Rethinking Wireless for the Developing World

Lakshminarayan Subramanian
Intel Research, Berkeley and
New York University

Sonesh Surana, Rabin Patra,
Sergiu Nedevschi, Melissa Ho, Eric Brewer
University of California, Berkeley

Anmol Sheth
University of Colorado, Boulder

ABSTRACT

Many rural regions in developing and developed countries with low user densities do not have good connectivity solutions. To date, networking research has largely focused on urban areas of the industrialized world with high user densities. In this paper, we make the case for research on new appropriate wireless technologies that can provide low-cost, rapidly deployable connectivity solutions for low user-density regions. To this end, we compare and contrast the connectivity requirements that arise in the two domains and pinpoint the new research challenges that arise in low user-density environments. We describe our research efforts in this space and also share our initial experiences in deploying low-cost Wifi-based Long Distance (WiLD) networks in India, Ghana and the San Francisco Bay area.

1 INTRODUCTION

Today, the evolution of networks in the developing world is taking quite an alternate route from the traditional networks we observe in the industrialized world. Many large cities in East Africa have currently deployed a large number of towers supporting a wide range of different long-range wireless technologies such as microwave, long-distance WiFi, WiMax and other commercial wireless broadband solutions. African countries see better opportunity in wireless options for regions that have low penetration of fiber and other wire-line connectivity solutions; many of these countries have a higher cellphone penetration rates than fixed-line penetration [6]. The primary reasons for the boom in the use of long-range wireless networks within developing countries are:

Low-cost and decentralized evolution: In developing countries, wire-line connectivity solutions are not economically viable in low-user density areas [7]. Satellite links, a common mode of Internet connectivity in much of Africa, is also very expensive and not widely affordable (typically US$2,000 per month for 1 Mbps). Establishing wireless distribution networks (microwave, WiMax, WiFi-based or CDMA450) to extend coverage within a region requires a much lower capital investment. This allows for decentralized rapid evolution of such networks by local entrepreneurs. Among different wireless options today, WiFi-based networks are *currently* much more economically viable than WiMax, CDMA450 and microwave.

Ease of deployment: Wireless networks are relatively easy and quick to deploy, particularly in cases where we do not need new towers. Networks in unlicensed spectrum further benefit because they can be set up by grass-roots organizations as needed, avoiding dependence on a telecom carrier. This is particularly important for rural areas, which are less enticing to carriers due to the low density and income of potential consumers.

Intranet usage: Providing network access does not necessarily have to be associated with Internet access. In many developing regions, basic local communications infrastructure is absent. A wireless network within a city or a district can enable a wide range of applications including telephony, essential services and health care. For example, we have deployed an intranet network in South India between hospitals and rural vision centers that supports rural telemedicine [8].

Despite such a phenomenal growth in the adoption of long-range wireless networks in developing regions, there have been very few research efforts that take a concerted view towards analyzing how to build such networks. A primary metric that distinguishes urban environments in developed countries with a majority of regions in the developing world (with the exception of highly populated cities) is the *density of users*. We argue that prior work on wireless mesh networks [4] is best suited for urban environments with high user densities. At lower user densities, the type of wireless network best suited to provide coverage is significantly different from the mesh networking model; such a network would consist of nodes with directional/sector antennas and point-to-point wireless links. Hence, the research challenges that arise in such an environment also significantly differ from those of mesh networks.

In this paper, we outline the research challenges that arise in building low-cost, long-range wireless networks for low density regions. Our research has primarily focused on WiFi-based networks given that WiFi is much cheaper than other wireless technologies and also operates in the unlicensed spectrum. Some of the early works by Bhagwat *et al.* [2] and Raman *et al.* [9] in this space focus on the specific aspects of tailoring the 802.11 MAC protocol to work in such settings; while this is indeed relevant, it represents a small portion of a much larger puzzle. In this paper, we take an end-to-end systems perspective at the overall challenge: how does one engineer a large-scale long-distance wireless network that can provide predictable coverage and good end-to-end performance in the face of competing traffic (from other sources using the same network) and over potentially highly lossy environments (induced by multi-path and external interference) and systemic link/node failures? Answering this question involves addressing research challenges at var-

