

A Low-Cost Efficient Wireless Architecture for Rural Network Connectivity

1 Introduction

Many rural regions around the world, especially in developing regions, do not have good connectivity solutions which are economically viable. As a result, many of these regions remain disconnected from both the rest of the world and from progress in general. In this proposal, I will describe the design of *WiFi-based Rural Extensions (WiRE)*, a new wireless network architecture that can provide connectivity to rural regions at extremely low costs. The WiRE architecture is tailored for the typical rural landscape in several developing regions, in which the population is spread across small but scattered rural regions (less than 1-2 sq kms) within 100-200 kms of the city. WiRE is designed to be a wireless distribution network that extends connectivity from the city to each village.

The WiRE architecture has largely been inspired by my prior work on WiFi-based Long Distance (WiLD) Networks [42, 62, 35, 54, 64, 34], a low cost point-to-point network connectivity solution that provides very high bandwidth (typically 6 – 10 Mbps) over very long-distances. While prior work on WiLD networks [48, 5, 42, 62, 35] has made significant progress in the design of high-performance MAC layer solutions, we still lack a vision of how to design a comprehensive, low-cost, rural connectivity architecture that can efficiently support a wide-range of applications. It is this goal that I wish to achieve in the WiRE network architecture design. To realize this architectural vision, we need to address several challenges at various protocol layers including the MAC, network, transport and the application layers. We will first motivate the need for low-cost connectivity before we outline these challenges in greater detail.

Motivation: Need for Low-Cost Rural Connectivity

As of Internet World Stats 2007 [28], the Internet penetration in North America is 69.7% of the population compared to 10.7% in Asia and 3.6% in Africa primarily restricted to urban areas. The fundamental problem in connecting rural regions is economics [34, 9, 8]. None of the traditional wire-line connectivity solutions (fiber, broadband and dial-up) are economically viable for such regions over at least the next decade due to low user densities [34, 15, 9]. Satellite networks provide great rural coverage but at very high costs: the ISP rate for 1 Mb/s of satellite connectivity in Africa exceeds \$3000/month [3].

In recent years, many developing countries have undergone a cellular revolution with a significant penetration of cellular networks in rural areas [26, 27, 23]. Commercial wireless broadband networks based on GPRS [55], WiMax [70, 22] and CDMA [36] technologies are also being widely deployed [36, 27, 26]. While a sizable fraction of the rural population owns cellphones for telephony services in Africa and Asia [46, 26, 71] the network usage is limited due to exorbitantly high usage costs, roughly ranging from 10 cents to \$1/min [2, 24, 23, 25]. Given that a large majority in rural areas earns less than a few dollars/day, these costs are unaffordable.

For any connectivity solution to be economically viable in rural regions with low-user densities, it is essential to have small per-user setup cost and minimal recurring costs [62, 34]. Networks with a base-station model, such as WiMAX, and cellular networks like GPRS and CDMA, have an asymmetric design philosophy where expensive base stations (costing roughly \$10K - 100K depending on range and capacity) are amortized by large number of cheap client-devices over many users [62, 34]. Operational costs of these networks in rural areas are also high [34, 64] due to: (a)

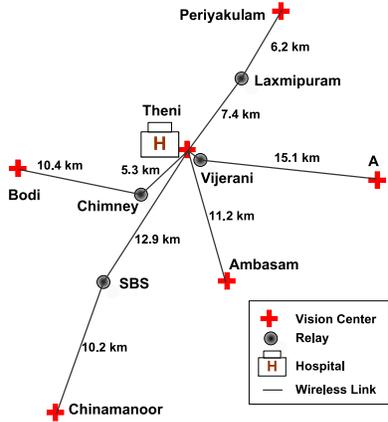


Figure 1: Aravind network

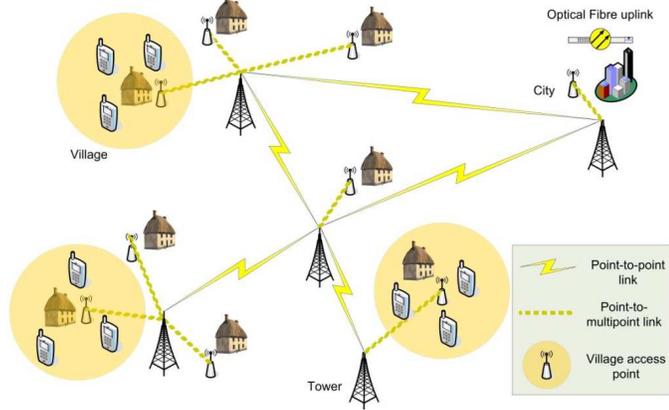


Figure 2: The WiRE network architecture

significant power consumption to cover large areas; (b) the need for backup power due to lack of reliable grid power; (c) high cost of physical security for expensive equipment. Together, these costs make existing cellular and wireless broadband services not viable in regions with low user densities. Hence, the expectation that cellular solves the connectivity problem for rural developing regions is thus somewhat of a myth!

Prior Experiences on Rural Connectivity: Prior to this proposal, I was involved in the design, implementation and deployment of WiLD networks [42, 35, 62, 64], a point-to-point WiFi connectivity solution that can provide 6 – 10 Mbps over 50 – 100 kms at very low costs. WiLD networks are extremely low-cost due to the use of unlicensed WiFi spectrum and leverage off-the-shelf low-cost and high available commodity hardware. To achieve high throughput in WiLD networks, we designed a new MAC protocol called WiLDMAC that addressed many of the critical shortcomings of the conventional 802.11 protocol in long-distance environments. WiLDMAC also improved over 2P [48], the only previously known protocol for WiLD environments, by being able to achieve high throughputs over highly lossy network environments (20-60% loss rates). Recently, we developed JazzyMAC [35] that significantly improves over both 2P and WiLDMAC to achieve near-optimal throughput in multi-hop settings. WiLD networks have become increasingly popular in the last few years with deployments in nearly 15 – 20 developing countries. Our WiLD network deployment for Aravind Eye Hospitals [21] in South India (illustrated in Figure 1), the largest eye hospital in the world with over 2 million patients per year, provides telemedicine services to over 50000 patients per year in 13 rural vision centers [64, 63]. Aravind Eye Hospitals recently obtained a Gates Foundation grant to expand their network to cover 500000 patients per year. We also broke the world record for the longest point-to-point wireless link achieving 6 Mbps over 384 kms in Venezuela [64]. Other WiLD deployments include the Digital Gangetic Plains project [48], Fractal [11], the AirJaldi [37], Aravind networks [64] and the Akshaya network [66, 41].

Other Connectivity Approaches: There have been a few recent WiFi-based solutions [51, 43] which have proposed MAC extensions for point-to-multipoint networks. We discuss these in greater detail in Section 3. An alternative economically viable connectivity approach is to use Delay Tolerant Networks (DTN) [16, 31, 30, 19, 44, 68] which leverage physical transportation systems to transport bits to and from the rural regions. DTNs by definition are not suited for interactive applications which is an important focus of the WiRE network architecture.

