FEASIBILITY STUDY OF SPATIAL REUSE IN AN 802.11 ACCESS NETWORK

A.R. Harish(1), **Sreekanth Garigala**(2), **Bhaskaran Raman**(3), **and Phalguni Gupta**(4)

(1) ***Deparmtment of Electrical Engineering, Indian Institute of Technology Kanpur - 208 016, India, Email: arh@iitk.ac.in***

(2) ***Godrej Castlemaine, Veritas Software India Ltd., Pune-1, India, Email: sreekanth.garigala@veritas.com***

(3) ***Deparmtment of Computer Science and Engineering, Indian Institute of Technology Kanpur - 208 016, India, Email: braman@cse.iitk.ac.in***

(4) ***As (3) above, but Email: pg@cse.iitk.ac.in***

INTRODUCTION

Despite a decade and a half of communication revolution, much of the rural population in developing countries is yet to see its benefits. The underlying reason for this is that communication technology (wired and cellular) is *value-priced* for western markets. Thus these technologies have found widespread deployment in metropolitan pockets (of developing nations) with a higher per-capita income, but not in rural areas where the density of users as well as their paying capacity are relatively low.

In the Digital Gangetic Plains (DGP) project [1] we are exploring the use of 802.11 [2] as a long-distance access technology to provide data and VoIP connectivity to rural villages. Although 802.11 was designed to be an indoor technology, it has attractive cost-economics – the equipment is *cost-priced* due to an open inter-operable standard and mass production. Our envisioned use of 802.11 is as depicted in our testbed in Fig. 1 – we have a multi-hop mesh network consisting of point-to-point 802.11 links built using directional antennas. This covers several tens of kilometres. Our testbed currently has 9 nodes at different villages, and 10 point-to-point links built using parabolic grid antennae for directional gain. We use off-the-shelf equipment of the 802.11b variant of the technology.

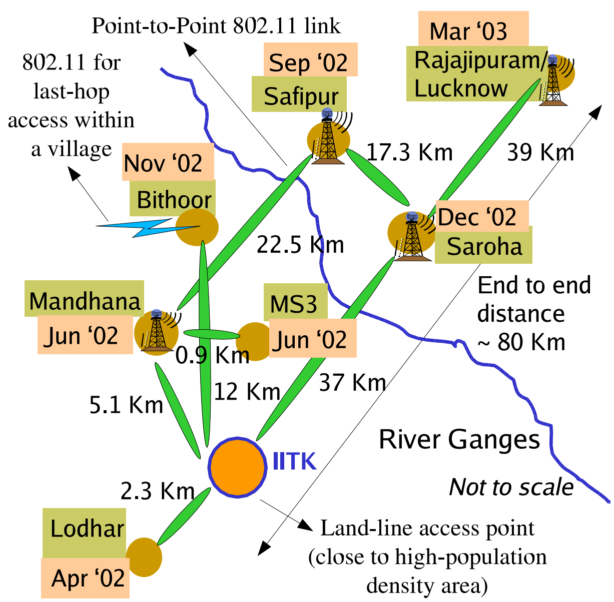

Figure 1: The Digital Gangetic Plains testbed

Spatial reuse referes to the scheduling of multiple (mutually non-interfering) transmissions simultaneously when all the links are operating in the same channel. In this paper, we explore the amount of spatial reuse possible with respect to the multiple directional links at a node. In this paper, we show that *simultaneous reception* along the links at a node, and *simultaneous transmission* from the node along these links are possible in our setting. We collectively term these two as *simultaneous synchronous operation*. We show that the power levels of various transmitters can be engineered such that the interference is rejected due to the directionality of the antennas.

SIMULTANEOUS SYNCHRONOUS OPERATION

Past work in Spatial-reuse Time-Division Multiple Access (STDMA) scheduling in multi-hop packet radio networks (e.g. [3]) has considered scenarios where a node can receive from only one neighbour at a time. With respect to transmission, a node may be able to broadcast to all neighbours, but not transmit independent information to its neighbours simultaneously. In our case, the network is static, and each node has multiple directional antennae, each pre-aligned towards a particular neighbour. For example, Fig. 2 shows node N with two directional antennae, each pointed towards a neighbour. This motivates us to consider the possibility of simultaneous operation of the links at a node.