Characteristic	High User Density	Low User Density
Connectivity requirements	Full coverage required	Islands connected to each other
End Devices	Individual, mobile, low power budget and non-LOS	Shared, fixed, high power and LOS
Topology	Star-topology	Point-to-point with end points within the network
Applications	Mainly Internet access	Internet as well as peer-to-peer Intranet access

Table 1: Characteristics of Low Density and High Density networks

ious layers of the networking stack. In this paper, we elaborate on these challenges and describe our initial efforts towards addressing some of these challenges. We also briefly describe our deployment experiences in building three such WiFi-based long distance networks in India, Ghana and the Bay Area.

2 LOW VS HIGH USER DENSITY REGIONS

In this section, we begin by contrasting low user density (rural and semi-urban) and high user density environments (urban) and make the case for point-to-point long distance wireless networks using directional antennas in low-density environments. We do so by pinpointing why other well-known wireless technologies (VSATs, cellular, mesh networks) are not economically viable in low-density environments. Next, given the distinction between these two environments, we describe the primary differences in the technical challenges that arise in point-to-point wireless networks in comparison to wireless mesh networks, which have received a lot of attention recently.

2.1 The Case for Point-to-Point Wireless

Figure 1 lists some of the fundamental differences between providing wireless connectivity in high user density and low user density environments. These differences mainly stem from the constraints of providing *low cost* wireless connectivity with small per-user cost and minimum or no recurring cost. In low density environments people are usually clustered around small localities (e.g. villages), with large distances among these clusters. Even within villages the user density is low compared to urban areas. In addition, the typically lower incomes lead users to share computer terminals (e.g. Internet kiosks) to amortize the relatively high cost of the devices and network connection.

Satellite networks provide fantastic coverage, but are very expensive. VSAT equipment installation costs over US$10,000 and the recurring monthly costs are over US$2,000 for an 1 Mbps downlink. In low user-density regions, VSAT is affordable only for businesses or wealthy users.

Networks with a base-station model such as WiMAX, and cellular networks like GPRS and CDMA, have an assymetric design philosphy where expensive base stations are amortized by large number of cheap clients over many users. In low-density regions, such base stations simply do not cover enough users to be economically viable. The expectation that cellular solves the connectivity problem for developing regions is thus somewhat of a myth: cellular success in developing countries is an urban phenomenon, with a few exceptions. Bangladesh has good rural coverage because it is actually a very high density country, and base stations that cover roads and rail lines also cover many villages. China has dictated good coverage as policy, despite the economic issues. Other countries either subsidize rural users through taxation, much like the US universal access tax, or require some rural coverage as part of spectrum allocation. In its intended deployment model, with expensive basestations covering many users, WiMax also shares the shortcomings of other cellular technologies.

Finally, 802.11 mesh networks [4], also assume high user density. Moreover, mesh networks suffer from two basic problems when scaled to larger areas. First, as the network grows, an increase in the number of APs with omnidirectional antennas leads to increased interference in overlapping cells. Second, the use of low-gain omni-directional antennas increases the hop length, and as a result throughput decreases. Bicket *et al.* [3] show that in Roofnet, longer routes (traversing multiple wireless hops) are disproportionately slower mainly due to inter-hop collisions.

Thus, we argue that for low density of users, approaches that provide full coverage are not feasible. The alternative would be to cover only those few places where connectivity is required, by employing long-distance point-to-point wireless links. Such links can rely on WiFi, point-to-point WiMax, or other technologies that support long-distance links offering reasonable bandwidths. In choosing such a technology, the most important factors are cost and configurability. An interesting case are environments that have a mix of low and high user density regions. Here, a combined approach where the mesh network is augemented by point-to-point links as required can also be considered ([5]).

