

Link Layer Multicasting with Smart Antennas: *No Client Left Behind*

Souvik Sen, Jie Xiong, Rahul Ghosh, Romit Roy Choudhury
Dept. of Electrical and Computer Engineering, Duke University, Durham, NC 27708
Email: {souvik.sen, j.x, rahul.ghosh, romit.rc}@duke.edu

Abstract—Wireless link layer multicast is an important service primitive for emerging applications, such as live video, streaming audio, and other content telecasts. The broadcast nature of the wireless channel is amenable to multicast because a single packet transmission may be received by all clients in the multicast group. However, in view of diverse channel conditions at different clients, the rate of such a transmission is bottlenecked by the rate of the weakest client. Multicast throughput degrades severely. Attempts to increase the data rate result in lower reliability and higher unfairness. This paper utilizes smart beamforming antennas to improve multicast performance in wireless LANs. The main idea is to satisfy the stronger clients with a high-rate omnidirectional transmission, followed by high-rate directional transmission(s) to cover the weaker ones. By selecting an optimal transmission strategy (using dynamic programming), we show that the multicast throughput can be maximized while achieving a desired delivery ratio at all the clients. We use testbed measurements to verify our main assumptions. We simulate our protocol in Qualnet, and observe consistent performance improvements over a range of client topologies and time-varying channel conditions.

I. INTRODUCTION

Emerging wireless applications, such as MobiTV [1], electronic classrooms [2], and WiFi telecasts in smart homes [3], are demanding link layer support for group communication. An access point (AP) is expected to disseminate content to all members of a group that subscribe to a common service. Serving these members through individual unicasts is a feasible but inefficient option. An ideal approach is wireless multicast, wherein, a packet may be delivered to all members of the group through a single transmission. Such an apparently simple multicast service involves various research challenges. (1) Clients scattered around an AP experience dissimilar channel conditions, resulting in different data rates that each can support. Network measurements have shown that such scenarios are pronounced due to shadowing and wireless blind-spots in indoor environments [4]. As a result, a single transmission to all the clients is bottlenecked by the data rate of the weakest client [5]. The multicast throughput can severely suffer [5], [6] due to this restriction. (2) The time-varying nature of the wireless channel causes the bottleneck data rate to change over time. A multicast protocol needs to adapt to this variation by identifying the bottleneck link first, followed by suitable rate control. In the absence of per-client acknowledgment,

bottleneck identification may not be trivial. (3) Even if bottleneck rate is suitably identified, packet losses are still possible due to fading and interference. The protocol will need to recover from such losses so that clients achieve an application-specified reliability. This paper aims to design a link layer multicast service that addresses these challenges in the context of WiFi networks. A practical solution is of interest that can accomplish high multicast throughput, while meeting a required per-node delivery ratio. Increasing transmission rates does not resolve the challenges, since some weak clients will fail to receive transmissions at higher data rates, and thus, be “left behind”. We believe smart antennas offer new opportunities to augment the state of the art in link layer multicast. We motivate the applicability of smart antennas, and present our main ideas next.

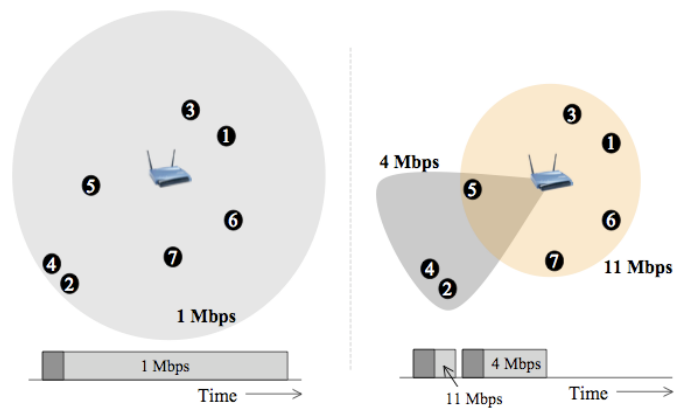


Fig. 1: Comparing a single 1 Mbps omnidirectional transmission against multiple high data rate transmissions.

Recent advances in signal processing and antenna design are enabling small, cost-effective smart antennas [7], [8]. Briefly, smart antennas offer a variety of beamforming capabilities that can support transmissions at higher data rates, without degrading reliability. However, unlike omnidirectional antennas, beamforming antennas are impaired by a smaller angular coverage, allowing a single transmission to only serve a subset of the multicast group. This inspires the possibility of multiplexing between smart and omnidirectional antennas to achieve a balance between throughput, reliability, and coverage.

Our main idea is to cover the strong clients with a high data-rate omnidirectional transmission, and then, service the weaker ones with high data-rate beamformed transmissions (Figure 1). We argue that the time consumed by multiple high data rate transmissions (provided beams and rates are chosen carefully), can be smaller than the time of a single omnidirectional transmission at the bottleneck rate. The reasons are two fold – (i) testbed measurements show that weak clients are typically a minority, and (ii) they tend to be spatially clustered in shadowed areas or wireless blind-spots. Covering all these weak clients may not require too many beamformed transmissions, facilitating performance improvements with smart antennas. Of course, translating this intuition into a complete system solution raises several research questions. Addressing them efficiently is the goal of this paper.

