

Experimental Evaluation of a Scalable WiFi Multicast Scheme in the ORBIT Testbed

Yigal Bejerano*, Jaime Ferragut[‡], Katherine Guo*,
Varun Gupta[†], Craig Gutterman[†], Thyaga Nandagopal⁺, Gil Zussman[†]

* Bell Labs, Alcatel-Lucent, Murray Hill, NJ, USA.

[‡] Mobile Networks Department - CTTC, Barcelona, Spain.

[†] Electrical Engineering, Columbia University, New York, NY, USA.

⁺ National Science Foundation, Arlington, VA, USA.

Abstract—IEEE 802.11-based wireless local area networks, referred to as WiFi, have been globally deployed and the vast majority of the mobile devices are currently WiFi-enabled. While WiFi has been proposed for multimedia content distribution, its lack of adequate support for multicast services hinders its ability to provide multimedia content distribution to a large number of devices. In earlier work, we proposed a dynamic scheme called *AMuSe* that selects a subset of the multicast receivers as feedback nodes. The feedback nodes periodically send information about channel quality to the multicast sender and the sender in turn can optimize multicast service quality, e.g., by dynamically adjusting transmission bit-rate. In this paper, we discuss several experimental results for the performance evaluation of *AMuSe*. Our experiments use more than 250 nodes placed in a grid topology in the ORBIT testbed. We consider different experimental scenarios: with and without the presence of external noise. Our focus is on studying the performance of WiFi nodes in WiFi multicast and establishing the conditions that make *AMuSe* an attractive scheme for feedback in WiFi multicast.

I. INTRODUCTION

Recent years have witnessed a rapid growth of mobile devices such as smart-phones and tablets equipped with wireless local area network (WLAN) interfaces that use WiFi standards [1]. While these devices allow users to access the Internet anywhere anytime, it is not straightforward to serve rich multimedia content, such as video streams, when users are clustered in crowded areas, due to a combination of high bandwidth requirements and a shortage of wireless spectrum. The inability to serve this growing demand for multimedia content using limited resources in crowded areas has prompted several solutions by both industry and academia. Many of these solutions [2]–[4] are typically based on dense deployment of Access Points (APs) for providing dedicated content delivery to each user. Such solutions, besides requiring considerable capital and operational expenditure, may not meet user expectations, due to extensive interference between adjacent APs.

A. Limitations of Existing Solutions

Current state of the art solutions use IEEE 802.11 (WiFi), leveraging either unicast or multicast data delivery. Commercial solutions [3], [4] rely on streaming content to individual

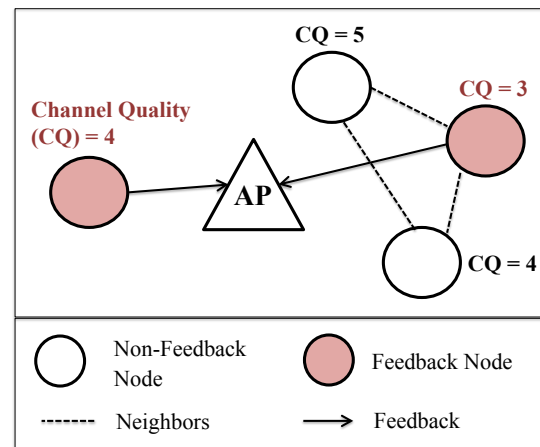


Fig. 1. Feedback node selection by *AMuSe* scheme. A node with the poorest channel quality in every neighborhood is selected as a feedback node. Each feedback node periodically sends updates about its channel quality to the AP.

users. With standards such as 802.11ac promising speeds up to 800 Mbps per user using multi-user MIMO, it is theoretically possible to serve video streams to hundreds of users. However, recent studies [5], [6] throw cold water on this promise. Large number of neighboring APs lead to hidden terminal problems and this, coupled with increased interference sensitivity stemming from channel bonding, makes the entire solution interference limited. Extrapolating from studies on 802.11n [5], [6], 802.11ac-based unicast for streaming a single video to multiple clients may not be able to support more than a hundred users, assuming all of them have 802.11ac capable devices.

Standard WiFi broadcast/multicast frames are transmitted at a fixed and low bit-rate without any feedback. This raises several known reliability and efficiency issues. High packet loss due to interference and hidden node problem can significantly degrade service quality, while transmission at low bit-rates leads to low network utilization.

