

# **Connecting the Last Billion**

Copyright 2015  
by  
Yahel Ben David

## 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 3.6 discusses the often overlooked computer-security concerns that are unique to rural communities, especially in developing countries.



# Contents

<b>CONTENTS</b>	<b>i</b>
<b>LIST OF FIGURES</b>	<b>ii</b>
<b>I RURAL INTERNET ACCESS – THE PREMISE OF WISPs</b>	<b>1</b>
<b>1 INTRODUCTION – WHAT ARE WISPs AND WHY SHOULD WE CARE?</b>	<b>2</b>
1.1 WHY WISPs . . . . .	2
1.2 MANY SMALL WISPs VS. GIANT TELCOS . . . . .	4
1.3 WIRELESS MEETS FIBER . . . . .	4
1.4 TOPOGRAPHY, ANTENNAS, FREQUENCIES AND HOW IT ALL FITS TOGETHER . .	12
1.5 WHO ARE WISPs ENTREPRENEURS AND SOME OF THEIR CHALLENGES . . . .	17
<b>2 FURTHERREACH</b>	<b>22</b>
2.1 INTRODUCING FURTHERREACH . . . . .	22
2.2 BUSINESS MODEL AND FINANCES: . . . . .	23
2.3 DESIGN PRINCIPLES . . . . .	31
2.4 DESIGN MEETS PRACTICE . . . . .	34
2.5 THE FURTHERREACH NETWORK . . . . .	36
2.6 THE CELERATE ARCHITECTURE . . . . .	57
2.7 RELATED WORK . . . . .	63
2.8 CONCLUSION AND FUTURE WORK . . . . .	64
<b>3 HOW CAN SDN HELP ENABLE RURAL INTERNET</b>	<b>66</b>
3.1 INTRODUCTION . . . . .	66
3.2 RURAL WIRELESS NETWORK OPERATORS . . . . .	68
3.3 OPPORTUNITIES FOR SDN . . . . .	70
3.4 TOWARDS A VIRTUAL WISP . . . . .	73
3.5 RELATED WORK . . . . .	74
3.6 CONCLUSION . . . . .	75
<b>BIBLIOGRAPHY</b>	<b>76</b>

# List of Figures

1.1	LEVEL-3 COMMUNICATIONS FACILITIES IN CALIFORNIA WITHOUT IP SERVICE - PARTIAL LIST. . . . .	10
1.2	LEVEL-3 COMMUNICATIONS FACILITIES IN NORTHERN CALIFORNIA WHERE IP SERVICES ARE AVAILABLE - PARTIAL LIST. . . . .	11
1.3	LINE-OF-SIGHT AND FRESNEL CLEARANCE ZONE. <i>Image use with permission - Glyn Roylance, Associated Broadcast Consultants</i> . . . . .	13
1.4	SIMILAR COVERAGE WITH DIFFERENT CHOICE OF ANTENNAS . . . . .	17
1.5	SMALL DISTRIBUTED RELAY SITE WITH LOW VISUAL IMPACT. . . . .	18
1.6	A SMALL NEIGHBORHOOD RELAY SITE. . . . .	19
2.1	FURTHERREACH SUBSCRIBERS, PER PLAN, OVER TIME. . . . .	25
2.2	FURTHERREACH REVENUES FROM SUBSCRIPTION AND INSTALLATION FEES. . . . .	27
2.3	FURTHERREACH PROFIT AND LOSS GRAPH ALONG WITH SUBSCRIPTION GROWTH. . . . .	28
2.4	FURTHERREACH'S CUMULATIVE (NEGATIVE) CASH-FLOW GRAPH. . . . .	29
2.5	DC-CURRENTS FOR AN OFF-GRID SITE — HYBRID SOLAR AND WIND. . . . .	35
2.6	FURTHERREACH SUBSCRIBERS MAP. . . . .	37
2.7	FURTHERREACH SUBSCRIBERS IN IRISH-BEACH (A SUBDIVISION OF MANCHESTER, CA). . . . .	39
2.8	FURTHERREACH MAP OF RELAY SITES - FALL 2015. . . . .	40
2.9	FURTHERREACH NETWORK DIAGRAM AND LINK CAPACITIES — SEPTEMBER 2015. . . . .	42
2.10	SCREEN CAPTURE OF A BASE STATION SUPPORTING 31 SUBSCRIBER CONNECTIONS. . . . .	43
2.11	MES SITE — THE ROOF OF OUR 2ND FIBER-FED SITE. . . . .	44
2.12	FURTHERREACH'S NETWORK SWITCH'S GUI AT THE MANCHESTER ELEMENTARY SCHOOL. . . . .	46
2.13	FURTHERREACH'S RELAY ON THE ROOF OF THE POINT ARENA HIGH-SCHOOL'S GYM. . . . .	47
2.14	FURTHERREACH'S NETWORK SWITCH'S GUI AT THE POINT ARENA HIGH-SCHOOL'S GYM. . . . .	49
2.15	FURTHERREACH'S TENMILES TOWER, WITH A LARGE AT&T TOWER IN THE BACKGROUND. . . . .	50
2.16	FURTHERREACH'S "BEAN" RELAY SITE. . . . .	52
2.17	"SMOKEPING" GRAPH TO UNITED STATES COAST GUARD'S (USCG) STATION. . . . .	53

2.18 <i>Navarro</i> SITE — FURTHERREACH’S TOWER POWERED BY A HYBRID OF RENEW- ABLE ENERGY SOURCES. . . . .	54
2.19 <i>Boomer</i> SITE — AN ATTEMPT TO TURN A TREE INTO A CORE RELAY. . . . .	56
2.20 AN OVERVIEW OF THE SUBSYSTEMS IN CELERATE. . . . .	59

## **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*<sup>1</sup> 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.<sup>2</sup>

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

---

<sup>1</sup>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.

<sup>2</sup>Often referred 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,<sup>3</sup> 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,<sup>4</sup> 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<sup>5</sup> 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 than those feeding a single subscriber.

<sup>3</sup>Based on recent years of both the state of California (CPUC) or federally funded grants for rural broadband projects.

<sup>4</sup>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.

<sup>5</sup>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 patterns are particularly high or otherwise indicative of resale, the ISP would either employ automated measures to slowdown and harm the connection (a practice often referred 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.<sup>6</sup> 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

---

<sup>6</sup>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).<sup>7</sup> 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.<sup>8</sup> 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 "burstable" capacity is brought to the fore.<sup>9</sup>
- **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.

---

<sup>7</sup>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.

<sup>8</sup>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.

