7-2009

Poster abstract: BeamCast: Harnessing beamforming capabilities for link layer multicast

Souvik SEN
Duke University

Rahul GHOSH
Duke University

Jie XIONG
Singapore Management University, jxiong@smu.edu.sg

Romit Roy CHOUDHURY
Duke University

Citation

Souvik SEN; GHOSH, Rahul; XIONG, Jie; and CHOUDHURY, Romit Roy. Poster abstract: BeamCast: Harnessing beamforming capabilities for link layer multicast. (2009). *ACM SIGMOBILE Mobile Computing and Communications Review*. 13, (3), 34-37.
Available at: https://ink.library.smu.edu.sg/sis_research/2809

Poster Abstract: BeamCast: Harnessing Beamforming Capabilities for Link Layer Multicast

Souvik Sen *souvik.sen@duke.edu*
Rahul Ghosh *rahul.ghosh@duke.edu*
Jie Xiong *j.x@duke.edu*
Romit Roy Choudhury *romit.rc@duke.edu*

ECE Department, Duke University, Durham, NC 27708, USA

Wireless multicast is an important service primitive for emerging applications such as live video, streaming audio and content telecasts. Transmission rate in such a link layer multicast is bottlenecked by the rate of the weakest client, leading to channel under-utilization. Attempts to increase the data rate results in lower reliability (due to higher bit error rate) and higher unfairness. This poster utilizes smart beamforming antennas to improve multicast performance in wireless LANs. The main idea is to satisfy majority of the (strong) 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, 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 wide range of client topologies and time-varying channel conditions.

I. Introduction

Applications like MobiTV [1], electronic classrooms [2], and WiFi telecasts in smart homes [3] are demanding link layer support for group communication. An ideal solution 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. The multicast throughput can severely [5, 6] suffer 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. 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. Increasing transmission rates does not resolve the challenges, since some (weak) clients will fail to receive transmissions at these data rates. In this context, smart antennas may offer new opportunities to augment the state of the art in link layer multicast. 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; testbed measurements show that – (i) 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. This paper presents *BeamCast*, an antenna-aware protocol that

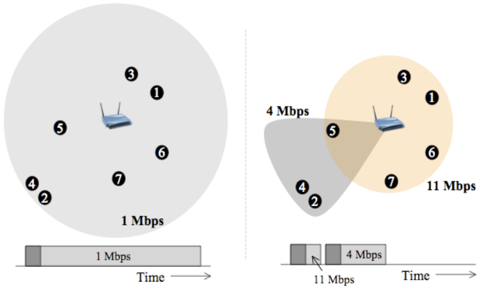

Figure 1: Comparing a single 1 Mbps omnidirectional transmission against multiple high data rate transmissions.

maximizes multicast throughput under specified reliability requirements. Performance results show consistent improvements over omnidirectional schemes, especially when the channel quality varies over time.

II. System Setting

We consider IEEE 802.11 based WLANs. Clients are scattered around the AP, and remain stationary in the time scale of packets. Such an environment is characterized with multipath and shadowing effects, resulting in wireless blind spots (particularly in indoor environments). Each access point is equipped with a smart beamforming antenna, while all the clients have simple, omnidirectional antennas. Beamforming antennas can regulate the radiation and reception patterns such that SINR can be maximized for a given interference environment. 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 $C = Wlog_2(1 + SINR)$, where C is the capacity and W is the bandwidth in use. Although the improvement in data rate is logarithmic, commercial antennas [7] offer more than $16dB$ mainlobe gains [8], leading to 4 times increase in data rates. We show that such rate improvements are sufficient for our case.

III. Motivation and Problem Formulation

This section reports measurement results to validate that real WLANs are typically characterized with a few spatially clustered weak clients. Servicing these few weak clients will require a few beamformed transmissions, suggesting the potential of performance improvements.