Research Agenda

WiRE uses a network structure (illustrated in Figure 2) that is significantly different from the traditional cellular, WiMAX, WiLD, and wireless mesh network models. For comparison, an example WiLD network is illustrated in Figure 1. Unlike the cellular network philosophy of providing broad network coverage, WiRE provides focused coverage within rural regions with little coverage outside. The network structure of a WiRE deployment is optimized based on the topography and the spread of rural regions. To efficiently reach out to sparsely spread out rural regions, WiRE uses a combinational network structure with four important components: (a) point-to-point network links; (b) point-to-multipoint network links; (c) local distribution mesh networks; (d) cellphones as end-devices in addition to PCs and kiosks. Given that land and tower costs are expensive, this network structure explicitly attempts to achieve maximum distribution with a small set of towers. While point-to-point links with highly directional antennas provide a high bandwidth backhaul that can cover long distances, point-to-multipoint links with sector antennas provide efficient distribution capabilities within shorter regions and mesh networks with omni-directional antennas are primarily used in small localities to provide coverage.

The use of cellphones in WiRE as end-devices is very important especially given the mass penetration of these devices in rural regions [23, 20, 26]. Many rural deployments which uses PCs, kiosks and other types low-cost computing devices have miserably failed due to complex user interfaces and the sheer lack of need [32, 33]. Nearly 100,000 kiosks deployed in rural India are being sparingly used [33, 17]. Cellphones, on the contrary, have gained significant acceptance among rural communities as they are simple to use, and are also accessible to the illiterate user community. Cellphones, in rural communities, are predominantly used as a voice interface [20], which is the main reason why telephony services remains the most important killer application in these environments [41]. Also, according to a recent study [56], most phones in the near future should have inbuilt WiFi capability [56].

To realize the WiRE network architecture, we need to address several challenges across different network layers:

MAC layer Challenges: Designing a unified high-performance MAC layer for the WiRE architecture is a challenging problem. WiRE uses three different types of network links (point-to-point, point-to-multipoint, omni) each with completely different network characteristics all of which operate on the same frequency band. To achieve this, we need to address several specific challenges: (a) develop new MAC protocols for point-to-multipoint; (b) handle complex interference interactions; (c) adapt to highly lossy links; (d) intelligent channel assignment; (e) adapt to fluctuating traffic demands.

Robust Network Design Challenges: Designing robust and reliable rural wireless networks is an arduous task due to a variety of factors: failure of cheap devices, lack of good and stable power sources, lack of good local support, nodes in hard to reach locations. We intend to build a suite of solutions including: (a) efficient topology design and routing to handle frequent outages while minimizing the need for new towers; (b) efficient low-cost power solutions including solar power solutions; (c) network management tools that can aid in configuration, fault diagnosis, monitoring and remote management/upgrades.

Application Specific Challenges: Apart from traditional set of Internet applications, we require WiRE to support specific applications of prime importance to the rural sectors including telephony, telemedicine, distance learning and mobile banking services. To enable these applications, we need to address two broad set of challenges. First, for telephony, video-streaming and video conferencing applications, we need to address the challenge of providing statistical end-to-end QoS guarantees in the face of fluctuating loss and available bandwidth variations. Second, to enable telephony and

secure mobile transactions, we need to be able to support mobility of cellphones within the WiRE network and also provide a secure naming mechanism based on the unique identity of cellphones.

Intellectual Merit: Computer science, as a field, has paid very little attention to important technical challenges that arise in the developing world. This proposal will significantly advance the understanding of networking challenges across all protocol layers in the developing world. This proposal will also advance the understanding of several fundamental aspects of wireless network design including interference, high-throughput, QoS, routing and transport issues.

Broader Impact: The WiRE architecture has the potential to significantly transform the rural landscape by providing network connectivity at very low costs and impact billions in rural regions who still remain disconnected from the rest of the world. We intend to do pilot deployments of the WiRE architecture in India, Ghana and South Africa where we work with well established local partners who have the capacity to reach out to millions in rural communities.

2 WiRE Network Architecture

In this section, we describe the WiRE network architecture and discuss important real-world challenges in deploying rural wireless networks based on our experiences. Figure 2 describes the basic WiRE network architecture. Unlike the traditional cellular model of providing broad coverage, the design philosophy of WiRE is to provide focused coverage within specific rural regions where connectivity is most required. The WiRE architecture has six important network components:

1. *wireless nodes* which are low-power single board computers that have the capability to support multiple wireless cards for different network links.
2. *point-to-point* links using highly directional antennas to provide network connectivity over long distances in the range of 50 – 100 kms.
3. *point-to-multipoint* links using sector antennas to distribute connectivity to multiple endpoints within relatively short distance lasting a few kilometers.
4. *multi-radio mesh* links using omni-directional links to extend wireless coverage within small local regions.
5. *cellphones* or *low cost computing devices* with WiFi-enabled interfaces that can act as end-devices.
6. *large local storage* of at least a few GB at each local wireless node to perform in-network optimizations for applications as well as store-and-forward intermittent operations in the event of a network outage.

The basic network structure of WiRE is a natural extension of WiLD networks, which I had worked on for the last three years in collaboration with Prof. Eric Brewer and his students at UC Berkeley. The focus of the WiRE network architecture is much broader in scope than WiLD networks. WiRE focuses on challenges across different protocol layers to build a complete solution for rural connectivity including support for a wide range of applications. Even from the MAC layer perspective, WiRE operates in a combinational wireless environment of point-to-point, point-to-multipoint and omnidirectional links, each of which have completely different MAC needs and interference characteristics.

WiRE has the flexibility to operate in any frequency spectrum. However, for practical and cost-related constraints, we choose all network links in WiRE to operate in the WiFi frequency band space (802.11 a/b/g). WiFi cards are cheap and highly available, enjoying economies of scale. The typical cost of a network link excluding the cost of the tower can be brought down to approximately \$600 (excludes the cost of tower) with no recurring cost. Since WiFi is classified as unlicensed spectrum in most countries, a WiRE network provider does not need to pay spectrum

costs which can be significantly high for other licensed frequency bands. The use of WiFi also makes WiRE easy to deploy and experiment with given that the entire network is composed of cheap off-the-shelf components. 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. All these factors promotes decentralized evolution of WiRE where a grass-roots organization can easily deploy a WiRE network without any dependence on a telecom carrier.

In WiRE, every wireless router uses a low power single board computers (SBC). The current typical configuration of a SBC has a 466 Mhz processor, 256 MB RAM and can support upto 4 wireless cards; in addition, we equip each wireless node with a large local storage to enable in-network application-level optimizations and also perform store-and-forward routing in the face of network disruptions. For radios, we use off-the-shelf high power 802.11 a/b/g Atheros cards with up to 400 mW transmit power. For long-distance point-to-point links that can traverse between 20 – 150 kms, we use high power radio cards with high-gain parabolic antennas with a gain factor of upto 30dBi. The highly directional nature of the wireless beam allows us to have several point-to-point links at a given node given multiple radios. For connecting many specific locations within a certain distances of upto 20 kms, we use a point-to-multipoint topology where a single wireless router can serve as a base station for several clients. Depending on the bandwidth requirements for each client, each node can serve upto 30 clients. The multi-radio mesh nodes within a local region are used for extending the connectivity within a specific region; these links in outdoor settings with 200mW cards can cover between 0.5 kms to 1 km. If necessary, depending on the topography of a region, we may require several multi-radio mesh nodes to completely cover a region.

