Turning 802.11 Inside-Out*

Pravin Bhagwat† Bhaskaran Raman‡ Dheeraj Sanghi+

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 metro 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 are exploring the use of 802.11 as a long-distance access technology. 802.11 is currently cost-priced due to competitive mass production 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 multi-hop directional 802.11 links, the testbed spanning up to 80km at its longest. 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. In this paper, we describe the novel technical challenges that lie ahead in using 802.11 to bridge the digital divide.

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 reveals that the growth rate has been phenomenal especially in the developing world. However, in terms of absolute numbers the growth is starkly behind that of the developed world. The difference is even more stark if we examine the *density* of access to communication technology (Tab. 1). Parallel with the growth, the *digital divide* has grown too.

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 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%: www.censusindia.net). Hence the low tele-density as given in Tab. 1.

Country/ Region	Telephone (users/1000)	Cell Phone (users/1000)	Internet (users/1000)
U.S.A.	664.5	450.8	501.5
Japan	596.9	587.8	439.4
Taiwan	573.4	968.8	349.0
Germany	634.8	682.3	373.6
Brazil	217.8	167.3	46.6
China	138.1	110.3	25.7
Africa	26.2	32.2	8.5
India	33.8	6.3	6.8

Table 1. Density of comm. technologies (2001) (source [4])

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. This translates to a tremendous cost for a large-scale deployment (U.S.$80 billion for a target deployment of 400 million lines – 40% of India's current population). 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 rural population density is much lower than in 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. In rural settings, the density of population as well as its paying capacity is small and hence a deployment cannot pay for itself. The same is true of data-based wireless access networks such as 802.16 WirelessMAN (wirelessman.org) or Ricochet (www.ricochet.com).

In this context, consider the 802.11 family of wireless technologies [3]. Since its inception in 1994, 802.11 WiFi has shown tremendous growth and acceptance (in US/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

*This work was supported in part by Media Labs Asia.
†: pravin@acm.org Adjunct Faculty, Dept. of CS&E, IIT Kanpur, INDIA; Adjunct Faculty, Winlab, Rutgers University, USA. ‡: braman@iitk.ac.in Corresponding author, Asst. Prof., Dept. of CS&E, IIT Kanpur, INDIA. +: dheeraj@iitk.ac.in Assoc. Prof., Dept. of CS&E, IIT Kanpur, INDIA.

mass production, the equipment and chip sets are *cost priced* and are hence inexpensive. (Approx. costs: $25-30 (chip-set), $120-700 (Access-Point), $60-110 (PCMCIA card)).

In the Digital Gangetic Plains project (DGP: http://www.iitk.ac.in/mladgp/), 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 outdoor long-distance rural networking – "turning" 802.11 "inside-out". Our vision for such use of 802.11 is as in our testbed (Fig. 1).

