A Radio Multiplexing Architecture for High Throughput Point to Multipoint Wireless Networks

Ramakrishna Gummadi
Massachusetts Institute of Technology, CSAIL
Cambridge, MA
ramki@csail.mit.edu

Rabin Patra
UC Berkeley
EECS Department
Berkeley, CA
rkpatra@cs.berkeley.edu

Sergiu Nedevschi
UC Berkeley
EECS Department
Berkeley, CA
sergiu@cs.berkeley.edu

Sonesh Surana
UC Berkeley
EECS Department
Berkeley, CA
sonesh@cs.berkeley.edu

Eric Brewer
UC Berkeley
EECS Department
Berkeley, CA
brewer@cs.berkeley.edu

ABSTRACT

WiFi-based point-to-multipoint systems are a cost-effective solution for providing high-bandwidth connectivity to remote rural regions. However, current point-to-multipoint deployments are hampered by several challenges. Their capacity to support a large number of clients simply by increasing the number of radios at the base station is limited because space restrictions on radio towers bound the number of antennas that can be installed at one physical location. Also the limited availability of wireless channels restricts the number of clients that can be simultaneously supported by a base station at any point of time.

In this paper, we explore high-throughput architectures for point-to-multipoint networks. We show how we can increase the number of radios at the base station without increasing the number of antennas simultaneously. We propose a simple yet practical multiplexing design that uses cheap RF combiner/splitter devices for multiplexing several radios onto a single antenna. We also examine a more general design that uses RF switches, and which allows us to allocate radios to antennas dynamically based on client traffic demands.

As a proof of concept, we demonstrate and evaluate the simple case of combining up to three radios operating on different channels onto only one antenna, using off-the-shelf combiner/splitters and attenuators. We show that not only is such a design feasible, but also that the achieved link throughput in both directions is as good as the one obtained by using separate antennas for each radio, as long as we provide sufficient RF isolation between the multiplexed radios.

WiNS-DR'08, September 19, 2008, San Francisco, California, USA.

Categories and Subject Descriptors

C.2.1 [**Computer Systems Organization**]: Computer-Communication Networks—*Network Architecture and Design*

General Terms

Performance, Design

Keywords

Rural wireless, Point-to-multipoint, Antenna combining

1. INTRODUCTION

Due to their cost-effectiveness and ease of deployment, wireless technologies centered around WiFi have become the preferred alternative technology for providing network coverage in developing regions. Such technologies have proved effective in bridging the wide connectivity gap between industrialized nations and developing regions [4], particularly in rural areas.

Modified WiFi hardware and protocols [11, 8] have enabled several real-world rural network deployments [1, 2, 3, 5], resulting in the emergence of a practical and cost-effective network deployment model. Figure 1 illustrates this model, which uses a combination of *a)* long-distance (5–50 km) *point-to-point* high-bandwidth backhaul WiFi links connecting towns and villages, and *b)* medium-range *point-to-multipoint* access links distributing the connectivity to schools, hospitals, kiosks and individual users. Both types of links use inexpensive, high-gain directional or sector antennas in order to increase the range of outdoor WiFi communication effectively.

This paper explores architectural design issues pertaining to the point-to-multipoint component of rural networks. Our goal is to examine point-to-multipoint configurations that deliver high throughput while reducing the number of required antennas, which we found from prior experience to be an important practical constraint in increasing throughput.

It is imperative to design point-to-multipoint access networks that afford high throughput to a large number of clients, even if the external Internet connectivity to such

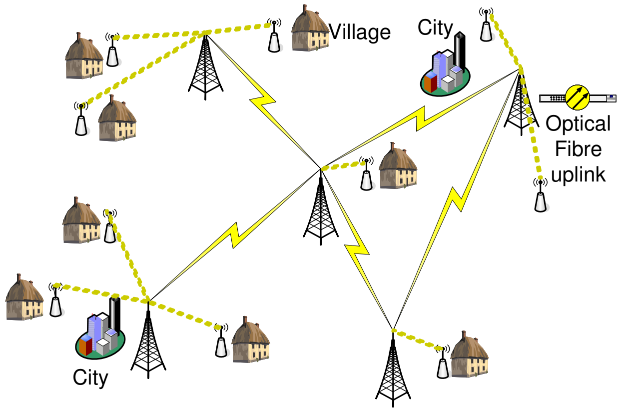

Figure 1: Example of a WiFi rural network featuring a combination of long-distance backhaul links and medium-range point-to-multipoint access links.

networks is modest. The reason it is important to provision access links with high throughput is that a single access link is usually shared by several users in tele-centers and schools, and users can rarely afford individual network connections and personal devices. Moreover, we have recently seen an increase in interactive applications such as tele-medicine [13] and remote education, which require high-quality video feeds and low-loss VoIP connectivity.

