Digital Gangetic Plains[*](DGP): 802.11-based Low-Cost Networking for Rural Areas 2001-2004: A Report

The DGP Media Labs Asia Team

Principal Investigator: Pravin Bhagwat
Co-Investigators: Dheeraj Sanghi, Bhaskaran Raman

Report edited by: Bhaskaran Raman

Department of Computer Science and Engineering
Indian Institute of Technology, Kanpur

July 2004

Abstract

The past decade has seen communication revolution in the form of cellular telephony as well as the Internet, but much of it has been restricted to the developed world and metropolitan pockets in the developing world. While the use of cellular technologies can cut down on the time to deploy access networks, the cost economics make this non-viable in growing telecom economies. In the Digital Gangetic Plains (DGP) project, we have been exploring the use of 802.11 as a long-distance access technology. 802.11, popularly known as WiFi, is currently cost-priced since it is an open/inter-operable standard, and hence is attractive for low cost and rapid deployment in rural areas.

We have built an extensive testbed in a rural setting consisting of directional 802.11 links, spanning up to about 80km at its longest. The testbed is set in the Kanpur-Lucknow corridor of the Gangetic Plains. To our knowledge such a long-distance, multi-hop testbed based on 802.11 is unique thus far. While 802.11 is attractive in terms of cost economics, it was inherently designed for indoor use. Our novel use of the technology for outdoor, long-distance access links presents several challenges. Our experience with the testbed has brought several research as well as operational issues to the fore.

Our contributions in this project thus far include: (1) demonstration of technical feasibility of using 802.11 for rural Internet connectivity, (2) an understanding of several operational and cost issues, (3) various technical contributions, (4) demonstration of and experience with various services on the testbed network, as well as (5) identification of various technical and operational issues going forward. In this report, we document our experiences with this technology, our contributions in this project, as well as the novel technical and operational challenges that lie ahead in using 802.11 to bridge the digital divide.

[*] A project supported by Media Labs Asia

Contents

1 Introduction and Motivation

The past decade has seen steady exponential growth of the Internet as well as cellular telephony. A region-by-region study of this growth is instructive. Figure 1 illustrates cellular telephone growth in terms of the Compound Annual Growth Rate (CAGR) for various countries and regions. The data shown is for the years 1995-2001 [10][1]. The growth rate has been phenomenal, especially in the developing world.

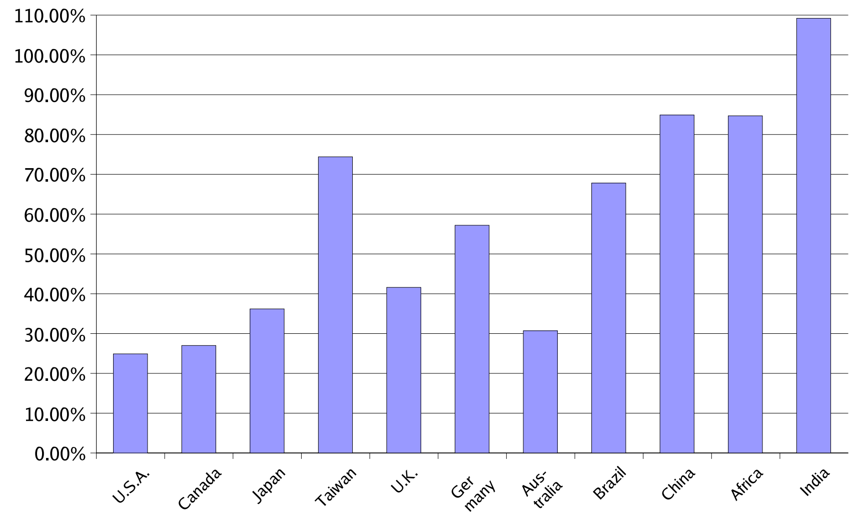

Figure 1: Compound annual growth rate of cellular phones: 1995-2001 (source [10])

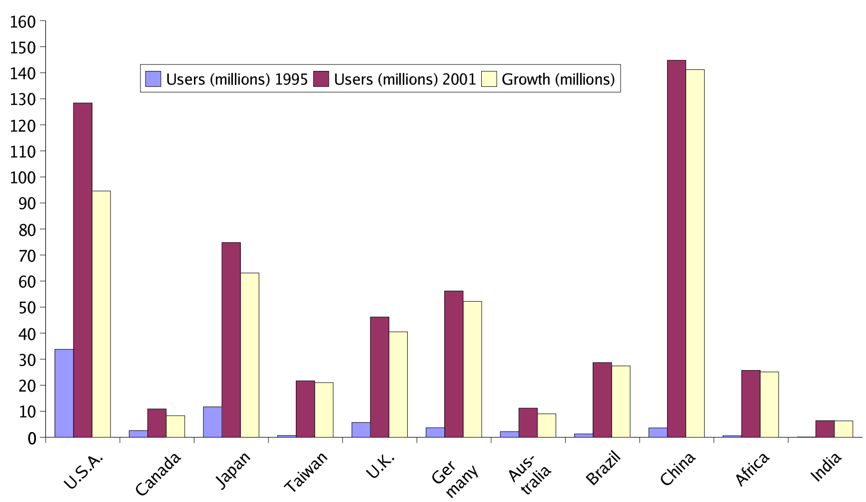

Figure 2: Absolute growth of cellular phones: 1995-2001 (source [10])

However, if we look at the growth in terms of absolute numbers, the growth in developing nations is starkly behind that of the developed world. This is shown in Figure 2. This means

[1] Also available from `http://www.cse.iitk.ac.in/~bhaskar/dgp/stats/`

that the CAGR is high for the developing regions of the world only since the growth started from a low number. The difference is even more stark if we examine the *density* of access to communication technology. Figure 3 compares the accessibility of various technologies to people in different countries. The density of telephone, cellular phone, and Internet users in the developing world still lags considerably behind that in the developed world. Hence, parallel with the growth the *digital divide* has grown too.

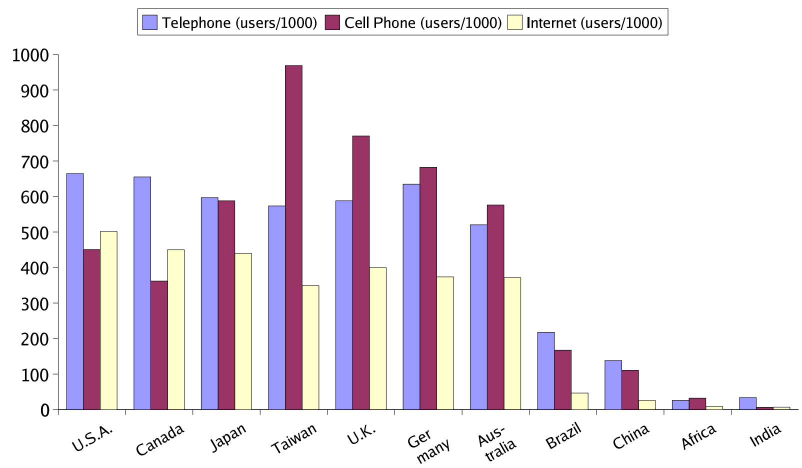

Figure 3: Density of communication technologies (2001) (source [10])

The data we have shown is for the period 1995-2001, since this period represents the communication/Internet-boom period. The data for the period 1998-2003 [10][2] shows a similar trend. The reasons for this are not hard to find. Communication evolution has been shaped by technical innovations and market forces in the western world. The business model thrives where the average per-capita income is high (U.S. $20,000 or more). For developing nations, the growth in Table 2 has been concentrated in metro pockets. This is especially unfortunate since the majority of the developing world population is in rural areas (e.g., in India, rural population constitutes 74% [1]). Hence the low tele-density as given in Figure 3.

The cost of large-scale communication deployment is out of reach of developing countries. The cost per land-line telephone connection today is about U.S.$400, and can be expected to drop to about $200 in the next decade. To see what this implies, consider the case of India. If we were to target 400 million telephone lines (40% of India's current population), this works out to U.S.$80 billion. This is about 25 times India's entire budget allocation for the entire department of telecommunications [5]. Also, 60-70% of the deployment cost is in the access network and not the core network. This acts as further disincentive for rural deployment since the density of users as well as their paying capacity is very less as compared to cities.

While cellular wireless technologies may help in quick deployment, like land-lines, their cost structure too is suited for the developed world. The service, and more importantly the equipment, are *value-priced* for markets where the users are willing to pay a high price. This dictates that the technology is economically non-viable for rural settings in developing telecommunication economies – a rural deployment cannot bring adequate returns on investment. The same is true of data-based wireless access networks such as 802.16 Wireless-MAN [16] as well as commercial packet radio systems such as Ricochet [11]. While 802.16 is

[2] Also available from `http://www.cse.iitk.ac.in/~bhaskar/dgp/stats/`

a technology specifically designed from the ground-up for long-distance point-to-point wireless networking, it remains to be seen when, and if at all, it will reach the scale of competitive mass production to lower costs significantly.

In this context, consider the 802.11 family of wireless technologies [8]. Since its inception in 1994, 802.11 WiFi has shown tremendous growth and acceptance (in U.S. and Europe) as a last-hop wireless solution in corporate/enterprise settings as well as in home networks. With widespread acceptance of the technology, open/inter-operable standard, and competitive mass production, the equipment and chip sets are *cost priced* and are hence inexpensive. Table 1 summarizes the costs of various relevant aspects of the technology.

Item	Cost (U.S.$)
Chip sets	$25-30
Access Points	$120-700
PCMCIA Cards	$60-110

Table 1: 802.11 equipment cost (approximate)

To our knowledge, the only other wireless technology which has been designed with the explicit goal of low-cost is CorDECT [28], developed by the TENET [15] research group[3]. CorDECT has been developed as an extension of the European digital cordless standard, DECT. And the reasons for the low cost of CorDECT are similar to that for 802.11 – mass production of the equipment. The two technologies are similar in terms of their requirements of line-of-sight – and hence the height of the antenna tower required (the antenna tower cost is a significant component of the system cost). While it is possible that both technologies will have their niche, in comparison with CorDECT, we believe that 802.11 has at least the following advantages.

- 802.11 is fundamentally a data-based standard, while CorDECT is fundamentally a voice-based standard [28]. With the wireless and telecommunication world moving towards a data-centric, Internet-based model, we believe that an 802.11-based approach has fundamental advantages in terms of leveraging protocols, standards, products applications, as well as deployed services in the future.
- 802.11 can provide peak data rates of up to 54 Mbps (albeit at the cost of using a larger spectrum), while CorDECT can provide only 70 Kbps peak data rate [28]. This potentially means fundamental limitations in the set of applications which can be run on CorDECT. This is especially so as wired bandwidths and the corresponding application requirements have been showing a steady increase.
- Finally, and importantly, given the growing popularity of 802.11 in enterprise and home networks, as well as a general shift towards data-centric (as opposed to voice-centric) networks, 802.11 will likely have continued higher mass production and hence further fall in prices. This trend is likely to be much stronger for 802.11 than for the voice-centric CorDECT.