The end-devices in WiRE can be either static computing devices such as PCs/kiosks or mobile devices such as cellphones. It is essential for cellphones to form an integral part of the WiRE architecture due to three factors. First, cellphones have such high penetration levels in rural developing regions that make them natural candidates for end-devices. Many of the new generation of low-cost cellphones come with inbuilt WiFi capabilities making them suitable for WiRE. Second, cellphones are extremely simple to use and are accessible to even illiterate users in rural areas. Finally, the open source movement for cellphones [65, 18] has radically transformed the set of new applications that can be deployed for these devices.

Rural Specific Applications: Apart from the traditional set of Internet applications (web browsing, Email etc), we require the WiRE architecture to enable specific services which are very important in the rural context. Many rural regions have remained disconnected from the rest of the world that the state of several essential services such as education, healthcare and financial services have remained abysmal in these regions. Providing connectivity alone is not sufficient; we require WiRE to support the appropriate set of applications to enhance essential services in rural areas. We have identified four such applications which we deem as essential for WiRE to support: (a) telephony services for cellphones; (b) telemedicine and teleconsultation services for improving rural healthcare; (c) interactive distance learning to improve rural education; (d) mobile banking to promote rural financial services. In order to enable each of these application, we need to address specific challenges in the network and transport layer. We outline these challenges and our initial approach to address them in Section 5.

2.1 Challenges Building Rural Wireless Networks: Lessons Learnt

In our experiences in deploying wireless networks in rural areas, we faced several challenges due to the ground realities of these regions. We document our experiences in our prior work [64]. We illustrate the important lessons we learnt from our deployment and their implications for the WiRE

architecture:

High loss rates: In many of our deployments, we found WiFi links to have high-loss rates ranging from 2% to as high as 50 – 60% due to poor signal quality, antenna misalignment or external interference (in semi-urban areas). Hence, the MAC protocols should be designed to handle high loss-rates.

Tower costs: Tower costs typically are much more than network equipment costs; the subsidized cost of a 30-40m tower in India was \$2500. Renting space from existing towers is also an expensive proposition. Hence, we need to minimize the number of towers needed in the topology.

Unreliable power: Many rural regions have interrupted and erratic power supply with significant voltage fluctuations which often cause network components to regularly fail. We also found the use of batteries and UPSs to be ineffective. Certain rural tower locations do not have a nearby grid supply. While we have developed preliminary solutions for the power problem such as a low-voltage disconnect and a microcontroller-based solar power controller, much work needs to be done to improve the stability and reliability of power.

Network Management: Network faults were a common occurrence in many of our deployments primarily due to the failure of network components. In addition, the local operators lacked the expertise to repair faults and international travel is prohibitively expensive to repair specific faults. The lessons for network management are three-fold. First, it is essential to design the network with some redundancy to tolerate node failures (many existing WiFi deployments use tree topologies). Second, we need simple configuration and management tools to aid the operator to locate the source of faults. Third, we require backchannels (using cellular links) and remote management tools to perform remote upgrades and repairs in the face of faults.

3 MAC Layer Challenges

The overarching MAC challenge in WiRE is to develop a unified MAC protocol that is configurable to different network settings and which can provide high throughput and predictive performance in multi-hop settings. While there have been several advances in these individual networks [6, 14, 38, 50, 4, 7, 49, 5, 48, 42, 62, 69, 1, 29, 51, 43], a unified approach has not been explored. Achieving high throughput in WiRE is a challenging problem due to a variety of factors:

Variable network characteristics: WiRE operates in three different types of network settings (point-to-point, point-to-multipoint, omni-directional) with completely varied physical and MAC layer characteristics. In addition, given the limited number of non-overlapping channels in 802.11b and inherent limitations in using the 802.11a frequency band over long distances, interference across these network links within WiRE is unavoidable.

Limitations of 802.11: The conventional 802.11 protocol is known to have fundamental shortcomings when applied to long-distance environments [48, 5, 12, 47, 42, 62, 35, 54]. First, CSMA/CA is a fundamentally flawed idea over long-distance links since one end-point cannot quickly sense packet transmissions from the other end-point thereby resulting in high packet collision rates. Second, 802.11 MAC uses a simple stop-and-wait protocol that substantially decreases channel utilization. If the ACK timeout is lesser than than the link RTT, the sender unnecessarily retransmits the packets.

Multiple Link Interference: Inter-link interference occurs when two adjacent 802.11 point-to-point links operating in the same channel or over-lapping channels interfere with each other despite transmitting in different directions.

Channel Loss Variability: In real world deployments, we found that WiFi links (both directional and omni) observe very high channel loss rates that fluctuate significantly with time. We observed

sustained high loss-rates of 50 – 60% on certain long-distance links [42, 62, 54].

Hidden Interference: We observed a peculiar type of interference in our network due to combinational nature of the WiRE architecture; a packet transmission from an omnidirectional antenna can interfere with a neighboring directional/sector antenna despite the receiver’s inability to sense the transmission [54]; the only way to detect such interference patterns is to correlate sending times of neighboring nodes with spikes in error rates. This form of interference is different from the standard hidden terminal problem.

3.1 MAC Design for Point-to-Point and Point-to-Multipoint Links

Point-to-Point: I was involved in the development of WiLDMAC, a modified MAC protocol that addresses the limitations of the conventional 802.11 MAC to achieve high throughput in long-distance settings. WiLDMAC [42] also addressed an inherent limitation in the previous proposal 2P [48, 5] that was not tailored to handle highly fluctuating channel conditions. To address the CSMA limitation of 802.11, WiLDMAC uses a TDMA based approach which is based on fixed timeslots coupled with an implicit echo-based protocol across each link to synchronize transmissions and receptions between the end-points. To improve the channel utilization on each link at longer distances, we replace the stock 802.11 stop-and-wait protocol with a sliding-window based flow-control approach in which we transmit a batch of packets together in a TDMA slot without waiting for individual ACKs. For a node having multiple point-to-point links sharing the same channel, we implement an inter-link synchronization mechanism similar to 2P. This protocol ensures that adjacent links either a) send simultaneously in the transmit TDMA slot. or b) receive simultaneously in the receive TDMA slot. We can achieve simultaneous transmit if carrier sensing is disabled; and simultaneous receive if the signal separation between the two receivers is sufficient.

To achieve predictable multi-hop performance in the face of fluctuating loss conditions, it is essential to have a loss recovery mechanism that can hide the loss variability in the underlying channel. Achieving such an upper bound q on the loss-rate is not easy because the loss distribution that we observed on our links is non-stationary. We use a combination of two mechanisms - retransmissions and FEC to deal with losses. A retransmission based approach can achieve the loss-bound q with minimal throughput overhead but at the expense of increased delay. An FEC based approach incurs additional throughput overhead but does not incur a delay penalty especially since it is used in combination with TDMA on a per-slot basis. The retransmissions based approach uses bulk acknowledgments (bulk ACKs). A bulk ACK is sent from the receiver for a window of packets as an aggregated bit-vector acknowledgment for all the packets received within the previous slot. The FEC-based recovery mechanism requires the sender to proactively perform FEC based encoding across all the packets in a slot.

