

# 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

Most of the current research efforts on wireless networks have focused on enhancing network connectivity within the urban areas of the developed world, and thus assume relatively high user densities. Many regions around the world with low user-densities do not have good connectivity solutions today, including rural areas in both developed and developing countries. In this paper, we make the case for research on new appropriate wireless technologies that can provide low-cost, rapidly deployable connectivity solutions for the 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 Wifi-based Long Distance (WiLD) networks in India, Ghana and the Bay Area in the US.

## 1 INTRODUCTION

Today, the evolution of networks in the developing world is taking quite an alternate route from the traditional types of networks we observe in the industrialized world. If one were to take a stroll around Nairobi or Dar-es-Salaam in East Africa, one would be amazed to see 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. Many African countries see better opportunity in wireless options for regions that have low penetration of fiber and other wireline connectivity solutions. For example, most of them now have much higher cellphone penetration rates than fixed-line penetration [5]. Also, many African countries, including all in East Africa, still rely on satellite links for Internet connectivity; here, satellite-based ISPs (*e.g.*, SimbaNet [9]) use a wireless distribution network to extend coverage from the satellite tap across a specific region (city or district).

The primary reasons for such a boom in the use of long-range wireless networks within developing countries are:

**Low cost:** Excluding tower costs, the cost of establishing a fully functional point-to-point WiFi link that can deliver over 1 Mbps across a 10–50 km distance is roughly US\$1000, compared to US\$1000/km just for *dark* fiber. Satellite solutions are also quite expensive, typically US\$2000 per month for 1 Mbps connectivity. WiMax and traditional microwave technologies are much more expensive than WiFi but more affordable than fiber.

**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.

**Decentralized evolution:** The initial capital expenditure (CapEx) essential for a local entrepreneur to establish a wireless distribution network (using WiFi) for access purposes is typically less than US\$30,000 (excluding spectrum costs for WiMax which can be high). This allows for decentralized rapid evolution of such networks by local entrepreneurs. Fiber-based solutions require a high CapEx with the risk of a very low return of investment in low income regions.

**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 [6].

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. Some of the early work by Bhagwat *et al.* [2] and Raman *et al.* [8] 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

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

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 various layers of the networking stack. In this paper, we elaborate on these challenges and describe some of our initial efforts towards addressing these challenges. We also document some of 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, to reduce the cost per user. Low density environments, by definition, have users spread around in a region over long distances except in specific pockets where users may be clustered around a small locality (*e.g.*, a village). Even in clustered environments, the density of users (devices) will be lower than urban environments especially because shared terminals are often used in Internet cafes or kiosks to amortize the relatively high cost of computer terminals. In such environments, point-to-point wireless connectivity solutions are much more viable than traditional connectivity solutions.

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, but remains the most common solution.

Cellular networks, including GPRS and CDMA, depend on expensive base stations that are amortized over many users. In low-density regions, such base stations simply do not cover enough users to be economical. The expectation that cellular solves the connectivity problem for developing regions is thus somewhat 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. Thus, many cellular providers incur losses in low user-density regions and partially recoup these losses by either charging very high usage rates or imposing a universal service charge on all users.

Finally, 802.11 mesh networks that assume high user density [4] have received a lot of research attention lately. They are a flexible, low-cost means of providing network coverage by using a dense deployment of access points (APs) with omni-directional antennas, but only for small areas (few  $km^2$ ). In providing coverage to larger areas, mesh networks suffer from two basic problems. First, as the network grows, an increase in the number of APs with omni-directional 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.

The fact that existing connectivity solutions are inappropriate for low user densities motivates the need for research in determining appropriate connectivity solutions for such environments. Two specific solutions have been proposed in this context: WiFi-based Long Distance (WiLD) networks [6, 8] and WiMax [10].