[3] At the Indian Institute of Technology, Madras

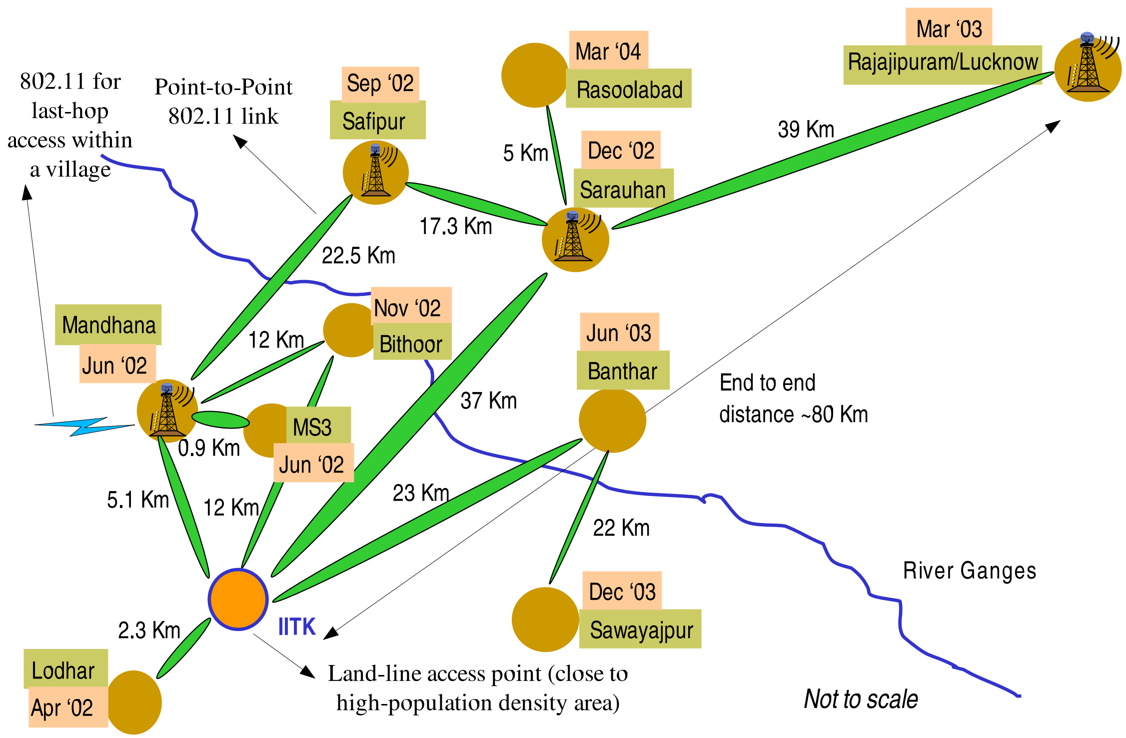

Figure 4: The Digital Gangetic Plains testbed

In the Digital Gangetic Plains (DGP) project[4], our goal is to enable *low cost* and *rapid* deployment of portable/mobile voice and data communication services in rural areas. Although 802.11 was primarily designed for indoor operation, given its attractive cost economics, we are exploring its use in long-distance rural networking. Our vision for such use of 802.11 is illustrated in Figure 4. We envision several tens of kilometers of (relatively) sparsely populated rural areas being covered by long-distance 802.11 links constituting a multi-hop wireless access network. This access network is connected to the rest of the Internet through a *landline* node, which has a wired Internet connection – e.g. through an Optical Fibre Cable (OFC) dropout. Such a landline will usually be an urban/semi-urban location where it makes economic sense to lay wired cables. We also envisage the use of 802.11 for the last-hop access within rural villages.

Figure 5 depicts the current and intended use of 802.11 in terms of the communication range. We also present other wireless technologies for comparison. Here, the x-axis gives the approximate distance for which the technology is designed. And the vertical axis gives the approximate throughput achievable. Note that we do not give the spectrum bandwidth in this picture. This is intentionally so – we wish to stress the point that the low-cost aspect of 802.11 is the more important factor which has driven our choice of this technology, as opposed to its spectral efficiency.

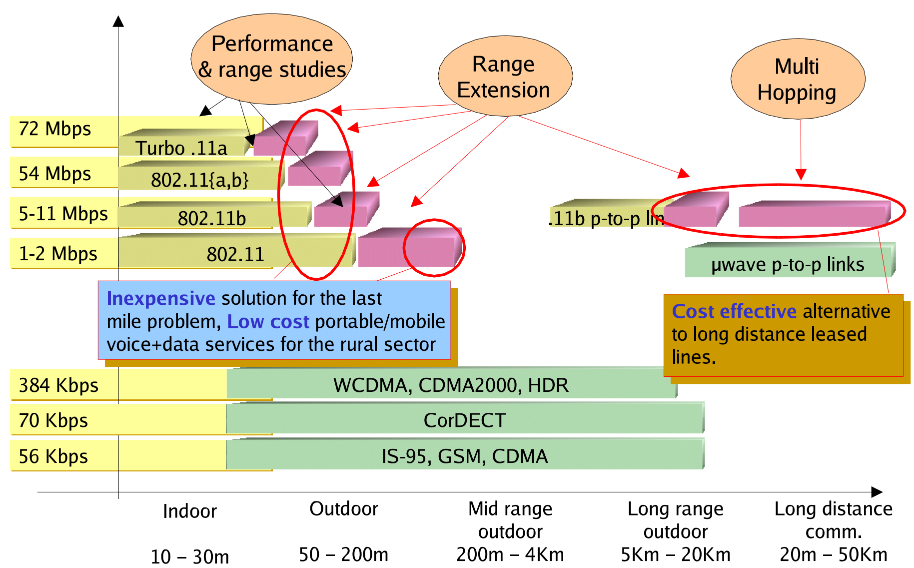

Figure 5: 802.11: current and intended use, and the value proposition

To understand the various operational and research challenges that arise in such a scenario, we have built an extensive testbed spanning several tens of kilometers. Figure 4 depicts our testbed. We use directional antennae for achieving long-distance communication using 802.11, as well as multi-hopping to achieve further range extension. We use the currently popular 802.11b variant in our testbed. As of this writing, we have a network of eight nodes and eight point-to-point links. The longest point-to-point link spans over 38km. To our knowledge, such a multi-hop, long-distance testbed using 802.11 is unique thus far. Through this testbed, we have been able to demonstrate low cost Internet and VoIP to villages that don't have a telephone to this day.

[4] At the Media Labs Asia, Indian Institute of Technology, Kanpur

We have “turned” 802.11, an indoor communication technology, “inside-out” to use it in a long-range outdoor setting. In this report, we first describe, in **Section 2**, the various technical and operational challenges in the use of an indoor technology for outdoor purposes. These challenges include those with which we started out at the beginning of this endeavour, as well as those which we have uncovered during our experience in this effort. On the operational side, we started with the fundamental question of “Is it possible to use 802.11 technology, with off-the-shelf equipment, for long-distance rural Inter-networking?”. The technical issues uncovered span the various layers of the OSI networking stack. At the physical (PHY) layer are issues such as radio propagation model, outdoor multi-path effects, etc. At the MAC (Medium Access Control) layer are issues such as the design of an appropriate protocol for the long-distance point-to-point links.

The testbed we have built has helped us tremendously in understanding these various issues. We now understand the procedures for establishing the long-distance links, and regularly use the testbed for various experiments. The description of the testbed is the topic of **Section 3**. Here we describe how we setup a link in the testbed, and discuss the various costs involved.

Throughout the project, we have addressed a range of technical challenges and several M.Tech as well as B.Tech theses have emerged from the topics addressed. We have addressed issues related to operations, security, site and network planning, MAC protocol design, among others. In terms of software tools, we have developed and released (a) a simulator for design studies of such 802.11-based networks, and (b) an open-source Linux-based sniffer software for capturing and analyzing 802.11 packets. We document our technical research contributions in **Section 4**.

We have been fortunate to work with various other Media Labs Asia projects such as Infothela, Telemedicine, Digital Mandi, etc., all at IIT-Kanpur. Interaction with these groups have made us better understand various application services which can be run on top of such a rural network. One specific service we have been successful in testing is the use of this technology for Voice-over-IP (VoIP). The description of our experiences with the different services is given in **Section 5**. We also discuss in detail our experience with the VoIP application.

Like any research and development effort, the Digital Gangetic Plains project has helped in uncovering many unsolved issues as well. These issues span technical as well as operational domains. We articulate these challenges which lie ahead, in **Section 6**. Finally, we summarize out contributions in this project in Section 7.

2 Operational, Technical, and Application Challenges

In DGP, we started out with the goal achieving low-cost rural connectivity. At the inception of the project, there were a wide range of unanswered questions. The fundamental reason for our choice of 802.11 is the cost benefits it offers. However, since 802.11 was designed for a very different operational scenario, there were several questions to be answered. Specifically, 802.11 was designed for a *indoor*, *short-distance (100-200m)*, *broadcast-based*, *local-area network*. Our intended use differs from these in all the four dimensions. We need *outdoor*, *long-distance (few km)*, mostly *point-to-point link-based*, *access network*.

These differences are the fundamental reason why there were a variety of feasibility questions to be answered. The first dimension of *outdoor* usage meant having to deal with different propagation characteristics, as well as very different network setup mechanisms. The second aspect of *long-distance* usage meant many feasibility questions as well as several questions related to appropriateness of the 802.11 protocol design for our purpose. The third factor of *point-to-point link-based* network also called for a re-look at the 802.11 protocol design in our setting. Finally, the last aspect of *access network* meant that the network now had a very different functionality to serve, and its operational requirements are different.

The various questions we started out with at the inception of the project fall under three broad categories: operational, technical, and applications – with the distinction being not so clear in some cases. We now state the challenges in these three categories, in Section 2.1, Section 2.2, and Section 2.3 respectively. We then outline which of these issues have been addressed effectively and which ones remain to be addressed, in Section 2.4. The details of how we have addressed these issues is the topic of later sections.

2.1 Operational Issues

The operational questions we started out with can be summarized as follows.

- Do off-the-shelf 802.11 equipment work for outdoor long-distance? If so, which of them work? What are the interoperability issues between equipment from different vendors for the long-distance links?
- What is the range which can be achieved? Range is a function of effective transmitter power, path loss, receive antenna gain, and receiver sensitivity. Receiver sensitivity for most commercially available radio chip-sets is about -85 to -90dBm while supported transmit power is about 15-20dBm. Some products (Orinoco, Cisco Aironet cards) provide external connectors to which high gain external antennas can be attached.
- How to plan, setup, and operate a long-distance wireless link? What is the required height clearance for a given link distance?
- What are the choices in terms of possible configurations? What are the operational settings for making the various equipment work in long-distance?
- What are the requirements in terms of weather-proofing the various equipment?
- Once operational, how stable is a link? What does stability depend upon and what is the dependence relation?
- What are the operational requirements in terms of power?
- What are the various costs involved and what are the cost/performance trade-offs?
- What are the issues in terms of deregulation and/or licensing of the spectrum? The 802.11b and 802.11g variants of the technology operate in the 2400-2483MHz range of the ISM (Industrail, Scientific, and Medical) band. The 802.11a variant operates in the 5725-5850MHz range of the ISM band. Currently in India, outdoor usage of this spectrum is licensed.

