

Helping the Lone Operator in the Vast Frontier

Thomas Pötsch
NYU-AD

thomas.poetsch@nyu.edu

Paul Schmitt
UCSB

pschmitt@cs.ucsb.edu

Jay Chen
NYU-AD

jchen@cs.nyu.edu

Barath Raghavan
ICSI

barath@icsi.berkeley.edu

Abstract

While the networking literature is replete with work on managing and operating networks—from the specifics of protocols to the design of management tools and architectures—there is comparatively little work on planning a network to be rolled out. In part this is because the task of network planning typically falls to carriers (for backbones) and cloud providers (for datacenters), which have the resources and the control to meet their specific needs. Here we consider network planning in situations that are quite different: resource poor and highly constrained.

Frontier networks are often planned by lone operators, and while these networks are small in size, they are large in number: they individually serve relatively few users but in aggregate serve millions of users. The key challenge here is the mismatch between the resources of these small operators and the complexity of their network planning. In this paper we detail the difficulties in this context, building on our team's first-hand knowledge of many networks of this type. We then present our initial efforts at frontier network planning, discuss next steps, and outline several open problems.

1. INTRODUCTION

A number of years ago we set out on a mission to identify and eliminate barriers to the universal adoption of Internet access. As is still true today, we knew that Internet access is as much about economics as it is about technology. Therefore, despite ISPs being unwilling to build expensive infrastructure to serve regions with low user densities, we were certain of our inevitable success so long as cheaper, faster, longer range, and more rugged wireless equipment continued to become available [23].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

HotNets-XV, November 09-10, 2016, Atlanta, GA, USA

© 2016 Copyright held by the owner/author(s). Publication rights licensed to ACM. ISBN 978-1-4503-4661-0/16/11...\$15.00

DOI: <http://dx.doi.org/10.1145/3005745.3005757>

Today, commodity wireless equipment is cheap, Internet is a basic human right, and major companies have joined the effort. However, despite buzz about high-cost, high-complexity, high-tech solutions to the problem, we have made only slow progress toward universal access. Near highly-connected cities there are communities connecting via dialup and their connections are getting *slower*—now crawling along at 9600 bps. Such neglected rural areas are home to 45% of the world's population.

In this paper we explore how to meet the challenges faced by the lone operator in the vast unconnected frontier. Building basic infrastructure in this frontier, even in wealthy nations, is an enormous endeavor. If universal connectivity is to be achieved, it will be not through the few, large operators connecting the last billions. Instead, connectivity will flow through the thousands of lone operators, often one-person outfits, who have a stake in bringing access to their own communities.

The key challenge for these operators is not one of hardware—commodity hardware is widely available and easy to set up—nor is it of management, as there are a number of free systems to aid them once up and running. Instead, it is a mismatch between the skills of to-be operators and the task at hand: planning a frontier network often requires a combination of extensive knowledge and practical expertise seldom found in one individual. As we discuss in Section 2, even for our expert team it was difficult to build such a network quickly, at low cost, and with few missteps; for unskilled lone operators who do not have our resources the difficulty is far greater.

Our remedy is a not purely technological, but a mix of a technological approach—automating the planning of frontier networks—and a social approach—identifying lone operators as the key enablers of universal access, and meeting their needs. The task of network planning typically falls to large carriers (in the case of backbones) and cloud providers (in the case of datacenters), both of which have the financial and political resources to overcome physical obstacles (e.g. dig trenches, acquire spectrum, build large towers, buy land). In contrast, frontier network operators must plan and operate *within* existing constraints and cope with the complex myriad of network planning tasks.

