AMuSe: Adaptive Multicast Services to Very Large Groups – Project Overview

Yigal Bejerano*, Varun Gupta†, Craig Gutterman†, Gil Zussman†
∗ Bell Labs, Nokia, Murray Hill, NJ, USA.
† Electrical Engineering, Columbia University, New York, NY, USA.

***Abstract*—WiFi multicast to *very large groups* has gained attention as a solution for multimedia delivery in crowded areas. Yet, most recently proposed approaches do not provide performance guarantees. In this paper, we describe the AMuSe system, whose objective is to enable scalable and adaptive WiFi multicast services. AMuSe includes a light-weight feedback mechanism that allows monitoring channel quality of a large number of users. This feedback allows the system to dynamically optimize the multicast transmission rate at the AP. We implemented AMuSe on the ORBIT testbed and evaluated its performance in large groups with approximately 200 WiFi devices in different scenarios. We show that AMuSe supports high throughput multicast flows to hundreds of receivers while meeting quality requirements and that it outperforms other systems.**

I. Introduction

Recent years have witnessed a rapid growth of mobile devices equipped with an IEEE 802.11 (WiFi) interface, which allow users to access the Internet anywhere and any time. Yet, due to a combination of high bandwidth requirements and a shortage of wireless spectrum, it is challenging to serve rich multimedia content (such as video streams) to users clustered in crowded areas. The growing need to support larger demands for multimedia content using limited resources in dense areas has prompted the design of several solutions by both industry and academia. Many of these solutions [1]–[3] are based on dense deployments of Access Points (APs) in order to provide 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.

With standards such as 802.11ac promising total speeds up to 1.3 Gbps using multi-user MIMO, it is theoretically possible to serve video streams to hundreds of users. However, recent studies [4]–[6] throw cold water on this promise. A large number of neighboring APs leads to hidden terminal problems and this, coupled with increased interference sensitivity due to channel bonding, makes the entire approach highly susceptible to interference. It seems that 802.11ac-based unicast to multiple receivers may not be able to support more than a hundred users, assuming all of them have 802.11ac capable devices.

On the other hand, some commercial products [2] are experimenting with WiFi multicast deployments for crowded environments. However, there remain several challenges to its widespread adoption. A recently published IETF Internet Draft highlights several open technical problems for WiFi multicast [7]. High packet loss due to interference and the

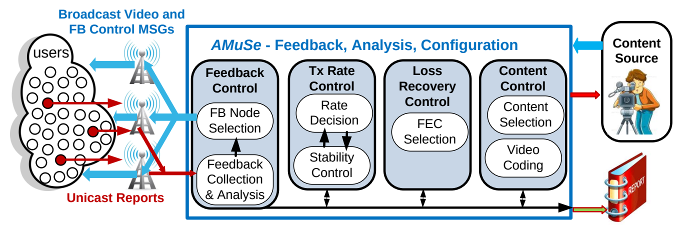

Fig. 1. The AMuSe System Description. The marked users are feedback nodes which periodically send updates about the service quality to AMuSe server.

hidden node problem can significantly degrade service quality. Multicast transmission at low bitrates leads to low network utilization. As described in Section II, there are numerous studies that propose solutions for overcoming these limitations from two angles. One direction of research aims to reduce the overhead of collecting feedback information at the multicast sender. The other aims to improve message reliability based on available feedback information. All the existing schemes, however, suffer from one or more issues including lack of scalability, inability to guarantee high service quality, or compliance with existing standards. Furthermore, *none of the schemes have been tested experimentally at scale.*

A. The AMuSe System

This paper presents an overview of the *AMuSe* (**A**daptive **Mu**lticast **Se**rvices) system. AMuSe provides a wireless multicast based solution for scalable and efficient delivery of multimedia content to a very large number of WiFi nodes in crowded venues (e.g., sport arenas, lecture halls, and transportation hubs). AMuSe does not require changes to the IEEE 802.11 protocol or wireless hardware. Therefore, it can be deployed as an overlay network on existing wireless infrastructure. This overlay network is comprised of AMuSe server on the network side and light-weight application-layer software on the mobile devices. This makes AMuSe attractive for delivering live video content to a dense user population that shares common interests (e.g., providing simultaneous video feeds of multiple camera angles in a sports arena).

We present the AMuSe system architecture in Fig. 1. Unlike our previous papers [8]–[11], where each one concentrates on a specific system aspect, this paper provides a complete overview of our solution.

