

Experiences in Using WiFi for Rural Internet in India

Bhaskaran Raman and Kameswari Chebrolu, Indian Institute of Technology, Kanpur

ABSTRACT

Access to communication can play a pivotal role in the socio-economic development of rural regions in the third world. For affordability, the choice of technology to achieve this is a significant aspect. We have chosen IEEE 802.11 as a cost-effective technology to provide rural connectivity in the context of two projects in India: Digital Gangetic Plains (DGP) and Ashwini. This article presents our experiences with these two projects and discusses five important aspects in the use of WiFi for rural connectivity: network planning and deployment, network protocols, network management and operations, power savings, and applications and services.

INTRODUCTION

Research and technology are most needed where development is lacking, that is, in the developing regions of the world. However, technology meant for developed telecom economies is often too high cost, or otherwise unsuitable, for use in rural regions of the third world. Thus, despite the enormous benefits enjoyed through the cellular revolution by people in metropolitan areas of developing countries like India, penetration of the technology in rural areas is poor or nonexistent. Two reasons for this are the low population density in rural areas (in comparison with cities), and the low paying capacity of users there. The cost of technology is an important factor in deciding its applicability for rural use.

In this article we present two projects in India in which we have used IEEE 802.11 [1] (WiFi) as a cost-effective technology to provide wireless access to rural areas:

- Digital Gangetic Plains (DGP)
- Ashwini

WiFi is not only cost-effective in its equipment, but also because of its operation in license-free spectrum. Further, WiFi offers the possibility of broadband speeds. Based on experience in the two projects, we describe various design considerations, lessons learned, technical issues we have addressed, and open issues in the use of WiFi for rural access.

The DGP project¹ was initiated in 2002 at the

Indian Institute of Technology, Kanpur (IITK), Uttar Pradesh, to explore the technical feasibility of establishing long-distance 802.11 links. Figure 1 shows the DGP network as it has evolved over time. We have a central location called the landline node, which has wired Internet connectivity. In the DGP testbed, this is the IITK location. We have many long-distance (up to several tens of kms) 802.11 links formed using high-gain directional antennas. A long-distance link is depicted in Fig. 2. Such links connect various surrounding villages to the landline through a multi-hop mesh network. The DGP network has primarily been used as a testbed for aiding various protocol studies, although we do have people using communication services at many locations.

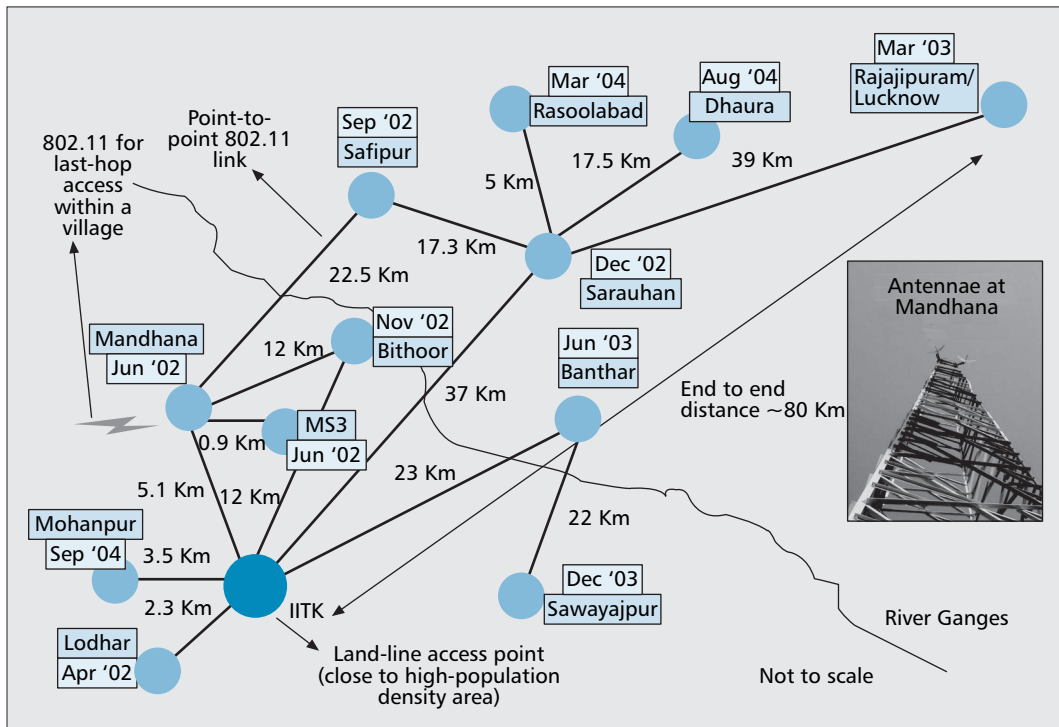
The Ashwini project² is a network deployment effort by the Byrraju Foundation to provide broadband access and services to a collection of villages in the West Godavari district of Andhra Pradesh, India. This is a service-centric deployment and currently a variety of interactive video-based applications (e.g., distance-education, tele-medicine, etc.) are being run on the network. Figure 3 shows the currently active links in the Ashwini network. Here, the Bhimavaram location is the *landline* node. It is worth noting that a further 17 nodes will be included in the network in the near future.