Prior work has focused on designing protocols to achieve high throughput. Protocols such as 2P [11] and WiLDNet [8] use a TDMA protocol in place of WiFi's normal CSMA/CA protocol in order to build high-throughput long backhaul links. Proposals such as SRAWAN [12] and WiFiRe [10] do the same for achieving higher aggregate throughput in point-to-multipoint settings.

While these protocol advances are significant, there is an important architectural issue in the design of point-to-multipoint networks that has received relatively little attention to date. This issue is how to design point-to-multipoint networks with a large number of radios that can provide high throughput, while using a smaller number of physical antennas to send and receive signals from these radios. This desire to design configurations that provide high throughput but require fewer antennas to be installed and maintained is driven by a practical constraint on the number of antennas that can be installed at the base station in a point-to-multipoint setting. The antenna constraint arises because of the following two considerations on the total number of towers, and the number of antennas per tower. We describe these two issues in turn.

First, there is a need for reducing the total number of towers that must be installed per deployment. With tall communication towers being the dominant factor in deployment costs [6], serving many clients with a small number of communication towers is essential. For example, in Akshaya [2], one of the largest rural wireless networks in the world, a key limiting factor in terms of deployment costs is the number of tall towers that can be erected reasonably.

Given this constraint on tower costs, the natural trend is to maximize the number of antennas per tower. But there is an upper limit to the number of antennas that can be reasonably installed per tower. This limit forces us to accept lowered throughput delivered to clients, and trade off capacity for coverage. To give a data point, while the number of client sites that can be supported by the Akshaya deployment per tower is a reasonable number that ranges between 10 and 40, each site only enjoys a typical throughput of less than 1 Mbit/s. Thus, while the existing model decreases network deployment costs, it also decreases the throughput delivered to each client and upper-bounds the number of supported client sites supported per tower to the number of antennas that can be mounted on a given tower.

The reason that existing architectures have constraints on the number of antennas that can be supported is that each additional antenna increases both the average *interference* at clients (even assuming a good implementation of TDMA protocol across all radios) and the *cost* of installation. In order to increase both capacity and tower coverage, the standard approach taken by network operators is to increase base station capacity by *sectorization*: the co-location of multiple base station radios on the same tower, and the use of multiple sector antennas, one for each radio. Assuming the availability of as many as K non-overlapping wireless channels, and assuming that S sector antennas are deployed for each channel, in theory, the per-tower capacity scales with $K \times S$. However, in practice, such scaling is limited due to interference (because the S sector antennas operating on the same channel can cause interference at both the client and the base station) and cost reasons, as explained in section 2.

Thus, given the difficulty and cost of installing and aligning many sector antennas on one tower, the number of antennas becomes the effective bottleneck in increasing the per-tower capacity. Our paper addresses this issue by proposing architectures aimed at *maximizing network capacity* while using *the minimum number of antennas per tower*. To this end, we leverage a set of techniques allowing multiple radios to connect to the same physical antenna, or allowing radios to dynamically switch the antennas they are connected to, adapting to the dynamic client traffic demand.

The remainder of the paper is structured as follows: In Section 2, we present the motivation for exploring radio multiplexing architectures. In Section 3, we describe a static multiplexing as well as a dynamic switching radio architecture. In Section 4, we evaluate the throughput achieved in practice by the static multiplexing architecture. We then conclude in Section 5.

2. MOTIVATION

In this section, we discuss the factors that motivate our architectural framework for improving capacity allocation by multiplexing radios among fewer antennas for interference, cost, and flexibility reasons.