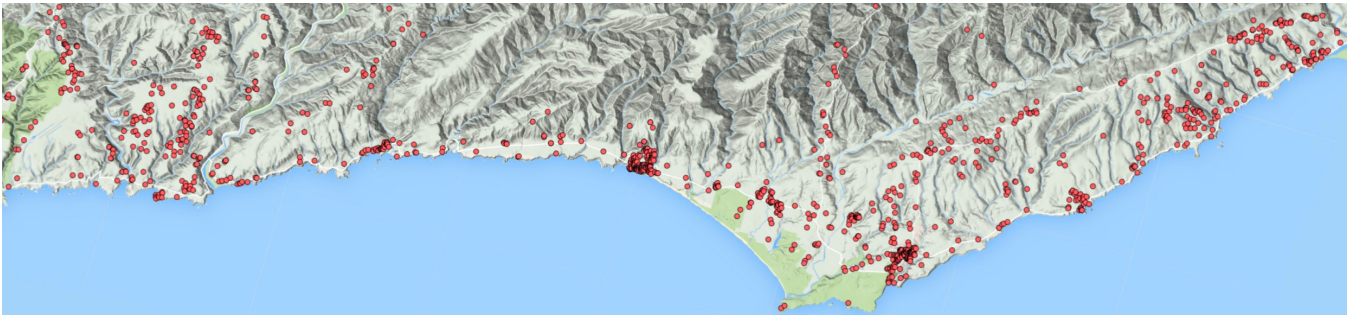


Figure 1: Map of 50 km \times 20 km frontier network region and locations of a subset of (potential) customers.

No one task dominates others in importance when planning frontier networks, but the accumulation of poorly made decisions can easily bring down a network, leaving users in the dark once again as is the history of numerous rural operators. In this paper we motivate the problem by describing one of our own experiences planning and deploying a frontier network. We highlight the practicalities of building such networks, what distinguishes them from other types of networks, what can and cannot be automated, our initial work on this challenge, and what remains to be done.

2. PLANNING FRONTIER NETWORKS

To make the challenges of planning frontier networks apparent, we present a shortened case study on a network we deployed in a previously-unconnected region in rural Northern California [11]. Our experience illustrates some of the problems faced by a team of networking professionals when designing and deploying a frontier network infrastructure. While all networks are unique to their circumstances, the challenges are broadly similar, as found when we spoke with many dozens of rural network operators in North America, Asia, and Africa and given our firsthand knowledge of many of these networks. Our aim is to give the reader a sense of the challenges faced in frontier network design and the operational issues that should be taken into account.

2.1 Case Study Context

We learned of a region, about 50 km by 20 km, that was without broadband Internet connectivity. Local users who wanted Internet access either used a small regional dialup Internet provider or used slow satellite Internet. Figure 1 is a rotated map of the region populated with data of a subset of (potential) customer sites.¹

The region had no coaxial infrastructure and poorly-distributed twisted-pair copper infrastructure, and thus no cable or DSL service. Cellular coverage was spotty, with no 4G service and unreliable 3G service; incumbent telcos had

expressed no interest in improving service to the region, and even left backup generators in disrepair, resulting in frequent outages due to unreliable grid power. Several rivers and creeks cross the main road, which flood frequently cutting off road access. The region as a whole was economically depressed, including a local tribal community, with about a quarter of households living in poverty; however, there were pockets of affluence. Over the past decade at least three other operators have provided service to the region for a time, only to fail due to poor network planning and infrastructure and other challenges, resulting in poor network reliability and performance and leading to eventual business failure. Given this context, our challenge was this: how do we build a cost-effective, performant network to provide connectivity to the population depicted in Figure 1?

2.2 Deployment Process

Our deployment team consisted of several engineers and technicians. Despite being a skilled and experienced group we encountered numerous complex issues in planning, deploying, and managing the network and had to grapple with these issues with few tools at our disposal. As a result, while we were careful to weigh the decisions we made in designing and deploying the network, many decisions were still ad-hoc, and some decisions we made turned out to be mistakes that took time and money to undo.

Our initial task was to identify a source of upstream bandwidth. No universal map of this information exists, and large telecoms (that are the usual providers of such service) do not publicize locations of their fiber facilities in such regions. After hearing local reports of a facility in the region, we contacted a large provider who, after months of our effort following up with them, confirmed for us that they would be able to sell us upstream bandwidth.

The lack of wireline infrastructure and the cost of building cell infrastructure and acquiring spectrum made microwave links (e.g., directional WiFi) a natural choice [18]. This hardware is cheap, low power, and easy to set up. However, such links require line-of-sight, have distance limitations, and can struggle with reflections, intermittent obstructions (i.e., severe weather), and is spectrum constrained.