Feedback Control [8], [11]: The core difficulty in enabling multicast is collecting limited yet sufficient feedback from the users for analyzing the network performance. *AMuSe leverages a low-overhead feedback mechanism for tuning the network* *parameters, i.e., optimizing the network utilization while preserving Quality of Service (QoS) requirements.*

The AMuSe feedback mechanism is based on the following hypothesis, which was proposed in [12] and was validated in our previous papers [8], [11].
Main Hypothesis: A cluster of adjacent nodes experience similar channel quality and suffer from similar interference levels. Hence, a single node can represent the service quality observed by the nodes in the cluster.
AMuSe dynamically divides the nodes in a network into clusters based on the adjacency of nodes and maximum cluster size (D m). In each cluster, the node with the worst channel condition is selected as a *Feedback (FB) node*. The FB node updates the AMuSe server about their experienced service quality, e.g., channel quality, which is then used for meeting the system objectives.
(ii) **Rate Adaptation** [10]: AMuSe enables the APs to transmit multicast traffic at the highest possible bitrate while meeting constraints set by a network operator, i.e. ensuring high *Packet Delivery Ratio* (PDR) for a vast majority of the nodes.
(iii) **Loss Recovery:** In large multicast groups, even a small percentage of packet losses can lead to large number required packet retransmissions. In such situations, application-level forward error correction (FEC) code is a more suitable option. Service reports collected by the feedback mechanism can be used to adjust the amount of FEC dynamically.
(vi) **Dynamic Content Control:** To ensure high quality of experience (QoE) to the users, AMuSe properly adapts the transmitted content and its video coding to the available bandwidth of each AP.

B. Our Contribution

We provide a comprehensive description of the AMuSe system for scalable and efficient wireless multimedia content distribution in crowded venues. The contributions of the paper can be summarized as follows:

- We first describe the commercial opportunities of the AMuSe system based on our interactions with over 100 executives and engineers across different industries, in Section III.
- After defining our model and goals in Section IV, we describe our experiment setup on the ORBIT testbed with hundreds of WiFi nodes and present our key observations in Section V.
- In section VI we presents the system components and discuss the design challenges.
- We show a potential application of AMuSe and its feature in Section VIII and consider a large scale deployment aspects in Section VII.
- We provide typical experimental results under various operation scenarios in Section IX.
- We conclude in Section X with a discussion on potential future steps.

II. Related Work

Various methods have been proposed for multimedia content dissemination to multiple receivers by using wireless multicast. This brief overview describes the most relevant studies to our paper (comprehensive overview on wireless multicast appears in [13]).

TABLE I
MULTICAST: FEATURES OF RELATED WORK

	Scalable (a)	**QoS Guarantees (b)**	**High Util. (c)**	**Standards Compatible (d)**	**Low Cost (e)**
Unicast	x	√	x	√	x
Basic multicast	√	x	x	√	√
Individual Feedback	x	√	x	x	√
Pseudo Broadcast	√	x	x	√	√
LBP	√	x	x	x	√
AMuSe	√	√	√	√	√

Multicast Feedback Mechanisms: Solutions for improving multicast service quality are based on collecting feedback from the receivers and adapting the sender rate. They integrate Automatic Repeat Request (ARQ) mechanisms into the protocol [14]–[21], add Forward Error Correction (FEC) packets [22]–[25], and utilize RA methods [16], [26]–[28]. The feedback mechanisms can be classified into four categories:
(i) Collecting *Individual Feedback* from all users for each received packet [15], [21], [25], [29]–[31]. Although this provides reliable feedback, it does not scale for large groups.
(ii) The *Leader-Based Protocol with acknowledgements (LBP-ACK)* method [16], [18], [23], [31], [32] selects a few receivers to provide feedback, typically the receivers with the lowest channel quality.
(iii) *Pseudo-Broadcast* [17], [32], [33] converts the multicast feed to a unicast flow and sends it to one leader. The leader acknowledges the reception of the unicast flow while the other receivers receive packets by listening to the channel in promiscuous mode.
(iv) The *Leader-Based Protocol with negative acknowledgements* (LBP-NACK) [14], [28], [34] method improves Pseudo-Broadcast by allowing the other receivers to send NACKs for lost packets.
The leader based approaches (ii)-(iv) cannot provide guarantees on the feedback accuracy [8], [21]. Moreover, the LBP-ACK and LBP-NACK methods require changes to the standard.

Additionally, [22], [24], [35] propose using strong FEC for overcoming losses without specifying any feedback mechanism. Others [17], [21], [28], [32] balance between the accuracy requirements and low overhead by using a combination of methods (e.g., Pseudo-Broadcast with infrequent reports from the other receivers).
Multicast Rate Adaptation (RA): In [16], [17], [26], [27], [36] the sender uses feedback from leaders (nodes with worst channel conditions) for RA. In [28] when the channel conditions are stable, RA is conducted based on reports of a single leader. When the channel conditions are dynamic, feedback is collected from all nodes. Medusa [32] combines Pseudo-Multicast with infrequent application layer feedback reports from all nodes. The MAC layer feedback sets backoff parameters while application layer feedback is used for RA and retransmissions of video packets.