<sup>9</sup>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<sup>10</sup> or Internet exchanges - but it's unlikely to find these in rural areas where WISPs would benefit from such low bandwidth prices.<sup>11</sup> 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.<sup>12</sup> To summarize, while residential-grade bandwidth has many limitations, its 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.

---

<sup>10</sup>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.

<sup>11</sup>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).

<sup>12</sup>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 WISP's 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 choose 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?

It's 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.

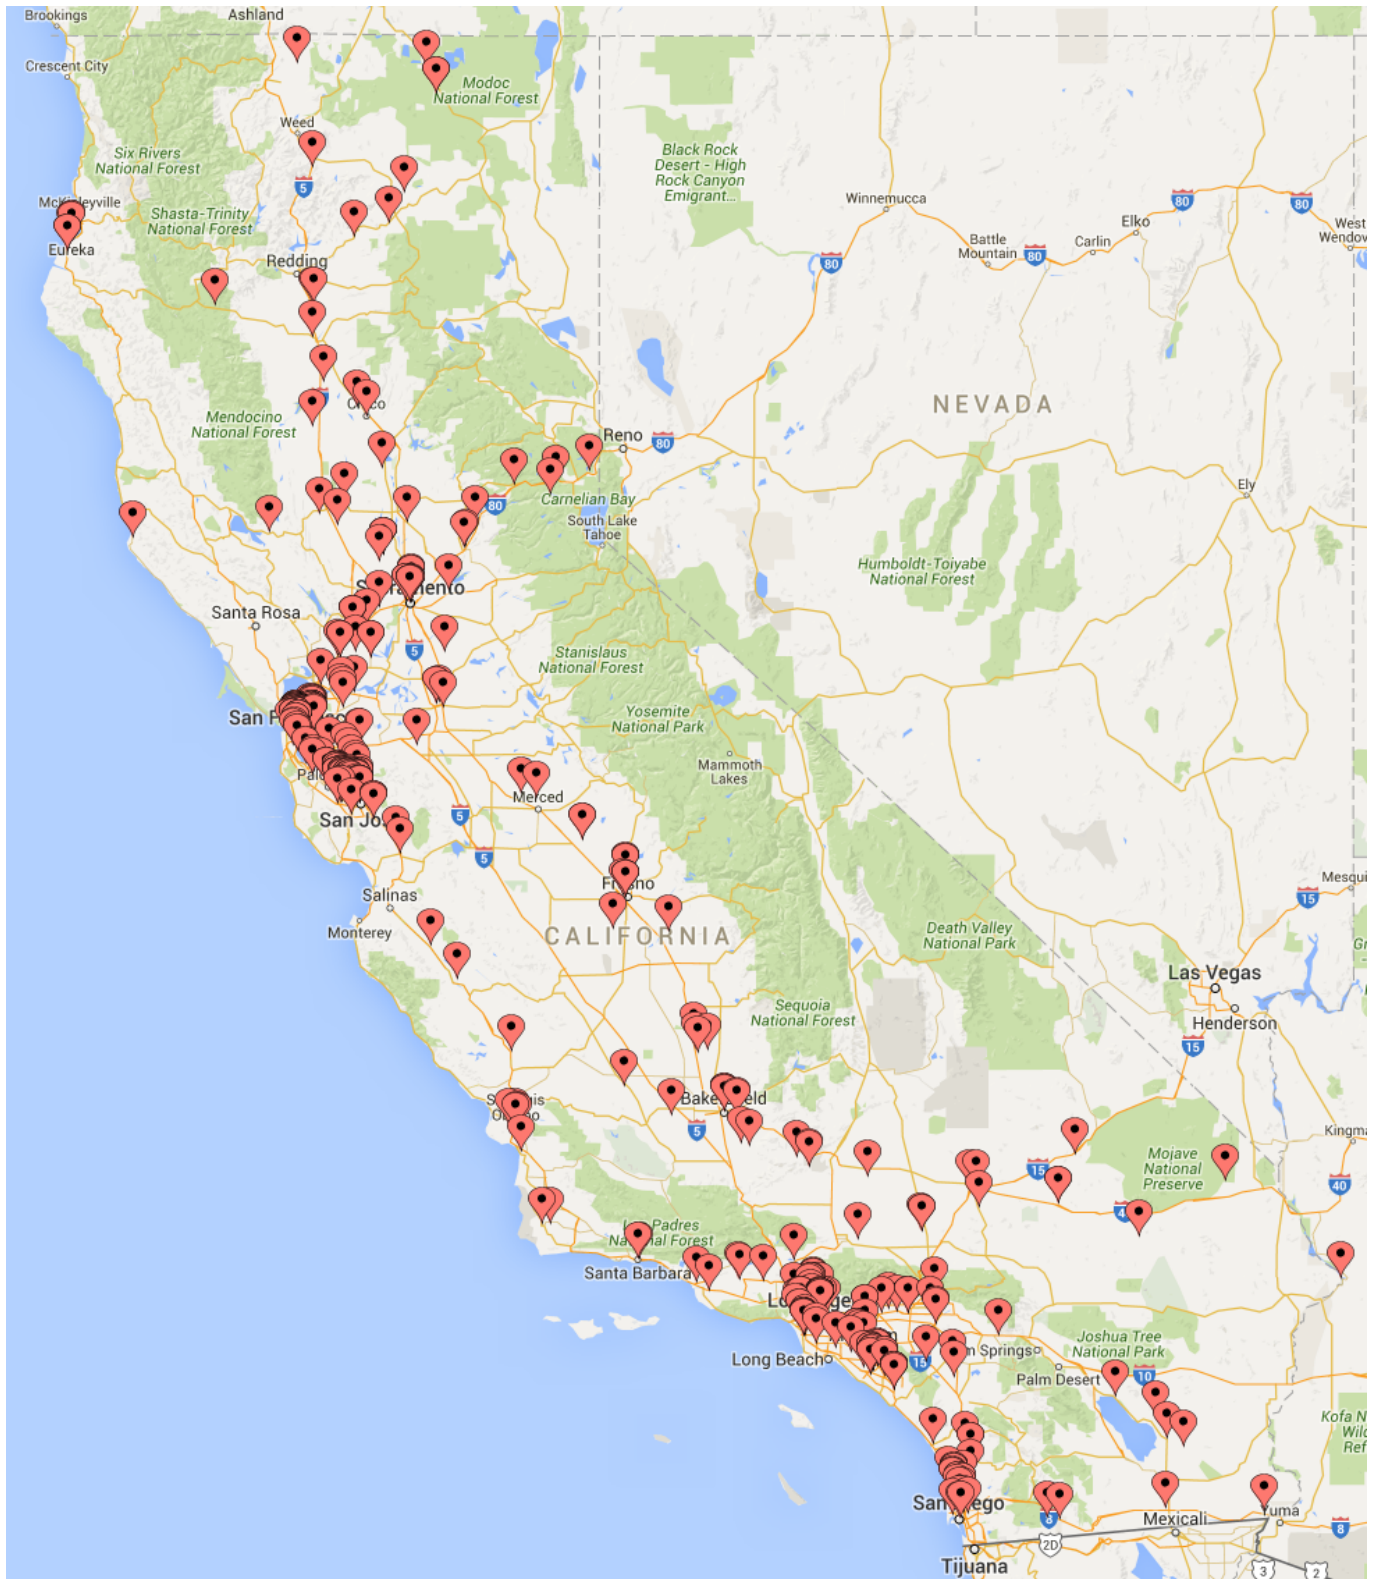


Figure 1.1: Level-3 communications facilities in California without IP service - partial list.