¹This data is a subset because these sites are only those who directly reported a desire to be connected, and therefore does not represent all households, businesses, civic buildings, and other locations that might need connectivity.

Our first challenge was to determine how to distribute connectivity from the upstream gateway site. The telecom rejected our proposal to mount gear at their facility at low cost, leaving us with no option but to trench fiber from their site to another location nearby where this could be distributed. An ideal nearby site was a large, empty hillside near the facility. After another two months of tracking down and negotiating with the reclusive, elderly owner of the empty land, we were told that we could use the hill only for an exorbitant monthly fee. In parallel we considered several other neighboring sites, all of which were further away and none of which had any elevation. After the hillside was eliminated from consideration, we opted to trench fiber further to an alternative, low-lying location, from which we then had to set up backhaul links to a more distant hilltop location we secured, which would serve as a major distribution hub. In each area that we aimed to expand connectivity, we first spent many weeks using existing rudimentary planning tools [13] to manually identify multiple locations in concert that had line-of-sight and were located with good proximity to user populations. This was ultimately a guess-and-check approach. Once we had narrowed the list of sites, we then spent additional time to negotiate with land owners, businesses, and civic institutions.

The topography of the region—a narrow stretch of land between ocean and mountains that rise 1,000 m—dictated where we could place relay sites. Our constraints were further modulated by additional factors: where we could get power², where line-of-sight existed, where we had access to sites, and where potential users were situated. Existing tools only serve to compute line-of-sight between pairs of nodes, something available in many GIS planning tools. Since such networks have been built for a number of years, we expected that existing tools might be capable of doing semi-automated planning, but we found that the state of the art has scarcely advanced over the last decade.

The choice of frequencies for our radios at our sites was also decided manually and after many considerations. Spectrum contention was commonplace; despite our heavy use of unlicensed 5 GHz spectrum, in which there are numerous non-overlapping channels, we were forced to use other unlicensed bands as well due to contention at major sites. One other network began serving the region after us, in primarily different areas and providing significantly lower speeds.

After over six months of extensive planning, negotiation, and rollout efforts, our modest network consisted of about six sites and provided coverage to perhaps fifty users; it eventually took *years* for our network to expand to serve the majority of the region’s userbase. When unthrottled, many subscribers could receive 30-60 Mbps symmetric throughput to the Internet with less than 5ms latency within our network. At major infrastructure sites we also deploy batteries

and networked power monitors, and power all key network devices using Power-over-Ethernet (PoE).

Several aspects of our network still did not perform well, leading us to attempt to use certain sites only to give up later. We deployed multiple parallel high speed backhaul links between major sites using different bands. However these different bands exhibited dramatically different performance characteristics during different types of severe weather. Another link in our network was from a peninsula to a point up the coast and had poor link quality in certain seasons of the year due to the sea.

2.3 Other Networks

As another example, we worked with one operator in North America, ‘Bob’, who had built his network over the course of over a decade and serves a few hundred users. Bob is a one-man operation who serves a region of about 60 km × 50 km. His network has been built in an entirely ad-hoc fashion, and to date he has no definitive map of all his sites and hardware, of spectrum use, of planning constraints, of capacity/traffic engineering considerations, of address allocation, or of other important information. Consequently, Bob’s network suffers frequent outages that can last for days at a time, but due to the lack of alternatives, customers continue to subscribe to his service.

Bob is a diligent and sincere individual, and the longevity of his network where others have failed is evidence of this fact, but he is not a networking expert. Bob’s network is constantly on the brink of failure. He requires money and time to repair and upgrade his network, which suffers frequent outages, but has little of either. There is a pressing need to design a network planning tool for non-experts like Bob. Such a tool would allow unskilled network operators to run reliable, high-performance frontier networks, potentially providing thousands of communities and millions of users in these communities with network access.