III.A. Measurements

We used Soekris boards and laptops, running MadWiFi drivers on 802.11b Atheros interfaces, to measure channel quality in a multicast setting. The AP was made to transmit broadcast packets at different data rates; clients measured the delivery ratio. Table 1 summarizes results from 4 sample 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. Evidently, topologies are characterized with few weak clients. Moreover, weak clients were frequently co-located in shadowed regions and blind spots in our measurements. Figure 2 shows a few identified spots. Additional measurements (not reported here in the interest of space) are reasonable evidence that indoor WiFi environments are often characterized with spatially-clustered, weak clients [4, 9].

Table 1: 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%

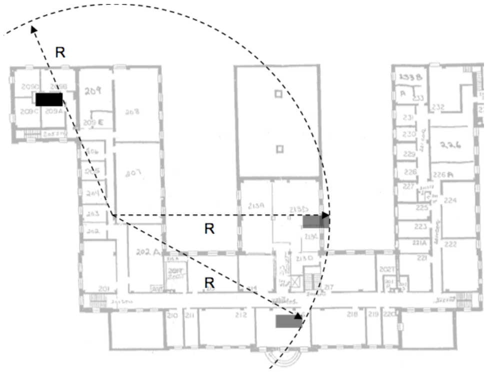

Figure 2: Shaded rectangles indicate blind spots for WLAN environments. The AP is positioned in the center of the circle.

III.B. Problem Formulation

We ask, given a group of clients and their individual data rates from the AP, what should be the optimum transmission strategy that maximizes multicast throughput under a required delivery ratio. Clients, who do not receive the packet through omnidirectional communication, get serviced through beamformed transmissions. So, determining the optimal assignment of clients (into omnidirectional and directional beams) is the main problem of interest. This 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 the same set of clients with different sets of (overlapping) beams. 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 3 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. Thus, beams may not require to transmit at the rate of the weakest client peresent on it. Choosing the optimal beam-set, and assigning corresponding rates to each of these beams, is non-trivial. This poster develops a multicast protocol that optimally partitions clients into omnidirectional and directional beams, and accomplishes transmissions at optimal data rates. The objective is to maximize multicast throughput while meeting a specified delivery ratio.

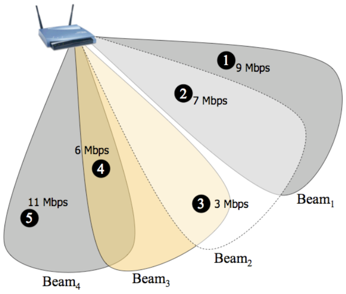

Figure 3: Problem of choosing optimal beams and rates among spatially overlapping beams, such that multicast throughput is maximized.

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.

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). Estimated rates are then fed into a dynamic program, which outputs an optimal set of $< beam_i, rate_i >$ tuples. Using these tuples, the scheduler iterates over a set of strategies which consist of atmost one omnidirectional transmission and corresponding beamformed transmissions. The optimal strategy 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.

V. Performance Evaluation

We implement *BeamCast* in Qualnet 4.0, 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. Unless specified, MinDR is 90%. We evaluate BeamCast using metrics of multicast throughput, minimum and average delivery ratios, and fairness. Due to space constraints, we present only throughput and delivery ratio results. For both Rayleigh and Rician fading, BeamCast surpasses 802.11 for all topologies. Throughput performance under Rayleigh fading has been shown in Figure 4. Figure 5 shows the minimum delivery ratio achieved under Rayleigh fading conditions and in presence of hidden terminals. While BeamCast meets the minimum requirements in most of the topologies, 802.11 is found to fail often. BeamCast reasonably tradeoff throughput for reliability. If reliability is critical, BeamCast can be made to perform multiple rounds of retransmissions. Figure 6 shows, with varying node density, performance degrades with increasing client base because the number of weaker clients increases. However, with increasing weaker clients, the performance gap between BeamCast and 802.11 increases. We observed this trend for all topologies, across wide variety of fading and interferences. Figure 4, 5, 6 present the results for $45°$ beamwidth and 4 times rate gain. Impact of varying beamwidth and rate gain is shown in Figure 7. Higher beamwidths offer moderate improvements because of the possibility to cover more (weak) clients with a single transmission. The benefit is expectedly more when the same beamwidth