CHAPTER 1. INTRODUCTION – WHAT ARE WISPS AND WHY SHOULD WE CARE? 11

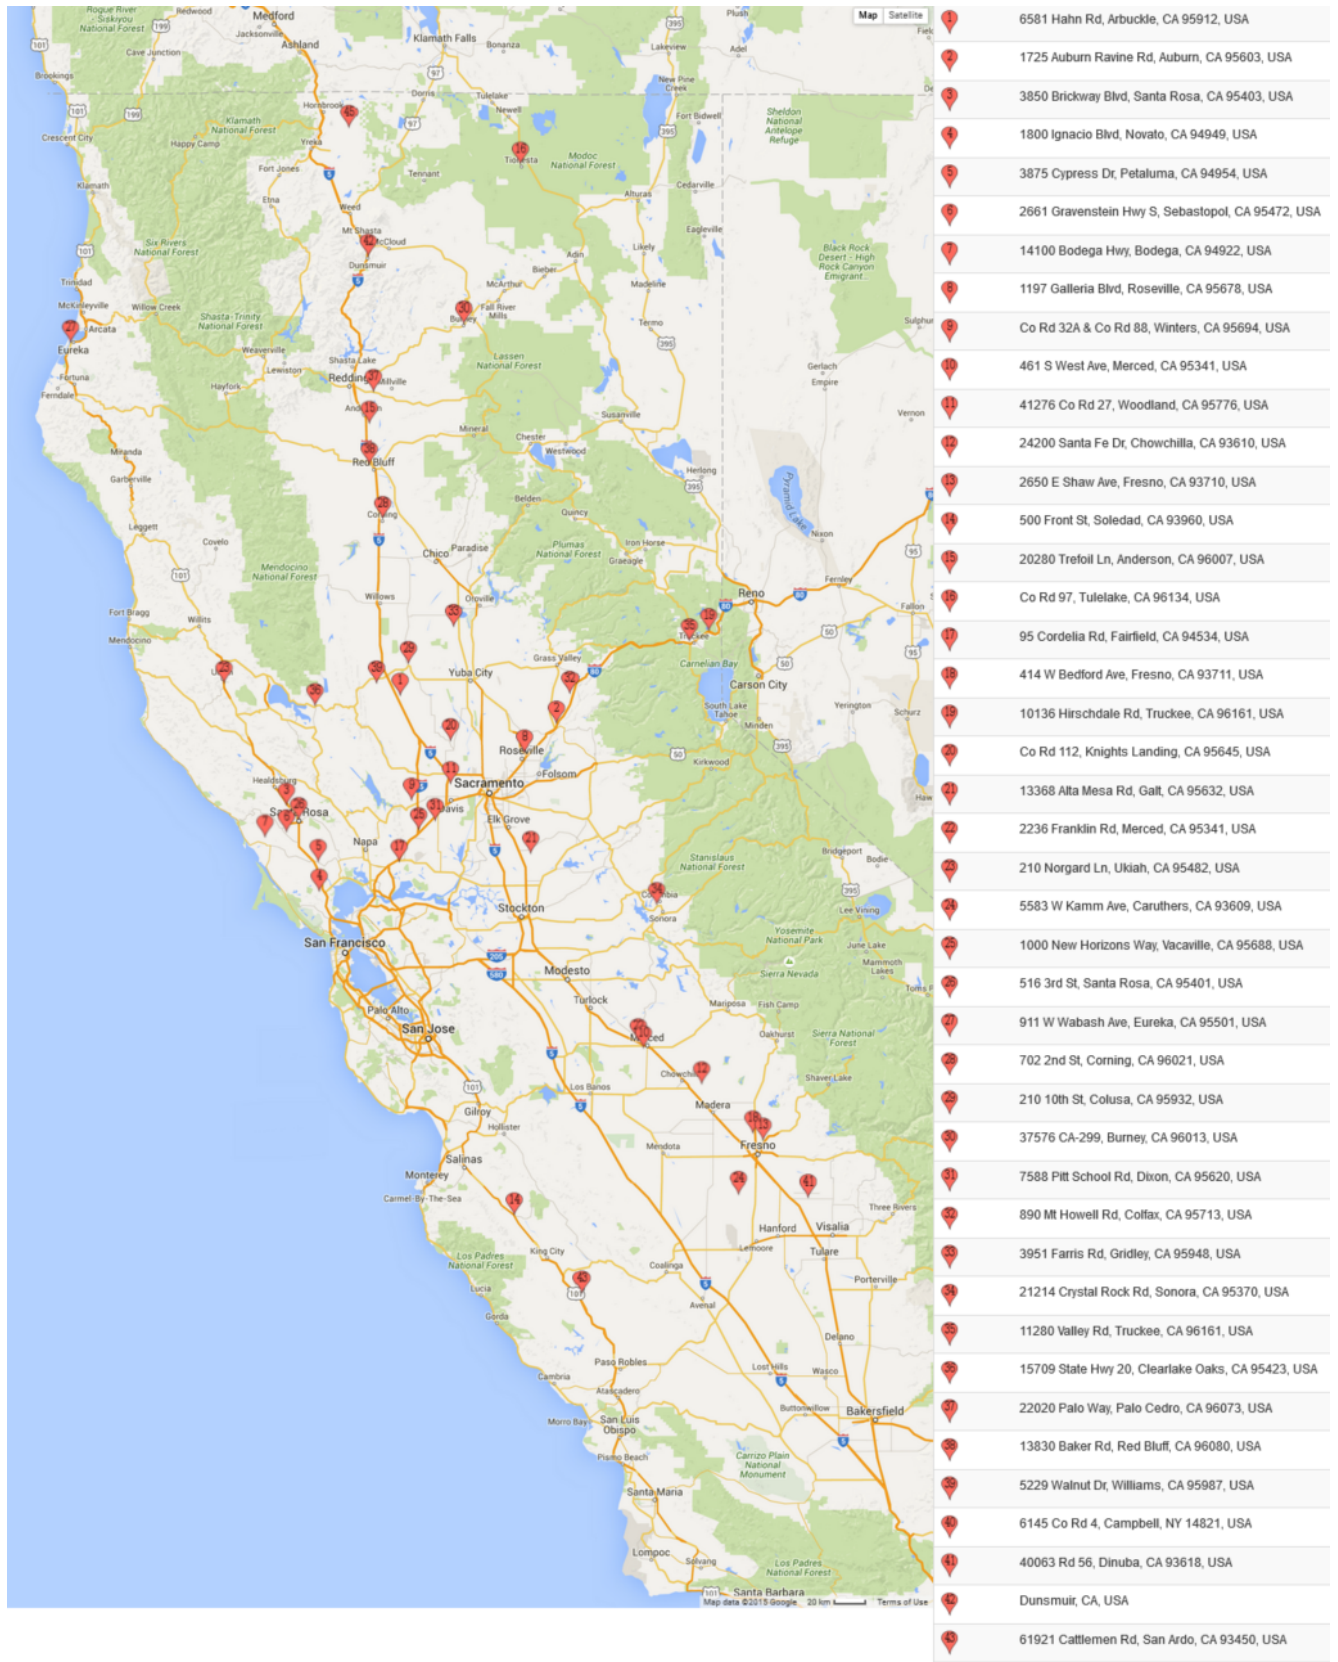


Figure 1.2: Level-3 communications facilities in Northern California where IP services are available - partial list.

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 least a CLEC.<sup>13</sup>

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)
- $\lambda$  is the signal's wavelength (in meters)
- d is the distance between receiver and transmitter (in meters)

---

<sup>13</sup>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).

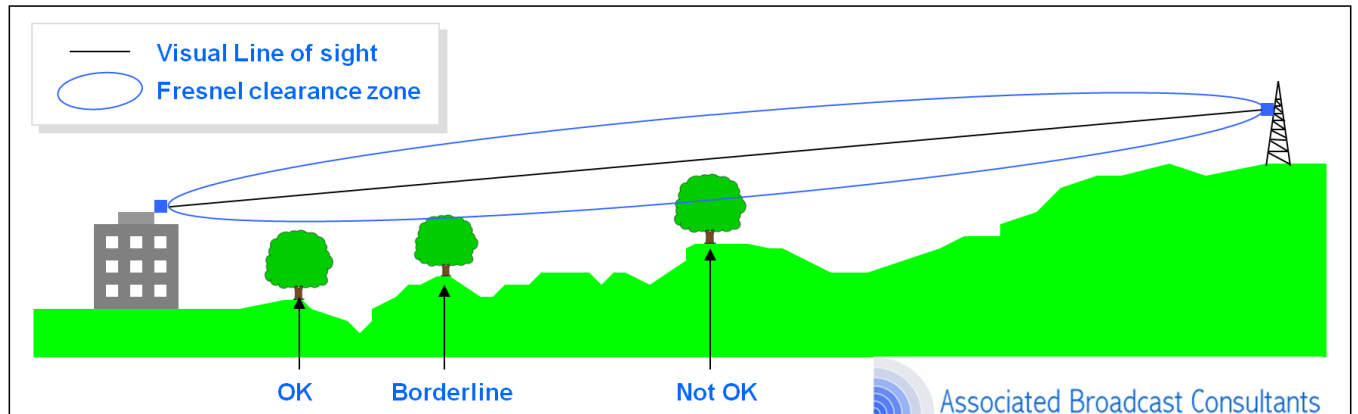


Figure 1.3: Line-of-sight and Fresnel clearance zone.

*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 Fresnel 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.

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

- $\lambda$  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, world-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.<sup>14</sup>
- 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

---

<sup>14</sup>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 direct 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.<sup>15</sup> 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 in 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.<sup>16</sup> 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,<sup>17</sup> 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

---

<sup>15</sup>Newer access points, such as the recent 802.11AC devices support larger numbers of subscribers.

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

<sup>17</sup>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 than 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.




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 than 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 possess 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 WISPs 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<sup>1</sup> 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

---

<sup>1</sup>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 Esplanade'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 cease 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<sup>2</sup> 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 have 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 \$110,000 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<sup>3</sup> 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 De-



ember, 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 began 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.<sup>4</sup> 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.

<sup>4</sup>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 (cumulative)	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
<b>Total monthly expenditures</b>	<b>19,991</b>	<b>25,087</b>	<b>55,720</b>	<b>29,813</b>	<b>30,195</b>	<b>35,756</b>	<b>26,389</b>	<b>41,856</b>	<b>29,319</b>	<b>22,272</b>	<b>20,242</b>	<b>31,697</b>	<b>26,565</b>	<b>28,683</b>	<b>32,482</b>	<b>25,125</b>	<b>481,191</b>
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,<sup>5</sup> 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 didn't 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 protect 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 reasons 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<sup>6</sup>. 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 is 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<sup>7</sup> 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

<sup>6</sup>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.