WiLD networks represent a very cheap off-the-shelf solution to the connectivity problem where existing 802.11 protocols can be tailored to work over long-distances using high-gain directional antennas instead of omnidirectional antennas. The highly directional transmission reduces inter-neighbor interference, and also extends the range of the links over long distances. WiLD networks maintain the same cost advantage and unlicensed spectrum as mesh networks, but

do not succumb to the degraded throughput of regular 802.11 mesh networks due to a large number of intermediate hops.

An alternative to WiLD networks is WiMax [10] which has been primarily designed for long distances. The WiMax architecture is point-to-multipoint and uses a single high-capacity base station, serving several users within an area that spans 10–50 km in radius. The primary limitation for low-density regions is the relatively high cost of a base station, as discussed above.

Despite these shortcomings of WiMax, it does present many strengths over a WiFi-based approach at both the PHY and the MAC layers. WiMax has a better physical coding layer and can better deal with multipath effects. We believe that WiMax is appropriate in urban environments in developing regions with medium user densities. Even in such environments, WiLD networks present a natural evolutionary step before deploying a WiMax network. However, we believe that WiMax is inappropriate for relay networks where the need is to cover much longer distances than 50 km, the current reach of either of these technologies.

Although we believe WiMax has merit, our current research work has primarily focused on WiLD networks. Next, to motivate the research challenges that arise in WiLD settings, we briefly describe how the characteristics of WiLD networks differ from mesh networks.

## 2.2 WiLD vs Mesh networks

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 a world where there are no external WiFi interfering sources, one can engineer a WiLD network such that there is no hidden terminal problem, which is not the case for mesh networks. However, in a world where a WiLD network is deployed in the presence of external interfering sources (access points within the neighborhood), the hidden terminal problem can be much worse in a WiLD network than in a mesh network. This is due to two factors: 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 transmission. However, in a mesh network with overlapping transmission regions among neighbors, the fraction of external interfering sources that act as hidden termi-

nals is much smaller. Due to long propagation delays, even external interfering sources within the range of a directional transmitter can interfere by detecting the conflict 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.<sup>1</sup> 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 high-throughput point-to-point links and routing in WiLD networks resembles traditional routing protocols.

<sup>1</sup>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.





**Figure 2:** WiLD deployment for an eye hospital in rural South India. Theri is the main hospital, all endpoints are rural eye care clinics, and the rest are relay points. All link distances are in kilometers

### 3 EXISTING DEPLOYMENT

Currently, we have deployed several WiLD networks in India (a 9-link topology shown in Figure 2)), Ghana (5 links) and the Bay Area in the US (7 links). We use these testbed deployments to understand the different research issues and also to implement and evaluate the solutions to those challenges. The WiLD network in India (Figure 2) 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 between its different campuses. Distances of our WiLD links vary from 10–80km, with relays installed where there is no line of sight due to geographical limitations.


The nodes of our networks are based on a small low power single board computer (SBC) which have a 266 MHz Pentium-based chip and 128 MB RAM and support up to 3 wireless cards. 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 also 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.

### 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) Handling high-loss variations; (2) MAC layer; (3) QoS Provisioning; (4) Routing and fault-tolerance; (5) Remote upgrade and maintenance; (6) Network planning. Associated with each of these challenges, we describe some of our early efforts to address them.

#### 4.1 Handling high-loss variations

Across all of our WiLD networks, the presence of external WiFi interference results in very high loss rates on WiLD



**Figure 3:** Loss variation over time across channels 1 and 11

links. Furthermore, due to the long distances, the extent of interference could be very different at the two ends, making WiLD links asymmetric in nature. Also, in presence of external WiFi interference it is common to have links with loss rates fluctuating between 5–80% over short time scales. The primary research challenges that arises is: In the face of high loss variations, how can one achieve predictable and good performance along a single WiLD link?

