenge, yet we had the fortunate ability to make an initial investment larger than what most small rural WISPs are capable, and thereby sustain continuous growth without letting people go.

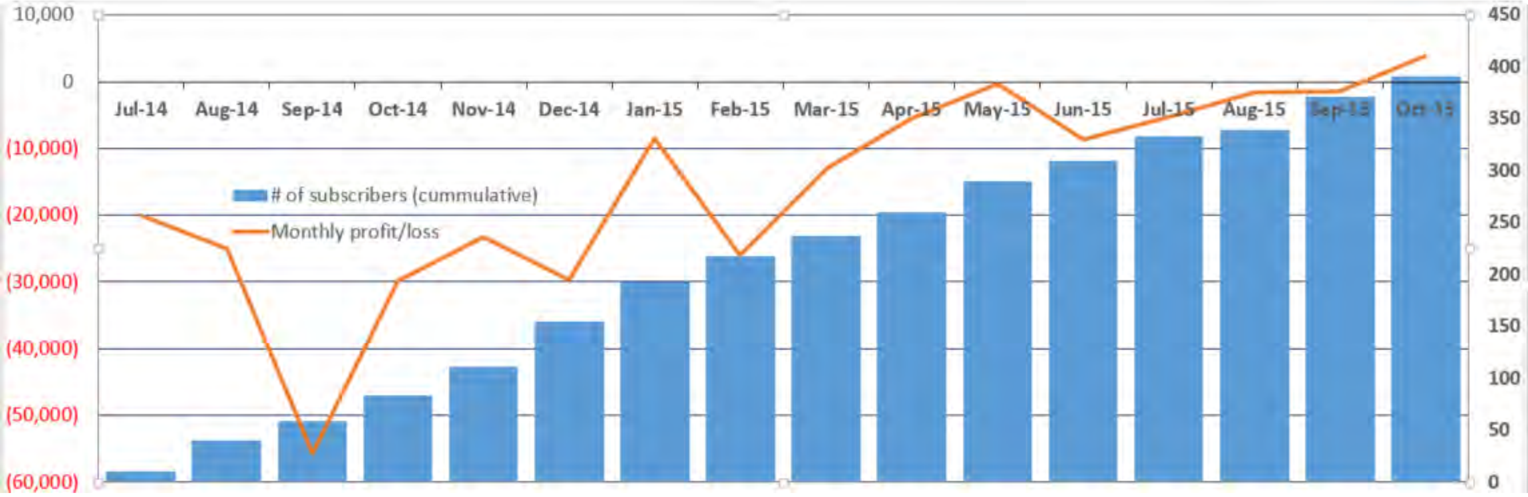

Figure 2.3: FurtherReach Profit and Loss graph along with subscription growth.

As shown in figure 2.3 the first month in which we've seen profit is October 2015. That said, we are far from ROI and we intend to attempt further growth before we begin paying back the CAPEX and achieve ROI. There are also other indirect expenses incurred by De Novo Group for the FurtherReach project which are not shown in these figures at all. Some are quite substantial, including my own salary and that of other De Novo Group officers. The numbers presented here are strictly FurtherReach's direct expenses.

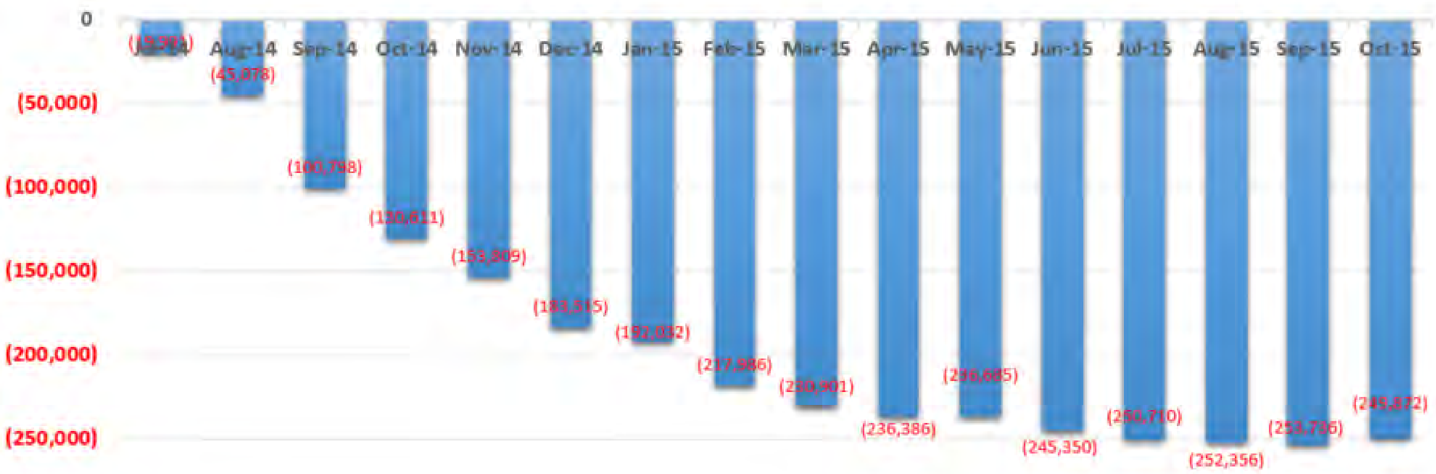

Figure 2.4: FurtherReach's Cumulative (negative) cash-flow graph.

As for CAPEX, the WISP business is quite different than other infrastructure projects. There are little initial expenses before actual roll-out to subscribers. Unlike traditional approaches to broadband delivery, there is no need for a preliminary phase of permitting, and no complicated and time consuming trenching for fiber-optics cables is required.[4] The CAPEX for most WISP operations could be thought of as all the expenses during it's initial growth phase, and OPEX begins once growth has stopped. Alternatively, CAPEX could be thought of as the investment for a particular service area — once served there is a switch to OPEX and growth in adjacent areas would be accounted separately, just like a separate project. With FurtherReach we have not reached these turning points yet — we continue organic growth without slowing down and without declaring a particular area completed. If we are to make such distinction and draw a line at the end of October 2015, for example, we could extract our CAPEX for the initial service area from figure 2.4 — it would be $250,000. It appears however, that we did reach an important financial milestone — financial sustainability. As seen in figure 2.4, the 3rd quarter of 2015, as well as the month of October, show that FurtherReach no longer increase its debt while maintaining continued growth. We expect that through this growth FurtherReach would no longer be depended on aid from De Novo Group and will be able to pay for the time put-in by the parent organization's officers, not to mention pay for other indirect costs incurred by De Novo Group, such as bookkeeping and accounting. At some point FurtherReach is expected to begin paying back for it's investment by De Novo Group, which could allow the parent organization to explore replication of the project in other non-adjacent areas.

[4] FurtherReach also laid fiber, but our first was only a couple of thousand feet long, needed approval from two landowners with no public hearings or permitting; the work was done in a single day and equipment cost was less than $10,000 overall.

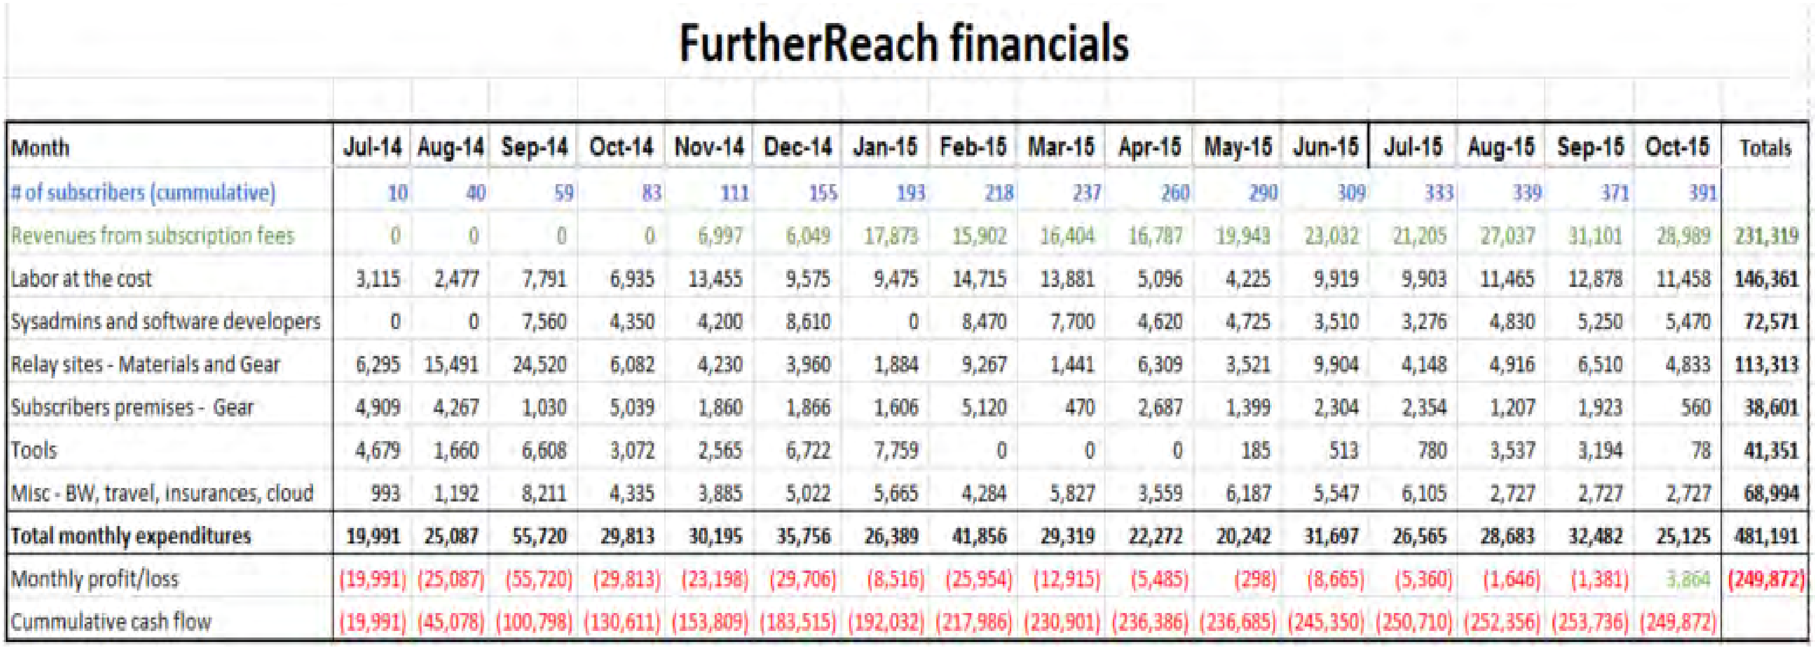

FurtherReach financials

Month	Jul-14	Aug-14	Sep-14	Oct-14	Nov-14	Dec-14	Jan-15	Feb-15	Mar-15	Apr-15	May-15	Jun-15	Jul-15	Aug-15	Sep-15	Oct-15	Totals
# of subscribers (cummulative)	10	40	59	83	111	155	193	218	237	260	290	309	333	339	371	391	
Revenues from subscription fees	0	0	0	0	6,997	6,049	17,873	15,902	16,404	16,787	19,943	23,032	21,205	27,037	31,101	28,989	231,319
Labor at the cost	3,115	2,477	7,791	6,935	13,455	9,575	9,475	14,715	13,881	5,096	4,225	9,919	9,903	11,465	12,878	11,458	146,361
Sysadmins and software developers	0	0	7,560	4,350	4,200	8,610	0	8,470	7,700	4,620	4,725	3,510	3,276	4,830	5,250	5,470	72,571
Relay sites - Materials and Gear	6,295	15,491	24,520	6,082	4,230	3,960	1,884	9,267	1,441	6,309	3,521	9,904	4,148	4,916	6,510	4,833	113,313
Subscribers premises - Gear	4,909	4,267	1,030	5,039	1,860	1,866	1,606	5,120	470	2,687	1,399	2,304	2,354	1,207	1,923	560	38,601
Tools	4,679	1,660	6,608	3,072	2,565	6,722	7,759	0	0	0	185	513	780	3,537	3,194	78	41,351
Misc - BW, travel, insurances, cloud	993	1,192	8,211	4,335	3,885	5,022	5,665	4,284	5,827	3,559	6,187	5,547	6,105	2,727	2,727	2,727	68,994
Total monthly expenditures	**19,991**	**25,087**	**55,720**	**29,813**	**30,195**	**35,756**	**26,389**	**41,856**	**29,319**	**22,272**	**20,242**	**31,697**	**26,565**	**28,683**	**32,482**	**25,125**	**481,191**
Monthly profit/loss	(19,991)	(25,087)	(55,720)	(29,813)	(23,198)	(29,706)	(8,516)	(25,954)	(12,915)	(5,485)	(298)	(8,665)	(5,360)	(1,646)	(1,381)	3,864	**(249,872)**
Cummulative cash flow	(19,991)	(45,078)	(100,798)	(130,611)	(153,809)	(183,515)	(192,032)	(217,986)	(230,901)	(236,386)	(236,685)	(245,350)	(250,710)	(252,356)	(253,736)	(249,872)	

Table 2.2: FurtherReach's expenditures, along with revenues and number of subscribers per month.

Table 2.2 details FurtherReach's direct program expenditure for the deployment at the coast, as well as labor costs for system administrators and the majority of software development for the Celerate controller discussed in 2.6.

Not surprising for an operation in North America, the expenditures are dominated by labor costs, and that is without including the compensation for De Novo Group's officers. Such a labor intensive operation is quite attractive for developing nations where labor costs are substantially lower — an important observation as these emerging regions are the prime target for WISPs and where the majority of the *last billion* resides.

At the time of this writing, it is obvious that long-term financial sustainability for FurtherReach has been achieved. Moreover, the operation continues to grow organically and the high risks we initially faced are mitigated slowly.

AT&T has a central office (CO) right in center of Point Arena. This CO is fiber fed and there is no shortage of capacity. It is also in proximity with most businesses in Point Arena (about 20-30) as well as 50-60 households. The proximity of the CO to subscribers makes it technically ideal to offer ADSL services to this community. However, AT&T chose not to do so, regardless of the many official requests made over the years, both by the city, the county as well as by multiple groups of activists. I believe the reason behind AT&T's decision is purely business related, although AT&T routinely cites poor infrastructure and capacity challenges. Before FurtherReach, the majority of the businesses in Point Arena used T1 lines from AT&T. Typically, a business would purchase a bundle of two T1 lines to gain reasonable speed,[5] and to enjoy better availability. The cost of a single T1 line is $420 per month and a bundle of two goes for $700. ADSL service for businesses cost $50 and residential goes for $30, this means that 8 to 14 times more ADSL subscribers are needed for AT&T to generate the same revenue from ADSL as they do from T1. Point-Arena, especially nearby the AT&T CO, does not have such a growth potential and thereby ADSL is not offered. There is a good likelihood that now, as AT&T has lost most of their T1 lines to FurtherReach, the telco giant will decide to offer ADSL services in Point Arena — a long overdue, healthy competition and choice.

2.3 Design principles

The design principles of FurtherReach were quite different from those of a novice WISP entrepreneur. While for De Novo Group, the ultimate goal is to learn from the network and use it for research and development, we pledged not to harm the subscribers and the service — therefore long-term sustainability took precedence in decision making over research and development activities. The principle of sustaining the network for the long-run is well aligned with that of a typical WISP entrepreneur who embarks on such initiative in order to generate profits (or at least support his family — some operators I have met had strikingly limited expectations for profit making, unlike most young entrepreneurs whose dreams are much more inflated). Overall, while the goals for sustainability are similar, the design principles are not.

With FurtherReach, we strove to innovate and challenge every common practice that is acceptable and prevalent in the WISP industry. We ended up using much good that existed, but not without challenging everything first. Finally, we're proud to have designed and built a network that I believe is superior in many ways, developed innovative software tools to support and maintain it, and produced refined methodologies, procedures and best practices for designing and building a WISP.

From the beginning, we wanted to deliver high bandwidth to subscribers. In contrast, most WISPs deliver 0.5 to 2 mbps plans, which are received warmly by the deprived subscribers (in our area there are still WISPs that offer 256Kbps speeds). Moreover, they, like the rest of the industry, use the term "up to that speed", without giving any lower-bound guarantees. We set as a minimum the new FCC definitions of broadband: 25Mbps down and 3Mbps up. Moreover, we wanted to be able to offer plans on par with the best plans available in the cities, in the order of 40-60 mbps and even faster. We didnt expect to offer these plans at the same price point of similar plans in the city, nor do we need to, as there is no competition — not for these fast speeds, although there are other WISPs operating in our vicinity who offer "up to" 4mbps download plans and "up to" 512Kbps up. The goal was to demonstrate the feasibility of delivering such speeds in these sparsely populated, rural settings, as well as to future proof our technology. The WISP industry has an unfortunate reputation of providing slow as well as unreliable service, but it's accepted given there are no feasible alternatives in these rural areas. We insisted on breaking both of these stigmas; FurtherReach was designed to be fast and reliable, as well as support future speeds.

> The Manchester, California cable landing station was constructed in 1956. The official story about the choice of location is for being the closest direct point to the Hawaiian islands — the far end of the coaxial cable put to use in 1969 and provided long-distance telephone service to Hawaii. The less told story, perhaps nothing more than a local legend for the choice of location, is that the cable was dropped into a deep underwater canyon to protected it from Russian spy submarines. Indeed, the large San Andreas fault goes offshore at Alder Creek, less than a mile north of the cable landing station, and it's plausible that it continues into the ocean in some form of a canyon or rift (I never researched this further - why spoil a good story). The coaxial cable was later upgraded to a fiber-optic cable in 1988, with an additional cable added in 1989 and yet another cable ring *Japan-US* constructed to connect with Vancouver and Japan in 2000.

One of the reason for choosing the South Mendocino coast as the service area for FurtherReach, was the availability of high capacity and high quality bandwidth, at wholesale prices, locally at a Level-3 Communications data center in Manchester, CA (there have been other criteria for choosing this locality, and these are discussed later).

As with most ILAs (In-Line Amplification Units), the Manchester facility has no IP-transit services, so FurtherReach purchases from Level-3 two products: A 10GE transport (aka wavelength) to their data center in San Francisco and a gigabyte of IP-transit (at their SF facility). We have negotiated an excellent deal if one looks at the price per megabyte, but our desire was only to buy 100Mbps and not ten times that much, not to mention the costly transport. (The details of the deal with Level-3 are under an NDA and cannot be disclosed, this goes to exemplify the competitive nature of this market, even in such a rural area, where nobody else is likely to ever buy such services from Level-3 or from a competitor).

The negotiations with Level-3 were quite interesting as well. Although we are under a NDA and cannot expose the details, here are some anecdotes:
We negotiated with Level-3 for the purchase of four different products:

1. Wavelength to San-Francisco
2. IP transit at the SF facility.
3. Co-location for our gear at their Manchester facility.
4. Roof access/usage for our antennas on their facility in Manchester, CA.

In the negotiations, we pressed hard to lower the cost of bandwidth and transport, while we allowed the ridiculously high costs of colocation and roof usage to remain. Once a contract had been drafted, we asked to remove the two lines items describing roof use and co-location and only leave the transport and IP-bandwidth. The Level-3 sales team could not envision we would deliver our fiber into their Manchester facility and would thereby not need co-location nor roof usage. This is exactly what we did. Today, we have three fiber cables going into that facility, we own these cables and we've laid (buried) them with our own hands. We do not have a single antenna on the Level 3 roof and instead we have three core sites, all fiber fed from Level-3, where we have our antenna farms.

The nature of the deal with Level-3 strengthened our design goal for a high-speed network. We were forced to buy so much more bandwidth than we wanted, and we are committed to this deal for three years, so we might as well try and deliver as much of this (paid-for) bandwidth as possible to our subscribers.

In addition to high speed, the next design principle was on availability and continuity of service. WISPs in general, and local operators specifically share a bad reputation when it comes to their service uptime. The alternative of Internet via satellite is similarly pretty bad, and often goes down with fog and especially during heavy rains. Even the aging AT&T-owned copper at the coast is not doing well, and many of the T1 lines offered to local businesses, as well as simple land-line phones, go down frequently, especially after heavy rains or strong storms. At FurtherReach we pledged to counter these stigmas and provide a highly available, resilient service that is not affected by the harsh weather conditions of the area. Finally, let us not forget, we had to accomplish this while emphasizing affordability of the service, at least on-par or below the price point of satellite Internet, and while ensuring the long-term sustainability of our operation.

2.4 Design meets practice

To achieve the desired availability we had to design every part of the core network with redundancy in mind. Its acceptable to have a single site go down, taking with it subscribers that are directly fed from local base stations, but the rest of the network should not be affected (or at least not go down completely) if any single site goes down. The desired network topology therefore, would be of a ring, with each relay site linked to at least two other relay sites. Moreover, increasing the availability of any single site is important as well. Therefore, where a ring topology is unfeasible due to topography constraints, a relay site would be linked to its single neighbor using at least two parallel wireless links. On the face of it, this seems like a costly choice that might risk financial sustainability, but in reality, the cost of modern outdoor radio gear and antennas, such as those from *Ubiquiti Networks*[5], is surprisingly low. When taking into account other costs associated with site acquisition and development, the actual radio gear and antennas end up being a small piece of the pie. This realization maintains that we should use as many radios as makes sense at each and every relay site, to ensure dense interconnection, higher capacity and high availability, all while not inflating the costs relative to the overall cost of the relay site. Naturally, being able to load-balance the traffic over these multiple links was a challenge, and a novel solution needed to be developed for that.

Moreover, co-locating multiple radios presents the challenge of spectrum scarcity as we need to avoid overlapping frequencies from mutually interfering. Contrary to common belief, spectral scarcity may become a serious problem also in rural areas, and our approach of dense radios collocation makes this more challenging. This problem gets worse the more relay sites are within overlapping range, yet we desire such dense topology as we enjoy higher capacity over shorter links, and antennas are cheaper and smaller when used for shorter distances. Finally, the physical space required for mounting of multiple antennas often ends up affecting the costs significantly. Especially if large antennas are needed for long range the "tower real-estate" becomes quite an issue. I later discuss our solutions for these challenges.

Although it may seem intuitively wise, perhaps even obvious, most small WISPs do not distribute their relay sites and often relay upon a small number of critical towers (typically a single one). Moreover, most WISPs do not use resilient topologies nor setup parallel links between sites. I believe this has to do with the "low profile" mentality of most WISP operators, coupled with the tolerance for low quality service by starved subscribers, and the desire to minimize investment and speedup ROI. Being cash-strapped and needing to support one's family does not help long-term thinking — a typical "rural", ad-hoc, business approach. At FurtherReach, we were not to compromise these design principles yet we did not anticipate the magnitude of the challenge to acquire relay sites, especially at the beginning.

Wireless links aside, other components in a given site can fail as well. Power supply systems and wired networking gear are good candidates, not to mention the masts and towers themselves might be damaged during storms and other calamities. Nevertheless when weighing the low likelihood of failures for such components (relative to radio links) against their cost, the decision was to forego redundancy for these systems. Instead the emphasis was on availability of spare parts (sometimes on site), rigorous monitoring, and clear procedures for repair and recovery. This is a known trade-off between MTBF (Mean Time Between Failures) and MTTR (Mean Time To Repair) — both affect downtime, but we have more control over MTTR once MTBF gets to a certain level. That said, we decided to ensure ample battery backup for each site. Even at sites where grid-power is available and considered reliable, we knew that minor surges and spikes on the grid would get equipment wedged, while major surges would damage our equipment[6]. Surges and spikes aside, we knew to expect prolonged, albeit rare, power outages due to storms. The area is known to experience at least one major outage every winter, and in our first winter we had two multi-day power outages, yet our system remained fully operational, unlike every other communications network in the area, including copper-based land-lines. DC-powering a site is an involved practice that adds substantial cost, typically more than the cost of all the networking equipment at the site, and therefore a painful decision from a business perspective that prioritizes long-term quality over faster ROI. In addition to large batteries, there is little value if their state of charge if not remotely monitored to alert about their status. At the very least we need to know if a battery bank is being charged or discharged, and at best we like to know the remaining capacity as well as other power parameters, especially important for off-grid sites. Finally, we desire smart and robust chargers for our core sites, to both prolong the life of the costly batteries[7] and to reduce failures by overheating or AC power surges.

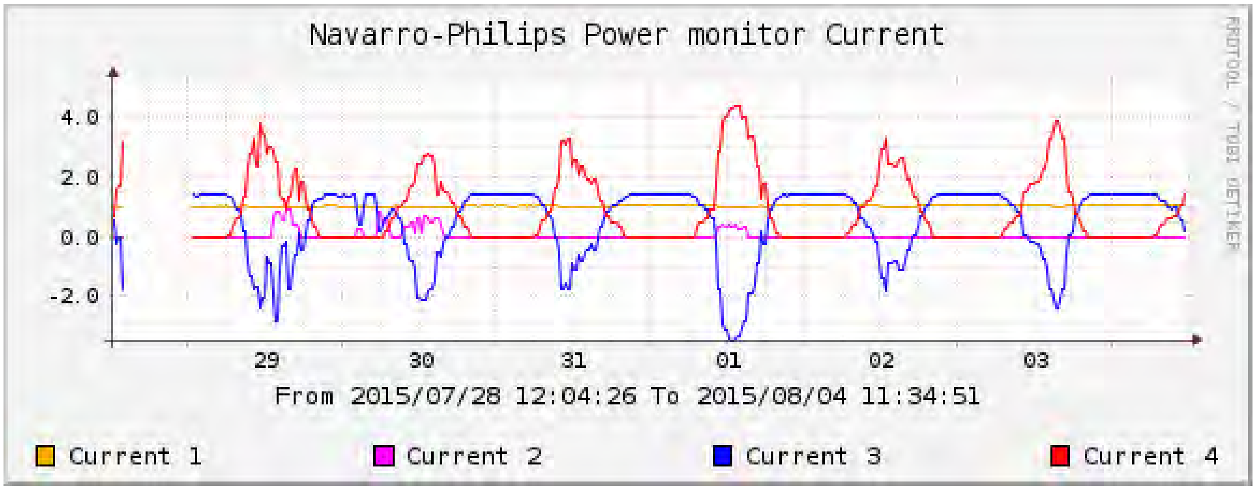

Figure 2.5: DC-currents for an off-grid site — hybrid solar and wind.

Figure 2.5 graphs the DC currents at one of our off-grid towers. The vertical axis is DC-Current in Amperes and the horizontal is time, shown as days of the month. The site is powered by a hybrid of renewable energy sources — solar and wind. Current #1 is the load, which appears

[6] We learned at first hand just how big of a problem this is in sites where batteries were delayed or temporarily switched to AC - somewhat unexpected of North-American power grid.

[7] Our batteries of choice are deep-cycle, sealed, AGM (Absorbed Glass Mat) — very costly (A 200AH 12V unit, like we commonly use at FurtherReach goes for $360 each and a typical 48V core site would have 4 of them).

constant at this coarse time granularity of a week, current #2 is the product of a wind-turbine, which has been minimal during this unfortunate time frame and mostly during daytime where it had not contributed much since solar production was much higher at that time. Current number #3 is at the battery - negative values indicate charge and positive discharge, finally current #4 is the product of the solar panels array. The missing data on the 28th (of July 2015) is due to upgrades on our central monitoring server, which while not harming the service to subscribers did lead to loss of valuable research data (and it's not the only time).

2.5 The FurtherReach network

Service area and subscribers maps

As shown in figure 2.6-a the general area where FurtherReach operates, is at the south part of Mendocino county's coast — the area between Mendocino in the north and Gualala in the south. Most of the population is along the coast with the hills and mountains to the east being mostly underpopulated forests. By fall 2015, FurtherReach has grown to serve subscribers from Albion in the North to Anchor Bay in the south. Our average growth rate, at the time of this writing, is about one new subscriber each day, and we plan to increase this rate. Figure 2.6-b shows the area around the city of Point Arena, and the adjacent native American community — Manchester Band of Pomo Indians of the Manchester Rancheria. The center of Point Arena is probably the most densely populated area under FurtherReach's coverage. At the zoom level of map 2.6-b many of the pins overlap and are not visible.

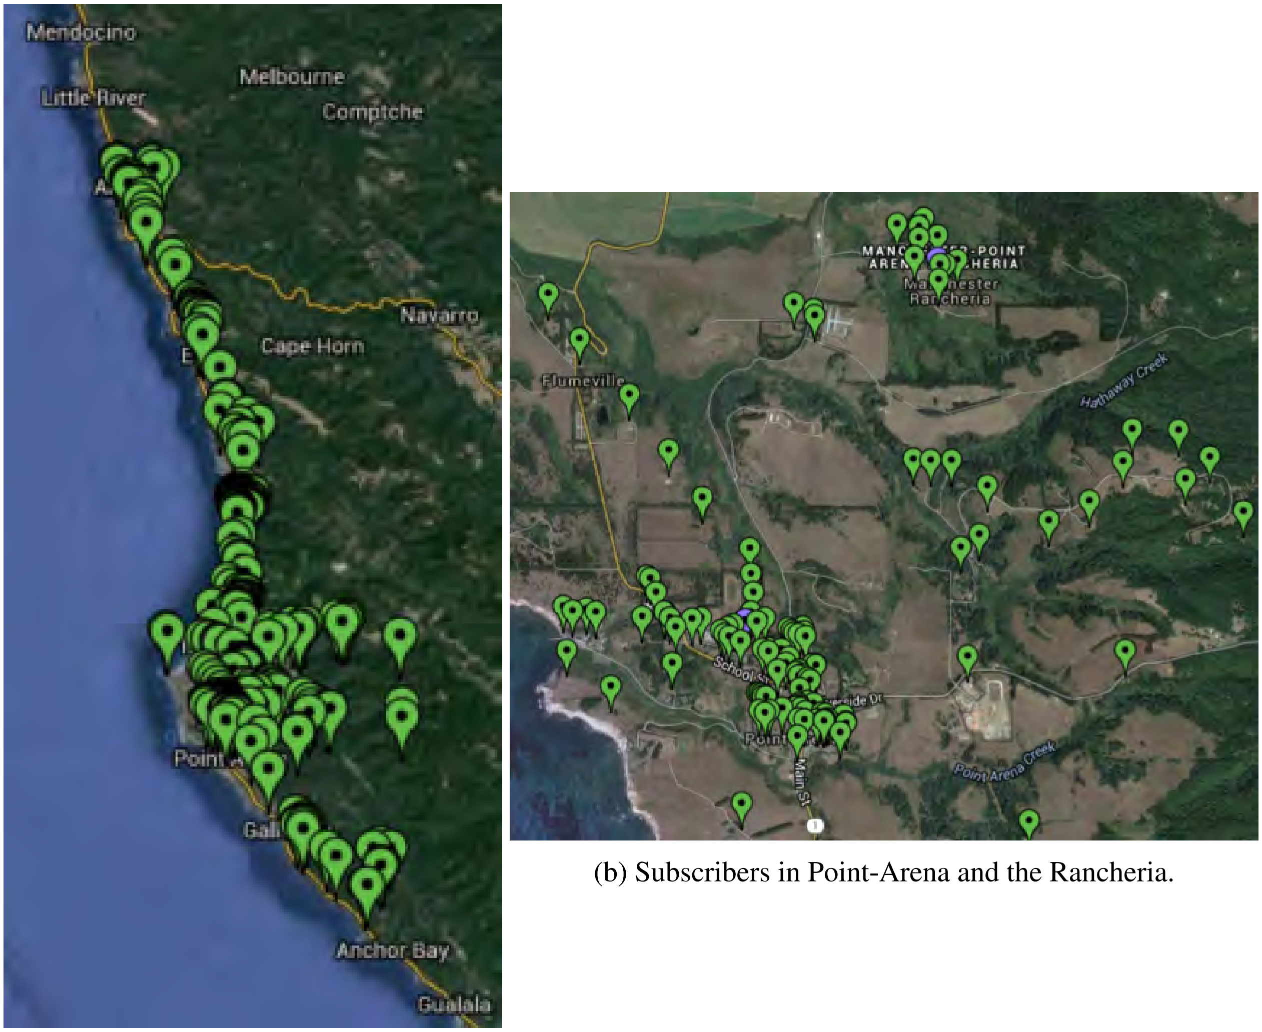

(a) General coverage overview.

(b) Subscribers in Point-Arena and the Rancheria.

Figure 2.6: FurtherReach subscribers map.

The 2nd most densely populated area under FurtherReach's coverage is the Irish Beach subdivision of Manchester, CA. A low-altitude areal photo of Irish-Beach is shown in figure 2.7 with green pins for connected households.[8] The majority of this community is served using three relay sites around it. One on a hill north of this community, using two base-station antennas[9] one facing South-West and the other South-East. They each serve 33 and 18 subscribers respectively. The 2nd site is at the South of this community, across the river, and it is using a similar base-station antenna facing north, which serves 34 subscribers. The 3rd site is on a hill North-East of the community and is using a small panel antenna (*Ubiquiti Nanostation M5*) serving 10 subscribers.

[8] The single purple pin represents a subscriber ”on hold” - probably a part-time resident who asked to reduce his monthly fees by temporarily disconnecting the service until they return in the summer.

[9] These are sector antennas housing *Ubiquiti Rocket Titanium* radios.