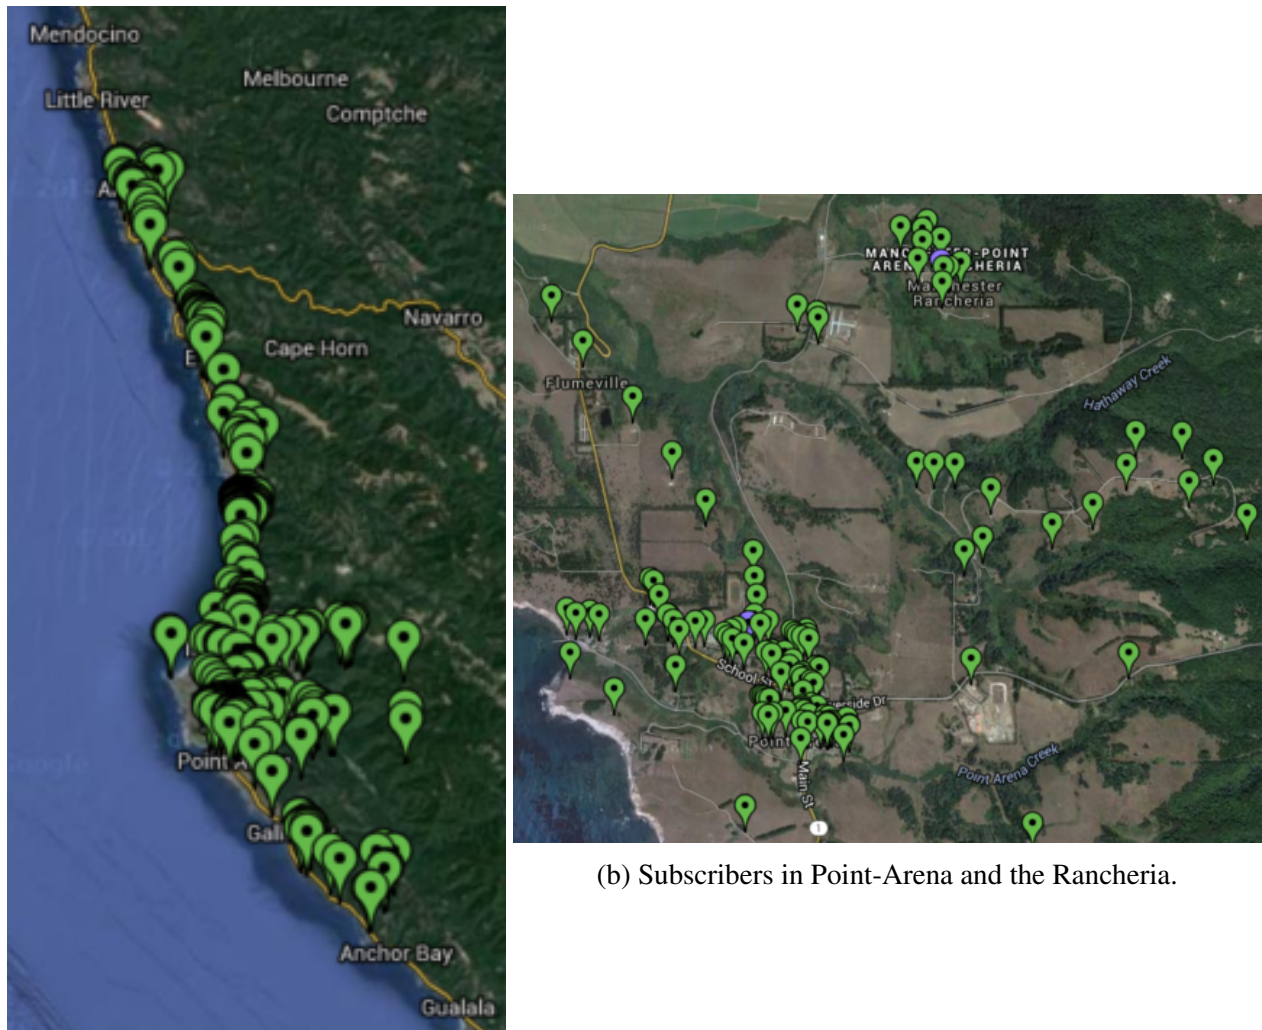
<sup>7</sup>Our batteries of choice are deep-cycle, sealed, AGM (Absorbed Glass Mat) — very costly (A 200AH 12V unit, like we commonly use at FurtherReach goes for \$360 each and a typical 48V core site would have 4 of them).

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.<sup>8</sup> 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<sup>9</sup> 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.

<sup>8</sup>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.

<sup>9</sup>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 (Uninterruptible 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<sup>10</sup>. 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.

<sup>10</sup>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,<sup>11</sup> 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.

---

<sup>11</sup>At dense locations we use a *Ubiquiti Rocket Titanium* BTS which has a gigabit Ethernet port, and faster CPU.