2.4 Complexity Challenges

Frontier networks, unlike ad-hoc networks, require complex planning to ensure high performance, robustness, and cost efficiency. As such, network operators must invest significant effort in planning, or deal with the consequences later. Almost all networking research, however, focuses on network complexity problems *internal* to the network (e.g., wiring, cooling, protocols, management, etc.); networking researchers and engineers are typically insulated from the many *external* planning problems (e.g., facility siting, tower siting, fiber path planning, power management, etc.) that other well-resourced teams are responsible for handling in most large organizations. In frontier networks, when all of these problems are borne by a single individual or a very small team, the task becomes overwhelming. Furthermore, frontier networks do not have the financial or political capital to mitigate the sources of complexity and therefore must address them directly. Here we briefly highlight the differ-

²The devices also tend to require DC power as they are sensitive to fluctuations in AC power.

ences in external complexity between a few network types.

Datacenter Networks. In a datacenter network, the number of possible options for deployment is limited by physical constraints within the region being considered, including where the planning department will permit construction and where sufficient power is available from the grid. While significant effort must go into obtaining a list of these potential sites, there are few degrees of freedom and only one location needs to be selected after evaluation. Within a conventional datacenter, considering a typical topology such as a Clos, there are a key design choices but they are *internal* (e.g., m and n for the Clos and its number of stages, rack design, transport protocols).

Enterprise Wireless Networks. In an enterprise wireless network that must provide wireless service across a number of buildings, we might consider each site to be a floor of a building, all of which must have service. Across a company's campus, for example, this may be on the order of 100 sites. Thus the number of potential sites is equal to the number of sites for the network deployment. Within each site—each floor of a building—there are only a modest number of choices for equipment locations. While coverage must be established despite physical obstructions and hardware choices, the complexity of these considerations is still relatively low as the antennas are omnidirectional.

Regional Wireline ISP Networks. In a regional wireline ISP network, which is most comparable to a frontier network in terms of the number of network nodes and in terms of its service goals, much of the network's design is predetermined by existing wireline copper infrastructure (twisted-pair for DSL and coaxial for cable). Network hardware, such as DSLAMs, is relegated to sites where existing head-end equipment can be placed, such as the regional telco's central office or existing neighborhood-level cabinets.

Frontier Networks. In contrast to the above categories of networks, a typical frontier network is deployed across a large and topographically diverse area (e.g., 50 km \times 20 km, or 1000 square km). In such an area, considering, crudely, that sites are typically parcels of rural land on the order of a couple of hectares each, there are about 50,000 potential sites. Even if we immediately aggregate or discard as non-viable 80% of these potential sites using various heuristics, some 10,000 possible site options remain. At each site, the number of constraints to be considered for placement of devices (which are directional, not omnidirectional) is on the order of twenty, including power availability, tree cover, slope of terrain, orientation, type of radio, type of antenna, type of tower or mast, type of hardware, and more.

Across these potential sites, the network only needs on the order of a dozen sites to serve the area, and such sites must be selected jointly, as the best set of sites (and their configurations) out of the thousands of options. This selection of the best small set from a large set of options results in combinatorial explosion, yielding many orders of magnitude greater

design complexity than in other types of network design. It is the inability to cope with the combinatorics of the problem that frequently pushes network operators to make many ad-hoc design decisions that result in networks that are unreliable and slow—and thus expensive and short-lived.

3. TOWARD AUTOMATED PLANNING

Our goal is to enable the semi-automated design of a frontier network. We envision the process involving a would-be network operator (who may or may not have any network design or management expertise) articulating the geographic locations to be served, policy aims, and other limitations or criteria, and being given a fully-specified network design by the system, including the relay/backhaul locations and the network hardware to deploy, device configurations including spectrum allocation, and physical deployment specifics including elevation and power considerations. Such a design could then be improved through iteration with the design system—for example, as land use is negotiated—and a final design could be used as a blueprint for deployment.

To enable this, the system we aim to build must translate constraints *from the wild* along with user-specified policies into a cohesive model that then enables the construction of a network design by a solver. Eventually, the design generated by the system's solver must be re-represented to enable the operator to refine and converge upon a network design. At the core of this iterative design process is a solver that combines the operator's design specifications and physical models to produce a network design. Thus, we must design an appropriate representation for the solver, write translation tools from the varied and large datasets of constraints, and design a specification language for the user. Here we take the first steps toward automated planning by modeling and representing constraints.