These three sites, however, were not enough to provide full coverage of this community albeit the seemingly favorable topography. Subscribers at the east part of this colony, right under our North-Eastern BTS would have their signal obstructed, we therefore had to install another small relay nearby to serve another three households that could not otherwise get service. This is a unique case which we try to avoid as the ratio of three subscribers per BTS is hard to justify financially. Nevertheless, exclusion of some subscribers, especially who are a part of a community that is fully served, creates an undesirable social inequalities that we wanted to avoid. A middle path to achieve these conflicting goals was to forgo the installation of battery backup at this relay site. The three subscribers have been made aware of that and do not expect service to continue during power-grid outages. Overall, for small community relays the principle of fate-sharing typically applies — if the local relay site loses power, it is very likely the subscribers around it also loses power and may not care anyway. However, many of our subscribers invest in UPS systems (Uninterpretable Power Supplies) to ensure their service remain operational during power failures. Most households make use of many battery-powered computing devices and would benefit from service especially during power outages. Furthermore, many subscribers have surrendered their AT&T land-lines and use VoIP telephony service over our network instead. These subscribers are typically the tech-savvy ones who would also ensure they have power backup. Most of our coverage area does not have cellular service, which further emphasizes the importance of a reliable network. Many subscribers also drop their satellite-TV service and instead stream video content over our network. The "cord cutters" enjoy big financial savings through these practices. Although this small community is served using four relays sites and five base-station sector antennas, it was still not enough, a handful of households at the North-West corner of the colony could not get a good signal from any of these five base-stations. Fortunately, a remote site to the north, which has a BTS facing south, was available to fill this gap, and although 8km away these subscribers enjoy the full capacity their links capabilities.

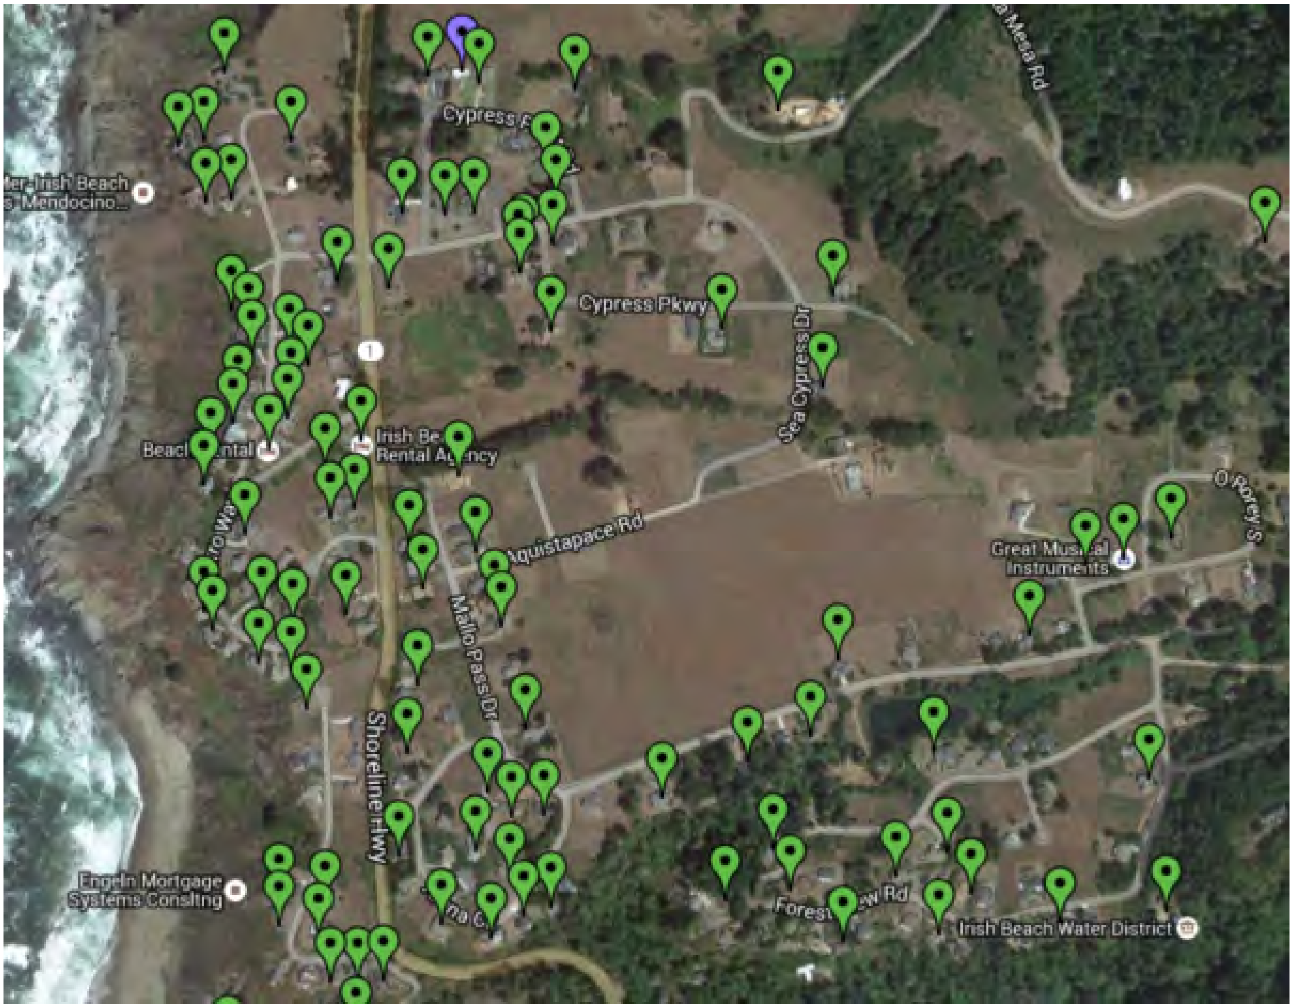

Figure 2.7: FurtherReach subscribers in Irish-Beach (a subdivision of Manchester, CA).

Network map - relay sites

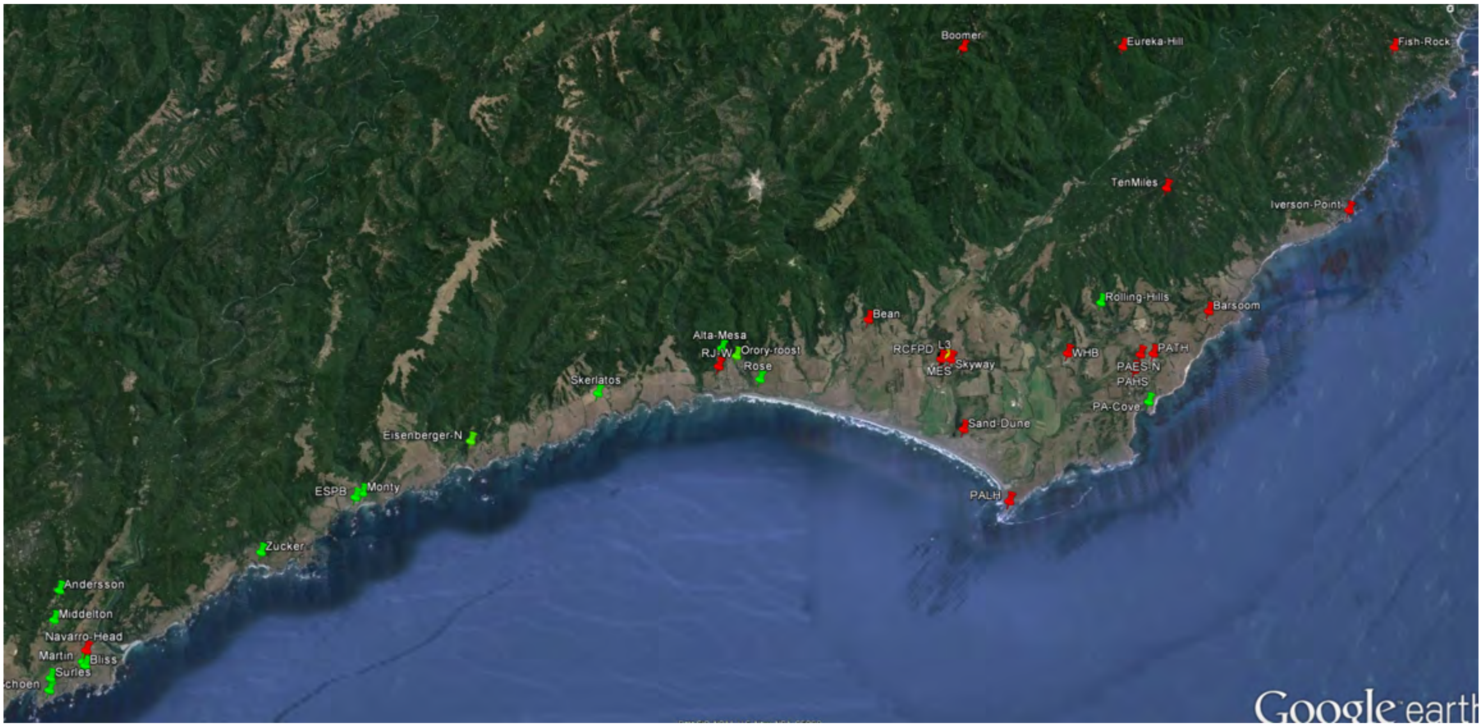

Figure 2.8: FurtherReach map of relay sites - Fall 2015.

Figure 2.8 shows FurtherReach's relay sites, as of September 2015. Red pins are core relays and green are smaller community relays.

Network topology and link capacities.

As seen in figure 2.9, as of September 2015, the FurtherReach network is comprised of 13 core relay sites, and 24 smaller relay sites that we call "neighborhood relays". We define a core site as one that has at least two back-haul links, typically to two or more other relay sites, but at times simply two parallel links to a single other site. Also, a core relay site would have minimal capacity of 250mbps, typically more, and much more when the two or more back-haul links are aggregated together. In figure 2.9 the rectangular shapes are core sites, and the circles represent neighborhood relay sites. A small site does not enjoy redundant back-haul links, and the capacity of its single link would be in the range of 90 to 150Mbps[10]. At each relay site, small or big, there are a number of base-stations (BTSs) to provide the "last mile" connection to subscribers' rooftops. Typical relay site would have three such BTSs, although some extra-small relays only have a single BTS.

[10] The vast majority of point-to-point links we use as back-haul to interconnect small relay sites comprise of devices that have gigabit Ethernet ports and their typical wireless link capacities are 300/300mbps. These are the wireless burst rates of a half-duplex links, thereby the actual TCP throughput is much lower.

The vast majority of the BTSs we use are either a *Ubiquiti Rocket* device with a sector antenna, or a *Ubiquiti NanoStation M5* device in which the router is in-built into a panel antenna. Both of these BTS types use a wired Ethernet port that supports only 100Mbps,[11] although the typical wireless burst speed is 300Mbps. Often, these 100Mbps Ethernet ports are the bottleneck of our deployment. Our most dense BTS currently serves 36 subscribers, and it is a device with a gigabit Ethernet port (*Rocket-Titanium*), see figure 2.10 for a screen-shot of it's administrator's interface . The most dense BTS that has only a 100Mbps port currently serves 22 subscribers and we find that 100mbps is sufficient even when some of these are business-level subscribers.

In figure 2.10 is a screen capture of a highly utilized BTS. As of September 2015, this device has 31 connections, which is below its maximum stated capacity. The device, *Rocket-Titanium*, has a gigabit-Ethernet port and is powered using 48v PoE. As can be seen in figure 2.10 most subscribers enjoy the maximum link quality and are capable of 300/300Mbps. Most have also been up since the BTS was last rebooted 41 days at the time of screen-capture, but strangely, some have rebooted their devices more recently. This area is dominated by vacation rentals who often disconnect the power when they leave - that may explain the shorter connection times.

[11] At dense locations we use a *Ubiquiti Rocket Titanium* BTS which has a gigabit Ethernet port, and faster CPU.

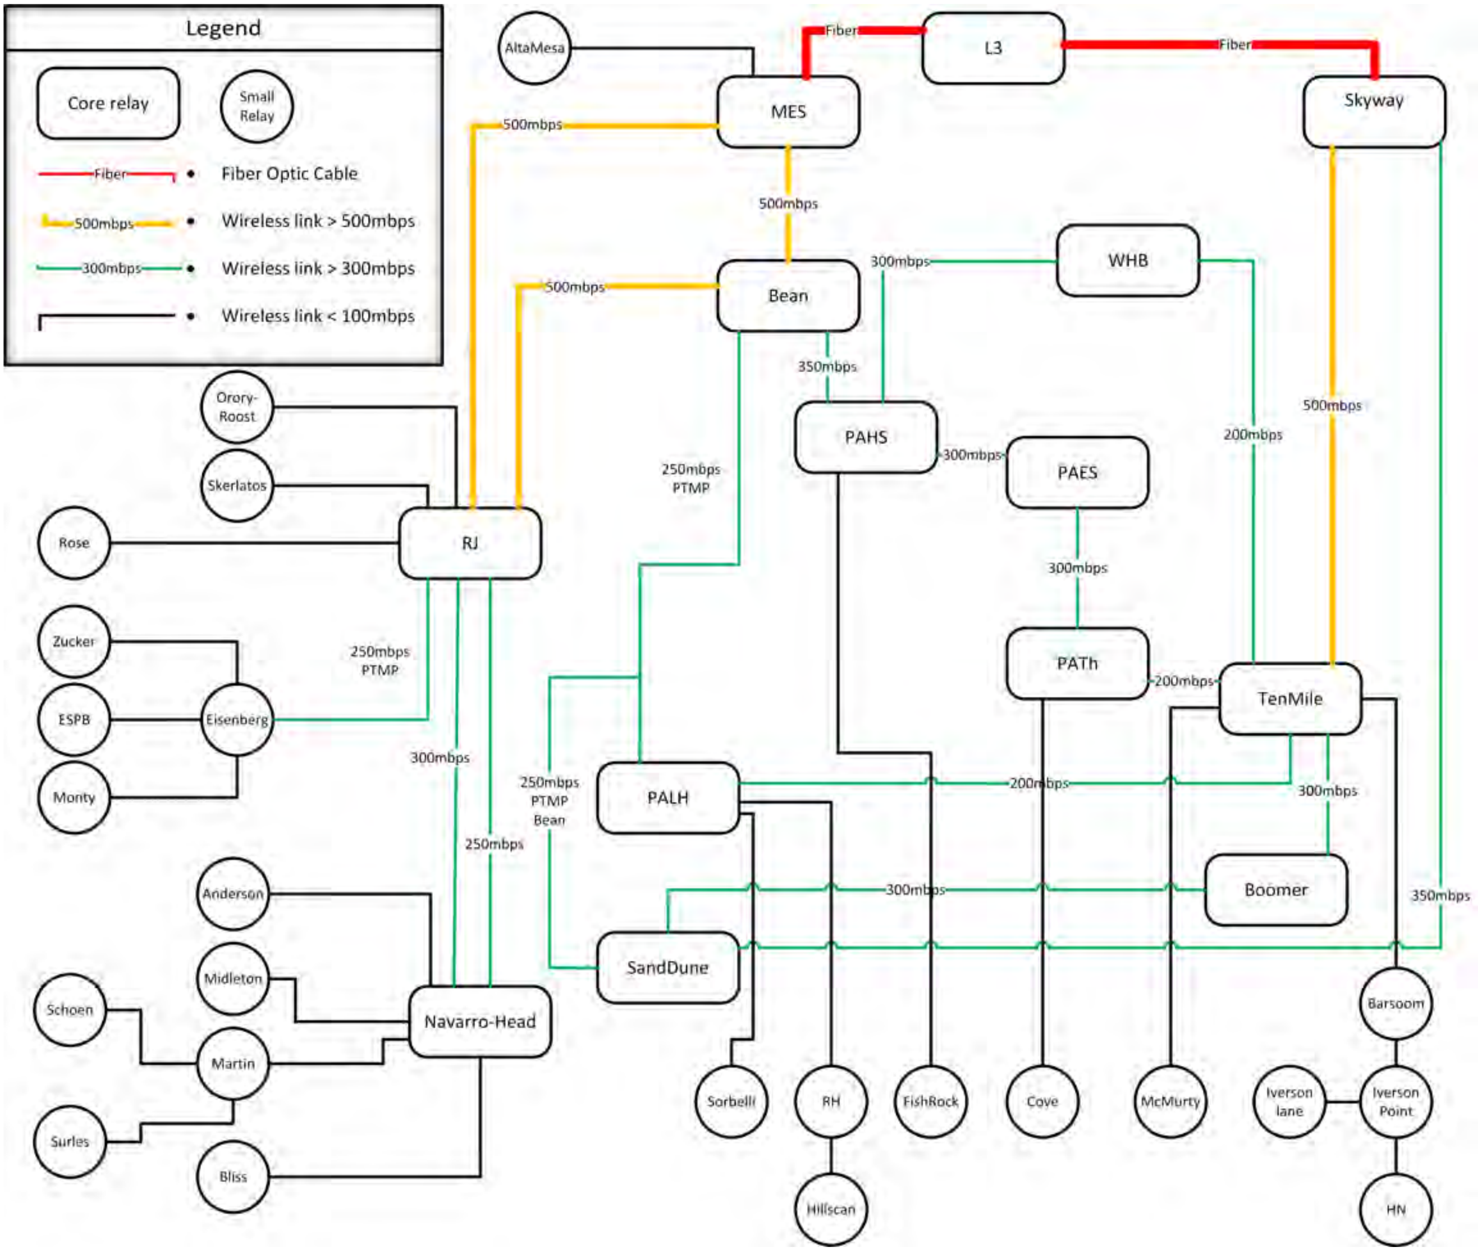

Figure 2.9: FurtherReach network diagram and link capacities — September 2015.

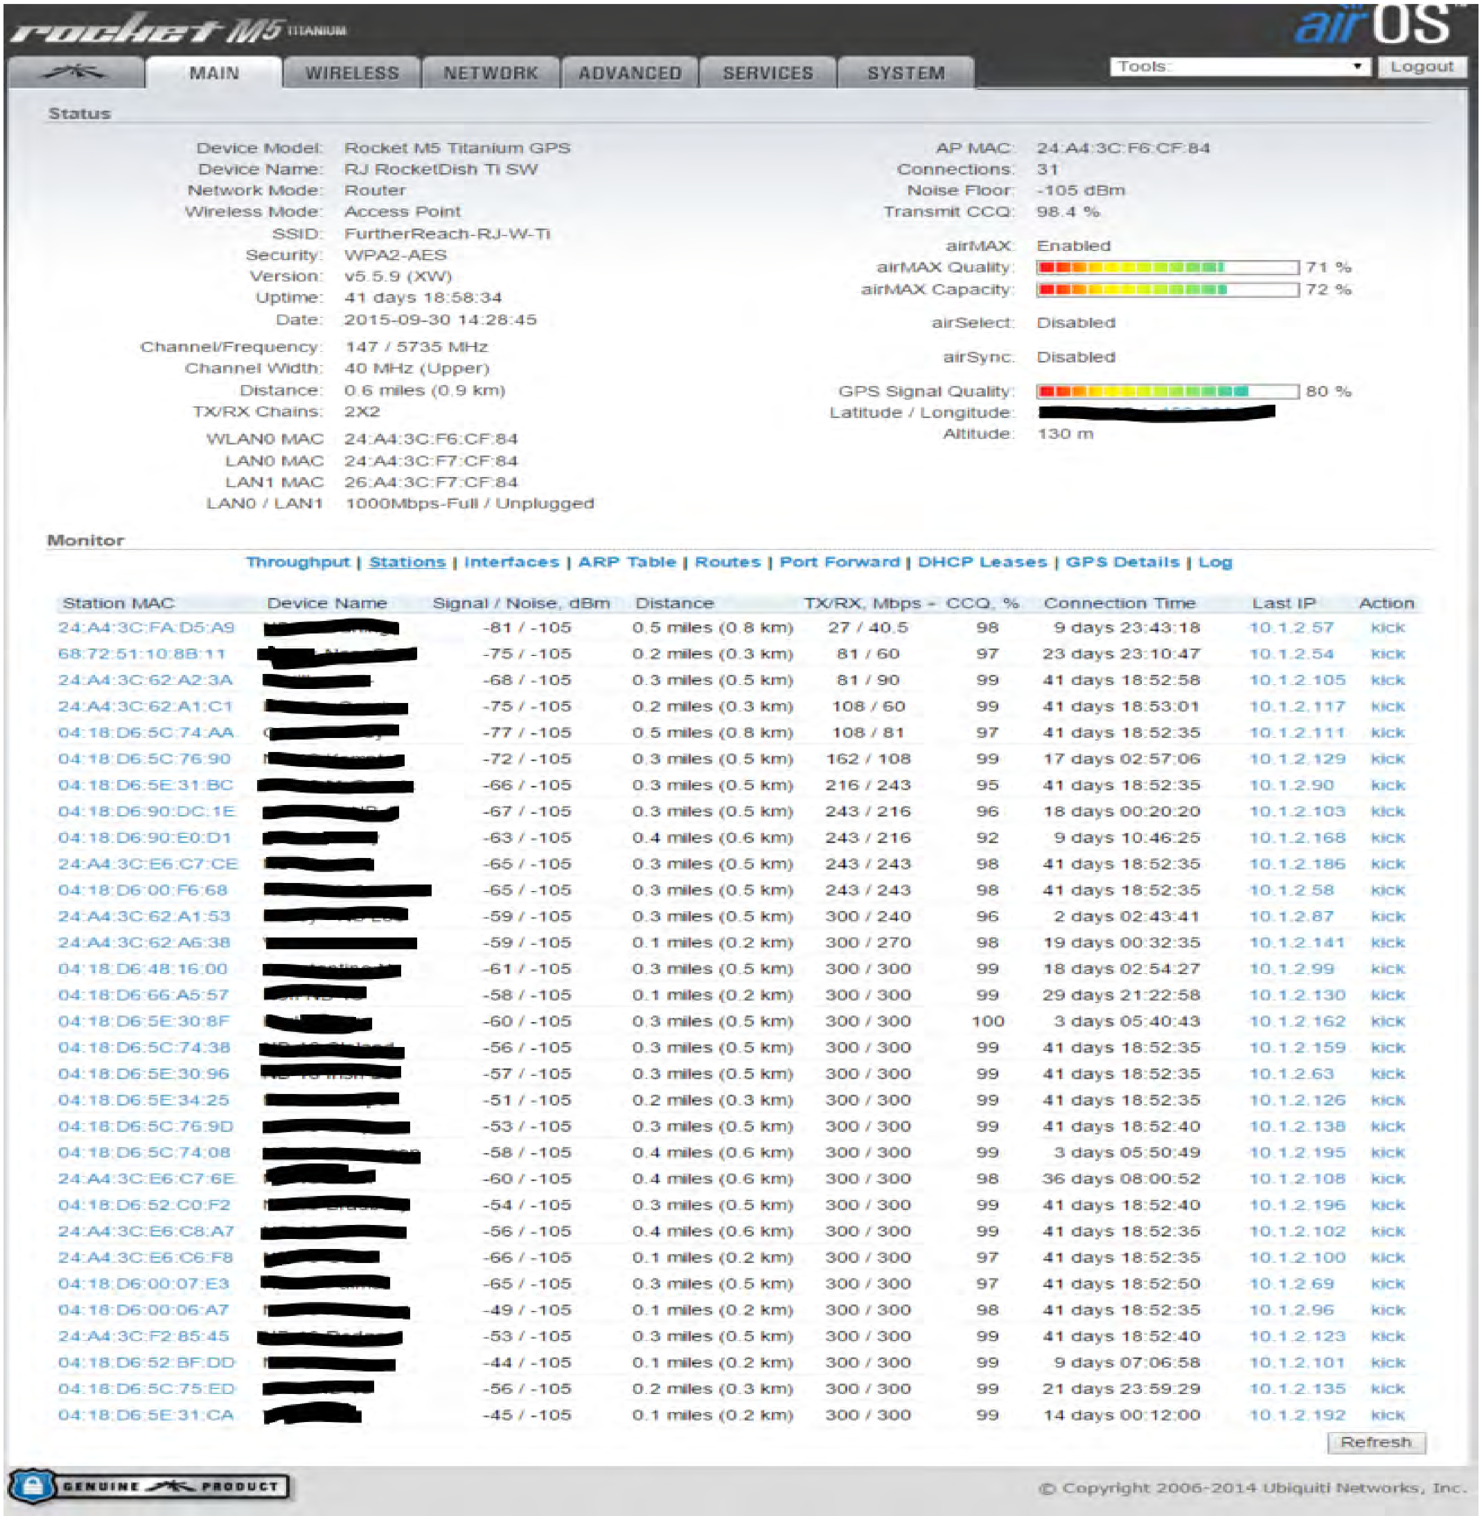

Status

Device Model: Rocket M5 Titanium GPS
Device Name: RJ RocketDish Ti SW
Network Mode: Router
Wireless Mode: Access Point
SSID: FurtherReach-RJ-W-Ti
Security: WPA2-AES
Version: v5.5.9 (XW)
Uptime: 41 days 18:58:34
Date: 2015-09-30 14:28:45
Channel/Frequency: 147 / 5735 MHz
Channel Width: 40 MHz (Upper)
Distance: 0.6 miles (0.9 km)
TX/RX Chains: 2X2
WLAN0 MAC 24:A4:3C:F6:CF:84
LAN0 MAC 24:A4:3C:F7:CF:84
LAN1 MAC 26:A4:3C:F7:CF:84
LAN0 / LAN1 1000Mbps-Full / Unplugged

AP MAC: 24:A4:3C:F6:CF:84
Connections: 31
Noise Floor: -105 dBm
Transmit CCQ: 98.4 %
airMAX: Enabled
airMAX Quality: 71 %
airMAX Capacity: 72 %
airSelect: Disabled
airSync: Disabled
GPS Signal Quality: 80 %
Latitude / Longitude: [redacted]
Altitude: 130 m

Monitor

Throughput | Stations | Interfaces | ARP Table | Routes | Port Forward | DHCP Leases | GPS Details | Log

Station MAC	Device Name	Signal / Noise, dBm	Distance	TX/RX, Mbps -	CCQ, %	Connection Time	Last IP	Action
24:A4:3C:FA:D5:A9	[redacted]	-81 / -105	0.5 miles (0.8 km)	27 / 40.5	98	9 days 23:43:18	10.1.2.57	kick
68:72:51:10:8B:11	[redacted]	-75 / -105	0.2 miles (0.3 km)	81 / 60	97	23 days 23:10:47	10.1.2.54	kick
24:A4:3C:62:A2:3A	[redacted]	-68 / -105	0.3 miles (0.5 km)	81 / 90	99	41 days 18:52:58	10.1.2.105	kick
24:A4:3C:62:A1:C1	[redacted]	-75 / -105	0.2 miles (0.3 km)	108 / 60	99	41 days 18:53:01	10.1.2.117	kick
04:18:D6:5C:74:AA	[redacted]	-77 / -105	0.5 miles (0.8 km)	108 / 81	97	41 days 18:52:35	10.1.2.111	kick
04:18:D6:5C:76:90	[redacted]	-72 / -105	0.3 miles (0.5 km)	162 / 108	99	17 days 02:57:06	10.1.2.129	kick
04:18:D6:5E:31:BC	[redacted]	-66 / -105	0.3 miles (0.5 km)	216 / 243	95	41 days 18:52:35	10.1.2.90	kick
04:18:D6:90:DC:1E	[redacted]	-67 / -105	0.3 miles (0.5 km)	243 / 216	96	18 days 00:20:20	10.1.2.103	kick
04:18:D6:90:E0:D1	[redacted]	-63 / -105	0.4 miles (0.6 km)	243 / 216	92	9 days 10:46:25	10.1.2.168	kick
24:A4:3C:E6:C7:CE	[redacted]	-65 / -105	0.3 miles (0.5 km)	243 / 243	98	41 days 18:52:35	10.1.2.186	kick
04:18:D6:00:F6:68	[redacted]	-65 / -105	0.3 miles (0.5 km)	243 / 243	98	41 days 18:52:35	10.1.2.58	kick
24:A4:3C:62:A1:53	[redacted]	-59 / -105	0.3 miles (0.5 km)	300 / 240	96	2 days 02:43:41	10.1.2.87	kick
24:A4:3C:62:A6:38	[redacted]	-59 / -105	0.1 miles (0.2 km)	300 / 270	98	19 days 00:32:35	10.1.2.141	kick
04:18:D6:48:16:00	[redacted]	-61 / -105	0.3 miles (0.5 km)	300 / 300	99	18 days 02:54:27	10.1.2.99	kick
04:18:D6:66:A5:57	[redacted]	-58 / -105	0.1 miles (0.2 km)	300 / 300	99	29 days 21:22:58	10.1.2.130	kick
04:18:D6:5E:30:8F	[redacted]	-60 / -105	0.3 miles (0.5 km)	300 / 300	100	3 days 05:40:43	10.1.2.162	kick
04:18:D6:5C:74:38	[redacted]	-56 / -105	0.3 miles (0.5 km)	300 / 300	99	41 days 18:52:35	10.1.2.159	kick
04:18:D6:5E:30:96	[redacted]	-57 / -105	0.3 miles (0.5 km)	300 / 300	99	41 days 18:52:35	10.1.2.63	kick
04:18:D6:5E:34:25	[redacted]	-51 / -105	0.2 miles (0.3 km)	300 / 300	99	41 days 18:52:35	10.1.2.126	kick
04:18:D6:5C:76:9D	[redacted]	-53 / -105	0.3 miles (0.5 km)	300 / 300	99	41 days 18:52:40	10.1.2.138	kick
04:18:D6:5C:74:08	[redacted]	-58 / -105	0.4 miles (0.6 km)	300 / 300	99	3 days 05:50:49	10.1.2.195	kick
24:A4:3C:E6:C7:6E	[redacted]	-60 / -105	0.4 miles (0.6 km)	300 / 300	98	36 days 08:00:52	10.1.2.108	kick
04:18:D6:52:C0:F2	[redacted]	-54 / -105	0.3 miles (0.5 km)	300 / 300	99	41 days 18:52:40	10.1.2.196	kick
24:A4:3C:E6:C8:A7	[redacted]	-56 / -105	0.4 miles (0.6 km)	300 / 300	99	41 days 18:52:35	10.1.2.102	kick
24:A4:3C:E6:C6:F8	[redacted]	-66 / -105	0.1 miles (0.2 km)	300 / 300	97	41 days 18:52:35	10.1.2.100	kick
04:18:D6:00:07:E3	[redacted]	-65 / -105	0.3 miles (0.5 km)	300 / 300	97	41 days 18:52:50	10.1.2.69	kick
04:18:D6:00:06:A7	[redacted]	-49 / -105	0.1 miles (0.2 km)	300 / 300	98	41 days 18:52:35	10.1.2.96	kick
24:A4:3C:F2:85:45	[redacted]	-53 / -105	0.3 miles (0.5 km)	300 / 300	99	41 days 18:52:40	10.1.2.123	kick
04:18:D6:52:BF:DD	[redacted]	-44 / -105	0.1 miles (0.2 km)	300 / 300	99	9 days 07:06:58	10.1.2.101	kick
04:18:D6:5C:75:ED	[redacted]	-56 / -105	0.2 miles (0.3 km)	300 / 300	99	21 days 23:59:29	10.1.2.135	kick
04:18:D6:5E:31:CA	[redacted]	-45 / -105	0.1 miles (0.2 km)	300 / 300	99	14 days 00:12:00	10.1.2.192	kick

Figure 2.10: Screen capture of a base station supporting 31 subscriber connections.

FurtherReach in photos — what does a WISP looks like?

As discussed in length, WISPs deployments are quite different from "traditional" communication infrastructure and the physical nature of these deployments is quite difficult to envision. The diversity of installations and the substantial differences between each relay site, not to mention subscribers' installations, makes photos an exceptionally valuable aid in depicting the roll-out. Following is a select list of relay sites, their photos and description of the unique features of each site. Through these photos I wish to convey the high degree of diversity required for WISPs deployment, and therefore the challenges in planning and forecasting until a site has been acquired.

MES relay site

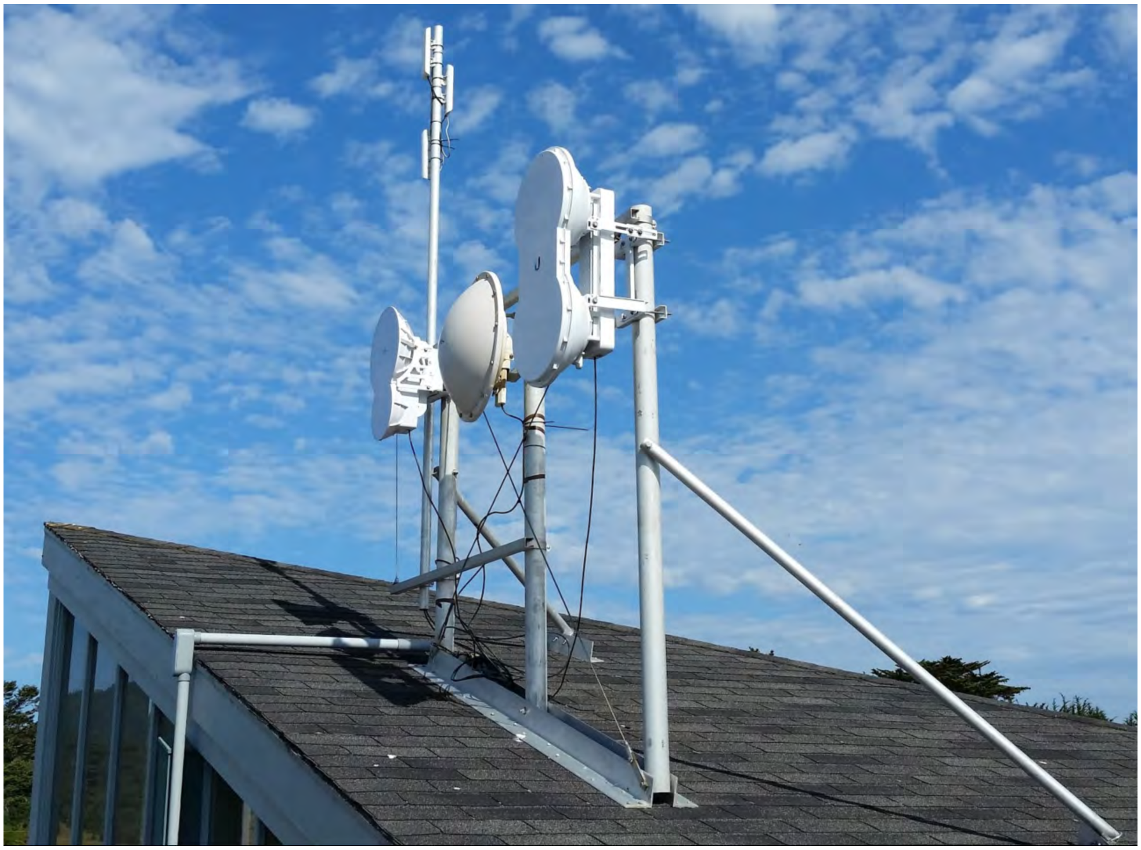

Figure 2.11: *MES* site — the roof of our 2nd fiber-fed site.