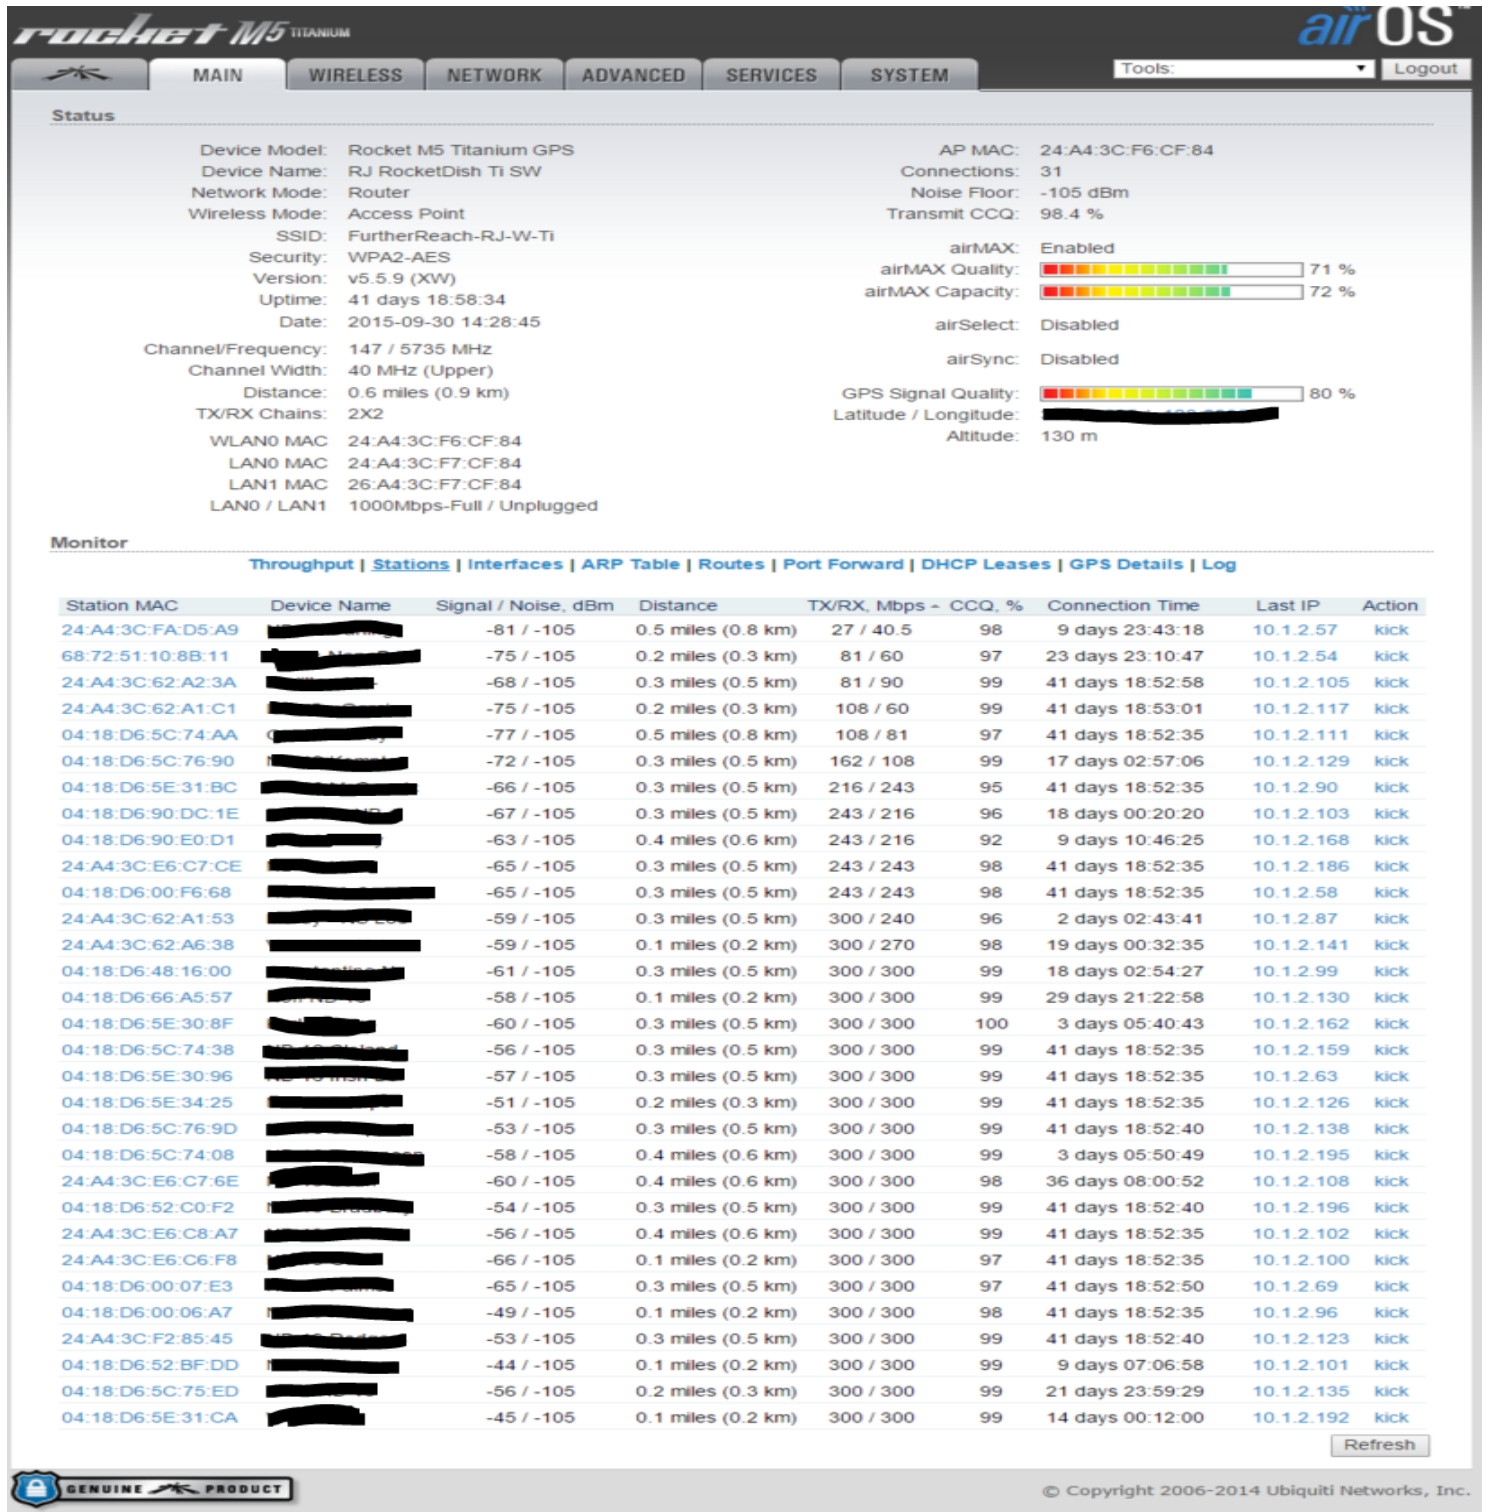


Figure 2.10: Screen capture of a base station supporting 31 subscriber connections.