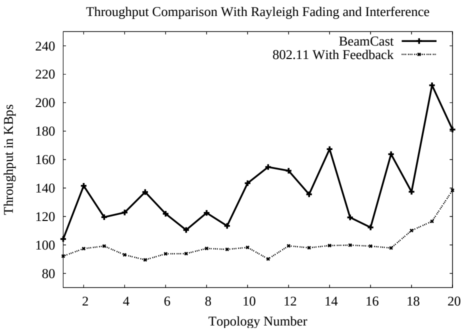

Figure 4: Multicast Throughput with Rayleigh Fading and interference

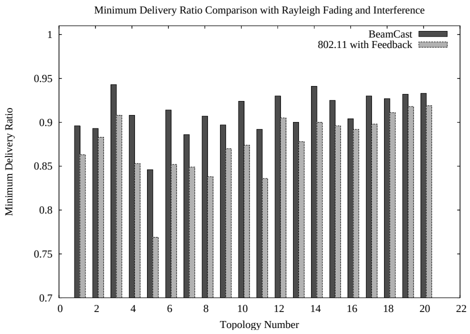

Figure 5: Minimum DR for a range of topologies in Rayleigh Fading

can support a higher rate.

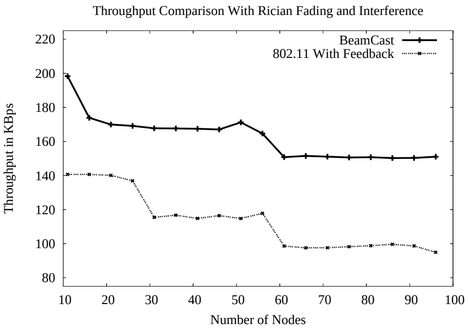

Figure 6: Performance with number of nodes for a given topology.

VI. Discussion and Future Work

We assumed that switching between different beams incurs negligible latency. We plan to account for this delay while designing the optimal schedule. BeamCast responds to a packet loss by retransmitting it at a lower data rate. Link layer loss discrimination is an open research problem [10], and solutions to it will benefit BeamCast. Optimizing multicast throughput in face of emerging PHY layer capabilities is a topic of our ongoing work.

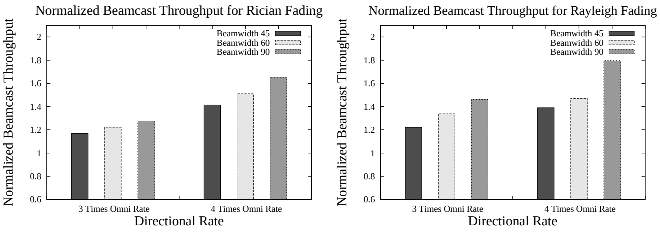

Figure 7: Normalized throughput for Rician and Rayleigh fading

References

[1] “Mobitv,” Mobile television and radio service provider.

[2] R. Anderson et. al, “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 et. al, “Connectivity in the presence of shadowing in 802.11 ad hoc networks,” in *IEEE WCNC*, 2005.

[5] Ai Chen, Gayathri Chandrasekaran, Dongwook Lee, and Prasun Sinha, “HIMAC: High throughput mac layer multicasting in wireless networks,” in *IEEE MASS*, 2006.

[6] H. Won, H. Cai, D. Eun, K. Guo, A. Netrevali, I. Rhee, and K. Sabnani, “Multicast scheduling in celluar data networks,” in *IEEE Infocom*, 2007.

[7] “Phocus array system,” http://www.fidelity-comtech.com.

[8] 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.

[9] D. Aguayo et. al, “Link-level measurements from an 802.11b mesh network,” in *ACM SIGCOMM*, 2004.

[10] 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.