We propose *BeamCast*, an antenna-aware protocol that maximizes multicast throughput under specified reliability requirements. The protocol consists of 3 components, namely (1) a measurement based Link Quality Estimator, (2) a MultiCast Scheduler, and (3) a Retransmission Manager. In steady state, the *multicast scheduler* periodically consults the *link quality estimator* to obtain per-client beam-directions and data rates. This information is used to feed a dynamic program that selects an optimal transmission strategy. Once a batch of packets are transmitted using this strategy, client feedbacks are assimilated at the *retransmission manager*. Based on the distribution of packet losses across different nodes, a subset of packets are retransmitted to meet the reliability requirements. We implemented BeamCast in Qualnet, and experimented over a broad range of scenarios. Performance results show consistent improvements over omnidirectional schemes, especially when the channel quality varies over time. We believe that our protocol is practical, efficient, and cost-effective for real-life environments. Our main contributions are summarized as follows.

(1) Validation of the challenges and opportunities through measurement and analysis. Measurements using Soekris boards and laptops verify our assumption about bottleneck clients. Theoretical analysis shows that addressing bottleneck clients individually can improve performance.

(2) An optimal rate and beam selection algorithm that maximizes throughput for a given delivery ratio. An $O(n^2)$ dynamic program yields the optimal strategy.

(3) A link layer multicast protocol executes this strategy, coping with channel variations and transmission losses. A subset of the lost packets are retransmitted to achieve the required delivery ratio. A

heuristic is used to select a subset that reduces the total time of transmissions.

(4) Performance evaluation through Qualnet simulations, using different metrics. We evaluate multicast throughput, delivery ratios, and fairness, under varying fading models. We show that except in rare occasions, “no client is left behind”.

The rest of this paper expands on each of these contributions. Issues and limitations are discussed in Section VI, followed by related work in Section VII. Finally, we conclude the paper with a brief summary in Section VIII.

II. SYSTEM SETTING

We consider IEEE 802.11 based WLANs. Each access point (AP) is equipped with a smart beamforming antenna, while all the clients have simple, omnidirectional antennas. The clients are scattered around the AP, and remain stationary in the time scale of packets. The environment is characterized with multipath and shadowing effects, resulting in wireless blind spots (particularly indoors). We assume that the link layer supports multicast addressing, and hence, only clients that subscribe to the multicast service can receive the packets.

Antenna Models and Assumptions

The term smart antennas represents antennas ranging over a wide spectrum of capabilities, complexity, and cost [7]–[10]. Two categories are popular, namely MIMO and beamforming. While both these antennas offer improvements in transmission rates, this paper focuses only on the regime of beamforming antennas. Nonetheless, we believe that our basic ideas can be extended to MIMO systems as well.

Beamforming antennas regulate the radiation and reception patterns such that SINR can be maximized for a given interference environment. The antenna electronically guides most of its energy in a software-specified direction, called the mainlobe. Some energy leaks out in other directions, called sidelobes. The higher energy intensity along the mainlobe improves the SINR at the receiver, resulting in improved data rates over omnidirectional antennas. The improvement is asymptotically bounded by $C = W \log_2(1 + SINR)$, where C is the capacity and W is the bandwidth in use. Although the improvement is logarithmic, commercial antennas [7] offer more than $15dB$ mainlobe gains [11], that in turn result in higher data rates. Figure 2(a) and (b) validates this through some of our measurements with the Phocus Array Antenna.

With adaptive beamforming, beams can be reshaped dynamically from omnidirectional to directional patterns,

and vice versa. Moreover, beams can be steered near-continuously, causing high spatial overlap between adjacent directional beams. In this paper, we assume realistic beam patterns with beamwidths between 45° to 90° . We assume that switching delay is negligible, although a non-negligible delay can easily be incorporated into BeamCast.

Performance Metrics

We evaluate BeamCast with 3 main metrics as follows.

(1) **Multicast Throughput** is defined as the average number of packets received by the multicast clients per unit time. More formally, let us assume that an AP multicasts M packets over a T_m time window. Let n denote the number of multicast clients, and let m_i denote the number of packets received by i^{th} user. Multicast throughput, MT , is then defined by

$$MT = \frac{\sum_{i=1}^n m_i}{nT_m} \quad (1)$$

(2) **Fairness** is used to compare the performance between strong and weak clients, when using our scheme. We use Jain's Fairness Index below, where $f(\cdot) \in [0, 1]$ is the network's fairness, x_i is an individual node's throughput ($\frac{m_i}{T_m}$), and n , the total number of clients.

$$f(x_1, x_2, \dots, x_n) = \frac{(\sum_{i=1}^n x_i)^2}{n * \sum_{i=1}^n x_i^2} \quad (2)$$

(3) **Minimum Delivery Ratio** is the fraction of transmitted packets that every client must receive. The delivery ratio for client i is defined as follows

$$DR = \frac{m_i}{M} \quad (3)$$

where M is the total number of transmitted multicast packets. Minimum delivery ratio of a network, $MinDR$, is the minimum DR over all clients in that network. We assume that the multicast application will specify a $MinDR$ to be attained by each client.

III. MEASUREMENT, ANALYSIS, AND FORMULATION

Our main observation is that real WLAN scenarios are characterized with multiple weak clients, many of which tend to be spatially clustered. Based on this assumption, we show that grouping clients into multiple transmissions can improve multicast performance. As mentioned earlier, the idea is to perform high rate omnidirectional transmissions to the stronger clients, followed by high rate beamformed transmissions to the weaker ones. *While the number of transmissions increases, each high-rate transmission can finish earlier, adequately compensating for the overhead of multiple transmissions.* While selecting appropriate beams and rates is the objective of this paper, we first need to verify our main observation. This section reports measurement results and analysis to validate that (1) real WLANs are typically characterized with a few,

spatially clustered, weak clients, and (2) that servicing the weak clients through beamforming holds potential of performance improvements.