Interference: Using several sector antennas on the same channel can lead to both client-side and base station-side interference, given the imperfect nature of the directional signal amplification obtained by using sector antennas. On the downlink direction (i.e., from base station to clients), simultaneous transmission from different sector antennas on the same channel can interfere at a client. Similarly, on the uplink direction (i.e., from clients to the base station), clients in different sectors operating on the same channel can cause mutual interference at their respective antennas. While this problem can be mitigated to some extent through careful

Figure 2: Picture of a radio tower in the Akshaya [2] deployment featuring multiple antennas for point-to-multipoint connectivity. The antennas need to be aligned on the vertical axis and aimed carefully.

aiming and positioning of sector antennas, the large number of antennas required by today's designs decreases the amount of interference that is mitigated by careful antenna alignment and positioning. Another technique for alleviating interference is to use dynamic power adaptation, as described in [9], but there are still client configuration scenarios where throughput must be sacrificed.

Cost: Deploying a large number of sector antennas at each tower is both expensive and physically challenging. Let us illustrate this point by examining Figure 2, which shows the top part of a real communication tower in the Akshaya networks, featuring 9 directional or sector antennas. While the setup shown here seems easy and feasible to the uninformed eye, the Akshaya network managers were very unhappy with the large deployment costs and even larger maintenance costs demanded by such an installation. In order to insure enough spatial isolation among several sector antennas, the communication tower was built much taller and heavier than it would have been required to ensure line of sight conditions, making it by far the most expensive tower in the entire network. Installing and aligning the antenna proved a challenging task as well.

Flexibility: In addition, since today's architectures drive each antenna through only a single radio, they're unable to adapt to *dynamically changing traffic patterns*. The most challenging part is maintaining and adapting the installation to accommodate an increasing number of users. In Akshaya, as more clients were connected, the antennas were realigned several times to accommodate the changes in workload and spatial distribution of the clients. Recurring interference problems also demanded countless rounds of antenna shifting and realignment, making it a nightmare for maintenance technicians involved. After this experience, the tower remained the only example when the installation of that many antennas was ever attempted in the network.

3. ARCHITECTURE

In this section, we discuss various ways to connect radios to antennas in order to increase system capacity while maintaining low system cost and complexity.

Consider the prevailing status quo. In this architecture, presented in figure 3, we allocate one radio per antenna. This architecture is simple, but has the drawbacks mentioned in section 2: *a)* leads to a large number of antennas – at most $K \times S$ (number of channels times number of sectors per channel); *b)* it reduces the flexibility of capacity allocation. As an example, if all active clients happen to be in the range of a single sector antenna, only one radio can be used to serve all of these clients, while the other radios remain idle. Conversely, if an antenna loses alignment or experiences connectivity failures, the radio connected to it and its corresponding throughput is effectively lost.

In order to overcome these drawbacks, we examine the following alternatives:

I) Combining several radios for each antenna: We argue that it is possible to design a multiplexing system that allocates multiple radios, operating on different channels, to a single sectorized antenna, by using a single n-port splitter/combiner device (Figure 4). This device acts as a combiner when one or more of the radios it is multiplexing are transmitting, and as a splitter when its radios are receiving. The main constraint is that, when radios are transmitting, the splitter/combiner must provide enough inter-port isolation to shield one radio from receiving the high-energy transmission of a neighboring radio, even if the two radios are on different non-overlapping (i.e., orthogonal) channels. We will evaluate this architecture in detail in section 4.

The main advantage of this approach is the fact that it only uses S sector antennas to achieves the same capacity as the default solution with $K \times S$ antennas. Since the number of orthogonal channels K is three (channels 1, 6, and 11) in the 2.4 GHz band, and up to 12 in the 5 GHz band, we can obtain significant savings in the number of required antennas. In turn, reducing the number of antennas decreases the deployment costs and simplifies installation and maintenance.

The disadvantage of this architecture is the fact that it is not load aware, with radios being statically tied to sectors, regardless of whether there is enough traffic demand for each sector. Another disadvantage is related to the loss introduced by the splitter on the receive path, when the incoming signal from the antenna is split across all radios, thereby reducing the signal received at the relevant radio by $\approx 10\log(n)$ dB, where n is the number of multiplexed radios.