The board of directors of the Manchester Elementary School had graciously allowed us to use their roof as a key relay site. We trenched about 1500 feet long ditch from the data center of *Level-3 communications* to the school, choosing a path that traverses only two other properties. We buried a two inch conduit, at an average depth of five feet, and pulled an armored fiber optic cable through it (12 fiber pairs). For a few months this has been our single gateway between wired (fiber) and our wireless network. The challenges with this site were numerous; First, the topography and tree coverage is such that this site has no value towards any direction other then North, as can be seen in figure 2.11 in which all major antennas are facing in a single direction. We therefore had to shoot northward to a hill over Irish Beach, and from there shoot back South to Point-Arena, or South-West to the lighthouse. This dependency on a single site, as well as the challenge of having the main gateway site facing in a single direction (and not the direction where most demand for bandwidth is from) caused elevated anxiety to the whole team, until we managed to trench another fiber to a south-facing gateway site. The slanted roof at this site presented interesting mechanical installation challenges for our rather large and heavy antennas. Strong winds, proximity to the ocean, and kids playing right below in their schoolyard all added to the challenge. Krispin, our local area lead, designed the complex mount shown, based on Aluminum profiles and thick-walled Aluminum pipes tighten together with stainless steel wire ropes and fasteners such as used on boats, coupled together using *Super-Stuart* channels made of hot-galvanized steel. Seen in figure 2.11 are two *AirFiber* full-duplex radios — the far one is a 24Ghz unit and the closer one works on the 5Ghz band. In the middle is an *AirFiber-X* radio inside a *Rocket dish* antenna. The AirFibers each offer capacity of about 780Mbps at these link distance (about 5 miles), and the AF5X gives about 500Mbps (half-duplex). At the top of the pole above these dishes, are 3 small panel antennas to provide the “last mile” connection to nearby subscribers.

Figure 2.12 is a screen capture from the network switch at the MES site. This POE (Power Over Ethernet) switch has advance power management capabilities which we mostly put to use at sites powered from renewable energy sources. In this figure, we can see the power consumption for each device, the total throughput graph, as well as the bandwidth distribution pie-charts. The charts show that the vast majority of the RX bandwidth comes via port number 14 which is the fiber-optic cable to Level-3. The TX pie shows that most of that bandwidth is divided between port 4 and port 1, the AirFiber links to Bean and to RJ’s respectively. From these two major core sites it gets distributed further. Interesting to note that ports number 5 and 6 are disabled and not used for network traffic. These devices are simply being powered by this switch, while their network traffic travels through an independent fiber pair to Level-3. These are a server and another network switch that are used for non-critical services (network monitoring) which we may want to power down in times of power grid outages to prolong our battery backup at this site.

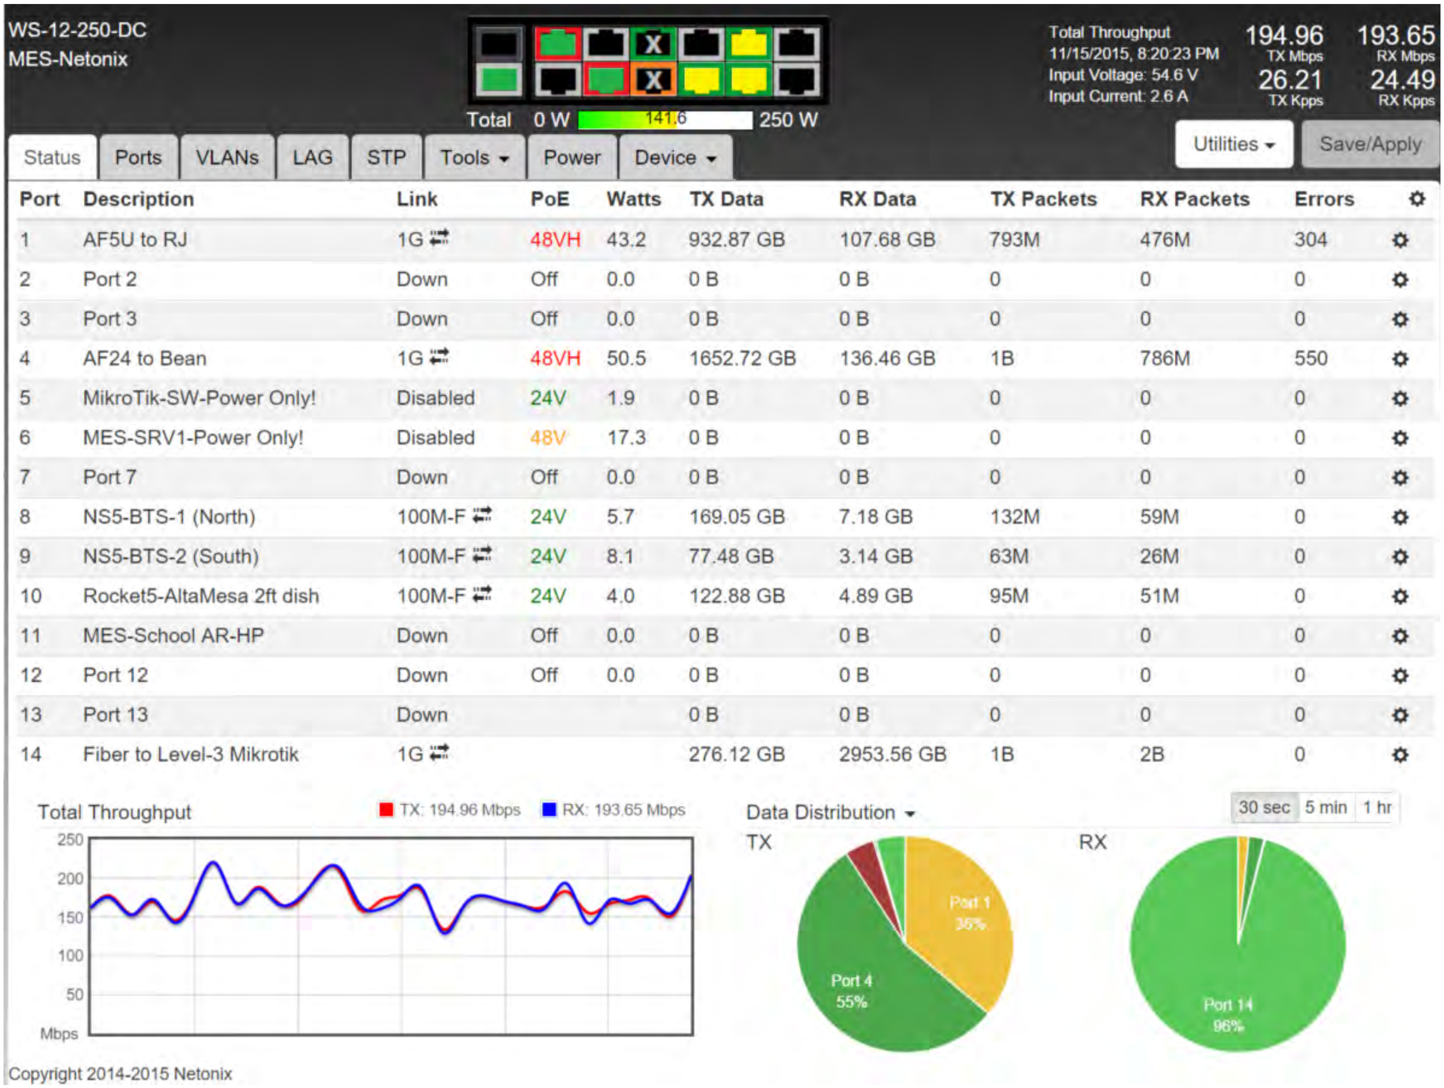

Total Throughput
11/15/2015, 8:20:23 PM
Input Voltage: 54.6 V
Input Current: 2.6 A

194.96 TX Mbps
193.65 RX Mbps
26.21 TX Kpps
24.49 RX Kpps

Port	Description	Link	PoE	Watts	TX Data	RX Data	TX Packets	RX Packets	Errors	⚙
1	AF5U to RJ	1G ⇄	48VH	43.2	932.87 GB	107.68 GB	793M	476M	304	⚙
2	Port 2	Down	Off	0.0	0 B	0 B	0	0	0	⚙
3	Port 3	Down	Off	0.0	0 B	0 B	0	0	0	⚙
4	AF24 to Bean	1G ⇄	48VH	50.5	1652.72 GB	136.46 GB	1B	786M	550	⚙
5	MikroTik-SW-Power Only!	Disabled	24V	1.9	0 B	0 B	0	0	0	⚙
6	MES-SRV1-Power Only!	Disabled	48V	17.3	0 B	0 B	0	0	0	⚙
7	Port 7	Down	Off	0.0	0 B	0 B	0	0	0	⚙
8	NS5-BTS-1 (North)	100M-F ⇄	24V	5.7	169.05 GB	7.18 GB	132M	59M	0	⚙
9	NS5-BTS-2 (South)	100M-F ⇄	24V	8.1	77.48 GB	3.14 GB	63M	26M	0	⚙
10	Rocket5-AltaMesa 2ft dish	100M-F ⇄	24V	4.0	122.88 GB	4.89 GB	95M	51M	0	⚙
11	MES-School AR-HP	Down	Off	0.0	0 B	0 B	0	0	0	⚙
12	Port 12	Down	Off	0.0	0 B	0 B	0	0	0	⚙
13	Port 13	Down			0 B	0 B	0	0	0	⚙
14	Fiber to Level-3 Mikrotik	1G ⇄			276.12 GB	2953.56 GB	1B	2B	0	⚙

Copyright 2014-2015 Netonix

Figure 2.12: FurtherReach’s network switch’s GUI at the Manchester Elementary school.

PAHS relay site

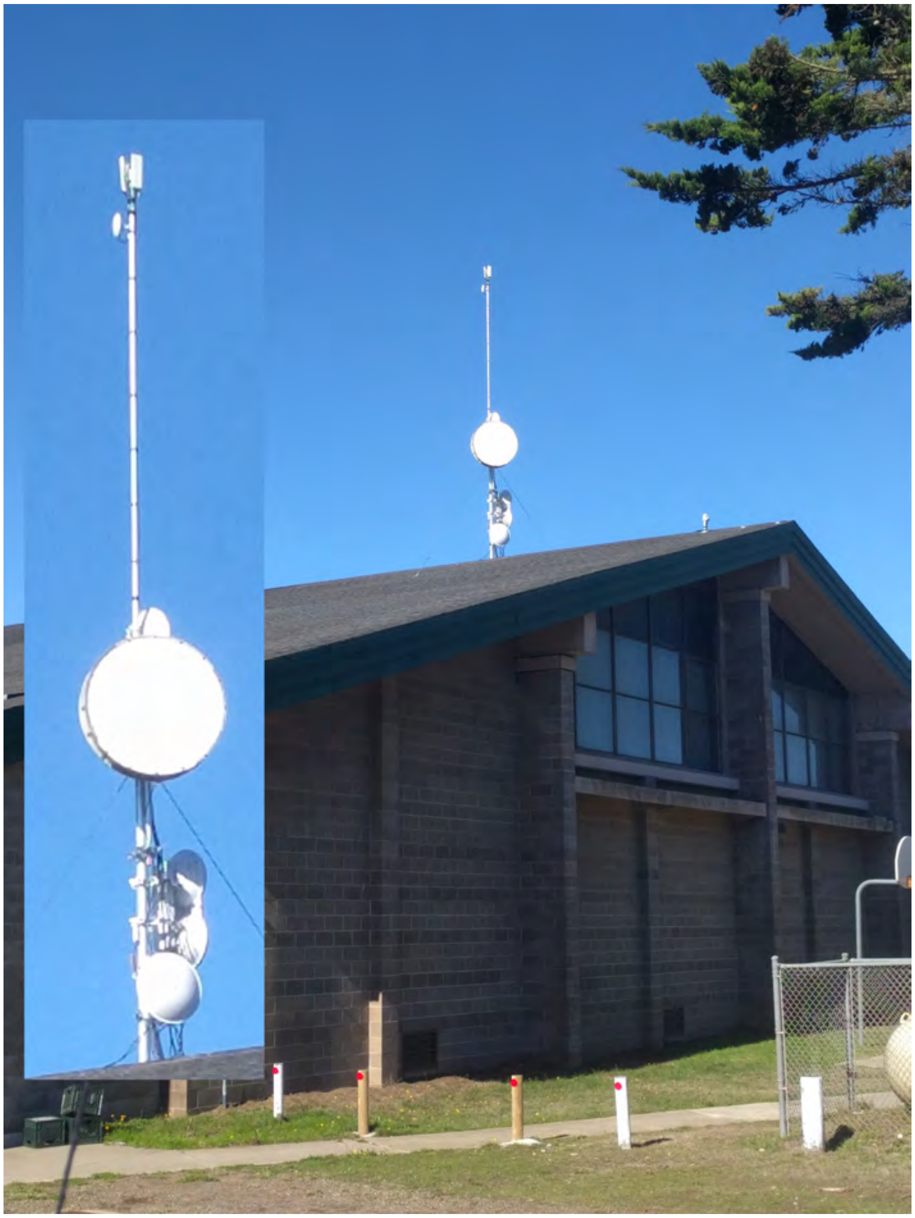

Figure 2.13: FurtherReach's relay on the roof of the Point Arena high-school's gym.

The gym of the Point Arena high-school is uniquely located to interconnect a number of core relay sites, as well as to serve a relatively dense population of subscribers nearby. At the top of the mast, as seen in figure 2.13 are three small sector antennas that forms the last-mile connectivity to subscribers around the high-school. The proximity to subscribers provide for a strong signal-to-noise ratio, which makes interference immunity better, enables higher link capacities, and allow the use of small CPE antennas at subscribers' rooftops. The small sector antennas at the top of this mast are of high beam-width and therefore works well even if strong winds shakes the mast. These antennas have low wind-load making them ideal to be placed at the top of the mast from where more subscribers are visible over the trees. Lower on that mast are the back-haul links interconnecting multiple other core sites.

Figure 2.14 shows the GUI of the network switch at the PA-High-school.
Port 1 connects to the large dish which hosts an AirFiber-5X radio to feed our Fish-Rock site (in Gualala).
Port 2 feeds the AirFiber unit which is the main back-haul link to this site from Manchester (Bean).
Ports 3-4 are not in use.
Ports 5-7 are the sector antennas at the top, that are aimed in these general directions: North, South-East, and South-West.
Port 6, in addition to feeding the South-East sector also powers a local Pico-station through daisy-chained POE pass-through from the Nanostation unit. The Pico is used for local 2.4Ghz access, mostly by our own staff.
Port 8 has an 801.11AC unit (PowerBeam AC-500) that connects to our Windy-Hollow site.
Port 9 connects to another AC-500 unit which links to the Point-Arena elementary school.
Port 10 is an NB19 unit pointing to our Cove site, which is not in use (it's a manual "backdoor" to get to the cove in case of a main link failure).
Pot 11 is where our networked power-monitor is connected.

It's interesting to point out the power consumption of these units - the AC units (ports 8 and 9) use much lower power than the power hungry AirFibers (see port 2 where the AF5 units draws 45W of power vs. ports 8 and 9 that each draws about 3.5W). The AirFiber unit is a full duplex device which result in links with about 1ms lower latency (RTT- Round Trip Time). This may sound insignificant but in the case of FurtherReach's excellent Internet gateway that offers less than 6ms to Google and most SF Bay-Area data-centers, an additional millisecond is quite harmful. That said, for solar-power sites, this privilege translates into an often unacceptable power budget that offsets the finances in a harmful way, and thereby an AC-type devices are preferred. The AC units, are also much cheaper and yet are capable of delivering about 500mpbs of TCP traffic or even more, in many situations, making these highly desirable units (when 80Mhz channel widths are used).
AirFiber-X devices, such as on port 1 - are half-duplex devices with latencies similar to the AC-units, yet with higher power (and therefore higher power consumption) and they use RF-connectors allowing them to be installed in larger antennas as needed.

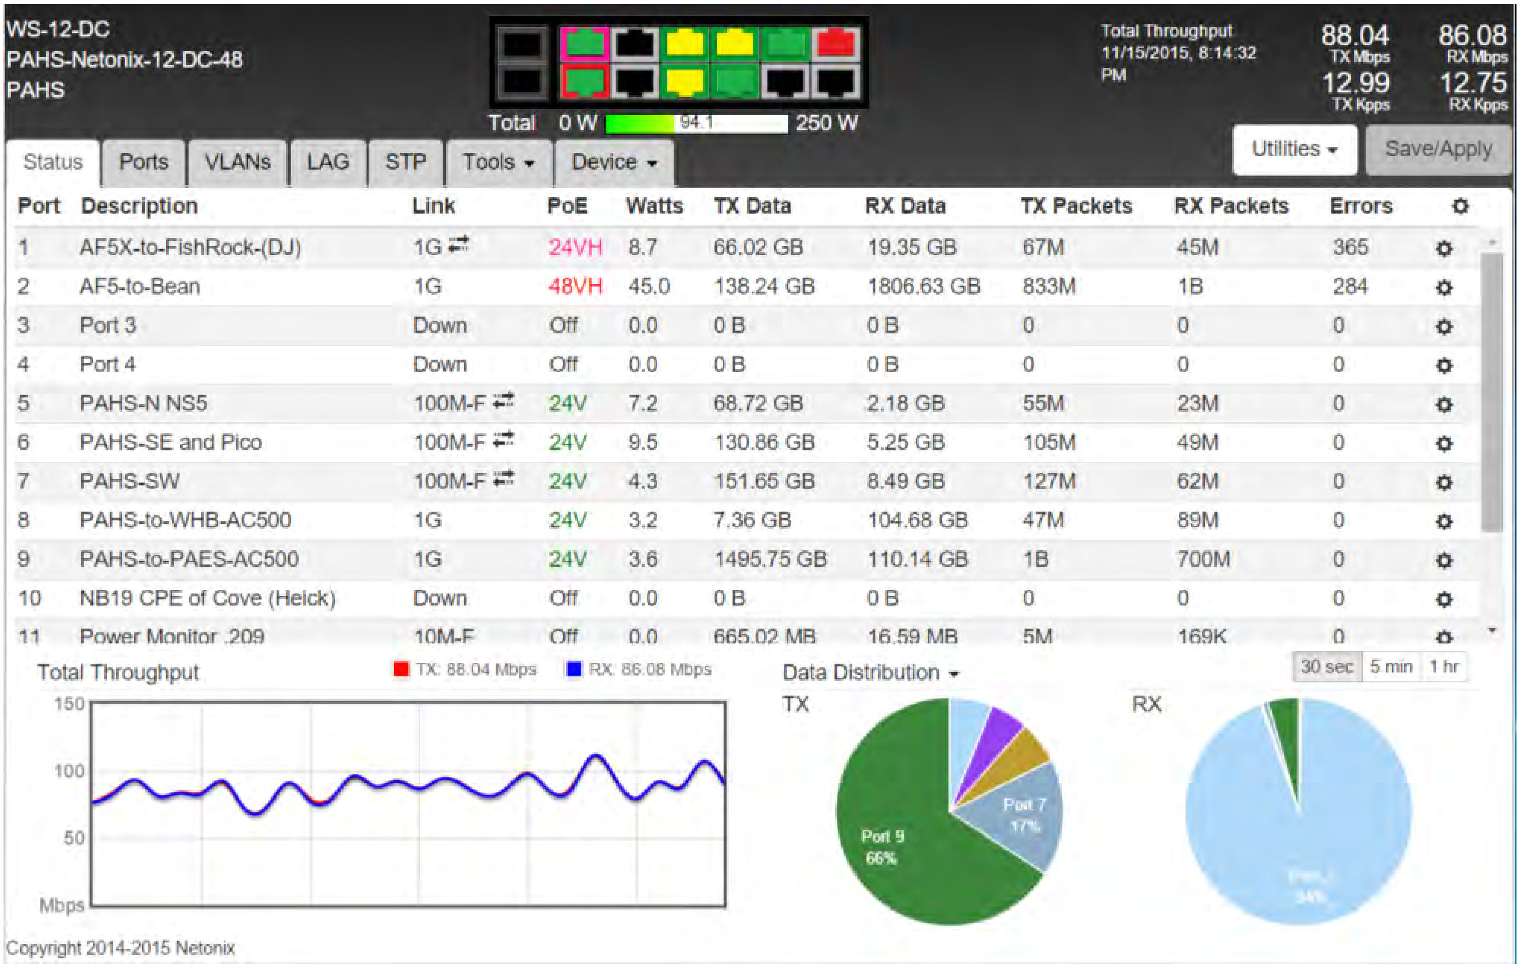

Port	Description	Link	PoE	Watts	TX Data	RX Data	TX Packets	RX Packets	Errors	⚙
1	AF5X-to-FishRock-(DJ)	1G ⇄	24VH	8.7	66.02 GB	19.35 GB	67M	45M	365	⚙
2	AF5-to-Bean	1G	48VH	45.0	138.24 GB	1806.63 GB	833M	1B	284	⚙
3	Port 3	Down	Off	0.0	0 B	0 B	0	0	0	⚙
4	Port 4	Down	Off	0.0	0 B	0 B	0	0	0	⚙
5	PAHS-N NS5	100M-F ⇄	24V	7.2	68.72 GB	2.18 GB	55M	23M	0	⚙
6	PAHS-SE and Pico	100M-F ⇄	24V	9.5	130.86 GB	5.25 GB	105M	49M	0	⚙
7	PAHS-SW	100M-F ⇄	24V	4.3	151.65 GB	8.49 GB	127M	62M	0	⚙
8	PAHS-to-WHB-AC500	1G	24V	3.2	7.36 GB	104.68 GB	47M	89M	0	⚙
9	PAHS-to-PAES-AC500	1G	24V	3.6	1495.75 GB	110.14 GB	1B	700M	0	⚙
10	NB19 CPE of Cove (Heick)	Down	Off	0.0	0 B	0 B	0	0	0	⚙
11	Power Monitor .209	10M-F	Off	0.0	665.02 MB	16.59 MB	5M	169K	0	⚙

Figure 2.14: FurtherReach's network switch's GUI at the Point Arena high-school's gym.

TenMiles relay site

It's nice to inherit an unused communication tower. That said, the trees around that 50+ years old structure have outgrown its height and forced us to extend it using a structure based on thick Aluminum poles held together using galvanized steel strut channels. The integrity of this extension, and the strength of the old structure itself are quite questionable and will require changes in the long run. Figure 2.15 our TenMiles site with the large AT&T tower far in the background and a zoom on the top of our tower with the antennas numbered. The two wooden poles on each side of the tower and the antennas mounted on them are defunct reminders of the old TV-translator site and should be removed for safety.

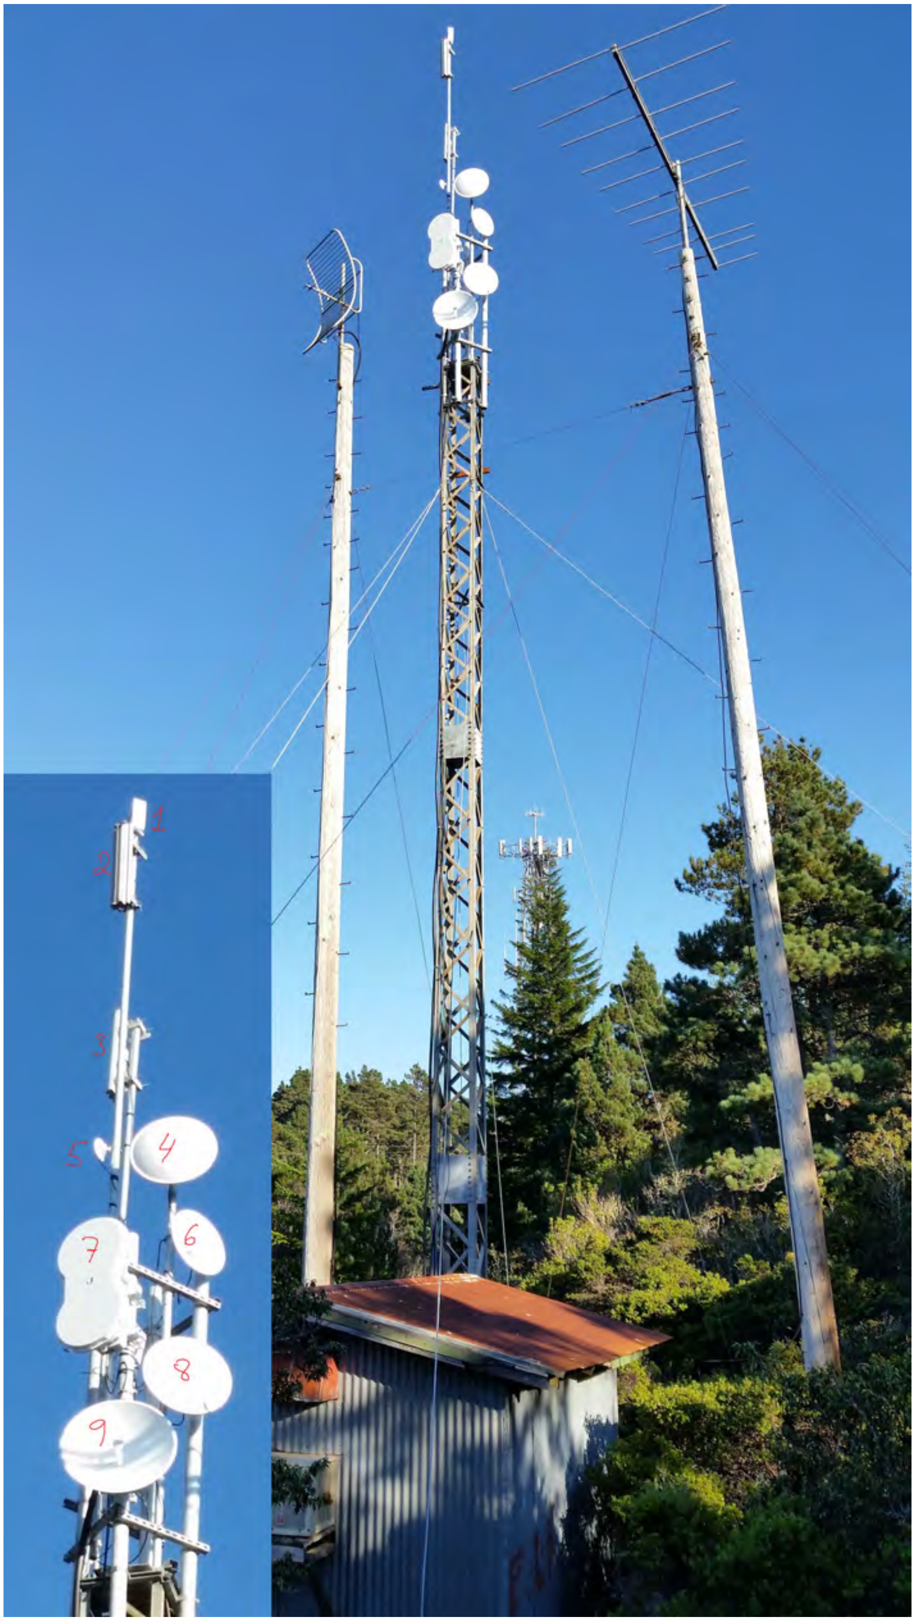

Figure 2.15: FurtherReach‘s TenMiles tower, with a large AT&T tower in the background.

At the time of writing, we had nine antennas mounted on this tower:
Antennas number 1 - An NS-3.65Ghz sector towards Curly Lane.
Antenna 2 - A 5Ghz Rocket sector towards Mountain-View area.
Antenna 3 is a 5Ghz Rocket sector serving North.
Antenna 4 is a 5Ghz ”AC500” dish (500mm diameter) towards Windy-Hollow (WHB).
Antenna 5 - is a no-longer in-use, a 16dBi nanobeam at 5Ghz that is mounted too high to take down without harming other operational units mounted on the same pole.
Antenna 6 is a 5Ghz Nanobeam 400 which is a backup link to *RedRoof House* relay.
Antenna 7 is the main back-haul link for this site, an AirFiber at 5Ghz band.
Antenna 8 is an AC500 unit that feeds the Point-Arena theater relay site.
Antenna 9 is an 620mm dish that links to the PA-Lighthouse.

Bean’s relay site

Shown in figure 2.16 is our largest solar-power site. It’s on a remote hill that has excellent views and allow to interconnect multiple core sites. The combination of multiple power hungry AirFibers and solar power is challenging. As such we keep this site powered down most of the day, and only power up the AirFiber units from 6PM to 1AM every night. These are the most busy hours and the only hours in which we benefit from the added capacity. During other hours of the day, this site is only on stand-by as backup in case of a failure at one of our major core sites. One of the AC500 units is always on, to feed the local sectors and allow our controller to power up the AirFibers if backup is needed. Obviously, the power up process is quite long (AirFibers may take up to 2 minutes to establish a link), so such backup procedure will result in a short downtime. However, this sort of backup is designed to be use only in an extreme cases of major failure at other core sites, and therefore is rare. We never had to use this site for backup due to failure. We did power it up during the day, at a number of occasions, to allow for planned maintenance at other core sites, such as for firmware upgrades or replacement of core switches. In the future, we may consider to beef-up the solar capacity of this site, and keep it running 24/7. The current size of the solar array at this site is 1KW, which is not sufficient for continuous operation of three AirFiber units (each draws about 50W). The low footprint of the antenna mounts at this site, coupled with the camouflage paint, makes it nearly invisible to the naked eye from most places that are accessible by road.

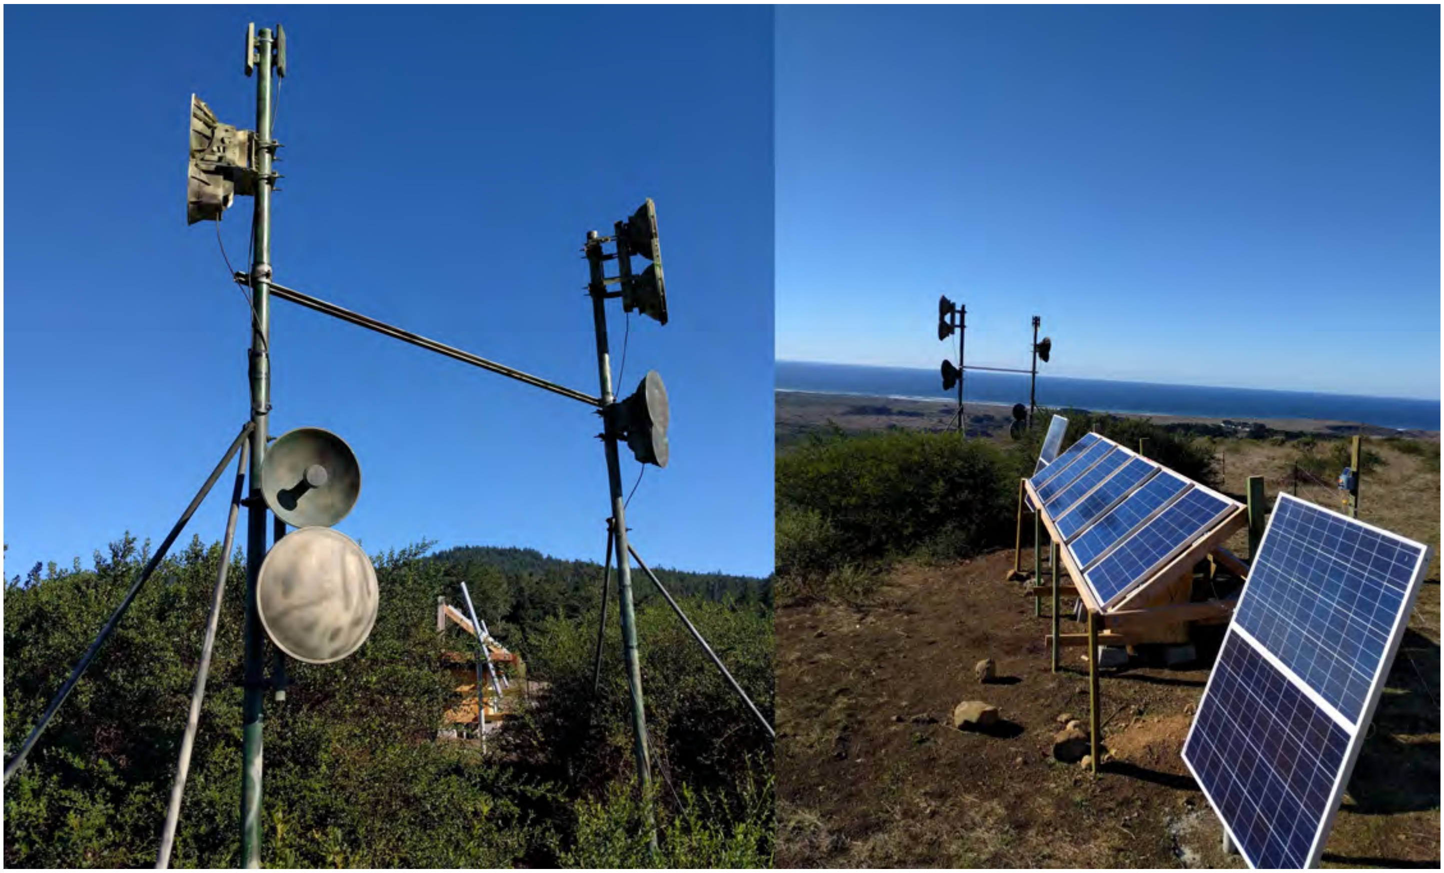

Figure 2.16: FurtherReach‘s “Bean” relay site.