A. Measurements

We used Soekris boards [12] and laptops, running MadWiFi drivers on 802.11b Atheros interfaces, to measure channel quality in a multicast setting. Clients were scattered at different positions around an AP, resembling topologies like labs, classrooms and cafes. The AP was made to transmit broadcast packets at different data rates; clients measured the delivery ratio (using the sequence number in each received packet). *Tcpdump* was used to gather data rates, RSSI, and SNR values from *radiotap* headers. Figure 2(c) shows the delivery ratios (DR) at each client for increasing transmission rates. The graph is derived from a single representative topology. Table I summarizes results from 4 other topologies, with 25 clients each. The table shows the fraction of clients that experiences a maximum of 1, 2, 5.5, or 11 Mbps data rates.

TABLE I: Max. data rates for client fractions.

Topology#	1Mbps	2Mbps	5.5Mbps	11Mbps
1	10%	5%	5%	80%
2	15%	5%	20%	60%
3	15%	10%	10%	65%
4	5%	5%	0%	90%

Evidently, topologies are characterized with weak clients. Moreover, weak clients were frequently collocated in shadowed regions and blind spots in our building. Figure 2(d) shows a few identified spots. Our measurements, along with others in [4], [13], [14], are reasonable evidence that indoor WiFi environments are characterized with few spatially-clustered, weak clients.

B. Analytical Model

While few weak clients may exist in WLANs, it's important to show that removing them can offer benefits. For this, we model multicast performance analytically, and study the impact of gradually removing weak clients. Our model assumes that each client has a packet error probability dictated by the quality of its link to the AP. For different error probabilities, we compute the expected number of transmissions, $E[T]$, for successful multicast. Then, by removing weak clients incrementally (starting from the weakest), we show that $E[T]$ decreases non-linearly. Throughput can be derived from the behavior of $E[T]$, theoretically confirming the opportunity for improvement. We present this analysis next.

Let p_i denote the probability of successful packet reception at client i ; the error probability q_i is then $(1 - p_i)$. The expected number of transmissions for unicast, $E[T_{uni}]$, is clearly $\frac{1}{(1 - q_i)}$. However, for multicast, the expected




Fig. 2: (a) Phocus Array antenna (b) Delivery ratio and rate improvement due to Phocus beamforming (c) Delivery ratio at each client for increasing transmission rate (d) Shaded rectangles indicate blind spots in ECE building.

number of transmissions, $E[T_{multi}]$, is derived as follows¹. Let us assume that error probabilities are independent across different nodes [15]. We compute the probability for the case in which all clients successfully receive the packet within j transmissions. If n is the total number of clients, probability that the i^{th} client will successfully receive the packet in at least one of j transmissions is $(1 - q_i^j)$. Hence, the probability that all the clients will successfully receive the packet within j transmissions is $\prod_{i=1}^n (1 - q_i^j)$. Similarly, the probability can be computed for $j - 1$ transmissions. The difference between these two probability mass functions (pmf) represents the probability that all nodes have successfully received the packet at the j^{th} transmission. Mathematically,

$$Pr[T_{multi} = j] = \prod_{i=1}^n (1 - q_i^j) - \prod_{i=1}^n (1 - q_i^{j-1}) \quad (4)$$

Thus, the expected number of transmissions for wireless multicast, $E[T_{multi}]$ is:

$$E[T_{multi}] = \sum_{j=1}^{\infty} j \times \left(\prod_{i=1}^n (1 - q_i^j) - \prod_{i=1}^n (1 - q_i^{j-1}) \right) \quad (5)$$

We use per-client error probabilities from our testbed measurements, and use them for computing $E[T_{multi}]$ in Figure 3. The X axis reflects the number of remaining clients ($X = i$ implies the removal of $(20 - i)$ clients from the complete multicast group). Evident from the graph, $E[T_{multi}]$ decreases sharply after the removal of few of the weakest clients. The improvements saturate when further weak clients are removed. This motivates the need to service a suitable group of weak clients separately (through beamforming). Figure 4 shows the corresponding throughputs, when each client is removed individually, and serviced with a high rate, beamformed transmission. Observe that the benefit of beamforming (in comparison to a single omnidirectional transmission) is substantial. Further benefits may be feasible if a beam is used to serve more than one client. The problem can be rich, as elaborated in the next subsection.

¹Observe that $E[T_{multi}]$ is not $\frac{1}{\min(1-q_i)}$.




Fig. 3: Change in $E[T_{multi}]$ with removal of weak clients




Fig. 4: Throughput analysis with smart antenna multicast

C. Problem Formulation

Our measurement and analysis indicate that WLANs are characterized with weak clients, and beamforming to them individually can offer performance improvements. However, the improvement is a function of the number of (weak) clients covered through beamforming. Figure 4 demonstrates the existence of an optimal. Beamforming to too few or too many clients produces sub-optimal results. Hence, the first problem is of client partitioning. Specifically, given a set of clients and their individual data rates from the AP, which group of clients should be serviced through omnidirectional communication? The data rate of the omnidirectional transmission will be governed by the weakest client in this group. Remaining clients can be serviced through individual beamforming. The above problem is simple when each beamformed transmission satisfies only one client (i.e., narrow

beamwidths). In reality, antenna beamwidths are reasonably large, and may be exploited for satisfying multiple clients in one transmission. Moreover, beamforming antennas can be steered near-continuously, resulting in significant spatial overlap between adjacent beams. Hence, it may be feasible to cover a given set of clients with different sets of (overlapping) beams. Observe that not all these beam-sets will achieve identical performance. The transmission rate of a beam will vary based on which other beams are included in its beam-set. The optimal choice of beam-sets and (corresponding) data rates will maximize multicast throughput. We present an example to illustrate this better.