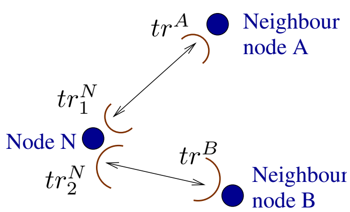

Figure 2: Multiple directional antennae at a node

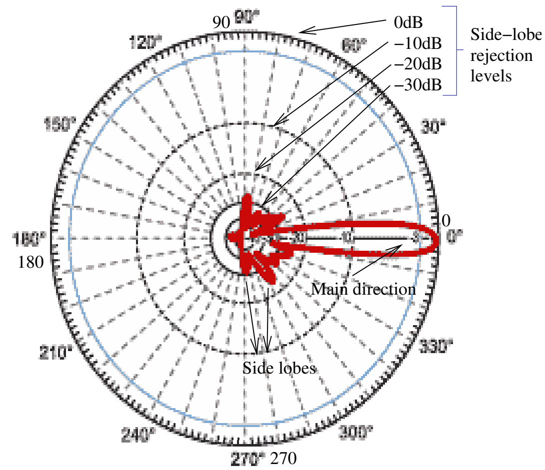

Figure 3: Spatial radiation pattern of parabolic grid antennae

Consider the two links in Fig. 2. There are three possible situations: (1) node N receiving along both transceivers tr_1^N and tr_2^N *(simultaneous reception)*, (2) N transmitting along both transceivers *(simultaneous transmission)*, and (3) N transmitting along one and receiving along the other. Now, the third scenario is not feasible since when node N is transmitting, the transmission power will be quite high near it, interfering with the receiving link. It remains for us to consider the first two scenarios, which we collectively term *simultaneous synchronous operation* (two links in synchrony with each other). We explore this in detail now.

Despite the directionality of the antennas, simultaneous synchronous operation needs careful consideration due to the following reason. The directional antennas have side-lobes (away from the main direction) along which transmission or reception can "leak". The radiation pattern for the parabolic grid antennas used in our testbed [4] is shown in Fig. 3. Given this, consider the scenario of simultaneous reception at N in Fig. 2. While tr_1^N is receiving from tr^A, it also hears the transmission from tr^B (as interference) due to tr_1^N's side-lobes. Likewise, tr_2^N sees interference from tr^A. Similarly, there is mutual interference during simultaneous transmission as well. To see if this mutual interference can be tolerated, we experimentally measure the required Signal-to-Interference Ratio (SIR) for error-free operation, using off-the-shelf 802.11b equipment. We denote this as SIR_{reqd}. We then compare SIR_{reqd} with the side-lobe rejection level of our directional antennae to argue that simultaneous synchronous operation is possible.

MEASURING THE REQUIRED SIR

Our experimental setup to measure the required SIR for error-free operation (SIR_{reqd}) is shown in Fig. 4. We have a "main" link between an 802.11 Access Point (AP) and a "main" 802.11 client. We also have an interfering link between an interfering AP and an auxiliary client. To achieve control over the various power levels, we use RF cables for the two links. The two links are physically isolated in two separate rooms. A controlled amount of interference power is fed from the interfering AP to the main client, through the use of two directional couplers (DC_1 and DC_2). Two attenuators are used: Att_1 to vary the signal power level seen by the main client, and Att_2 to vary the interference level seen by the main client.

The two clients are laptops, with 802.11b PCMCIA cards – the cards have external antenna connectors for attaching cables. Each of the APs is connected to a PC, which acts as a traffic generator. We use a continuous (saturating) UDP stream for the main link traffic, and use a continuous TCP stream on the interfering link. The nature of the interfering traffic (TCP vs. UDP) is immaterial, as long as it is continuous. The interference traffic is intended to cause packet errors in the main UDP traffic. We measure the received UDP throughput at the main client, as a function of the interference level, and thus measure SIR_{reqd} for error-free operation (of the main link).