II) Multi-port switch connecting radios to antennas: To address the flexibility issue, we propose a second alternative (described in Figure 5), that uses dynamic switching in order to connect multiple sectorized antennas and implement a two-stage switching. In Stage 1, each radio uses a RF-multiplexer to switch its power to a particular antenna port. In Stage 2, the output of all the RF multiplexers is combined into an antenna.

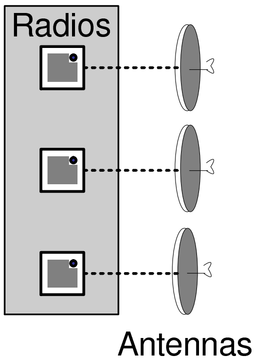

Figure 3: Basic design where each radio connects to only one antenna. We achieve a throughput of $K \times S$, where K is number of channels and S is number of sectors, only by using $K \times S$ antennas.

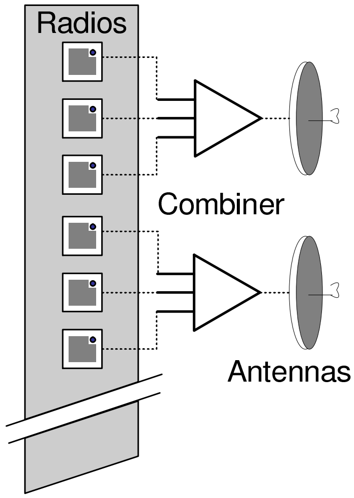

Figure 4: Each antenna is connected to $K = 3$ radios using a splitter/combiner. We achieve a throughput of $K \times S$ using only S antennas but $K \times S$ radios.

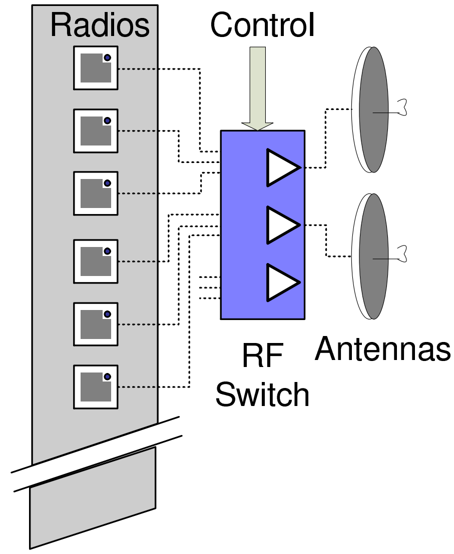

Figure 5: All the radios are connected to the antennas through a RF switch that is controlled by the host. The switch can route a radio to any of the antennas. We achieve a throughput of $K \times S$ using only S antennas and radios.

This architecture is more flexible, as radios can be steered to antennas based on dynamic demand as well as in order to tolerate remote link failures. While electronically steerable antenna systems, such as those proposed in [7], also have the potential to increase throughput by taking client traffic demands into account, electronically steerable antenna systems remain expensive (ranging several thousands of dollars upward), which decreases their attractiveness for low-cost wireless networks.

However, this architecture also has two significant disadvantages. First, there is increased cost and complexity due to the use of RF multiplexers. Second, each additional RF component introduces losses ranging from 0.7 dB up to 3 dB. One way to combat these losses is to use power amplifiers and low-noise amplifiers (LNAs) on the transmit and receive side respectively of the antennas. In order to understand how severe these isolation and signal loss problems are in practice, we evaluate the static multiplexing architecture in detail in the next section.

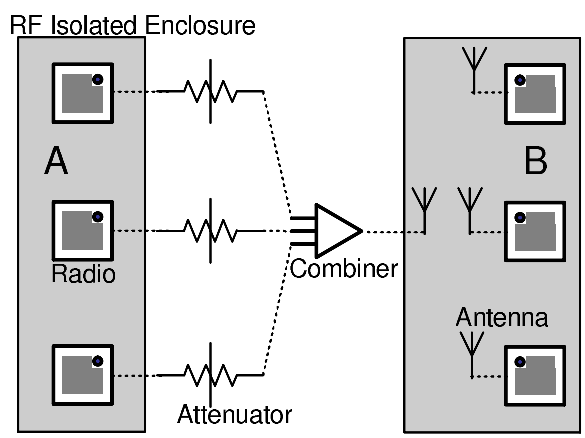

Figure 6: Experimental setup for combining multiple radios using the same base station antenna.