Figure 5 shows four overlapping beams B1, B2, B3, B4 covering client sets $\{1,2\}$, $\{2,3\}$, $\{3,4\}$, and $\{4,5\}$ respectively. Each client is annotated with data rate that it can sustain. Observe that different beam-sets, $\{B1, B2, B3, B4\}$, $\{B1, B3, B4\}$, $\{B1, B2, B4\}$, etc., can cover all the clients. However, the optimal choice is $\{B1, B3, B4\}$ with rates of $\{7, 3, 11\}$ Mbps respectively. The other beam-sets achieve sub-optimal rates of $\{9, 7, 3, 11\}$ and $\{9, 3, 6\}$ Mbps respectively, resulting in lower throughput. Choosing the optimal beam-set, and assigning corresponding rates to each of these beams, is non-trivial. This paper aims to develop a multicast protocol that will optimally partition clients into omni and directional beams, and accomplish transmissions at optimal data rates. The objective is to maximize multicast throughput while meeting a specified delivery ratio.

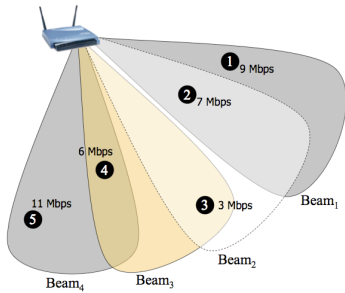


Fig. 5: Problem of choosing optimal beams and rates among spatially overlapping beams, such that multicast throughput is maximized.

Importantly, time-varying channel fluctuations and collisions affect data rates and delivery ratios. An ideal multicast protocol should be able to adapt to such changes. Suitable retransmission schemes need to be designed to recover from failures, and thereby, meet requisite delivery ratios. The problem is harder than unicast because multicast services typically do not expect per-packet client feedbacks in the form of acknowledgments. This paper addresses the above problems through the proposed protocol, BeamCast.

IV. PROTOCOL DESCRIPTION

BeamCast consists of 3 main modules: (1) a Link Quality Estimator, (2) a Multicast Scheduler, and (3) a Retransmission Manager. The protocol executes in rounds, each round corresponding to a batch of packet transmissions. At the beginning of a batch, the estimator estimates the data rates for different clients (using feedbacks from the previous batch). Using the estimated rates, the scheduler computes the optimal set of $\langle beam_i, rate_i \rangle$ tuples that maximizes multicast throughput for a pre-specified minimum delivery ratio (MinDR). Packets are disseminated according to this schedule. Clients receive (subsets of) these packets, and send batch-wise PHY/MAC layer feedbacks. The retransmission manager assimilates all the client feedback, and retransmits a minimal subset of the lost packets (to satisfy MinDR at all clients). The feedbacks are also forwarded to the link quality estimator, which in turn prepares the scheduler for the next batch of packets. We describe the functionalities of each of the module next.

A. Link Quality Estimator (LQE)

At network initiation, the AP broadcasts a batch of HELLO packets at every data rate. Clients record these hello packets along with channel-related information, including RSSI, SNR, etc. Each client then computes the delivery ratio, the average SNR, and average RSSI values, for each data rate. The summary is sent back to the AP using an AP-specified TDMA schedule. Once the network is operational, a similar feedback mechanism is exercised for every batch of data transmissions.

The LQE's job is to process the feedbacks, and estimate the maximum transmission rate that each client can support. These rates determine the transmission strategy for the next batch of transmissions. For this, the LQE uses a combined theoretical and learning approach. Specifically, continuous client feedbacks are assimilated in a database. The database consists of the average SNR that achieves the minimum delivery ratio (MinDR) for a given data rate. For example, at 11 Mbps, the average SNR that achieves 90% DR may be 19dB. At the end of a batch, if a client's DR decreases, the LQE extracts its SNR values, and consults the database to obtain the largest smaller data rate that satisfies the MinDR requirement. However, when the SNR increases again, LQE suitably increases the transmission rate. The client is assigned this selected data rate. Over time, the database values are updated with new client feedbacks. This may be possible even when the AP sends unicast traffic to the client. If SNR and DR information at some data rate were not updated for a long duration (perhaps because the client did not receive packets at that rate), LQE resorts to theoretical values as described next.

We present the theoretical relationship between delivery ratio (DR) and SNR, for different rates. On average, DR is the probability of correct reception of a packet. If L is the length of a packet, and BER, the bit error rate, then DR can be expressed as:

$$DR = (1 - BER)^L \quad (6)$$

In 802.11, due to different modulation schemes, different data rates experience different BERs for the same SNR. While 1 and 2 Mbps employ DBPSK and DQPSK respectively, both 5.5 Mbps and 11 Mbps employ CCK modulation. In the interest of space, we only present the BER expression for CCK [16], [17] as follows.

$$BER_{CCK} = 1 - \frac{1}{\sqrt{2\pi}} \int_{-K}^{\infty} \left(\frac{1}{\sqrt{2\pi}} Z \right)^{\frac{N}{2}-1} \exp\left(-\frac{v^2}{2}\right) dv \quad (7)$$

where

$$Z = \int_{-(v+K)}^{(v+K)} \exp\left(-\frac{y^2}{2}\right) dy \quad (8)$$

and N is the number of possible transmitted signal vectors, and $K = \sqrt{2E_b/N_0}$. Of course, $\frac{E_b}{N_0}$ denotes the ratio of average energy per bit to the noise power spectral density at the receiver input (for Additive White Gaussian Noise (AWGN) channel). By substituting for BER in equation (6), we can compute the values of DR for different modulation schemes. The theoretical curves are plotted in Figure 6.