2.2 Technical/Research Issues

The technical issues can be classified under the various layers of the OSI networking stack. We start with a discussion of the physical layer issues and move up the OSI stack subsequently. The discussion here is necessarily technical, and we expect the reader to have a background on wireless networking and know the basics of 802.11, such as that given in [29]. The non-technical reader may skip this subsection. Further details can be found in the following publication [19].

Issues at the Physical (PHY) Layer

- *Empirical path-loss model:* One of the first challenges we faced was to determine how far 802.11 transmission would go. An important aspect of this is the *path loss* involved. While standard path-loss models exist for propagation of radio waves in free-space, the issue here is to understand the propagation under the given channel conditions. The characteristics of the channel include: a mostly flat terrain with a few scattered trees, atmospheric dust/moisture, and effects of the earth's curvature.

- *Equalizer design:* 802.11 PHY has been designed to operate under indoor channels where the Root-Mean-Square (RMS) delay spread is of the order of 100ns. Delay spread of 1μs or more (which is likely in long-distance outdoor scenarios) can cause severe multipath. Equalizers built in current 802.11 radio chip-sets are not designed to cope with such high delay spread channels. Adding equalizers for outdoor channels would potentially solve this problem. This problem must be addressed in the context of 802.11 PHY, and involving minimal changes to the chip-set, to preserve the cost advantages of 802.11.

- *Power efficiency:* One of the challenges we expect to face in our network is the availability of power. Rural locations rarely have reliable power. Hence the use of solar panels is an attractive option. A typical solar panel (costing about $200) could generate about $35W$ at peak operation. With an average efficiency of about 0.7 during periods of sunlight, and with about 7-8 hours ($\frac{1}{3}$ of the day) of sunlight, a round-the-clock average of about $35 \times 0.7 \times \frac{1}{3} \simeq 8W$ can be achieved. Hence optimizing the power requirements of operating the equipments at a village node is an important issue here.

Issues at the Medium Access Control (MAC) Layer

- *MAC issues for long-distance links:* The 802.11 MAC is based on Carrier Sense Multiple Access, with Collision Avoidance (CSMA/CA), and was designed for operation when a small number of people share a channel in an indoor setting. First, the value used for the ack timeout is too small to work over long-distance links (packet propagation delay over 15km is 50μs). Similarly, in the 802.11 contention-based MAC, the contention window slot-time requires adaptation to long distance lines (default slot-time value is 20μs).

  Further, the case of a node with multiple point-to-point links is a classic case of hidden nodes. However using an RTS/CTS (round-trip 100-150μs) for a packet transmission (1,000 bytes at 11Mbps is about 700μs) is very inefficient. Also, while RTC/CTS may solve the hidden-node problem, there is also the exposed-node problem in a multi-hop network.

  Apart from these, a more interesting observation is that a generic contention resolution is *not required* in our setting since we are not dealing with arbitrary contention. A time-division multiple access (TDMA) based may be more appropriate.

  An additional aspect we should consider here is the frequency dimension: 802.11b defines at least eleven channels (1-11), with overlapping frequency ranges. Figure 6 shows the power density (schematic) of the various channels as a function of the frequency. There are a maximum of three mutually non-overlapping channels: 1, 6, 11. We need to address the issue of how to allocate the available channels for each link in the network.

- *MAC issues in last-hop access in villages:* Apart from the use of point-to-point links, we also intend to use 802.11 for last-hop access to cover a region such as a village

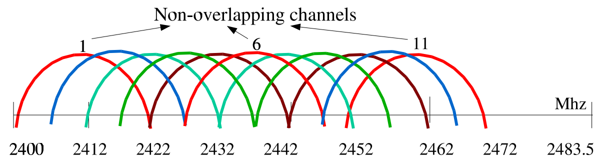

Figure 6: Overlapping 802.11 channels: schematic

(see Figure 4). This raises several issues. First, the system can have high contention, something that 802.11 MAC was not optimized for (e.g. see [20]). Further, we have the issue of CSMA-MAC performance in a geographically spread out region (0.5-1km) within a village. There is a need to understand the performance of 802.11 MAC under such conditions and subsequently improve upon it.

- *Co-existence of MAC protocols:* If we use different MACs for the long-distance links and the last-hop access (e.g., STDMA- and CSMA-based), there is a need to understand the interaction between the two. Channel allocation for both the long-distance links and last-hop access, performance under adjacent and co-channel interference, as well as fairness and bandwidth allocation issues require a careful study.

Routing and Reconfigurability Issues

A routing protocol has two related purposes in a network: (1) conveying reachability information, and (2) routing around congestion/failures (traffic engineering). Routing and traffic engineering depend on (a) the network topology, as well as (b) link capacities. In our network, both these are *dynamically reconfigurable, at will*. Link capacities can be varied by either changing the channel allocation, or by changing the various power levels of transmission. The network topology itself can be varied dynamically by turning on/off links as needed – we may turn off some links to save power, or to reduce interference at some other locations.

These various dimensions of reconfigurability can be used for fault tolerance and traffic engineering purposes. This is especially important in our network because of the following reason. In wireline ISP networks, as we get closer to the core of the network, the links are provisioned with higher capacities. This is not possible in our 802.11 network since the maximum achievable link capacity is fixed. (This also has the positive side-effect of alleviating scaling issues in traffic measurement and modeling). Note that the dynamic reconfigurations in our network are quite different as compared to ad-hoc networks since our network can be reconfigured at will. There are thus two related sets of issues: (a) the routing algorithm, and (b) algorithms and mechanisms for network reconfiguration.

Topology Construction Issues

While the routing and reconfiguration discussed above operate in sub-second time-scales (routing), or time-scales of a few minutes to a few hours (network reconfiguration), there is a related third issue that involves more static decisions. This concerns the network topology construction itself. That is, given a set of geographic locations that have to be networked, how do we determine the set of potential links in an operational network? We will have to setup antenna and equipment for a set of links. This will possibly be a superset of the set of links that are likely to be turned on at any given time (to be determined during network operation based on traffic engineering decisions).

TCP Performance Issues

While we have elaborated on issues at the physical, MAC, and routing layers, these have implications on other issues as well. With respect to TCP, its performance over multi-hop ad-hoc networks has been a topic of recent research. However, our network is not ad-hoc, and we are using long-distance point-to-point links. Also, a MAC protocol different from the default 802.11 MAC may be appropriate in our setting. There is a need to understand the *performance of TCP* under such conditions.

Regulatory Framework

There are also operational issues such as the regulatory framework for such use of 802.11. Our current transmission power and long-distance ranges are under existing FCC rules. But clearly, the FCC regulations that were designed for indoor use of 802.11 are inappropriate in our setting. There is currently a "fear of the unknown" in drafting a different regulatory policy for 802.11 in the developing world. The right regulatory framework for our setup requires further study.

2.3 Applications for Rural Areas

One of the goals in the project was to understand the application issues set in the rural context. The use of concrete applications helps us better understand the operational, technical, and cost issues. There is a need to understand the application requirements as well as cost/benefit factors. This is especially so since applications in rural areas are likely to be different, and low-cost operation is one of the main goals.

Our testbed is set in a rural area. This allows us to network several villages and run applications for real use by people as the testbed evolves.

2.4 Summary of Issues Addressed

We have been able to successfully address many of the operational, technical, as well as application issues stated above. In terms of operational issues, we have been able to demonstrate feasibility of using 802.11 for long-distance rural connectivity. We have now experimented with various equipment and have built expertise in setting up long-distance links on top of pre-built towers, cellular-phone towers, as well as make-shift masts. We have been able to achieve ranges up to about 40km for a single link, and about 80km through multi-hopping.

We have operated several links on a continuous basis, using weather-proof casings for equipment where necessary. Finally, our prototype has also helped us determine the costs of various parts of the system.

In terms of technical issues too we have addressed several. These have been in terms of student projects (MTech, BTech, as well as course projects). At the physical layer, we have developed a path loss model as well as verified it on the DGP testbed. At the MAC layer, we have designed new protocol which is maximally efficient for our DGP network. We are currently in the process of implementing it. We have also designed various other mechanisms such as one for the topology construction, for a given set of villages. In addition, we have also developed a simulator to enable future studies of various design issues. Based on our operational model and the MAC protocol, we also have proposed a regulatory framework for the 802.11 spectrum.

Section 3 and Section 4 document our solutions to the various operational and technical issues respectively. And Section 5 describes our experiences with application services.

3 The Digital Gangetic Plains Testbed

Our approach to understanding the three categories of issues, especially the operational issues, has been to build a testbed and experiment with it extensively. Our testbed is set in a rural location. This is ideal since the target deployment of this technology is in a rural setting where there are few high rise buildings in the terrain for several kilometers (the physical channel behaviour is determined by this).

We are interested in finding how coverage, throughput, and system cost would be affected when (a) 802.11 transmitters are mounted on top of pre-existing tall towers (electricity, TV broadcast, microwave, or cellular telephony), (b) Receivers are placed on rooftops in sparsely populated areas, (c) High gain directional antennas are used, (d) Transmitters and receivers are separated by tens of km, (e) Atmospheric conditions (temperature, humidity, fog, rain, etc) change, and (f) regulations for effective transmit power are relaxed.

Predicting system performance under such scenarios is difficult due to several reasons. Empirical models for radio channels for these environments do not exist. Existing simulation tools (such as ns-2 [14]) do not model the RF channel, directional antennae (side lobes and back radiation), antenna height effects, or the effect of large propagation delay on MAC. Further, modeling interference between microwave or cellular links and 802.11 is difficult. Also, equipment from different vendors perform differently and it is hard to model the differences among vendor equipment using a simulator.

The DGP testbed has been built with the following three goals.

1. *Quantify 802.11 performance outdoors:* To conduct signal coverage and performance experiments under a variety of outdoor channel conditions, build empirical path loss models for outdoor 2.4GHz channels, understand link performance under different channel conditions and under adjacent/co-channel interference.

2. *Range extension:* To test 802.11 radios beyond the prescribed limits by mounting radio transmitters and receivers on top of tall towers, and by joining multiple point-to-point links.

3. *Cost reduction:* To experiment with techniques which can reduce overall system cost through judicious choice of antennae, cable length, tower height, etc., and through better network planning and engineering.

The construction of DGP testbed was fraught with technical, operational, and regulatory challenges. When the original idea was conceived, we did not know if such a network was technically feasible. We also had no idea about what it will take to erect towers, mount antennae, and align those without aid of proper tools. The first step involved obtaining experimental licenses from the wireless regulatory body. We also secured permissions to use pre-existing tower infrastructure in the region for installing 802.11 equipment and antennae.

In this section, we describe how our testbed has been setup, and how we have handled different operational issues. We start with a description of the overall setup, in Section 3.1. We discuss how a single link is setup as well as how the overall mesh network is setup. Next, in Section 3.2, we document our experiences with various off-the-shelf equipment in the context of outdoor use. In Section 3.3, we discuss the important topic of the various costs involved in the network setup. Finally, we list the different testbed nodes and links and describe the specifics for each case, in Section 3.4.

3.1 Overall Setup

In this subsection, we first describe how a single long-distance point-to-point link is setup. Subsequently, we explain how the network as a whole is setup.