Figure 2.17 is a *SmokePing* graph showing latencies over 24 hours period to the United States Coast Guard (USCG) station that we feed. At the date shown, the Bean site was powered most of the day, and only got switched off from 1AM to 7AM which is visible in this graph as a notable increase in latency. The alternate link to the USCG station is half-duplex setup, while the AirFibers at Bean are full-duplex and thereby offer lower latency. Also notable are short surges in latency at the beginning, and especially at the end of the Bean shut-off period. This is due to network convergence events associated with these topology changes, yet no packet-loss was registered (colors other than green).

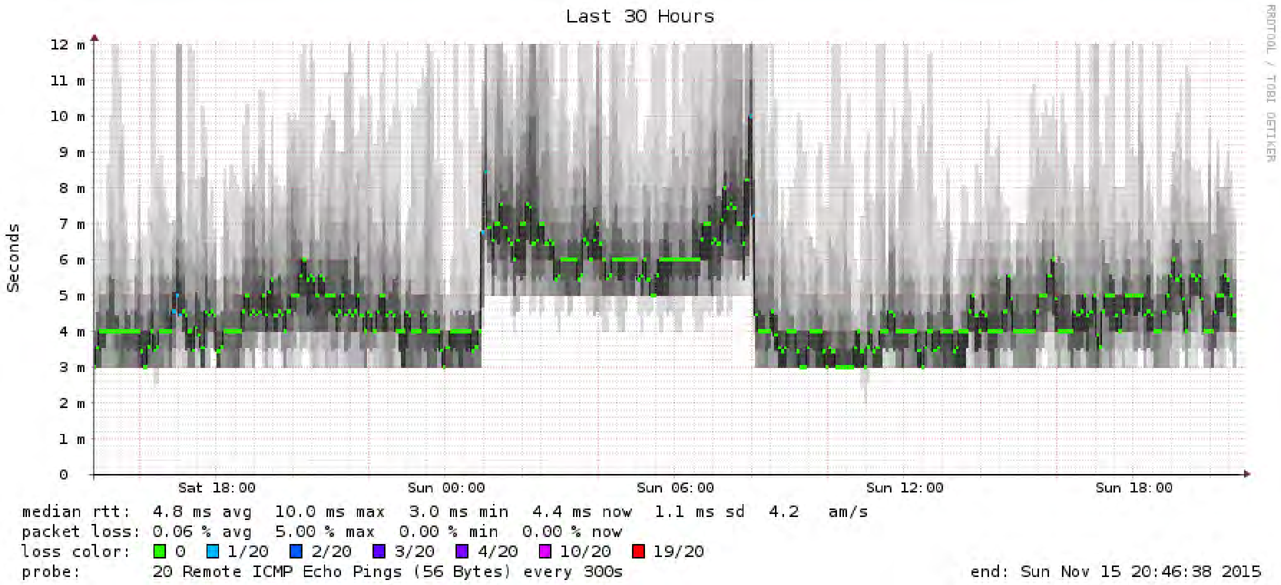

```
median rtt:   4.8 ms avg   10.0 ms max   3.0 ms min   4.4 ms now   1.1 ms sd   4.2   am/s
packet loss:  0.06 % avg   5.00 % max   0.00 % min   0.00 % now
loss color:   🟩 0   🟦 1/20   🟦 2/20   🟪 3/20   🟪 4/20   🟪 10/20   🟥 19/20
probe:        20 Remote ICMP Echo Pings (56 Bytes) every 300s                  end: Sun Nov 15 20:46:38 2015
```

Figure 2.17: “SmokePing” graph to United States Coast Guard’s (USCG) station.

Navarro-head tower

Navarro-head tower, is the tallest structure that FurtherReach had built from scratch. It is 45 feet tall, comprised of five tower sections each 10 feet high, and about half of the lower section is buried under the surface in reinforced concrete foundations. Shown in figure 2.18 is the completed tower, with two of the solar panels slightly hidden at the bottom, the cabinet for batteries and gear at the base of the tower (we fondly call these cabinets, “dog houses”), the antennas, and the wind-turbine at the very top. Right below the turbine, are three small panel antennas *Ubiquiti NanoStation M5* devices to serve nearby subscribers. The two large dish antennas form the redundant back-haul from Manchester. Just above the larger dish is a weather station.[12] To save energy and prevent depletion of the batteries on dark and non-windy days, we recently began to shutdown one of the back-haul links, the more power-hungry of the two, from 1am until 7:30am every night. During these hours, should the secondary link to fail, the whole site would go down as the main link is down and we’ll loose the ability to re-power it. Given the unlikelihood of a any back-haul link to fail and coupled with the relatively unused hours of the night, we find this an acceptable compromise. In the future, we plan to re-power the main link upon loss of connectivity to the network’s core, but we lack the software ability to achieve this now.

[12] Accessible at: https://www.rainwise.net/weather/albionca and also at: http://www.wunderground.com/cgi-bin/findweather/getForecast?query=pws:KCAALBIO5 .

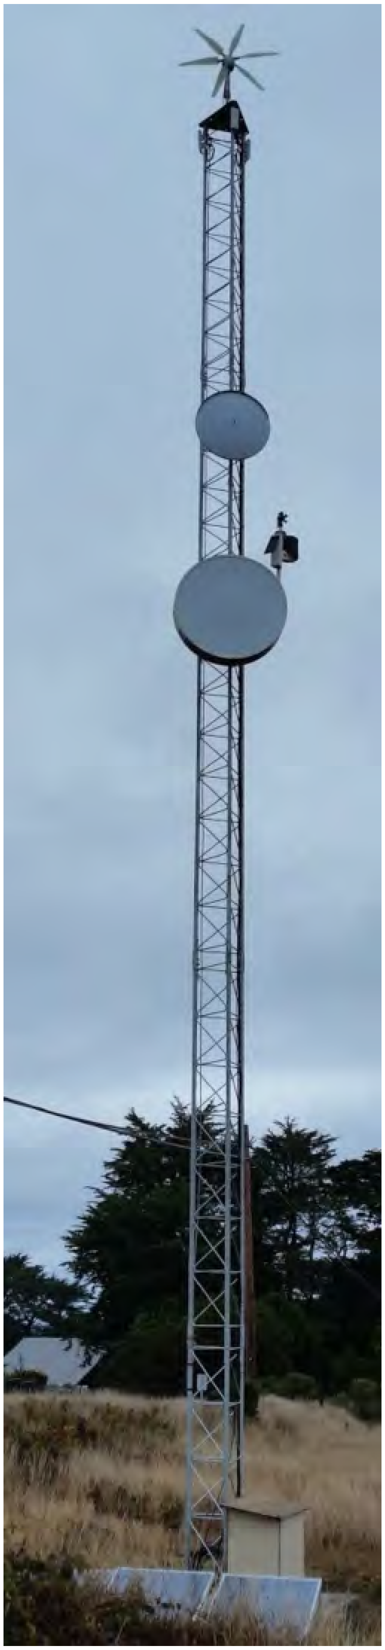

Figure 2.18: *Navarro* site — FurtherReach's tower powered by a hybrid of renewable energy sources.

Boomer's site

This site is our attempt at turning a large tree into a communications tower. The incentives for this attempt are quite obvious — the trees is already standing and might not be as costly as the erection of a tower, a tree does not require permits, it is perhaps visually more appealing than a tower, it is quite solid even at heights, does not require guy-wires, no cement-truck access is needed, etc. That said, turning this tree into a safe, tower-like work environment for our team, has proven quite challenging and costly — certainly far more involved than we expected. Visible in figure 2.19 are the shiny Aluminum steps we have fixed onto the tree using long screws that penetrate beyond the external shell and but avoid the delicate core of the trunk. These steps are fixed at a couple of inches away from the trunk, using PVC spacers, to provide safe footing. Also visible left to these steps is steel wire-rope, similar to the safety system common on communication towers — the climber clips a special ascender-like device onto the steel cable which allows free movement, both up and down, unless there is a sudden fall, in which the device gets wedged and arrests the fall. With this system, just like at communication towers, should the climber loose grip or a step would get dislodged this safety system would prevent a fall. We already invested in such tower-climbing safety gear and thereby desired to use it for the tree, instead of purchasing equipment used by arborists which is overall less safe. Less visible at the top, right to the steps, is a large dish antenna for the back-haul. There is a person standing to its right, also a bit hard to see, but demonstrate the size of that dish which is almost as big as the person. There is also a large sector antenna above it which is very difficult to see in that photo with the limb in the background. This site is not yet fully operational — it has been seven months so-far since we purchased the materials for this site. We needed to hire professional arborists to help trim large limbs that were in the way, and mount the safety cable as well as the steps. This site is also off-grid and the thick woods presents a challenge for solar-panels placement. Additionally, it will be a distributed site, in which the main back-haul antennas are at the south-West side of the mountain-top property and the ongoing wireless links would be mounted on another tree at the north part interconnected with a fiber-optic cable, but essentially two separate banks of solar panels and batteries.

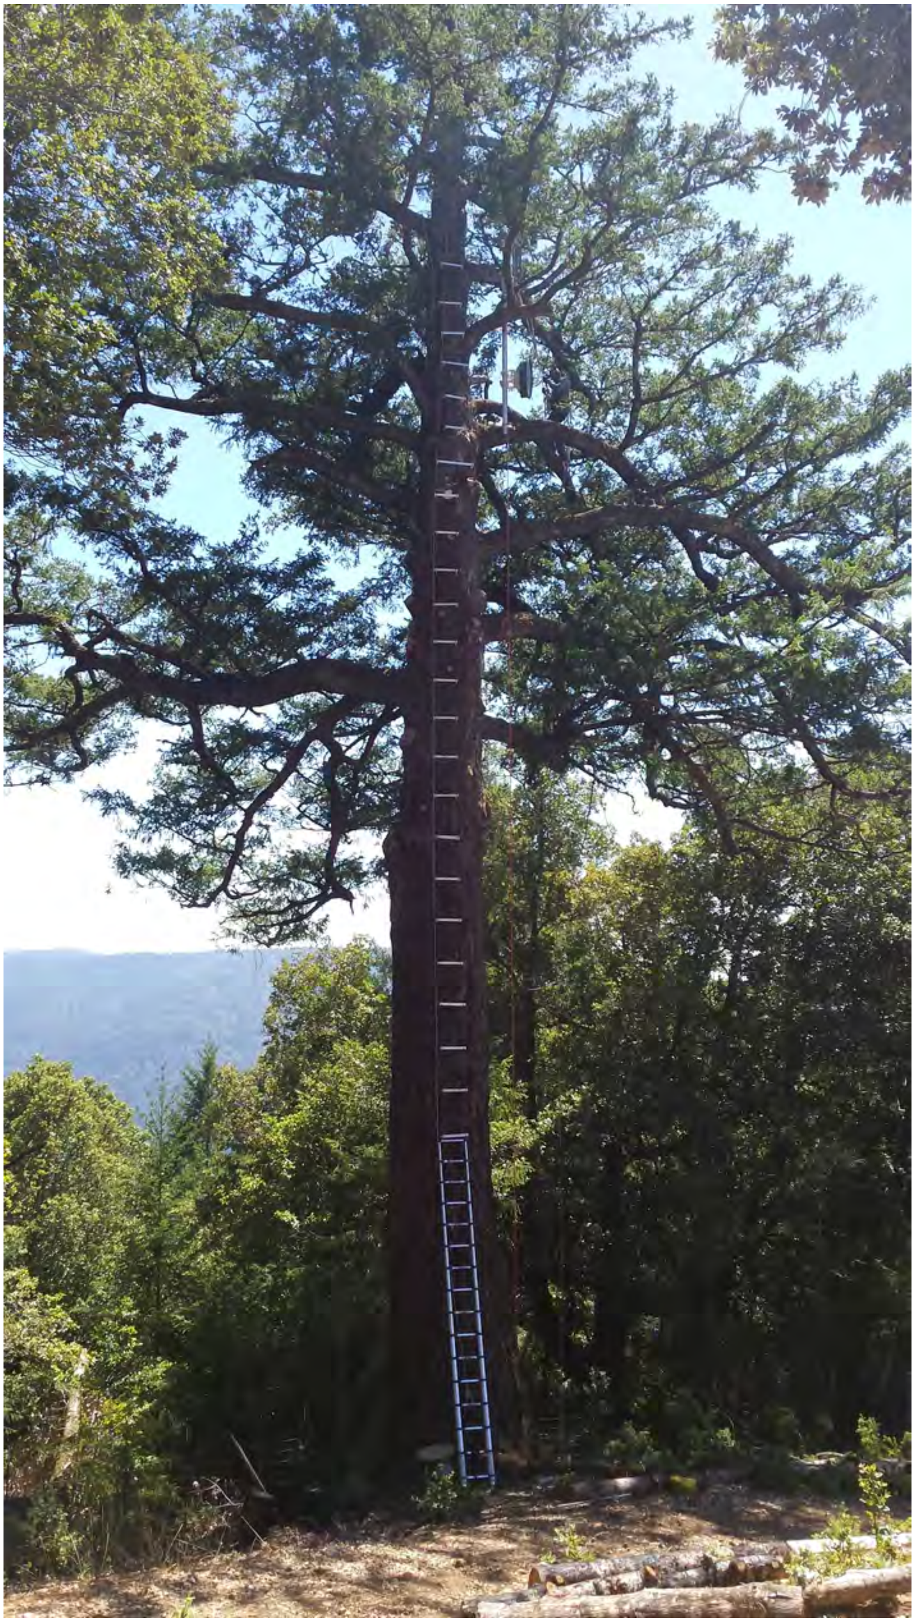

Figure 2.19: *Boomer* site — an attempt to turn a tree into a core relay.

2.6 The Celerate Architecture

Next we describe the design and implementation of Celerate, an architecture for easing the startup and management of WISPs. We use Celerate to manage all aspects of our deployment. Celerate is a work in progress, and will likely require a few years of field deployment and additional refinement before it reaches maturity.

Challenges

Since Celerate is directly informed by the challenges we have encountered during our deployment and the findings of our WISP studies, we briefly describe the specific challenges that WISPs face that Celerate is designed to respond to. We emphasize that all the challenges listed in this section have both been borne out by our experience and were noted by our study participants. In this we keep squarely in focus our finding that the key way to increase rural broadband connectivity is to help create more WISPs rather than scale any individual WISP; these new WISP operators are unlikely to have the knowledge or training of existing WISP operators.

Skill Sets. Building a new WISP requires diverse skill sets. Many rural WISPs are operated by a single person, perhaps with part-time contract help for infrastructure work. At small scale, rural WISPs simply can't support a large team of specialists. Thus a single individual is tasked with challenges as diverse as tower and tree climbing, network architecture, carpentry, IP security, negotiating land use, spectrum management, customer support, billing, and reviewing legal agreements. Indeed, it is exactly because of this diversity of required skills—not commonly found in any one individual—that we hypothesized that new tools were needed for such rural WISP networks, to simplify the management of the network and give the WISP operator room to focus on physical infrastructure upgrades, maintenance, and customer support.

Commodity Gear. Almost all WISPs rely upon low-cost commodity wireless hardware and existing software; building custom hardware and implementing custom software is too time consuming, expensive, and requires skills beyond the average WISP operator. In turn commodity wireless hardware vendors operate with thin margins, and as such provide meager software support or flexibility for their hardware. As a result, network management systems, such as SDN systems, cannot leverage support in individual devices, but instead must manage a heterogeneous network of SDN-oblivious commodity devices each with their own quirks.

Operational Issues. The nature of and approaches to resolving operational issues differentiate WISPs from conventional operators. It is not that conventional ISPs or large network operators do not have many of the same operational challenges, but the issues they face are of a different scale: WISPs have a far lower ratio of human expertise and resources relative to the challenges faced in network operation and a far higher ratio of physical and operational challenges relative to the size of the network. This is largely due to dispersed, less-reliable infrastructure and more difficult deployment environments.

Geography and Land Use. In terms of performance, fixed wireless is almost always an inferior option to wired service, and thus is typically used when wired infrastructure is unavailable or is too expensive to deploy. Often, the causes for a lack of wired infrastructure is a lack of population density, rough terrain, or, more often, both. Thus WISPs begin with multiple disadvantages: they must deliver service to a sparse region (thus making it hard to recoup fixed costs), using wireless instead of wired infrastructure, and do so over challenging terrain. Just because land is not developed in rural areas doesn't mean the landowner will make it available to a (cash-poor) WISP. Lack of existing tower infrastructure with good coverage of the region, especially line-of-sight coverage, means the WISP must build it or find an alternative. Indeed, the best prospects for placement of directional wireless gear is on existing structures such as water tanks, masts, sheds, barns, and the like.

Public Policy. Policy constraints tend to come in two forms: restrictions on spectrum, and restrictions on physical deployments. Spectrum restrictions are not new: WISPs are generally limited to operating on unlicensed spectrum only due to lack of licensing fees, but that spectrum is shared with a variety of other non-WISP devices, complicating spectrum planning. Physical restrictions typically take the form of regulations on tower construction and placement; for example, we found that placement of towers over a certain height under certain zoning required complex permitting and public approval, a process that can take a very long time and can easily end a WISP deployment before it begins.

Summary. The primary challenge posed by geography, land use, public policy, and similar non-technical issues is not the time it takes to resolve them—in due course, WISPs are able to resolve each issue in some way. Instead, the primary challenge is that these issues place non-technical constraints on the design of a WISP network and of the resources that can be deployed to parts of the network; they constrain how WISPs can manage networks on the ground and affect even what hardware they can deploy. It is within these constraints that we aim to address the above challenges to meet the needs of rural WISPs.

Subsystems

Many of the operational challenges in running a WISP are left out of traditional systems that aim to provide network management support. To make this clearer, next we differentiate between four subsystems of network management that make up Celerate: the device subsystem, the network subsystem, the management subsystem, and the operations subsystem, which we depict in Figure 2.20. Each of these subsystems corresponds to a plane that communicates with subsystems above and below. As we move up the hierarchy of subsystems from the device subsystem to the operations subsystem, we deal with a broader technical scope, greater complexity of mechanisms and policy, and slower timescales. Below we provide a sketch of our design for each subsystem and describe initial steps we have taken to realizing these subsystems. We do not claim that the overall structure or the individual pieces of Celerate are novel; indeed, much of it builds upon the canon of networking research over the past two decades. The key difference between Celerate and other network management systems is that it aims to be holistic, since we built it to address the real challenges we faced in building and running a WISP, and we use and improve our Celerate implementation on a daily basis. As mentioned earlier, we have found that the management and

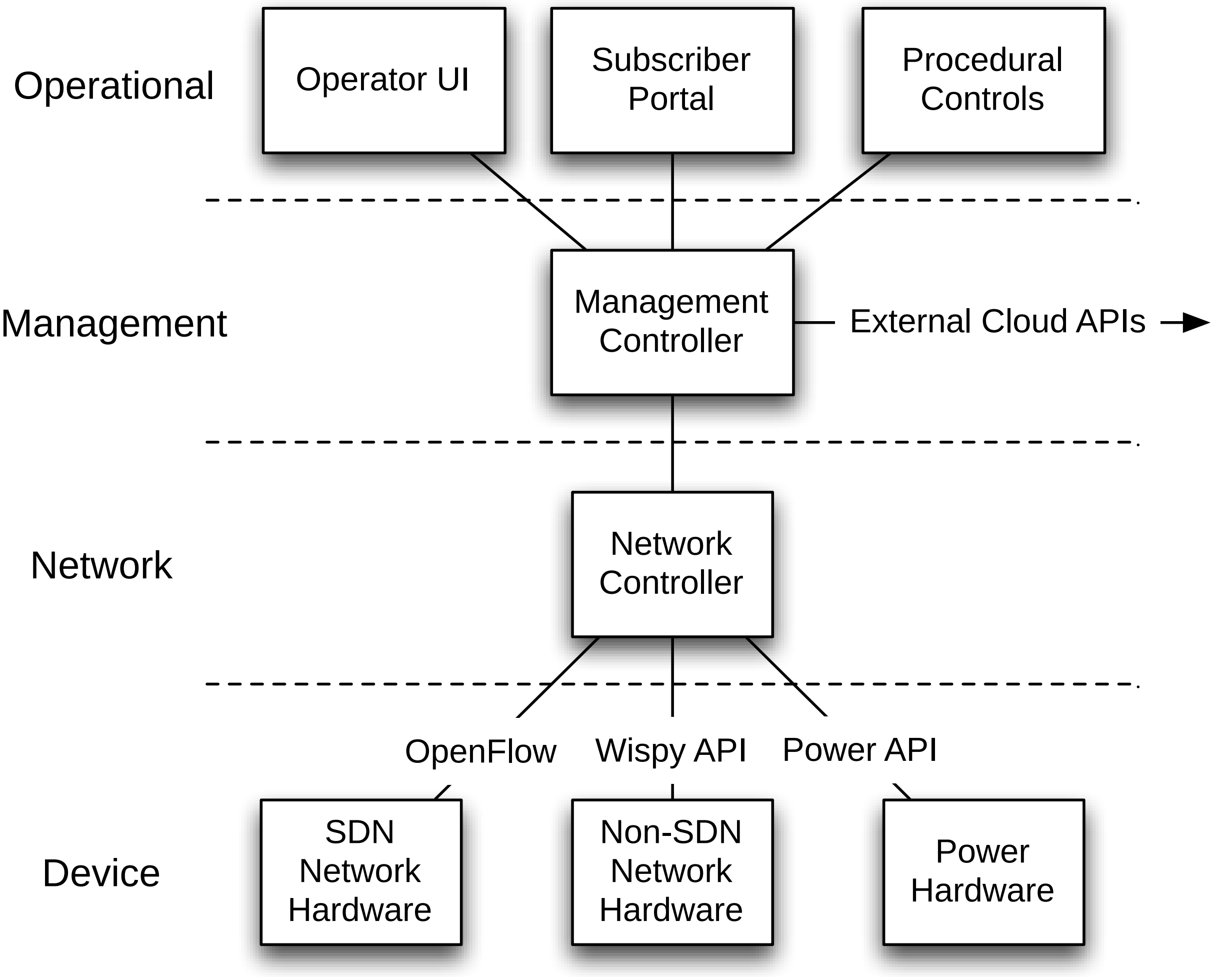

Figure 2.20: An overview of the subsystems in Celerate.

and operations subsystems to be most valuable in day-to-day operations.

Device Subsystem

The device subsystem corresponds to the traditional data plane in a networking context. In our context, however, the data plane is just a subset of functionality exposed. For example, in our deployment, the device subsystem includes SDN-enabled networking hardware, commodity non-SDN networking hardware, Power-over-Ethernet (PoE) switches with proprietary interfaces, power monitors, and battery systems. All of these devices provide interfaces for both control and monitoring, but they vary significantly in the interfaces they present. We are in the process of refining and developing APIs to manage these devices; the details of these APIs are relatively mundane, but are exactly what is needed to uniformly control devices that expose different types of functionality. At the moment we expose three APIs to the network controller above: OpenFlow, the Celerate API, which exposes common non-OpenFlow networking functionality such as advanced traffic shaping, SNMP, and wireless control, and the Power API, which exposes the ability to control and monitor power.

Most of the commodity networking devices actually run a forked version of OpenWrt [8] and can therefore run standard Linux networking tools. However, the manufacturers have added their own layer of configurations on top, generally in an effort to provide an easy to use GUI. Either way, attempting to connect the network controller directly to these devices would leave a messy, tangled architecture.

Network Subsystem

The network subsystem includes the traditional control plane in a networking context, but extends slightly further, since in our context the network subsystem has to communicate with a higher subsystem above it in addition to a more diverse subsystem below (that includes several non-traditional devices). Our network subsystem consists of a traditional SDN controller (currently built upon POX [9]) suitably modified to address three of the key challenges we have in managing our dataplane. First, since the device subsystem is highly diverse, the SDN controller has to have support for APIs beyond OpenFlow. Second, the controller receives its ground truth not solely from the network, but as a synthesis of live status of the network from these diverse devices and management state from the subsystem above. Third, the controller requires the ability to speak with the device subsystem, but unlike many SDN deployments (such as in datacenters), we have no easy means of out-of-band communication between the controller and the devices it controls. Thus we have been developing an approach for broadcast-based in-band control for SDN, to ensure the network controller can always communicate with devices in the field.[13]

These two subsystems are, roughly speaking, the standard two subsystems that SDN is concerned with. Beyond them, in industry and academia, a large array of diverse and incompatible approaches have been applied to higher subsystem management, usually tailored to specific applications and settings. To our knowledge, none of these higher-layer subsystem approaches are applicable to the unique needs of WISPs, and so we have designed our own, as we discuss next.

Management Subsystem

The management subsystem corresponds to the systems that track a range of information necessary for WISP management: network topology, physical deployments, network site documentation, and subscriber data including information about their physical installations, billing, and configurations. The management subsystem contains the ground truth about the network topology and about users of the network, and conveys this information downward to the network subsystem and upwards to the operations subsystem. We augment this ground truth with periodic automated discovery of new devices that are installed by technicians. We also integrate inventory tracking and management into Celerate.

Our management subsystem implementation is non-traditional for networking software, but an approach that is both flexible and easy to integrate with other systems: we use the Meteor Node.js

[13] We believe that this in-band control approach may have broader applicability, but is out of scope for this paper.

(Javascript) framework, and our current implementation consists of over 5000 lines of Javascript, 2000 lines of front-end templating code, and hooks into numerous third-party modules and services that provide functionality, such as financial transaction processing, that must remain external to our system.

By separating the management and networking subsystems, we allow for the management subsystem to become disconnected or intermittently connected and not affect the operation of the network. In doing so, we can host and operate the management subsystem elsewhere (using cloud hosting), removing the burden from the WISP operator while not hurting the delivery of Internet service; setting up and running a cloud instance is, to our surprise, hard for some WISP operators we have spoken with, and so this separation allows us to offer it as a hosted service for other WISPs to use. In addition, the management controller can make calls to external cloud-based APIs, something that would be dangerous to do from the network controller.[14]

Operations Subsystem

Finally, the operations subsystem corresponds to the operators and the operational tools used to manage the entire WISP. The operations subsystem is largely human, but also consists of tools for managing and tracking operational and billing issues in the network, monitoring the status of the network, notifying operators about failures, and procedural controls to ensure that operator actions are constrained. This subsystem is a crucial aspect of SDN as applied to WISPs. Indeed, keeping track of subscriber relationship information is the core functionality a WISP needs, and thus is a core of Celerate. In addition, the operations subsystem enables an integrated view of all information needed by a WISP operator, from new-user signups to billing to link status and capacity to power status, and can present a network manager with an entirely new perspective on their network.[15] In addition, we provide a portal to subscribers similar to that of most large ISPs, so subscribers can see their plan information and pay their bill.

The Operator UI presents a user-friendly interface for technicians in the field to easily search and update this information. Indeed the user-friendliness of the interface, while far from the traditional challenges that are faced by SDN systems, is actually one of the most important, in our experience, to the overall usefulness of the system. Some of the most urgent software fixes we have made are when the operator interface has a bug that prevents field technicians from using it easily on their devices in the field.

While a WISP may have a very professional installation crew, depending upon them to configure hardware in the field—wireless spectrum, routing, services including logging and monitoring, etc.—can be putting too much weight on their shoulders. One of our long-term goals it to enable the operator—with the help of the network management system—to preconfigure devices which can be "plug-and-play". Especially as a WISP starts and begins to grow, not requiring a networking

[14] One of these external APIs is Stripe, a PCI-compliant credit card processing system which enables us to process credit-card payments without handling credit card data directly [10].

[15] This can be crucial in responding to weather. Our deployment region suffers from heavy storms, which can render our otherwise strong 24GHz backhaul link entirely unusable. We plan to leverage Celerate to enable an operator to anticipate this in advance based upon forecasts and adjust network flows accordingly.

expert to administrate the network can make the difference between success vs failure, or high vs low performance. We hope to deploy our prototype implementation of such support in our network in the coming months.

Status

We have been developing each of these subsystems independently, have built systems for each, and are integrating them, a process that we expect will require an additional feedback from deployment experience and iteration. Contrary to what we expected at the outset, we have found that we have reaped the most benefit from our work on the management and operations subsystems, and as such much of our development effort has gone into these tools. As a result, we currently use only the management and operations subsystems in our field deployment. Our field technicians have come to rely upon these tools and use them daily when planning, expanding, and debugging the network.

Our challenge has been to define what each subsystem provides to and needs from other subsystems. For example, now that we have delineated for what information the management subsystem holds the authoritative copy (e.g., network topology, power monitoring configurations) vs the network subsystem (e.g., currently-active links based upon routing decisions, power-cutoff thresholds), there is the potential to fully integrate them.

2.7 Related work

For at least a decade researchers have advocated developing networking to meet the needs of poor and rural regions around the world [11]. The potential for WISPs to provide rural Internet access has been recognized for at least that long [12]. Prior work shows low-cost wireless hardware can deliver Internet access affordably and how to modify the MAC and PHY to improve performance [6, 13, 14, 15]. Similar techniques have been adopted by vendors of WISP hardware, making them available to the average WISP operator.

Surana et al. provided an early look at the operational challenges faced in rural wireless networks [16]. More recently, Rey-Moreno et al. describe lessons from the deployment of a community wireless mesh network [17]; this work however focuses on the particular challenges of operating a mesh network and in building a bottom-up community network. Similarly, Gabale et al. describe experiences from managing and a system for monitoring another rural mesh network [18]. While some of our findings and experiences overlap with those of prior work, we focus on the particular challenges around the *business* of operating rural WISPs.

Many WISP-specific management tools exist. HeyWhatsThat [19] is a mapping and link planning tool. Powercode [20], Swiftfox [21], and Azotel [22] are commercial WISP management systems that provide both subscriber, device and network management tools; TowerDB is a open-source project in the same space. While exact feature sets vary, all of these and Celerate solve similar problems for WISP operators. Celerate differs from these by supporting deeper integration with network devices and is designed to be a modular, open-source, and SDN-capable platform for WISP management.

2.8 Conclusion and Future Work

While our deployment and exploration of the benefits of improved management to WISPs is still young, we believe that our deployment has already showed, to us at least, that these approaches can help new WISPs start up and operate more smoothly. A local partner WISP has repeatedly asked for us to set up our systems for him as he sees that it will have immediate benefits for his operation, and we are in the process of doing so.

We expect that tying the wireless physical layer to SDN is likely to yield benefits for rural WISPs—especially when coordinated across multiple WISP operators—and will further simplify and automate these networks. Our long-term vision is of a WISP network that manages itself—a network that actively diagnoses failures and informs the local operator what to fix. While this vision is not a new one in networking, the context makes it particularly applicable, since WISPs are usually run by very small teams with limited skill diversity. Using SDN to globally coordinate the physical layer and network layer—taking into account RF connectivity, control of electronically steerable antennas, link bitrates, real-time workloads, and multiple backhaul paths through the network—would be beneficial for WISP networks (and would also be an attractive avenue of study, one that we plan to pursue).

That said, this work shows that the low-hanging fruit for WISP deployments comes from easing rather straightforward management burdens. WISPs are largely small operations constrained by their physical and economic environment, with stages of growth marked by jumps in capital expenditure. While the industry as a whole would benefit from regulatory changes like more allocation of spectrum for unlicensed use (which, happily, is a policy priority for the U.S. regulator), every small WISP that starts must to some extent re-invent and re-discover the best practices and systems necessary to run a sustainable network. To this end, we have developed Celerate as a modular and extensible system for WISP management that addresses the full stack of business concerns of the WISP. In doing so, we believe Celerate can serve to lower the difficulty of starting a WISP, increasing their number and improving access to Internet in rural areas.

We are also considering using electronically steerable antennas, which are becoming available and affordable, to further enable new deployment and operational models under a controller that enjoys a global network view—indeed, without such a global view, steerable antennas in such a deployment would become wholly unmanageable. Exotic bands, such as TV whitespaces, the result of recent regulatory changes, may also be of interest, especially when controlled globally to avoid interference on one hand, while benefiting from better wireless propagation qualities on the other. We are hopeful that these innovations can be rapidly productized—making commodity devices and open software available—and replicated across many rural communities especially in developing countries as it is inspiring to see how the communities where we deploy our testbed networks are benefiting from these network deployments.

In addition, remote management services could enable external experts to help WISPs with broader network management issues. This would provide a stepping stone towards service decoupling, where local WISPs are infrastructure providers over which multiple IP service providers can offer service.

Needless to say, clearly documenting our operational and business models, as well as our deployment strategy, and making these easily accessible to WISPs, is a part of our mission which we hope would prove beneficial.

Attracting additional researchers and developers to test and experiment with innovative new solutions over the FurtherReach network is another desired goal.

Though at the end of our work we will not be able to address many of the non-technical challenges involved in starting a rural WISP (many of which we describe earlier), we hope to significantly ease the process of building WISPs to connect the last billion users to the Internet.

Chapter 3

How can SDN help enable rural Internet

3.1 Introduction

Almost four billion people – two thirds of the world's population – lack access to the Internet. Of those, over 90% are in the developing world, many of whom either live in rural locations far from existing Internet infrastructure, are too poor to afford connectivity, or face a combination of both [23, 24]. This situation is unfortunate, as Internet access can enable economic growth and enhance individual freedoms [25]. Businesses without reliable broadband Internet service cannot take advantage of cloud services and the efficiency benefits such services provide, nor can they engage with customers and suppliers electronically – practically a requirement for any business hoping to operate beyond their locality. Broadband Internet is being used to deliver cost-effective social services such as telemedicine [14] and distance education [26]. Even cell phone networks, arguably the most successful technology in terms of global adoption, rely on high-capacity backhaul to base stations to provide service. These use cases and others demonstrate that increasing global broadband availability in these areas is a key international development challenge.[1]

Although Internet penetration is increasing, a disparity in access exists between comparatively wealthy, urban areas and poor, rural ones that is unlikely to change for the foreseeable future without a change in the underlying network technology. The latter are fundamentally difficult for Internet service providers to serve profitably: sparse population densities reduce opportunities for oversubscription, low purchasing power of potential customers implies small profit margins, and resource constraints make providing acceptable service quality hard. The history of rural wireless networks is rife with "pilot projects" that never reached meaningful scale or slowly fell apart when the (often US-university-affiliated) team who installed them left the area. Realizing the benefits of Internet access requires profitable Internet service providers whose customers trust them to provide reliable service over the long term. This requires innovation to drive down the costs of network

[1] We do not claim Internet access is a panacea for ending global poverty; clearly meeting basic human needs such as food, clean water, and healthcare are more immediate challenges. However, once basic needs are met, as is increasingly the case in all but the least developed countries, Internet connectivity can enhance existing capacity and drive further development.

operation and enable service providers to operate sustainable businesses, even in the most rural and poor areas.