B. The *AMuSe* Scheme

In [7], we proposed a scalable and adaptive interference mitigation scheme for WiFi multicast referred to as *Adaptive*

multicast Services (AMuSe). Our solution is based on accurate receiver feedback and incurs a small control overhead. It enables the APs to transmit multicast traffic at the highest possible bit-rate, while ensuring high *Packet Delivery Ratio* (PDR) for a large fraction of the nodes. The design of *AMuSe* is based on the observation that adjacent nodes experience similar channel quality and interference patterns [8]. *AMuSe* dynamically divides the nodes in a network into a few clusters based on adjacency of nodes and maximum cluster size (D m). In each cluster, one node is selected as the *Feedback (FB) node* and it updates the AP about its service quality, e.g., channel quality (e.g., see Fig. 1). The AP, in response, may adjust the bit-transmission rate, retransmit lost packets, etc. to mitigate losses. *AMuSe* can be implemented as a light-weight application on any WiFi enabled device and consequently, does not require changes in existing standards. *AMuSe* can yield different subsets of feedback nodes with different measures of channel quality. Thus, choosing the right channel quality measure is of significant importance.

C. Our Contribution

In [7], the performance of *AMuSe* was evaluated in the ORBIT testbed [9]. This paper outlines several additional experiments conducted in the ORBIT testbed, using more than 250 WiFi nodes, to evaluate *AMuSe*. We discuss the observations about the performance of WiFi nodes in a multicast setting. Our focus is on easily measured channel quality metrics like packet delivery ratio, signal strength, etc. We study variation of these metrics in different scenarios: with and without external noise. We discuss the interplay between these metrics and observe that packet level statistics are the most accurate representatives of channel quality. We study the similarity of channel quality in clusters of receivers of different sizes. This helps us establish the tradeoffs involved in selecting fewer feedback nodes through *AMuSe* to ensure a lower control overhead vs. selecting more feedback nodes to ensure more accurate feedback. These experiments serve as a guide for the selection of optimal channel quality measure and for choosing the appropriate number of feedback nodes for *AMuSe*.

In the rest of the paper, we discuss related work on WiFi multicast in Section II. The ORBIT testbed and our experimental design and settings are explained in Section III. We present the detailed experimental results in Section IV before concluding in Section V.

II. RELATED WORK ON WIRELESS MULTICAST

Nearly all of the solutions to improve multicast services rely either on integrating automatic repeat request (ARQ) mechanisms into the protocol architecture [10]–[14], or adding FEC to the multicast stream [15], [16], or both [17]. Other studies propose rate adaptation mechanisms for improving network utilization [18]. In all cases, a key requirement is having appropriate feedback from receivers regarding the quality of the experienced multicast services. These feedback gathering mechanisms can be classified into four main categories:

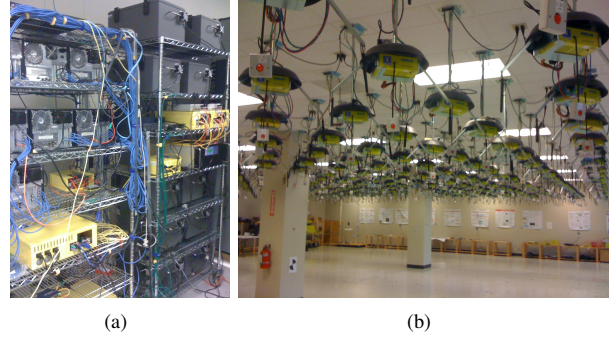


Fig. 2. ORBIT Testbed (a) Sandboxes for development and testing. The wireless interfaces are connected via RF cables to prevent interference with the main grid. (b) Main grid with 400 WiFi enabled nodes arranged in a grid topology.

(i) **Individual Feedback:** These mechanisms require all receivers to send acknowledgements of received packets, e.g., *More Reliable Groupcast* (MRG) [1], [19] from IEEE 802.11 working group. This approach offers reliability but incurs high feedback overhead with large groups.

(ii) **Leader-Based Protocol with Acknowledgements (LBP-ACK):** The LBP-ACK approach [13]–[15], [17], [19], [20] selects a subset of the receivers to provide feedback. This approach naturally provides more scalability than the Individual Feedback approach.

(iii) **Pseudo-Broadcast:** The Pseudo-Broadcast approach [10], [11], [20], [21], converts the multicast feed to a unicast flow and sends it to one leader, typically, the receiver with the weakest channel. The leader acknowledges the reception while other receivers receive packets by listening to the channel in promiscuous mode. Feedback from a single leader in LBP-ACK and Pseudo-Broadcast can lead to poor service reliability.

(iv) **Leader-Based Protocol with Negative Acknowledgements (LBP-NACK):** The LBP-NACK approach [12], [18], [22] improves Pseudo-Broadcast by allowing the non-leaders to send NACKs for lost packets. The LBP-NACK approach not only requires changes to the standard, but also suffers from lack of reliability as a non-leader cannot reply with a NACK if it cannot identify a corrupted packet.

While there have been a few extensive experimental studies of WiFi [23]–[26], WiFi multicast has not been experimentally evaluated on a large scale.