Our evaluation showed that WiLDMAC significantly outperformed the conventional 802.11 MAC even with best possible choice of parameters. Figure 3 shows the cumulative throughput of TCP flowing simultaneously in both directions for a single long-distance link (emulated using a channel emulator) illustrating the effectiveness of WiLDMAC with increasing link distance. In fact for a 65 km link in Ghana, WiLDMAC’s throughput at 5.5 Mbps is about 8x better than standard CSMA. To quantify the improvements of WiLDNet from inter-link synchronization, we perform TCP throughput measurements over a multiple hop topology. We can see from Figure 4, for same channel operation, the cumulative TCP throughput in both directions with WiLDMAC (4.86 Mbps) is more than twice the throughput observed over standard 802.11 (2.11 Mbps). In a recent result [35], we showed the fixed time-slot approach of both 2P and WiLDMAC results in significantly lower throughput than the optimal achievable throughput in multi-hop WiLD networks. To address this, we developed JazzyMAC [35], a variant of WiLDMAC with adaptive time-slots based on traffic

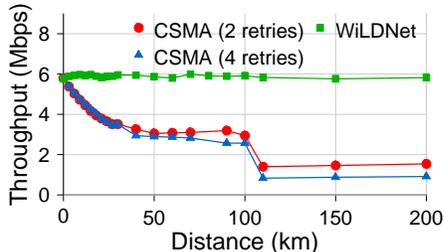


Figure 3: TCP flow in both directions for WiLDNet vs 802.11 CSMA. Each measurement is for a TCP flow of 60s, 802.11b PHY, 11Mbps.

Description (Mbps)	One direction	Both directions
Standard TCP: same channel	2.17	2.11
Standard TCP: diff channels	3.95	4.50
WiLD TCP: same channel	3.12	4.86
WiLD TCP: diff channels	3.14	4.90

Figure 4: Mean TCP throughput (flow in single direction and cumulative for both directions simultaneously) comparison for WiLDNet and standard 802.11 CSMA over a 3-hop outdoor setup Averaged over 10 measurements of TCP flow for 60s at 802.11b PHY layer datarate of 11Mbps.

conditions to achieve close to optimal throughput in multi-hop settings.

Point-to-Multipoint: Similar to point-to-point links, conventional 802.11 MAC style protocols are inappropriate for point-to-multipoint settings due to long distances. Recently, the SRAWAN [51] and the WiFiRe [43] projects proposed TDMA-based MAC protocols for point-to-multipoint WiFi networks. SRAWAN primarily explores a single-radio base station model and uses beacons from the basestation to synchronize the clients and a mix of Round-Robin and Weighted Fair Queueing to achieve QoS. WiFiRe seeks to increase spatial usage by synchronizing TX and RX from multiple radios but does not address the issue of how to achieve optimize allocation to clients. The WiMAX standard [70] also proposes a TDMA based MAC for supporting multiple clients but does not support any sort of synchronization across different radios. In addition, none of these approaches are well-suited to handle fluctuating channel conditions which can disrupt the protocol by dropping important protocol synchronization packets.

Based on our design of WiLDMAC and JazzyMAC, we are exploring the design of a point-to-multipoint MAC protocol that achieves near-optimal per-client throughput in the face of variable traffic demands and loss-rate fluctuations. The basic idea of our approach is to use adaptive time-slots in the TDMA-based protocol which continuously vary with traffic demand. To achieve synchronized transmissions over lossy channels, we use an implicit coarse-grained time-synchronization mechanism similar to WiLDMAC rather than exchanging synchronization packets or beacons. We also leverage the adaptive loss recovery mechanisms to efficiently recover from high packet loss rates. Other possible directions we are considering to explore in the future are: (a) use of multiple overlapping sector antennas for fault tolerance; (b) adaptive control of power, beam direction and beam width; (c) use of fast-switching electrically steerable antennas.

3.2 Towards a Unified, Adaptive and Auto-configurable MAC Layer

While one may envision designing specific MAC protocols for specific environments, in practice, we require a unified MAC that can be installed as a single software in all the nodes that can adapt and be configured to specific environments. Part of the challenge is that, operators who install networks in rural areas are not sophisticated enough to properly configure these networks. In large scale deployments, the network configuration and management becomes a much harder challenge as we have experienced in prior deployments [42, 21, 37]. The preliminary design of our unified MAC layer borrows ideas from several existing MAC layer protocols [50, 42, 35, 48, 69, 10, 51, 43] including WiLDMAC and JazzyMAC. We outline the key design choices:

Auto-configuration of network type: Each node tries to automatically detect the type of network on each of its radio interface based on three parameters: (a) number of neighbors; (b) RTT to each neighbor; (c) loss characteristics to each neighbor. Reliably estimating RTT at sub-millisecond scales is not simple; one way to estimate distance is to vary the ACKTimeout parameter, to determine at what point auto-retransmissions are triggered at the driver level. Using these parameters, we determine the network type and correspondingly disable or enable CSMA (only for omni-directional antennas).

Local time-synchronization + Distributed TDMA: Similar to WiLDMAC, we use a local time-synchronization mechanism combined with a distributed TDMA protocol to synchronize transmissions and reception at every node that uses a point-to-point or point-to-multipoint link. In WiRE, multiple nodes can be installed in the same tower within close proximity; to prevent inter-link interference we need to locally synchronize nodes within close proximity and use a synchronized TDMA slotting across all these nodes.

Inferring Hidden Interference: Hidden interference is not easy to determine. To do so, at a node with a directional or a sector antenna, we need to determine if the periods of high error rates with the transmission times of other nodes within a 2-hop neighborhood (direct neighbors cannot cause hidden-interference); if the correlation is high, then those nodes are candidate choices for hidden interference. In our WiLD deployments, hidden interference caused link loss rates of upto 80% [54].

Conflict Map: To deal with inter-link interference and hidden-interference (on directional links), we use the idea of Conflict Maps (CMAP) a recent work of Vutukuru et al [69] to determine the interference map and use this information to determine a potentially non-overlapping local transmission schedule.

Intelligent Channel Assignment: Non-overlapping channels are a scarce resource in the WiFi spectrum world; recent work by Chandra et al. [10] showed how one could easily derive channels with “adaptive frequency width”. We intend to explore using this approach in WiRE to achieve two properties: (a) significantly increase the number of non-overlapping channels in the system; (b) create different channels with variable widths and use “fatter width” channels for important point-to-point links and “leaner width” channels for point-to-multipoint links.

4 Robust Network Design Challenges

The design and operation of rural wireless networks raises many challenges which cannot be solved by just using high-performance equipment [64]. The key challenges we want to address are: (1) optimal design of network topology to decrease deployment cost, (2) increased component failure due to low quality power, (3) difficulty in doing fault diagnosis because of non-expert local staff and limited connectivity for remote experts, and (4) difficulty of frequent maintenance because of remoteness of node locations. All of these problems can be fixed by having higher operating budgets that can afford highly trained staff, stable power sources, and robust high-end equipment. But the real challenge is to find solutions that are sustainable and low-cost at all levels of the system.

4.1 Network Design

The key network design challenge is: *given a topography and the location of rural areas in a region, how do we design an optimal network topology that minimizes the number of towers and achieves a certain minimum level of network redundancy?* In addition, we need to consider the line of sight as an important issue since point-to-point links require line of sight for operation; this usually implies towers of a minimum height at each end. A variant of this problem was studied by Sen and Raman [52] for long-distance WiFi networks. Our problem varies from their problem definition in

two ways: (a) We need to optimize the network structure for a combinational network as opposed to a point-to-point network case; (b) We need to add redundancy into the network design to improve network robustness.