3.1 Modeling

The interactions between elements that compose frontier networks make the problem of automated design particularly complex. There are numerous design elements that must be considered in a frontier network design, including site topography, upstream bandwidth availability, line-of-sight occlusion, spectrum, hardware choice, cost, tower/mount choice, power availability, and more. It is these physical elements we aim to model.

To build the models we require, we must first begin with two sources: raw data (e.g., land topography datasets) and existing research on the behavior of specific elements (e.g., on wireless signal propagation). We must identify high-quality data sources for all of the elements we consider. Many of these data sources are incomplete, vary in granularity in different regions, and can disagree with each other. Since such variability in the data sources can confuse the solver, we must clean these sources.³ We must then ex-

³Ordinarily this is not a research challenge, but no existing stakeholder has both the incentive and means to do so.

press uncertainty explicitly in the models we build for each of these elements.

Each element itself has numerous important properties, and some are inherently difficult to model accurately. However, unlike work that is fundamentally about modeling, our aim in modeling them is not for their own sake, but towards the goal of automated design, and incomplete but practical models can be valuable to this end. Thus we must identify those properties of each design element most salient to the design task at hand, and focus our modeling work on those.

3.2 Solver Representation

Design specifications include three key components: 1) operator goals (e.g., bisection bandwidth, reliability, served geographic areas), 2) operator constraints (e.g., financial limits, deployment time, spectrum/hardware limits, known land availability), and 3) physical models.

The network planning problem naturally lends itself to a graph optimization representation where the vertices represent potential sites and edges represent connectivity (based on line-of-sight and transmission range). Within this optimization problem further specification by the network operator or dictated our physical models may be incorporated into edge weights. The remaining optimization problem of connecting the vertices with the minimum total weight edges can be performed by a solver.

We do not view the theoretical development of a solver in this context as a novel research contribution, nor do we expect to definitively solve this problem from a theoretical perspective. Given that multiple aspects of the design problem are NP-Hard, such as site selection and spectrum allocation [19], the underlying constraint datasets are rife with error, and the policies expressed by the operator are ambiguous, there is limited room for or value in developing an “optimal” algorithm. Instead, it is sufficient to develop a practical system that is capable of producing a result that is substantially better than the status quo today. The resulting decrease in life-cycle cost over the course of the network’s design, deployment, and management will make all the difference for the viability and longevity of frontier networks.

4. PROTOTYPE IMPLEMENTATION

Our prototype planning system is built on top of the GRASS GIS library [10], which provides routines to access, modify and analyse geospatial data. Furthermore, the library provides useful functions for frontier network planning such as terrain analysis, network analysis, data visualization, etc. We use elevation data from the latest SRTM dataset [9] post-processed using GRASS GIS. We model the planning problem as a graph problem that can be solved by the calculation of a Steiner Tree; since this is NP-Hard, we leverage an approximation algorithm. Our system currently consists of the following stages:

Rasterize possible tower locations. Since the number of potential relay tower locations within the region of interest

is large, we define a fixed raster of possible locations on the elevation map. That is, the raster defines possible tower locations so that they are equidistant and narrows the number of possible locations to a finite set (vertices). Although these locations are an approximation and have a lower resolution compared to actual possible locations, this is a reasonable initial approximation before the actual mounting point of the antenna is known.

Determine visibility to other locations. For every possible site location on the pre-defined raster and also given their surrounding terrain, the geographical area that is visible from every location, or *viewsheds*, can be calculated to determine visibility. These viewsheds include all surrounding locations that are within line-of-sight and exclude points that are beyond the transmission range. If two nodes are visible to each other, the potential network link is represented by the addition of an edge between the two vertices. We apply this process pairwise to test the visibility of all tower locations with each other for all combinations of equipment transmission ranges to produce a basic visibility network graph for all possible equipment combinations.

Weight Edges with Costs. The edges in the visibility network are assigned weights that represent the cost associated with the link such as hardware costs, etc. In our current implementation, we use uniform weights for simplicity.