III. OBJECTIVES AND TESTBED ENVIRONMENT

A. Objectives

Our objectives are the following:

- (i) Understanding variation of channel quality metrics under different conditions such as: varying bit-rates at the AP, presence of external noise, etc. We want to identify the best metric for *AMuSe*.
- (ii) *AMuSe* is based on the property that adjacent nodes experience similar channel quality. We study this property by grouping adjacent nodes in clusters of different sizes and observing the similarity of different channel quality

metrics in these clusters. Larger cluster size implies fewer feedback nodes in *AMuSe* and consequently, lower control overhead. However, fewer feedback nodes might lead to poorer representative feedback. Our aim is to understand this trade-off.

B. ORBIT Testbed

The ORBIT testbed [9] is a dynamically configurable grid of 20×20 (400 overall) nodes equipped 802.11 Network Interface Cards (NIC). The separation between adjacent nodes is 1 meter. We label each node in the grid according to its location (x, y) with x and y indicating its column and row location respectively. The testbed includes a single noise generator attached to 4 noise antennas at the 4 corners of the grid, namely (1,1), (1,20), (20,1), and (20,20). Note that these noise generators are located right next to receiver nodes at these grid positions. Each noise antenna is attached to an attenuator whose attenuation can be independently controlled, permitting the emulation of a richer topology. In addition to the grid, the testbed contains smaller sandboxes for testing purposes which emulate nodes on the grid. Each sandbox has 2-4 nodes which are connected through a cable that emulates the RF channel to avoid interference with the main grid. The sandbox and the grid are shown in Fig. 2(a) and Fig. 2(b) respectively.

C. Experimental Settings

In our experiments, the node at the corner (1, 1) serves as a single multicast AP, configured in master mode. All the 802.11 radios are configured in channel 40 of 802.11a. We observed that channel 40 at the 5Ghz band suffers from less external interferences on the ORBIT grid than the channel at 2.4Ghz band regardless of the time frame of the experiments. The AP sends multicast UDP flows with each UDP packet of payload size 1400 bytes. The farthest node from the AP in the testbed is roughly 28 meters away which is less than a quarter of the transmission range of an AP. Thus, transmission power is set to $1mW = 0dBm$ to compensate for the relatively small size of the testbed. The other nodes are configured in managed mode and act as receivers. In order to avoid performance artifacts stemming from a mismatch of WiFi hardware and software, we only choose nodes equipped with *Atheros 5212/5213* wireless cards with *ath5k* wireless driver for our experiments. The experimental parameters are summarized in Table I.

Every node i keeps track of the parameters listed in Table II. The AP records the broadcast and multicast transmission bit-rate TX_{AP} . The nodes keep Link Quality (LQ_i), Received Signal Strength Indication ($RSSI_i$), a vector of packet sequence numbers for packets received by the node P_i^{vec} , node location in its column position in the grid x_i , and row position y_i . These parameters are transferred to a local machine for off-line processing after each experiment. To enable the measurement of packet statistics, we add packet sequence numbers in the UDP payload. The *Packet Delivery Ratio* (PDR) value of each node i is calculated from its P_i^{vec} vector as a percentage of correctly received packets. The LQ and RSSI are read

TABLE I
EXPERIMENTAL PARAMETERS

Parameter	Setting
Mode	802.11a
Channel	40
Transmit Power	0 dBm
Wireless Driver	ath5k
Wireless Cards	Atheros 5212/5213
Transport Protocol	UDP
UDP payload size	1400
AP Location	(1,1)

TABLE II
EVALUATION PARAMETERS

Parameter	Definition
LQ_i	Link Quality at node i
$RSSI_i$	RSSI at node i
P_i^{vec}	Vector of the packets received by node i
x_i, y_i	Location of node i
TX_{AP}	Broadcast/Multicast transmission bit-rate at the AP

directly from the card. These three parameters are the only practically measurable *channel quality* metrics on a typical wireless card. We are interested in the behaviour of these channel quality metrics in wireless multicast setting. From our experiments, we observed that LQ can be obtained by adding a fixed threshold value to the RSSI metric. LQ measurements do not have a unit and should not be confused with Signal to Noise Ratio (SNR). For the rest of the paper, we will only focus on the behavior of LQ.

We now describe the two experiment scenarios that are important to evaluate the performance of WiFi multicast.