Recent technical innovations such as modifying commodity 802.11 WiFi equipment for use at long distances (100+ km), coupled with novel deployment methodologies [27] have reduced costs at the physical layer, as have regulatory decisions to allocate microwave spectrum for unlicensed use. A new class of Internet service providers has developed as a result, which utilize point-to-point, "fixed wireless" access technology to provide service to remote, sparsely populated areas. We refer to such an organization as a *Wireless Internet Service Provider* (WISP).

Yet infrastructure cost is only one component of the cost structure for rural ISPs; profitable operation depends on controlling management and support costs. In this regard, a fundamental problem remains: rural wireless networks are highly variable, heterogeneous environments that are very difficult to manage, yet WISPs must do so while understaffed and under severe resource constraints. Rural wireless network operators thus need a fundamentally new paradigm for network management; without solving this problem, WISPs will not become widespread.

We believe that this challenge presents an important opportunity for the software-defined network community. SDN offers a principled approach to managing rural wireless networks and provides opportunities for making their operation simpler and more efficient, and may enable them to alter their fundamental business models. Specifically, by decoupling the control and data plane, SDN enables a WISP to decouple *construction* of physical infrastructure, which must be done locally, from *configuration* of their network, which can be done remotely. Going further, this decoupling enables the infrastructure deployment business and the ISP business to be operated by *completely separate entities*. As a result, software-defined rural wireless networks can decrease costs and lower technical and business barriers to entry, thereby enabling profitable operation of rural WISPs and expanding access to the Internet. For the SDN community, the rural wireless environment represents a radical departure from typical data center deployment environments. Rural networks are highly resource constrained, rely on aggressive traffic engineering, and present a dynamic physical environment that is difficult to manage. Thus, applying software-defined networks to rural wireless networks will "push the limits" of what is possible in SDN.

We acknowledge that SDN is not the only model that could enable the decoupling we discuss in this paper. That said, a key challenge for doing so in existing networks is the lack of a consistent global view of network state, as well as the lack of a standard interface a third party could use to configure an operator's network. Modern software-defined networks provide both of these. Existing SDN controllers and control protocols may even be adequate for WISPs, though the rural wireless network environment is more heterogeneous and resource-constrained than the data center environment for which existing systems were designed.

In this paper, we explore the opportunity for software-defined networking in rural wireless networks. Section 3.2 describes the environment of a rural wireless network through a brief case study of a large rural wireless network in the Indian Himalayas. In Section 3.3, we discuss the role that SDN could play in such networks. We consider the implications of fully virtualized rural wireless networks in Section 3.4. We then present related work in Section 3.5 before concluding.

3.2 Rural Wireless Network Operators

Rural wireless network operators have several defining characteristics and challenges that set them apart from other networks. In this section, we describe AirJaldi, a large rural wireless network operating in the Indian Himalayas. We also describe several actual operational experiences from the AirJaldi network. These illustrate the important yet unexpected issues faced by WISPs in the developing world.

Case study: AirJaldi

AirJaldi is a social enterprise whose goal is to empower rural communities through provision of affordable, wireless, Internet access. Started in 2005 as a single-person operation in the Himalayan town of Dharamsala in India, the headquarters of the Tibetan government-in-exile, AirJaldi now operates multiple profitable networks spread across a number of Indian states. The largest network, in and around Dharamsala, serves approximately 10,000 users within a radius of 120km. AirJaldi uses outdoor 802.11 microwave radios in the 2.4 and 5.8GHz bands to link subscribers' rooftops to a central Internet gateway using relay stations on mountain tops. Some of the wireless links are 50km or longer, though the majority of subscribers are within 15km of a relay station. AirJaldi operates its own wireless backhaul links to connect to its upstream Internet providers in nearby cities or towns. The rapidly evolving upstream ISP market in India causes a relatively high churn in the number of upstream ISPs and total upstream capacity used by AirJaldi; at the time of this writing the upstream capacity of the Dharamsala network totals 15Mbps, primarily provided by two incumbent ISPs with additional ADSL lines from a 3rd ISP as backup and for congested hours.

The business unit in Dharamsala employs 20 permanent employees, half of whom are technical operators and installers. With the exception of a handful of top-tier subscribers, most clients enjoy a small subset of Internet protocols and applications; apart from web browsing, email and VoIP, most other ports are blocked for the majority of users. Moreover, most subscribers' web browsing is subjected to content filtering proxies to reduce bandwidth used for services such as pornography, rich media, and P2P file sharing. In addition, subscribers are subject to downstream bandwidth shaping based on the purchased plan (256kbps, 512kbps, *etc.*). These limits allow short-lived bursts to enjoy higher bandwidth and thereby satisfy most users better than competing non-burstable services such as ADSL. Network management is based on availability monitoring using pings from the central NOC using tools like Nagios [28] and a set of automated scripts to push configuration changes to remote routers. At the central Internet gateway, tools such as `ntop` [29] are used to observe network load and identify misbehaving flows. Manually configuring the network, with its high node churn, dynamic workloads, high failure rate, and frequent security incidents is a challenging undertaking, particularly when serving novice Internet users who are unfamiliar with the various failure modes of the network.

Tales from a WISP

The experiences described below reflect actual operational events faced by AirJaldi network operators, providing a flavor of the type of technical and non-technical issues that further complicate network management for a WISP. A large and well-trained staff would make most of these problems are trivial, but that is cost-prohibitive for a WISP, and current techniques make automated solutions difficult.

Congestion and oversubscription. A subscriber calls the Network Operations Center (NOC) complaining about poor throughput on their connection. The operator notices a large amount of bandwidth being consumed by a major subscriber, company B. Because the NOC operator knows that company B is closed due to a holiday, this arouses suspicion; further investigation reveals a user at company B is accessing bandwidth-intensive pornographic content. The operator calls the CEO of the WISP to approve routing company B's traffic through a content filter to reduce the volume of traffic. The CEO approves, and after implementing filtering, congestion at the WISP's border link is greatly reduced.

Environmental issues, unreliable upstream, and tolerant users. It's monsoon season, and part of the WISP's service area is hit by thunderstorms. Many subscribers in a particular district are disconnected, indicating failure of the local power grid (most of the WISP's customer premises equipment is solar powered, so it can observe that customers' internal Ethernet links are down). Some hours later, one of the WISP's two upstream providers goes down: it's a cellular operator with a central office in that district, and their batteries must have depleted after hours of power failure. Unfortunately, this failure occurs at a peak usage time, forcing the WISP to scramble to cope with the loss of half its upstream bandwidth. All subscribers' traffic is routed through aggressive content filtering proxies, and all but a small number of well-known ports (e.g., 80, 443, 456) are blocked. The lowest tier of subscribers – free users who cannot afford to pay for services – are temporarily disconnected. As congestion worsens, more aggressive filters are put in place, blocking videos and security updates. Measures are necessarily heavy-handed, as the WISP has no mechanisms in place to selectively throttle specific types of traffic or users. Despite this, no one calls to complain; in fact, users are surprised the service is up at all given the weather.

RF interference. A village the WISP serves receives its water supply for two hours every Monday and Wednesday morning. The WISP receives sporadic complaints of poor service quality in this village, and dispatches a technician to investigate. The technician determines the problems are due to RF interference, but the source of the interference is unknown, and even changing frequencies does not seem to help. After several weeks of investigation, they discover that an electrically ungrounded rooftop water pump was causing the interference; properly grounding the pump solves the problem. Despite its simple solution, a lack of monitoring tools prevented the WISP from correlating failures and thus more quickly identifying the source of interference.

3.3 Opportunities for SDN

Given the challenging management environment that rural wireless network operators face, we consider opportunities for software-defined networking to improve upon the status quo and facilitate the spread of Internet access.

Decoupling skills

Rural wireless network operators like AirJaldi perform two core operational tasks: construction of physical wireless infrastructure and configuration of that infrastructure. The skill sets required for each have little overlap, and thus a WISP must either maintain separate staff for each or provide training in both areas to all their staff. Both options are expensive and inefficient. Technicians who install physical network and radio equipment must understand RF propagation and microwave link planning as well as construction techniques like tower climbing and safe wiring installation. These skills are completely different than those needed by technicians in the network operations center, whose job it is to ensure policy-compliant configuration and operation of the network. Yet for today's rural wireless operators, configuration and status monitoring is tightly coupled to the physical infrastructure they deploy. Adjusting configuration parameters on individual routers and access points is commonplace, and troubleshooting link failures requires understanding the full networking stack.

Software-defined networking enables network virtualization [30], which allows network operators to treat their physical network as an abstract pool of resources, specify management policies against this abstraction, and let the SDN controller handle the configuration of individual network components. In doing so they decouple the physical network from network policy. Decoupling these tasks also enables a rural network operator to decouple the staff responsible for each and increase specialization among their employees.

Specialization is generally more efficient for any organization, but it has particular benefits for rural network operators. Both physical and network configuration require specialized training, but all types of networks need network administrators, not just rural wireless networks. A capable network administrator can seek out jobs in both urban and rural areas, knowing their skills will be in demand anywhere. This is a serious problem for network operators in poor and rural regions: after investing the resources to train network administration staff, those staff could apply their skills in an urban area where increased business opportunities enable firms to offer higher wages. Rural-to-urban migration and "brain drain" is a problem in rural areas globally, and it particularly hurts rural parts of the developing world by sapping talent and expertise from areas that need both. In order to retain their network administrators, the rural WISP would need to offer wages competitive with those of firms in urban areas, a financially untenable prospect.

In contrast, training physical installation technicians is a less risky investment for rural network operators. Physical installation skills are not easily transferable, nor is there a significant job market for such technicians in urban areas. The long-distance microwave links used by rural WISPs are not practical in dense cities as the unobstructed line of sight such links require is not generally available. More fundamentally, the market for physical installation technicians is limited to organizations that operate radio equipment, as opposed to the wide range of organizations who need network administrators.

Decoupling these tasks enables a novel solution to this training and staffing challenge: outsourcing network management. Specifically, given a complete global view of network state and an abstract, logical model for the underlying physical network, control plane management can be conducted from anywhere, even urban areas. Architecturally, rural wireless network operators would run one or more SDN controllers within their own network, but the policy description for those controllers would be crafted by the operator's own network administrators or by a third-party network management consultancy to translate business needs and service agreements into a logical network configuration.

Outsourcing network management presents new – though arguably more tractable – issues for the WISP. Network management outsourcing should not significantly impact performance or cause outages. For instance, the management outsourcing provider could also operate the WISP's actual SDN controller, but this would require all policy decisions to incur WAN RTTs and expose control traffic to a higher likelihood of failure. Even if the WISP operated their SDN controller in their own network, building a separate, reliable, "control channel" network is impractical; all control traffic must be transmitted in-band over links that may experience significant and non-uniform loss, congestion, and delay. As a result, deciding where in their network to place control logic is a non-trivial decision.

Research challenge 1 *Can software-defined networks be constructed to be resilient to unreliable control channels?*

Research challenge 2 *How does a WISP verify their network has been configured as intended by the third-party management provider?*

Virtual circuits as policy abstraction

Internet service providers are in the business of providing service to their customers according to agreed upon service-level expectations. Regardless of resource sharing, customers expect their service agreement to be met. This is indeed the key challenge for an ISP: oversubscribe the infrastructure sufficiently to operate profitably while being able to deliver the expected level of service to clients.

ISPs today rely on the benefits of statistical multiplexing and heavy-handed throttling techniques to achieve this balance. Implementing effective traffic engineering requires configuration of individual elements of the operator's network. This situation is ludicrous; it is difficult to accurately specify a client's service requirements using these ad-hoc and limited-expressibility techniques, which are the only mechanisms current networks allow. Moreover, techniques such as bandwidth capping and shaping do not always reflect the reality of an ISP's cost model; as long as their network is uncongested an ISP has little incentive to limit usage.

Research challenge 3 *How can high-level policy declarations, including the sophisticated traffic engineering policies needed by WISPs, be translated into concrete network configurations?*

Ideally, a WISP would be able to allocate a per-customer virtual circuit from their network and then specify a set of service requirements for that *customer circuit*. Such requirements could include network properties such as bandwidth, latency, and jitter, as well as service level agreements specifying how frequently service requirements could go unmet. In addition to service requirements, a network operator could allocate middlebox processing services to such a circuit. This is essential in many bandwidth-constrained networks: in order to adequately serve any number of clients, extensive caching and content filtering must be employed.

Standardization of tools

A third key opportunity for software-defined networks in rural wireless networks arises from the fragmented ecosystem of currently available tools. WISPs require a suite of tools for network monitoring, configuration, billing, and user authentication. Complicating the situation, tools from different vendors do not necessarily interoperate or expose common configuration interfaces. A similar situation exists for debugging and troubleshooting: with no unified or automated mechanisms for reasoning about the status of the whole network, operators are forced to rely on ad-hoc techniques for identifying and fixing faults. Individual WISPs develop institutional knowledge to cope with this situation over time through experience, often learned the hard way.

This situation is akin to having unique programming languages and architectures for every organization that develops software. Developers could discuss best practices, but these would be of limited generality between shops. Sharing common libraries would be impossible – clearly an undesirable situation. Yet this is the status quo for WISPs today; best practices are encoded in people, and the implementation of a best practice in a network is specific to its particular environment, precluding sharing. A solution naturally arises in an SDN: the controller presents a global view of network state, a well-defined API and programming model for accessing and modifying that state, and implicitly a standard abstraction for monitoring and managing equipment from multiple providers.

Research challenge 4 *How can management tools leverage an SDN controller's global network view to expose relevant network state to human operators, including those with limited expertise or who do not have easy physical access to network hardware?*

Research challenge 5 *How can operators develop "libraries" for fault identification and correction that are generic across networks?*

These are, in truth, engineering and product development challenges in addition to research ones. Yet solutions would be of immediate practical benefit to existing WISPs. Thus, this is a point of leverage to drive adoption of software-defined networking in such networks.

3.4 Towards a Virtual WISP

The opportunities for SDN described in Section 3.3 are each practical innovations that would directly impact rural wireless network operators as they build and manage their networks today. Yet the decoupling between construction and configuration that SDN provides also enables new business models for WISPs. In particular, these two tasks can be conducted by *completely separate entities*. Such decoupling lends itself well to the franchise business model that has been successful for AirJaldi, in which regional operators focus on construction and maintenance of physical infrastructure while the parent company oversees the network management.

The logical extension of this idea is that WISPs would, rather than acting as an ISP themselves, "rent out" their network to an established ISP. The task of the WISP, then, becomes simply one of building wireless infrastructure and ensuring it can be managed by an SDN controller. This model radically changes the way WISPs interact with existing telecoms. Rather than competing with incumbent telecoms, which often have monopoly status, government subsidy, or other strong competitive advantages, WISPs are able to cooperate with an incumbent provider. In this arrangement, the WISP provides an incumbent telecom access to new customers outside their existing service range. In return, the incumbent telecom brings their business expertise (and, if applicable, regulatory licenses) to the rural market. For example, billing customers with small and irregular incomes (as is common in the developing world [31]) is challenging. Large cellular service providers that serve remote areas face a similar problem, and have already developed payment infrastructures to cope with it [32, 33]. The WISP can allow such a cellular service provider to offer Internet service over the WISP's infrastructure and thus take advantage of the payment infrastructure the telecom may already have in place. The WISP itself only needs to bill and interact with its ISP "customer".

The problem in this scenario is that end users in rural environments are subject to a non-competitive market for Internet service. Whatever ISP has an agreement with the local WISP has, in effect, monopoly power, as building a competing wireless infrastructure is a significant barrier to entry for competitors. In the long term, we envision *virtual rural wireless network operators* that rent their infrastructure to multiple ISPs rather than only a single one. A fully virtualized WISP would be an *infrastructure service provider*, analogous to the role that cloud providers such as Amazon's EC2 plays in the server hosting market. Rather than directly providing service to subscribers, the rural network operator would provide infrastructure to existing telecom and Internet service providers, its "client ISPs". The WISP would present these service providers, its clients, with a virtualized abstraction of its network as presented by the SDN controller. Crucially, these ISP clients would be able to modify their slice's configuration without interaction with the VWISP, just as users of cloud-hosted virtual machines require no interaction with their hosting provider to deploy new services. While the WISP would still be responsible for building the physical infrastructure to connect their own network to potential clients, the subscriber would interact directly with the client ISP for billing and support. This model of service provision is also beneficial to consumers as it enables multiple Internet service providers to utilize the same physical infrastructure, thus increasing competition.

Once a client ISP obtains a virtual slice from the VWISP they should be able to configure it using arbitrary control mechanisms. Moreover, client ISPs should be able to specify their own network configuration policies in arbitrary ways to suit their individual business needs. In order to reason about the degree of service they can offer to their clients, subscribers also desire consistent performance and network behavior – without these, the client ISPs cannot enter into meaningful service agreements with their customers. These needs lead to the following two key research challenges:

Research challenge 6 *How can multiple, independent network configuration policies co-exist on the same physical network infrastructure while ensuring safety properties?*

Research challenge 7 *To what degree can isolation between tenants' virtual network slices be guaranteed given a resource-constrained underlying physical network with little redundancy?*

3.5 Related Work

As far as we can tell, our work represents the first application of SDN technologies to resource-constrained wireless networks. The intellectual contribution of our work is an analysis of the technical challenges involved in adapting SDN to networks with bandwidth constraints, unreliable control- and data-channels, high churn rates, and a shortage of human technical expertise.

Network virtualization. Our concept of virtualized WISPs, described in §3.4, is essentially an application of the network hypervisor pioneered by Casado et al. [30]; client ISPs define policies over their own logical network slice, which is mapped onto a single physical network. Similarly, the controller synchronization platform described in Onix [34] is directly relevant to achieving fault-tolerance and scalability in rural networks. The challenge in applying these technologies to the developing world is that unlike in the datacenter environment, rural wireless networks are loosely coupled, unreliable, and resource constrained.

Cellular networks. The business aspect of virtual WISPs is similar to roaming agreements in the cellular market: customers receive service from third-parties when outside their home service area, yet it appears to the customer that they are still receiving service from their own provider. Tower sharing schemes, in which multiple providers make use of shared tower sites and infrastructure, is used to reduce capital expense particularly in developing and rural markets. Incumbent providers may share infrastructure in ad-hoc arrangements, or through joint ventures into so-called "tower companies" [35]. While Mobile Virtual Network Operators (MVNOs), small providers who sell service directly to customers but do not own their own cellular infrastructure, are common, there is no equivalent concept of network virtualization in the cellular market, though multi-operator cellular base stations have been proposed [36].

SDN for wireless networks. The OpenRoads platform at Stanford University provides hardware, slicing among multiple tenants, and open APIs for experimenting with OpenFlow-based wireless networks [37]. Dely et al. further demonstrate the applicability of SDN concepts to wireless mesh networking [38] by using rapid OpenFlow updates to solve the host mobility problem. Our work expands this line of research into resource-constrained wireless environments and considers the operational implications of SDN for rural operators.

Managing rural wireless networks. Most work on rural wireless networks has focused on the forwarding plane rather than the management plane. WiLDNet [14] and JaldiMAC [27] respectively propose MAC protocols for point-to-point and point-to-multipoint long-distance wireless networks for WISPs, but do not consider how to manage large networks consisting of these links. Surana et al. [39, 40] describe a number of techniques for identifying faults in rural wireless networks as well as observed failure modes and practical solutions from field experience. Meraki offers a centralized, cloud-hosted network monitoring and management tool, though their offering is targeted towards well-provisioned enterprise environments as opposed to WISPs [41].

3.6 Conclusion

Our vision for applying software defined networks to the rural wireless environment may seem unconventional on its surface, but in fact our proposal is quite modest in the context of the broader SDN community. Just like operators of large data center networks, WISPs are not well served by the status quo for network management. Indeed, our agenda dovetails with current research trends in SDN. Consolidating control and management of a rural wireless network will simplify their operation, as will decoupling the tasks of infrastructure construction and network configuration. This decoupling further enables new cooperative business models for rural wireless networks. Taken together, we believe SDN has an important role to play in spreading sustainable, reliable Internet access to people worldwide.

Chapter 4

Rural computer security

4.1 Introduction

Computing security is an important concern for researchers working on issues of technology and development. From casual concerns about viruses in email or infections via USB key drives used in cybercafes, to serious issues with identity theft on shared networks, computing security poses a diverse set of challenges.

Although we do not argue that the threat in the developing world is necessarily *greater* than elsewhere, we do find that security issues in the developing world are *different*, and as such are underrepresented in research and practice. We argue the poor security situation is not simply due to a lack of technical tools: the root causes stem from a combination of technical and non-technical issues, and thus require a multidisciplinary approach for solution design. Examining these problems is not only valuable in providing a more secure computing environment in the developing world, but also holds the potential to re-examine issues of mainstream security studies, given the diversity of computing environments, networks, and risk behaviors.

Rapid economic growth and increasing technology usage, coupled with the continued existence of the five forces described in the next section, create a security landscape where attackers have both the incentive and the ability to inflict significant harm to technology users. This hypothesis is supported by evidence from the computer security world: the Norton CyberCrime Report attempts to quantify the time and money for recovery from cybercrime across 14 countries [42]. Juxtaposing these results with GDP reveals that these costs form a much higher percentage of GDP for developing countries (i.e., Brazil, China, India). We can also use proliferation of botnets as a metric for the impact of poor security, and again we see that India, Russia and Brazil are top sources of spam, with China being on par with the USA [43]. Similarly, we see that Brazil (10.5%), India (9.3%), Russia (7%), the USA (5.8%), and China (5.1%) lead the global botnet activity for dictionary attacks [44].

Aside from these measurable impacts of security issues, the indirect impact of poor security is equally damaging. In particular, compromised computing infrastructure is often unreliable and can negatively affect the confidence placed in technology rather than fostering its adoption. We observed this distrust in technology ourselves as part of our experiences with computer users in India. As a result, we formulate two hypotheses.

Developmental Impact: There is significant evidence that poor security practices prevail in the developing world, yet the extent to which this hinders development is unknown. Indeed, understanding the impact of poor security practices is challenging since often it is not manifested directly. Beyond the direct cost to consumers and indirect effects on technology adoption, we suspect poor security prevents technology-based development from achieving its full potential.

Risk: The second open issue is the relative level of risk in developing regions, where we define risk as the probability of being exploited multiplied by the cost of a security breach. The probability of exploitation is a function of both the quality of security practices and the incentives for attack. Although users in developed countries tend to have more resources for defending their computing environment, they also present a more lucrative target, due to higher incomes and more powerful computing infrastructure. The relationship among these factors is complex and we expect that there exist environments where all three (ability to defend, incentives for attackers, and impact of exploitation) are aligned to promote attacks. It seems likely that novice middle-income users, a large and growing group, are both easy to reach and worth attacking.

4.2 Security Landscape

In exploring this space, we identify five core forces that shape the security landscape in the developing world. We discuss the potential developmental impact they imply, and propose a multidisciplinary research agenda towards addressing the broad range of challenges.

Poor Security Hygiene

A key factor in assessing the security of a system is its “security hygiene”, i.e., the degree to which it runs with up-to-date software patches and recent malware protection. For example, systems with a high security hygiene regularly update their operating system and anti-virus software, have recent versions of security-critical components such as Flash or Java, and employ browsing blacklists (e.g., Google Safe Browsing). The concept of security hygiene also applies to the firmware of embedded devices, such as wireless routers. Naturally, systems with high security hygiene exhibit a smaller attack surface and therefore face fewer potential hazards.

Security hygiene in the developing world is generally poor. For example, one study found only 30-40% of systems in the AirJaldi [45] network performed anti-virus and OS updates [46]; another study of an Indian telecenter found high infection rates and wasted bandwidth due to Portuguese spam [26]. One of the reasons for this situation is the fact that Internet access in the developing world is characterized by scarce bandwidth and frequent failures. Even when an individual connection is fully operational, end-to-end traffic flows are subject to bottlenecks at upstream links or high packet loss rates.

These network conditions make obtaining security updates, patches, and malware signatures (which often require regular and lengthy downloads) a tedious and failure-prone process. The combination of slow download speeds with frequent failures and congestion often leads to failed downloads that need to be attempted again, leading to a self-reinforcing state of poor performance that is difficult to escape, and which may ultimately have detrimental effects on the security hygiene. Indeed, Maier et al. [46] find that the fraction of OS and anti-virus updates of the total HTTP traffic is larger in the AirJaldi network compared to the large European ISP, which could be explained by the intermittent connectivity that causes more updates to fail. To support this hypothesis, we analyze the TCP state of connections to `windowsupdate.com` and `update.microsoft.com` in the AirJaldi network and compare it to two research sites in Berkeley, California: the Lawrence Berkeley National Laboratory (LBNL) with more than 12,000 hosts, and the International Computer Science Institute (ICSI). Comparing connection logs from the whole month of February 2011, we find that 56% of all Windows Update connections at AirJaldi are terminated with a RST by the originator, indicating a failed update attempt. This is an order of magnitude higher than for ICSI and LBNL.

In industrialized nations, bandwidth is often seen as a commodity with relatively small cost. Not surprisingly, operating systems, software, and security applications alike take it for granted that patches and updates for their products will be delivered to clients online. With the prevalence of risk that is promulgated online, and given the near ubiquity of reliable Internet in the primary markets of software producers, this is arguably the most practical means of providing updates. However, the assumption of easily available, inexpensive, and reliable connectivity does not necessarily hold in the developing world, and may even be counter-productive as a security practice.

Further complicating this view of security hygiene is the fact that many types of malware spread without any online connectivity. Offline infection vectors such as USB storage devices have been shown to be very effective and often harder to protect against [47]; the disinfection and cleaning processes are significantly more complex without Internet access. These attack vectors can be very unexpected. For instance, while surveying rural branch offices of an NGO in Rajasthan, India, the author found several offline PCs running Windows XP that had been infected with viruses. Apparently, the viruses were reaching these machines via memory cards of digital cameras used by the NGO for fieldwork: there were very few external backup devices in these offices, but up to 70 memory cards were floating in circulation per office (for less than 15 PCs) Interestingly, the NGO reported that mainstream anti-virus programs, which offered adequate security in their Internet-connected urban offices, could not fend off all viruses in these rural locations.

Unique Usage Patterns

The unique resource constraints faced by technology users in the developing world has given rise to new computing environments and applications not commonly seen in the developed world. For example, people in developing regions are beginning to rely on mobile technology for conducting financial transactions even in places where credit cards and the web have not penetrated. The M-Pesa service in Kenya and its numerous clones throughout the developing world [48] enable people to manage savings accounts, make electronic payments and now even avail themselves of loans and micro-insurance [49]. Some of these services have recorded daily transaction volumes of over hundreds of millions of dollars [48].

Although these services are not new, the security of the underlying technology is poorly understood: unlike ATMs and credit cards in the West, there are no industry standards for building secure mobile banking technology; in fact, at least three such services have been successfully attacked in different ways in the past year alone [50, 51, 52]. The core challenge in designing secure solutions for this application stems from the fact that a majority of the phones available in developing regions are still of a very basic nature: they are either not programmable at all or offer only limited programming capability. This makes it difficult to adapt security solutions designed for other forms of electronic commerce into the current context, and practitioners are forced to resort to ad-hoc approaches, as done in current deployments. The situation is further complicated by the fact that these systems have been shoehorned on top of the GSM signaling channel (using protocols like SMS or USSD) which were never designed or intended to facilitate secure transactions. Given the increasing penetration of such services and the recent spate of attacks against them, it is of urgent importance that security researchers come up with an innovative remedy for this situation.

Another source of security concerns is shared resource computing. PCs in developing countries are predominantly a shared resource [53]. Most families do not own a PC and therefore use cybercafes, kiosks and community centers where PCs are shared with others. Extensive literature exists regarding the risks involved with the use of public terminals, as well as suggestions for mitigating these problems [54]. Although the use of public computers in the West is diminishing — and with it academic interest from computer scientists — securing this type of computing environment remains a common challenge in the developing world. Additionally, differences in elements such as purchasing power, literacy levels, and maturity of privacy and identity concepts may render existing solutions ill-suited to rural developing regions. There is much literature on the prevalence of both outgoing security threats through phishing activity and virus infection of individual machines in cybercafes and shared computer centers [55]. Shared machines at computer centers suffer from a "tragedy of the commons" scenario with regard to security, because individual users have no real incentive to keep the devices safe for other users. At the same time users at cybercafes are frequently driven by what may be termed as risky activities, such as pornography or online gaming [56]. A cybercafe manager who denies such services to his or her clientele risks losing customers. Shared computer centers are often the primary location not just for Internet access, but also general computer access. Users typically use the machines for a range of activities, like desktop publishing or writing a resume, for which they frequently use external USB storage. These are frequently infected, and given that for many users the storage is the only computing artifact they actually own, they carry the infection wherever they go. As with censorship, prohibiting USB keys is often not a viable business decision.

Novice Users

According to the World Bank[57], more than half of the Internet users in low and middle income countries have joined since 2005; the number of users in these countries has increased by over 190 million between 2007 and 2008 alone. Despite the obvious need, disseminating security educational material and tools is extremely challenging. Many of these users have their first experience with the Internet on a mobile device, a platform with its own unique security implications. A general lack of language support and localization of applications and technical instructional material can undermine user awareness of security risks. Content in the online environment is primarily in English and Chinese, further complicating the situation for the many users not literate in these languages.

Inexperienced users may exhibit a variety of behaviors that may adversely impact security, such as engaging in risky online behavior [58] or forwarding emails from questionable sources. While there is little existing work on the comfort of users with potential attack vectors such as pop-ups, online adware, or installing untrusted software, there has been some work on the spread of viruses through email. The damage caused by email-based virus spreading has been an important area of work in both engineering [59] and economic studies [60] in terms of the cost of such activities to individual productivity and to the network [61]. There have also been studies examining the motivations of people in forwarding emails [62, 63], though there is little structured empirical knowledge on the demographics of "potentially unsafe" email behavior [64]. Though there are significant variances in data, the general trend seems to suggest that there is a greater likelihood of email infection in the developing world, especially in China and India. The MessageLabs network infection report [65] finds that India and China are both leading sources as well as intended destinations of attacks.

Several news and marketing reports discuss the incidence of phishing and the consequent lack of user trust on Internet-based activity in India. A survey commissioned by VeriSign revealed that at least 76% of the web users in India are at risk from online fraud due to the inability to identify different forms of phishing. Recognizing the potential biases of marketing research, this data nonetheless suggests that low user awareness of security issues increases risk. Although there is no conclusive evidence that cultural factors make the developing world more susceptible to phishing, there is plenty of research that shows cultural differences in online behavior and purchasing practices [66]. Research at Symantec, for instance, showed that spammers and creators of malicious software are both aware of user behavior issues and target these specifically in their effort to expand botnets by exploiting religious festivals, such as Diwali in India [67]. Attitudes towards potentially dangerous online material also appear to vary by region. One study found that despite higher use of anti-virus software, users in rural India tend to click on Google search results that are clearly marked as dangerous and which most European users avoided [46]. Given the large influx of novice users, these issues remain a significant cause for concern.

Piracy

Our experience suggests that most proprietary software in use throughout the developing world is pirated. While use of pirated software is not necessarily a security risk, it is challenging to verify that such software is not malicious. Pirated operating system disk images, for instance, are difficult to scan due to their size; additionally malware can be easily hidden from the user once the operating system itself is compromised.

Moreover, verifying the source of pirated software is impossible, and the motivations of the creators of such pirated software are unclear. The pirated operating system distributions the author came across clearly required a non-trivial amount of work to produce, and many even include a variety of additional (pirated) application software pre-configured and ready to use. Despite this, the proliferation of these all-in-one pirated images makes some forms of recovery from infection easier: users simply re-image their machines when they have been compromised. As data loss is almost expected, the practical cost to re-image for recovery is low.

Use of pirated software also appears to decrease security hygiene. Older versions of pirated software are generally considered easier to obtain in the gray market; this is especially true for Windows XP, first released in 2001, compared to more recent versions of Windows [68]. Additionally, software vendors take measures to protect against software piracy, and while mostly unsuccessful it is far easier to confirm the authenticity of software when it requests security updates. As expected, many pirated copies cannot obtain security patches and are therefore left vulnerable to attacks. Since companies like Microsoft monitor computers making online requests for updates, users who are concerned about the legal implications of piracy may voluntarily avoid the update process.

The proliferation of pirated software is a multifaceted phenomenon. Despite a negative relationship between software piracy rates and economic development, software piracy cannot be explained by economic factors alone — a range of policy and cultural issues also plays a role [69, 70]. Legacies of weak regulatory controls and the network effects of widespread piracy contribute to a casual attitude towards piracy [71]. For many countries, the regulation of piracy is politically sensitive as it ties in with a range of other trade issues [72]. A tangentially related concern is the widespread piracy of music and movies, which often involve quasi-legal websites [73]. Both the content and the websites themselves are likely to pose malware risks. Because the regulatory and enforcement environments are unlikely to change dramatically within the next few years, piracy will likely continue to influence the security landscape.

One possible means of mitigating the security risk of pirated software is increasing awareness of open-source alternatives to popular software, which are generally less often targeted for attack than proprietary alternatives, if not actually more secure. As the authenticity of open-source software can be easily verified, the risk of "pre-compromised" software from unknown, potentially malicious sources is dramatically reduced. Open-source software can be easily bundled together for distribution, much like the all-in-one bundles of pre-configured pirated software currently in circulation; thus users can continue to enjoy the benefits of quick, low-cost recovery currently provided by all-in-one pirated software distributions without the associated risks.