While we have used WiFi as the technology of choice in the DGP and Ashwini projects, two other technologies which are worthy of comparison are CorDECT [2] and WiMAX [3]. CorDECT is also a low-cost technology designed for rural access (TENET group, IIT-Madras). However, the bandwidths achievable with CorDECT are much lower (about 70 kb/s). While WiMAX promises broadband speeds, it is not clear if it will achieve the same cost scale as WiFi. WiFi operates in the ISM (Industrial, Scientific, Medical) band, which is license-free in many parts of the world. It is already well established, has captured the U.S. and European markets, and has achieved low-cost mass production. In contrast, WiMAX is still in trials and lacks worldwide acceptance.

In the sections to follow we describe five different aspects relevant to the deployment of long-distance WiFi networks:

¹ Supported by Media Lab Asia, India:
<http://www.cse.iitk.ac.in/users/braman/dgp.html>

² Supported by Media Lab Asia, India and the National Institute for Self Governance (NISG), India: <http://www.byrraju.foundation.org/ashwini/overview.htm>



■ **Figure 1.** *The Digital Gangetic Plains testbed.*

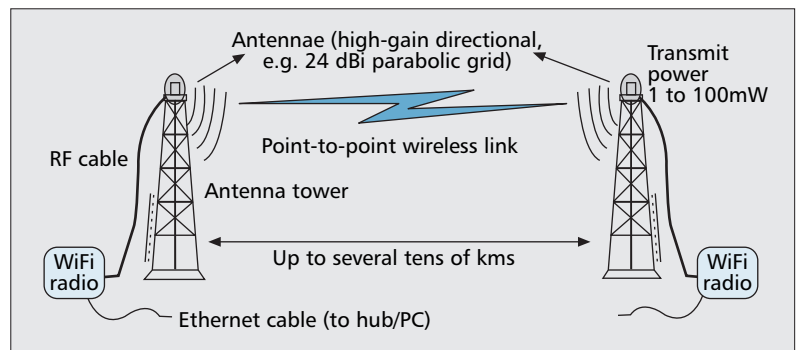
- Network planning
- MAC protocols
- Network management
- Power saving
- Applications and services

These issues have a unique flavor due to the use of the technology in rural areas. The uniqueness is due to three reasons: the primal need for cost optimization, considerations of power consumption, and the need for targeted applications for use in rural communities.

NETWORK PLANNING

For any rural system, cost is the primary consideration in every step of the design. Next to the choice of technology (WiFi), the context in which cost optimization is significant is in *network planning*. We elaborate on this issue below. A long-distance WiFi link requires line-of-sight to get sufficient signal strength for reception; otherwise, the attenuation at 2.4 or 5 GHz is too high beyond a few hundred meters. This in turn implies that network deployment involves significant infrastructure in terms of antenna towers. The approximate cost of antenna towers/masts is given in Table 1. In comparison, note that WiFi radio cost can be about US\$50 or less. Tall towers are one or two orders of magnitude costlier than the radio equipment!

Hence in long-distance WiFi deployments, network planning to optimize the infrastructure cost will play a significant role. This is quite unlike WiFi-based *community networks* in metropolitan areas, where an ad hoc deployment model is workable. Apart from cost reduction, a secondary concern requiring network planning in our setting is to guarantee a minimum level of performance to each village node. A simple way



■ **Figure 2.** *WiFi long-distance link: an illustration.*

to quantify this performance is in terms of the achievable network throughput to each village while operating at capacity. For instance, in the Ashwini project, we have a requirement of providing at least 384 kb/s to each village to enable high-quality video-conferencing.