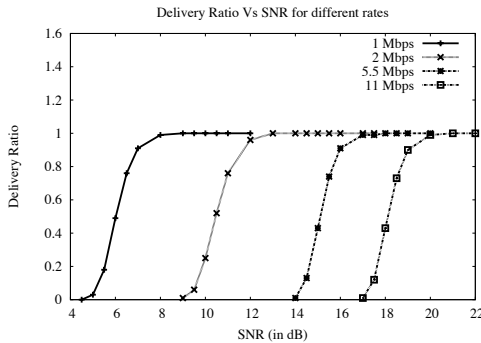


Fig. 6: Delivery ratio (DR) of different rates in 802.11b

B. Multicast Scheduler

The LQ Estimator provides the estimated per-client data rates, to be used for the next round of transmissions. From this list of n data rates, where n is the number of clients, the multicast scheduler extracts all $m \leq n$ distinct data rates. The m data rates are sorted in increasing order, denoted as R_{omni} . With beamforming, we assume that each rate is increased by a multiplicative factor, $K > 1$.

$$R_{omni} = \{r_i^o : i \in \{1, 2, \dots, m\} \text{ and } r_i^o < r_j^o \text{ if } i < j\}$$

The scheduler iterates through all the m distinct transmission rates, each iteration corresponding to a transmission strategy. Strategy i consists of two components: (1) an omnidirectional transmission at rate r_i^o , and (2) one or many beamformed transmissions (at suitable rates) that cover all clients not covered by the omnidirectional transmission at rate r_i^o . For the second component, the optimal choice of beams and data rates is non-trivial due to the overlapping nature of beams. The multicast scheduler uses a dynamic program (DP) to generate the optimal beams and corresponding data rates that maximize multicast throughput for a given MinDR (detailed in the next subsection). For strategy i , let us denote the optimal set of beams and (corresponding) rates as β_i and ρ_i . Note that rates in ρ_i are for beamformed transmission, hence, scaled by the factor K . Hence, the optimal multicast schedule, S_{opt} , can be expressed as:

$$S_{opt} = \min\left\{\frac{L}{r_i^o}I + \sum_{\forall r_j \in \rho_i} \frac{L}{r_j}\right\}, \quad \forall i \in \{1, 2, \dots, m\} \quad (9)$$

where L is the size of a data packet. I is an indicator variable which is set to 0 if pure beamformed transmission is used, otherwise it is set to 1. The terms $\frac{L}{r_i^o}$ and $\frac{L}{r_j}$ denote the time of omnidirectional and beamformed transmissions, respectively. In 802.11, each packet is preceded with a PLCP header, H , transmitted at the base data rate, r_{base} . To account for this header, the scheduler can be rewritten as:

$$S_{opt} = \min\left\{\left(\frac{L}{r_i^o} + \frac{H}{r_{base}}\right)I + \sum_{\forall r_j \in \rho_i} \left(\frac{L}{r_j} + \frac{H}{r_{base}}\right)\right\}, \quad \forall i \in \{1, \dots, m\}$$

At the beginning of every round, the AP executes this scheduler and selects the optimal transmission strategy, S_{opt} . To understand the complexity of the scheduling scheme, we present the details of the Dynamic Program.

Dynamic Program:

The inputs to the dynamic program is a set of clients that must be covered with an optimal selection of (potentially overlapping) beams. For this, we first perform a translation from the clients' cartesian coordinates to radial coordinates, (R, θ) . Now, for all clients that are located at the same angle, θ , the weakest one is chosen; the rest are discarded from the client set. This is because any transmission that covers the weakest client in a given radial direction, will obviously cover the other clients in that same direction. The clients are then sorted in increasing order of their angular coordinate. Figure 7 illustrates this operation for a simple scenario. The sorted list of clients is $\{8, 2, 7, 4, 5, 1\}$, each associated with an angle θ_i , and a rate r_i (assigned by the LQE). Figure 7 also shows the set of beams covering these clients.

The problem is then to compute an optimal subset from these beams, and the corresponding transmission rates, such that transmission time is minimized while all nodes are covered. The dynamic program sweeps across the clients angularly (starting from θ_0), and optimally covers increasing sizes of conical sectors.




Fig. 7: Radial client distribution, beam B3 covers 4 and 7




Fig. 8: Dynamic programming recursive solution

Denote $C[\theta_i, \theta_j, B]$ as the minimum cost of covering clients within a conical sector bounded by θ_i and θ_j ($\theta_i < \theta_j$), using beams from set B . We also define a function $M(\theta_i, B_k, r)$, that returns the angle of the first uncovered client, when beam B_k has covered a client at θ_i , using a rate r . In Figure 7, if client 8 is covered by B_1 using 11 Mbps, then the next uncovered client is 2, which needs transmissions at a relatively lower rate. Thus M can be expressed as:

$$M(\theta_i, B_k, r) = \psi \quad \text{s.t. } \forall \theta \in (\theta_i, \psi) \quad C'(\theta) \leq (1/r) \\ = \text{null if no such } \psi \text{ exists}$$

where $C'(\theta)$ is the cost of transmission to the client at angle θ . M will return null when there exists no uncovered client at a greater angle. In Figure 7, when B_5 covers node 1, M will return null. Now, using these functions, the cost to cover all nodes from θ_i to θ_j can be recursively computed as follows.

$$C[\theta_i, \theta_j, B] = \begin{cases} \frac{1}{r}, \\ \text{if } \forall B_k: \theta_i \in B_k, \forall r: r \leq \frac{1}{C'(\theta_i)}, M(\theta_i, B_k, r) \text{ is null} \\ \min. \{ \forall B_k: \theta_i \in B_k, \forall r: r \leq \frac{1}{C'(\theta_i)} \\ \{ \frac{1}{r} + C[M(\theta_i, B_k, r), \theta_j, B - B_k] \} \}, \\ \text{otherwise} \end{cases}$$