Setting up a long-distance point-to-point link

The first main set of equipments in a particular long-distance link is the 802.11 radio, one at each end of the link. Off-the-shelf 802.11 radios can be categorized into the following: (a) client devices, (b) Access-Points (APs), and (c) bridges. Clients are designed to form client-AP or client-bridge links in infrastructure mode, or client-client links in adhoc mode (see [29]). Bridges are nothing but APs, which can run the ethernet bridging protocol in a bridged-LAN. The main reason behind our experiments is that it is not necessary that these off-the-shelf equipment work for our outdoor long-distance links.

The most common setup in our testbed uses a bridge/AP at either end, since this is the easiest and most flexible to use. We document our experience with other setups in the next subsection.

In addition to the radio, we also have two antennae, one at each end. Both these antennae are usually high-gain directional antennae. One common high-gain directional antenna in use is the parabolic-grid antenna [7] with 24dBi directional gain. With such antennae placed at either end, we get an effective gain of 48dBi. For shorter links, use of other kinds of antennae, such as sector antennae or isotropic antennae, is also possible.

We have towers at either end, with the height being dependent on the distance as well as any height difference between the two ends. There are two possibilities at either end, as follows:

- The bridge/AP is placed on the ground, and an RF cable is used from the bridge/AP to the antenna atop the tower.
- Alternatively, the bridge/AP is placed on top of the tower, and connected directly using connectors (or a short RF cable) to the antenna; and we have an ethernet cable coming from the top of the tower, to provide connectivity at the bottom (where a PC or hub may be placed).

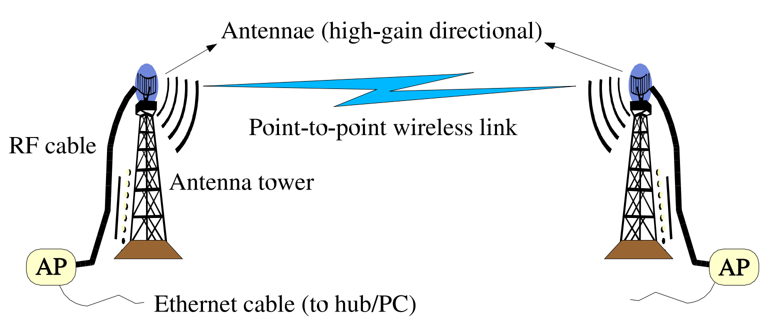

Figure 7: Long-distance link setup using bridge/AP at the bottom

The two possibilities are shown in Figure 7 and Figure 8 respectively. The pros and cons of these two alternatives are as follows. The main advantage of the second configuration is the elimination of the cumbersome as well as expensive RF cable. Our experience has shown that long RF cables are prone to breakage and wear-outs, especially at the connecting ends. Further,

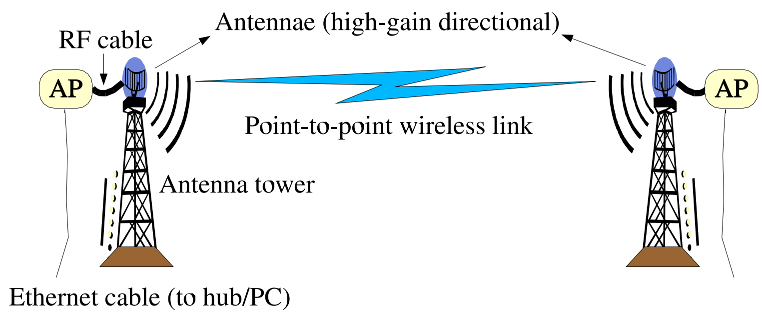

Figure 8: Alternative long-distance link setup using ap/bridge at the top

RF cables cause significant amount of fall in signal strength (about 5-6dB for about 40-50m). In comparison, in the second setup, ethernet cables are cheap, flexible, easy to handle. Also by placing the bridge/AP close to the antenna, any additional cable loss is prevented. This second setup is also helped by the fact that the bridge/AP can potentially take in power through the ethernet cable (e.g. Cisco Aironet 350 series bridge/APs [3]).

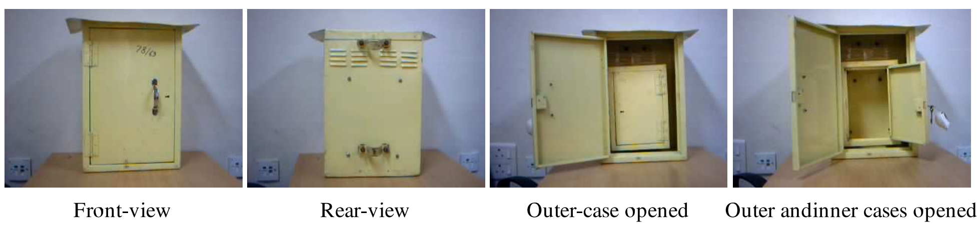

Figure 9: Weather-proof casing for equipment atop an antenna tower

The primary disadvantages of the second setup are two-fold. First, we need a weather-proof casing for the bridge/AP to stay atop the tower. We have effectively addressed this issue by designing and testing a weather-proof casing. This is shown in Figure 9. We have successfully run such a configuration for extended periods of time (several days at a stretch). The second disadvantage is something we still need to take a closer look at. This has to do with the fact that an bridge/AP placed atop the tower is not only close to the intended antenna, but also other antennae – to other directional links from the node. This may potentially cause interference problems when two links at a node are operated at the same channel. This is however not a problem at least with the current CSMA/CA protocol in 802.11, since there is the natural back-off mechanism. This factor however does need to be kept in mind for our proposed modified MAC protocol for the DGP network.

Antennae alignment procedure for a long-distance link

One of the important operational procedures in a point-to-point link setup is the antenna alignment procedure. The procedure has two main steps, as described below.

1. We first take GPS readings at either end of the link. Using this, we then calculate the angles at which the antenna atop the tower at either location should be oriented. We then mount the antennae and orient them along these approximate directions, using a compass. If we are using a parabolic-grid antenna, we usually mount it such that it has good directionality in the horizontal plane (as opposed to the vertical plane).

2. The second step involves a finer alignment procedure as follows.

- We have the 802.11 radio equipment at either end, and establish a link. To be on the safe side, we use maximum power settings at this stage. This is to account for any minor misalignments to begin with. At one end, we have a transmitting 802.11 bridge/AP. At the other end, we have a receiver. We then monitor the received power level by setting up a sniffer device. A sniffer captures the various 802.11 packets in the vicinity, and records/reports various parameters, including the received power level.

  For our sniffer, we setup a laptop with a Cisco Aironet PCMCIA card, and use the Airopeek software to make it function as a sniffer. The sniffer itself is setup using an RF power splitter, so that it accurately estimates the power received by the actual receiver.

- The antenna at the client is then fixed in the direction in which it receives maximum signal.

- Then using the same setup, we perform fine-grained alignment of the antenna at the transmitter side until we get the maximum possible received power.

Recently we acquired an antenna alignment kit, which essentially replaces transmitter/receiver setup we have. The kit has a specially-built transmitter and receiver; with the receiver fitted with a power reading meter. This eliminates the need for the laptop sniffer device as well. A procedure similar to the above procedure can be performed to align the antennae at either end. This is shown in Figure 10.

Figure 10: Antenna alignment kit

Once the alignment has been done, we switch the transmission power to a level such that we get "good" reception. Normally we use 100 mW transmit power that is the maximum level provided by the Cisco A350 series bridges.

Sometimes it so happens that the initial approximate alignment procedure is not good enough, and the point-to-point link fails to form initially. We have experienced this especially for some of very long links in our testbed. The challenge faced here is the "out-of-band" communication required before the link has been established, to establish the link itself. We have a team of people at either end of the link. When one team has to communicate to the other to move the antenna either way, or such, we require an out-of-band mechanism. When cellular phone connection is available at both the locations, we use it for such out-of-band communication. When such cell-phone coverage is absent, as in the case of remote villages, our solution is the following. We simply rely on the *SSID* feature provided by the 802.11 MAC management protocol. We encode any message to be communicated in the SSID as an English text and have the other end pick up the message. Of course, once the point-to-point link has formed, any communication between the two ends, required for fine-grained alignment, can use the wireless link itself.

Forming the mesh network from the point-to-point links

The above discussion pertains to a single link in the network. The multi-hop mesh network is made up of such point-to-point links. In our current setup, connectivity beyond a link is currently achieved through standard the ethernet bridging protocol, implemented by the bridge/AP equipments at each node of the network. The various bridge/AP equipments at a node are connected with each other via an ethernet hub. The same ethernet hub can also be used to connect a local PC at a node, for connecting it to the network. Such a setup is schematically shown in Figure 11.

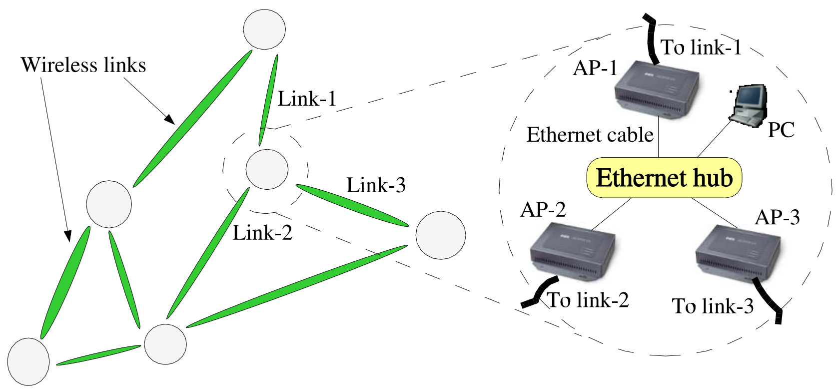

Figure 11: Using bridging to form the mesh network from the collection of wireless links

In this bridged network setup, one of the bridges is designated as a root-bridge, and performs the function of the root in the ethernet bridging protocol. We have used a bridged configuration for the primary reason of ease of setup. An alternative is to use a routing protocol. For this, we would require a router at each node of the network. We are working towards such a solution in the future, since routing provides much more flexibility than bridging. We have already experimented with devices which implement both the wireless radio as well as routing functionality (Lucent COR/ROR).

3.2 Experience with Various 802.11 Radio Equipments

Above, we have described in abstract as to how the mesh network is setup. Here we provide specifics of the 802.11 equipment used.

As mentioned, the use of 802.11 bridges is the most common setup for our point-to-point links. We have experience with equipments from various vendors: Cisco, Lucent, D-Link, Hotware, and X-Lin. We have however found that these radio devices in bridge mode are not inter-operable. That is, Cisco bridges work only with Cisco bridges, Lucent only with Lucent, etc..

Note that this inter-operability problem is only with bridge-bridge links, and not with bridge-client links – for which inter-operability is actually quite good. We have found the reason behind the bridge-bridge non-inter-operability to be the following. The commercial APs in our testbed use a proprietary (non-inter-operable) polling-based MAC protocol for sharing the bandwidth of the half-duplex point-to-point link. The polling-based scheme is more efficient than the contention-based CSMA/CA protocol. (Note that this however does not handle multiple links – while the MAC protocol modification we propose can handle multi-link, multi-hop scenarios also.)

Cisco bridges support clients while in bridged mode; other bridges do not have such support. Cisco bridges can also be setup as an AP, or as a client device. There is a separate Cisco AP device also. Cisco bridge/APs allow to set a "distance" parameter which they use to set the timeouts appropriately. Interestingly, we have found in our testbed that if this distance parameter is underestimated in configuration by more than 5km, the AP fails to form the link. This is likely because of the incorrect estimate of the ACK timeout.