## 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 than 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.

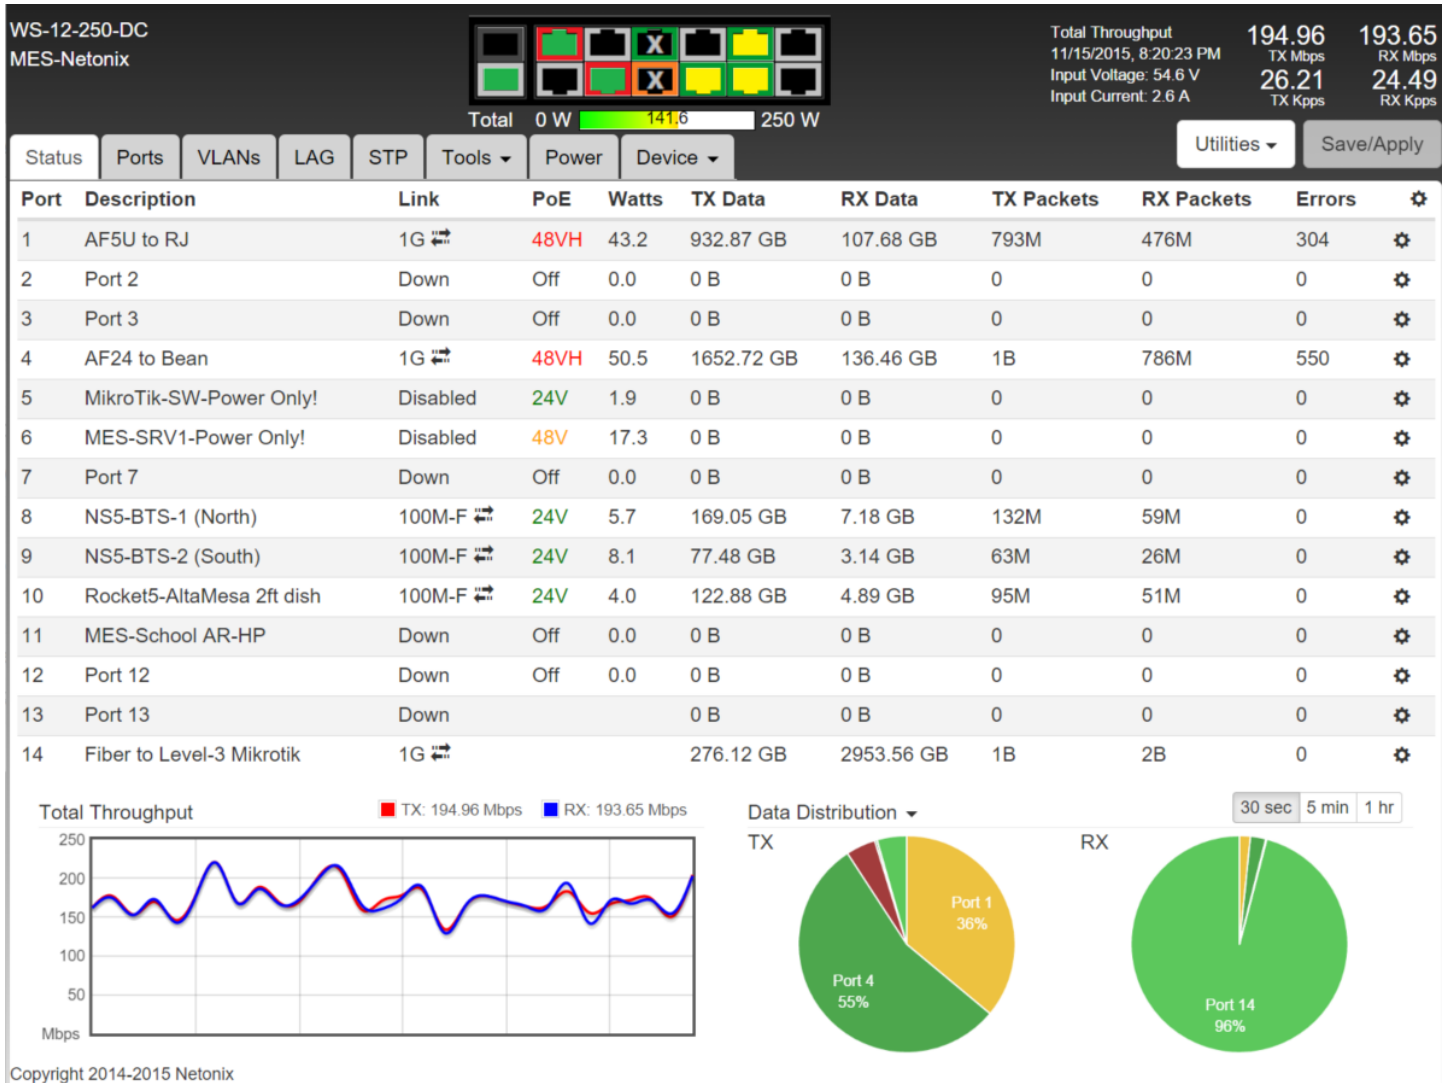


Figure 2.12: FurtherReach’s network switch’s GUI at the Manchester Elementary school.