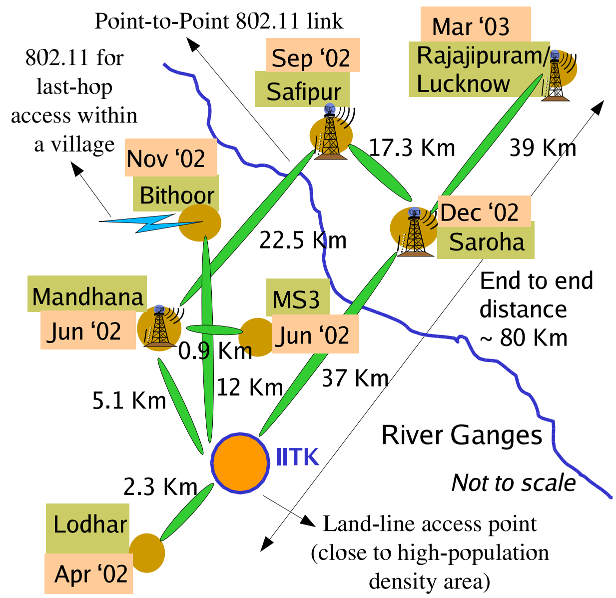

Figure 1. The Digital Gangetic Plains testbed

Although several 802.11 community networks [1] are popping up, our intended use is quite different. We envision several tens of kilometers of (relatively) sparsely populated rural areas being covered by long-distance point-to-point 802.11 links (using directional antennae), constituting a multi-hop wireless access network (Fig. 1). To our knowledge, there has thus far been no systematic effort to understand the various operational and research issues in this setting. The point-to-point links are a cost-effective alternative to wired links. Apart from the long-distance mesh network, we also envisage the use of 802.11 for the last-hop access within rural villages (Fig. 1). 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. We have been able to bring low cost Internet and VoIP to villages that do not have a telephone to this day.

Our outdoor long-distance use of 802.11 compels a revisit to the protocols at various layers of the OSI stack, as well as various system design issues. In this paper, we describe our testbed (Sec. 2), and the various technical issues that lie ahead in making 802.11 a cost and performance effective solution for rural networking (Sec. 3).

2 The Digital Gangetic Plains Testbed

Our testbed is set in a rural area. 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). This also allows us to network several villages and run applications for real use by people as the testbed evolves. 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) FCC 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) 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 effort was started in Apr'02 and the testbed extent has grown steadily since then. We currently have eight nodes and eight links, with the longest point-to-point link going as far as 38km. The testbed is depicted in Fig. 1. We recently added the 38km link – with this, the maximum (air) distance between two points in the network is about 80km. Over the last one year period 28 bachelor and master's students have carried out experiments on this network. All locations except 'IITK' and 'Rajajipuram/Lucknow' 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'.

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 multipath. The heights of all towers are about 40m, except the one at 'MS3' where the height is about 15m. We use off-the-shelf parabolic grid antennae for directional gain and off-the-shelf 802.11b Access-Points (APs) at all the locations. The APs are currently setup in bridge-mode. In the future, we also plan to explore network-layer routing in the wireless access network.

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.

3 Technical Challenges

Our use of 802.11 for rural networking is primarily motivated by cost considerations. The technical issues raised in this section fall under two categories: (a) issues due to outdoor, long-distance use of 802.11 (something it was not designed for), and (b) more generic issues in the design of protocols/systems for rural use. We start with a discussion of the physical layer issues and move up the OSI stack subsequently.

3.1 Physical layer issues

Empirical path loss models

The first challenge we faced was to figure out how far 802.11 transmission would go. 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.

Assuming free space propagation model, our theoretical path loss calculations indicated that with high gain antennas on both transmit and receive side it was possible to set up 11Mbps links up to distances of 40km. However, we were not sure whether free space path loss model would be applicable in flat riverbed, agricultural land with scattered tall trees. The effect of fog, rain, and temperature oscillations was another unknown which we did not know how to account for in the link budget analysis. Antenna mast oscillations due to high wind speed is another reason because of which signal strength fluctuations can happen at the receiver.

The longest link in our network so far is 38km long. It has been built using a pair of 100mW transmitters connected to 23dBi parabolic grid antennae placed on top of the 40m high towers on both ends. There is approximately 6dB difference between the observed signal strength and theoretically predicted value using free space path loss model. This 4-6dB difference is consistent across all seven links we had built as of Feb 2003 [6]. According to our experiments so far, free space path loss model with 4-6dB correction is a good predictor for path loss over long distance point-to-point links.

Our immediate next challenge is to estimate how much area can be lit up by a single transmitter, for last-hop access. In each village site, we have a high gain omni directional antenna mounted slightly above the ground. Measurements at each site have revealed different coverage patterns – path loss exponent typically varies between 2.2 to 3 [6]. Our aim now is to build a suite of empirical path loss models for a variety of scenarios so that coverage area of any 802.11 transmitter can be predicted. Through use of such accurate coverage prediction models, if a given area can be lit up with fewer transmitters, overall system cost can be reduced.