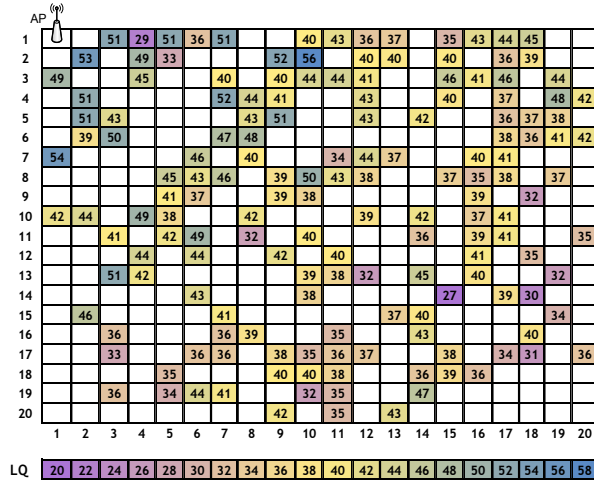
- (i) Different bit-rates: We fix the AP multicast transmission bit-rate, denoted as TX_{AP} , to different values allowed by the card (6, 9, 12, 18, 24, 36, 48, 54 Mbps), for a duration of 10 seconds. We conduct this set of experiments without adding external noise.
- (ii) Varying noise: We fix the AP multicast transmission bit-rate to 12 Mbps and vary the power at the noise generator. All the noise attenuators are set to maximum attenuation of 63 dBm except near node (1,20) which has 0 attenuation. The noise generator is configured to provide *Additive White Gaussian Noise* (AWGN) noise for the entire spectrum of channel 40. Starting with -70dBm (low noise), we vary noise power in steps of 5dBm up to -35dBm (high noise)¹.

The experimental results of the above mentioned scenarios are described in the next section.

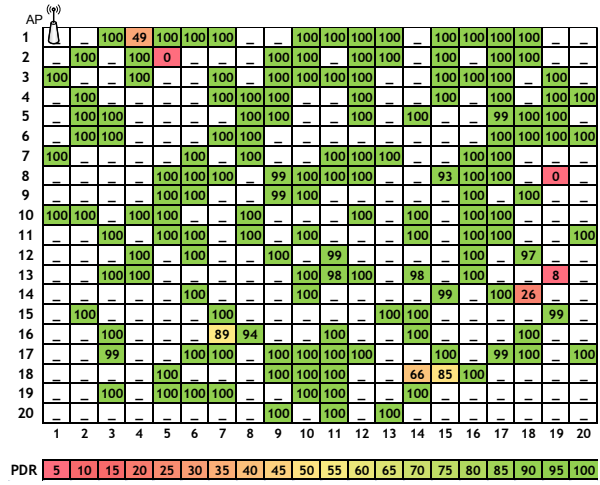
IV. EXPERIMENTAL RESULTS

We first explore the variation of different channel quality metrics under the two experimental scenarios described in the previous section.

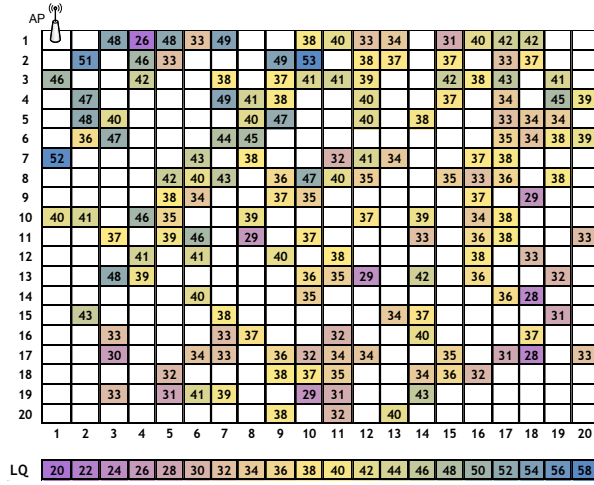
¹The results of noise levels higher than -35dBm are not presented as with higher noise powers, a significant number of nodes disconnect from the AP.