Until now, for practical and cost-related reasons, we have chosen to examine the possibility of using WiFi-based Long Distance (WiLD) links. WiFi cards are cheap and highly available, enjoying economies of scale. In our existing WiLD deployments, the cost of a WiLD link is approximately $800 (excludes the cost of tower) with no recurring cost.[1] Because they operate in unlicensed spectrum, WiLD links are easy to deploy and experiment with, and spectrum license costs are eliminated. Manufacturers of WiFi chipsets (e.g. Atheros) often support open-source drivers, allowing us to completely subvert the stock 802.11 MAC protocol and tailor the protocol to meet our needs.

An alternative would be to use point-to-point WiMax

---

[1] We are also deploying solar cells in our WiLD deployments

links; such links would have a few important advantages over WiFi: configurable channel spectrum width (and consequently datarate), better modulation (especially for non-line of sight scenarios); operation in licensed spectrum would permit higher transmit power, and thus longer distances and better signal strengths. However, existing commercial WiMax products are only tailored for cellular providers and do not support point-to-point mode of operation. Existing WiMax hardware is more expensive than WiFi (about \$10000 for basestations), and the high spectrum license costs in most countries dissuage grassroot style deployments. Currently it is also very difficult to obtain licenses for experimental deployment and we are not aware of open-source drivers for WiMax basestations and clients (Wavesat offers a mini-PCI based WiMax client development kit [10]).

Consequently we advocate the use of WiLD links as the currently preferred solution; however, research investigating long-distance point-to-point wireless networking should be (for the most part) agnostic to the specific underlying wireless technology being used, allowing for other solutions to be used as they become available. We formulate our research challenges accordingly.

2.2 WiLD vs Mesh networks

We continue by discussing how the characteristics of WiLD networks differ from those of mesh networks, and thus lead to very different research agendas. We point out three key aspects that significantly differ between 802.11 deployments in low-density settings (WiLD networks) and high-density settings (mesh networks): external WiFi interference, multipath characteristics and routing protocol characteristics.

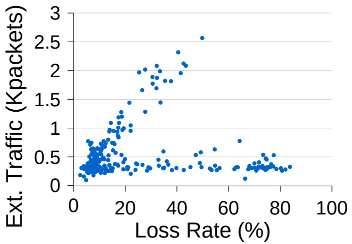

Figure 1: Loss rate vs. ext. traffic observed on WiLD link

External WiFi Interference: In settings where WiLD links co-exist with other external omni-directional WiFi transmitters (access points within the neighborhood), the hidden terminal problem is exacerbated. This is due to two features of WiLD links: directional transmissions and links with long propagation delays. Due to the highly directional nature of the transmission, a large fraction of interfering sources within range of the receiver act as hidden terminals since they cannot sense the directional transmission. However, in an omni-directional mesh network with overlapping transmission regions among neighbors, the fraction of external interfering sources that act as hidden terminals is much smaller. Due to long propagation delays, even external interfering sources within the range of a directional transmitter can interfere by detecting the medium to be busy too late. Hence in WiLD settings, *any* external source can act as a hidden terminal.

Therefore, external WiFi interference can be a very important source of loss in WiLD environments; this is much less so in mesh networks. Figure 1 shows a scatter plot between the loss rate and the absolute number of external WiFi traffic frames received on an urban link over a period of 6 hours. The figure shows that a subset of the loss rate samples are strongly correlated with the external traffic.[2] This result is very different from the measurements reported in Roofnet [4] where the authors show the correlation between loss rate and external WiFi traffic to be very weak. Although these measurements are collected in urban links, they also directly apply in low-density networks where one of the endpoints is in an urban environment.

Multipath characteristics: In Roofnet [1], the authors conclude that multipath interference was a significant source of packet loss. However, in WiLD networks, we observe quite the opposite. This is primarily because the delay spreads in WiLD environments are an order of magnitude lower than that of mesh networks. The two factors contributing to lower delay spreads in WiLD networks are the long distance of WiLD links, and the line-of-sight (LOS) deployment of the nodes. The strong line-of-sight component in WiLD deployments ensures that the attenuation of the primary signal is only due to path loss, and most of the secondary paths are due to reflections from the ground. Furthermore, the long distance between the endpoints ensures that the primary and the secondary reflection travel almost the same distance, and hence reduces the delay spread. In comparison to our WiLD deployment, the Roofnet deployment has shorter links and non-LOS deployments, which significantly increases the delay spread.