Pigtail, with Lucent card

Rabbit-ears, with Cisco card

Figure 12: Pigtail and rabbit-ear connectors

Among client devices, currently two PCMCIA cards provide connectors for external antennae. One is Cisco Aironet, and the other is Lucent Orinoco. We use "rabbit-ear" connectors for Cisco cards, and "pigtail" connectors for the Lucent Orinoco cards. These are shown in Figure 12. We have found that Lucent Orinoco cards work exceptionally well with any bridge/AP for long distances.

Apart from the bridge-bridge links, we have experimented with two other configurations, as below.

AP-Client Link

We testbed an AP-client link mainly for the reason that AP/Client 802.11 devices are currently far cheaper than the 802.11 bridges (such as Cisco A350). While testing the AP-client link, we also included experimentation with antennae at various heights. In the test setup, a parabolic grid antenna was used at the AP end (Mandhana), and an omni directional antenna was used at the client end. The height of the antenna at the client end was varied with a variable height mast. Also, the client was placed at different distances, along the direction of Bithoor, from the AP at Mandhana. We have the following results so far.

- If clear line-of-sight (LOS) is available, the link works upto 4Km, with client antenna height 5-15ft.
- The link works upto 1Km with slight vegetation in LOS, and with client antenna being 5-15ft.
- The link works for as long as 6.5Km. if the client antenna height is set at 15m.

Use of Host-AP

Host-AP is an open-source project [6], which has developed a Linux-based device driver to make a client 802.11 device work as an AP. This is possible since the difference between an AP and a client is essentially in software. Using the Host-AP driver and an Intersil Prism-based card (`www.intersil.com`), we have been able to test bridging, routing, as well as the Wireless Distribution System (WDS) mode of 802.11 operation [29] – while working indoors. We have also tested this configuration for a long-distance Host-AP-Client link – with the Host-AP at the IITK site, and the client at Lodhar. We have found that this configuration is possible to setup with only Lucent Orinoco cards as clients to connect to the Host-AP.

3.3 A Discussion of the Various Costs Involved

Since low cost networking is one of the primary goals of DGP, an understanding of the current cost break-ups is important. This would then help in further improvements in terms of system cost reduction.

We can distinguish between cost factors which are per-node, and those which are per-link. We discuss these in turn.

Per-node costs

The main per-node cost currently is the *antenna tower*, which is required for line-of-sight clearance. The current costs of towers at various heights are tabulated in Table 2. Our experience has been that about 30m clearance is usually required for 7-8km links. This is the rough figure as seen in our testbed links, as well as the experiment with AP-client link given in Section 3.2. This is also the approximate value given in [28]. We note that 7-8km is the average village-village distance in India.

Item	Cost (approx. U.S.$)
Antenna tower	
15m	$1,500
20m	$1,750
25m	$2,250
30m	$3,250
40m	$5,750
Antenna mast	
10m	$85
15m	$130
20m	$170
Ethernet hub	$60
Power backup	$1,150

Table 2: Per-node costs (approximate)

Where possible, we can use a mast instead of a fully-built tower. Masts can easily reach heights of up to 15m, and can be especially effective when mounted on already existing buildings in village locations.

The other per-node costs for the communication network currently include: (a) an ethernet hub, and (b) a power backup if required. The cost of these, also given in Table 2, is minimal as compared to the antenna tower cost.

Per-link costs

The per-link costs include the 802.11 radio, RF cables, and ethernet cables at either end. The current costs of various possibilities for the 802.11 radio are as given in Table 3. We see that Cisco bridges are the most expensive – perhaps because of the specialized functionality they perform. This is one of the reasons we have also been experimenting with other possibilities for the point-to-point links (Section 3.2).

For the Host-AP configuration described earlier, a Prism card and a Single-Board-Computer is required. Although we have experimented Host-AP only with a PC, we include the cost of the Sokaris SBC in Table 3.

Item	Cost (U.S.$)
Bridges	
Cisco	$1,000
Lucent	$900
D-Link	$85
Pheenet	$800
X-Lin	$425
Access Points	
Cisco	$600
Lucent	$1,100
Client devices	
Cisco	$125
Lucent Orinoco	$60
Prism	$90
D-Link	$120
Sokaris Single Board Computer (SBC Net 4521)	$270
RF Cable	$4 per metre
Ethernet cable	$0.50 per metre

Table 3: Per-link costs (approximate)

We also see in Table 3 that RF cable cost is higher than that of ethernet cable. Going forward, we intend to experiment further with the setup in Figure 8 (Section 3.1), which uses long ethernet cables and avoids long RF cables.

Finally, we also note that the per-link costs are also small as compared to the antenna tower cost.

Overall system costs

The overall system cost also depends on how many links are used to connect a set of nodes. We have explored the technical aspects of how a network may be planned such that we get good performance while at the same time minimizing the number of links provisioned. The details of these are given in [31], and the same is summarized in Section 4.

Spectrum costs

While spectrum costs are not part of our non-commercial, experimental testbed, these costs would become a significant factor when this technology is deployed commercially. The current spectrum costs are X for up to 5km, and 4X for 5-60km, for outdoor use of the spectrum in India (X is about Rs. 5,000 per 25KHz). In this respect, we advocate two things: (a) that a single 802.11 channel be licensed for operation of the long-distance links; and the spectrum costs decided based on village location and the current economy there, and (b) operation in the last hop be left license-free, at least for village locations, even though such operation is technically outdoor and hence licensed currently.

The licensing of the channel for long-distance links would ensure that the network operator does not see unexpected or malicious interference from other users or operators. The delicensing in the last hop would allow operation of 802.11 as in the 802.11 community networks (see [2]).

3.4 Status of the Various Nodes in the DGP Network

The DGP testbed effort was started in Apr'02 and the testbed extent has grown steadily since then. We have so far setup up to 14 links between 12 different locations, with the longest point-to-point link going as far as 38km. The testbed is depicted in Figure 4. With the addition of the Sarauhan-Rajajipuram link, we have been able to demonstrate multi-hop connectivity up to about 80km. Over the last one and half year period 32 bachelor and master's students have carried out experiments on this network.

All locations except 'IITK', 'Rajajipuram/Lucknow', and IISTEM/Banthar are rural villages. None of the villages had prior Internet connection. The villages have little or no cell-phone coverage, and the village at 'Saroha' had no land-line telephone as well. Currently, the only node in the network that has wired Internet connection is site 'IITK'.

Link	Date of activation	Length (km)	Remarks
IITK-Lodhar	Apr 2002	2.3km	Lucent ROR at Lodhar, Currently not active
IITK-MS3	Jun 2002	5km	Client card in PC at MS3 Currently not active
IITK-Mandhana	Jun 2002	5.1km	–
Mandhana-Safipur	Sep 2002	22.5km	Cisco bridge at both ends, Weather proof box at both ends
IITK-Bithoor (Dhruv Teela)	Nov 2002	12km	Currently inactive
Safipur-Sarauhan	Dec 2002	17.3km	Cisco bridge in weather-proof boxes on both sides
IITK-Sarauhan	Jan 2003	37km	100mW AGC amp at IITK, 1W AGC amp at Sarauhan
Sarauhan-Rajajipuram	Mar 2003	39km	Cisco bridges on both sides 250mW AGC amp at Sarauhan 100mW AGC amp at Rajajipuram
IITK-Banthar	Jun 2003	23km	Lucent ROR with 100mW AGC amp at both ends
Mandhana-Bithoor (FAB lab)	Nov 2003	12km	Cisco bridges on both sides,
Banthar-Sawayajpur	Dec 2003	22km	Hotware bridges on both sides
IITK-Bithoor (FAB lab)	Jan 2004	12km	Cisco bridges on both sides
Sarauhan-Rasoolabad	Mar 2004	5km	X-Lin bridges on both sides, Currently not active

Table 4: Details of links in the DGP testbed

Most of the long distance links are over flat terrain. The heights were chosen so as to avoid any obstructions such as buildings or trees in the Freznel zone and thus avoid multi-path. The heights of all towers are about 40m, except the one at 'MS3' where the height is about 12m. We recently upgraded this to a 30m tower. We use off-the-shelf 802.11b Access-Points (APs) at all the locations. For high directional gain, we use off-the-shelf parabolic-grid antennae [7].

The APs are currently setup in bridge-mode. In the future, we also plan to explore network-layer routing in the wireless access network.

We document the details of the different nodes/locations and links in the DGP testbed. This is given in the form of a table – Table 4 for the different links. The rows correspond to the various links, and the columns correspond to their various properties. The node-specific details are given in Table 5. The rough geographic locations of the various nodes are as given in Figure 4 (Section 1).

Node	Date of activation	Latitude	Longitude	Remarks
IITK	Apr 2002	$26^o30'46.944$”	$80^o13'59.166$”	Faculty building top, Wired connection to Internet from here
Lodhar	Apr 2002	$26^o31'15.000$”	$80^o12'44.000$”	6ft mast on rooftop, Currently inactive
MS3	Jun 2002	$26^o33'41.700$”	$80^o14'06.312$”	36ft mast, 30m tower as of Jun 2004
Mandhana	June 2002	$26^o33'31.638$”	$80^o13'37.920$”	40m BSNL tower, Site has power inverter with timer control
Safipur	Sep 2002	$26^o44'26.076$”	$80^o20'54.978$”	40m BSNL tower
Bithoor (Dhruv Teela)	Nov 2002	–	–	20ft mast on rooftop, Currently inactive
Sarauhan	Dec 2002	$26^o43'23.772$”	$80^o31'19.992$”	40m tower, has DG set and a solar power pack; usually does not use DG set for day-time operation, Site has VoIP (FXS) for telephony
Rajajipuram (Lucknow)	Mar 2003	$26^o50'29.172$”	$80^o52'53.358$”	40m BSNL tower
Infokiosk at Mandhana	Apr 2003	–	–	Client card in kiosk catches leakage from point-to-point link, Currently not operational *Not shown in Figure 4*
Infokiosk at Safipur	May 2003	–	–	Client card in kiosk catches leakage from point-to-point link Currently not operational *Not shown in Figure 4*
Banthar	Jun 2003	$26^o28'56.300$”	$80^o27'48.700$”	25m tower, Lucent ROR at site also works as router; site located at IISTEM
Bithoor (FAB lab)	Nov 2003	$26^o35'57.200$”	$80^o16'26.400$”	25m tower on rooftop, Site has VoIP (FXS) for telephony
Sawayajpur	Dec 2003	–	–	30m tower
Rasoolabad	Mar 2004	–	–	10ft mast, Site has VoIP (FXS) for telephony, Currently not operational

Table 5: Details of nodes in the DGP testbed

4 Research Issues Addressed

In Section 3, we had discussed how we have addressed the various operational issues. Throughout the project, we have also addressed technical issues which have needed attention, as described in Section 2.2. Our contributions have been documented extensively in various writeups and theses (see Appendix B) – we summarize these briefly here.