Find optimal tower locations. Given the visibility network graph $G = (V, E, w)$ and desired service locations $S \subseteq V$, the Steiner Tree calculation aims to produce the minimum cost tree that spans the vertices of S where extra vertices from V may be included in the tree so as to reduce the total weight. We leverage a heuristic Steiner Tree solver that produces an approximate minimum-weight Steiner Tree for S .

Figure 2 shows a sequence of steps from our implementation finding the best interconnect among several, arbitrarily chosen customer sites (indicated by red markers). Figure 2(a) shows the elevation data as colored regions where darker indicates a higher altitude of the terrain. The pre-defined possible tower locations (vertices) are projected on the elevation map as orange crosses. For illustration purposes, the horizontal and vertical distances between two locations are chosen to be ~ 750 m and only a restricted region (~ 170 km²) of the map is shown in this example. Figure 2(b) depicts the visibility range of each location in a different color and represents the viewsheds of each location. The viewsheds are further processed and locations that are visible to each other are connected by edges. The resulting visibility network is shown in Figure 2(c) and indicated by gray lines. Finally, intermediate nodes (blue markers) are added to construct a Steiner Tree (red lines) from the 8 initial customer sites.

5. DISCUSSION

Many obvious improvements can be made to our current system including: reducing the raster size, incorporating variable mounting height, antenna sizes, power output,

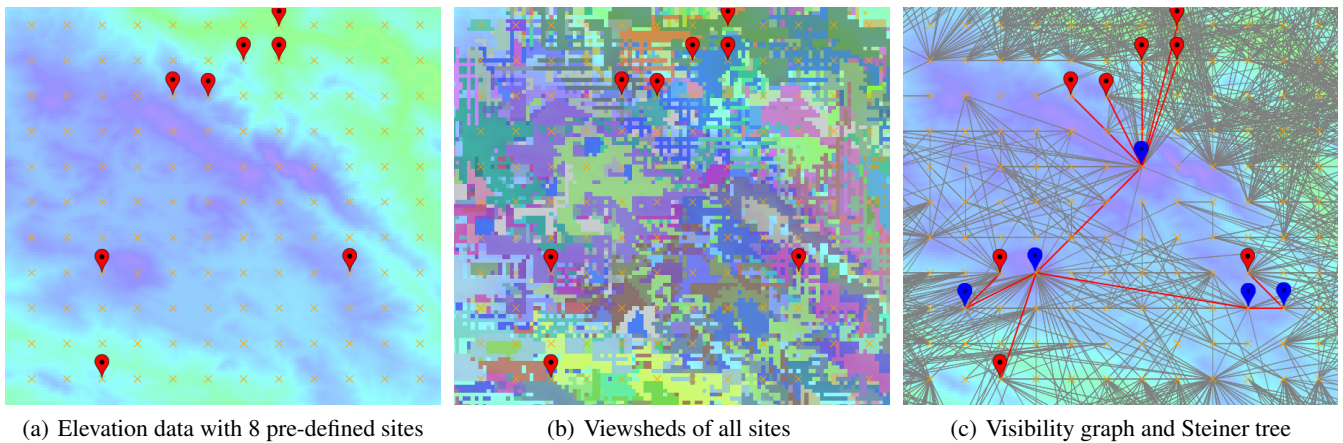


Figure 2: Finding the best relay tower locations (blue markers) and links to interconnect 8 sites (red markers).

tree cover, Fresnel zones, spectrum, equipment costs, reliability metrics, bandwidth, etc. The incorporation of these constraints (at the resolutions we require) will increase the problem size by at least two orders of magnitude. Currently, our implementation—not considering the additional variables above—takes approximately 40 minutes to compute the visibility graph for a region on a single standard machine. Given that the compute time of our algorithms increase quadratically with the problem size, a high fidelity optimization will take about five orders of magnitude more computation, and adding the above variables will increase the complexity by another four orders of magnitude.