PHY performance under outdoor channel conditions

Few months ago we were attempting to build a link between 'IITK' and 'MS3' (5km link; 'MS3' has 15m tower). We came across a scenario in which signal strength was high, yet the packet-error rate was high. Currently the link does not work at all. We suspect that the reason for these is the following. 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.

The above problem does not surface in the other long-distance links since they are free of multipath. Mounting antennae 40m above the ground surface ensures that there is clear line-of-site and no obstruction in Freznel zone between the two ends.

Link performance under a variety of channel conditions – analytical as well as experimental studies of QPSK, CCK, and OFDM modulations, in 2.4GHz as well as 5Ghz bands are topics for further study.

Power efficiency: a new perspective

Power efficiency at the end client has received a great deal of attention thus far. In our setting, power efficiency of APs and routers also become issues. One of the operational challenges we face in our testbed is the presence of power to run the equipment at the various locations. Remote locations may not have readily accessible power links to tap from, and even otherwise, power is rarely reliable. It is not uncommon to have power outages in remote villages for several days at a stretch.

For these reasons, 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.

However, in our current setting, the APs consume about 30W each. Thus even a node with a single link (single AP) cannot be operated using a single solar panel. In future, we also plan to use routers at each node. This would present further problems. Hence in our setting, there is enormous value to optimizing the power consumption of APs and routers. This is a topic for further exploration.

Spectral efficiency Vs. Cost Tradeoff
Taking a perspective broader than just the use of 802.11, consider the following. Improving spectral efficiency has been a key driver for the wireless communication research for the past 10-20 years. This focus has resulted in more and more complex channel coding and modulation methods which aim at improving spectral efficiency at the cost of throwing more signal processing power at the problem. This makes sense in spectrally crowded regimes such as US and Europe where spectrum is considered a precious resource and its use must be optimized at any cost.

In the developing world spectrum resource is not as precious. We argue that the focal point of optimization therefore must be shifted from spectral efficiency improvement to cost reduction. WiFi despite its poor spectral efficiency is more attractive than other available options because it can provide connectivity at low cost. Our goal in DGP is to work with this technology while still achieving good performance.

Summary: *In summary, PHY layer issues need a revisit in the context of outdoor use of 802.11. More generally, future design of communication technologies for rural use should give primary consideration to cost and system power consumption.*

3.2 MAC issues

MAC issues for long-distance links
802.11 MAC 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)[1]. 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.

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[2]. A time-division based scheduling approach (TDMA) is more appropriate. That is, scheduling the various links of the network for transmission, taking into account what simultaneous transmissions, or *spatial-reuse* is possible. This is known by the generic term Spatial-reuse TDMA (STDMA) scheduling (e.g. see [7]).

Apart from the time dimension (scheduling), we also have the frequency dimension: 802.11b defines at least eleven channels (1-11), with overlapping frequency ranges. Fig. 2 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.

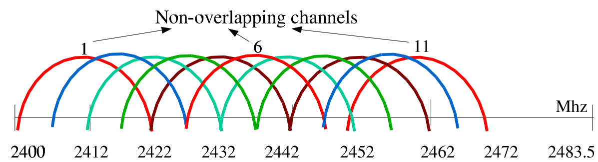

Figure 2. Overlapping 802.11 channels: schematic

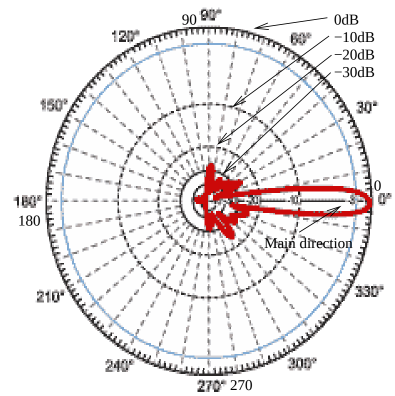

Figure 3. Spatial radn. pattern of parabolic grid antennae

To understand the nature of channel reuse possible at a single node, across its multiple links, we performed two experiments. First, we used the two links 'IITK-Mandhana' and 'Mandhana-Safipur' of our testbed, with only these two links turned on, and the power level of all transmitters fixed at 100mW. One of the links ('IITK-Mandhana') was fixed at channel 11, while the channel for the other link was varied from 6 through 11. Packet-level traces revealed that there was little or no interference when using channel pair $(11, 6)$ or $(11, 7)$. However, with greater channel overlap (e.g., $(11, 8)$, $(11, 9)$), there was significant interference. Although we have directional antennae, interference still happens because of interference power leaking in through "side-lobes" away from