**Mitigating losses:** To mitigate the effect of packet losses triggered by external WiFi interference, we propose a combination of two simple mechanisms: (a) channel hopping; (b) adaptive FEC. Figure 3 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. By sensing external interference on a channel, the end-points in the network can collaboratively decide to switch to an alternate less congested channel. However, as seen from the above figure, even in absence of long bursts there still exists a residual 5–8% loss. To mitigate this loss, FEC could prove to be a useful technique. 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, especially given that loss-rate distribution is time-varying. This makes the problem of determining the appropriate FEC recovery mechanism a challenging one.

#### 4.2 MAC-layer Challenges

The 802.11 protocol is known to suffer from fundamental limitations when used in long-distance multihop networks [8]. These problems 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 orig-

inating from a single node operate on the same channel or even within two 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 Bulk ACKs:** The above challenges of the stock 802.11 MAC protocol motivate the need for a TDMA-based MAC protocol that synchronizes the transmissions from the endpoints. In addition to a single point-to-point link, in presence of a point-to-multipoint topology Raman et al. [8] propose having *simultaneous send* and *simultaneous receive* to eliminate interference.

Once using TDMA, the stop-and-wait recovery mechanism of 802.11 is unsuitable. We implement a sliding-window based flow-control approach, in which the receiver acknowledges a set of frames at once (bulk ACKs).

Figure 4 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 channel emulator that allows us to precisely specify the delay between the sender and receiver. We can see 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 provides sustained high throughput even at very long distances.

**TDMA Slot Challenges:** Allocating slots in a multihop network is non-trivial, but slot allocation to achieve optimal throughput across all the WiLD network links, with the constraint of simultaneous transmit and receive, is a challenging problem for general graph topologies. However, it can be shown that for bipartite graphs, we can always find such a slot schedule.

As we see later in section 4.3, the TDMA slot schedule can be driven by traffic demands. The challenge then is to devise either a distributed or centralized mechanism for configuring slot schedules in the network.

### 4.3 Supporting Quality of Service (QoS)

We envision a WiLD network to be used by a variety of different applications with different QoS guarantee requirements. One such example is the video conferencing application that we use in our rural telemedicine system in India.




Figure 4: Comparison of WiLD MAC and stock 802.11 MAC

Unlike the Internet, which only supports a best-effort model, here we have the flexibility to deploy any QoS mechanism proposed in the literature. Currently, we are investigating two problems in supporting QoS:

**Optimal TDMA slot schedule:** The use of the TDMA MAC protocol with fixed slot sizes provides us a way to bound parameters like delay, jitter and bandwidth. However, in reality, most of the applications have an asymmetric traffic pattern which in turn would impose additional constraints on the TDMA slot size selection and allocation in the network. An example of an asymmetric traffic pattern is our telemedicine deployment; high resolution video from all the patients has to be transmitted to the doctor, however only a low resolution video stream is required in the opposite direction. This requires the use of larger slot sizes in the patient to doctor direction for all the links in the path. However in a single channel network, such asymmetric allocation is possible only for single hop networks. With additional non-overlapping channels, as proposed in [7] this constraint can be relaxed. This points us to an optimization problem where given a set of possibly asymmetric traffic demands from all the nodes and link synchronization constraints, we have to jointly determine the slot size and the slot schedules for each link such that application demands are met.

**Traffic priority classes:** The low processing power (266 MHz) and memory constraints (128 MB) of the router boards rules out any form of strict/statistical QoS guarantees which would require intermediary nodes to maintain per-flow state, track per-flow usage and needs an admission control process. We introduce the notion of *traffic priority classes* where at each hop, we implement mechanisms to provide per-hop QoS guarantees for every class. Additionally, we use the routing protocol to compute the achievable QoS parameters on a per-class granularity. Based on the achievable QoS properties for every class, application traffic can be tagged with the corresponding class that meets its QoS requirements. This notion of traffic priority classes is a simple but powerful QoS mechanism that provides QoS guarantee information at a coarse granularity without the need for per-flow state and per-flow traffic accounting.