We presented a very clean representation of the problem as a set of fixed specifications to optimize against known constraints. However, the real world is messy and there are many unpredictable events. Consider what happens when an operator builds two sites of a five site backhaul network segment, only to lose land rights to the third site, or to find that the third site is in fact blocked by a new obstruction. While the planner’s constraints can be changed to reflect this, this change might otherwise cause our system to produce a wholly different design without the first two sites that were already built. In this situation, the operator needs the ability to express the cost of abandoning work that has already been done (i.e., if there is truly no other way to build the network without abandoning the already-built sites, then our system should find those options, but should prefer leveraging already-completed work).

Another problem is that the Steiner Tree formulation (and other off-the-shelf solvers) may not capture the goals of the network designer. In this initial implementation we used the Steiner Tree optimization because it is a natural fit for the basic planning task, but we did so knowing it would prove inadequate. For example, the output of the algorithm produces a tree, which by definition has no redundant links. Given the general unreliability of wireless links in these settings, it is important to consider backup and parallel links.

6. RELATED WORK

Networking in developing and rural regions is a topic that has attracted attention over the last decade [6, 17]. This focus has intensified in recent years and researchers have recognized that the challenges presented in these networking regimes are substantially different and require different solutions [2, 7, 8, 12, 16, 18, 20, 24]. While there has been work on protocols and techniques for faster and more reliable networks, research often does not extend far enough to keep such networks alive after the research is done [25].

Determining the topology of wireless networks was a significant topic [5, 21, 29] and there are numerous patents on planning cellular networks [1]. Within wireless, variations of topology planning considered hardware factors such as transmission power and directional antenna [14, 21]. The closest work to ours is an algorithm by Sen and Raman that attempts to minimize the overall network cost by considering tower height, antenna type, and transmit power [22].

When deploying their networks today, rural operators use a mix of incomplete planning tools. Some tools build upon terrain data to estimate line-of-sight between two locations, enabling an operator to perform rudimentary topography planning for relay sites [13], other tools [26, 28] provide numerous tools to manually plan, understand, and deploy wireless networks. Other systems, such as TowerDB and Celerate, attempt to simplify network management by juxtaposing geographic locations of devices with network information (e.g. IP-address, frequency, and SSID) [11, 27].

A few alternative planning models have also been proposed. IncrEase [4] is planning paradigm that incrementally introduces sets of additional transmission sites. In [15], the authors describe a mathematical model for automated network planning that considers economic and technical constraints. Several approaches have been proposed in cellular network planning for the placement of base stations (e.g. Andrews et al. [3]), but these generally focus on spectrum and interference rather than physical topography.