4. EVALUATION

In this section, we evaluate the static multiplexing design of combining multiple radios to a single antenna using off the shelf splitter/combiners (henceforth referred to simply as splitters).

In order to recreate outdoor long distance environments, we use the experimental setup described in Figure 6. We use two RF-isolated enclosures, one for the base station radios (enclosure A), and another for the client-side radios (enclosure B). Each of these enclosures contains three 802.11b/g mini-PCI Atheros-based radios (Ubiquiti SR-2) with a maximum TX power of 25 dBm, driven by WRAP wireless routers. The radios in the first box are combined using a 3-way commercial splitter/combiner, offering a 20 dB port-to-port isolation. As we will see shortly, this isolation between the radios is not enough: even if configured on non-overlapping channels, the transmit energy of one radio is too large to be fed directly to the other radio, with only 20 dB attenuation. We therefore simulate a splitter featuring more port-to-port isolation by adding additional attenuators at every base station radio.

In the other enclosure B, corresponding to client-side radios, each radio is equipped with an independent "rubber-duck " antenna.

No. of radios	Channels	All radios TX	All radios RX	Mix TX/RX
2	1,6	14.51	14.89	14.77
2	6,11	14.74	13.98	14.16
2	1,11	13.90	13.79	13.70
3	1,6,11	21.14	20.34	

Figure 7: UDP throughput (sum on all radios) achieved for various scenarios when we combine multiple radios to a single antenna.

Using this setup, and with 20 dB attenuators at each base station radio, we first test that combining two or three radios on a single antenna is feasible. Therefore, we measure the UDP throughput achievable by combining the transmissions between the three pairs of radios under three scenarios: *a) simultaneous TX* from the base station, *b) simultaneous RX* at the base station, and *c) mix of TX and RX*. We investigate several scenarios with either two or three pairs of radios used simultaneously on non-overlapping channels.

The results of this experiment are presented in table 4, and show that the sum of the throughput achieved by combining N radios to one antenna is indeed roughly equal to the throughput achievable by using N independent antennas, one for each radio (the individual UDP throughput is $\approx$ 7.5 Mbps). This total throughput is sustained when using either 2 or 3 radios connected to one antenna.

This experiment shows that combining multiple radios is feasible, but our experiment required additional attenuation at each radio in order to insure the necessary inter-radio isolation. In the next experiment, we vary this inter-radio isolation, and measure its effect on the total achieved throughput. We begin with testing the scenario where two combined radios transmit simultaneously, and we measure their combined UDP throughput.

Figure 8 shows the aggregate throughput achieved by the two radios when they transmit simultaneously, as we vary the isolation between the two radios by changing the attenuation of each radio from 10 dB to 25 dB. We are using 10 dB attenuators at both the basestation radios. The x-axis is the total port-port isolation between the two simultaneous transmitters, and the y-axis is the aggregate throughput achieved by both the radios.

To get maximum isolation, we set the transmitting radios to miniumum TX power (at 10 dBm). This corresponds to a total isolation of 55 dB between the two transmitters (the rightmost point on the x-axis) and our system achieves the maximum possible system throughput of around 15 Mbps for a channel separation of at least 4. As we decrease the effective attenuation inserted between each transmitter and the splitter/combiner by increasing the transmit power of each radio, the aggregate throughput begins to suffer. For a total isolation of 40 dB (leftmost point on the x-axis), which is achievable by commercial splitter/combiners, we find that the achievable aggregate throughput drops by a factor of 2/3, to around 8–10 Mbps.

We attribute this decreased throughput to transmission back-offs caused by CCA (Clear Channel Assessment) as well as lost acknowledgments due to lower isolation, even if the two transmitters are separated by several channels. We have two potential ways of dealing with this problem. One relies on using lower-powered transmitters, thereby providing good isolation (i.e., 50 dB or more), and using a power amplifier closer to the antenna in order to boost the signals to their normal power levels. The other option is to set transmission parameters to disable CCA and Acks and extend the current TDMA protocols to work even across channels. We are currently exploring both possibilities.