Both the APs are configured for a transmit power of $0dBm$. The measured power level of the main AP at point B1 is $-20dBm$. The received signal power level at the client is thus $-20 - Att_1$. The measured power level from the interfering AP at B2 is $-28dBm$. The interference level at the main client is thus $-28 - Att_2 - 20$ (additional $20dB$ drop due to directional coupler DC_1). These power levels are chosen such that we are able to vary the signal/interference levels over 5 orders of magnitude, including power levels near the threshold of operation of the commercial 802.11b clients.

We note that the setup needs to ensure that one AP does not "back-off" due to the other's transmission, as will happen in the 802.11 CSMA/CA MAC. The back-off will avoid interference, which we are trying to create and measure. Since we cannot disable "back-off" in the commercial APs that we used, we isolate the two APs. Isolation is ensured by the way the attenuators and directional couplers are connected. For instance, the power level of the main AP seen at the interfering AP is $-20 - Att_1 - 40 - Att_2 - Loss_{cable2} - 20$ (start from point B1, and account for all drops). This is extremely low – at least $-120dBm$ in all our experiments – much less than typical noise level at room temperature. Additionally, we also verified that the two APs do not see each other by checking that the TCP throughput on the interfering link remains the same even when both links are up simultaneously.

We conducted four sets of experiments, each with a different setting for the main power level seen at the client (i.e., with four different settings for Att_1). In each experiment, we varied the interference level seen at the client (by varying Att_2). We measure the (UDP) throughput seen at the client, averaged over a 25sec interval. The throughput as a function of the Signal-to-Interference Ratio (SIR), for the various signal power levels is hown in Fig. 5. Note that the SIR decreases along the x-axis – this is meant to capture increasing interference along the x-axis.

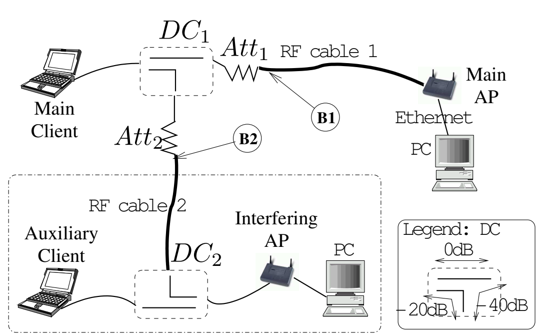

Figure 4: Experimental Setup

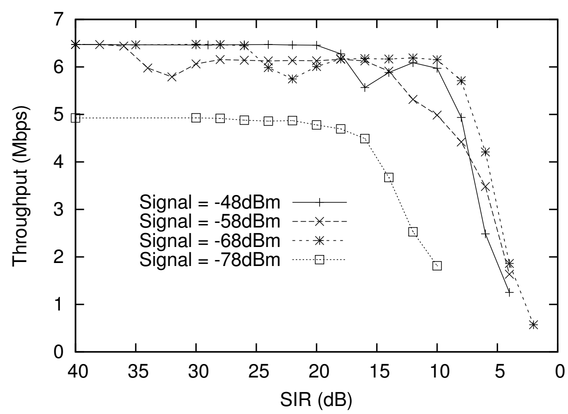

Figure 5: Throughput variation with decreasing interference

In all the four plots of Fig. 5, an SIR of $40dB$ gives the same throughput as the case where there is no interference. The throughput value in this scenario is about 6.5Mbps. This is less than the raw physical rate of 11Mbps due to the various PHY/MAC/LLC overheads. The absolute throughput is however not important for our study here, only the drop in throughput due to interference.