Adversaries' Perspectives

An attacker's motivations for compromising a system determines in large part the types of risk faced by technology users. These motivations can vary widely, from individuals seeking personal financial gain,[1] to organized criminals offering for-hire botnet infrastructure, to governments seeking to achieve geopolitical goals [47]. Because the computing infrastructure in the developing world presents a set of advantages and disadvantages for cyber-criminals, the incentives for attack vary as well.

[1] A common example is the Nigerian letter scam, (also known as "419 fraud").

As a result of these complex relationships, the value of a compromised machine varies across nations. One way to measure the economic value of a compromised machine to cybercriminals is the "pay-per-install" (PPI) price, which is the cost on the global black market to buy rights to install malware per compromised machine. PPI prices for computers in the developing world are significantly (in some cases a full order of magnitude) lower than those in the developed world [74]. The market value could be lower for a number of reasons, including greater ease of infection (i.e., increased supply) or lower functionality of developing world hardware (i.e., decreased demand). A better understanding of how compromised machines are used in the developing world would shed light on the forces behind these costs.

Another issue faced in many developing countries is pervasive domestic monitoring, censorship, and other restrictions on the flow of information. While this issue is neither solely a problem of developing nations nor is it ubiquitous among them, freedom of expression is regarded as an important development objective [75]. Improved security techniques and practices, in turn, can help alleviate this problem.

Understanding the motivations behind attacks is a complicated issue, clearly in need of more study. Nevertheless, the perspectives of adversaries must be included to understand the broader security landscape.

4.3 Challenges for Future Research

Computing security in the developing world is clearly a multifaceted and multidisciplinary problem, and the five factors we identify as defining the security landscape in developing regions offer fertile land for further research. Although these represent only a small subset of the issues affecting computing security in the developing world, they highlight several important and pressing challenges in the area.

Policy

National and international policy plays a major role in influencing security practices, both positively and negatively. Most countries have legacy information and communication technologies (ICT) laws that are inadequate in the face of the rapid and complex changes occurring today, often leading to inadequate or counterproductive policies. For example, India prohibits the use of encryption, a policy justified (ironically) by national security, as it simplifies eavesdropping by authorities. In practice, this exposes its citizens to attacks by criminals, both domestic and foreign. This policy received significant attention recently, as many major international companies such as RIM, Google, and Skype offer encrypted services in violation of the ban; this is essential to protect their users from identity theft, financial fraud, and other security problems.

Technical

Several of the unique facets of the security landscape in developing regions call for new technological innovation. The problem of offline infection vectors calls for new virus-protection or virus-cleansing mechanisms that are predominantly designed for the online world. System administrators and kiosk owners need to be better-equipped in handling data loss and corruption that results from sharing PCs across multiple individuals.

The problem of ensuring secure communication on developing-world cellular phone networks is a fascinating research area as well, offering a rich variety of open problems for security researchers. As already pointed out, the mobile phone landscape in developing regions makes it difficult to simply import developed-world solutions for this problem. Given that this landscape is likely to persist for several years [48] and that the demand for security-sensitive applications in developing regions is on the rise, it is a good opportunity for security researchers to do interesting research that can have wide-scale impact. Even the fundamental notions of privacy and authenticity become non-trivial to achieve when operating under these constraints.

Finally, a fundamental problem plaguing the developing world is that of providing digital identities to citizens. In places like India, there is growing emphasis on deploying digital mechanisms to uniquely represent and verify the identity of each individual. It is hoped that such mechanisms will redress some of the problems that disenfranchised citizens have faced in the past due to the inability to uniquely represent themselves in government and other official transactions. Although there are obvious benefits from such systems, if carelessly implemented, these systems could damage individuals' privacy and human rights. Implementing a privacy-preserving unique identity and identification system under developing world constraints like limited infrastructure, intermittent connectivity and poor educational standards, is a problem that requires immediate attention from security researchers.[2]

Business Aspects

Technical procedures and business decisions made by ISPs directly affect subscribers' security. Unfortunately, these decisions are driven by business factors and often pay little concern to user security. For instance, an ISP's decision to deploy ADSL modems that do not support NAT exposes subscribers to malicious incoming connections. A study by Cui et al. [76] identified at least 540,000 publicly accessible embedded devices set with a default password, and their findings suggest that the actual numbers are much higher. Perhaps unsurprisingly, 80% of the vulnerable devices identified were in Asia. Our own observations from India confirm that BSNL, India's largest telecommunications company, also installed ADSL routers that were configured with default passwords, and were not configured to support NAT. What factors drive leading Asian ISPs to make these poor security decisions while their North American and European counterparts choose more wisely? We suspect the sheer volume and fast growth rates of subscribers in Asia play a

[2] The UID project in India is attempting to build a biometric-based identification system for citizens that would operate even in locations with poor connectivity. A security analysis of the privacy preserving mechanisms used by this project is missing in the literature.

role, as do considerations about maintenance and support. The length of the purchasing process is likely another factor and is affected by purchasing volumes as well as political and regulatory environments.

Awareness and Education

Increased awareness of security risks and practices is crucial to improving the computer security situation throughout the developing world, and educating users about these is itself a challenging problem. Affecting this change requires solid grounding in the current state of security education; this present state is largely unknown and unstudied in many developing regions. Indeed, while many of our points about security are based on only anecdotal reports, they suggest potentially significant variations in the way developing world users perceive computing security. In particular, there are many open questions regarding what the actual impact of poor security practices has been on development outcomes. Another related question is how these issues impact organizations such as businesses and governments; indeed, poor computing security practices in these environments can have significant direct and indirect consequences.

4.4 Related Work

It appears that computing security in the developing world has been mostly overlooked by computer scientists. Apart from exceptions discussed earlier [46, 50, 51], addressing specific concerns, other disciplines have dominated this space.

Several regional studies have surveyed policy issues related to security, but there are distinct gaps in the data available. Studies on countries in the Asian region have fairly good data on network use as well as threats [77] and make a strong case for cross-national cooperation on security[78, 79]. There is comparatively less work on Latin American and African nations. Work on the African continent has found that with the exception of a few countries in North Africa and the Republic of South Africa, there is limited legislation or institutional capacity regarding cybersecurity issues [80]. There has been a particular spike in the interest in secure networks due to concern about terrorism-related issues[81, 82], the political issues surrounding privacy and security[83].

Schmidt et al. [58] consider the perspective of user perception. Their study compares awareness and understanding of malware and spyware among Chinese and US firm workers, and finds Chinese workers to be significantly less aware of risks due to malware than US workers.

4.5 Conclusion

At its core, the goal of computing security is to ensure users of computing systems can trust and rely upon those systems to accomplish their desired tasks. Failures in computing security present direct impacts, but more fundamentally they undermine the utility of a society's computing infrastructure. Increasing access to, relevance of, and reliance upon ICTs in developing regions promises significant development benefits. While the problems that poor computing security presents to these users may be merely inconveniences today, without the trust and reliability that are provided by strong computing security the promise of ICTs for development will not be fully realized.

We have reached an unusual paradox: most of the people affected by computer security are outside the focus of its research and design, and every year the disparity grows. Although different, the problems of users in developing regions are at least as interesting, and increasingly important – for users, for development, and for overall global security.

Issues such as shared computers, limited training and literacy, and piracy all require a combination of disciplines to achieve real improvements in security. The problems are only part technical, but nonetheless require both new research and technical leadership to drive policy, education, and the deployment of more secure systems.

Part II

Dissent Networking

Chapter 5

How should we build Dissent Networks

5.1 Introduction

In the wake of the 2011 Arab Spring, international attention focused on the role that the Internet and social media services such as Facebook and Twitter can play in supporting popular uprisings against repressive regimes. At the same time, the actions of these regimes demonstrated the fragility of the infrastructure that connects people to these services, as well as their willingness to use the full power of the state to engage in large-scale censorship of the Internet and other communication networks. In response, researchers and technically-minded activists around the world have started projects which aim to build censorship-resistant communication networks. Their goals vary, ranging from building an alternative Internet infrastructure outside the control of corporate or government interests to building emergency communications infrastructures for times of crisis. Yet for the most part, they all share a common goal—building networks that can survive serious disruption to existing communications infrastructure while ensuring free expression among their users.

We concur that this is a worthy goal. However, we believe that much of the work in this space suffers a disconnect from reality which stems from a lack of a clearly defined set of properties for such networks. To this end, we propose and define "dissent networking", discuss the desired properties, and address the suitability of proposed technologies and solutions (or lack thereof). Dissent networks aim to allow free expression even in the face of censorship and communications blackouts. Dissent networks are:

- *Resilient against communications blackouts.* Should be challenging for any entity to disable.
- *Resistant to monitoring and tracking of users.* Both who is using the network and any sensitive messages they send should be secret.
- *Able to be built from innocuous components.* Should only require readily available hardware, and the possession and use of required hardware shouldn't be illegal or suspicious.
- *Able to run at meaningful scales.* Should be more effective at disseminating information than people with megaphones; more broadly, given a level of service, should be able to run at non-trivial scales.

The most common proposal to meet the needs of this space are wireless mesh networks. We argue that traditional mesh networks face an inherent tension between their ability to fulfill the first three facets of our definition, which they can do at small scale, and the last, which they can only do by compromising on one (or more) of the first three. As a result, we do not believe that current proposals for such networks constitute an effective countermeasure to Internet censorships or blackouts. We emphasize that we do not aim to dismiss wireless mesh networks out of hand, but instead focus our criticism on common assumptions in proposed mesh-based solutions and present design-level approaches for getting around these shortcomings.

The core contributions of this work are a taxonomy of wireless mesh networks and their relationship to dissent networking, and a set of requirements for effective countermeasures to blackout circumvention.

5.2 Related Work

Many systems for blackout circumvention have been proposed recently. The Commotion Wireless project [84] is building a customized firmware to enable WiFi access points and other devices to form mesh networks, with a focus on ease-of-deployment. Serval has developed a WiFi mesh mobile telephony system [85]. The Free Networking Foundation [86] aims to support the development of community-owned censorship-resistant networks. Rangzen [87] is a privacy-preserving mobile mesh network that leverages social ties for routing. These projects all leverage WiFi-based mesh networks to varying degrees and each carry the explicit goal of building censorship-resistant networks.

Beyond these projects that aim to build independent network infrastructure, several others focus on circumventing other forms of Internet censorship. Tor is an overlay network for secure and anonymous communications on the Internet using a peer-to-peer network of onion routers [88]. Ultrasurf [89] and Freegate [90] likewise enable secure and anonymous Internet access, though these rely on centralized proxy servers. VPNs and proxy servers are also commonly used to bypass censorship. While these projects fill a similar need to the one we discuss in this paper, they all assume the existence of some underlying form of connectivity and thus provide no resistance to blackouts.

5.3 What is a mesh network?

The basic idea of a wireless mesh network is relatively universal: multiple devices (“nodes”) each communicate directly with their neighbors, and messages from one node to another are forwarded through the mesh via intermediate nodes. This contrasts with “infrastructure” wireless networks, such as cellular phone networks, where client devices (e.g., cell phone) communicate with a master device (e.g., a cell phone tower), which is typically connected via a wired link to the rest of the provider’s network. While “infrastructure” networks are best thought of as a hierarchical tree, mesh networks are often thought of as a well-connected graph. Akyildiz et al. [91] provide an overview of the space.

Beyond this basic definition, however, mesh networks can take a variety of forms. Consider the following three definitions of mesh networking taken from a few projects in the blackout circumvention space:

> Mesh networks afford an alternative to [the] centralized "hub-and-spoke" WLAN model: rather than relying on the ISP for Internet connectivity, mesh technologies can produce ad hoc networks that allow distributed nodes to act as the senders, receivers, and conduits of information. In the mesh model of networking, "each user has the capability to receive and send information and to relay information on behalf of other connected computers." **Berkman Center**

> [A] mesh wireless network offers the ability of users to connect directly to each other and facilitate a distributed network infrastructure that provides multiple paths for communication to the network and does not require a centrally-located towers [. . .]They can bypass obstacles, [. . .]have no single point of failure, and are easily expandable. **Commotion Wireless**

> A mesh network is one where any device can be connected to one or more other neighbor devices in an unstructured (ad-hoc) manner. Mesh networks are robust and simple to configure because the software determines the routing of data automatically in real-time based on sensing the network topology. **FabFi**

The above examples illustrate the variety of (sometimes conflicting) attributes that characterize mesh networks. In general, these networks fall across a design space defined by three main trade-offs.

Planned vs. Organic growth. The growth of a planned mesh network is intentionally designed and laid out. Such a network may use antennas which require careful alignment, implement strict policies regarding which devices can and cannot join the mesh, or rely on careful management of radio spectrum use. In contrast, organically-grown mesh networks grow without a particular goal for their topology without needing to coordinate the placement of new nodes. These networks typically utilize non-directional, low-gain antennas and rely on automated routing protocols to allow the mesh network to grow without explicit human involvement.

Centralized vs. Decentralized management. The human organization that operates a mesh network can be centralized or decentralized. In an organizationally-centralized network, a single person or group is responsible for a network's operation and management. In a decentralized network, multiple independent groups cooperate in some way to build a single mesh network, and no one entity has control over an entire network.

Static vs. Mobile topology. In a static mesh network, the mesh nodes are fixed and immobile. Typically, static networks utilize dedicated wireless routers as mesh nodes. Mobile mesh networks use mobile devices as mesh nodes, such as smartphones or laptops. In general, the dynamically changing conditions of a mobile mesh network make them harder to manage than a static mesh

Project	Characteristics
Freifunk [92]	Planned, Centralized, Static
Meraki [41]	Organic, Centralized, Static
Serval [85]	Organic, Decentralized, Mobile
Freedom Tower [86]	Organic, Decentralized, Static

Table 5.1: Major examples of mesh networking projects and their space in our taxonomy.

network. Note that we refer to the mobility of the *infrastructure* from which a network is built; the fact that a mobile device can connect to a network (such as a smartphone using a WiFi access point) does not make the network a mobile one.

Table 5.1 shows how various major mesh networking systems fit into this taxonomy. This taxonomy highlights a key tension for projects that wish to use mesh networks to overcome censorship. To successfully resist communications blackouts, a networking technology should grow organically, be mobile, and employ decentralized management—widely available radio direction finding equipment can identify the location of mesh nodes, and any centralized management system represents a single point of failure for the whole network. Unfortunately, as we'll see in §5.4, building mesh networks that scale and function efficiently is challenging without being planned, static, and centrally managed. The incompatibility of these two goals places serious constraints on the viability of wireless mesh networks as an effective blackout circumvention tool.

5.4 Scaling Mesh Networks

The capacity scaling of wireless mesh networks has been well-studied in the literature. Gupta and Kumar's foundational result [93] proved that the per-node capacity of a multihop wireless network approaches zero as the number of nodes increases. Li et al. provided experimental validation of this result for 802.11-based networks [94]. This point bears repeating—under reasonable and practical assumptions, the capacity of a mesh network provably tends to zero as it grows. Both of these results, however, are primarily theoretical, and make strong assumptions about properties of the network such as link rates, external interference, coverage radius, and node layout. While we emphasize these results are nonetheless quite general (the Gupta/Kumar result, for example, holds for arbitrary networks), an intuitive understanding of how and why mesh networks scale is useful for practical situations.

Capacity of Mesh Networks

Channel contention is the primary factor that prevents per-node capacity in mesh networks from scaling. Mesh nodes carry traffic on behalf of other nodes in the network; critically, each node can transmit and receive from multiple other nodes. Mesh networks typically use omnidirectional antennas ("omnis") to support communication regardless of the relative orientation of nodes. Antennas are passive devices that concentrate RF energy; omnis have radiation patterns resembling spheres or disks. Other radiation patterns are possible using directional antennas, but again these can only focus a node's energy over a smaller area; these are less useful for mesh networks since they limit the degree of each node. Using omnis is a design decision to prioritize unplanned deployment over efficiency: most of the energy transmitted by each node is wasted by being radiated away from the recipient.

Yet poor efficiency is not the real problem. Note that the radiation pattern of an antenna applies to both what it transmits and what it receives, and rather than just two nodes, consider a regular lattice of nodes that is evenly spaced. For simplicity, we assume that each node has a fixed radius over which it can successfully transmit and receive messages, and that nodes are spaced by less than this radius[1]. When node A transmits to node B, none of A's neighboring nodes can receive any transmissions due to collisions. To these nodes, A acts as source of interference at node A's location, no different than government jamming equipment. This highlights a key property of nodes with omnis—not only do they cause interference in all directions when they transmit, they are susceptible to interference from *any* direction. If nodes use a carrier-sense MAC protocol such as 802.11, the problem is more insidious—even if one of A's neighbors wanted to transmit to a node outside A's transmission radius, it must wait until A's transmission ended.

The problem is further compounded by the fact that most commonly available radio equipment used for mesh networks has only a single transceiver, which is a half-duplex device. While multiradio equipment is available today, laptops, mobile phones,[2] and consumer-grade access points rarely have more than one. The multiradio equipment that is available is specifically designed for mesh networks; we argue that such purpose-built hardware makes targeting dissidents easier (§5.5). Finally, although we assume that nodes themselves generate the network's traffic, nodes can also serve as access points for devices like laptops or phones; this (suboptimal) design can lead to increased contention.

Channel contention carries two implications. First, mesh networks have poor performance because of time wasted waiting for opportunities to send traffic. Second, mesh networks have highly variable performance [95] since the scale of contention varies significantly based on workload (along with environmental factors that affect radio propagation).

Application Support

Despite these challenges, meshes can provide a useful degree of service—*if* applications running on them can tolerate their unique shortcomings. For example, so-called "smart meters" use mesh networks to report customers' usage to their utility companies; messages are forwarded across the network to "gateway" nodes connected to the Internet. This is an application particularly well-suited for mesh networks. First, it is highly *delay tolerant*—as long as the utility company receives its billing data within a few minutes or even hours the data is still useful. Secondly, it requires

[1] This is essentially the model used by Li et al., though here we assume the transmission and reception radius are equal. Of course, real-world RF behavior is much more complex.

[2] While phones and laptops often *do* have multiple radios (e.g., WiFi, Bluetooth, and cellular) typically only the WiFi radio is used for mesh due to support for "ad-hoc mode" and legal constraints.

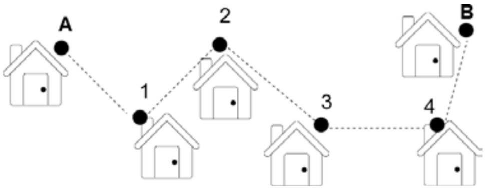

Figure 5.1: Multihop communication across a mesh network.

little bandwidth—even with low absolute efficiency the mesh is still able to meet the application's performance requirements.

In contrast, web traffic is a workload that performs poorly in these high-contention environments. Consider the basic task of a user sending a TCP request and receiving a response over a multi-hop mesh network as depicted in Fig. 5.1. We assume "oracle routing" that determines the optimal path for all traffic, though doing so in practice is challenging. TCP requires the network to send bidirectional traffic: packets from A to B will generate acknowledgements upon receipt. This is a problem for typical mesh networks that use single-radio nodes which share the same channel. When A in Fig. 5.1 transmits to B, each packet must be received by nodes 1-4 and can only be retransmitted when the channel is free, halving effective bandwidth at each point. Each time a node in the path transmits a packet, none of its neighboring nodes may transmit or receive, lest they create collisions. Synchronizing transmissions is a challenging problem; WiFi-based networks utilize a mechanism known as RTS/CTS to announce their intention to transmit. While this mechanism reduces collisions, it increases the amount of time the channel is idle: each RTS/CTS exchange between nodes requires at least two transmissions before sending actual data. Every ACK that B sends back to A undergoes the same process. The end result is that each wireless hop substantially decreases effective bandwidth and increases latency and loss, even in this simple case. Multiple pairs of communicating nodes exacerbate the problem.

Mobile nodes enable capacity to scale linearly under certain assumptions [96] but introduce new opportunities for loss and delay (e.g. nodes not being in range of each other). Highly variable latency and loss due to collisions are standard conditions in a mesh network, and since these violate assumptions of TCP congestion control mesh networks tend to be ill-suited for TCP-based applications. Mesh networks present a challenging environment for voice traffic (which requires low jitter) for similar reasons. Directional antennas also improve capacity [97], though networks using them topologically resemble non-mesh networks and rely on purpose-built hardware, a problem addressed in §5.5.

This discussion suggests two strategies for building mesh networks that scale. The first is to reduce the degree of channel contention in the mesh network by carefully planning how nodes can interfere with each other and where new nodes are added to the network. Such a network provides a high level of service, but wouldn't be a "dissent network". The second strategy is to accept the limitations of mesh networks and build applications that can work under those regimes. For example, applications that leverage delay-tolerant networking [98] principles can cope with such limitations [99, 100], as can very low bandwidth applications. Such applications could prove quite useful for dissent networks [87].

We finally note that our discussion ignores several key unsolved problems in scaling wireless mesh networks. Most notably, routing across ad-hoc mesh networks continues to be an area of active research and engineering effort. We've chosen to ignore this for two reasons. First, this paper focuses on real-world networks in real-world environments. Few mesh routing protocols have seen the level of sustained development and testing necessary to fairly judge their ability to function in such environments. Second, and more importantly, our criticism of mesh networks for blackout circumvention is an *architectural* one, and is orthogonal to the routing protocol used. Even with a "perfect" routing protocol, mesh networks cannot overcome the fundamental physics of radio from which their scaling properties derive.

5.5 Supporting Dissent

Our objective here, of course, is not simply to tear down wireless mesh networks. There are several examples of mesh networks that have scaled well and serve large numbers of users, such as Freifunk or the Athens Wireless Network. Yet the bar for dissent networking is higher—such networks will be used in environments where even the act of using such a network puts the user at risk. Centralized and planned networks can't work in this environment, as they have a single point of failure, and static mesh nodes are easy targets for a government with even the most basic electronic surveillance equipment. Not only can a repressive government shut down a network by attacking the technology itself, it can also attack the organization and people behind it.

The goal of work in this space is to promote freedom of expression under oppressive regimes—in short, to support political dissent. At their core, censorship and suppression of communication are *non-technical* problems; while technical solutions may alleviate their direct impacts, the root issue is one of unjust governance. Technology doesn't produce political movements. A key idea from the technology for international development literature states that technology only amplifies human intent [101]. Put differently, technology plays a *multiplicative* role, not an additive one. Moreover, technology amplifies both positive and negative intentions [102]. Any anti-censorship tool can thus only build upon existing social movements and simultaneously carries the potential to amplify the efforts of repressive regimes (e.g., by providing another mechanism to track dissident activity).

This presents a pair of related challenges to dissent-oriented projects. First, such projects should leverage existing social trust networks. Doing so simultaneously builds upon pre-existing social infrastructure while using that infrastructure to reduce risk to users.Secondly, such projects should minimize the extent to which the systems they are developing could be used for harm. We emphasize two particular elements of this second challenge—the need to use "innocuous" hardware that doesn't raise suspicion and the need to provide anonymity (not pseudonymity) guarantees to users.

Innocuous Hardware

Projects that propose illegal or restricted hardware face challenges for sourcing equipment, may put activists and users at increased risk, and provide an easy excuse for government crackdowns. Import restrictions on radio equipment are enforced the world over; attempts to smuggle such gear could end not only in confiscation of the equipment, but even arrests and severe punishment. Exceptions are typically made for WiFi devices and similar gear that operates on unlicensed spectrum, yet some countries have not deregulated the use of such spectrum. While customs tend to overlook importation of innocuous equipment like laptops and smartphones, they reserve the right to confiscate equipment and harass citizens who attempt such import when the regime feels it might be used "inappropriately".

Additionally, using licensed spectrum makes it easy for a regime to identify and locate who is using it (particularly if the user has applied for a license). A network using such spectrum without a license provides an easy excuse for a government to terminate operation and dole out punishment to the network's operators. Moreover, illegally using licensed spectrum carries the risk of disrupting operations of legitimate license holders, who have an incentive to report such activity to authorities. This is a plausible outcome for activists who, for example, set up unlicensed cellular systems (even low-power ones) as has been proposed by some groups [84].

In countries which allow use of deregulated spectrum, setting up rooftop WiFi antennas may not be outright illegal, yet may be a cause for government scrutiny and harassment. Due to the conspicuous nature of such equipment, its existence can raise suspicion from neighbours or agents and collaborators of the regime. The governments of such countries incentivize citizens to report any suspicious dissent activity and thereby turning them into informants and collaborators. Fear of persecution therefore renders widespread adoption of such rooftop technology unlikely, even if the actual gear and spectrum use is within the boundaries of the law.

Given the above, we believe the use of unconventional, purpose-built, or otherwise uncommon equipment is likely to be shut down quickly, limiting the impact of projects using such hardware. Worse, such equipment could put users at risk or aid an oppressive regime in tracking users. We conclude that the only devices that can be used in a meaningful dissent network are innocuous, ubiquitous devices such as smartphones and preexisting indoor WiFi access points.

The Need for Anonymity

Adversaries can analyze communication content and patterns to identify dissidents. While tools like encryption ensure communication security, dissent networks also require *privacy*. While security protects communications from eavesdroppers, privacy aims to limit the information revealed by legitimate communications. Such communications may involve malevolent agents, so it is critical to avoid leaking condemning information. We wish to prevent persecution of individuals based on their involvement in such a network.

We contend that the only truly safe solution to this problem is anonymity. Users should ideally be unlinkable with their true names, but this may be impossible in practice due to surveillance. A potentially sufficient alternative is deniability—users should be able to make a plausible case for

innocence. This requires two levels of protection: 1) The network should be useful for non-dissent purposes, so usage is not incriminating in and of itself. 2) Activity on the network should be impossible or difficult to trace. To achieve the latter point, systems may need to satisfy a number of notions of anonymity:
Author anonymity: It is impossible to link a message with its author.
Reader anonymity: It is impossible to link a document with its readers.
Document anonymity: Servers do not know which documents they are storing.
Query anonymity: A server does not know what client request it is filling.

These varieties of anonymity are defined in [103], and many dissent networking solutions address subsets thereof via pseudonymity—i.e., users are associated with network identities disjoint from their true identities. However, pseudonymity is not safe enough for dissent networking, since attributing profile information to individuals facilitates identification. It has been shown repeatedly that personal information in social networks can be correlated with external information to deanonymize users [104]. Pseudonymity can also be implicit, enabling similar threats. For instance, fixed-infrastructure networks can lead to localization and deanonymization of users [105, 106]. In the allegedly anonymous Bitcoin network, researchers learned information about individual users by observing transaction patterns [107]. Decentralized mesh networks are more robust to traffic analysis because interaction records are difficult to trace, so the main concern is avoiding explicit pseudonymity.

Ideally, a dissent networking solution should have all the above anonymity properties, but the network must still function as a communication tool. Theoretical results from other domains have demonstrated fundamental tradeoffs between privacy and system utility [108], suggesting that similar tradeoffs may exist for communication networks. For instance, some networks rely on user trust graphs (useful for deanonymization) to defend against sybil attacks. The relation between privacy and communication efficiency is an important research question, but strong privacy should nonetheless remain a conscious design goal for this space.

5.6 Moving Forward

This work takes a critical view of proposed blackout circumvention systems; we acknowledge we offer few explicit solutions. We nonetheless believe that there is good work to be done in this space.

Dissent-oriented mesh networks can improve by leveraging mobility, directional antennas, and limitation-tolerant applications, while providing strong anonymity. There are several examples of work that partially meets these requirements for a successful dissent network. For instance, the Dissent and TOR projects incorporate notions of deniability and anonymity into the system functionality [109, 88]. Projects like Commotion and Serval exploit mobility and delay-tolerance in a mobile mesh setting, while avoiding exotic hardware [84, 85]. Ideally, systems should aim to address *all* the requirements; this is attempted in [87], though the practical scalability of such a solution is yet unproven. Along these lines, we hope the community will consider "communications" broadly while pushing to build workable dissent networks. Though mesh networks will likely encounter scalability problems for applications like telephony or even point-to-point communications, other models (e.g. one-to-many communication) have yet to be explored.

5.7 Conclusion

Developing effective countermeasures to communications blackouts involves requirements beyond what most existing projects have set out to meet. Mesh networks, the most commonly proposed solution, suffer a fundamental tension between scale and safety for use under a repressive regime. Such networks can reach meaningful scales by adopting centralized management, planned growth, and a static topology, making them more susceptible to government interference. Networks can retain a decentralized nature at the cost of lower quality of service, requiring applications tailored to their limitations.

More than this, we feel that prior work has not paid enough attention to the fact that building alternative network infrastructure is itself a subversive act. Those who build such systems do so with the full awareness that the design choices they make can have grave consequences for their users. At the same time, given the public resources that have been directed to this space, these blackout circumvention systems should be able to scale to meaningful sizes—beyond just demonstration deployments. We believe that our definition of dissent networking captures these two goals, and that projects that attempt to meet our definition will produce more effective countermeasures to communications blackouts.

Chapter 6

Rangzen - designing a dissent network

6.1 Introduction

Recently, tyrants have attempted to subdue political turmoil and civil uprising by imposing large-scale communication blackouts. These blackouts consisted of shutting down Internet access, cellular and wired telephone systems, and at times even the power grid. Our goal is to provide an alternate communication solution that is independent of government- and corporate-controlled infrastructure. We coin the term *dissent networking* to describe communication or networking solutions that facilitate civil dissent.

Anonymity

Anonymity is a critical property for dissent networking due to fear of persecution; we consider intimidation and punishment of the authors of subversive content as part of our threat model. Studies in this area typically focus on secrecy and authentication[110], properties that conflict with anonymity. Some studies suggest pseudonymity[1] to preserve user privacy[111], but pseudonymity has been shown relatively easy to deanonymize, especially through correlation with external information[112]. This approach has been moderately fruitful in purportedly anonymous systems like the Bitcoin network [113], and the dangers are even greater in a dissent networking context.

Anonymity-preserving prioritization

A key challenge for any community-owned, decentralized, communication network is that of resource allocation and control. The network's finite resources must be shared among citizens in a manner that mitigates the effects of unwanted traffic and abusive users while allotting higher capacity to desired content. Intuitively, the way to achieve such a prioritization mechanism is by means of community-reputation systems—content from reputable users should receive higher priority.

[1] Decoupling of real users from their network identities.

However, such mechanisms require the network to collectively store information about individuals in order to pass judgment, thereby reducing the degree of anonymity within the system.

Balancing the need for anonymity with prevention of network abuse and attack resiliency is the principle design challenge that we tackle in this paper. We see this as the first of many pertinent research challenges on the road to implementing a Rangzen network.

Risks of exotic hardware

Several systems have addressed similar challenges in the past by proposing hardware-dependent solutions like rooftop antennas or improvised towers [114, 115, 116]. However, given the harsh restrictions that are typically enforced by oppressive governments, there are significant dangers in the setup and operation of solutions relying on alternate infrastructure elements. To avoid these dangers, we maintain that a dissent networking solution should be solely comprised of regular smartphones loaded with an enabling software application. In doing so, Rangzen leverages existing communication capabilities within phones (e.g. WiFi, Bluetooth) while removing further dependencies on infrastructure like the cellular network.

Fundamentality of DTN

We acknowledge the numerous pilots and vast body of research on mobile-mesh networks—projects that failed to scale beyond the lab. The majority of failures have stemmed from attempting to support Internet-like, online, end-to-end connectivity that conflicts with the store-and-forward communication paradigm. The store-and-forward paradigm, on which peer-to-peer ad-hoc meshes are based, increases resource contention exponentially with every added node, thereby extending latency and limiting scalability. Moreover, it is unlikely that our target localities will be dense enough to provide the desired end-to-end coverage, even if we had a way to avoid contention. High node mobility and churn also detract from the network's ability to establish end-to-end connectivity.

These fundamental challenges in supporting Internet-like connectivity over a mobile mesh led us to focus on a disruption- and delay-tolerant network (DTN) paradigm. Although ill-suited for many Internet applications, it provides a robust packet delivery fabric that is grounded on an extensive body of work, primarily from the sensor network research community.

In a DTN-mesh, phones exchange traffic when they opportunistically come within radio range of each other and collaboratively relay messages on behalf of other members. Such epidemic-like diffusion of content embraces mobility to overcome wide geographic gaps[2] and does not depend on high node density for delivery. The DTN-mesh framework is naturally conducive to broadcasting messages, which motivates our choice to design a microblogging tool.

[2] A smartphone that travels on a bus may link remote locations or even countries.

Threats and Goals

We assume throughout this paper that the proportion of government agents in the system is very low compared to the number of citizens ($\approx 0.1\%$). We speculate that no system dominated by adversarial agents could reliably protect its legitimate users. Despite this low ratio of agents, Rangzen is designed explicitly to circumvent government censorship, so we anticipate a number of unique threats.

Radio Jamming: Traditionally, the first threat that comes to mind when discussing any sort of wireless communications is of radio jamming. We believe that Rangzen is resilient to such an attack, given the close physical proximity that is required for two mobile phones to establish a direct wireless connection. Generating a jamming signal that is strong enough to overshadow a transmission of a nearby sender is unfeasible at scale. The government may apply powerful and focused jamming at key locations where people congregate, which indeed would disturb and spoil the plentiful opportunistic data exchanges that would otherwise occur in such locations. However, given the mobility of devices, such focused disturbances are likely to be insignificant for the overall Rangzen network and costly for the government.

DoS, Information Poisoning, and Sybils: Another threat is denial of service attacks that flood the messaging system. These may come from an oppressive government or a more mundane attacker like spammers. This attack involves the malicious use of devices that spread nonsense or misleading messages. Malicious devices might be active government agents or artificial Sybil (fake) devices. The government can set up wireless routers in various places as a means of interacting with citizens' devices. These routers may impersonate agents, essentially making them omnipresent throughout the country.

We have designed our message prioritization algorithm in Section 6.3 with these attacks in mind. As we introduce features and complexity to our design, we will note where we see other possible threats and how our design provides means of resistance.

Goals: Any solution to this problem should exhibit the following essential properties: It should enable communication with low latency while providing resilience to the aforementioned threats. It should scale gracefully to support hundreds of thousands of nodes. It should function in a distributed, infrastructure-independent, and delay-tolerant fashion without sacrificing anonymity of users or leaking information about whom they trust. Various systems in the literature address subsets of these requirements, but Rangzen is unique in that it addresses all of them. Since stringent anonymity constraints pose the main challenges to our design, our primary technical contribution is an algorithm that prioritizes messages in a privacy-preserving and decentralized manner.