Routing: From a topology perspective, two distinguishing factors between mesh and WiLD networks are that mesh networks are unplanned while WiLD networks are planned, and that the quality of links in mesh networks is time-varying and nodes have several neighbors to potentially forward packets. Hence, in mesh networks, routing is more opportunistic where nodes forward packets based on the quality of the link at a given time. Roofnet's routing protocol, Srcr, chooses routes with a minimum "estimated transmission time" (ETT) as a route selection metric [3]. In contrast, WiLD networks consist of a few dedicated point-to-point links and routing in WiLD networks resembles traditional routing protocols.

3 EXISTING DEPLOYMENT

Currently, we have deployed several WiLD networks in India (a 9-link topology), Ghana (5 links) and the Bay Area

---

[2] Based on experiments performed in a wireless channel emulator we observed that at a channel separation of 2, the receiver is not able to receive the frames from the external interference source. However, the signal spillage of the interference source in the primary channel is sufficient to cause frame corruption. This explains why a subset of loss rate is not correlated with external WiFi traffic.

in the US (7 links). We use these testbed deployments to understand the different research issues and to implement and evaluate the solutions to those challenges. The WiLD network in India connects several village-based vision centers to the local Aravind Eye Hospital, and supports remote eye care as well as distance learning through interactive video conferencing. In Ghana, the links are used by the University of Ghana to share Internet access, for distance learning, and to exchange electronic library information among its different campuses. Distances of our WiLD links vary from 10–80km, with relays installed where there is not line of sight due to geographical limitations.

We use low power single board computers (SBC) with a 266 MHz x86-based chip, 128 MB RAM and up to 3 wireless cards for our wireless routers. For radios, we use off-the-shelf high power 802.11a/b/g Atheros cards with up to 400 mW of transmit power output. The platform runs a stripped down version of Linux from a 256 MB CompactFlash card. To form long distance links we use high gain parabolic directional antennas (24 dBi, 8 degree beam-width). In multihop settings, nodes can use multiple radios with one radio per fixed point-to-point link to each neighbor.

The above choice of hardware enables us to design routers that are low cost(less than $400), consume less power (5–10W) and are of low weight (10–15 kg for a node with two antennas). While the small size and weight allows us to use less expensive guyed-wired towers, the low power consumption means that we can use small solar panels, which reduce the operating cost and increase reliability when uninterrupted grid power supply is not available in developing regions.

4 RESEARCH CHALLENGES

In this section, we elaborate on the research challenges that arise in engineering large-scale WiLD networks to achieve predictable end-to-end performance in the face of competing traffic from other sources and highly lossy links (induced by external interference). We classify the research challenges into the following categories: (1) MAC layer challenges; (2) Loss recovery mechanisms; (3) QoS Provisioning; (4) Troubleshooting, reconfigurability and management; (5) Network planning and deployment. Associated with each of these challenges, we describe some of our early efforts to address them.

4.1 MAC Layer Challenges

The first challenge in running 802.11 on long-distance multihop links is to adapt the 802.11 MAC protocol [9] to overcome its fundamental limitations which can be summarized as:

- ACK timeouts: The simple stop-and-wait recovery mechanism of the stock 802.11 protocol requires each packet to be independently acknowledged. This recovery mechanism is ill-suited for long propagation delays, as it limits utilization and thus bandwidth. Worse, if the time taken for the ACK to return exceeds a card-specific maximum timeout, the sender will retransmit unnecessarily and waste bandwidth .
- Collisions due to bidirectional traffic: The CSMA/CA channel-access mechanism is not suitable for long distance links; listening at the transmitter reveals little about the state of the receiver, due to the long distance and stale carrier sense information due to propagation delays.
- Multi-link Interference: When multiple WiLD links originating from a single node operate on the same or overlapping channels, the transmission of one link can interfere with packet reception on other links, because local side lobes are of similar strength to the signal received from afar.