Empirical Path Loss Model for Outdoor 802.11b Wireless Links

One of the first things we worked on was to experimentally measure and empirically model the path loss behaviour of 802.11b links. When we started out, there are various models available for path-loss in micro-wave frequency range. However they do not cater specifically to the 2.4 GHz range of 802.11b. They are either for frequency range below 2 GHz or they cater to long distances (above 1km). In the MTech thesis taken up by Rajesh Gandhi [22], empirical path-loss models were provided for the following scenarios: (a) campus wide networks, (b) along the roads in straight line, and finally (c) long distance point-to-point links (1-40kms). SNR versus throughput curves were studied experimentally, and a comparison of these was made with standard SNR versus BER curves available for the modulation techniques used.

The main results from the study are the following:

- For the long-distance links, the path loss PL is given by $PL = FPL + 3 + 0.15 \times D$, where FPL is the free-space path-loss, and D is the distance covered by the link, in kilo-metres.
- For 802.11b links, the path loss along a straight "road" is typically given by $PL = FPL + 15 + 20 \times D$, where FPL and D are as mentioned earlier.
- Similar expressions for the path loss were obtained for the cases where the 802.11b transmitter was placed at a varying heights (see [22] for the details).

For the long-distance point-to-point links, the first of these results above is the most significant. If effectively means that the free space path loss model with a 4-6dB correction is a good predictor for the path loss over such links.

WLAN Watch: A Step Towards The Study Of 802.11b WLANs

In DGP, there is a need to study 802.11b in-depth, to suggest any possible performance enhancements and improvements at the MAC and/or application layers. This requires appropriate software tools for trace collection, performance analysis, and such purposes. However, there are no such adequate software tools available to carry out these experiments. As part of his MTech thesis, Imtiaz Ur Rahaman built tools that can assist the study of the 802.11b networks and lead to possible performance enhancements [30].

The tools can provide the signal and noise level of each packet which in turn could be used, for example, to find the signal level range in which most of the packets over air get corrupted and the signal level below which most of the packets get lost. To assist the study of the 802.11b networks, the tools also help in establishing a relationship among packet size, packet corruption ratio, throughput, power of transmitter, etc.

The tool built as part of the thesis, called WLAN Watch, provides all this information among other things. WLAN Watch is generic and currently supports Cisco Aironet 350 Series and Intersil Prism-II NICs.

SynOp: Simultaneous Synchronous Operation

The 802.11 MAC protocol was originally designed for an indoor LAN setting. It is well known that it performs poorly when used in multi-hop settings [33, 21]. The main problem is one of *exposed nodes* in the wireless setup. Multiple parallel transmissions are not possible, since the carrier-sense-based CSMA/CA protocol "thinks" there are other interfering transmissions.

In the DGP network setup, we use multiple directional antennae, and multiple radio devices per node (one per link at the node). Ideally, it should be possible to operate these links simultaneously. However, since each interface's transmission is exposed to the other interfaces at a node, a problem similar to the exposed node problem occurs. As a result, only of the links can operate at a time – unless they are assigned different, non-overlapping channels.

Now, there are three non-overlapping channels in 802.11, and each is 25MHz wide. It is spectrally inefficient to use many channels. Ideally, we would like to achieve maximum parallelism among the point-to-point links using just a single channel of operation.

We have come up with and explored a key idea in this regard: we operate the links at a node such that all the radio interfaces at the node are either all transmitting (simultaneous transmission: SynTx), or all receiving (simultaneous reception: SynRx). We have shown theoretically as well as experimentally on our testbed that such operation is indeed possible [23].

2-P: A MAC Protocol Based on SynOp

Based on the simultaneous synchronous operation flexibility outlined above, we have designed a MAC protocol as a replacement for the 802.11 CSMA/CA MAC. The 802.11 CSMA/CA protocol is designed for a case where there are many nodes contending at random for access to a shared channel. However, for the long-distance point-to-point links, there is no *arbitrary* contention. In this case, a time-division based scheduling approach (TDMA) is more appropriate.

In addition, due to the flexibility offered by SynOp, *spatial-reuse* is possible. This is known by the generic term Spatial-reuse TDMA (STDMA) scheduling (e.g. see [32]). This motivates us to propose the following STDMA-based MAC protocol for the point-to-point links in the DGP network. We simply switch each node between a transmit phase, when it transmits along all its links, and a receive phase, when it receives from all neighbours along its links. When a node is in transmit phase, its neighbours are in receive phase, and vice versa.

We term the above protocol as 2-P, to indicate the two phases. We have designed mechanisms for implementing this protocol with minimal changes to existing 802.11, to preserve the cost-economics of the equipment. The detailed design of the protocol can be found in [31].

TeNs: The Enhanced Network Simulator

Network Simulator (NS) is a discrete event simulator targeted at networking research, and has emerged over the years as the most popular research tool in networks. IEEE802.11b has been implemented in NS; however the implementation is not quite realistic or complete. In the Bachelor's project done by Sabyasachi Roy and Ashwini Kumar [27], we have improved upon the implementation of IEEE802.11b in NS, and also added additional novel features to improve its modeling capabilities. This project built on the initial work done by Ankur Khandelwal [26]. We call the improved Network Simulator as The Enhanced Network Simulator (TeNs) [17].

Specifically, the enhancements to NS include:

- *Adjacent and co-channel interference:* 802.11b defines eleven overlapping channels. Hence adjacent and co-channel interference are possible, and this is not modeled in NS. In this project, we have implemented as well as experimentally validated such modeling in the simulator.

- *Directional antennae:* NS has support only for omni-directional antenna, and with the use of directional antenna in DGP, there is a need for directional-antenna support in NS.

Various kinds of directional antennae have been added to TeNs, and it is flexible enough to allow for further additions too.

- *Gray region and temporal variation:* Gray region refers to a region, which is far enough from the signal source, so that the signal strength received is near to, but above the receive threshold of the receiver. It is observed that the signal strength (which determines the throughput) does not go to zero abruptly, rather it shows a steady decline accompanied by minor fluctuations in that region. Wireless communication also shows the phenomena of random temporal variations. These refer to the random variations in signal strength (and throughput) at a point, which is observed in real life WLANs. This is natural to expect, because of wireless mode of communication, which is subject to interferences like multipath, diffraction, etc. Gray Region has been found not to be modeled adequately in NS. We have modeled as well as validated gray region and temporal variations in our enhanced simulator.

- *Adaptive data-rate transitions:*. Adaptive data-rate transitions, is a mechanism by which a signal sender (like an access point), can switch between standard data-rates (11, 5.5, 2, 1 Mbps) to reduce bit error, and hence improve performance. This feature has been implemented and validated in the enhanced version.

- *Multiple interfaces in wireless nodes:* The current NS implementation allows only one wireless network interface per node. In fact, the node entity and the interface are tightly coupled in the existing implementation. We have implemented multiple interface support in TeNs.

- *Multihop support:* Multihop support is also required for wireless nodes in DGP. This requires forwarding of packets from a node to another in a mesh kind of network. Note that current ad-hoc routing protocols incorporate this kind of forwarding, but it involves lots of overhead, which is not what is encountered in a static wireless multi hop mesh network like DGP. This has also been added in TeNs.

Topology Construction Issues

One of the operational issues which requires significant technical analysis is that of topology construction. By this, we mean – given a set of village nodes to be networked, what is the most appropriate network topology to construct. There are many factors which go into the choice of the network topology. These include (a) the overall cost of the network should be minimized, and (b) the network should have some properties like fault tolerance (connected in the face of some link or node failures). While these two factors are common for any network, the factor of wireless interference is an additional aspect to consider in our DGP network. We seek to minimize interference among nodes while setting up the point-to-point links. Such interference can occur despite the directional links since the directionality is never perfect and there are leakages in directions other than the main direction.

We have designed algorithms for such topology construction [31]. The overall mechanism consists of choosing links to be formed based on a set of heuristics which seek to minimize cross-link interference. We also include mechanisms by which only those set of links which are required for connectivity at a particular time are turned on. The overall topology thus constructed can also enable the proposed modified MAC protocol, 2-P, as well as the related SynOp flexibility. We have simulated the topology construction mechanism successfully on the set of villages in the Durg district in Uttar Pradesh. The set of villages we considered had a total 128 locations.

Power Allocation for Interference Minimization

Apart from topology construction, the other operational consideration which requires an in-depth technical analysis is that of power allocation. By this we mean the decision of the transmit power levels at each of the nodes for each of the links. Here again, the goal is to minimize interference, and this mechanism goes hand-in-hand with the topology construction mechanism. For a given topology, we wish to determine the transmit power levels of the various transmitters in the network. In the Bachelor's project done by Paul Ipe [24], we have modeled this problem as a set of linear equations to be satisfied. The set of linear equations is then solved using a standard LP (linear programming) solver software. Through this mechanism, we have been able to successfully allocate transmit powers (without mutual interference) for the 128-node network topology for the Durg district, as generated by our network topology construction mechanism.

5 Experience with Application Services

In the DGP project, there has been a good forcing function throughout for the usage of our network connectivity services. Other projects at the Media Labs Asia, IIT-Kanpur, have been using our network at various locations. These different projects deal with different rural application scenarios. The details of these are as follows.

- **Infothela:** The Infothela project [9] aims at providing mobile Internet services through the design of an appropriate light-weight mechanical "thela", or "rickshaw" which can house a computer and related equipment. It banks on local village coverage as provided by the DGP network. Currently there are two village locations where Infothela operates: Goraha and Bithoor. In both cases, the connectivity is through an 802.11 client device at the Infothela, which operates through leakage off an existing long-distance link. The client device has omni-directional antennae. Such a setup is shown in Figure 13. The Infothela at Goraha operates using leakages off the IITK-Mandhana link, and the one at Bithoor uses leakage off the IITK-Bithoor link.

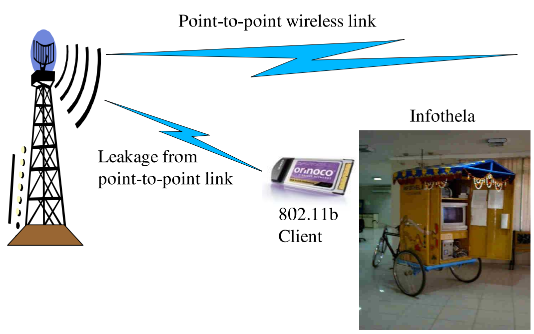

Figure 13: Infothela, and the connectivity setup using DGP

Note that an alternate setup is also possible – which uses a separate omni-directional antenna at a village for local access, instead of using the leakage from the point-to-point link.

The main use of the connectivity thus achieved at the Infothela, has been mainly for Internet browsing so far.

- **Digital Mandi:** The goal of this project [4] is to provide a web-based portal for e-trading of agricultural products. Buyers and sellers can find each other using the portal. In addition to the portal, trading and information access are also possible through e-mail and Short-Message-Service (SMS). This project has partially used Infothela for Internet connectivity. In addition to access of the web-portal developed by the group, the Internet connectivity provided by DGP has also been used for instant messenger-based mandi-database access.

- **Tele-medicine:** The aim of this project [12], as the name suggests, is to provide health services through the Internet. The network is used for interactive video conferencing as well as for access of the web-portal developed by the group [13]. The tele-medicine group has so far primarily operated from the Mohanpur village. Connectivity at the village is through a client card with an omni-directional antenna atop a 17 feet mast. The antenna picks up leakage from the IITK-Lodhar point-to-point link – just as in the setup in Figure 13.

