Optimal Scheduling and Power Control for TDMA based Point to Multipoint Wireless Networks

Rabin Patra, Sonesh Surana, Sergiu Nedevschi, Eric Brewer
Department of Electrical Engineering and Computer Science
University of California, Berkeley
{rkpatra,sonesh,sergiu,brewer}@cs.berkeley.edu

ABSTRACT

In TDMA-based point-to-multipoint rural wireless deployments, co-located base station radios and sector antennas are used to increase base station capacity. To achieve maximum capacity with limited availability of non-overlapping wireless channels, we need to operate as many radios as possible from different sectors on the same channel. However, operating co-located radios on the same channel can result in substantial interference especially with the current practice of operating all radios at maximum power. We investigate techniques that increase network throughput by eliminating this interference.

To this end we formulate an LP optimization problem that maximizes throughput by computing optimal transmit schedules, optimal allocation of clients to base station radios, and optimal radio power levels. Our results suggest that there is a large gap between currently-used and optimal strategies, creating opportunities for simple, practical algorithms to address these issues. Our techniques are equally applicable to both WiFi based networks as well as other point-to-multipoint technologies such as WiMax.

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, Power control

1. INTRODUCTION

Wireless technologies are the most promising options to bridge the wide connectivity gap between industrialized nations and developing regions, particularly rural areas. Due to its cost-effectiveness, IEEE 802.11 [?] (WiFi) technology has seen widespread use beyond its originally intended purpose of short-range indoor networks. Given the sparse distribution of potential users in rural areas, WiFi has been modified by 2P [?] and WiLDNet [?] to work in medium- to long-distance point-to-point settings. Additionally, there is ongoing work such as SRAWAN [?] and WiFiRe [?] that are modifying WiFi to be higher performing in medium- to long-distance point-to-multipoint settings as well. Another promising technology in this context is IEEE 802.16 [?] (WiMAX), providing medium-range point-to-multipoint connectivity in denser urban areas.

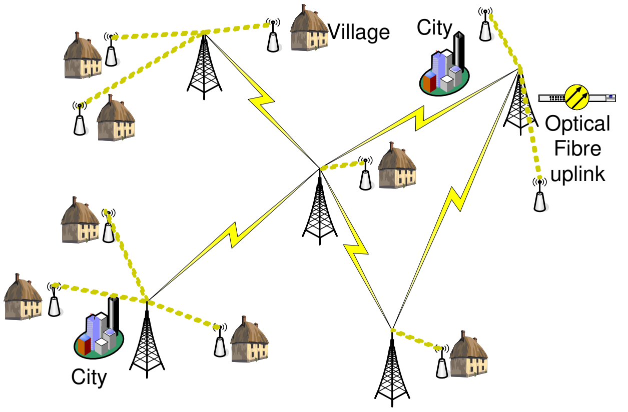

Figure 1: Example of a WiFi rural network featuring a combination of long-distance backhaul links and medium-range point-to-multipoint access links.

In the midst of all this innovation, and with growing real-world rural deployments [?, ?, ?, ?], we are seeing the emergence of a practical cost-effective deployment model: a combination of high-bandwidth, long-distance *point-to-point* (5-50 km) backhaul links connecting cities and villages, and *point-to-multipoint* access links (upto 20 km) distributing the connectivity to smaller locations such as schools, hospitals and kiosks with static clients (example network illustrated in Figure 1). The Akshaya network [?] in southern India and the AirJaldi network [?] are real examples of such networks. Often, the clients of these point-to-multipoint networks such as schools and kiosks further share the connectivity locally using WiFi or ethernet.

In this paper, we explore the design challenges for the

NSDR'08, August 18, 2008, Seattle, Washington, USA.

point-to-multipoint component of rural wireless networks where each base-station provides high bandwidth access to multiple static clients. A key challenge in such point-to-multipoint access networks is to provide high bandwidth services to a large number of clients. We require high bandwidth not only for connecting to the Internet but also for services such as video-conferencing for tele-medicine and remote education. With tall communication towers being a dominant cost factor [?], serving many clients with few communication towers is essential. For example, in Akshaya [?], one of the largest rural wireless networks in the world, the ratio of clients to each base-station is between 10 and 40. This model decreases network deployment costs, but also decreases the throughput delivered to each client.