Figure 8 illustrates the recursive nature of this equation. Suppose the optimal solution assigns a client (located at angle θ) to a beam B_k with a rate r_k . Let the next uncovered client be at an angle ϕ . Then, the minimum cost of covering all clients between θ_i to θ_j is the sum of the minimum cost of covering θ_i to ϕ , i.e. $(1/r_k)$ and that from ϕ to θ_j . Thus, in Figure 8, if the assigned rate for beam B_1 satisfies both nodes 8 and 2, then the total cost of covering clients 8 to 1, is the cost of B_1 added to the cost of covering clients from 7 to 1. The above recursion is solved in a bottom-up manner, and may benefit from simple techniques that prune the subproblems. For example, if a client is not covered by more than one beam (e.g., node 1 in Figure 8), then that beam has to be selected. Moreover, if this client is the weakest among all others in that beam, then the client's rate must be assigned to that beam. We reduce the AP's processing time using such simple techniques. The complexity of the dynamic program, however, remains $O(n^2 m^2)$, where n is the number of clients, and m is the number of overlapping beams.

C. Retransmission Manager

The AP receives feedback from clients at the end of every batch transmission. To ensure minimum delivery ratio across all clients, the AP prepares to retransmit some of the lost packets. Observe that it is possible to choose an optimal subset among the lost packets, such that they satisfy the MinDR constraints, while incurring least retransmission time. Our proposed multicast scheduler can be re-applied with appropriate modifications to select this optimal set of packets. In case processing time at (cheap) APs is a concern, we propose a simpler heuristic to reduce this time. We begin by removing all clients that have satisfied MinDR. The heuristic is based on the observation that the utility of retransmitting a packet, f_i , depends on the set of unsatisfied clients, U_{f_i} , that had not received f_i . We denote the number of members in this set with $N_i = |U_{f_i}|$. Further, the transmission time of the packet, T_{f_i} , is computed by grouping unsatisfied clients that were originally assigned to the same beam. Now, for each group, the weakest unsatisfied client determines the transmission duration to this group. T_{f_i} is computed as the sum of transmission durations over all the groups.

The heuristic then computes a score, $S_{f_i} = \frac{T_{f_i}}{N_i}$, for each packet. The AP orders these scores in increasing order, and retransmits the packet with the least score. Clients that may satisfy MinDR due to this additional retransmission are discarded. The scores for remaining packets are recalculated, and reordered. The least-score packet from this new order is retransmitted. This is repeated as long as there is at least one remaining client. Of course, one round of retransmission may not guarantee that MinDR is met. Multiple retransmissions may be necessary to cope with stricter guarantees and greater channel fluctuations.

V. PERFORMANCE EVALUATION

We implement *BeamCast* in Qualnet 4.0 [18], and compare its performance with a variant of omnidirectional 802.11. This variant – called *802.11 with Feedback* – assimilates periodic client feedbacks, and estimates the bottleneck rate using the same mechanism as *BeamCast*. Reliability requirements are considered while performing this operation. However, *802.11 with Feedback* does not retransmit packets when it identifies transmission losses; it only responds through rate control. The main simulation parameters are presented in Table II. Clients are distributed randomly around the APs. Unless specified, MinDR is 90%. We evaluate *BeamCast* using metrics of multicast throughput, minimum and average delivery ratios, and fairness. We report the effects on different topologies, client density, co-channel interference, and a variety of wireless fading models.

TABLE II: Simulation Parameters

Tx Range of AP	250m
Client Base	10 to 100
Path Loss Model	two-ray
Shadowing Model	constant
Fading Model	None, Rayleigh, Rician
Antenna Beamwidth	45°, 60°, 90°
Rate Increase with Beamforming	3, 4

Figure 9(a) compares the multicast throughput with *BeamCast* and 802.11 across 20 different topologies, each with 50 nodes. Results show that *BeamCast* consistently outperforms 802.11. Of course, the improvement varies across the topologies because in some cases, the rate difference between the weak and strong clients is small. Hence, the benefits of servicing the weak clients separately is lower. While this experiment was performed under Rayleigh fading, Figure 9(b) shows results of the same experiment with Rician fading (Rician Factor=2). The performance of both *BeamCast* and 802.11 improve in several topologies, however, the performance gap remains consistent. From additional results (not reported in the interest of space), we observed that *BeamCast*'s improvement decreases with lower channel fading. Real channel conditions demonstrate significant fluctuations and fades over time, hence, we argue that *BeamCast* is a practical, deployable solution. When all clients

experience same channel condition, *BeamCast* optimally chooses a single omnidirectional transmission, ensuring that it never performs worse than 802.11.

To understand the impact of hidden terminals, we placed 4 interferers on the periphery of a circle with the AP at its center, and a radius of 500m. Each interferer transmits packets every 5ms, resulting in collisions at the multicast clients. Figure 9(c) and 10(a) present throughput results from the same scenarios described earlier. Evidently, throughput decreases due to numerous collisions, especially at the weaker clients located relatively closer to the interferers. However, for both the Rayleigh and Rician fading, *BeamCast* surpasses 802.11 for all topologies. More interestingly, *BeamCast* copes with the collisions by invoking effective retransmission strategies. Recall that the minimum delivery ratio requirement was specified as 90%. Figure 10(b) shows the minimum delivery ratio achieved by both these schemes. While *BeamCast* meets the minimum requirements in most of the topologies, 802.11 is found to fail often. While 802.11 may be augmented with a retransmission scheme, observe that its throughput will degrade proportionally. The average delivery ratio is presented in Figure 10(c). *BeamCast* surpasses the 90% threshold in all the scenarios, while 802.11 fails again in 25% of the cases. Offering deterministic guarantees on delivery ratio may be difficult. We believe *BeamCast* reasonably trades off throughput for reliability. If the need for reliability is critical, *BeamCast* can be made to perform multiple rounds of retransmissions. We plan to investigate such extensions in our future work.