Network planning considering cost and performance involves at least six related aspects:

- The tower locations, representing the nodes of the network
- Tower heights at each location
- What links to form, to have a connected network topology
- What antenna types to use for each link
- The transmit powers to use for each radio
- The channel of operation for each link

Typically, the broad *locations of the towers* are chosen to be the villages where network connectivity is desired. Further, the finer choice of where in each village to locate the towers is also typically done based on nontechnical reasons of convenience of operation. This has been the case both in DGP as well as Ashwini. In general,

Transmit power allocation is a non-obvious aspect of network planning, especially since it can be changed dynamically even after deployment. However, unless planned, it can result in interlink interference. This in turn can cause cases of hidden nodes, leading to packet errors, protocol inefficiency, and application performance degradation.

Tower/mast height (m)	10	15	21	24	27	30	45
Cost (\$)	100.00	150.00	800.00	950.00	1100.00	1850.00	5000.00

■ **Table 1.** Typical antenna tower/mast costs.

however, a commercial deployment may choose its towers to be located at places where the height above mean sea level (MSL) is higher in comparison to the others. This is to help reduce the tower height requirement.

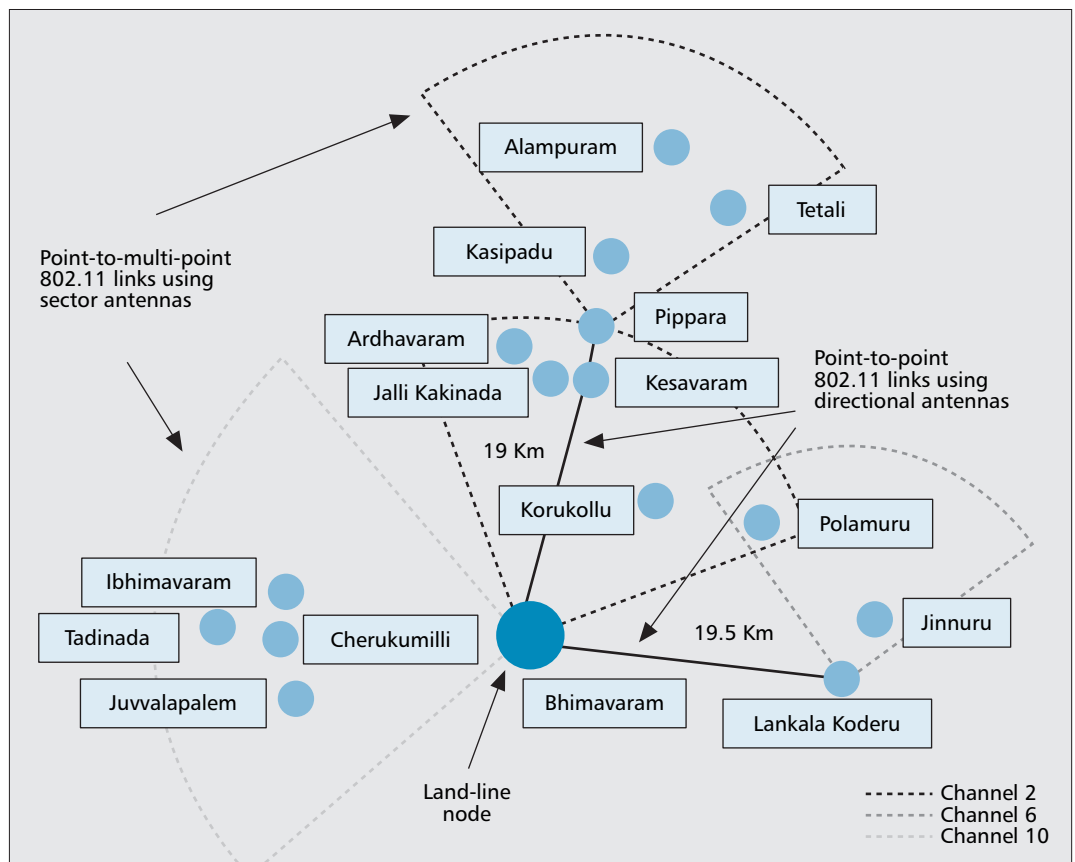
The next four aspects of tower heights, network topology, antenna types, and transmit powers are closely interrelated; the choice of each of these affects the other directly or indirectly. In [4] we have explored this interdependence in depth. Figure 4 summarizes the dependencies and our approach in addressing the above issues. We solve four interrelated problems which we term topology search, height assignment, antenna assignment, and power assignment, respectively. We have modeled height and power assignment problems as linear programming (LP) problems. And we have heuristic algorithms for the topology search and antenna assignment problems.