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. With K non-overlapping wireless channels and S sectors, the maximum theoretical capacity of a base-station scales up as $K \times S$. This implies that for each channel, we need to operate S base-station radios, one each for all the S sectors. Operating all these S radios on the same channel can result in interference at both the base-station and the clients. This is because base-station radios typically operate at maximum power by default. So while the isolation provided by sector antennas prevents interference for some clients, other clients remain affected by it. Moreover, given the fixed placement of clients, the same set of clients will continue to suffer from the interference effects, making the service unfair.

In this paper we specifically investigate the use of adaptive power control to address the problem of interference for all the base-station radios that operate on the same channel. Our techniques are complementary to other design challenges in scaling up base-station capacity - such as reducing the cost of installation and maintenance of antennas by combining multiple radios on to a single antenna. Our work is also orthogonal to the lower level design of current long-distance MACs; we require only a TDMA-based slot schedule,[1] as well as the ability to synchronize time slots among co-located radios. Such features are already supported by or can be easily added to the approaches listed previously (WiMax, WiFiRe, SRAWAN, 2P, WiLD).

Our objective is to maximize the total network bandwidth, while satisfying all the per-client minimum bandwidth constraints. To achieve this, we frame our problem as an LP optimization problem, and compute an optimal combination of *a) client to base-station allocation*, *b) link transmission schedule* and *c) radio transmit powers*. We also quantify the importance of performing power control by comparing the maximum bandwidth achievable with or without this feature.

Our evaluations assume realistic antenna patterns and signal propagation characteristics, and indicate that techniques performing smart transmission scheduling and power control have the potential to substantially increase the number of clients that can be accommodated, and also the total network bandwidth. Our results motivate the need to investigate practical point-to-multipoint approaches that can make use of techniques similar to the ones presented here.

The remainder of the paper is structured as follows: in section 2 we present our motivation; in section 3 we discuss the assumptions made about the point-to-multipoint network and the requirements of the MAC layer. We also describe the various power control strategies that we consider. Next we present our LP formulation in section 4 and our evaluation results in section 5. We conclude and present future work in section 6.

2. MOTIVATION

In this section we motivate the need for using adaptive power control to improve spatial reuse for single channel settings of point-to-multipoint networks. The use of multiple radios with sectorised antennas at the base-stations, sharing a single channel can introduce interference, which can be of one of the following types:

Rx-Tx interference at base-station: When a base-station radio, say R_a, receives a packet from a client, while another base-station radio, R_b at the same tower is sending a packet, interference can occur at the base-station radio, R_a which is trying to receive packets. Due to imperfections in the directional characteristics of the sector antennas (side lobes), R_a can easily overhear R_b's transmissions, and this almost always results in the corruption or capture of the packet being received at R_a.

Tx-Tx interference at a subscriber: When client c_i can simultaneously hear transmissions from two different base-station radios (out of which only one packet is intended for c_i), interference occurs if the transmissions have *similar signal strengths.* By similar signal strengths we mean that the signal difference is less than an *isolation threshold*. For Atheros based WiFi radios, we have experimentally determined this isolation threshold to be approximately $12dB$ in past work [?] .

Rx-Rx interference at base-station: This is symmetric to the previous case, but now the interference occurs at base-station radio R_a due to conflicting transmissions from two subscribers c_i and c_j. Again, a large enough isolation between these two transmissions is required to avoid interference.

Of these three types of interference, the first type (Rx-Tx) is difficult to handle when using sector antennas, and most existing TDMA-based solutions [?, ?] disallow the situation itself by synchronizing co-located radios such that they either all *transmit simultaneously*, or they all *receive simultaneously.* We therefore look at the remaining types of interference, and propose solutions to address them.