TDMA MAC Protocol with sliding window: The above limitations of the stock 802.11 MAC protocol motivate the need for a TDMA-based MAC protocol that synchronizes the transmissions from the endpoints of a single point-to-point link. For a node having multiple outgoing point-to-point links, Raman et al. [9] propose having *simultaneous send* and *simultaneous receive* to eliminate interference. In addition, the stop-and-wait recovery mechanism of 802.11 is unsuitable. We implement a sliding-window based flow-control approach with the TDMA slots.

TDMA Slot Scheduling: Given these constraint of simultaneous transmit and receive, finding a feasible TDMA slot schedule in a multihop network is non-trivial especially if we want to achieve optimal throughput across the whole network. However, it can be shown that for bipartite graphs, we can always find such a slot schedule.

4.2 Loss Recovery Mechanisms

Across all of our WiLD networks, the presence of external WiFi interference results in very high loss rates on WiLD links. Furthermore, due to the long distances, the extent of interference could be very different at the two ends, making WiLD links asymmetric. Also, it is common to have links with loss rates fluctuating between $5-80\%$ over short time scales.

Figure 2 shows the loss rate sampled every 1 minute across channel 1 and 11 for a 20 km WiLD link. The figure shows that both channel 1 and 11 have long bursts of high loss rate due to external interference. Even in absence of long bursts there still exists a residual 5–8% loss. Given the situation, an important challenge is to device appropriate link level loss recovery mechanisms that can achieve predictable performance in the face of high loss variations.

Retranmissions with Bulk ACKs: The first approach for loss recovery is where the receiver acknowledges a set of frames at once using bulk ACKs, in the sliding window setting proposed previously. The lost packets are then retransmitted accordingly.

Figure 3 shows the comparison of bidirectional TCP throughput achieved at various distances by the stock 802.11 MAC protocol (using CSMA) and by our implementation of the TDMA MAC protocol with bulk ACKs. To emulate long distances, we use a wireless channel emulator.. We can see

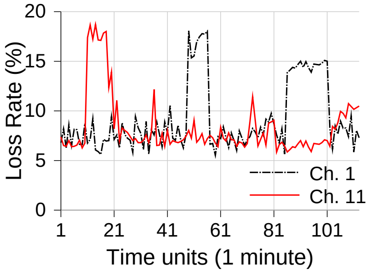

Figure 2: Loss variation over time across channels 1 and 11

that as the distance increases, the throughput of CSMA MAC decreases gradually until distance reaches 110 km, which corresponds with the maximum ACK timeout, and then it drops drastically. However, the TDMA MAC protocol using bulk ACKs provides sustained high throughput even at very long ranges.

Adaptive FEC: With such highly variable packet losses such as shown in figure 2, the retransmissions based approach would give us 0% loss but with highly variable delay and this is not suitable for audio and video traffic. We therefore propose an adaptive FEC based loss recovery mechanism which limits the delay experienced at each hop while guaranteeing a small loss rate. We are currently investigating appropriate FEC coding mechanisms for our WiLD setting. We observe that the loss variability of the WiLD links are very hard to predict, making the problem of determining the appropriate FEC recovery mechanism a challenging one.

4.3 Quality of Service

Many applications that use WiLD networks require QoS (*e.g.*, video-conferencing sessions in rural telemedicine). Unlike the case of the Internet architecture, in WiLD networks we have the flexibility of modifying routers to implement QoS mechanisms. However, many of the traditional QoS mechanisms do not blindly carry over due to peculiar constraints imposed by WiLD networks. First, unlike traditional wired links, WiLD links cannot be characterized by a fixed bandwidth value. In the presence of high loss variations, the available bandwidth (after recovery) is time-varying. Also, the need for synchronous packet transmissions and receptions at a node, creates a direct coupling between the available bandwidth on adjacent links; in other words, any variation in the slot size along one link, affects the one-way bandwidth on adjacent links. Second, WiLD networks experience highly-variable delays due to the TDMA nature of packet transmissions coupled with loss recovery. Hence, providing end-to-end bandwidth and delay guarantees for flows requires scheduling mechanisms that can take into account the variable link bandwidths and link delays. Traditional QoS mechanisms assume the concept of *flow isolation i.e.*, once a set of resources are allocated to a flow, this flow is unaffected by competing flows. This assumption does not completely hold in WiLD settings since the introduction of a new flow can potentially affect the resource allocation of competing flows (either along on links in the same path or adjacent links along the path).