With respect to the tower height assignment, an insight from the CorDECT deployment efforts is worth mentioning here. In forming a long-distance link, instead of having towers at each end, we could have a tall tower at one end. At the other, we could just use a mast consisting

of a water-pipe of 10 to 15 m height. The cost of a mast is much lower (see the entries for 10 and 15 m in Table 1). Such an arrangement would be advantageous if we amortize the cost of one tall tower among many other masts.

Transmit power allocation is a nonobvious aspect of network planning, especially since it can be changed dynamically even after deployment. However, unless planned, it can result in interlink interference. This in turn can cause cases of hidden nodes,³ leading to packet errors, protocol inefficiency, and application performance degradation. In the Ashwini project as well as the DGP network, we have observed cases of interlink interference. Although the problem is surmountable as of now, due to the relatively small scale of traffic, the issue is bound to require careful consideration in the near future.

It is worth noting that the DGP and Ashwini networks themselves have not used any sophisticated network planning. The topology was ad hoc for DGP, while it involved minimal manual consideration of achievable throughput in the case of Ashwini. In both cases, the heights were chosen with only a local view, individually for



■ **Figure 3.** The Ashwini network topology.

³ Hidden nodes: two wireless senders out of range of (hidden from each other) sending to a common receiver.

each link to achieve line of sight. The antenna types and transmit power used were based on intuition.

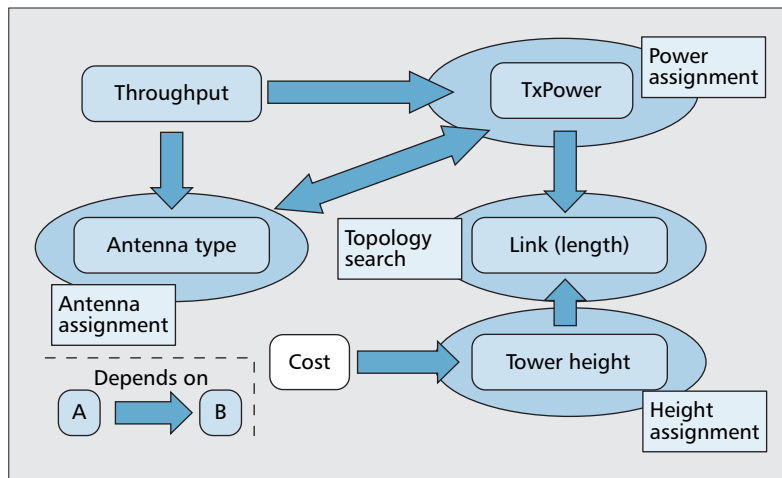
The last of the network planning issues listed above, that of *channel allocation*, is also an important one. It spans technical, business, and political domains. While 802.11a has up to 12 nonoverlapping channels, 802.11b/g has only three. Unfortunately, only the 802.11b/g 2.4 GHz band is available for license-free outdoor operation in India. Channel allocation becomes more significant if local 802.11 operation also has to be accommodated. For instance, apart from the backbone networks shown in Figs. 1 and 3, we may desire to use 802.11 for local coverage within a village.

The issue becomes even thornier if we consider the fact that there could be interference across independent deployments. While the possibility of this happening is less in rural environments, it could be common at least in the parts of the network where larger towns or cities are involved. There are already several outdoor 802.11 links being used in metropolitan areas. We have observed this significant trend in our DGP network. Whereas there were hardly any other 802.11 deployments at the inception of the project, we have now found several such deployments interfering with our links at various locations in the network.

The above discussion suggests that while rolling out a large scale network deployment, say to provide connection to the villages in a district, we may need legal or semi-legal mechanisms to protect the deployment from others which may come up in the future. Of course, such protection becomes easier if the deployment were to use as few channels as possible. We have explored mechanisms in which the backbone network uses only one channel, with the other channels being free for other deployments, or for local 802.11 coverage at each village. The protocol support for such operation leads us to the next topic of discussion.