Let us begin by analyzing one of these types of interference (Tx-Tx) in an example scenario featuring three base-station radios, each connected to a sector antenna; we assume these radios are synchronized as described above. Figure 2 shows the polar propagation plot of the base-station radios with their antennas, in the default case (used in current deployments) when they all transmit at maximum power. The propagation region of each antenna is defined as the region where the received signal strength is higher than the receive signal threshold Th_{rx}, over which a client can successfully receive the packet. For example, clients c_1, c_2, c_3 and c_4 fall in the receive lobe corresponding to radio R_1. However, if all the radios are transmitting at the same time, a few clients

---

[1] Due to long propagation delays and the use of directional or sector antennas, a carrier-sense approach to medium access control – the preferred approach for short-range wireless – is inefficient. Instead, the majority of long-distance MACs (2P, WiLDNet, WiFiRe, SRAWAN, WiMax) are TDMA-based.

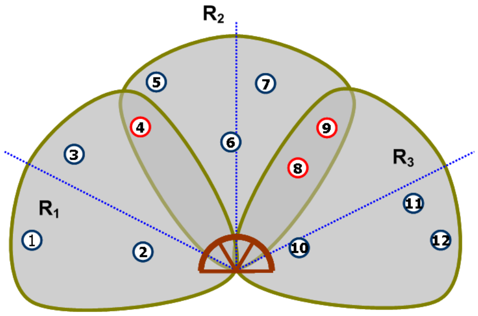

Figure 2: Basestation with 3 radios and 12 clients. All radios are using maximum power. Clients c_4,c_8,c_9 are in a region with potential interference.

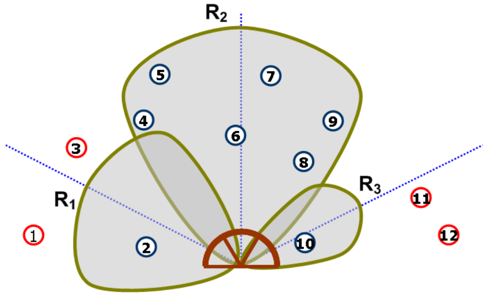

Figure 3: Basestation with 3 radios and 12 clients. While radio R_2 is using maximum power, radios R_2 and R_3 are using power control to reduce interference. Clients c_4,c_8,c_9 are now in interference-safe regions.

like c_4, c_8 and c_9 will fall in the intersection of the antenna receive lobes of two neighboring radios. At these clients, Tx-Tx interference can occur if the signal strength difference between the primary radio and the interfering radio is lower than the *isolation threshold* (Th_{cap}). The set of clients that can receive a clean signal from a radio is referred to as the *safe set* of the radio. Thus, when radio R_1 is transmitting, it's safe set consists of the clients c_1, c_2 and c_3.

One way to eliminate interference and ensure that the minimum isolation threshold is achieved is to change the transmit power at the base-stations. For example, if we want to eliminate interference at clients c_8 and c_9, we can reduce the transmit power of the radio R_3, such that clients c_8 and c_9 now fall only into the receive region of radio R_2. On the other hand, now radio R_3 cannot reach clients c_{11} and c_{12} any longer (Figure 3).

Thus the problem of minimizing Tx-Tx interference (or maximizing the number of clients that receive), can be framed as finding the optimal radio power levels, and the optimal allocation of clients to base-station radios, such that it enables the base-station to serve the most clients. The power levels and client allocation could be fixed, or they could change during every time slot, with the latter approach providing additional gains. The solution to this problem can then be used to regulate transmissions during the time slots in which the base-station radios transmit to clients.

The problem of minimizing Rx-Rx interference is very similar, and entails finding optimal transmit power levels for the client radios, together with optimal client allocation, that would enable the most clients to transmit successfully to the base-station. The solution to this optimization problem would be used during the slots when the base-station receives packets from clients.

3. ASSUMPTIONS AND STRATEGIES

In this section we describe our assumptions of the point-to-multipoint setup and MAC-layer requirements, and then present different strategies to improve spatial use.

3.1 Assumptions

A base-station has multiple co-located radios, each connected to a sector antenna. As mentioned earlier, we are specifically interested in the situation where all these radios are operating on the same channel. If additional non-overlapping channels are available, additional radios could also be used, however we would still like all the radios on a particular channel to operate simultaneously.