6.2 Design Principles

The microblogging network consists of citizens with smartphones. Whenever two citizens encounter, their phones automatically exchange stored messages, leading to epidemic message distribution. This setup facilitates both decentralization and independence of infrastructure, and microblogging is inherently delay tolerant. Additionally, epidemic routing serves to minimize latency in a DTN setting. Thus we are left to address three remaining goals: anonymity, resistance to flooding/misinformation, and scalability. We handle these requirements by means of a prioritization algorithm, which is the core of Rangzen.

The two limited resources most relevant to our system's performance are the storage capacity of every node (flash memory on each smartphone) and the capacity to exchange messages during opportunistic encounters between devices. Storage considerations depend on the type of messages transmitted (e.g. multimedia vs. text), but we expect the storage on modern smart-phones to suffice for supporting a vast microblogging network.[3] However, the typically short opportunistic encounters between mobile devices may prove problematic,[4] and each device's mobility and battery power are likely to additionally restrict the available bandwidth. Therefore, Rangzen emphasizes the transmission and storage of trusted messages, where trust is determined by our prioritization algorithm. Devices transmit messages with high trust rating first upon establishment of an opportunistic connection. These messages are then most likely to propagate through the network. Similarly, messages with low trust ratings are less likely to get transmitted, except when the duration of the opportunistic encounter is long. Least trusted messages would be deleted from a device's message pool first to make space for incoming trusted messages. The point of prioritization is therefore to assign trust ratings in a reliable yet privacy-preserving way.

Message prioritization through trust

In the Rangzen network, pairs of users establish trust relationships, which are intended to reflect real trust between the devices' operators. To accomplish this, establishing trust relations should rely on out-of-band verification. For example, users might need to read each others' screens or establish recognition over voice telephony. In this paper, we assume that trust relationships are symmetric, and both parties must confirm trust to establish a link. Borrowing a term from social networking, we will sometimes refer to devices that trust each other as friends.

The network of devices and their trust relations have an implicit graphical structure that we call the *trust graph*. We refer to nodes separated by a path of length ℓ as ℓ-hop friends. Messages can flow between any two network nodes that opportunistically encounter one another, even if the nodes are not single-hop friends; Rangzen therefore depends on inferred, imperfect trust to prioritize message flow. Our objective is to build trustworthiness scores from devices' trust relations in a way that discounts messages generated by agents and/or associated Sybil identities.

We rely on a key assumption that is often exploited in Sybil defense literature [117, 118]: We assume that agents have difficulty establishing trust with citizens. This limits the number of links on the trust graph between agents and citizens. These links are often referred to as *attack edges* in the trust graph. Even with arbitrarily many Sybil identities in the network, the number of

[3] Modern smart phones have 64GB of storage. If we use half of that for Rangzen and expect an average message size to be 1000 Bytes (for comparison, SMS is limited to 190 Bytes), then each node could store 32 million messages, or 320 messages per user in a network of 100,000 nodes.

[4] The lowest bitrate supported by WiFi is 1 Mbps, which means a theoretic transfer rate of 125 messages per second (for 1000 byte messages). Two smart-phones that are relatively static and in proximity may exceed 50 times that bandwidth.

attack edges is still small compared to the number of citizens, as shown in Figure 6.1. Our design leverages this limited resource to filter messages from attackers.

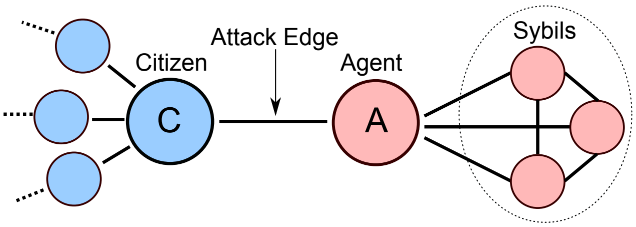

Figure 6.1: Trust graph structure.

To preserve anonymity, Rangzen eliminates the notion of authorship. Without this information, messages are instead prioritized by the trustworthiness of a traversed route.[5] Because of the limited number of attack edges, citizens and agents will have relatively few friends in common. The prioritization algorithm exploits this property by assigning trustworthiness between two nodes proportionally to the number of their shared friends. Intuitively, the more common friends two people have, the more reason they have to trust each other's information. Using friendship to infer trust has been proposed by Trifunovic *et al.* as a defense against Sybil attacks [119]; however, in Rangzen, we extend these ideas to address our strong anonymity requirements.

6.3 Algorithms

Core algorithm

This section introduces the fundamentals of Rangzen's prioritization algorithm. We demonstrate how Rangzen securely computes trustworthiness based on the trust graph. Later in this paper, we expand on this simplified algorithm to improve performance.

Each node in the social network maintains a list of its trusted friends (first-hop neighbors on the trust graph) as well as a set of stored messages. Every message m has an associated priority $p_m \in [0, 1]$, which defines the order in which messages are transmitted during opportunistic encounters. Messages with $p_m = 1$ are transmitted first; messages with low p_m are transmitted last and deleted first if the message cache is full. Each message enters the network with rank $p_m = 1$, since authors value their own messages highly.

Two arbitrary nodes, Alice (A) and Bob (B), have associated trustworthiness scores $T(A, B) \in [0, 1]$ indicating how much Alice trusts Bob, and $T(B, A)$ indicating the opposite. This trustworthiness score is a function of the number of mutual friends between them, and it is used to determine the priority of incoming messages. Whenever a message passes from Bob to Alice, she multiplies the incoming message's priority score by the value of $T(A, B)$; thus if she finds Bob untrustworthy,

[5] We expect to give users the option of signing and endorsing messages, but at its core, Rangzen supports purely anonymous communication.

the incoming message will have low priority. The priority of a message is therefore a nonincreasing function of the number (and the trustworthiness) of links it has traversed. In practice, we want to allow a user to manually upvote a message so that popular messages can propagate quickly.

Let $T_1(A, B)$ denote Alice's trust for Bob using only information about single-hop friends, and let $\mathcal{A}_1$ denote the set of Alice's single-hop friends. We define

$$T_1(A, B) = \max\left(\frac{|\mathcal{A}_1 \cap \mathcal{B}_1|}{|\mathcal{A}_1|}, \epsilon\right) \tag{6.1}$$

where $|\cdot|$ denotes the size of a set and $\epsilon > 0$ is a threshold that ensures nonzero priority even for untrusted messages. Giving a nonzero value to each trustworthiness link preserves the ordering of incoming untrusted messages instead of collapsing them all to priority zero; this allows users to prioritize even among untrusted messages.

Given our assumptions on the trust graph structure, messages from agents are given low priority. Suppose Bob is an agent with a single attack edge to Charlie (C), and Alice is an arbitrary honest citizen. When messages pass from Bob to Alice, we have

$$T_1(A, B) = \begin{cases} 1/|\mathcal{A}_1| & \text{if } C \in \mathcal{A}_1 \\ \epsilon & \text{if } C \notin \mathcal{A}_1 \end{cases}$$

where the latter case is more probable. This example suggests why messages originating from agents are likely to have their priorities downgraded and malicious messages are unlikely to ever dominate an honest device's message pool.

Note that $T_1(\cdot, \cdot)$ is an asymmetric function; if an attacker were to add Alice as his only friend in the network, then he would trust Alice completely, but she would trust him only a little bit because he is one of many friends. Thus the best way for the attacker to fool Alice is by making as many friends as possible. Out-of-band trust validation makes it difficult to establish trust links in Rangzen, so the agent must either truly befriend nodes or coerce their cooperation.[6]

Because leaking private information is a threat to users' personal security, Alice and Bob must compute trustworthiness without revealing details about whom they trust. To this end, Rangzen employs a form of private set intersection (PSI) that allows parties with distinct information sets to learn the number of common elements without revealing either party's information.

We note the scalability of a network using this algorithm. New nodes enter the network organically by establishing trust with other Rangzen devices. Each device can update its list of friends when trust is established, and the network continues to function as expected.

Private Set Intersection-Cardinality

Private set intersection-cardinality (PSI-CA) allows two nodes to learn how many friends are held in common without learning which friends are common. It is conducted whenever two devices opportunistically meet; each node's private data set consists of its own friend list.

[6] We discuss plausible deniability as a means to address this threat in section 6.6.

We assume government agents will misbehave in any way possible to learn information about the other party; PSI-CA cannot provide guarantees against opponents who arbitrarily choose their private friend sets or refuse to participate. In our setup, the adversary (an agent) is trying to earn the trust of citizen nodes, thereby disincentivizing non-participation. On the other hand, the agent has an incentive to falsely augment the size of his friend list to appear more trustworthy. It is therefore difficult by design to learn the "friend IDs" of nodes without actually being friends; these IDs are protected by the PSI-CA protocol and by the lack of authorship tracking.

Suppose a node somehow obtains a large set of friends. We institute an upper limit F on the number of friends that can be compared in a single PSI-CA exchange to prevent the node from appearing trustworthy to everyone. This protects against both agents who coerce many trust links and citizens with lax trust standards. Equivalently, we posit that legitimate citizens will have at most F friends for some $F > 0$. If a party submits more than F elements to the comparison, that party is automatically mistrusted. This forces people with excessive numbers of friends to select only a subset thereof for the comparison. Similarly, if a node has few friends, it pads its list with randomly drawn values to obtain a list of length F. The probability of a randomly chosen filler ID coinciding with another node's friend set is low, and fixing F prevents adversaries from learning the size of a friend set during an encounter.

There are several PSI-CA algorithms in the literature, including [120, 121, 122]. To the best of our knowledge, only [121] addresses malicious adversaries, and our approach is similar to their Cardinality-Mal protocol. It relies on concepts such as homomorphic evaluation of polynomials and zero-knowledge proofs, which we will cover briefly.

An *additively homomorphic cryptosystem* has the property that

$$\mathcal{E}(a+b) = \mathcal{E}(a) \cdot \mathcal{E}(b) \tag{6.2}$$

where $\mathcal{E}(\cdot)$ denotes encryption using said cryptosystem. This key property implies that for any constant c,

$$\mathcal{E}(ca) = \mathcal{E}(a)^c. \tag{6.3}$$

Additively homomorphic cryptosystems are commonly used in private set intersection algorithms because they facilitate computation in the encrypted domain. In particular, they make it easy to evaluate polynomials in the encrypted domain given the encrypted polynomial coefficients. In this PSI-CA protocol, we will utilize an important example of such a cryptosystem called the Paillier cryptosystem [123].

We use zero-knowledge proofs and privacy-preserving protocols to deal with malicious adversaries. A zero-knowledge proof is a method for proving that a party knows a secret without revealing the secret to the verifying party. Efficient implementations for the Paillier cryptosystem rely on proving knowledge of discrete logarithms [124]. We will utilize two such functions. Proof of plaintext knowledge (PK$\{v \mid \mathcal{E}(v)\}$) shows that the prover knows the plaintext identity r given that the encryption $\mathcal{E}(v)$ that is visible to the verifier [125]. Proof of correct polynomial evaluation (PE$\{(v,r) \wedge r \neq 0 \mid \mathcal{E}(r \cdot f(v))\}$) shows that the prover knows the plaintext values v and r and $r \neq 0$, given the encrypted polynomial f and the encrypted evaluation of the polynomial $\mathcal{E}(r \cdot f(v))$ [125, 126].

Protocol Description Suppose Alice and Bob each possesses a set of friend keys, denoted $\mathcal{A} = \{a_1, ..., a_{|\mathcal{A}|}\}$ and $\mathcal{B} = \{b_1, ..., b_{|\mathcal{B}|}\}$ respectively. Let a_i denote the ith element of set $\mathcal{A}$. The algorithm consists of two rounds of PSI-CA. In the first round, Alice learns the number of shared friends, and in the second, Bob does. Since the iterations are identical except with switched roles, we will only explain the case in which Alice is trying to learn the number of common friends. The steps are as follows:

1. Alice performs the following:

   a) She generates a secret-key/public-key pair for the homomorphic encryption scheme.

   b) She generates a polynomial $f_{\mathcal{A}}(x)$, with the elements of $\mathcal{A}$ as roots:

   $$\begin{aligned} f_{\mathcal{A}}(x) &= \prod_{k \in \{1,2,...,|\mathcal{A}|\}} (x - a_k) \\ &= \eta_0 + \eta_1 x + \ldots + \eta_F x^F \end{aligned}$$

   The degree of $f_{\mathcal{A}}$ is F because there are exactly F elements in each private set by construction.

   c) She sends the encryption of each coefficient of $f_{\mathcal{A}}$ (except η_F) to Bob, along with proof of plaintext knowledge (PK$\{\eta_i \mid \mathcal{E}(\eta_i)\}$) for each coefficient. η_F is always assumed to equal 1.

2. Bob executes the following:

   a) Using homomorphic encryption properties, he evaluates the polynomial F times: once for each of his entries, giving $\mathcal{E}(f_{\mathcal{A}}(b_i))$.

   b) For each set element b_i, he multiplies the encrypted evaluation of $f_{\mathcal{A}}(b_i)$ by a distinct, randomly drawn number r_i, giving $\mathcal{E}(r_i \cdot f_{\mathcal{A}}(b_i))$.

   c) He generates proof of correct polynomial evaluation PE$\{(b_i, r_i) \wedge r_i \neq 0 \mid \mathcal{E}(r_i \cdot f_{\mathcal{A}}(b_i))\}$.

   d) He returns the F randomized polynomial evaluations and proofs of correct construction to Alice.

3. Alice decrypts the F polynomial evaluations. The number of zeros is the number of common elements.

If both parties execute this procedure, each will obtain the number of common elements. In 2a, if B_i is a shared friend, the polynomial evaluates to $\mathcal{E}(0)$; otherwise it evaluates to the encryption of some nonzero value. Bob cannot determine if the result is an encryption of zero because he lacks the private key. Since the Paillier cryptosystem is randomized, multiple instances of $\mathcal{E}(0)$ will look different with high probability. Finally, step 2b prevents Alice from learning about Bob's friends. If the argument was initially a zero, it will remain a zero, indicating a mutual friend. Otherwise, the argument is scaled by the random quantity r_i. This scaling prevents Alice from solving the polynomial for $f_A(b_i)$.

The scheme requires three total rounds of communication: one transmission in which Alice sends her encrypted polynomial, one in which Bob returns the evaluated polynomials as well as his own encrypted polynomial, and one final transmissions as Alice returns her evaluations of Bob's polynomial. Each transmission is $O(F)$ in size.

Security and Correctness The proof of security for this scheme is analogous to that of Cardinality-Mal in [121]; it shows that for each participant in this scheme, there exists a participant G in the ideal model such that the views of the participants in the real and ideal models are indistinguishable. However, the danger in this scenario does not stem from the security of the scheme, since the two parties only transmit semantically secure encryptions of their data. Instead, we must ensure that the adversary cannot impact the correctness of the scheme.

The provided protocol protects against two types of misbehavior: Alice encrypting f_A improperly and Bob evaluating Alice's polynomial improperly. Step 1c forces Alice to set $\eta_F = 1$ to prevent her from misrepresenting her polynomial as $f_A(x) = 0$, which has every number as a root. In step 1c, inability to provide such a proof of knowledge corresponds to faking a friend set by using an incorrect polynomial f_A, from a previous encounter for instance. Faking a friend set without knowing the underlying IDs will only affect how much agent trusts the citizen, not the other way around; however, it does allow the agent to learn about trust relationships in the network. Bob is forced to evaluate the polynomial properly by providing proof of plaintext knowledge of r_i and b_i, and by ensuring that $r_i \neq 0$, which would always result in an encryption of zero. The zero-knowledge proofs therefore ensure correct execution of the protocol.

Multi-hop extension

In the basic form of our algorithm, each node only knows the identities of its own trusted friends (first degree neighbors on the trust graph). This high-privacy setting is good for protecting the trust graph, but it also diminishes the receiver's ability to prioritize messages relayed via distant nodes. Such nodes may be agents or honest citizens, but with no knowledge of the trust graph, the receiver cannot make such a distinction.

We address this issue by allowing each node to maintain a 'sketch' of the local trust graph. This sketch manifests itself as a list of all the friends within an ℓ-hop neighborhood on the trust graph, for some $\ell \geq 1$. In this case, trustworthiness becomes a function of multiple hops of friendship.

The number of hops in this neighborhood, ℓ, is a function of both the privacy desired by an individual user as well as the overarching anonymity settings in the deployment environment; a larger neighborhood gives a better estimate of incoming message reliability at the risk of reduced privacy. The size of a local neighborhood should be upper bounded by some hard threshold to preserve a certain minimum level of privacy (see section 6.3 for a discussion on how to decide and set such system-wide parameters).

If we store $\ell > 1$ hops of friends in the trust graph, then the greater the separation between two nodes on the trust graph, the less they should trust one another. Therefore, each node stores a list of IDs for each friend in the local trust graph, and we call this set a *neighborhood*. Concretely, Bob's neighborhood ID for a friend named Alice who is located i hops away corresponds to a cryptographic hash of the tuple (Alice, i). This ID, denoted $k(\text{Alice}, i)$, is generated by Alice, and it cannot be used to identify Alice without access to the cryptographic key she used to generate it. Bob will not store only $k(\text{Alice}, i)$, but also all tuples from i to ℓ, which we will refer to as the *sketch set* for (Alice, i):

$$S(\text{Alice}, i) = \{k(\text{Alice}, i), k(\text{Alice}, i+1), \ldots \\ \ldots, k(\text{Alice}, \ell)\}.$$

Note that every instance of $k(\text{Alice}, i)$ is identical, regardless of who possesses it, so Alice need only generate her full sketch set $S(\text{Alice}, 1)$ once upon joining the network. From a privacy standpoint, sketch sets do not provide complete information about a trust graph neighborhood because the trust relationships between neighbors are lost. Additionally, an attacker viewing a neighborhood set has no way of understanding which IDs belong to the same sketch set unless the node has only a single friend.

Link trustworthiness was previously a function of the number of mutual friends between two entities. Since neighborhoods describe classes of friends defined by separation distance on the trust graph, different classes should be weighted differently. The heuristic we consider is a weighted sum of the proportion of common elements in each class. Instead of computing $T_1(A, B)$ as before, we consider a weighted sum of $T_j(A, B)$, which finds the intersection proportion from the jth ring of friends, for $j \leq \ell$:

$$T_j(A, B) = \max\left(\frac{|\mathcal{A}_j \cap \mathcal{B}_j|}{|\mathcal{A}_j|}, \epsilon\right) \tag{6.4}$$

where ϵ is the same as in equation 6.1 and $\mathcal{A}_j$ is the set of Alice's j-hop friends. Therefore, the reliability of an edge from Alice to Bob is determined as

$$T(A, B) = \min\left(\sum_{j=1}^{\ell} \alpha_j \cdot T_j(A, B), 1\right) \tag{6.5}$$

where α is a system-wide vector of parameters that weights the different levels of separation in the trust graph. Therefore, $\alpha_j \in [0, 1]$ and $\sum_j \alpha_j = 1$; also, $\alpha_i > \alpha_j \;\; \forall i > j$, since closer friends should count more than distant friends. The two parties will execute ℓ rounds of PSI-CA in total.

This formulation clarifies why each node should store a sketch set for each member of its neighborhood: Not storing the sketch set would result in distant nodes on the graph appearing as attackers. For example, consider the scenario in Figure 6.2, with $\ell = 2$. Suppose Charlie is a 2nd-hop friend of Alice's and a 1st-hop friend of Bob's. Then if Alice and Bob were only storing $k(\text{Charlie}, 2)$ and $k(\text{Charlie}, 1)$ respectively, Charlie would not show up as a common friend. Therefore, storing the sketch sets of these respective keys allows two nodes to find common elements among their friend sets with some (slightly skewed) perception of separation distance, even if the generating keys are not the same.

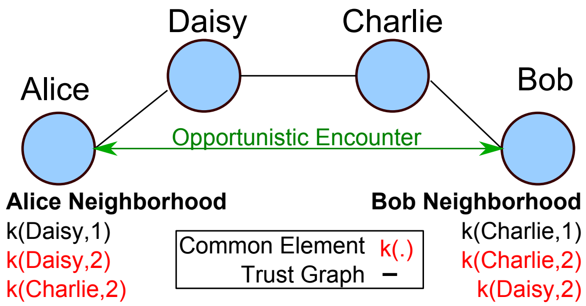

Figure 6.2: Sample encounter. Alice and Bob meet each other, as indicated by the green edge. After calculating their 2nd degree common friends, they each see that they have two 2nd-hop common neighbors.

Before Internet Blackout

During this phase, the initial social network of Rangzen is established. While Rangzen is designed to operate in the absence of communication infrastructure, we leverage the pre-blackout phase to speed up the creation of the network and the distribution of the Rangzen software application with the use of a central server. Through Internet access prior to the disconnection event, we overcome the delays introduced by the opportunistic DTN and allow for real-time propagation of messages. The server also selects some preliminary parameters, such as the default neighborhood size ℓ and the multi-hop weight parameters α. We envision the server being located outside the target country, and consequently beyond the control of the malicious government. The government may fully or partially block access to such a server, in which case blocked nodes can simply assume the disconnected phase has begun. Censored access could also be circumvented using traditional overlay networks or proxies.

Users join the Rangzen social network by generating a full sketch set and establishing friendships with people already in the network. Explicitly, for each new user v, the central server distributes appropriate subsets of v's full sketch set to every member of v's neighborhood. For instance, suppose Bob joins the network by becoming friends with Alice. The server will start by transferring Bob's full sketch set to Alice and vice versa. Let $\mathcal{A}_i^j$ denote the set of all Alice's friends that are between i and j hops away, inclusive. Then for each of Alice's friends $f_j \in \mathcal{A}_1^\ell$, if f_j is located $i < \ell$ hops away from Alice on the trust graph, the server will give to Bob the sketch set $S(f_j, i+1)$, and to f_j the sketch set $S(\text{Bob}, i+1)$. In doing so, the server gives Bob a sketch of his multi-hop neighborhood, and also updates the friend lists of everyone in Bob's neighborhood.

In terms of message dissemination, the central server will handle the entire prioritization pipeline. Having a global view of the trust graph, the server can avoid conducting a private set intersection for every pair of nodes.

After Internet Blackout

After the Internet blackout occurs, prioritization relies entirely on the private set intersection computations described earlier.

One of the difficult parts of the offline phase is scaling up the network. That is, if a new member joins the network, how does the system update the appropriate nodes' sketches? Due to reduced connectivity, the server can no longer take care of updating all the appropriate neighborhood lists, therefore all neighborhoods must be transmitted from device to device. Suppose that Bob joins the network by becoming friends with Alice. Bob starts by transferring his full sketch set to Alice and vice versa. Then for each of Alice's friends $f_j \in \mathcal{A}_1^{\ell-1}$, Alice will give to Bob the sketch set $S(f_j, i+1)$, which she possesses by construction. In doing this, Alice gives Bob his full multi-hop neighborhood (assuming that Alice knows her full neighborhood). The main imbalance in this scenario is that Alice cannot inform her neighborhood of the new addition, because the central server is not accessible and she is presumably not within transmitting distance of everyone in her neighborhood. Thus we wait for an opportunistic encounter. The next time Alice comes into contact with one of her neighbors f_j that is strictly fewer than ℓ hops away, she will transmit an appropriate subset of Bob's sketch set. This effectively informs f_j that Bob is now part of f_j's extended neighborhood. Note that f_j cannot deduce the identity of Bob from the received sketch set; the sketch set received by f_j will never be inserted into messages in any way—it will only be used in the context of the private set intersection protocol, which makes it difficult to correlate hashed keys with real identities.

6.4 Evaluation

In evaluating our prioritization algorithm we emphasize two key properties: malicious message infiltration and degree of message diffusion. Obtaining realistic datasets for evaluation of Rangzen is challenging because we need information on human mobility as well as interpersonal trust. Data from typical social networks has little relevance to our scenario of strict trust relations. Similarly, we expect common mobility traces, such as of vehicles or students on a university campus, to poorly represent our use cases. Finally, since trust relations and human mobility are correlated, it is unrealistic to model Rangzen by mapping unrelated social networks and mobility traces. For these reasons, we do not validate Rangzen using real datasets, but instead develop synthetic, conservative datasets to give a worst-case notion of system performance.

Simulated environment

To synthesize a social graph, we build a small-world network according to the construction by Watts and Strogatz, which relies on adding random edges to circulant graphs [127]. Small-world networks exhibit similar properties to social networks—namely, short average path length between nodes and high local clustering [127]. Our social graph is further augmented by adding a small number of agent nodes and adding edges from each agent node to uniformly-selected-at-random citizen nodes in the graph. We previously assumed the number of agents would be 0.1 percent of the number of citizens; to give a conservative estimate of system performance, we increase this proportion by an order of magnitude in simulation and set the number of agents to be two percent of the number of citizens. An instance of this graph construction with 2^7 citizen nodes is shown in Figure 6.3; the outer ring of red squares represents agent nodes, while the inner ring of blue circles consists of all the citizen nodes. Edges represent trust relationships. Although this graph size is unrealistic, it serves to demonstrate trends in the system.

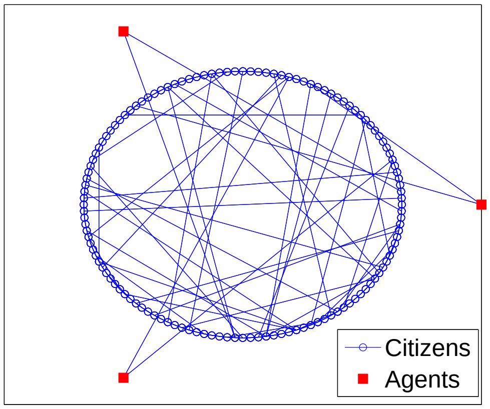

Figure 6.3: Sample social network graph, with 2^7 citizen nodes, and 3 agent nodes. The inner ring represents citizens and their trust relationships, while the outer nodes represent agents.

With regards to connectivity, we assume that any node in the network will meet some other node in a given time interval with a fixed probability. This encounter rate is higher for agents, as they will attempt to use the infrastructure under their control to impersonate citizens nodes. This Bernoulli random process model of encounters is a discrete approximation of a Poisson arrival process, which is often used to model memoryless random processes like human arrivals [128]. This connectivity model lacks a number of real-life dependencies, including time and location. However, by uniformly pairing nodes for opportunistic encounters, the model disproportionately favors encounters with untrusted nodes; this slows the propagation of messages and gives a conservative estimate of communication performance. We also assume that an agent's cache is always full of agent messages with priority 1. In contrast, each honest citizen will generate a new message at a given timestep with a small probability. This corresponds to authoring a new message or upvoting an existing one, with the end result that the node's cache contains an honest message of priority 1.

Malicious message infiltration

In measuring malicious message infiltration, we cannot completely eradicate malicious messages from citizens' caches, since the agents' caches are constantly filled with high priority malicious messages. The important notion is that malicious messages should be concentrated at the bottom of a citizen's message pool. Therefore, the first positions of a message pool should be full of honest messages, while the last positions (i.e. the lower priority ones) do not matter. To measure this, we simulated system operation over a number of time steps, and look at the average proportion of honest messages in each cache position. Cache index 1 is the most trusted, so we would like a high proportion of trusted messages at low cache indices. The result of this simulation for the single-hop algorithm at different time iterations is shown in Figure 6.4. As desired, the lower indices contain more honest messages on average, while the proportion of nodes with honest messages decreases as the priority decreases. Moreover, as time progresses, the proportion of honest messages in citizens' message pools throughout the cache converges to the shape at $t = 80$ iterations in Figure 6.4, modulo the randomness in the system.

This figure gives a conservative picture, since it only captures the ordering of messages. Due to the prioritization protocol, malicious messages have a lower priority score on average than honest messages; in our simulation after stabilization, malicious messages in citizens' caches had an average reliability of 0.080, while citizens' messages have an average reliability of 0.17; while the latter number may seem small, note that this includes messages coming from completely unknown nodes in the graph, which are mistrusted by the prioritization algorithm as much as messages from malicious nodes.

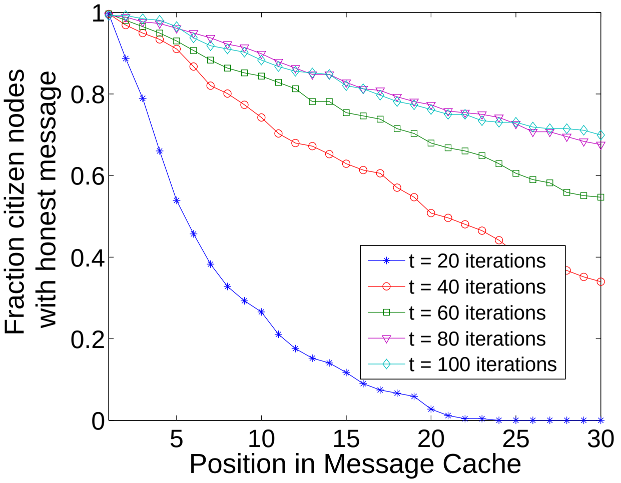

Figure 6.4: Average proportion of honest messages in citizens' message pools as a function of priority in the pool. Lower indices indicate higher priority. Values are averaged over 2^7 citizen nodes.

The shapes of the equilibrium curves in Figure 6.4 are dependent on a variety of parameters—one of the most important of these parameters is the agent encounter rate. We assume that agents will transmit messages at a higher rate than ordinary citizens to maximize the number of honest citizens reached. Figure 6.5 shows the equilibrium curves for various agent encounter rates ranging from 0.1 (same as citizen encounter rate) to 1.0 (constantly exchanging messages). As expected, this shows that as the agent encounter rate grows, the proportion of agent messages in citizens' message pools increases significantly. However, even in the worst case scenario of agents successfully transmitting messages all the time, the first indices in the cache are still primarily occupied by honest nodes.

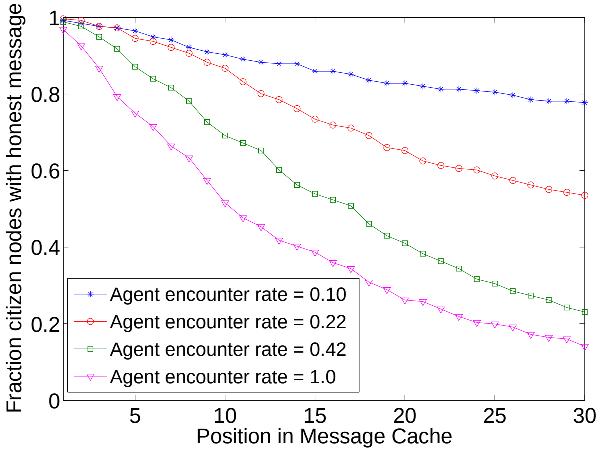

Figure 6.5: Average proportion of honest messages in citizens' message pools, indexed by priority in the pool. Lower indices indicate higher priority. Values are averaged over 2^7 citizen nodes.

Message diffusion

An equally important performance aspect is the degree to which messages are able to spread in the system. In practice, this will depend primarily on human mobility patterns, which we do not know. However, we can lower bound message spread using our pessimistic mobility model of uniform encounters. We simulate the diffusion of a single message, assuming no upvotes. We also assume the author posts the message in a highly trafficked area such as a shopping mall to encourage maximum dispersion; this is realistic if an individual wishes to reach many people.

With synthetic data, simulated message diffusion times mean little; however, we can observe the effect of the multi-hop extension on trust levels in the network. In some sense this is more fundamental than observed diffusion times because it is independent of mobility models. The trustworthiness function allows us to upper bound the priority of a received message, which determines how far the message can reach.

In Figure 6.6, we show the mean priority of the received message as a function of the receiver's distance (in number of hops) from the author. These curves confirm the algorithm intuition that using larger friendship neighborhoods enables greater message spread. Moreover, the parameter vector α, which determines the weight of various friendship levels in the trustworthiness function, allows us to shape the curves in Figure 6.6 as desired.

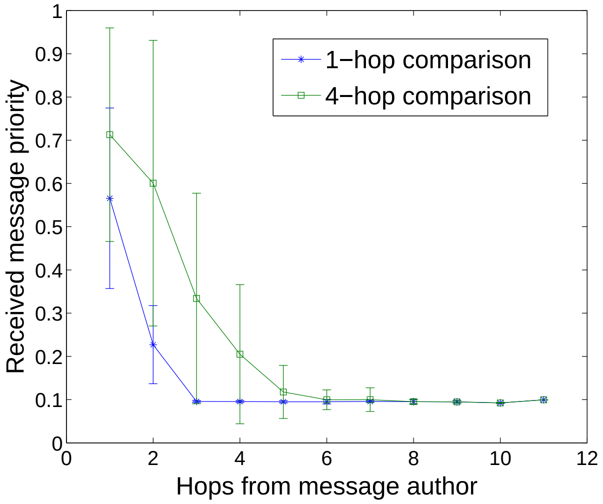

Figure 6.6: Incoming message priority as a function of separation distance (in hops on the trust graph) from the author.

6.5 Related work

Inspired by SumUp [129], GateKeeper [130], Sybilinfer [131] and Sybillimit [118], we base our defenses on the properties of the social network. The number of social links connecting agents to citizens is limited by the number of real friends the agents have—the *attack edges*. Moreover, since the citizens in Rangzen only create social links with trusted friends, the agents cannot create attack edges arbitrarily, resulting in a small number of attack edges compared to the number of citizens. A fundamental departure from these studies is our ability to perform similar defenses without leaking information about the social trust graph, thanks to privacy-preserving set intersections [121, 122].

Although security issues in mobile ad-hoc networks (MANETS) have been studied extensively, anonymity concerns remain relatively unexplored. Exceptions include [132], which quantifies common interests between MANET nodes in a privacy-preserving way, and [133], which emphasizes the anti-localization of authors in a communication network. Despite high-level similarities, [132] focuses on designing primitives for multi-party interest-casting, while we focus instead on inferring trust and resisting network attacks. Regarding [133], we are less concerned about localization since our strong anonymity properties make attribution of messages to authors difficult.