### 4.4 Routing and Fault Tolerance

Our initial WiLD network deployments have all been tree topologies making routing a trivial task. However, we recognize the need to make these network fault tolerant given that any antenna or link issue will disconnect that subtree. We experienced several hardware failures in our testbed which ranged from corruption of the flash memory, lightning strikes, rain water clogging around the RF cables and attenuating the signal strength etc. To balance the cost versus fault tolerance tradeoff, we are exploring two options. The first is to install additional redundant links in the network and use any standard routing protocol for determining routes within the network. Alternatively, a longer-term research plan is to augment the testbed with low-cost electronically steerable antennas (switched parasitic), which allow us to multiplex

one antenna for different physical links. Similar to phased-array antennas, these antennas help fault tolerance both by automating realignment and by creating back-up links dynamically. A related open problem is how to route traffic in a network with reconfigurable links.

#### 4.5 Remote upgrade and management

To reduce the operational cost of WiLD networks, it is critical to minimize the need for trained personnel. Besides the shortage of personnel, long distances of the links makes accessing the endpoints an all day process. Consequently, there is a critical need for remote management.

The first requirement for remote management is *continuous data collection* from the routers to a central database. We periodically initiate reverse ssh tunnels from our routers in all our deployments to our server in Berkeley. The tunnels work through firewalls and over satellite links that require initiation from the client side.

A *safe upgrade mechanism* is 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.

However, in many cases this is not adequate. For example, in one case a software bug caused the sshd daemon to crash, causing the end-point to be disconnected. In such a case, a completely *orthogonal communication channel* is required to reboot the router box remotely. Besides “sneaker net” we are looking into GSM/SMS as an expensive backup path, but this obviously assumes coverage.

#### 4.6 Planning, Deployment and Scalability

Because WiLD networks have long distance links that require line of sight, the site selection becomes very important. In addition to making sure that the selected sites are sufficiently high to achieve line of sight, we must also consider the level of external WiFi interference at each endpoint. Deployment of long distance links with directional antennas also requires careful alignment from both end points. On the other hand, site selection in urban mesh deployments need not be planned because they can use omnidirectional antennas at short distances that do not need explicit alignment.

Network scalability is also challenging for WiLD environments. As the size of the network increases, the hop count also increases. In a TDMA MAC protocol, this translates to increased end-to-end latency, which puts a premium on short slots (at least for voice traffic) and thus better slot synchronization across the network, which is an open problem.

To simplify these issues, we are also exploring the use of low-cost electronically steerable antennas, which eliminate the need for careful alignment and make it easier to add new links to the network without having to physically add new antennas and radios.

## 5 NON-TECHNICAL CHALLENGES

Besides the technical challenges already mentioned, we have encountered a variety of non-technical problems in deploying wireless networks in developing countries. Deploying wireless networks in developing regions presents 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 that are not really interested in progress. 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 from other technologies.

## 6 CONCLUSION

We argue that, albeit the increasing interest and need, there is a lack of cost-efficient networking solutions for connecting regions with low user densities. We believe that concerted research efforts towards developing such solutions is required at this point. To this end, we examined various wireless options and their suitability, and explored WiLD networks as a promising alternative. 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] International Telecommunications Union. World Telecommunications/ICT Development Report. 2006. [http://www.itu.int/ITU-D/ict/publications/wtdr\\_06/](http://www.itu.int/ITU-D/ict/publications/wtdr_06/).
- [6] R. Patra, S. Nedeveschi, S. Surana, A. Bakre, E. Brewer, K. Fall, and L. Subramanian. Achieving High Throughput in Long-distance Wireless Networks. *Intel Technical Report*, 2006.
- [7] B. Raman. Channel Allocation in 802.11-based Mesh Networks. In *IEEE INFOCOM*, April 2006.
- [8] B. Raman and K. Chebrolu. Design and Evaluation of a new MAC Protocol for Long-Distance 802.11 Mesh Networks. In *ACM MOBICOM*, 2005.
- [9] SimbaNet. <http://www.simbanet.com>.
- [10] WiMAX forum. <http://www.wimaxforum.org>.