MAC PROTOCOLS

IEEE 802.11 specifies a MAC protocol based on carrier sense multiple access with collision avoidance (CSMA/CA) contention resolution. Although such an approach is appropriate for indoor operation, it faces many inefficiencies during outdoor long-distance operation. First is the inefficiency due to the timing parameter settings such as ACK timeout, contention slot-time, interframe spacing, and so forth. This is because the round-trip-time is much higher in long-distance settings: 100 μ s for a 15 km link as compared to a few microseconds in indoor settings. More importantly, a contention-based approach is fundamentally not suited for a situation where we do not have arbitrary contention. There are inherent timing uncertainties involved with arbitrary contention resolution. The combination of the above: timing uncertainties and the larger timescales involved would mean that providing quality of service (QoS) guarantees to voice/video applications would be complicated in long-distance settings, with the use of CSMA/CA.



■ Figure 4. Dependencies in network planning.

In seeking an appropriate protocol for long-distance WiFi networks, it is useful to make a distinction between point-to-point (P2P) links and point-to-multipoint (P2MP) links. A P2P link typically uses high gain directional antennas (e.g., 24 dBi) and consists of one radio at each end. A P2MP link involves a single radio at one end connected to several radios at different other points. It typically uses a sector antenna at the central end. Such P2MP links are shown in Fig. 3. For neither of these types of links is CSMA/CA appropriate due to the reasons mentioned above.

For a network (or a part of a network) with P2P links, we have explored the possibility of having all the links operate simultaneously in the same channel [5] (recall the need for minimizing channel usage from the previous section). The key idea here involves loosely synchronizing the operation of the multiple links at a node: they will all transmit at the same time (SynTx), or will all receive at the same time (SynRx). Each node will switch between these SynTx and SynRx phases, and hence we have named this protocol 2P [5]. Such operation would avoid the arbitrary contention which would result in using CSMA/CA.

The performance improvement of 2P over CSMA/CA can be significant. For a simple network with two adjacent links, simultaneous operation of the two links means a twofold improvement in performance with 2P as compared to CSMA/CA. For larger networks, we have shown through simulation in [5] that the performance increase can be as high as a factor of about 20.

We have prototyped the 2P protocol using driver-level modifications to the HostAP driver as well as using hardware abstraction layer (HAL) modifications to the MADWIFI driver [6]. However, the inability to completely turn off the carrier-sense feature has impeded our efforts thus far. It would be useful if such an interface were made available in a WiFi chip platform.

For P2MP links too, we have developed MAC protocol modifications for operation in long-distance settings. Our protocol, called SRAWAN (Sectorized Rural Access Wide Area Network) [6] is time-division multiple access (TDMA) in

In some respects, the diagnosis of any poor performance should be easier in long-distance WiFi networks. This is because of the predictability of wireless performance in such line-of-sight settings. For instance, we can tell what the received signal strength should be for a given wireless link.

style. SRAWAN is in essence a simplified version of the WiMAX MAC protocol, adapted for operation on top of 802.11 PHY. The TDMA-based approach eliminates the contention and hidden-node issues which would arise in a CSMA/CA based approach. It allows us to provide QoS in P2MP settings for video/voice applications. It also improves the performance for other applications (e.g., file download, Web browsing). In our experimentation with our prototype implementation thus far, we have achieved satisfactory performance for medium-sized networks (up to 20 client connections), with effective MAC overheads similar to or better than normal 802.11 indoor operation [7]. We plan further testing at scale in the near future.

NETWORK MANAGEMENT

Network management consists of performance and fault diagnosis and repair. It is known that such management is more complicated in wireless networks in comparison with traditional wired networks. There are far more reasons for bad behavior of wireless performance: interference from other WiFi sources, from other non-WiFi sources, signal strength variation, and so on.

While work is in progress to address network management in enterprise WiFi networks (e.g., 802.11v, 802.11k), this is not the case for long-distance settings. The remote diagnosis and repair of these problems is especially important to address in long-distance WiFi networks for several reasons:

- Any physical visit required would involve significant cost, since the distances involved are large.
- Rural locations are relatively inaccessible.
- The availability of trained personnel in rural areas is relatively poor.

To present an extreme case, visiting the Saruhan site of the DGP testbed to diagnose any problem involves close to an entire working day of travel (six hours) of one or two trained personnel. Sadly, one of the prime reasons requiring our visit there has been a simple problem: the WiFi bridge resets its settings to factory defaults if it is rebooted due to a power outage. Clearly, the bridge management software was not designed for operation in rural regions where power outages are the norm rather than the exception.