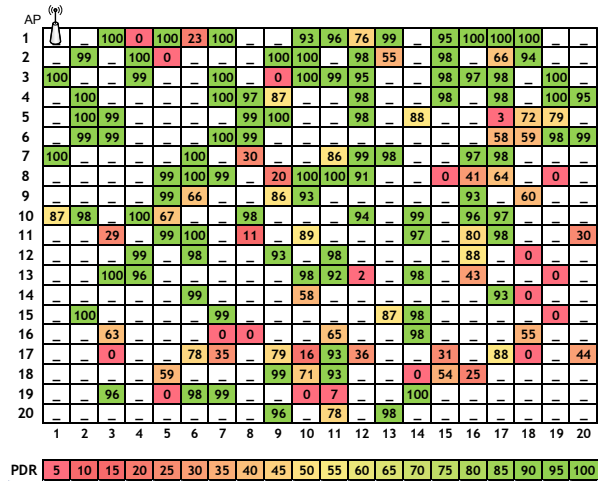
(a) LQ Heatmap, bit-rate = 36 Mbps



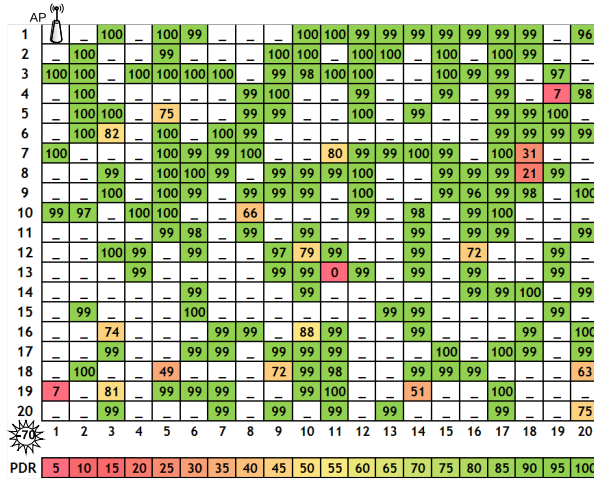
(b) PDR Heatmap, bit-rate = 36 Mbps



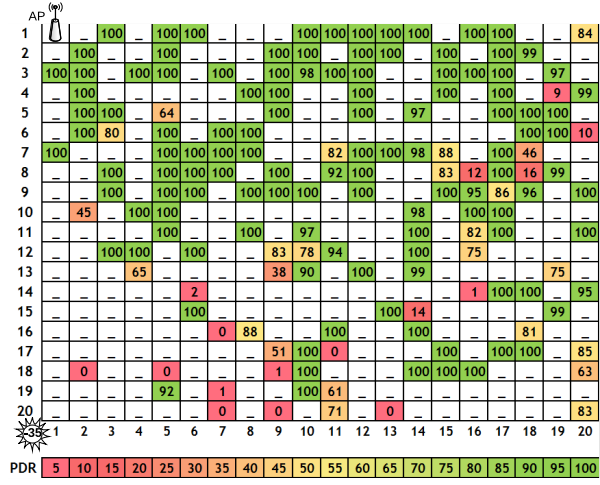
(c) LQ Heatmap, bit-rate = 48 Mbps



(d) PDR Heatmap, bit-rate = 48 Mbps



(e) PDR Heatmap, noise power = -70dBm, bit-rate = 12 Mbps



(f) PDR Heatmap, noise power = -35dBm, bit-rate = 12 Mbps

Fig. 3. PDR and LQ heatmaps at different AP bit-rates and external noise power levels for one experiment run.

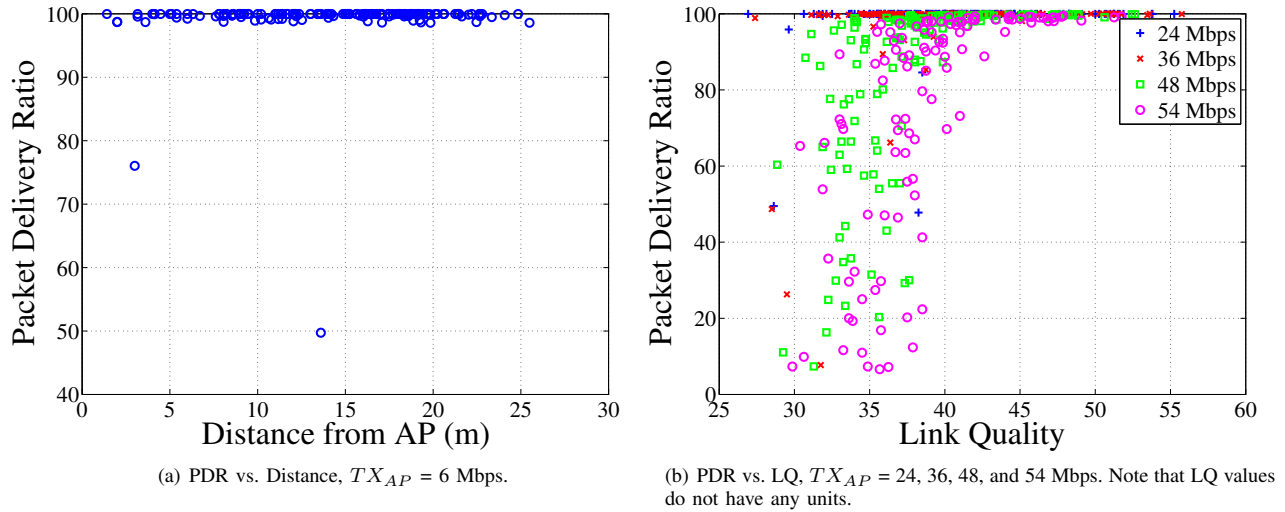


Fig. 4. Experimental results for variation of PDR with distance from AP and PDR with LQ.