Fixed clients are situated upto 20 km away and have directional antennas pointing to the base-station. Time is divided into downlink and uplink slots for base-station and client transmissions respectively. Base stations can synchronize transmissions and receptions among all their radios. All client traffic flows via the base-station to the outside world. Each client has a simple interference model where the difference in received signal strengths must be greater than the isolation threshold (Th_{cap}), which we set to be $20dB$.

3.2 Point-to-Multipoint Strategies

The different possible strategies for power control are described here as we increase the degree of freedoms in controlling the base-station radios.

Always maximum power (F-MAX): The base-station radios always transmit at the fixed and maximum power setting in all slots. This could be very inefficient as there could be many clients that can never be in the safe set of any of the radios in any slot.

Greedy scheduling with maximum power (G-MAX): The base-station s use a greedy algorithm to coordinate with each other and select a set of M clients for each slot, where M is the number of base-station radios. For each slot in a period, the strategy first tries to allocate a set of clients whose bandwidth bounds have not been satisfied. Other radios that are not allocated any client are switched OFF.

Optimal maximum power allocation (O-MAX): All the base-station radios can only be set to the maximum power setting or are switched OFF during a slot. We find the best possible allocation of clients to radios for each slot in a period so that we satisfy the bandwidth bounds for each client.

Optimal fixed power allocation (O-FIXED): All the base-station radios and clients have a fixed power setting which is determined to be the optimal power setting for the best possible allocation of clients to radios.

Optimal dynamic power allocation (O-DYN): All the base-station radios can change their transmit power in every slot. These power settings and the allocation of clients to radios in each slot is optimally decided every slot.

4. OPTIMAL LP FORMULATIONS

We define RS_j as the set of clients at which the received power from radio j is above the receive power threshold (Th_{rx}). IS_j is the set of clients where the difference between the power from radio j and the other neighboring radios is less than the capture threshold (Th_{cap}), i.e. it is the set of clients which do not receive successfully from radio j due to interference. Then the *safe set* is $RS_j - IS_j$.

The aim of the optimal formulation is to select for each radio, R_j, one client c_i from its safe set at every slot such that we satisfy the minimum bandwidth requirement for each client over the whole period. This problem can be formulated as a mixed integer LP problem.

Definitions:
Set of clients: $C = (c_1, c_2...c_N)$
Set of radios: $R = (R_1, R_2...R_M)$
Set of slots: $S = (s_1, s_2...s_P)$
Path loss between client c_i and radio R_j: L_{ij}
Maximum transmit power for a radio or client: P_{max}
Sum of antenna gain between client c_i and radio R_j: G
Receive power threshold: Th_{rx}
Capture threshold: Th_{cap}
Minimum BW limit for a client: BW_{min}
Large constant:P_{inf}

Variables:
X^s_{ij}, binary variable that is 1 iff c_i is in the *safe set* of radio R_j in the sth downlink slot.
PD^s_j, transmit gain of radio R_j in the sth slot.
PU^s_i, transmit gain of client c_i in the sth slot.

Objective Functions: Maximize the total bandwidth over all the slots such that all the clients satisfy a minimum bandwidth requirement.

$$\text{Maximize} \sum_{i=1}^{N} \sum_{s=1}^{P} \sum_{j=1}^{M} X^s_{ij} \tag{1}$$

Subject to:

$$\begin{aligned} & PD^s_j + G - L_{ij} > Th_{rx} + P_{inf}(X^s_{ij} - 1) \\ & \qquad \forall s \in S, \forall c_i \in C, \forall R_j \in R \qquad (2) \\ & (PD^s_j - L_{ij}) - (PD^s_k - L_{ik}) > \\ & \quad Th_{cap} + P_{inf}(X^s_{ij} - 1) \\ & \qquad \forall s \in S, \forall c_i \in C, \forall R_j, R_k \in R \text{ and } j \neq k \qquad (3) \end{aligned}$$

$$\sum_{j=1}^{M} X^s_{ij} \leq 1, \ \forall c_i \in C, \forall s \in S \tag{4}$$

$$\sum_{i=1}^{N} X^s_{ij} \leq 1, \ \forall R_j \in R, \forall s \in S \tag{5}$$

$$0 \leq PD^s_j \leq P_{max} \ \forall R_j \in R, \forall s \in S \tag{6}$$

$$\sum_{s=1}^{P} \sum_{j=1}^{M} X^s_{ij} \geq BW_{min} \ \forall c_i \in C \tag{7}$$