In some respects, the diagnosis of any poor performance should be easier in long-distance WiFi networks. This is because of the predictability of wireless performance in such line-of-sight settings [8]. For instance, we can tell what the received signal strength should be for a given wireless link. This would also be indicative of the expected application performance. We can then quickly report any deviation from this.

Support required for network management should be in terms of:

- Collecting the wireless configuration and network topology information centrally
- Measuring the signal strength and noise level in each link
- Detecting any sources of interference in the same channel as well as adjacent channels

- Passively and/or actively measuring the wireless error rate as well as application-level throughput

It should be possible to take such measurements for links in the network for any given transmit power, transmit rate, and channel of operation. There should also be support for performance testing separate links of the network simultaneously.

All the above information should be made available centrally for online/offline analysis. We are in the process of putting together a network management tool with such capabilities. There are open research issues in terms of how the information collected centrally may be used to diagnose and correct any performance problems.

Looking at to the issue of fault tolerance, this issue should be addressed in the network topology design itself (discussed earlier) by provisioning possible backup transit nodes. Dynamic support is also required in terms of detecting the fault and switching over to a backup link.

An easy way to do dynamic switchover from a current intermediate node to a backup node is of course at the routing layer. However, this may not be the right thing to do, since it would require that the underlying physical link be already present. That is, any transmissions from the village node to the intermediate node would also always reach the backup node. This would end up interfering with the normal operation of the designated backup node, whereas ideally, we want the backup node to be affected only when required (i.e., during failure of the original intermediate hop).

In the DGP project, we have explored the use of RF switches to address this issue: the physical transmissions are directed to the backup node only when a failure has been detected on the main intermediate node. The use of such RF switches will not add much to the cost of the system: we have prototyped with an RF switch as inexpensive as US\$3 (RSW-2-25-P).

POWER SAVINGS

Next to cost reduction, another important concern in systems for developing countries, especially rural areas, is the presence (rather, absence) of power to run the equipment at the various locations. In these places, power is rarely reliable and it is not uncommon to have power outages for several days at a stretch. To have reliable operation of services in these deployments, current practices involve the use of heavy capacity batteries to power the equipment. The batteries are recharged whenever power is available from the grid or through the use of solar panels. For example, in Ashwini, about US\$2000 is invested per village site in batteries which are charged off the grid when power is available. The battery is used not only to support networking equipment, but also the local PC and other peripheral hardware. In DGP, the cost of a solar panel just to power the networking equipment was about US\$1100. Apart from the cost, since the solar panels are deployed outdoors, there is the risk of theft: we have lost one of the solar panels at the Saruhan location of DGP.

To minimize the investment cost in

batteries/solar panels, there is the utmost need for power conservation in these settings. While there is as much need to conserve power for all of the equipment in use, we elaborate further on conserving power of the networking equipment, since this is the main focus of this article.

A popular platform for several WiFi mesh networks is a single board computer platform Soekris (www.soekris.com) that can use off-the-shelf WiFi cards. Our measurements indicate that the power consumption of Soekris can be as high as 5 to 8 W [9], only a small fraction of which is consumed by the WiFi interface itself.

A suitable approach to conserve power in these settings is to turn off the equipment (Soekris) when the links are idle and to turn it on only when connectivity is desired. A remote host that needs connectivity would have to turn on equipment at each intermediate hop to the wired gateway. This requires a mechanism for remote power-on.

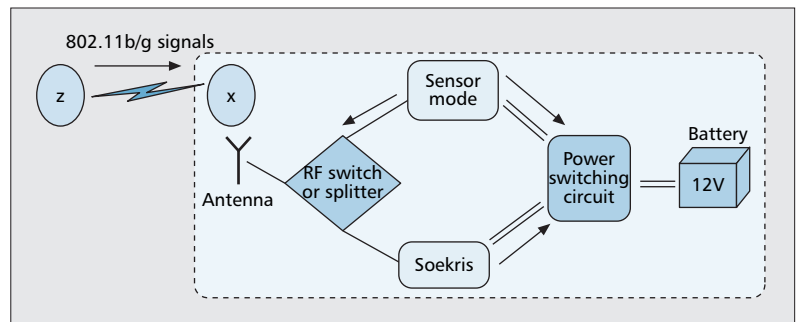
To achieve this, we have developed a mechanism called Wake-on-WLAN [9], which is depicted in Fig. 5, where a node Z wishes to remotely power-on node X. Wake-on-WLAN relies on a novel implementation using off-the-shelf 802.15.4 sensor motes (Moteiv's Tmote sky: www.moteiv.com). At X, we use a mote in conjunction with the WiFi radio setup. Unlike the WiFi setup, the mote itself consumes very little power (100 mW or less) and can potentially be left on all the time. To power-on X remotely, Z first sends continuous 802.11 transmissions. The mote's 802.15.4 radio chip (Chipcon's CC2420), on sensing the energy level of these transmissions, powers-on the Soekris board along with the WiFi interface.