In this study, we address primarily the challenges stemming from our unique threat model; as such, we leave underlying communication foundation issues for future work. Our design is nevertheless shaped by the fundamental constraints of opportunistic and delay-tolerant networking. These considerations were motivated in part by recent work on neighbor discovery in mobile ad-hoc networks [134, 135, 136] as well as canonical DTN literature [137]. Naturally, we are also inspired by recent studies to combat information censorship [138, 139].

6.6 Future Work

Implementing a viable Rangzen network is a complex challenge. In this work, we focus on the least explored problem: anonymity-preserving, social graph-based message prioritization. However, there remain many obstacles to be resolved.

User security: While our core algorithm is privacy-preserving, leakage of trust graph information could occur in the multi-hop scenario. In particular, if the government were to forcibly access the devices of many citizens, it could learn each device's neighborhood. Given enough devices in addition to correlation with external information, the government could make inferences about the trust graph structure. Of course, the smaller the neighborhood size (in hops), the more phones required to rebuild the graph. The degree of anonymity reduction in such a scenario needs to be further evaluated beyond our initial exploration in this study.

On a related note, plausible deniability is an important property for our network. We wish to provide a mechanism by which citizens can safely signal to the network that an establishment of trust has been made under duress. Traditionally, this is done by entering a special password different from the user's regular one, indicating that the user is being forced. Preventing agents from detecting the activation of such a hidden mechanism is challenging.

Resource awareness: Since the algorithm is run on mobile devices, it is important to consider resource costs. We anticipate that certain aspects of the algorithm, including neighbor discovery and duplicate message transmission, will prove particularly costly.

During network operation, devices must automatically and efficiently detect the presence of physically close nodes (and communicate with the central server in the pre-blackout phase). Modern smartphones support many modes of communication for doing so, including access to cellular infrastructure, WiFi, Bluetooth, and even physical transportation of memory cards. A multimodal connectivity-seeking networking layer should be designed to gracefully degrade across these modes of connectivity. The Hercules project addresses some of these issues [140], but it focuses on real-time modes while neglecting opportunistic DTN and real-time ad-hoc modes. These have been studied in depth elsewhere but also require heavy alterations for use in Rangzen [134, 136]. We also anticipate popular messages being circulated widely, causing redundant message transmissions. To save bandwidth during opportunistic encounters, nodes need only incrementally replicate messages that are already stored by the receiving node. Studies such as TIERstore [141] and Haggle [142] address these issues in similar environments.

Alternate use cases: We would like to explore uses for Rangzen beyond the specific application presented in this paper. In the microblogging application space, users might wish to distribute multimedia as well as text, which could present additional resource allocation questions. Additionally, for the transmission of confidential messages between friends, we envision the need for message encryption. Public keys could be exchanged during establishment of trust, which could then be used to encrypt messages intended for the corresponding friend; exchanging public keys also implies the ability to sign and authenticate tamper-proof messages. Communication between nodes that do not trust one another is more challenging because it requires key distribution over a DTN mesh without a centrally trusted certificate authority.

More broadly, we wish to address non-dissent use cases. This is important for two reasons: It would provide a cover story for the application, preventing it from being outlawed, and it would encourage the general public to download and use the application. The latter reason is important for studying system functionality at scale. We envision disaster preparedness as such a plausible use case. While the anonymity guarantees and attack resiliency properties of Rangzen might be less important for such scenarios (at least when natural disasters are the concern), the robust, opportunistic distribution qualities over a DTN are highly attractive.

Usability: Care should be taken to ensure Rangzen is user-friendly. It is particularly important for the system to guide users in making informed decisions that may affect their security. Open questions include how to establish trust relations among peers, especially when they are not physically close.

6.7 Conclusions

Dissent networking is a relatively unexplored territory that presents extreme challenges. Designing a communication network for citizens in the face of an adversarial government is a major undertaking that can lead to devastating consequences if done poorly. We have presented Rangzen: an anonymity-preserving microblogging tool designed for circumvention of government-imposed communication blackouts and censorship. Our goal was to present a decentralized, delay-tolerant communication system that is both resilient to network attacks and anonymity-preserving for users. We addressed this problem by designing an algorithm that exploits social graph structure and privacy-preserving set intersections to prioritize messages. We have simulated this algorithm on synthetic data and found that on average, it filters out malicious messages so that users see primarily honest messages in the top slots of their message caches.

Bibliography

[1] Shaddi Hassan, Yahel Ben-David, Max Bittman, and Barath Raghavan. "The Challenges of Scaling WISPs." In: *The sixth annual Symposium on Computing for Development (ACM DEV 2015)*. 2015.

[2] Matthew Luckie, Amogh Dhamdhere, David Clark, Bradley Huffaker, and kc claffy. "Challenges in Inferring Internet Interdomain Congestion". In: *Proceedings of the 2014 Conference on Internet Measurement Conference*. IMC '14. Vancouver, BC, Canada: ACM, 2014, pp. 15–22. ISBN: 978-1-4503-3213-2. DOI: 10.1145/2663716.2663741. URL: HTTP://DOI.ACM.ORG/10.1145/2663716.2663741.

[3] Qi Liao, Zhen Li, and A. Striegel. "Is more P2P always bad for ISPs? An analysis of P2P and ISP business models". In: *Computer Communication and Networks (ICCCN), 2014 23rd International Conference on*. 2014, pp. 1–6. DOI: 10.1109/ICCCN.2014.6911841.

[4] Robert Faris, Hal Roberts, Bruce Etling, Dalia Othman, and Yochai Benkler. "Score Another One for the Internet? The Role of the Networked Public Sphere in the US Net Neutrality Policy Debate". In: *Berkman Center Research Publication* 2015-4 (2015).

[5] *Ubiquiti Networks*. HTTP://WWW.UBNT.COM/.

[6] Yahel Ben-David, Matthias Vallentin, Seth Fowler, and Eric Brewer. "JaldiMAC: taking the distance further". In: *Proceedings of the ACM NSDR Workshop*. 2010.

[7] Gregory A. Jirak. *The Esplanade Lesson: High Speed Internet Access Essential not Optional*. 2011.

[8] Florian Fainelli. "The OpenWrt embedded development framework". In: *Proceedings of the Free and Open Source Software Developers European Meeting*. 2008.

[9] *POX, Python-based OpenFlow Controller*. HTTP://WWW.NOXREPO.ORG/POX/ABOUT-POX/.

[10] *Stripe*. HTTPS://STRIPE.COM/.

[11] E. Brewer et al. "The case for technology in developing regions". In: *IEEE Computer* 38.6 (2005).

[12] François Bar and Hernan Galperin. "Geeks, Cowboys, and Bureaucrats: Deploying Broadband, the Wireless Way". In: *African Journal of Information and Communication* 6 (2005), pp. 48–63.

[13] Kameswari Chebrolu, Bhaskaran Raman, and Sayandeep Sen. "Long-Distance 802.11b Links: Performance Measurements and Experience". In: *Proceedings of ACM MOBICOM*. 2006.

[14] Rabin K Patra, Sergiu Nedevschi, Sonesh Surana, Anmol Sheth, Lakshminarayanan Subramanian, and Eric A Brewer. "WiLDNet: Design and Implementation of High Performance WiFi Based Long Distance Networks". In: *Proceedings of USENIX/ACM NSDI*. 2007.

[15] Lynne Salameh, Astrit Zhushi, Mark Handley, Kyle Jamieson, and Brad Karp. "HACK: hierarchical ACKs for efficient wireless medium utilization". In: *Proceedings of USENIX ATC*. 2014.

[16] Sonesh Surana, Rabin K Patra, Sergiu Nedevschi, Manuel Ramos, Lakshminarayanan Subramanian, Yahel Ben-David, and Eric A Brewer. "Beyond Pilots: Keeping Rural Wireless Networks Alive". In: *Proceedings of USENIX/ACM NSDI*. 2008.

[17] Carlos Rey-Moreno, Zukile Roro, William D Tucker, Masbulele Jay Siya, Nicola J Bidwell, and Javier Simo-Reigadas. "Experiences, Challenges and Lessons from Rolling out a Rural WiFi Mesh Network". In: *Proceedings of the 3rd ACM Symposium on Computing for Development*. ACM. 2013, p. 11.

[18] Vijay Gabale, Rupesh Mehta, Jeet Patani, K Ramakrishnan, and Bhaskaran Raman. "Deployments Made Easy: Essentials of Managing a (Rural) Wireless Mesh Network". In: *Proceedings of the 3rd ACM Symposium on Computing for Development*. ACM. 2013, p. 10.

[19] *Heywhatsthat*. HTTP://HEYWHATSTHAT.COM/.

[20] *Powercode*. HTTP://POWERCODE.COM/.

[21] *Swiftfox*. HTTP://WWW.SWIFTFOX.NET/.

[22] *Azotel Technologies*. HTTP://WWW.AZOTEL.COM/.

[23] *International Telecommunications Union - Statistics*. http://www.itu.int/ITU-D/ict/statistics/.

[24] *Population Reference Bureau - World Population Datasheet*. 2011.

[25] Yael Valerie Perez and Yahel Ben-David. "Internet as Freedom - Does the Internet Enhance the Freedoms People Enjoy?" In: *Information Technology for Development* (2012), pp. 1–18. DOI: 10.1080/02681102.2011.643203. eprint: HTTP://WWW.TANDFONLINE.COM/DOI/PDF/10.1080/02681102.2011.643203. URL: HTTP://WWW.TANDFONLINE.COM/DOI/ABS/10.1080/02681102.2011.643203.

[26] Joyojeet Pal, Sergiu Nedevschi, Rabin Patra, and Eric Brewer. "A Multidisciplinary Approach to Open Access Village Telecenter Initiatives: the case of Akshaya". In: *E-Learning* 3.3 (Sept. 2006), pp. 291–316.

[27] Y. Ben-David, M. Vallentin, S. Fowler, and E. Brewer. "JaldiMAC: taking the distance further". In: *Proceedings of the 4th ACM Workshop on Networked Systems for Developing Regions*. ACM. 2010, p. 2.

[28] *Nagios Network Monitor*. http://www.nagios.org.

[29] `ntop` *Traffic Monitor*. http://www.ntop.org.

[30] Martín Casado, Teemu Koponen, Rajiv Ramanathan, and Scott Shenker. "Virtualizing the network forwarding plane". In: *Proceedings of the Workshop on Programmable Routers for Extensible Services of Tomorrow*. PRESTO '10. Philadelphia, Pennsylvania: ACM, 2010, 8:1–8:6. ISBN: 978-1-4503-0467-2. DOI: 10.1145/1921151.1921162. URL: HTTP://DOI.ACM.ORG/10.1145/1921151.1921162.

[31] D. Collins. *Portfolios of the poor: how the world's poor live on $2 a day*. Princeton University Press, 2009.

[32] *M-PESA*.

[33] N. Hughes and S. Lonie. "M-PESA: mobile money for the "unbanked" turning cellphones into 24-hour tellers in Kenya". In: *Innovations: Technology, Governance, Globalization* 2.1-2 (2007), pp. 63–81.

[34] T. Koponen et al. "Onix: A distributed control platform for large-scale production networks". In: *OSDI, Oct* (2010).

[35] *Capgemini - Mobile Tower Sharing and Outsourcing: Benefits and Challenges for Developing Market Operators*. 2009.

[36] *MultiRAN Virtual Base Station*. 2009.

[37] Kok-Kiong Yap, Rob Sherwood, Masayoshi Kobayashi, Te-Yuan Huang, Michael Chan, Nikhil Handigol, Nick McKeown, and Guru Parulkar. "Blueprint for introducing innovation into wireless mobile networks". In: *Proceedings of the second ACM SIGCOMM workshop on Virtualized infrastructure systems and architectures*. VISA '10. New Delhi, India: ACM, 2010, pp. 25–32. ISBN: 978-1-4503-0199-2. DOI: 10.1145/1851399.1851404. URL: HTTP://DOI.ACM.ORG/10.1145/1851399.1851404.

[38] Peter Dely, Andreas Kassler, and Nico Bayer. "OpenFlow for Wireless Mesh Networks". In: *WiMAN 2011*. 2010.

[39] Sonesh Surana, Rabin Patra, Sergiu Nedevschi, Manuel Ramos, Lakshminarayanan Subramanian, Yahel Ben-David, and Eric Brewer. "Beyond pilots: keeping rural wireless networks alive". In: *NSDI'08: Proceedings of the 5th USENIX Symposium on Networked Systems Design and Implementation*. San Francisco, California: USENIX Association, 2008, pp. 119–132. ISBN: 111-999-5555-22-1. URL: HTTP://PORTAL.ACM.ORG/CITATION.CFM?ID=1387598.

[40] Sonesh Surana, Rabin Patra, and Eric Brewer. "Simplifying Fault Diagnosis in Locally Managed Rural WiFi Networks". In: *ACM SIGCOMM Workshop on Networked Systems for Developing Regions (NSDR)*. 2007.

[41] *Meraki*. HTTP://MERAKI.CISCO.COM/.

[42] *Norton Cybercrime Report: The Human Impact*. 2010.

[43] *M86 security Labs, Spam source by country*. "http://www.m86security.com/labs/spam_statistics.asp". 2011.

[44] *Project Honey Pot Statistics*. "http://projecthoneypot.org/statistics.php". 2010.

[45] *AirJaldi.Org - Wireless Network, Dharamsala, India*. http://www.airjaldi.org.

[46] Gregor Maier, Anja Feldmann, Vern Paxson, Robin Sommer, and Matthias Vallentin. "An Assessment of Overt Malicious Activity Manifest in Residential Networks". In: *DIMVA'11: Proceedings of the Conference on Detection of Intrusions and Malware and Vulnerability Assessment*. to appear. Springer, 2011.

[47] B. Krebs. *"Stuxnet" Worm far more sophisticated than previously thought*. 2010. URL: HTTP://KREBSONSECURITY.COM/2010/09/STUXNET-WORM-FAR-MORE-SOPHISTICATED-THAN-PREVIOUSLY-THOUGHT/..

[48] Saurabh Panjwani. "Towards End-to-End Security in Branchless Banking". In: *Workshop on Mobile Computing Systems and Applications (HotMobile)*. ACM, Mar. 2011.

[49] Financial Access Initiative. *M-Kesho in Kenya: A new step for M-Pesa and mobile banking*. HTTP://FINANCIALACCESS.ORG/NODE/2968. May 2010.

[50] Michael Paik. "Stragglers of the herd get eaten: security concerns for GSM mobile banking applications". In: *Proceedings of the Eleventh Workshop on Mobile Computing Systems &#38; Applications*. HotMobile '10. Annapolis, Maryland: ACM, 2010, pp. 54–59. ISBN: 978-1-4503-0005-6. DOI: HTTP://DOI.ACM.ORG/10.1145/1734583.1734597. URL: HTTP://DOI.ACM.ORG/10.1145/1734583.1734597.

[51] Saurabh Panjwani and Edward Cutrell. "Usably secure, low-cost authentication for mobile banking". In: *Proceedings of the Sixth Symposium on Usable Privacy and Security*. SOUPS '10. Redmond, Washington: ACM, 2010, 4:1–4:12. ISBN: 978-1-4503-0264-7. DOI: HTTP://DOI.ACM.ORG/10.1145/1837110.1837116. URL: HTTP://DOI.ACM.ORG/10.1145/1837110.1837116.

[52] *Security Breach at M-Pesa*. HTTP://WWW.TELCO2.NET/BLOG/2010/02/SECURITY_BREACH_AT_MPESA_TELCO.HTML. Telco 2.0. 2010.

[53] Eric Brewer et al. "The Case for Technology in Developing Regions". In: *IEEE Computer* 38.6 (2005), pp. 25–38.

[54] Ramón Cáceres, Casey Carter, Chandra Narayanaswami, and Mandayam Raghunath. "Reincarnating PCs with portable SoulPads". In: *Proceedings of the 3rd international conference on Mobile systems, applications, and services*. MobiSys '05. Seattle, Washington: ACM, 2005, pp. 65–78. ISBN: 1-931971-31-5. DOI: HTTP://DOI.ACM.ORG/10.1145/1067170.1067179. URL: HTTP://DOI.ACM.ORG/10.1145/1067170.1067179.

[55] A. R. Garuba. “Computer Virus Phenomena in Cybercafé”. In: *Security and Software for Cybercafes* (2008), p. 186.

[56] O. B. Longe and F. A. Longe. “The Nigerian Web Content: Combating Pornography using Content Filters”. In: *J. Information Tech. Impact* 5 (2005).

[57] *World Bank, World Development Indicators*. 2009. URL: HTTP://DATA.WORLDBANK.ORG/INDICATOR.

[58] M. B. Schmidt, A. C. Johnston, K. P. Arnett, J. Q. Chen, and S. Li. “A cross-cultural comparison of US and Chinese computer security awareness”. In: *Journal of Global Information Management* 16.2 (2008), p. 91.

[59] C. C. Zou, D. Towsley, and W. Gong. “Email virus propagation modeling and analysis”. In: *Department of Electrical and Computer Engineering, Univ. Massachusetts, Amherst, Technical Report: TR-CSE-03-04* (2003).

[60] C. Kanich, C. Kreibich, K. Levchenko, B. Enright, G. M. Voelker, V. Paxson, and S. Savage. “Spamalytics: An empirical analysis of spam marketing conversion”. In: (2008), pp. 3–14.

[61] J. Goodman, G. V. Cormack, and D. Heckerman. “Spam and the ongoing battle for the inbox”. In: *Communications of the ACM* 50.2 (2007), pp. 24–33. ISSN: 0001-0782.

[62] J. E. Phelps, R. Lewis, L. Mobilio, D. Perry, and N. Raman. “Viral marketing or electronic word-of-mouth advertising: Examining consumer responses and motivations to pass along email”. In: *Journal of Advertising Research* 44.04 (2004), pp. 333–348. ISSN: 0021-8499.

[63] B Taylor. “Sender reputation in a large Webmail service”. In: (2006).

[64] M. D. Kibby. “Email forwardables: folklore in the age of the internet”. In: *New Media & Society* 7.6 (2005), p. 770. ISSN: 1461-4448.

[65] M. Sunner. “Developing World, Developing Problems”. In: *Risk Management & BCDR (MessageLabs)* 9 (2009).

[66] P. Y. K. Chau, M. Cole, A. P. Massey, M. Montoya-Weiss, and R. M. O’Keefe. “Cultural differences in the online behavior of consumers”. In: *Communications of the ACM* 45.10 (2002), pp. 138–143. ISSN: 0001-0782.

[67] Shanthi Kannan. “Social networking sites prone to virus attacks”. In: *The Hindu* (2009).

[68] Bin Gu and Vijay Mahajan. “The Benefits of Piracy - A Competitive Perspective”. In: (2004). WISE 2004: Workshop on Information Systems and Economics.

[69] K. Bagchi, P. Kirs, and R. Cerveny. “Global software piracy: can economic factors alone explain the trend?” In: *Communications of the ACM* 49.6 (2006), pp. 70–76. ISSN: 0001-0782.

[70] T. T. Moores. “An analysis of the impact of economic wealth and national culture on the rise and fall of software piracy rates”. In: *Journal of business ethics* 81.1 (2008), pp. 39–51.

[71] A. Katz. “A network effects perspective on software piracy”. In: *University of Toronto Law Journal* 55.2 (2005), pp. 155–216. ISSN: 1710-1174.

[72] P. K. Yu. “Still Dissatisfied After All These Years: Intellectual Property, Post-WTO China, and the Avoidable Cycle of Futility”. In: *Georgia Journal of International and Comparative Law* 34 (2005).

[73] Joe Karaganis. *Media Piracy in Emerging Economies*. Social Science Research Council. 2011. URL: HTTP://PIRACY.SSRC.ORG/ABOUT-THE-REPORT/.

[74] Kevin Stevens. *The Underground Economy of the Pay-Per-Install (PPI) Business*. Blackhat Conference 2010.

[75] Amartya Sen. *DEVELOPMENT AS FREEDOM*. Alfred A Knopf, 1999. ISBN: 0198297580.

[76] A. Cui, Y. Song, P. Prabhu, and S. Stolfo. “Brave New World: Pervasive Insecurity of Embedded Network Devices”. In: (2009), pp. 378–380.

[77] N. Thomas. “Cyber Security in East Asia: Governing Anarchy”. In: *Asian Security* 5.1 (2009), pp. 3–23.

[78] *State of Cybersecurity and the Roadmap to Secure Cyber Community in Cambodia*. IEEE, 2009, pp. 652–657.

[79] V. Godse. “Building an Ecosystem for Cyber Security and Data Protection in India”. In: *Ethics and Policy of Biometrics* (2010), pp. 138–145.

[80] K. Cole, M. Chetty, C. LaRosa, F. Rietta, D. K. Schmitt, S. E. Goodman, and G. A. Atlanta. “Cybersecurity in Africa: An Assessment”. In: (2008).

[81] J. R. Westby. “Countering Terrorism with Cyber Security”. In: *Jurimetrics* 47 (2006), p. 297.

[82] C. Wilson. “Computer Attack and Cyber Terrorism: Vulnerabilities and Policy Issues for Congress”. In: *Focus on Terrorism* (2007), p. 1.

[83] J. Gomez. “Dumbing down democracy: Trends in internet regulation, surveillance and control in Asia”. In: (2004).

[84] Andrew Reynolds, Josh King, Sascha Meinrath, and Thomas Gideon. “The Commotion Wireless project”. In: *Proceedings of the 6th ACM Workshop on Challenged Networks*. ACM. 2011, pp. 1–2.

[85] Paul Gardner-Stephen. *The Serval project: Practical Wireless Ad-hoc Mobile Telecommunications*. 2011.

[86] *The Free Networking Foundation*. HTTP://THEFNF.ORG.

[87] Giulia Fanti, Yahel Ben David, Sebastian Benthall, Eric Brewer, and Scott Shenker. *Rangzen: Circumventing Government-Imposed Communication Blackouts*. Tech. rep. EECS Department, University of California, Berkeley, 2013. URL: HTTP://WWW.EECS.BERKELEY.EDU/PUBS/TECHRPTS/2013/EECS-2013-128.HTML.

[88] Roger Dingledine, Nick Mathewson, and Paul Syverson. *Tor: The Second-Generation Onion Router*. Tech. rep. DTIC Document, 2004.

[89] *Ultrasurf*. HTTPS://ULTRASURF.US/.

[90] *Freegate*. WWW.DIT-INC.US/FREEGATE.

[91] Ian F Akyildiz, Xudong Wang, and Weilin Wang. "Wireless mesh networks: a survey". In: *Computer Networks* 47.4 (2005), pp. 445–487.

[92] *Freifunk Wireless Network*. HTTP://START.FREIFUNK.NET/.

[93] Piyush Gupta and Panganmala R Kumar. "The Capacity of Wireless Networks". In: *IEEE Transactions on Information Theory* 46.2 (2000), pp. 388–404.

[94] Jinyang Li, Charles Blake, Douglas SJ De Couto, Hu Imm Lee, and Robert Morris. "Capacity of ad hoc wireless networks". In: *Proceedings of the 7th Annual International Conference on Mobile Computing and Networking*. ACM. 2001, pp. 61–69.

[95] Daniel Aguayo, John Bicket, Sanjit Biswas, Glenn Judd, and Robert Morris. "Link-level measurements from an 802.11b mesh network". In: *ACM SIGCOMM Computer Communication Review* 34.4 (2004), pp. 121–132.

[96] Matthias Grossglauser and David NC Tse. "Mobility increases the capacity of ad hoc wireless networks". In: *IEEE/ACM Transactions On Networking* 10.4 (2002), pp. 477–486.

[97] Su Yi, Yong Pei, and Shivkumar Kalyanaraman. "On the capacity improvement of ad hoc wireless networks using directional antennas". In: *Proceedings of the 4th ACM International Symposium on Mobile Ad Hoc Networking & Computing*. ACM. 2003, pp. 108–116.

[98] Kevin Fall. "A delay-tolerant network architecture for challenged internets". In: *Proceedings of the 2003 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*. ACM. 2003, pp. 27–34.

[99] Michele Garetto, Paolo Giaccone, and Emilio Leonardi. "Capacity scaling in delay tolerant networks with heterogeneous mobile nodes". In: *Proceedings of the 8th ACM International Symposium on Mobile Ad Hoc Networking and Computing*. ACM. 2007, pp. 41–50.

[100] Uichin Lee, Soon Y Oh, Kang-Won Lee, and Mario Gerla. "Scaling properties of delay tolerant networks with correlated motion patterns". In: *Proceedings of the 4th ACM Workshop on Challenged Networks*. ACM. 2009, pp. 19–26.

[101] Kentaro Toyama. "Technology as Amplifier in International Development". In: *Proceedings of the 2011 iConference*. ACM. 2011, pp. 75–82.

[102] Evgeny Morozov. *The Net Delusion: The Dark Side of Internet Freedom*. PublicAffairs, 2012.

[103] Roger Dingledine, Michael J Freedman, and David Molnar. "The free haven project: Distributed anonymous storage service". In: *Designing Privacy Enhancing Technologies*. Springer. 2001, pp. 67–95.

[104] O. Goga, H. Lei, SHK. Parthasarathi, G. Friedland, R. Sommer, and R. Teixeira. "Exploiting Innocuous Activity for Correlating Users Across Sites". In: *World Wide Web Conference (WWW)*. 2013.

[105] Julien Freudiger. "When Whereabouts is No Longer Thereabouts: Location Privacy in Wireless Networks". PhD thesis. École Polytechnique Fédérale de Lausanne, 2011.

[106] Mudhakar Srivatsa and Mike Hicks. "Deanonymizing mobility traces: Using social network as a side-channel". In: *Proceedings of the 2012 ACM Conference on Computer and Communications Security*. ACM. 2012, pp. 628–637.

[107] Fergal Reid and Martin Harrigan. "An Analysis of Anonymity in the Bitcoin System". In: *Security and Privacy in Social Networks*. Springer, 2013, pp. 197–223.

[108] John C Duchi, Michael I Jordan, and Martin J Wainwright. "Privacy aware learning". In: *arXiv preprint arXiv:1210.2085* (2012).

[109] Henry Corrigan-Gibbs and Bryan Ford. "Dissent: accountable anonymous group messaging". In: *Proceedings of the 17th ACM conference on Computer and communications security*. ACM. 2010, pp. 340–350.

[110] P. Gardner-Stephen and S. Palaniswamy. "Serval mesh software-WiFi multi model management". In: *Proceedings of the 1st International Conference on Wireless Technologies for Humanitarian Relief*. ACM. 2011, pp. 71–77.

[111] Ben Dodson, Ian Vo, TJ Purtell, Aemon Cannon, and Monica Lam. "Musubi: disintermediated interactive social feeds for mobile devices". In: *Proceedings of the 21st international conference on World Wide Web*. ACM. 2012, pp. 211–220.

[112] Arvind Narayanan and Vitaly Shmatikov. "De-anonymizing social networks". In: *Security and Privacy, 2009 30th IEEE Symposium on*. IEEE. 2009, pp. 173–187.

[113] Fergal Reid and Martin Harrigan. "An analysis of anonymity in the bitcoin system". In: *Privacy, security, risk and trust (passat), 2011 ieee third international conference on and 2011 ieee third international conference on social computing (socialcom)*. IEEE. 2011, pp. 1318–1326.

[114] A. Reynolds, J. King, S. Meinrath, and T. Gideon. "The commotion wireless project". In: *Proceedings of the 6th ACM workshop on Challenged networks*. ACM. 2011, pp. 1–2.

[115] J. Glanz and J. Markoff. "US underwrites internet detour around censors". In: *The New York Times* 1 (2011).

[116] J. Dibbell. "The Shadow Web". In: *Scientific American* 306.3 (2012), pp. 60–65.

[117] Haifeng Yu. "Sybil defenses via social networks: a tutorial and survey". In: *SIGACT News* 42 (3 2011), pp. 80–101. ISSN: 0163-5700. DOI: http://doi.acm.org/10.1145/2034575.2034593. URL: http://doi.acm.org/10.1145/2034575.2034593.

[118] Haifeng Yu, Phillip Gibbons, Michael Kaminsky, and Feng Xiao. "SybilLimit: A Near-Optimal Social Network Defense against Sybil Attacks". In: *IEEE Symposium on Security and Privacy*. IEEE Comoputer Society, 2008, pp. 3–17.

[119] S. Trifunovic, F. Legendre, and C. Anastasiades. "Social Trust in Opportunistic Networks". In: *INFOCOM IEEE Conference on Computer Communications Workshops , 2010*. March, pp. 1–6.

[120] Emiliano De Cristofaro, Paolo Gasti, and Gene Tsudik. *Fast and Private Computation of Cardinality of Set Intersection and Union*. Cryptology ePrint Archive, Report 2011/141. 2011.

[121] L. Kissner and D. Song. *Private and threshold set-intersection*. Tech. rep. DTIC Document, 2004.

[122] E. De Cristofaro, J. Kim, and G. Tsudik. "Linear-complexity private set intersection protocols secure in malicious model". In: *Advances in Cryptology-ASIACRYPT 2010* (2010), pp. 213–231.

[123] P. Paillier. "Public-key cryptosystems based on composite degree residuosity classes". In: *Advances in CryptologyEUROCRYPT99*. Springer. 1999, pp. 223–238.

[124] J. Camenisch and M. Stadler. *Proof systems for general statements about discrete logarithms*. Technical report, ETH Zurich. 1997.

[125] Ronald Cramer, Ivan Damgård, and Jesper Nielsen. "Multiparty computation from threshold homomorphic encryption". In: *Advances in cryptologyEUROCRYPT 2001* (2001), pp. 280–300.

[126] Markus Jakobsson and Ari Juels. "Mix and match: Secure function evaluation via ciphertexts". In: *Advances in CryptologyASIACRYPT 2000* (2000), pp. 162–177.

[127] D. Watts and S. Strogatz. "An undirected, unweighted network representing the topology of the western states power grid of the united states". In: *Nature* 393 (1998), pp. 440–442.

[128] R.W. Wolff. "Poisson arrivals see time averages". In: *Operations Research* 30.2 (1982), pp. 223–231.

[129] Nguyen Tran, Bonan Min, Jinyang Li, and Lakshminarayanan Subramanian. "Sybil-resilient online content voting". In: *NSDI'09: Proceedings of the 6th USENIX symposium on Networked systems design and implementation*. Boston, Massachusetts: USENIX Association, 2009, pp. 15–28.

[130] Nguyen Tran, Jinyang Li, Lakshminarayanan Subramanian, and Sherman S.M. Chow. "Optimal Sybil-resilient node admission control". In: *The 30th IEEE International Conference on Computer Communications (INFOCOM 2011)*. Shanghai, P.R. China, Apr. 2011.

[131] George Danezis and Prateek Mittal. "SybilInfer: Detecting Sybil Nodes using Social Networks". In: *NDSS*. 2009.

[132] Gianpiero Costantino, Fabio Martinelli, and Paolo Santi. "Privacy-preserving interest-casting in opportunistic networks". In: *Wireless Communications and Networking Conference (WCNC), 2012*. IEEE. 2012, pp. 2829–2834.

[133] Xiaofeng Lu, Pan Hui, Don Towsley, Juahua Pu, and Zhang Xiong. "Anti-localization anonymous routing for Delay Tolerant Network". In: *Computer Networks* 54.11 (2010), pp. 1899–1910.

[134] Jo Agila Bitsch Link, Christoph Wollgarten, Stefan Schupp, and Klaus Wehrle. "Perfect Difference Sets for Neighbor Discovery: Energy Efficient and Fair". In: *Extremecom*. 2011.

[135] Gjergji Zyba, Stratis Ioannidis, Christophe Diot, and Geoffrey M. Voelker. "Dissemination in opportunistic mobile ad-hoc networks: The power of the crowd". In: *The 30th IEEE International Conference on Computer Communications (INFOCOM 2011)*. Shanghai, P.R. China, Apr. 2011.

[136] P. Dutta, D. Culler, and S. Shenker. "Asynchronous Neighbor Discovery: Finding Needles of Connectivity in Haystacks of Time". In: *Information Processing in Sensor Networks, 2008. IPSN '08. International Conference on*. 2008, pp. 531 –532. DOI: 10.1109/IPSN.2008.60.

[137] Kevin Fall. "A delay-tolerant network architecture for challenged internets". In: *Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications*. SIGCOMM '03. Karlsruhe, Germany: ACM, 2003, pp. 27–34. ISBN: 1-58113-735-4. DOI: HTTP://DOI.ACM.ORG/10.1145/863955.863960. URL: HTTP://DOI.ACM.ORG/10.1145/863955.863960.

[138] Sam Burnett, Nick Feamster, and Santosh Vempala. "Chipping away at censorship firewalls with user-generated content". In: *Proceedings of the 19th USENIX conference on Security*. USENIX Security'10. Washington, DC: USENIX Association, 2010, pp. 29–29. ISBN: 888-7-6666-5555-4. URL: HTTP://DL.ACM.ORG/CITATION.CFM?ID=1929820.1929859.

[139] Yair Sovran, Alana Libonati, and Jinyang Li. "Pass it on: social networks stymie censors". In: *Proceedings of the 7th international conference on Peer-to-peer systems*. IPTPS'08. Tampa Bay, Florida: USENIX Association, 2008, pp. 3–3. URL: HTTP://DL.ACM.ORG/CITATION.CFM?ID=1855641.1855644.

[140] Kok-Kiong Yap, Te-Yuan Huang, Masayoshi Kobayashi, Yiannis Yiakoumis, Nick McKeown, Sachin Katti, and Guru Parulkar. "Making use of all the networks around us: a case study in android". In: *Proceedings of the 2012 ACM SIGCOMM workshop on Cellular networks: operations, challenges, and future design*. ACM. 2012, pp. 19–24.

[141] Michael Demmer, Bowei Du, and Eric Brewer. "TierStore: a distributed filesystem for challenged networks in developing regions". In: *FAST*. Vol. 8. 2008, pp. 1–14.

[142] Jing Su et al. "Haggle: Seamless networking for mobile applications". In: *UbiComp 2007: Ubiquitous Computing* (2007), pp. 391–408.