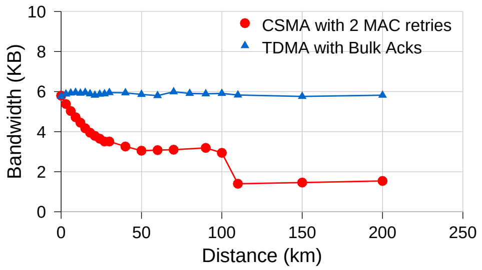

Figure 3: Comparison of WiLD MAC and stock 802.11 MAC

In addition to these differences, WiLD nodes have a low processing power (266 MHz) and stringent memory constraints (128 MB) that may rule out many fancy strict/statistical QoS mechanisms which would require nodes to maintain per-flow state and track per-flow usage. We are currently deploying simple QoS mechanisms based on *traffic priority classes* similar to Diffserv without supporting any form of strict guarantees. To provide statistical guarantees at a per-hop level, the primary link-layer parameters that we can manipulate are: (a) loss-recovery parameters (FEC, retransmissions); (b)varying the TDMA slot-size to reduce delay. Manipulating these parameters represents a trade-off spectrum between achieved loss-rate, delay characteristics, available bandwidth. As part of future work, we plan to analyze this trade-off spectrum and quantify the achievable QoS properties in WiLD environments. Another related problem is the *optimal TDMA scheduling problem*: Given a *traffic demand matrix* between various sender-receiver pairs, can we compute an *slot schedule* for every link in the network that can satisfy all the traffic demands?[3]

4.4 Troubleshooting, Reconfigurability and Management

A key aim in WiLD networks is to reduce the operational cost of maintaining the network. This is critical due to the lack of trained manpower in many developing countries, and long delays involved in accessing the endpoints of a link due to the distances and tower/pole deployments of the wireless routers.

Our experience with WiLD deployments shows that the network can malfunction in a number of ways ranging from complete failure of links (hardware board failure, corruption of the flash memory cards, lightening strikes), to performance degradation over time (from misalignment of antennas, signal attenuation from rain water clogging RF cables, interference from external sources).

Reconfigurability: One way to deal with complete failure of links or nodes is to design a redundant network topology, with more than one possible path between the wireless

---

[3] This problem assumes that all links are in the same channel. Given non-overlapping channels, one can imagine a similar problem coupled with the need for an appropriate channel allocation mechanism.

nodes. To reduce the cost of additional redundant links we are exploring the use of low-cost *electronically steerable antennas* instead. On a link failure, these antennas can dynamically realign themselves and reform the topology of the network to route around failed nodes or links such that network connectivity is maintained.

Safe Upgrades: A *safe upgrade mechanism* is also required for changing either the firmware or even the network configurations on the routers. Any failure during this process could lead to the endpoints being disconnected and out of reach. To avoid such failures, we use the built-in hardware watchdog timer to power cycle the router on a failed kernel change or erroneous configuration change and revert to a default “golden” version.

Monitoring: The challenge in network management is to continuously monitor the network with both passive and active measurements to test for anomalous behavior. Additionally, the data aggregated from the distributed end-points in the network should be automatically analyzed to pin-point the location of the fault as well as diagnose the root cause of the fault. This information should be provided to the semi-skilled network administrator in a human readable form with concrete troubleshooting steps to perform.

Currently, in our existing deployments, we periodically initiate reverse ssh tunnels from the wireless routers to our server in Berkeley to collect a high level periodic health summary of each router node in the network. An alternate solution is to have a completely *orthogonal communication channel* like GSM/SMS. They provide a backup path for rare situations where a remote reboot is required, but are expensive and assume some form of cellular coverage.