When the WiFi radio at X is powered on, the WiFi link with Z is established based on prior configuration. Thus, the WiFi link is established on-demand. Wake-on-WLAN takes advantage of the fact that both 802.11b/g and 802.15.4 work in the 2.4 GHz band. This also obviates the need for a separate antenna for communication during the power-on procedure. For the same reason, we cannot use such a mechanism when using 802.11a (5 GHz operation).

The additional hardware used in Wake-on-WLAN is inexpensive, too: the CC2420 radio can cost about US\$10. We have prototyped Wake-on-WLAN and tested it successfully outdoors on the 3.5-km-long IITK-Mohanpur link (Fig. 1). Preliminary analysis of power savings, considering link usage pattern observed on one of our testbed links that provides VoIP connectivity to a rural village, have shown that with the use of Wake-on-WLAN we can achieve over a tenfold reduction in power consumption [9].

There are a few aspects of the above mechanism that need further investigation. One is the effect random noise or interference has on the wake-up mechanism. One way to get around this is to look for a specific signature pattern in the 802.11 transmissions. A similar concept can also be applied to extend this concept from P2P links to P2MP links: to wakeup one particular radio of a P2MP link based on a signature pattern.

The possibility of being able to dynamically change the network topology using Wake-on-WLAN also opens up interesting possibilities:



■ Figure 5. Wake-on-WLAN architecture [9].

- Adapting the topology to match the traffic pattern efficiently
- Rebooting nodes remotely to aid in troubleshooting or in protocol debugging

APPLICATIONS AND SERVICES

Just addressing the above issues to provide Internet connection at villages, and expecting conventional applications to be used, would be naive. In our experience (as well as others'), typical "browsing"-based services will not find many buyers in rural settings. This is very much unlike metropolitan settings where Internet browsing centers have become commonplace. For rural areas where most people are still on the wrong side of the digital divide, *targeted application services* are required to gain the full benefit of digital communication access.

One such service which we have experience with is voice over IP (VoIP). We ran this specific service for about two years at the Sarauhan village. This service was especially popular, since the nearest regular phone from Sarauhan was about 6 km away at the time we began the service. We have extensively documented the various socio-economic aspects of this VoIP deployment in [10]. It is worth mentioning here that the service has been a loss-making venture in its current model: the net revenue earned saturated at about US\$9/mo, while the running expenses themselves were about US\$52/mo.

Further, the infrastructure cost was about US\$11,000, with the dominant components being the antenna tower at Sarauhan (US\$6600), the Cisco wireless bridges (US\$2000), and the solar power system (US\$1100). While optimizations are possible in each of these components [10], further trials are required. In any case, it is best to view VoIP as a service that can attract potential customers to other services. This is because voice telephony is a very well-understood service, even among rural communities. There is a range of other services which have been deployed in the Ashwini project. Most of these are based on two-way video communication. Examples of such services include the following:

- *Education enhancement*: Currently several supplementary education classes are conducted on subjects, such as English, mathematics, and sciences.
- *Health training*: Information and survey camps are held through the video sessions.

There are several technical issues in making WiFi work in long-distance rural settings. Despite the radio itself being low-cost, the other parts of the system may involve substantial cost and optimization is required to enhance the viability of WiFi.

- *Agricultural information*: This is also in demand, since the primary occupation of the populace around Bhimavaram is farming.
- *Business process outsourcing (BPO)*: A BPO centre is run at the Jallikakinada village for a Hyderabad-based software company. This service does not use video-based communication, however.

All of the video-based services are conducted by having a subject expert come to the central location (Bhimavaram). The subject expert communicates through video with the audience/trainees at the village end. At each of the villages, there is an Ashwini training centre with wireless connectivity and video-conferencing facilities for this purpose. Several of the video sessions are run one-to-many. That is, from the central Bhimavaram location to two or three other villages simultaneously. At times, two different video-based programs are run simultaneously from the central location to different sets of villages.