  Our experience with the video-conferencing application so far has been that the quality is not good enough for use for tele-medicine yet. We suspect the reason for this is that the connection goes via the IITK LAN, which has many users. Although the LAN is well provisioned in terms of bandwidth, there is no guaranteed Quality of Service (QoS) – which is essential for video applications. Also, we specifically think that the 802.11 wireless hop is not the problem. This is because the link offers plenty of bandwidth – 11Mbps as compared to about 200Kbps or so required for the video at maximum. Also, there is no high contention for the wireless link.

  We are in the process of confirming these conjectures. Going forward, interactive video conferencing for use by the tele-medicine group is one of the applications we intend to focus on in the DGP project.

Apart from the above services, we have also experimented extensively with the Voice-over-IP (VoIP) application. We see this application to have a good potential to provide a useful service at remote villages. The rest of this section focuses on this application. In Section 5.1, we describe the setup and how the service operates. Then in Section 5.2, we present the details related to the economics of this particular service, and the challenges in making it commercially viable.

5.1 Voice over IP (VoIP)

In DGP, we have experimented with some economic models for rural India based on the applications of the wi-fi technology. For this purpose connectivity has been given to some private parties at Mandhna and Safipur to establish Internet kiosks. DGP extended the multi-hop network to Sarauhan village in Dec 2002. As Sarauhan is a small village located in a very remote area of Unnao district, we have found that nobody is interested in Internet surfing. Villagers are found to be more interested in voice communication using telephone. In due course of time we realized this – a phone facility is the actual need of the villages rather than Internet and email. The DGP team then decided to open an extension counter of a nearby PCO at the village using a VoIP communication over a wi-fi link.

Proof of Concept: With the VoIP equipments installed at IITK (FXO unit which takes in the telephone line) and Sarauhan (FXS unit which serves the normal phone), we tried phone calls to outside world using the phone line of IITK and became successful. To check the calling potential of the villagers, we ran a one day PCO using the same line on 4 Jan, 2003 and found that the village strongly needs a PCO counter.

Implementation Challenges: Supporting the laws of human psychology, it has been observed that people love to call if it is free to them but most of calls were just for the sake of calling. Whenever they are charged for the calls, they were quite conscious about the bill. So it was necessary that they can have the bills of their calls.

The billing system is built in the VoIP boxes but they meter only the call from one VoIP box to other VoIP box. They do not meter the outgoing calls on the PSTN line. At the same time, implementing a billing server is quite a costly affair.

Our solution to this problem was as follows. We decided to use a PSTN line of nearby PCO to provide the service and to print the bills as usual using the PCO machine at the PSTN line of PCO.

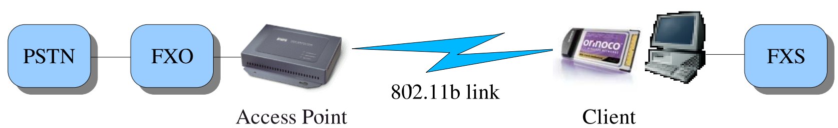

Figure 14: Setup for experimental PCO extension counter

When someone wishes to make a call from the experimental PCO extension counter, he/she will dial the number of the VoIP phone at the PCO end. The person at the PCO will then make a PSTN call on behalf of the caller and gives the line to the caller using a two line phone. At the end of call, the PCO machine prints the bill and the person at the PCO communicates the same to the person at the extension counter who collects the cash from the caller. At the end of day the people exchange the revenue collected and the bills. Such a setup is schematically shown in Figure 14.

Implementation: Ekta PCO of Rasoolabad village became ready to take part in this experiment and a WLAN link was made connecting the PCO to Sarauhan village and the experimental PCO extension counter was inaugurated on Jan 03, 2004.

Future Work

Two types of services are possible using this system:

Domestic Phones: This system will have one-to-one mapping of PSTN line and client connection hence it will be two way (incoming and outgoing). The subscriber will get the dial-tone of PSTN line on lifting the handset and be able to make calls on that PSTN line. In case of incoming calls, the FXO will pick-up the call and forward it to corresponding FXS. The billing system of the PSTN network has to be adjusted in a way that it should not bill a call until the end user picks it up (not the FXO).

Public Phones: In this system many users (FXS) can use same PSTN line as it is one way (outgoing only). This system is similar to VCC system of DOT. When a person lifts the handset of phone, he will get a dial-tone of FXS. He/She will dial a number to access the FXO. The software at the FXO will validate his/her login by asking his card number and allow/disallow making phone calls. The software for the FXO is yet to be developed.

We plan to explore these possibilities going forward. Some of the other issues we plan to address are as follows.

Power solutions are a requirement – practically Sarauhan village gets the AC mains power for only 2-3 hrs in a day, that too of voltage ranging from 90 to 120. Three things were planed to get rid of the problem: (1) A stabilizer (2) A solar power system with battery bank (3) A diesel generator set as backup. The service was non-operational due to this issue for a long time.

The latest status on this is that we have removed equipments from Rasoolabad and placed them at IITK, providing a connection from the Media Labs Asia lab. In addition we have Solar powered battery pack at Sarauhan. We also have a standard metering device at Sarauhan to allow proper billing there. This has been in operation for over a month since Jul 2004.

Hardware integration of Wi-Fi client and VoIP FXS modules into a single unit is to be done to make the device user friendly. Once this is done, the user will only require attaching an antenna and a standard analog telephone instrument to the device and powering it up.

5.2 Experience with the PCO extension counter at Sarauhan

Information kiosks that rely entirely on donors' funds for their operation have less chances of becoming viable mechanisms for development. Instead, there should be a focus to develop financially sustainable and commercially viable kiosks with clear business models. From this perspective we conducted an experiment to develop the PCO at Sarauhan.

Earlier, the villagers at Sarauhan used to walk about 5 kms away to Munshiganj or Rasoolabad to make a phone call from PCOs situated there. We set-up a wireless link between our center at Sarauhan and a PCO at Rasoolabad and connected them with VOIP phones. An operator at village now takes request from a customer at the village and let the PCO operator know the number he/she wishes to dial and connects them. Once the call is complete, the PCO operator tells the charges which is collected by the operator at village. As per the current arrangement, 20% of total charges are retained by the village operator and rest given to the PCO operator. This revenue sharing still allows the PCO operator to earn on each call.

We made a comparison of infrastructure cost apportioned over a period, operation cost and the income generated after paying to the PCO operator. Based on the data generated so far, the results are as shown in Figure 15.

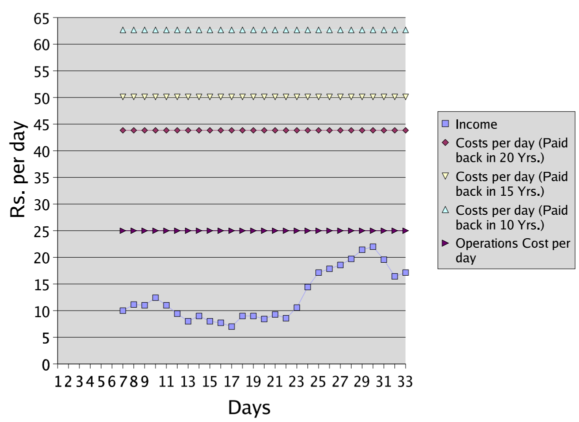

Figure 15: Experience with costs of running a PCO counter

In the figure, operation cost involves wages of operator and other miscellaneous expenses. Infrastructure cost includes the cost of VoIP and phone equipments, radio devices and mast etc. at both the ends. It also includes the cost of solar power system (the most expensive element of total cost) at village end. Costs-per-day includes both the above costs and is linearly apportioned without interest over the pay back period.

In our experience, there are several factors which affect income, as follows.

- *Weather conditions:* Villagers are unable to come to PCO during rain or extreme cold.

- *Power failure* at Village Counter, PCO, or at the BSNL exchange causes disruption in service.
- *Hours of operations of village counter:* Villagers make more phone calls during the night hours. Currently we are incapable to operate regularly during the nights.
- *Seasonal factors:* The village has many orchards. At the time of trading season villagers are likely to make more calls. Similarly other factors like marriages, etc. may also affect the income.
- *PCO like privacy and billing facility:* This will make the facility friendlier and trustworthy to use.

Our expectation of daily income is around Rs. 50 after stable operation, which shall make this solution economically self sustainable with a reasonable payback period of Infrastructure cost. Providing service through an Infothela with a regular schedule of visits to nearby weekly markets may further enhance the income generation capability as a PCO extension.

6 Challenges Ahead

In this section, we describe the various issues which we intend to explore in the near future. This is based on the findings in the DGP project so far. Like our earlier work, we subdivide our proposed work in three different categories: (a) operational issues, (b) technical issues, and (c) services. We discuss these in turn.

6.1 Operational Issues

The operational issues we intend to focus on in the next two years stem from the findings in the project so far. The specific areas of investigation we propose are as follows.

- *Network monitoring tools:* The 802.11-based mesh network we have built is quite different from wired networks. The fact that the network is wireless-based makes it quite different. In the 802.11 community in western markets, people are only beginning to understand the various operational issues. This is especially so for outdoor usages of the technology, such as in Community Networks. The need for monitoring and management is especially important for our network since it is intended to be an operated access network. We intend to develop various monitoring tools, as we work on the technical issues and application issues. This would be with the goal of making the task of operations and management easy. An example of this could be the following: a pocket-PC (iPAQ/Simputer) based sniffing tool which can be used to measure and report received signal strength at a village location – for diagnosis of any wireless network problem.
- *Low-cost antenna towers:* Our experience with the DGP project and our testbed prototype has shown that the antenna tower cost is likely to be dominating when we seek to build a multi-hop mesh network. Further, we may require antenna towers of about 25-30m height for line-of-sight clearance between two villages. The model of antenna tower we have been working with so far is akin to the cellular-phone towers (e.g. used by commercial companies such as Reliance, BSNL, etc.). We believe that this is a heavy-weight design for the intended purpose of simply mounting 2-3 antennae at a village location. We plan to explore how the antenna tower cost may be cut, but using lighter designs. The current design allows for humans to climb on top, for antenna

mounting/dismounting, alignment, etc.. If we can have mechanisms for antenna alignment on a lighter design such as a tower supported by guy wires, then lower cost can be achieved. We intend to explore such possibilities for antenna tower cost reduction.

- *Operational issues with the new MAC protocol:* We have designed a TDMA-based 2-phase MAC protocol as a replacement for the 802.11 CSMA/CA, since the latter is inefficient for operating the long-distance links. We intend to implement the protocol, and understand the operational issues involved with this. Here there are important details to understand, such as the required power levels of operation, head-room required in power-level settings to enable the 2-phase protocol, etc.. We also need to have a knowledge of what off-the-shelf equipments can be used for such operation.

- *Experience with 802.11g:* While we have so far focused mainly with the 802.11b variant of the technology, the emerging 802.11g variant promises to provide potentially higher bandwidths. However, it is unclear if it would be possible to achieve higher bandwidths for the long-distance links. It is also unclear as to which of the 802.11g equipments would even be able to operate in long-distance modes. We intend on experimenting with such cases using off-the-shelf 802.11g equipments.

- *Low-power PC:* In our experience so far, the availability of power at village locations is one significant factor which needs attention. To this end, we intend to experiment with solutions for low-power PCs. This can potentially be achieved by eliminating the high-end parts of the PC, and going for low-performance, but low-power components.