Equation 2 enforces the condition that the client selected for a slot has a receive power above the receive threshold. Equation 3 enforces the condition that every selected client can capture the transmission from the corresponding radio without interference. Equation 4 ensures that every client is included in the safe set of exactly one radio. Equation 5 ensures that each radio is sending to only one client in each slot. Equation 6 enforces the maximum power constraint for each radio. Finally, Equation 7 ensures that the bandwidth of each client is at least more than the minimum bound.

If there is no feasible solution that satisfies all the clients, we fold back to solving a different LP optimization, that has as an objective the maximization of the total number of clients for which the minimum bandwidth constraint is satisfied. For this we introduce an additional binary variable Z_i, that is 1 *iff* the total bandwidth of client c_i over all the slots is greater than BW_{min}. Then, our new optimization function becomes:

$$\text{Maximize} \sum_{i=1}^{N} Z_i \tag{8}$$

,and condition 7 (which proved to be unsatisfiable) is now replaced with the following:

$$\sum_{s=1}^{P} \sum_{j=1}^{M} X^s_{ij} \geq BW_{min} Z_i \ \forall c_i \in C \tag{9}$$

Equation 9 ensures that the binary variable Z_i is set to one iff client c_i satisifies its bandwidth requirement over all the slots.

To constrain the power output of the radios according to various power control strategies discussed in section 3.2, we add further constraints.

Always Maximum power (F-MAX): Since each radio is *always* set to the maximum power, the additional constraint (Eq. 10) is :

$$PD^s_j = P_{max} \ \forall s \in S, \forall R_j \in R \tag{10}$$

Optimal maximum power (O-MAX): Here each radio can only be set to the maximum power setting, but we also have the choice of switching off a radio in a slot. We need an additional constraint (Eq. 11) to enforce this. This constraint can be transformed into linear constraints by using additional binary variables.

$$\begin{aligned} PD^s_j &= P_{max} \text{ iff } R_j \text{ is ON} \qquad (11) \\ &= 0 \text{ iff } R_j \text{ is OFF} \\ & \quad \forall s \in S, \forall R_j \in R \end{aligned}$$

Optimal fixed power (O-FIXED): Now each radio's power output is fixed for all the slots. However, we still have the choice of switching off a radio in a slot. We add another constraint (Eq. 12) that can be transformed into linear constraints by using additional binary variables.

$$\begin{aligned} PD^s_j &= PD_j \text{ iff } R_j \text{ is ON} \qquad (12) \\ &= 0 \text{ iff } R_j \text{ is OFF} \\ & \quad \forall s \in S, \forall R_j \in R \end{aligned}$$

Optimal dynamic power (O-DYN): Each radio can change its power in every slot. Equation 6 suffices and no additional constraints are required.

Receive power constraints: In practice, we can also use similar power control strategies for uplink slots from the clients to the radios. In this case, we have to ensure that the signal from the selected client is not drowned out by other clients at the radio. We can add additional constraints to the LP formulation by including an additional set of uplink slots (same as the number of downlink slots). We will need an analogous set of variables (Y_{ij}^s) to specify whether a client uses an uplink slot to a radio. The new capture constraint is shown in equations 13 and 14. We have not included the full set of equations due to space constraints.

$$\begin{aligned}
&PU_i^s + G - L_{ij} > Th_{rx} + P_{inf}(Y_{ij}^s - 1) \\
&\qquad \forall s \in S, \forall c_i \in C, \forall R_j \in R && (13) \\
&(PU_i^s - L_{ij}) - (PU_k^s - L_{kj}) > \\
&\qquad Th_{cap} + P_{inf}(Y_{ij}^s - 1) \\
&\qquad \forall s \in S, \forall R_j \in R, \forall c_i, c_k \in C \text{ and } i \neq k && (14)
\end{aligned}$$