Figure 11 shows the impact of varying node density on multicast throughput, for a given topology. Performance degrades with increasing client base because the number of weaker clients increases. Hence, 802.11 reduces their transmission rates and *BeamCast* has to invest more transmissions. The curve in Figure 11 increases in one occasion with an increase in number of clients. This is an infrequent case resulting from fortunate channel quality improvements at the weaker clients. In general, throughput degrades with increasing weaker clients, and as expected, the performance gap between *BeamCast* and 802.11 increases. We observed this trend for all topologies, across wide variety of fading and interferences. Again, when the number of clients are few, we observed few cases when *BeamCast* was marginally better than 802.11. This is because, with few clients, the variance of the channel qualities is typically lower, leaving less room for improvement through beamforming.

We used Jain's Fairness Index to compare *BeamCast* with 802.11. Both schemes achieve comparable fairness. Table III reports results from an example topology of 50 nodes (other topologies exhibit similar trends).

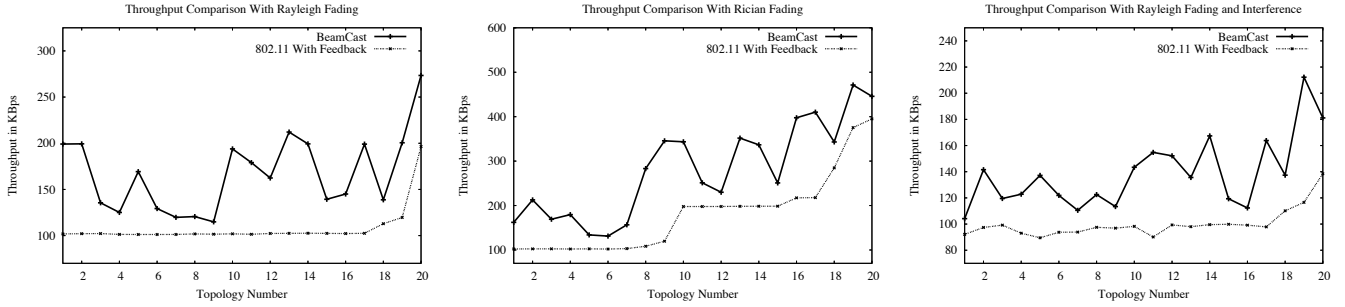


Fig. 9: Throughput with (a) Rayleigh fading, (b) Rician fading, and (c) Rayleigh fading and interference.




Fig. 10: (a) Throughput with Rician fading and interference, (b) MinDR for different topologies (Rayleigh fading), (c) AvgDR for different topologies (Rayleigh fading).




Fig. 11: Performance with increasing number of nodes.

TABLE III: Jain's Fairness Index in Different Scenario

Fading Type	BeamCast	802.11 with Feedback
Rayleigh w/o Interference	0.9997	0.99942
Rician w/o Interference	0.99979	0.99983
Rayleigh with Interference	0.99957	0.99923
Rician with Interference	0.99984	0.99983

Understanding the impact of beamwidth and rate gain is of interest. In the interest of space, we condense the average performance from 20 topologies into Figure 12. Figure 12(left) shows the normalized throughput (over 802.11) for Rician Fading, with different rate gains and increasing beamwidths. Higher beamwidths offer moderate improvements because of the possibility to cover more (weak) clients with a single transmission. The benefit is expectedly more for four times rate gain. Figure 12(right)




Fig. 12: Throughput for varying beamwidth.

shows average multicast throughput for the same 20 topologies, but with Rayleigh Fading. The improvements are pronounced. As explained earlier, higher channel fluctuations offer greater opportunity through multiple beamformed transmissions. Realistic channel conditions are proven to be time varying, making BeamCast practical for real WLAN networks.

VI. DISCUSSION AND FUTURE WORK

BeamCast responds to a packet loss by retransmitting it at a lower data rate. This may be suitable when fading is the cause of the loss. However, if collisions lead to losses, reducing the transmission rate is wasteful. Link layer loss discrimination is an open research problem [19], [20], and solutions to it will benefit BeamCast.

BeamCast has not been optimized to reduce packet collisions due to interference from nearby APs. Observe that it might be possible to coordinate APs such that

their omnidirectional transmissions occur serially, but their beamformed transmissions occur in parallel. This will improve spatial reuse among multicasting APs, while reducing the probability of collisions. Optimizing multicast throughput in multi-AP architectures is a topic of our ongoing work, especially in enterprise environments.

VII. RELATED WORKS

Multicast has been well studied at the network layer [21]. Only recently, there has been increased research attention towards challenges in link layer multicast. Chaporkar *et. al* [22] proposed algorithms for throughput optimality under constraints of network stability. While the ideas are useful, their propositions to use busy tones may not be practical in the context of 802.11 systems. Park *et. al* [23] propose a rate adaptation scheme that improves throughput by utilizing periodic SNR feedbacks from clients. The protocol is similar to the *802.11 with Feedback* scheme that we use to benchmark BeamCast. Chen *et. al* [6] use unary channel feedbacks (UCF) and unary negative feedback (UNF) to estimate channel quality information. The proposed ideas can be well integrated into the Link Quality Estimator in BeamCast. In [24], authors show the possibility to optimize multicast in a multi-AP environment. Won *et. al* [5] design a multicast scheduler that achieves proportional fairness under dynamic channel conditions in cellular data networks. In another work [25], authors devise a reliable multicast protocol through multiple CTS and ACK transmissions. Though reliable, per-packet control overhead can become excessive with large client bases. In a parallel thread of research, the opportunities of beamforming antennas have been well studied [26], [27]. Jaikao *et. al* [28] have investigated the benefits of beamforming in ad hoc network multicasting. To the best of our knowledge, BeamCast is the first attempt to exploit beamforming capabilities for wireless, link layer multicast.