7. REFERENCES

- [1] E. Amaldi, A. Capone, and F. Malucelli. Planning umts base station location: Optimization models with power control and algorithms. *IEEE Transactions on wireless Communications*, 2(5):939–952, 2003.
- [2] A. Anand, V. Pejovic, E. M. Belding, and D. L. Johnson. Villagecell: cost effective cellular connectivity in rural areas. In *Proceedings of ACM ICTD*, 2012.
- [3] J. G. Andrews, F. Baccelli, and R. K. Ganti. A tractable approach to coverage and rate in cellular networks. *IEEE Transactions on Communications*, 59(11):3122–3134, 2011.
- [4] G. Bernardi, M. K. Marina, F. Talamona, and D. Rykovanov. Increase: A tool for incremental planning of rural fixed broadband wireless access networks. In *2011 IEEE GLOBECOM Workshops (GC Wkshps)*, 2011.
- [5] S. Bosio, A. Capone, and M. Cesana. Radio planning of wireless local area networks. *IEEE/ACM Transactions on Networking*, 15(6):1414–1427, 2007.
- [6] E. Brewer, M. Demmer, B. Du, M. Ho, M. Kam, S. Nedeveschi, J. Pal, R. Patra, S. Surana, and K. Fall. The case for technology in developing regions. *IEEE Computer*, 38(6), 2005.
- [7] J. Chen, L. Subramanian, J. Iyengar, and B. Ford. Taq: enhancing fairness and performance predictability in small packet regimes. In *Proceedings of ACM EuroSys*, 2014.
- [8] M. Chetty, S. Sundaresan, S. Muckaden, N. Feamster, and E. Calandro. Measuring broadband performance in south africa. In *Proceedings of ACM DEV*, 2013.
- [9] Earth Resources Observation and Science (EROS) Center. U.S. Releases Enhanced Shuttle Land Elevation Data. Available at <http://eros.usgs.gov/>, 2013. Accessed: 2016-06-21.
- [10] GRASS Development Team. *Geographic Resources Analysis Support System (GRASS GIS) Software, Version 7.0*. Open Source Geospatial Foundation, 2016.
- [11] S. Hasan, Y. Ben-David, M. Bittman, and B. Raghavan. The Challenges of Scaling WISPs. In *Proceedings of ACM DEV*, 2015.
- [12] K. Heimerl, K. Ali, J. E. Blumenstock, B. Gawalt, and E. A. Brewer. Expanding rural cellular networks with virtual coverage. In *Proceedings of USENIX/ACM NSDI*, 2013.
- [13] Heywhatsthat. <http://heywhatsthat.com/>.
- [14] Z. Huang, C.-C. Shen, C. Srisathapornphat, and C. Jaikaeo. Topology control for ad hoc networks with directional antennas. In *Computer Communications and Networks, 2002. Proceedings. Eleventh International Conference on*, pages 16–21. IEEE, 2002.
- [15] S. Hurley, S. Allen, D. Ryan, and R. Taplin. Modelling and planning fixed wireless networks. *Wireless Networks*, 16(3):577–592, Apr. 2010.
- [16] D. L. Johnson, E. M. Belding, K. Almeroth, and G. van Stam. Internet usage and performance analysis of a rural wireless network in Macha, Zambia. In *Proceedings of ACM DEV*, 2010.
- [17] S. M. Mishra, J. Hwang, D. Filippini, T. Du, R. Moazzami, and L. Subramanian. Economic analysis of networking technologies for rural developing regions. In *1st Workshop on Internet and Network Economics, Dec 2005*, 2005.
- [18] R. K. Patra, S. Nedeveschi, S. Surana, A. Sheth, L. Subramanian, and E. A. Brewer. Wildnet: Design and implementation of high performance wifi based long distance networks. In *Proceedings of USENIX/ACM NSDI*, 2007.
- [19] C. Peng, H. Zheng, and B. Y. Zhao. Utilization and fairness in spectrum assignment for opportunistic spectrum access. *Mobile Networks and Applications*, 11(4):555–576, 2006.
- [20] B. Raman and K. Chebrolu. Design and evaluation of a new mac protocol for long-distance 802.11 mesh networks. 2005.
- [21] R. Ramanathan and R. Rosales-Hain. Topology control of multihop wireless networks using transmit power adjustment. In *INFOCOM 2000. Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, volume 2, pages 404–413. IEEE, 2000.
- [22] S. Sen and B. Raman. Long distance wireless mesh network planning: problem formulation and solution. In *Proceedings of the 16th international conference on World Wide Web*, pages 893–902. ACM, 2007.
- [23] L. Subramanian, S. Surana, R. Patra, S. Nedeveschi, M. Ho, E. Brewer, and A. Sheth. Rethinking wireless for the developing world. In *Proceedings of ACM SIGCOMM HotNets*, 2006.
- [24] S. Surana, R. K. Patra, S. Nedeveschi, M. Ramos, L. Subramanian, Y. Ben-David, and E. A. Brewer. Beyond pilots: Keeping rural wireless networks alive. In *Proceedings of USENIX/ACM NSDI*, 2008.
- [25] S. Surana, R. K. Patra, S. Nedeveschi, M. Ramos, L. Subramanian, Y. Ben-David, and E. A. Brewer. Beyond Pilots: Keeping Rural Wireless Networks Alive. In *Proceedings of USENIX/ACM NSDI*, 2008.
- [26] Swiftfox. <http://www.swiftfox.net/>.
- [27] Tower DB. <http://towerdb.inveneo.org/>.
- [28] WISPTools. <http://wisptools.net/>.
- [29] Y. Wu, P. A. Chou, Q. Zhang, K. Jain, W. Zhu, and S.-Y. Kung. Network planning in wireless ad hoc networks: a cross-layer approach. *IEEE Journal on Selected Areas in Communications*, 23(1):136–150, 2005.