Steerable Antennas: With electronically steerable antennas, we have more degrees of freedom for various antenna configuration parameters. To simplify analysis, we discretize the steerable antenna configuration: there are D directions (0 to 360 degrees) and L types of lobe patterns (sizes). Now we have to define the path loss for each configuration from each radio to each client. Let the path loss from client c_i from radio R_j in the antenna configuration with direction D_p and lobe size L_q be L_{ij}^{pq}. We also define ${}^sT_j^{pq}$, binary variable that is 1 iff R_j is in configuration with direction D_p and lobe size L_q, in the sth downlink slot. Then the loss on the path can be expressed as an additional linear constraint as equation 15. Finally we have to ensure that each antenna has exactly one configuration in every slot (equation 16).

$$\begin{aligned}
L_{ij}^s &= \sum_{p=1}^{D}\sum_{q=1}^{L} {}^sT_j^{pq} L_{ij}^{pq} && (15) \\
&\qquad \forall c_i \in C, R_j \in R, D_p \in D, L_q \in L, s \in S \\
\sum_{p=1}^{D}\sum_{q=1}^{L} {}^sT_j^{pq} &= 1 \quad \forall D_p \in D, L_q \in L, s \in S && (16)
\end{aligned}$$

5. EVALUATION AND METHODOLOGY

In this section, we present some preliminary evaluation of the various point-to-multipoint strategies presented in Section 3.2 for selected topologies. The objective is to quantify the gap between the naive strategies and the smarter optimal strategies which will enable us to look for practical implementations of the optimal strategies.

To that end, we first compare the different strategies in terms of the maximum possible bandwidth achievable while satisfying the minimum bandwidth constraints of all the clients in the network. In case if it is not possible to satisfy the constraints of all the clients, we try to maximize the total number of clients whose constraint can be satisfied.

We evaluate the following scheduling strategies: 1) Always maximum power (F-MAX), 2) Greedy scheduling with maximum power (G-MAX), 3) Optimal scheduling with maximum power (O-MAX), 4) Optimal scheduling with fixed power (O-FIXED) and 5) Optimal scheduling with dynamic power (O-DYN).

5.1 Methodology

Topology generation: We randomly place clients at a range of 1 – 20 km from the base station within an angular coverage of 180 degrees.

Radio antenna pattern: We place three radios at the base station. The radios are connected to fixed antennas that are placed at an equidistant angular separation. The maximum power of each radio is set to $P_{max} = 23$ dBm (200 mW). The antenna gains of the base station antennas and the client-side antennas are assumed to be 18 dBi at 2.4GHz. The capture threshold is set to $Th_{cap} = 20$ dB which is conservatively set with respect to measurements conducted by us previously [**?**]. The antenna lobe size is selected such that roughly one-third of the clients are reachable by more than one radio.

Mixed Integer Linear Program (MILP) solving: We formulate all the point-to-multipoint scheduling strategies except the greedy scheduling as MILP problems (Section 4). We solve the optimization for a fixed number of slots (referred to as a period) which gives us a periodic schedule. We set the number of slots in a period to be equal to the number of clients. This is to ensure that it is always possible to find a schedule where each client is served at least once during the entire period. This also bounds the maximum delay of each client to the number of slots in the period.

We then use the CPLEX LP solver [**?**] to optimize the mixed integer LP formulations. We use a maximum time limit of 600 seconds to find optimal solutions; this was generally sufficient in most cases.

5.2 Evaluation

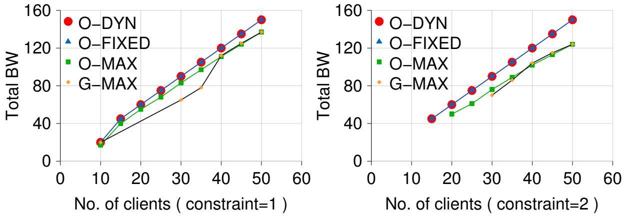

Figure 4: Maximum bandwidth with increasing clients. The naive maximum power strategies (F-MAX) did not find any feasible solution. For 3 radios.