[1] Commercial APs allow to set a "distance" parameter which they use to set the timeouts appropriately; in our testbed, if this distance is underestimated in configuration by more than 5km, the AP fails to form the link.

[2] The commercial APs in our testbed use a proprietary (non-inter-operable) polling-based scheme for sharing the bandwidth of the half-duplex point-to-point link. This however does not handle multiple links.

the main direction. The "side-lobe" patterns for the directional antennae used in our testbed [2] are shown in Fig. 3.

The above experiment was performed with fixed transmit power levels. Since interference is a function of power levels, it may be possible to reduce it significantly through careful power budgeting. So, we performed another experiment to determine the Signal-to-Interference power ratio required *at the common receiving node* for interference-free operation, when using overlapping channels. For this, we required fine-grained power control (e.g. using attenuators). Since we do not yet have this setup in our testbed, we examined this in an indoor lab. We had a "main link" between an AP and a client laptop, and an "interfering link" between another AP and a second laptop, with both links operating on the same channel (i.e., maximum channel overlap, e.g., $(11, 11)$). To achieve tight control over power levels, we used RF cables for the links, instead of a wireless channel. We fixed the power level of the main AP, and added attenuators such that the main signal power, at reception, was $-40dBm$. Some of the power from the interfering AP was fed (through connectors) as interference to the main client laptop. We varied the interference power level through attenuators, and found that the main link could operate, using the *same channel* as the interferer, if the power level *at reception* of the interferer was 10dB below the signal power *at reception*. (We observed significant packet errors on the main link at higher interference levels).

Now, this 10dB isolation can be achieved for two links at a node in our mesh network as follows. Since the setting is static, we can adjust the power levels at the transmitters such that the power *at reception* is the same along the *main direction* for both the links under consideration. Since the side-lobes in Fig. 3 are at least 25dB or more below the power level in the main direction (beyond an angular separation of $25 - 30^o$), such isolation is thus possible – reception on one link at a node will not interfere with reception on another.

The essence of all this is that the following nature of spatial/channel reuse is possible. *Simultaneous reception from different directions, on the same channel, is possible at a node* (with $25 - 30^o$ separation between links). Similarly, simultaneous independent transmission along multiple directions is also possible[3]. Channel allocation for the various links, and STDMA scheduling need to be explored under such a spatial/channel reuse flexibility. Although STDMA scheduling in packet radio networks is as old as two decades (see [7] and references therein), such a simultaneous reception (or transmission) flexibility has not been considered so far.

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 (Fig. 1). This raises several issues. First, the system can have high contention, something that 802.11 MAC was not optimized for (e.g. see [5]). 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.

Summary: *MAC issues in our setting need a revisit, under the constraint that we are to work within the scope of the 802.11 PHY, to retain its cost advantages*[4]. *We need to address issues of channel allocation and STDMA scheduling for the long-distance links. These need to be explored in light of the simultaneous reception flexibility in our network.*

3.3 Routing and reconfigurability

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*, as we explain below.

In the context of our discussion in the previous subsection, consider just the long-distance links. Suppose the power levels are such that a maximum capacity of 11Mbps is possible. This represents the sum of the raw link bandwidths along either directions of a link. We now observe the following. (1) The 11Mbps can be (dynamically) split in any fashion between the two directions. (2) More generally, define a *channel subgraph* to be a maximally connected subgraph of the original network where all the links are allocated the same channel. The links in the network are thus partitioned into various channel subgraphs. Such partitioning depends on the current channel allocation. Further, the scheduling of various links are *inter-dependent* within a channel subgraph (they are independent across two channel subgraphs with non-overlapping channels – say $(6, 11)$). Thus the capacities achieved in either direction on the various links in a channel subgraph are related to one another in any STDMA scheduling scheme.