**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.

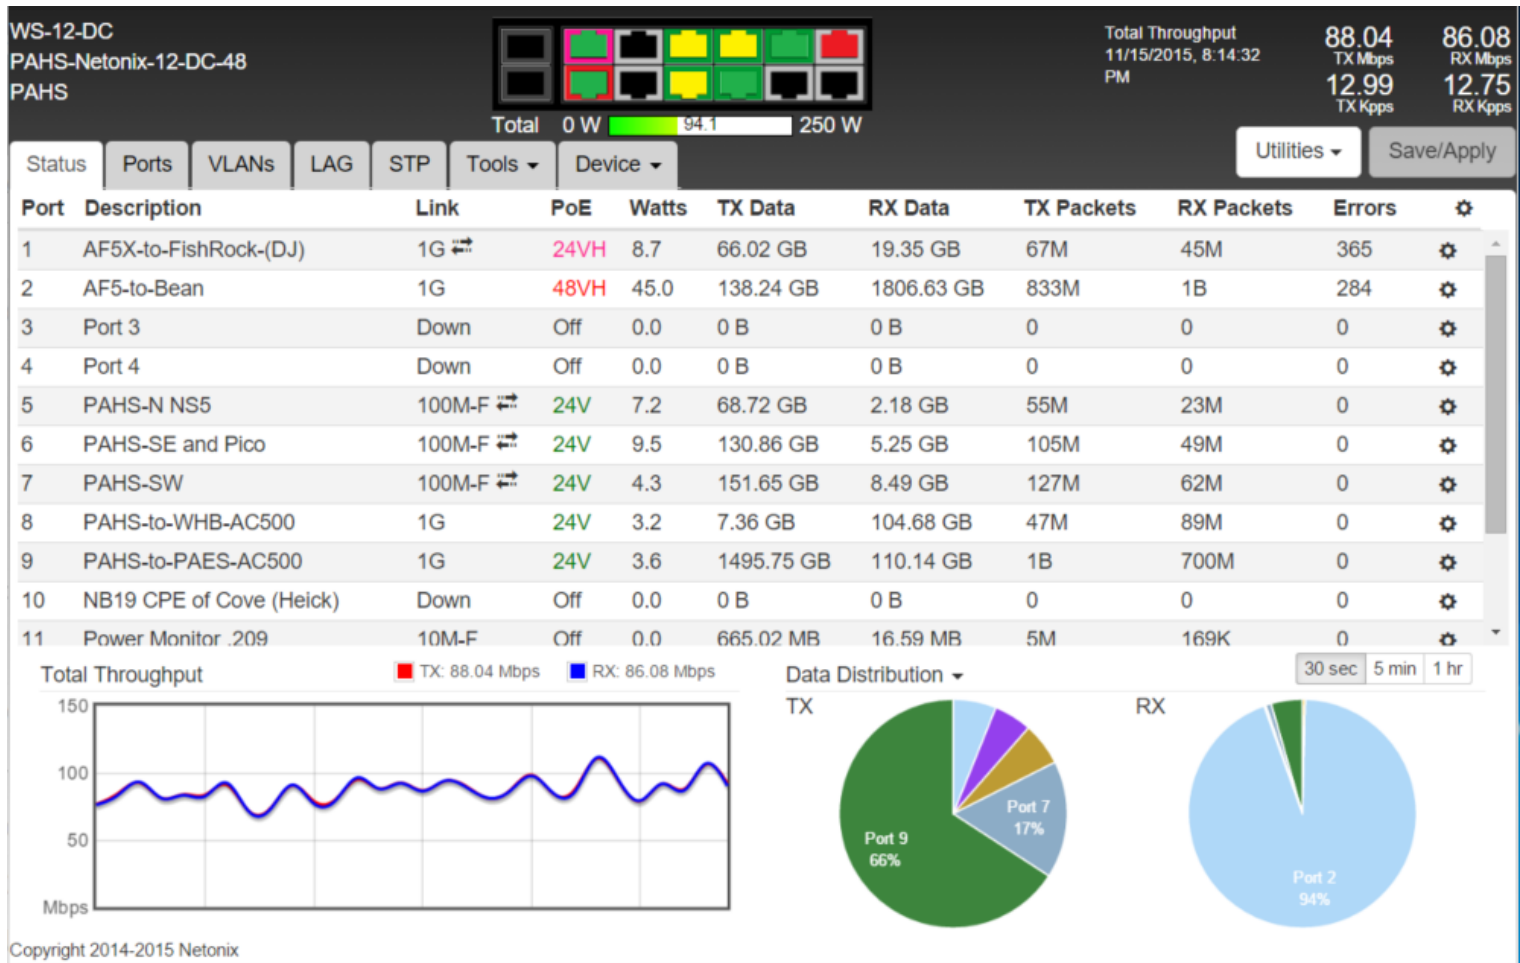


Figure 2.14: FurtherReach’s network switch’s GUI at the Point Arena high-school’s gym.

### 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).




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.<sup>12</sup> 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.

<sup>12</sup>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 tree 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 lose 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



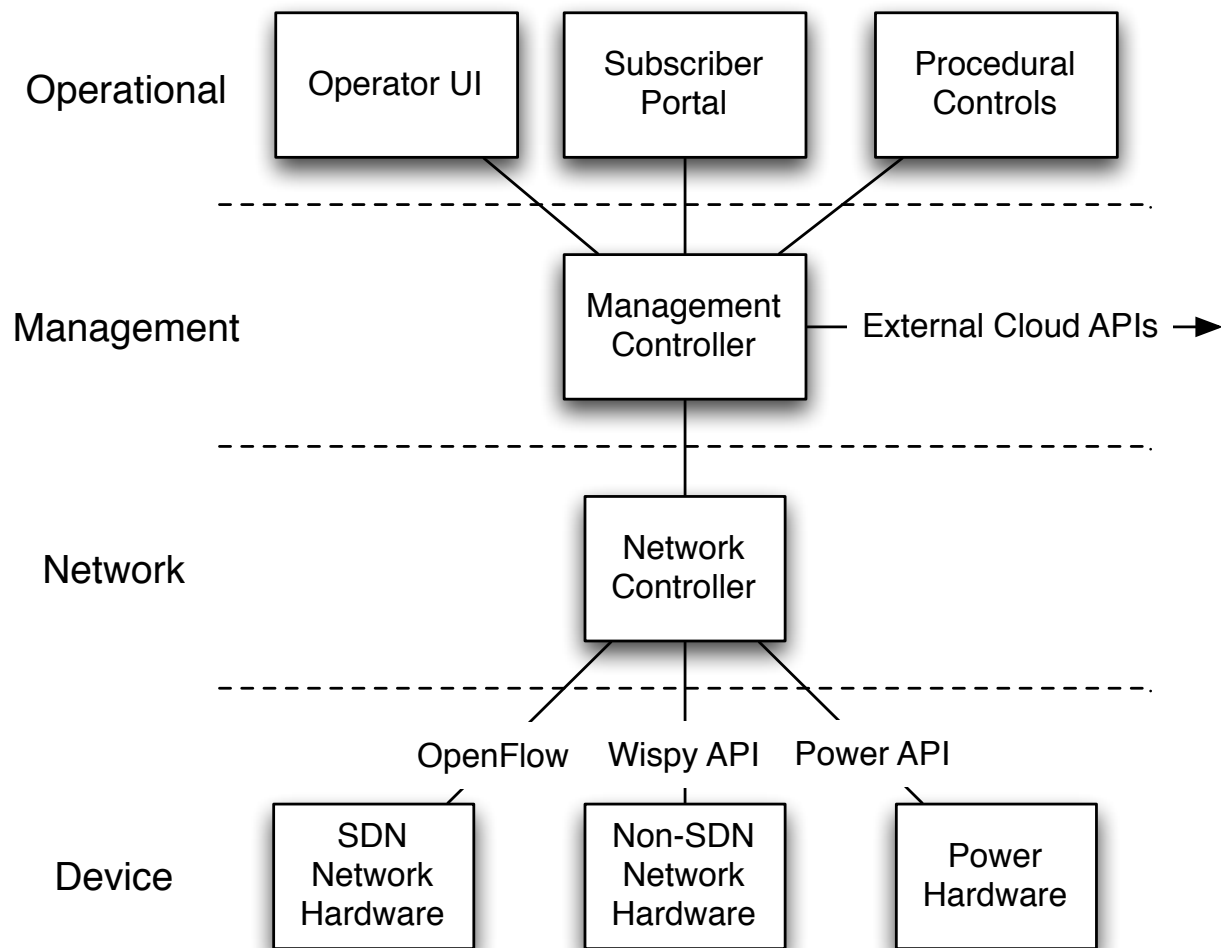


Figure 2.20: An overview of the subsystems in Celerate.

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.<sup>13</sup>

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

---

<sup>13</sup>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.<sup>14</sup>

### 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.<sup>15</sup> 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 is 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

---

<sup>14</sup>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].

<sup>15</sup>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 an 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 de-

ployment 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.<sup>1</sup>

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

---

<sup>1</sup>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 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.

## 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](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/](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/](http://www.noxrepo.org/pox/about-pox/).
- [10] *Stripe*. [HTTPS://STRIPE.COM/](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 Nedeveschi, 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 Nedeveschi, 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/](http://HEYWHATSTHAT.COM/).
- [20] *Powercode*. [HTTP://POWERCODE.COM/](http://POWERCODE.COM/).
- [21] *Swiftfox*. [HTTP://WWW.SWIFTFOX.NET/](http://WWW.SWIFTFOX.NET/).
- [22] *Azotel Technologies*. [HTTP://WWW.AZOTEL.COM/](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](http://WWW.TANDFONLINE.COM/DOI/PDF/10.1080/02681102.2011.643203). URL: [HTTP://WWW.TANDFONLINE.COM/DOI/ABS/10.1080/02681102.2011.643203](http://WWW.TANDFONLINE.COM/DOI/ABS/10.1080/02681102.2011.643203).
- [26] Joyojeet Pal, Sergiu Nedeveschi, 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](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](http://doi.acm.org/10.1145/1851399.1851404).
- [38] Peter Dely, Andreas Kessler, and Nico Bayer. “OpenFlow for Wireless Mesh Networks”. In: *WiMAN 2011*. 2010.
- [39] Sonesh Surana, Rabin Patra, Sergiu Nedeveschi, 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](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://MERAKE.CISCO.COM/](http://meraki.cisco.com/).