Maximum overall bandwidth: We first compare the strategies in terms of the total bandwidth while satisfying all client constraints. Figures 4(a)and 4(b) show the maximum bandwidth achieved by the different algorithms for bandwidth constraint of 1 slot per period (50 slots) and 2 slots per period (50 slots) respectively as we increase the number of clients in the network. None of the strategies were able to achieve perfect scheduling of 3 slots per period for all the clients (50 slots and 3 radios). The F-MAXstrategy was not able to find any feasible solution that satisfied all clients even for the smallest constraint. The greedy strategy (G-MAX), was only able to find feasible solutions for larger topologies (clients > 30) where there is more flexibility to choose slots. The optimal scheduling strategies (O-MAX, O-FIXED, O-DYN) were able to satisfy all clients but there was a gap (upto 17%) between the strategy that can use only the maximum power (O-MAX) and the others. We can see that this gap is higher as we increase the minimum constraint to 2 slots. It can also be seen that both the Optimal Fixed (O-FIXED) and the Optimal Dynamic (O-DYN) strategies achieve the maximum possible bandwidth.

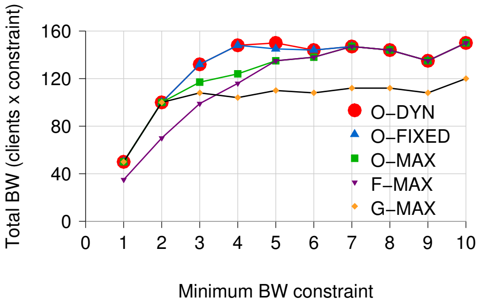

Figure 5: Maximum number of clients with increasing the bandwidth demand per client. Y-axis is the total bandwidth utilized in slots i.e. product of the no. of clients and the minimum constraint. For 50 clients and 3 radios.

Maximize clients with bandwidth constraints: When it is not possible to satisfy the constraints for all the clients, we want to maximize of the number of clients whose bandwidth constraint can be actually satisfied. Figure 5 shows the number of clients that satisfy the minimum bandwidth constraint as we increase the constraint for a topology with a fixed number of clients (50). On the Y-axis, we plot total bandwidth which is being utilized i.e. the product of the number of clients and the bandwidth constraint. We can see that there is a substantial gap between the optimal strategies and the fixed power strategy (about 50% at constraint of 3). This gap reduces as we make the constraint tighter. All the strategies can reach upto the maximum capacity of the network (150 for 50 slots and 3 radios) as it becomes easier to just select a small subset of clients that can be scheduled multiple times easily. It is interesting to note that even the strategy with always maximum power (F-MAX) also manages to reach a subset of the clients. The greedy strategy (G-MAX) does much worse than all the other solutions at higher constraints. We can also see that both the Optimal Fixed (O-FIXED) and the Optimal Dynamic (O-DYN) strategies achieve the maximum possible bandwidth all throughout. To summarize, the part of the graph which is relevant in practice is where the minimum bandwidth constraint is lower and here we see that the optimal power strategies outperform the greedy strategies.

6. CONCLUSION

Scaling up the capacity of point-to-multipoint networks using multiple co-located base-station radios connected to multiple sector antennas raises many design challenges. These include handling interference between the different base-station radios, time synchronization between base-station radios and clients, providing bandwidth and loss guarantees to clients and reducing the cost (both installation and maintenance) of using large number of antennas on towers.

In this paper, we have specifically looked at the problem of interference (both at the base-station and at the clients) resulting from simultaneous operation of multiple co-located base-station radios on different sectors but on the same channel. Our work investigates the use of transmit power control to alleviate this situation. We formulate the problem of using smart link scheduling and various link transmit power control strategies to maximize bandwidth across all clients subject to satisfying their minimum bandwidth requirements as a linear program. We also provide a formulation for the case of electronically steerable antennas, an area of active research for cost-effective deployments. With our formulations, solving the LP gives us an optimal allocation of clients to radios and also a schedule of transmissions. We show that we can schedule upto 50% more clients compared to naive strategies indicating that there is an opportunity to leverage transmit power control in real deployments. Our approach is complementary to any TDMA-based MAC. Given these early results, we are now looking at simple practical algorithms that can provide gains using transmit power control in such point-to-multipoint networks.