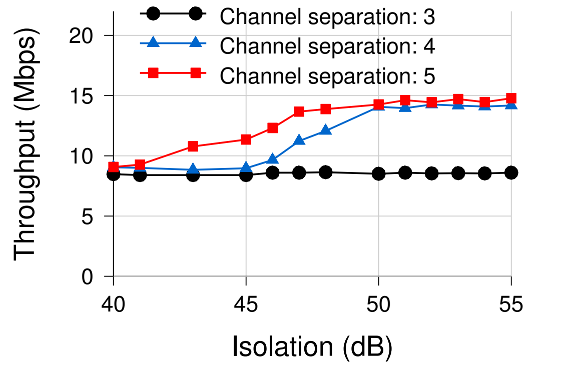

Figure 8: Aggregate throughput under simultaneous transmission for different isolation levels.

5. CONCLUSIONS

In point-to-multipoint settings, the current model of having a radio statically assigned to an antenna at the base station presents several limitations. First, in order to achieve high throughput, we need a large number of antennas, equal to the number of channels times the number of sectors. Second, we lack the flexibility today to allocate capacity according to dynamically changing demands and numbers of active clients, without increasing the number of antennas and associated radios. But increasing the number of antennas leads to high-cost towers and difficulties in maintenance, both of which represent significant problems in developing regions. The key is thus to use architectures that increase capacity while minimizing the number of antennas needed.

We presented two related architectures that assign several radios on different channels to an antenna. First, we showed a simple static approach of multiplexing radios to antennas using RF splitters and combiners. Second, we presented a dynamic two-stage approach for switching any radio to any antenna. Finally, we showed that static allocation using inexpensive off-the-shelf splitters is indeed feasible in practice, pointing to the need for further evaluation of these architectures.

6. REFERENCES

[1] AirJaldi Wireless Network. http://summit.airjaldi.com.
[2] Akshaya E-Literacy Project. http://www.akshaya.net.
[3] Aravind Eye Care System. http://www.aravind.org.
[4] E. Brewer, M. Demmer, B. Du, M. Ho, M. Kam, S. Nedevschi, J. Pal, R. Patra, S. Surana, and K. Fall. The case for technology in developing regions. *IEEE Computer*, 38(6):25–38, 2005.

[5] K. Chebrolu and B. Raman. FRACTEL: A Fresh Perspective on (Rural) Mesh Networks. In *ACM SIGCOMM Workshop on Networked Systems for Developing Regions (NSDR)*, August 2007.

[6] S. M. Mishra, J. Hwang, D. Filippini, T. Du, R. Moazzami, and L. Subramanian. Economic Analysis of Networking Technologies for Rural Developing Regions. In *WINE*, December 2005.

[7] V. Navda, A. P. Subramanian, K. Dhanasekaran, A. Timm-Giel, and S. Das. MobiSteer: Using steerable beam directional antenna for vehicular network access. In *MobiSys'07*.

[8] R. Patra, S. Nedevschi, S. Surana, A. Sheth, L. Subramanian, and E. Brewer. WiLDNet: Design and Implementation of High Performance WiFi Based Long Distance Networks. In *NSDI*, 2007.

[9] R. Patra, S. Surana, S. Nedevschi, and E. Brewer. Optimal Scheduling and Power Control for TDMA based Point to Multipoint Wireless Networks. In *ACM SIGCOMM Workshop on Networked Systems for Developing Regions (NSDR)*, 2008.

[10] K. Paul, A. Varghese, S. Iyer, and B. R. A. Kumar. WiFiRe: Rural Area Broadband Access Using the WiFi PHY and a Multisector TDD MAC. *New Directions in Networking Technologies in Emerging Economics, IEEE Communications Magazine*, 2006.

[11] B. Raman and K. Chebrolu. Design and Evaluation of a new MAC Protocol for Long-Distance 802.11 Mesh Networks. In *ACM MOBICOM*, August 2005.

[12] N. P. Reddy. The SRAWAN MAC Protocol to support Real-Time Services in Long Distance 802.11 Networks. Master's thesis, IIT Chennai, 2006.

[13] S. Surana, R. Patra, S. Nedevschi, M. Ramos, L. Subramanian, Y. Ben-David, and E. Brewer. Beyond Pilots: Keeping Rural Wireless Networks Alive. In *NSDI*, 2008.