A. Variation of PDR and LQ

Varying bit rate:

We observe that changing bit rate at the AP has a marginal effect on the LQ measurements. Fig. 3(a) and Fig. 3(c) show this effect for two values of bit-rate: 36 Mbps and 48 Mbps respectively. However, the PDR shows significant changes with increasing bit-rate as observed by comparing Fig. 3(b) and Fig. 3(d). Changing the bit-rate leads to changes in *Modulation and Coding Scheme (MCS)* of transmitted data which leads to changes in PDR values. High bit-rates also lead to large number of nodes reporting packet losses.

We plot the PDR vs LQ for different bit-rates in Fig. 4(b). The PDR values are close to 100% for almost all nodes for bit-rates up to 24 Mbps. Some degradation of PDR values is observed for bit-rates of 36 Mbps and even higher variance of PDR values are seen for 48 Mbps and above. The correlation between the PDR and LQ is not very strong, suggesting that nodes with the same LQ value may have significantly different PDR. Further, it is easy to see a threshold effect where nodes with LQ below a certain threshold have very low PDRs.

Varying noise power:

After considering experimental scenario (i), we move to scenario (ii) varying power at external noise generator and keeping AP bit-rate constant. We discover that LQ values do not change with increasing external noise power. The LQ reported by the nodes is a driver-specific function of the RSSI. The RSSI values in Atheros cards are a moving window average of signal strength obtained from PLCP (Physical Layer Convergence Protocol) preamble [24]. Higher external noise power results in PLCP decoding failure and RSSI from the packet not considered for the moving window. This is the reason why LQ values do not change with external noise. Thus, we conclude that RSSI and LQ are not effective metrics for measuring channel quality.

We turn our attention to PDR values in the presence of

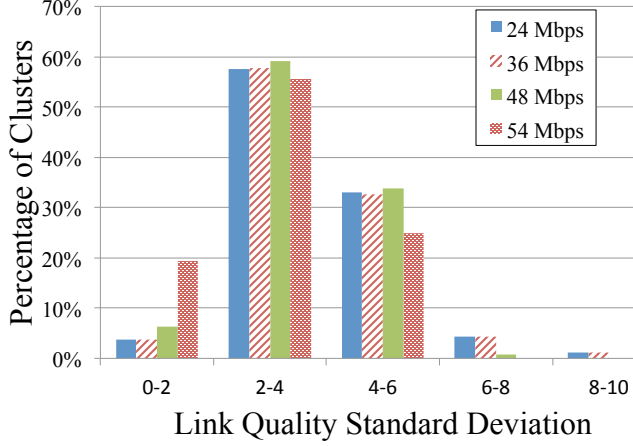
noise. As seen in Fig. 3(e) and Fig. 3(f), increasing noise power has a strong effect on PDR. The effect of noise is more pronounced near the noise antenna where the nodes start disconnecting from the AP. Increasing external noise power leads to nodes below a certain threshold of Link Quality disconnecting from the AP because of sustained loss of beacons from the AP. Since changing the bit-rate in the ath5k driver leads to changes in the MCS of beacon messages as well, this disconnection can be partially mitigated by modifying the driver to transmit beacons at 6 Mbps. Some nodes farther away from the AP also report low PDR values, e.g., node at position (18,4). This effect arises from the extremely static environment of the indoor testbed which results in standing waves and destructive interference at some locations. The effect of destructive interference is also visible without the presence of external noise, e.g., in Fig. 4(a) where we plot PDR vs distance from AP at low bit-rate of 6 Mbps. The nodes which are very close to the AP and report very low PDR are termed as **abnormal nodes** [7].

This variation of PDR with LQ as well as variation of PDR with distance to the AP is consistent with prior work, e.g., [23], [24] and [25].

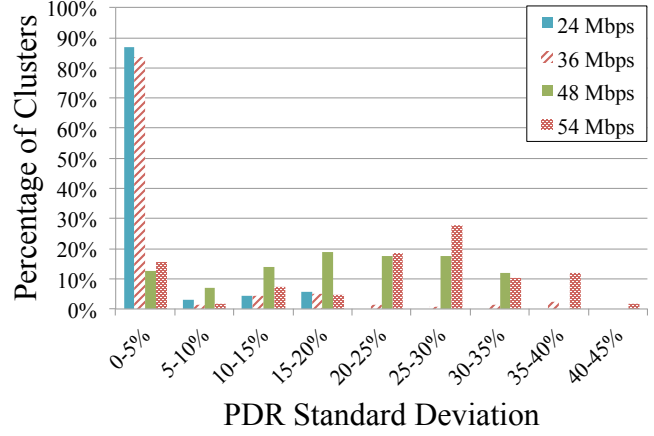
B. Adjacent nodes experience similar LQ and PDR

In [7], we compared the *Standard Deviation (STD)* of LQ and PDR with and without the presence of external noise. Here, we compare the effects of varying the cluster size on the STD of LQ and PDR. We measure the STD without adding external noise in each cluster with radius of 3 and 6 meters on the grid. We expect the STD across clusters to be a good measurement of similarity of PDR and LQ values for clusters in *AMuSe* scheme. Histograms of the distribution of the LQ and PDR STD in are shown in Fig. 5(a)–5(d).