Thus, *the link capacities can be varied dynamically, at will* by varying the scheduling and/or the channel allocation. Another way in which the link capacities can be varied is by changing power levels of transmission. A lower power level

[3]Note however that we cannot simultaneously transmit along one link, and receive along another, at a node. This is because the transmission power will be quite high near the transmitter.

[4]Firmware level changes to MAC/LLC are possible with minimal addition to system cost. We are exploring collaboration with the industry to get access to such firmware.

can potentially achieve longer range at the cost of lower capacity (802.11b defines 1, 2, 5.5, and 11Mbps transmissions).

Another dimension of reconfigurability is the network topology itself. The network may be constructed with a number of point-to-point links. Not all links may be turned on simultaneously for several reasons. (1) There may not be enough channels available for meaningful allocation to all links (in other words, channel allocation and STDMA scheduling may be easier/more effective if some links are simply turned off). (2) There may not be enough power at a node to turn on all APs and operate all links. Thus *the network topology itself can be varied dynamically, at will*.

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. While these 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).

Summary: Fig. 4 summarizes the set of issues discussed above in the context of *reconfigurability* and network planning. While each of the issues may have been addressed individually in different settings (e.g., traffic engineering in wired networks, channel allocation in cellular networks), the interrelated sets of issues is unique to our setting and represents both an opportunity as well as a challenge.

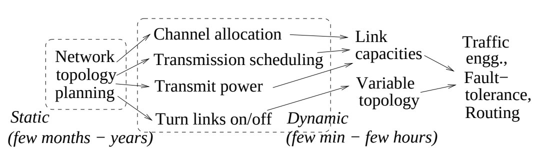

Figure 4. Reconfigurability: opportunities and challenges

3.4 Other 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 adhoc 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 (Sec. 3.2). We plan to use the testbed to understand the performance of TCP under such conditions.

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.

4 Conclusions

Despite a decade and a half of communication revolution, much of the world is yet to see its benefits. Most technologies are designed for developed telecom economies and equipments value-priced accordingly. We have taken a currently cost-priced technology (802.11b) designed for indoor use and "turned it inside out" for covering low-density rural areas. We have identified a plethora of technical issues that require a revisit due to various reasons. Cost reduction is the primary motivation in our choice of 802.11, and many of the performance-related issues arise due to our outdoor use of 802.11. In line with this logic, we also argue that future communication (or other) technologies for rural use should primarily consider system cost and system power consumption as points of optimization. Currently, our efforts are on expanding the testbed, addressing the various operational issues, and using it further to explore the different research issues.

Acknowledgments: We thank the entire DGP team for the various testbed related efforts. We also thank Kameswari Chebrolu for comments on earlier versions of this paper.

References

[1] Community Wireless / Rooftop Systems. www.practicallynetworked.com/tools/wireless_articles_community.htm.

[2] HyperGain HG2424G 2.4 GHz 24 dBi High Performance Reflector Grid Antenna. http://www.hyperlinktech.com/web/hg2424g.php.

[3] IEEE P802.11, The Working Group for Wireless LANs. http://grouper.ieee.org/groups/802/11/.

[4] International Telecommunications Union - Statistics. http://www.itu.int/ITU-D/ict/statistics/.

[5] F. Cali, M. Conti, and E. Gregori. IEEE 802.11 Wireless LAN: Capacity Analysis and Protocol Enhancement. In *INFOCOM*, Mar/Apr 1998.

[6] R. Gandhi. Empirical Path Loss Models for 802.11b Links. Master's thesis, Indian Institute of Technology, Kanpur, 2003.

[7] A. Sen and M. L. Huson. A New Model for Scheduling Packet Radio Networks. In *INFOCOM*, 1996.