6.2 Technical Issues

The broad technical issues which we intend to look at in the future can be summarized as follows.

- *The 2-P TDMA-based MAC protocol:* We have so far completed the preliminary design of the 2-phase TDMA-based MAC protocol, as a replacement for the 802.11-based CSMA/CA protocol. This needs to be implemented, and various issues such as the setting of the different parameters in the protocol need to be studied carefully for understanding the performance behaviour. We initially intend to use the ns-2 simulator we have built to study the modified MAC protocol, and subsequently also use the testbed to study the performance behaviour.

- *Topology construction:* While we have designed abstract algorithms for topology construction, these need to be validated by using the testbed. Again we intend to use the simulator initially, and then the testbed itself for performance studies.

- *TCP performance studies:* The performance of higher level protocols such as TCP needs to be studied, and especially using the modified 2-P MAC protocol.

- *Performance studies for VoIP and Video applications:* Apart from TCP performance, we also need to understand and optimize the performance of real-time applications such as VoIP and Video. These studies too are especially important under the proposed modified MAC protocol. We expect this study to be directly relevant to our experiments with the application services we intend on looking at.

6.3 Application/Service Issues

We intend to focus on two main applications/services in this second phase of the DGP project.

- *Voice-over-IP application:* Our experience so far with this service suggests that we can optimistically expect a daily income of around Rs. 50 after stable operation, which shall make this solution economically self sustainable with a reasonable payback period of infrastructure cost. However, this is only a conjecture at this point, and going forward, we intend to experiment with this further. Among the things which will make the VoIP application more useful are the following: (a) a reliable power solution, (b) integration of the VoIP application with a pocket-PC like device such as iPAQ or a Simputer – this will also drastically reduce the power requirements as compared to a full-fledged PC, and (c) a solution for automated billing, which is currently done using a manual procedure. Many of these issues fall into the category of technical/operational issues – we intend to address these with the VoIP application in focus.

- *Video application for telemedicine:* The main hurdle in the use of video for telemedicine has been the lack of reliable 802.11 coverage from our testbed, for the tele-medicine team, at the exact location where they work. We hope that tools such as the wireless network monitoring tool will help us manage such coverage with greater ease in the future. This is one thing we intend to focus on with respect to this application. Another problem has been that we have not been able run good quality video on the 802.11 wireless links so far. Preliminary experiments have revealed that there is nothing fundamental about 802.11 that prevents reception of good quality video, and the issue may be one of higher error rates and jitters due to poor link engineering. Again we expect that the network monitoring/sniffing tools may be of help here. We also intend on experimenting with different video codecs to understand which one suits best for the telemedicine application we have in mind.

7 Summary of Contributions

The contributions of the Digital Gangetic Plains project thus far can be summarized as follows.

- The main contribution so far is the establishment of technical feasibility of the 802.11 technology for long-distance usage for the purpose of reaching access networks to rural villages. We now understand the various procedures for setup of a long-distance link. We have also built a lot of experience in understanding which of the off-the-shelf 802.11 equipments work in which configurations, and which of them do not.

- Cost is the most important factor in the building of a network for rural use. In our prototype testbed, we now have an understanding of the various costs involved in building the network. The main cost factor as we understand so far is the antenna tower required. Going forward, we intend to experiment with low-cost antenna tower solutions.

- One of the technical issues we have a good understanding of now, is the propagation of 802.11 radio waves in long-distance conditions. We have developed propagation models and also matched these up with measurements on the testbed. The detailed results can be found in [22].

- We have experience in building tools such as "WLAN-Watch", which enables MAC-level sniffing of nearby 802.11 traffic. This is documented in detail in [30].

- In multi-hop 802.11 networks, an important factor of efficiency is the number of links which can operate in parallel (in the same channel) without mutual interference. We have developed a novel mechanism to improve this efficiency. We term this as "SynOp: Simultaneous Synchronous Operation" – it enables the multiple directional links at a node to operate simultaneously. We have shown through experiments that SynOp is feasible (without actually implementing SynOp). The details of these are documented in [23].

- Off-the-shelf 802.11 equipments use the standard CSMA/CA (Carrier Sense Multiple Access, with Collision Avoidance) Medium Access Control (MAC) protocol. This is known to be inefficient in multi-hop 802.11 networks. We have the preliminary design of a far more efficient alternative MAC protocol, based on the SynOp flexibility mentioned above. Going forward, we need to implement and experiment with this enhanced MAC protocol.

- One of the significant contributions so far has been the development of enhancements for the widely popular network simulator, ns-2. These enhancements allow us to simulate networks such as the DGP testbed. We intend on using this simulator for the various design studies in the future. The details of this simulator, including software download, can be found at `http://www.cse.iitk.ac.in/˜bhaskar/tens/`.

- In addition to the above, we have also designed algorithms for planning the 802.11 mesh network prior to deployment. We have also looked into algorithms for the right allocation of transmit powers at the various links in the network.

- Finally, the experience with the testbed has helped us articulate the challenges ahead [19], which is an important aspect of research and development.

Acknowledgment

The Digital Gangetic Plains project would not have been possible without the effort of a whole team of people. We thank everyone who has participated in the project. We are grateful to Media Labs Asia for providing support for the project. The testbed effort required a lot of help from people at each of the testbed sites. While it may not be feasible to list every person involved here, we specifically wish to thank BSNL, Manaviya Shiksha Sanskar Sansthan (MSSS/MS3), and IISTEM for help at their testbed sites. We would like to thank Mr. Om Shankar Tripathi for giving us land and logistical support at Sarauhan.

References

[1] Census of India 2001. `http://www.censusindia.net/`.
[2] Community Wireless / Rooftop Systems. `www.practicallynetworked.com/tools/wireless_articles_community.htm`.
[3] Data Sheet: Cisco Aironet 350 Series Access Points. `www.cisco.com/warp/public/cc/pd/witc/ao350ap/prodlit/a350a_ds.htm`.
[4] Digital Mandi. `http://www.digitalmandi.net/`.
[5] Government of India Union Budget, 2003-04. `http://indiabudget.nic.in/ub2003-04/glance.htm`.
[6] Host AP driver for Intersil Prism2/2.5/3 and WPA Supplicant. `http://hostap.epitest.fi/`.

[7] HyperGain HG2424G 2.4 GHz 24 dBi High Performance Reflector Grid Antenna. `http://www.hyperlinktech.com/web/hg2424g.php`.
[8] IEEE P802.11, The Working Group for Wireless LANs. `http://grouper.ieee.org/groups/802/11/`.
[9] Infothela. `http://www.iitk.ac.in/MLAsia/infothela.htm`.
[10] International Telecommunications Union - Statistics. `http://www.itu.int/ITU-D/ict/statistics/`.
[11] Ricochet Networks. `http://www.ricochet.com/`.
[12] Tele-medicine. `http://www.bimarijankari.org/`.
[13] Tele-medicine Web Portal. `http://www.sehatnama.org/`.
[14] The Network Simulator: ns-2. `http://www.isi.edu/nsnam/ns/`.
[15] The TENET Research Group, Indian Institute of Technology, Madras. `http://www.tenet.res.in/`.
[16] WirelessMAN. `http://WirelessMAN.org/`.
[17] TeNs: The Enhancement Network Simulator. `http://www.cse.iitk.ac.in/~bhaskar/tens/`, May 2004.
[18] A. Agarwal and A. Rungta. Antenna Design for WLAN Applications. Technical report, BTech Project Report, Indian Institute of Technology, Kanpur, May 2003.
[19] P. Bhagwat, B. Raman, and D. Sanghi. Turning 802.11 Inside-Out. In *HotNets-II*, Nov 2003.
[20] F. Cali, M. Conti, and E. Gregori. IEEE 802.11 Wireless LAN: Capacity Analysis and Protocol Enhancement. In *INFOCOM*, Mar/Apr 1998.
[21] Z. Fu, P. Zerfos, H. Luo, S. Lu, L. Zhang, and M. Gerla. The Impact of Multihop Wireless Channel on TCP Throughput and Loss. In *IEEE INFOCOM*, 2003.
[22] R. Gandhi. Empirical Path Loss Models for 802.11b Links. Master's thesis, Indian Institute of Technology, Kanpur, 2003.
[23] S. Garigala. Experimental Validation of Simultaneous Operation in an 802.11 Multi-hop Mesh Network. Master's thesis, Indian Institute of Technology, Kanpur, July 2004.
[24] P. Ipe. Power Allocation Issues in a Wireless Mesh Network. Technical report, BTech Project Report, Indian Institute of Technology, Kanpur, May 2004.
[25] A. Jain and Y. Mishra. Designing a Yagi-Uda Antenna for Wireless LAN. Technical report, BTech Project Report, Indian Institute of Technology, Kanpur, May 2003.
[26] A. Khandelwal. Performance of 802.11b in a Grid Network. Technical report, BTech Project Report, Indian Institute of Technology, Kanpur, April 2003.
[27] A. Kumar and S. Roy. Realistic Support For IEEE802.11b MAC in NS. Technical report, BTech Project Report, Indian Institute of Technology, Kanpur, May 2004.
[28] Midas Communication Technologies Private Limited, `http://www.midascomm.com/`. *CorDECT Wireless Access System*, Dec 2000.
[29] B. O'Hara and A. Petrick. *802.11 Handbook: A Designer's Comparison*. IEEE Press, 1999.
[30] M. I. U. Rahaman. WLAN Watch: A Step Towards The Study Of 802.11b Wireless LANs. Master's thesis, Indian Institute of Technology, Kanpur, April 2003.
[31] B. Raman and K. Chebrolu. Revisiting Network and Protocol Design for 802.11-based Rural Internetworking. Submitted for publication, draft available on request.
[32] A. Sen and M. L. Huson. A New Model for Scheduling Packet Radio Networks. In *INFOCOM*, 1996.
[33] S. Xua and T. Saadawi. Revealing the problems with 802.11 medium access control protocol in multi-hop wireless ad hoc networks. *Computer Networks*, 38, 2002.

A People in the Digital Gangetic Plains Team

- **Faculty:** Pravin Bhagwat, Dheeraj Sanghi, Sanjay G. Dhande, Bhaskaran Raman, A. R. Harish, M. Sachidananda

- **Project team:** Manoj Kumar, Mohan K. Mishra, Vivek Bhagwat, Rajesh Khare, Anish Bhatia, Akshay Mathur, Rajat Kumar Mishra, Rajneesh Tiwari, Prateek Saxena, A. K. Singh, Ram Chandra Prajapati, Yogesh Singh

- **Master's students:** Rajesh Gandhi, Imtiaz Ur Rahaman, Sreekanth Garigala
- **Bachelor's students:** Ankur Khandelwal, Sabyasachi Roy, Ashwini Kumar, Paul Ipe, Anuj Agarwal, Alok Rungta, Ankit Jain, Yagyadatta Mishra

B List of Publications and Theses

See the following citations under "references" above:

- **Publications:** [19]
- **Master's theses:** [22, 30, 23]
- **Bachelor's theses:** [26, 27, 24, 18, 25]

C Availability of Code and Other Resources

`http://www.iitk.ac.in/mladgp/`
`http://www.cse.iitk.ac.in/users/braman/dgp.html`