At a signal level of $-78dBm$, the packet error rate is high even without any interference, and hence the lower throughput. This represents the operational limit of the commercial 802.11b client used. A phenomenon that is repeatable, but one for which we do not yet have an explanation is the presence of "kinks" in the plot – where increasing interference causes increase in throughput in a small region. We are still exploring the reasons behind this.

Ignoring the "kinks", we see in Fig. 5 that the throughput shows a sharp decline when the SIR is below a certain threshold. These threshold knee points represent the required SIR value for error-free operation (SIR_{reqd}). This is about 10 to $15dB$ for the four different plots. Even taking into account the "kink" region, for a signal level of $-48dBm$, any drop in throughput happens only when the SIR is below $18dB$. We now use this measured value of SIR_{reqd}, to explain how simultaneous synchronous operation would be possible.

REQUIRED SIR AND SIMULTANEOUS SYNCHRONOUS OPERATION

Referring to Fig. 2 and the scenario where node N is receiving along both links, suppose tr_1^N receives a main signal level of P_{r_1} from tr^A, and similarly tr_2^N receives a signal level of P_{r_2} from tr^B. Now, suppose that for a given angular separation between the links, the side-lobe is S dB below the main direction. The interference seen from tr^B by tr_1^N is thus $P_{r_2} - S$, and the SIR is $P_{r_1} - P_{r_2} + S$. Similarly the SIR for tr_2^N is $P_{r_2} - P_{r_1} + S$. Now, suppose we adjust the power levels at tr^A and tr^B such that the received power levels P_{r_1} and P_{r_2} are the same, the SIR for both the links would be S.

Referring back to Fig. 3, we see that the side-lobe level is at least 25 to $30dB$ below the main lobe power level, beyond an angle of 30° or so from the main direction. This is higher than the SIR_{reqd} as measured earlier ($18dB$ or less). Thus the two links will not see the mutual interference. That is, simultaneous reception is possible.

Similarly, for simultaneous transmission, if we adjust the power levels such that the power of transmission P_{t_1} and P_{t_2} (at the transceivers at node N) are equal, simultaneous transmission is possible. Note that although we have stated $P_{r_1} = P_{r_2}$ (and $P_{t_1} = P_{t_2}$) for ease of explanation above, we strictly only require $|P_{r_1} - P_{r_2}| < S - SIR_{reqd}$ (and $|P_{t_1} - P_{t_2}| < S - SIR_{reqd}$). We have demonstrated simultaneous synchronous operation in our test bed by realizing 6.5Mbps throughtput.

We however note the following. (1) A much higher level of S can be achieved by ensuring greater angular separation – in Fig. 3, the side-lobe rejection level is $40dB$ or higher beyond an angle of $90°$. This would relax the transmission and reception power constraints. (2) Importantly, we can make use of multiple 802.11 channels. The above discussion on simultaneous synchronous operation pertains only to links operating on the same channel. If two links have different (non-overlapping) channels, they can operate independently altogether. One way to use this flexibility is to ensure that only links with sufficient angular separation are allocated the same channel. While we have considered a simple set up involving just two links, a detailed explanation of simultaneous synchronous operation in a network involving several links has been presented in [5].

CONCLUSIONS

The possibility of spatial reuse in an 802.11 access network using simultaneous synchronous operation has been demonstrated with the help of an experimental setup. The concept has been demonstrated on the DGP test bed.

REFERENCES

[1] Digital gangetic plain project, http://www.iitk.ac.in/mladgp

[2] IEEE P802.11, The working group for wireless LANs, http://grouper.ieee.org/groups/802/11

[3] A.Sen and M.L. Huson, A new model for scheduling packet radio networks, In INFOCOM, 1996.

[4] HyperGain HG2424G 2.4GHz 24dBi high performance ref-lector grid antenna, http://www.hyperlinktech.com/web/hg2424g.php

[5] B.Raman and K. Chebrolu, "Design and Evaluation of a new MAC Protocol for Long-Distance 802.11 Mesh Networks", 11th Annual Int. Conf. on Mobile Computing and Networking (MOBICOM), 2005, Germany.