As the towers compose a substantial part of the total cost of the network, the challenge is to select the location of sites and links so that the overall cost of the towers is minimized (determined by the height 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 by judiciously selecting the transmit power of the nodes. By over-provisioning the signal at the receiver, capture effect can be used to eliminate most of the WiFi interference.

An additional significant problem in the deployment of WiLD networks is the difficulty of performing accurate manual alignments of the directional antennas for each long distance link. This is exacerbated by the fact that factors like wind and wear and tear of towers can cause the antennas to further 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.

Other unexpected factors can also have an impact on network design. It turned out that omni-directional antennas attract lightning more when they are usually mounted on top of masts and have a sharper tip, compared to directional antennas that are typically mounted below the maximum height of the mast.

4.2 Robust Power Solutions

An important challenge to robust wireless network design is the lack of reliable power. From our past experience in the Aravind [21] and the AirJaldi [37] networks, we found out that lack of stable and quality power has greatly contributed to a substantial decrease in the robustness of system components that would otherwise work quite reliably. Although issues such as frequent power outages in rural areas are well known, we were surprised by the *degree of power quality* problems in rural villages even when power is available. Our measurements of the grid power supply in India showed that power spikes above 500V, often with reversed polarity, and some even reaching 1000V are common, and so are extended sags below 70V and swells above 350V.

The key to understand the power problem is that the real cost of power in rural areas is not the cost of grid power supply, but of cleaning it using power controllers, batteries and solar-power backup solutions. Also, due to short lifetime of batteries and ineffective UPSs, power cleaning is a recurring cost [64, 39]. Solar power, although still expensive, turns out to be more competitive than expected as it produces clean power directly.

While much work remains to be done in this space, we have been designing preliminary solutions to combat the power problem [64]. Our approach to this problem is to develop a combination of smart hardware components and better techniques to avoid damage due to lightning and power surges. We first designed a Low Voltage Disconnect (LVD) solution, which prevents both routers from getting wedged at low voltages and also over-discharge of batteries. Now we are working on a microcontroller based low-cost power controller that supplies stable power to the equipment by combining input from solar panels, batteries, and even the grid. It has several features such as maximum power point tracking, low voltage disconnect, trickle charging and very importantly, support for remote management via ethernet.

4.3 Network Management Tools

Network management plays a fundamental role in reducing the network downtime in the face of outages. In rural networks, network management is a challenging problem due to: (a) the lack of local support; (b) poor transportation to rural areas; (c) constant equipment failures; (d) lack of reliable power. We are currently developing different network management tools to ease configuration, fault diagnosis and network monitoring [64].

Accurate diagnosis of a problem can greatly reduce response time and thus downtime. For example, a remote host is running properly but is unreachable when an intermediate wireless link goes down, better diagnosis will prevent unnecessary travel to the remote location. Other challenges for remote monitoring are misunderstanding among local staff about equipment usage which often worsens the problem and lack of good connectivity to remotely login to these networks.

As a result, all aspects of system management require some level of monitoring. We built a *push-based* monitoring mechanism that we call “PhoneHome” in which each wireless router pushes status updates upstream to our US-based server. We collect both passive parameters and active measurements such as maximum link or path throughput and loss. PhoneHome proved to be helpful in understanding failures, diagnosing and predicting many faults. First, it helped maintain network reachability information, alerting the local staff when the network was down and action needed to be taken to recover.

It is also important to have out-of-band access or a backchannel to the nodes that is separate from the primary wireless path to it. Simple operations such as correcting a router misconfiguration, or rebooting the router remotely can be easily done using the back-channel. Backchannel access is also useful in getting information about battery status from a remote node.

Various types of backchannels are possible. GPRS based backchannels can be used to diagnose misconfiguration of routers in case of network partitions. To decrease costs, instead of using GPRS as the backchannel, a cheaper mechanism could be using SMS channels. With SMS, console access would need to be implemented from scratch. Instead of console access, one approach would be to just query the remote router over SMS. In general, failure-independent recovery mechanisms are essential for managing systems remotely. In situations where the main router itself is wedged or is in a non-responsive state, we need components that can reset or reboot the main router for recovery. The components should not be affected by the failure themselves. Both software and hardware based watchdog mechanisms can be used for this.

5 Application Specific Challenges

In this section, we briefly outline some of the application specific challenges and describe our initial design ideas towards addressing some of these challenges.

QoS Challenges: Many important applications in rural developing world such as telephony, telemedicine and distance learning would require QoS guarantees from the underlying network layer. While the underlying MAC layer, does provide a certain level of error-recovery, this is typically insufficient to achieve end-to-end QoS. In addition, the net available bandwidth on each link varies as a function of time due to the TDMA protocol used at the MAC layer. Providing QoS guarantees in a network where every link is lossy as well as has time-varying bandwidth is known to be a hard problem [61]; traditional QoS mechanisms have been designed for networks with fixed capacities. To solve this problem, we leverage ideas from my prior work on OverQoS [61], an overlay network based architecture that uses the basic concept of a *controlled loss virtual links (CLVL)* to provide statistical end-to-end QoS over bandwidth-varying lossy network links. A CLVL provides two guarantees: (a) the loss-rate of any flow within the virtual link is bounded by a small value q

with high probability; (b) A certain minimum bandwidth c_{min} can be guaranteed on a virtual link with high probability. The values of q and c_{min} are dependent on the characteristics of the underlying link. In WiRE, we intend to use the CLVL concept to provide link-level guarantees and use well-known QoS signaling mechanisms to provide end-to-end QoS.

Multicast services: The local store at every wireless node is critical to support a wide range of network services. To support efficient multicast and bulk content distribution for the distance learning application, the local store can be used as an in-network replication entity for enhancing end-to-end performance. Even in the case of video multicast over lossy network links, the local store can be used to quickly recover from packet losses downstream in the network. Given the vagaries of power supply in rural regions, network outages are the norm and not the exception. In such cases, the local store can be used provide several intermittent network services including simple store-forward. For web browsing, the local store can act as a proxy cache to enhance the system performance.

Cell Mobility: To support telephony on WiFi-enabled cellphones, WiRE needs to provide a phone-translation service that can translate from cell-phone numbers to IP addresses within the WiRE network. Our solution is similar to MobileIP [45]. We associate a central registry that maps each cellphone to a home region; the main wireless node with the local storage acts as the translation server for the home region. This node within the region maintains the current IP address of the cellphone. If a user moves to a new region, the visiting region's server updates the home region with the new IP address. WiRE does not support mobility across regions during a call session which significantly simplifies the design; therefore, updates of IP addresses to the home region server are very infrequent. To handle mobility across basestations within a region, we assign static IP addresses to home region devices based on the MAC identity; visiting nodes are assigned dynamics IPs. For fault tolerance purposes, the primary static IP address of a cellphone is propagated to the central registry.

Secure Identities: Every cellphone has a unique identity that is routable using the traditional cellular network. In our system, each cellphone has two independent routing channels: (a) using the WiFi network; (b) using the traditional cellular network. While WiRE may not be trusted, the cellular network provides a trusted channel between any two cellphones. If a cellphone has a unique identity I , then the cellphone can make I into a cryptographically strong identity using a self-certifying key [60]. I can locally generate a public-private key pair $(P(I), Q(I))$ and distribute the self-certifying name $(I, P(I))$ to any other cellphone within the WiRE network. The device can prove that its ownership of the identity I by any challenge response protocol using the two independent routing channel - send the challenge on the cellular network and the response on WiRE. Apart from this, any server with a cellphone interface can as a Public-Key Infrastructure and distribute secret keys over the trusted cellular network. The ability to provide cryptographically secure identities for cellphones within WiRE enables the system to support secure financial transactions and mobile banking services. The cellphone does not need to be always in the vicinity of the cellular network to achieve these security properties; once a trusted channel is established using the cellular network and secret keys are exchanged, then cellular connectivity is not essential.