4.5 Planning and Deployment

Planning of WiLD networks needs much more careful consideration compared to mesh networks with omnidirectional antennas. Since WiLD links traverse long distances, they require line of sight for operation and this usually implies towers at each end. As the towers take a substantial part of the total cost of the network, the challenge is to select the locations of sites and the links so that the overall cost of the towers is minimized (determined by the heights of the towers). Site selection is also influenced by the presence of external WiFi interference, as well as interference from the nodes which are part of the WiLD network. WiFi interference from the nodes within the network as well as from the external sources can be minimized be judiciously selecting the transmit power of the nodes. By over-provisioning the signal at the receiver, capture effect can be used to eliminate most WiFi interference.

An additional significant problem in the deployment of WiLD networks is in performing the manual careful alignment of the directional antennas that is essential for the working of every long distance link. This is complicated by the fact that factors like wind and wear and tear of towers can cause the antennas to misalign over time. In this respect, electronically steerable antennas can be used for automatic alignment. The open research challenge lies in devising efficient algorithms to discover peer nodes and maintain alignment using continuous adaptation over time.

5 NON-TECHNICAL CHALLENGES

While deploying wireless networks in developing countries we encountered a variety of non-technical problems. These deployments present much larger installation, maintenance and servicing costs, due to lack of local technical expertise, equipment availability and logistics. Consequently, there is a need for production-quality solutions, and not just research prototypes. The hardware and software must be robust, user friendly, and simple to install, maintain and manage. Local partners must be trained as well. Our group has learned these lessons the hard way in India and Ghana.

Another barrier is local telecommunication regulation, which is hindered by limited technical staff, “imperfect” government, and the presence of local incumbent monopolies. Some of the problems we encountered are: restrictions on using VoIP (favoring local telecom monopolies), licensed or even restricted frequency bands that are unlicensed everywhere else in the world, and unregulated wireless usage resulting in significant same-band interference.

6 CONCLUSION

We argue the need for concerted research efforts to develop cost-efficient networking solutions for providing connectivity to regions with low user densities. To this end, we examined various wireless options and their suitability, and explored WiLD networks as a promising option. By taking a broad view of the problem, we found challenges at essentially every layer of the network and thus a range of areas for new research.

REFERENCES

[1] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris. Link-level Measurements from an 802.11b Mesh Network. In *ACM SIGCOMM*, 2004.

[2] P. Bhagwat, B. Raman, and D. Sanghi. Turning 802.11 Inside-out. *Hotnets-II*, 2004.

[3] J. Bicket, D. Aguayo, S. Biswas, and R. Morris. Architecture and Evaluation of an Unplanned 802.11b Mesh Network. In *ACM MOBICOM*, 2005.

[4] S. Biswas and R. Morris. Opportunistic Routing in Multi-Hop Wireless Networks. *Hotnets-II*, 2003.

[5] Dharamsala Wireless-Mesh Community Network. http://www.tibtec.org.

[6] International Telecommunications Union. World Telecommunications/ICT Development Report. 2006. http://www.itu.int/ITU-D/ict/publications/wtdr_06/.

[7] S. M. Mishra, J. Hwang, D. Filippini, T. Du, R. Moazzami, and L. Subramanian. Economic Analysis of Networking Technologies for Rural Developing Regions. *Workshop on Internet and Network Economics*, 2005.

[8] R. Patra, S. Nedevschi, S. Surana, A. Bakre, E. Brewer, K. Fall, and L. Subramanian. Achieving High Throughput in Long-distance Wireless Networks. *Intel Technical Report*, 2006.

[9] B. Raman and K. Chebrolu. Design and Evaluation of a new MAC Protocol for Long-Distance 802.11 Mesh Networks. In *ACM MOBICOM*, 2005.

[10] Wavesat. WiMax 3.5GHz mini-PCI Reference Design Kit. http://www.wavesat.com/products/mini-pci.html.