Fig. 5(a) and Fig. 5(c) show that the LQ STD is very similar across all the bit-rates for both cluster sizes. Thus, adjacent nodes experience similar LQ and due to the invariance of LQ




(a) LQ STD: varying TX_{AP} without noise, cluster size = $3m$.



(b) PDR STD: varying TX_{AP} without noise, cluster size = $3m$.



(c) LQ STD: varying TX_{AP} without noise, cluster size = $6m$.



(d) PDR STD: varying TX_{AP} without noise, cluster size = $6m$.

Fig. 5. Experimental results for variation of LQ and PDR in a neighborhood.

with AP bit-rate, the LQ STD does not change. By comparing Fig. 5(a) and Fig. 5(c), we see that the higher percentage of clusters report higher LQ STD, especially in the range of 2-4 and 4-6, for the larger cluster size.

We now consider the distribution of the PDR STD values. Fig. 5(b) shows that with $TX_{AP} \leq 36$ Mbps, only very few clusters show significant deviations ($> 5\%$) in PDR, as most nodes have PDR above 99%. We see significant variability of the PDR at higher bit-rates. Comparing Fig. 5(b) and Fig. 5(d), we see that higher percentage of clusters report higher PDR STD (i.e. $> 5\%$), for larger cluster sizes. Further, from Fig. 5(d), higher bit-rates lead to higher PDR STD values for a significant number of clusters.

A multicast transmission system should ideally operate at a bit-rate where a large fraction of the nodes experience high PDR and the PDR STD is very low to guarantee reliable service to all the receivers. Increasing the cluster size reduces the number of feedback nodes (and consequently the control overhead) but leads to increased STD of channel quality in clusters, particularly the PDR STD at higher bit-rates. The

above observations serve as a good motivation to carefully set the cluster size parameter for *AMuSe*. The ideal cluster size will depend on service requirements of the particular multicast application employing *AMuSe*.

V. CONCLUSION

In this paper, we study the performance of a scalable WiFi multicast scheme referred to as *AMuSe* through large-experiments on the ORBIT testbed. *AMuSe* selects a subset of nodes as feedback nodes based on their channel quality which periodically update the AP about the network conditions. We study the reliability of several channel quality measures reported by WiFi nodes. We establish that even though there are no practical ways to directly measure signal and noise strength values at WiFi nodes, packet delivery statistics are reasonably effective indicators of channel quality. We quantify the potential tradeoffs in selecting fewer feedback nodes through *AMuSe* vs. ensuring reliable representative feedback.

In [7], we discussed a simple rate adaptation mechanism at the AP that utilizes *AMuSe* feedback to ensure reliable

multicast. Our ongoing work aims to study the impact of adjusting *AMuSe* parameters like cluster size, feedback frequency, etc. on the rate adaptation algorithm. Our end goal is to provide fast response to interference without compromising system stability.

VI. ACKNOWLEDGMENTS

The authors would like to express their gratitude to Ivan Seskar from WINLAB (Rutgers University) for his critical support in conducting experiments and several useful technical discussions. This work was supported by (while one of the authors was serving at) the National Science Foundation, and in part by NSF grant CNS-10-54856, NSF CIAN ERC under grant EEC-0812072, by the Spanish Ministry of Economy and Competitiveness under grant TEC2011-29700-C02-01, the Spanish Ministry of Education under grant FPU AP2009-5000, and the Generalitat de Catalunya under grant 2009-SGR-940. Any opinion, findings and conclusions or recommendations expressed here do not necessarily reflect the views of the National Science Foundation.