6 Work and Deployment Plan

Table 1 illustrates the tentative time-line and work plan for the next five years. We intend to deploy the proposed network architecture, both in rural developing regions as well as locally at NYU for testing purposes.

Doing any deployment in a rural developing country setting is a very challenging task due

Year	Goals
1	point-to-multipoint MAC, solar power controller, management tools (ver 1), improving deployment at Aravind
2	unified MAC protocol design, fault analysis and monitoring tools, QoS challenges, deployments in Ghana, South Africa
3	unified MAC testing, QoS testing, multicast design, cell mobility, review of deployment and management tools, expand deployments
4	unified MAC design (complete, code release), fault tolerant issues, secure identities using cellphones, version 1 of WiRE, upgrade all deployments to new solutions
5	assess deployment feedback, revise WiRE to version 2, complete WiRE deployment

Table 1: Work plan

to three factors: minimal local support, the expectation of perfect-working systems and the high travel costs for fixing network errors. There is very small room for error in these systems; local users expect whatever is being deployed to be simple to use and also function properly almost all the time - if not, the usability of the system significantly deteriorates. Therefore, it is essential to build a testbed within NYU to stress test the network architecture before venturing into any developing country. Currently, in collaboration with my colleague Jinyang Li, we have built a 24-node indoor multi-radio wireless testbed which provide point-to-point and point-to-multipoint network access across different buildings at NYU at very low-costs as opposed to renting cable modem services.

To make a deployment successful and have impact in a rural region, it is essential to identify the right local partner to work with. For this project, we have established relationships with partners and universities in India, Ghana and South Africa for potential deployments. We will continue to work with Aravind Eye Hospitals in expanding their existing WiLD network to support the new proposed architecture and the new software. Another deployment point within India is our ongoing collaboration with Amrita University and Amrita Institute of Medical Sciences (AIMS) in South India who have expressed interest in interconnecting with Aravind telemedicine network. AIMS and Aravind are among the largest hospitals in India serving over 2 million patients per year [40, 21]. In Africa, our work will primarily be based out of Ghana and South Africa, primarily centered around telemedicine services. In Ghana, we work with NYU in Ghana, Korlebu Hospital, University of Legon and West Africa AIDS Foundation. We recently worked with OneTouch (a large ISP in Ghana), to interconnect nearly 2000 physicians to use the OneTouch network for free teleconsultation services. In South Africa, we have strong working relationships with cellular service providers and hospitals and we have been approached to do a pilot project in Johannesburg.

7 Education Plan

Computer Science is relatively new to the Information and Communication Technologies for Development (ICTD) space. To improve the interest in this space as well as enhance its awareness, there is a great need for curricular development, multi-disciplinary education and cross disciplinary research. NYU, being in the heart of New York City, lays much emphasis on international developmental activities and is in a unique position to do research work in the ICTD space with the local presence of the UN, the Development Research Institute and the Earth Institute. NYU is one of the few international schools with a university campus in Africa and Abu Dhabi. My educational

plan consists of the following parts:

Interdisciplinary Collaboration: To work in the ICTD space, interdisciplinary collaboration is vital. At NYU, I lead the Cost-effective Appropriate Technologies for Emerging Regions (CATER) research group which is a joint effort from researchers in Computer Science, School of Medicine, Public Health, Economics and Public Policy. The core focus of the CATER group is to address the research challenges that arise in the development of appropriate low-cost technological solutions for developing regions. I also closely work with the TIER research group at UC Berkeley led by Prof. Eric Brewer, who has done pioneering research in the ICTD space. We co-advise a few students at NYU and UC Berkeley. I am also a part of the NYU Africa House focusing on African developmental activities.

Curriculum Development: I recently designed a new graduate course titled “Information and Communication Technologies for Developing Regions” which focused on how ICT can play an important role in addressing pressing problems in developing regions in the space of healthcare, education, finance, agriculture and supply chain management. Designing this course was very challenging since it brought together material from different areas: computer science, economics, public policy, global health and education. The course also had guest lectures from reputed experts from the School of Medicine, Economics and the United Nations. This course resulted in several interesting ICTD projects of which a few may be deployed in Africa in the upcoming year. I plan to revise this course to make it an inter-disciplinary course and make it accessible to students from other departments. I also have the approval of NYU to teach this class on a short term basis in the NYU Ghana and Abu Dhabi campuses. Apart from this course, I primarily teach two courses at NYU on “Networks and Distributed Systems” and “Security”. In both these courses, I constantly discuss networking, systems and security research challenges in the developing world.

Mentoring Students: Working with students and learning from them is something I thoroughly enjoy. I currently work with a set of highly talented PhD students. I advise four PhD students at NYU (one co-advised with Jinyang Li) and work closely with three students at UC Berkeley (advised by Eric Brewer). My prior work on WiLDNet was in joint collaboration with Eric and the students at Berkeley. All my classes have been project-oriented and I have advised many Masters students to successfully complete challenging research projects. As part of the CATER group, I do get the opportunity to regularly interact with students in other disciplines. I also currently advise one student in Amrita University in India as part of an ongoing collaboration in this space.

Field Work: A large fraction of students have very little exposure to the developing world. Field work is an essential part for any student who works in this space. I intend to develop a case-study program in collaboration with NYU School of Public Policy and the Global Health program where graduate students from different disciplines will do field work and needs-assessment studies in developing countries. This also increases the chances of successful deployments. Currently, two of my PhD students and three medical students are spending the summer in Ghana and India on case studies and deployment efforts. The largest educational value in this space comes from field work.

Community Outreach: I intend to give talks and tutorials in several venues to raise awareness about ICTD to fellow researchers and encourage them to work in this exciting new research space. I intend to co-organize conferences and workshops that specialize in this space. This year, I chaired the SIGCOMM 2008 workshop on Networked Systems for Developing Regions (NSDR) and the WWW 2008 track on developing regions. From next year, we plan to make NSDR, a premier publication venue in this space.

8 Prior Research and Educational Accomplishments

I will elaborate on my contributions in three research topics and educational accomplishments:

Technologies for Developing Countries: As I described earlier, I was involved in the design, implementation and deployment of WiLD networks [42, 62, 54, 35, 64, 34] which has been widely deployed in different countries. Apart from this, we recently developed PaperSpeckle [53], a system that uses a simple microscope and a pen to extract a unique signature for any piece of paper based on the inherent structural properties of paper. Paper Speckle is tamper-proof, extremely low-cost and has several applications in developing countries in supply chain management, offline paper authentication and financial services. We have also developed, RuralCafe [13], a system that addresses the problem of how to enhance web search to work efficiently over intermittent and low bandwidth networks. The SmartTrack project [67], a recently initiated project aims at building a cellphone-based distributed information system that can be used for tracking the flow of AIDS drugs in Africa; this is currently under deployment in Ghana.