The infrastructure costs in the Ashwini project have been substantial: about US\$5000/village site for the networking/computing equipment alone. Major contributions to the cost are the antenna tower (US\$2000), the off-the-shelf radio (US\$1000), and the power backup (US\$2000). The current operational expenses are very high: about US\$15 to US\$20/day/site. This includes hardware/software maintenance, Internet charges, power, manning, and the cost of bringing the subject experts.

Hence, although we chose WiFi for its inexpensive equipment (two to three orders of magnitude cheaper than cellular network equipment), we have found that the overall system cost still needs optimization in order to enable commercial operation. We believe that such optimization is indeed possible: much of the above costs are high, since it is only the initial prototype. We have suggested several such optimizations in [10]. Further, on the positive side, the various video-based services in Ashwini are very popular and are taking root in the community. We hope that these will help in recovering the capital expenses and help sustain the network in the long run.

CONCLUSION

WiFi is unique in that it is an inexpensive data-centric broadband access technology. This makes WiFi a high-potential technology for rural communication services. There are several technical issues in making WiFi work in long-distance

rural settings. Despite the radio itself being low-cost, the other parts of the system may involve substantial cost and optimization is required to enhance the viability of WiFi. Apart from cost optimization, minimizing power consumption is also an important aspect to pay attention to in rural settings. Orthogonally, providing the right kind of services packaged for rural communities is essential to gaining the benefits of information and communication technology.

REFERENCES

- [1] IEEE P802.11, The Working Group for Wireless LANs: <http://grouper.ieee.org/groups/802/11/>.
- [2] A. Jhunjhunwala, B. Ramamurthi, and T. A. Gonzales, "The Role of Technology in Telecom Expansion in India," *IEEE Commun. Mag.*, vol. 36, no. 11, Nov. 1998, pp. 88–94.
- [3] IEEE 802.16, Wireless MAN, <http://www.ieee802.org/16/>.
- [4] S. Sen, Topology Planning for Long Distance Wireless Mesh Networks, Master's thesis, IIT-Kanpur, May 2006; with B. Raman.
- [5] B. Raman and K. Chebrolu, "Design and Evaluation of a New MAC Protocol for Long-Distance 802.11 Mesh Networks," *11th Annual Int'l. Conf. Mobile Comp. and Networking*, Aug./Sept. 2005.
- [6] P. Kumar, Design, Implementation, and Evaluation of New MAC Protocols for Long Distance 802.11 Networks, Master's thesis, IIT-Kanpur, May 2006; with B. Raman.
- [7] N. R. Puli, *The SRAWAN MAC Protocol to Support Real-Time Services in Long Distance Networks*, Master's thesis, IIT-Kanpur, Aug. 2006; with B. Raman.
- [8] K. Chebrolu, B. Raman, and S. Sen, "Long-Distance 802.11b Links: Performance Measurements and Experience," *MOBICOM*, 2006.
- [9] N. Mishra *et al.*, "Wake-on-WLAN," 15th Annual Int'l. World Wide Web Conf., May 2006.
- [10] S. Sen, S. Kole, and B. Raman, "Rural Telephony: A Socio-Economic Case Study," *Int'l. Conf. Info. and Commun. Technologies and Development*, May 2006.

BIOGRAPHIES

KAMESWARI CHEBROLU (chebrolu@iitk.ac.in) is an Assistant Professor with the Department of Electrical Engineering of the Indian Institute of Technology, Kanpur. Her research interests are in the areas of wireless network architecture, protocol design, and analysis. She is specifically interested in technology for developing regions. She obtained a B.E. degree in electronics and communication engineering from Andhra University (1998), and M.S. and Ph.D. degrees in electrical and computer engineering from the University of California at San Diego in 2001 and 2004, respectively.

BHASKARAN RAMAN (braman@cse.iitk.ac.in) received a B.Tech. in computer science and engineering from Indian Institute of Technology, Madras in 1997. He received M.S. and Ph.D. in computer science from the University of California, Berkeley, in 1999 and 2002, respectively. He is currently an Assistant Professor at the CSE department at Indian Institute of Technology, Kanpur. His research interests and expertise are in communication networks, wireless/mobile networks, large-scale Internet-based systems, and Internet middleware services. His current focus is on appropriate communication and computing technology for rural use.