VIII. CONCLUSIONS

This paper identifies the opportunity to exploit beamforming antennas for wireless link layer multicast. The main idea is to execute multiple high data rate transmissions using a combination of omnidirectional and beamformed antenna modes. Such a strategy can outperform a single omnidirectional multicast at the bottleneck data rate. Through periodic link estimation, optimal beam-rate scheduling, and judicious retransmissions, we demonstrate consistent performance improvements. We believe *BeamCast* could be an early step to meet multicasting demands of next generation wireless networks.

ACKNOWLEDGMENT

We would like to thank Samir Das, Vincent Conitzer, and Kamesh Munagala for their insightful comments on the paper. We are grateful to NSF for the CAREER CNS-0747206 grant, and to Cisco Research, for partially funding this work.

REFERENCES

- [1] "Mobitv," Mobile television and radio service provider.
- [2] R. Anderson, R. Anderson, O. Chung, K.M. Davis, C. Prince, V. Razmov, and B. Simon, "Classroom Presenter: a classroom interaction system for active and collaborative learning," in *WIPTE*, 2006.
- [3] "Smarthome," <http://www.smarthome.com/index.aspx>.
- [4] P. Stuedi, O. Chinellato, and G. Alonso, "Connectivity in the presence of shadowing in 802.11 ad hoc networks," in *IEEE WCNC*, 2005.
- [5] H. Won, H. Cai, D. Eun, K. Guo, A. Netrevali, I. Rhee, and K. Sabnani, "Multicast scheduling in cellular data networks," in *IEEE Infocom*, 2007.
- [6] A. Chen, G. Chandrasekaran, D. Lee, and P. Sinha, "HIMAC: High throughput mac layer multicasting in wireless networks," in *IEEE MASS*, 2006.
- [7] "Phocus array system," <http://www.fidelity-comtech.com>.
- [8] T. Ohira and K. Gyoda, "Electronically steerable passive array radiator ESPAR antennas for low-cost adaptive beamforming," in *IEEE International Conference on Phased Array Systems*, 2000.
- [9] C. A. Balanis, "Smart antennas for future reconfigurable wireless communication networks," Tech. Rep., Arizona State University, April 2000.
- [10] "Paratek microwave inc.," <http://www.paratek.com/rfid.htm>.
- [11] V. Navda, A. P. Subramanian, K. Dhanasekaran, A. Timm-Giel, and S. Das, "MobiSteer: Using steerable beam directional antenna for vehicular network access," in *ACM MobiSys*, 2007.
- [12] "Soekris engineering," <http://www.soekris.com>.
- [13] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris, "Link-level measurements from an 802.11b mesh network," in *ACM SIGCOMM*, 2004.
- [14] H. Lundgren, E. Nordstrom, and C. Tschudin, "The gray zone problem in ieee 802.11b based ad hoc networks," *ACM SIGMOBILE Mobile Computing and Communications Review*, vol. 6, no. 3, pp. 104–105, 2002.
- [15] S. Biswas and R. Morris, "ExOR: Opportunistic multi-hop routing for wireless networks," in *ACM SIGCOMM*, 2005.
- [16] M. Borgo, A. Zanella, P. Bisaglia, and S. Merlin, "Analysis of the hidden terminal effect in multi-rate ieee 802.11b networks," in *International Symposium on Wireless Personal Multimedia Communications, (WPMC)*, 2004.
- [17] J. G. Proakis, "Digital communications, 4th ed.," *New York, NY: McGraw Hill*, 2000.
- [18] "Qualnet," <http://www.scalable-networks.com/>.
- [19] S. Rayanchu, A. Mishra, D. Agrawal, S. Saha, and S. Banerjee, "Diagnosing wireless packet losses in 802.11: Separating collision from weak signal," in *IEEE Infocom*, 2008.
- [20] Jaehyuk Choi, Jongkeun Na, Kihong Park, and Chong kwon Kim, "Adaptive optimization of rate adaptation algorithms in multi-rate w lans," in *IEEE ICNP*, 2007.
- [21] P. Paul and S. Raghavan, "Survey of multicast routing algorithms and protocols," 2002.
- [22] P. Chaporkar, A. Bhat, and S. Sarkar, "An adaptive strategy for maximizing throughput in mac layer wireless multicast," in *ACM MobiHoc*, 2004.
- [23] Y. Park, Y. Seok, N. Choi, Y. Choi, and J.-M. Bonnin, "Rate-adaptive multimedia multicasting over ieee 802.11 wireless lans," in *IEEE CCNC*, 2006.
- [24] Ai Chen, Dongwook Lee, and Prasun Sinha, "Optimizing multicast performance in large scale w lans," in *IEEE ICDCS*, 2007.
- [25] S. Jain and S. R. Das, "Mac layer multicast in wireless multihop networks," in *IEEE COMSWARE*, 2006.
- [26] R. R. Choudhury, X. Yang, N. H. Vaidya, and R. Ramanathan, "Using directional antennas for medium access control in ad hoc networks," in *ACM Mobicom*, 2002.
- [27] R. R. Choudhury and N. H. Vaidya, "Deafness: A mac problem in ad hoc networks when using directional antennas," in *IEEE ICNP*, 2004.
- [28] C. Jaikao and C. C. Shen, "Multicast communication in ad hoc networks with directional antennas," in *IEEE ICCCN*, 2003.