Secure routing protocols: My thesis work [57, 59, 60] developed decentralized security mechanisms for Internet routing protocols which do not rely on a central authority or a PKI. Prior solutions for securing Internet routing protocols relied on a PKI-based approach which had significant hurdles to deployment. This work won the *Best student paper award* at NSDI 2004 and my thesis received the *C.V. Ramamoorthy award* at UC Berkeley. We showed how these ideas could be extended to secure DNS in a decentralized manner [57]. We leverage these ideas in the design of the secure naming system for cellphones described in this proposal.

Network Architecture: Two notable network architectures that I have developed are OverQoS [61] and HLP [58]. OverQoS proposed an overlay-based QoS architecture that can be incrementally deployed on the Internet which unlike all previous QoS architectures, did not require any modifications to the routers in the Internet. HLP is a next-generation inter-domain routing protocol that was designed as a replacement to the existing Border Gateway Protocol (BGP) used in the Internet today. HLP addressed several fundamental shortcomings in the design of BGP including poor scalability, convergence, security, stability and diagnosis support.

Educational accomplishments: In the last two years at NYU, I have designed and taught three different graduate courses: (a) Networks and Distributed Systems (co-taught with Jinyang Li); (b) What if a Computer Lies? (a course on Computer Security); (c) ICT for Developing Regions. All these courses have been project-based involving continuous interactions with each student group on a weekly basis. Overall, many course projects have been successful and have taken shape into becoming long-term research projects. In the Networks and Distributed systems class, the students performed two assignments, of which, one of them focused on analyzing the performance of wireless networks deployed within NYU and how to improve them; we later used the results to reconfigure the channel allocation in one of the buildings within NYU. In the security class, each student also performed a detailed survey of a specific topic within computer security. At Berkeley, I have regularly given guest lectures in both undergraduate and graduate classes in networking. I was a TA for the Digital Logic Design course which had over 200 students.

References

- [1] ADYA, A., BAHL, P., PADHYE, J., WOLMAN, A., AND ZHOU, L. A Multi-Radio Unification Protocol for IEEE 802.11 Wireless Networks. In *IEEE BroadNets* (Oct. 2004).
- [2] AFRICA COUNTS COST OF MAKING A CALL. <http://news.bbc.co.uk/2/hi/business/4277477.stm>.
- [3] AFRICAN SATELLITE ISPS. <http://www.satsig.net/ivsatsat-africa.htm>.
- [4] AGUAYO, D., BICKET, J., BISWAS, S., JUDD, G., AND MORRIS, R. Link-level Measurements from an 802.11b Mesh Network. In *ACM SIGCOMM* (Aug. 2004).
- [5] BHAGWAT, P., RAMAN, B., AND SANGHI, D. Turning 802.11 Inside-out. *ACM SIGCOMM CCR* (2004).
- [6] BICKET, J., AGUAYO, D., BISWAS, S., AND MORRIS, R. Architecture and Evaluation of an Unplanned 802.11b Mesh Network. In *ACM MOBICOM* (Aug. 2005).
- [7] BISWAS, S., AND MORRIS, R. Opportunistic Routing in Multi-Hop Wireless Networks. *Hotnets-II* (November 2003).
- [8] BREWER, E., DEMMER, M., DU, B., HO, M., KAM, M., NEDEVSCHI, S., PAL, J., PATRA, R., SURANA, S., AND FALL, K. The Case for Technology in Developing Regions. *Computer* 38, 6 (2005), 25–38.
- [9] BREWER, E., DEMMER, M., HO, M., HONICKY, R. J., PAL, J., PLAUCHÉ, M., AND SURANA, S. The Challenges of Technology Research for Developing Regions. *IEEE Pervasive Computing* 5, 2 (April-June 2006), 15.
- [10] CHANDRA, R., MAHAJAN, R., MOSCIBRODA, T., RAGHAVENDRA, R., AND BAHL, P. A Case for Adapting Channel Width in Wireless Networks. *ACM SIGCOMM* (2008).
- [11] CHEBROLU, K., AND RAMAN, B. FRACTEL: A Fresh Perspective on (Rural) Mesh Networks. In *ACM SIGCOMM Workshop on Networked Systems for Developing Regions (NSDR)* (August 2007).
- [12] CHEBROLU, K., RAMAN, B., AND SEN, S. Long-Distance 802.11b Links: Performance Measurements and Experience. In *ACM MOBICOM* (2006).
- [13] CHEN, J., AND SUBRAMANIAN, L. Redefining Web Search in Intermittent Networks. In *NYU Technical Report* (2008).
- [14] DOERR, C., NEUFELD, M., FILFIELD, J., WEINGART, T., SICKER, D. C., AND GRUNWALD, D. MultiMAC - An Adaptive MAC Framework for Dynamic Radio Networking. In *IEEE DySPAN* (Nov. 2005).
- [15] DU, B., DEMMER, M., AND BREWER, E. Analysis of WWW Traffic in Cambodia and Ghana. *15th International WWW Conference* (May 2006). Edinburgh.
- [16] FALL, K. A Delay Tolerant Network Architecture for Challenged Internets. In *Proceedings of ACM SIGCOMM* (2003).
- [17] FONSECA, R., AND PAL, J. Computing Devices for All: Creating and Selling the Low-Cost Computer. *1st International Conference on ICT and Development* (May 2006).

- [18] GOOGLE ANDROID. <http://www.google.com/android>.
- [19] GUO, S., FALAKI, M., OLIVER, E., RAHMAN, S., SETH, A., ZAHARIA, M., AND KESHAV, S. Very low-cost internet access using KioskNet.
- [20] HONICKY, R. J., BAKR, O., DEMMER, M., AND BREWER, E. A Message Oriented Phone System for Low-Cost Connectivity. *Workshop on Hot Topics in Networks (HOTNETS)* (November 2007).
- [21] HOSPITALS, A. E. <http://www.aravind.org>.
- [22] IEEE 802.16. IEEE 802.16 WirelessMAN Standard for Wireless Metropolitan Area Networks.
- [23] IN AFRICA, C. P. S. <http://www.revivaltimes.com/cellphonestats.html>.
- [24] IN KENYA, C. C. P. <http://www.ke.celtel.com/en/personal-plans/prices/index.html>.
- [25] IN SOUTH AFRICA, C. P. U. S. <http://www.markettree.co.za/south-african-statistics.php?statistics=cellphone>.
- [26] INTERNATIONAL TELECOMMUNICATION UNION STATS. <http://www.itu.int/ITU-D/ict/statistics/ict/>.
- [27] INTERNATIONAL TELECOMMUNICATIONS UNION - STATISTICS. <http://www.itu.int/ITU-D/ict/statistics/>.
- [28] INTERNET WORLD STATS. <http://www.internetworldstats.com>.
- [29] JAIN, K., PADHYE, J., PADMANABHAN, V., AND QIU, L. Impact of Interference on Multi-hop Wireless Network Performance. In *ACM MOBICOM* (Sept. 2003).
- [30] JAIN, S., DEMMER, M., PATRA, R., AND FALL, K. R. Using Redundancy to Cope with Failures in a Delay Tolerant Network. *ACM SIGCOMM* (August 2005).
- [31] JAIN, S., FALL, K., AND PATRA, R. Routing in a Delay Tolerant Network. In *ACM SIGCOMM* (New York, NY, USA, 2004), ACM, pp. 145–158.
- [32] KAM, M., RAMACHANDRAN, D., SAHNI, U., AND CANNY, J. Designing Educational Technology for Developing Regions: Some Preliminary Hypotheses. *3rd International Workshop on Technology for Education in Developing Countries (TEDC)* (July 2005).
- [33] KURIYAN, R., RAY, I., AND TOYAMA, K. Integrating Social Development and Financial Sustainability: The Challenges of Rural Kiosks in Kerala. *1st International Conference on ICT and Development* (May 2006).
- [34] MISHRA, S. M., HWANG, J., FILIPPINI, D., DU, T., MOAZZAMI, R., AND SUBRAMANIAN, L. Economic Analysis of Networking Technologies for Rural Developing Regions. *Proceedings of 1st Workshop on Internet and Network Economics (WINE)* (December 2005).
- [35] NEDEVSCHI, S., PATRA, R., SURANA, S., RATNASAMY, S., SUBRAMANIAN, L., AND BREWER, E. WiLDNet: Design and Implementation of High Performance WiFi Based Long Distance Networks. *USENIX NSDI* (2007).