REFERENCES

- [1] "IEEE draft standard for information technology telecommunications and information exchange between systems local and metropolitan area networks - specific requirements, part 11: Wireless LAN medium access control (MAC) and physical layer (PHY) specifications - amendment: MAC enhancements for robust audio video streaming," IEEE, July 2011.
- [2] Y. Tanigawa, K. Yasukawa, and K. Yamaoka, "Transparent unicast translation to improve quality of multicast over wireless LAN," in *Proc. IEEE CCNC'10*, 2010.
- [3] "Cisco, white-paper, cisco connected stadium Wi-Fi solution," 2011. [Online]. Available: http://www.cisco.com/web/strategy/docs/sports/c78-675064_svcs.pdf
- [4] "Yinzcam," <http://www.yinzcam.com/>.
- [5] K. Pelechrinis, T. Salonidis, H. Lundgren, and N. Vaidya, "Experimental characterization of 802.11n link quality at high rates," in *Proc. ACM WITTECH'10*, 2010.
- [6] N. Hajlaoui and I. Jabri, "On the performance of IEEE 802.11n protocol," in *Proc. ACM WITTECH'12*, 2012.
- [7] Y. Bejerano, J. Ferragut, K. Guo, V. Gupta, C. Gutterman, T. Nandagopal, and G. Zussman, "Scalable WiFi multicast services for very large groups," in *Proc. IEEE ICNP'13*, 2013.
- [8] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris, "Link-level measurements from an 802.11b mesh network," *Proc. ACM SIGCOMM'04*, 2004.
- [9] "ORBIT testbed," <http://orbit-lab.org/>.
- [10] R. Chandra, S. Karanth, T. Moscibroda, V. Navda, J. Padhye, R. Ramjee, and L. Ravindranath, "DirCast: a practical and efficient Wi-Fi multicast system," in *Proc. IEEE ICNP'09*, 2009.
- [11] N. Choi, Y. Seok, T. Kwon, and Y. Choi, "Leader-based multicast service in IEEE 802.11v networks," in *Proc. IEEE CCNC'10*, 2010.
- [12] J. K. Kuri and S. Kumar, "Reliable multicast in multi-access wireless LANs," *ACM/Kluwer Wirel. Netw.*, vol. 7, pp. 359–369, 2001.
- [13] M.-T. Sun, L. Huang, A. Arora, and T.-H. Lai, "Reliable MAC layer multicast in IEEE 802.11 wireless networks," in *Proc. IEEE ICNP'02*, 2002.
- [14] X. Wang, L. Wang, Y. Wang, Y. Zhang, and A. Yamada, "Supporting MAC layer multicast in IEEE 802.11n: Issues and solutions," in *Proc. IEEE WCNC'09*, 2009.
- [15] O. Alay, T. Korakis, Y. Wang, and S. Panwar, "Dynamic rate and FEC adaptation for video multicast in multi-rate wireless networks," *ACM/Springer Mobile Netw. and Appl.*, vol. 15, no. 3, pp. 425–434, 2010.
- [16] H.-T. Chiao, S.-Y. Chang, K.-M. Li, Y.-T. Kuo, and M.-C. Tseng, "WiFi multicast streaming using AL-FEC inside the trains of high-speed rails," in *Proc. IEEE BMSB'12*, 2012.
- [17] M. Wu, S. Makharia, H. Liu, D. Li, and S. Mathur, "IPTV multicast over wireless LAN using merged hybrid ARQ with staggered adaptive FEC," *IEEE Trans. Broadcast.*, vol. 55, no. 2, pp. 363–374, 2009.
- [18] W.-S. Lim, D.-W. Kim, and Y.-J. Suh, "Design of efficient multicast protocol for IEEE 802.11n WLANs and cross-layer optimization for scalable video streaming," *IEEE Trans. Mobile Comput.*, vol. 11, no. 5, pp. 780–792, 2012.
- [19] Z. Feng, G. Wen, C. Yin, and H. Liu, "Video stream groupcast optimization in WLAN," in *Proc. IEEE ITA'10*, 2010.
- [20] S. Sen, N. K. Madabhushi, and S. Banerjee, "Scalable WiFi media delivery through adaptive broadcasts," in *Proc. USENIX NSDI'10*, 2010.
- [21] Y. Park, C. Jo, S. Yun, and H. Kim, "Multi-room IPTV delivery through pseudo-broadcast over IEEE 802.11 links," in *Proc. IEEE VTC'10*, 2010.
- [22] Z. Li and T. Herfet, "HLBP: a hybrid leader based protocol for MAC layer multicast error control in wireless LANs," in *Proc. IEEE GLOBECOM'08*, 2008.
- [23] H. Rahul, F. Edalat, D. Katabi, and C. G. Sodini, "Frequency-aware rate adaptation and MAC protocols," in *Proc. ACM MOBICOM'09*, 2009.
- [24] A. Vlavianos, L. Law, I. Broustis, S. Krishnamurthy, and M. Faloutsos, "Assessing link quality in IEEE 802.11 wireless networks: Which is the right metric?" in *Proc. IEEE PIMRC'08*, 2008.
- [25] S. Kaul, M. Gruteser, and I. Seskar, "Creating wireless multi-hop topologies on space-constrained indoor testbeds through noise injection," in *Proc. IEEE TRIDENTCOM'06*, 2006.
- [26] K. Ramachandran, S. Kaul, S. Mathur, M. Gruteser, and I. Seskar, "Towards large-scale mobile network emulation through spatial switching on a wireless grid," in *Proc. ACM SIGCOMM E-WIND '05*, 2005.