- [36] NEDEVSCHI, S., SURANA, S., DU, B., PATRA, R., BREWER, E., AND STAN, V. Potential of CDMA450 for Rural Network Connectivity. *IEEE Communications Magazine* (Jan 2007). Special Issue on New Directions In Networking Technologies In Emerging Economies.
- [37] NETWORKS, A. J. <http://www.airjaldi.com>.
- [38] NEUFELD, M., FIFIELD, J., DOERR, C., SHETH, A., AND GRUNWALD, D. SoftMAC - Flexible Wireless Research Platform. In *HotNets-IV* (Nov. 2005).
- [39] NORTHERN ARIZONA WIND AND SUN. Deep Cycle Battery FAQ, 2006. http://www.windsun.com/Batteries/Battery_FAQ.htm.
- [40] OF MEDICAL SCIENCES, A. I. <http://www.aimshospital.org>.
- [41] PAL, J., NEDEVSCHI, S., PATRA, R., AND BREWER, E. A Multidisciplinary Approach to Open Access Village Telecenter Initiatives: The case of Akshaya. *E-Learning* 3, 3 (2006), 291.
- [42] PATRA, R., NEDEVSCHI, S., SURANA, S., SHETH, A., SUBRAMANIAN, L., AND BREWER, E. WiLDNet: Design and Implementation of High Performance WiFi Based Long Distance Networks. *Proceedings of NSDI 2007* (2007).
- [43] PAUL, K., VARGHESE, A., IYER, S., AND KUMAR, B. R. A. Wi-FiRe: Rural Area Broadband Access Using the WiFi PHY and a Multisector TDD MAC. *New Directions in Networking Technologies in Emerging Economics, IEEE Communications Magazine* (2006).
- [44] PENTLAND, A., FLETCHER, R., AND HASSON, A. DakNet: rethinking connectivity in developing nations. *Computer* 37, 1 (2004), 78–83.
- [45] PERKINS, C. IP Mobility Support. *Request for Comments* (1996).
- [46] PHONE, G. <http://www.grameenphone.com>.
- [47] RAMAN, B., AND CHEBROLU, K. Revisiting MAC Design for an 802.11-based Mesh Network. In *HotNets-III* (2004).
- [48] RAMAN, B., AND CHEBROLU, K. Design and Evaluation of a new MAC Protocol for Long-Distance 802.11 Mesh Networks. In *ACM MOBICOM* (Aug. 2005).
- [49] RANIWALA, A., AND CKER CHIUEH, T. Architecture and Algorithms for an IEEE 802.11-based Multi-channel Wireless Mesh Network. In *IEEE INFOCOM* (March 2005).
- [50] RAO, A., AND STOICA, I. An Overlay MAC Layer for 802.11 Networks. *Technical Report, UC Berkeley* (2004).
- [51] REDDY, N. P. The SRAWAN MAC Protocol to support Real-Time Services in Long Distance 802.11 Networks. Master’s thesis, Indian Institute of Technology, Kanpur, Aug. 2006.
- [52] SEN, S., AND RAMAN, B. Long Distance Wireless Mesh Network Planning: problem Formulation and Solution. *World Wide Web 2007* (2007).
- [53] SHARMA, A., SUBRAMANIAN, L., AND BREWER, E. Secure Rural Supply Chain Management using Low Cost Paper Watermarking. In *Proceedings of ACM SIGCOMM NSDR workshop 2008* (2008).

- [54] SHETH, A., NEDEVSCHI, S., PATRA, R., SURANA, S., SUBRAMANIAN, L., AND BREWER, E. Packet Loss Characterization in WiFi-based Long Distance Networks. *IEEE INFOCOM* (2007).
- [55] STANDARD, G. <http://www.gsmworld.com/technology/gprs/index.shtml>.
- [56] STUDY: WIFI CELLPHONES WILL HIT IT BIG. http://news.cnet.com/2100-1039_3-6088484.html.
- [57] SUBRAMANIAN, L. Decentralized security mechanisms for routing protocols. Ph.D. thesis, University of California, Berkeley.
- [58] SUBRAMANIAN, L., CAESAR, M., EE, C. T., HANDLEY, M., MAO, M., SHENKER, S., AND STOICA, I. HLP: A Next Generation Inter-domain Routing Protocol. In *Proceedings of ACM PODC 2005* (2005).
- [59] SUBRAMANIAN, L., ROTH, V., STOICA, I., SHENKER, S., AND KATZ, R. H. Listen and Whisper: Security Mechanisms in BGP. In *Proceedings of ACM/USENIX NSDI 2004* (2004).
- [60] SUBRAMANIAN, L., ROTH, V., STOICA, I., SHENKER, S., AND KATZ, R. H. Reliable Broadcast in Unknown Fixed Identity Networks. In *Proceedings of ACM PODC 2005* (2005).
- [61] SUBRAMANIAN, L., STOICA, I., BALAKRISHNAN, H., AND KATZ, R. OverQoS: An Overlay Based Architecture for Enhancing Internet QoS. In *USENIX/ACM NSDI* (March 2004).
- [62] SUBRAMANIAN, L., SURANA, S., PATRA, R., HO, M., SHETH, A., AND BREWER, E. Rethinking Wireless for the Developing World. *Hotnets-V* (2006).
- [63] SURANA, S., PATRA, R., AND BREWER, E. Deploying a Rural Wireless Telemedicine System: Experiences in Sustainability. *IEEE Computer* (June 2008).
- [64] SURANA, S., PATRA, R., NEDEVSCHI, S., RAMOS, M., SUBRAMANIAN, L., BEN-DAVID, Y., AND BREWER, E. Beyond Pilots: Keeping Rural Wireless Networks Alive. *Proceedings of NSDI 2008* (2008).
- [65] SYMBIAN OS. <http://www.symbianos.org>.
- [66] THE AKSHAYA E-LITERACY PROJECT. <http://www.akshaya.net>.
- [67] THE SMARTTRACK PROJECT. <http://cater.cs.nyu.edu>.
- [68] UNITED VILLAGES. <http://www.unitedvillages.com>.
- [69] VUTUKURU, M., JAMIESON, K., AND BALAKRISHNAN, H. Harnessing Exposed Terminals in Wireless Networks. *USENIX NSDI* (2008).
- [70] WiMAX FORUM. <http://www.wimaxforum.org>.
- [71] YUNUS, M. The Grameen Bank. *Scientific American* (November 1999).