Table I summarizes the main features of existing approaches. In summary, at least one of the following weaknesses hinders their performance: (i) Requirement of feedback from a large number of receivers. (ii) Ignorance of AP to interference-related packet loss, (ii) Low network utilization to compensate for lack of feedback information or due to abnormal nodes, (iv) Requirement of changes to standard WiFi protocol, or (v) Expensive deployment of numerous APs.

III. Industry Needs

To identify the needs in the domain of wireless multicast, we interacted with over 100 executives and engineers in a variety of relevant industries such as telecom, sports, entertainment, etc. Based on our interactions we identified several scenarios where AMuSe can be beneficial.

Sports Stadiums: Sports teams and stadium owners are looking for new ways to engage fans at the game. The in-stadium experience of fans lags the in-home experience where the audience can watch enhanced multimedia content such as replays, views from different angles, etc. Recent fan surveys have indicated that for a large percentage of the audience, the poor multimedia experience at stadiums negatively affects the game experience. There is also a growing interest from the fans for next generation content such as virtual reality and 3D video during a game. To meet the growing needs, stadiums are heavily investing in dense WiFi deployments. The cost of such deployments can be as high as $10 million for a large stadium. Even with this capital expenditure, supporting high quality video delivery to thousands of users may not be feasible. AMuSe is particularly attractive for serving the needs of this industry.

Concert Halls and Theaters: Several concert halls and theaters are seeking innovative ways to connect with their audience and monetize their existing video content. This includes giving visitors access to behind the scenes content during intermissions, zoomed-in views of performers, subtitles in different languages, etc. Besides the aforementioned problem of prohibitive capital expenditure associated with dense wireless deployments, many such venues cannot deploy additional network infrastructure due to the historical nature of the buildings. These limitations make AMuSe an attractive solution to provide content to users in such venues.

Public Events: These events include but are not limited to parades, university commencements, and music festivals. A common problem for the organizers of such events is a lack of existing network infrastructure. Deploying extensive infrastructure for such sporadic events is not economical and cost-effective. At the same time, the audience in such events can benefit from enhanced views of the live event and additional venue specific content. In such a scenario, AMuSe can provide real-time video updates to users with much less capital expenditure.

Emergency Services: Quick dissemination of information to the public during emergencies is of critical importance. During such events, the network infrastructure could be especially strained or non-functional. AMuSe can allow for fast information broadcast using commodity hardware.

Fig. 2. A photo of the ORBIT Testbed [37].

IV. Network Model and Objective

A. Network Model

We consider a WiFi LAN with multiple APs and frequency planning such that the transmissions of adjacent APs do not interfere with each other. Thus, *for our feedback node selection and rate adaptation we can consider each AP separately*. We assume low mobility (e.g., users watching a sports event). Although we consider a controlled environment, the network may still suffer from sporadic interference, as shown in Section V.

B. Objective

Our *objective is to develop a practical and efficient wireless video distribution system to very large groups of users in crowded places while meeting several device quality requirements.* As explained in Section VI-C a user experiences high service quality if its *Packet Delivery Ratio* (PDR) is above a given *PDR threshold* L, (e.g., $L = 85\%$). Given a *Population-Threshold* X (e.g., $X = 99\%$) the system should satisfies the following requirements:

(R1) High throughput – Operate at the highest possible rate, i.e., the *target rate*, while preserving SLAs.

(R2) Service Level Agreements (SLAs) – Given a *PDR Threshold* L and a *Population-Threshold* X, the system should guarantee that at least $X\%$ of the nodes experience PDR above L (i.e., are normal nodes). Except for short transition periods, this provides an upper bound of $A_{max} = \lceil n \cdot (1 - X) \rceil$ on the number of permitted abnormal nodes.

(R3) Scalability – Support hundreds or thousands of nodes per AP.

(R4) Stability – Avoid rate changes due to sporadic channel condition changes.

(R5) Fast Convergence – Rate adaptation operations should converge fast to the target rate after long-lasting changes (e.g., user mobility or network changes).

(R6) Standard and Technology Compliance – No change to the IEEE 802.11 standard or operating system of the nodes.

V. ORBIT - Challenges and Key Observations

We evaluate AMuSe on the ORBIT testbed [37], shown in Fig. 2, which is a dynamically configurable grid of 20×20 (400) 802.11 nodes where the separation between nodes is 1m. It is a good environment to evaluate AMuSe, since it provides a very large and dense population of wireless nodes, similar to the anticipated crowded venues.

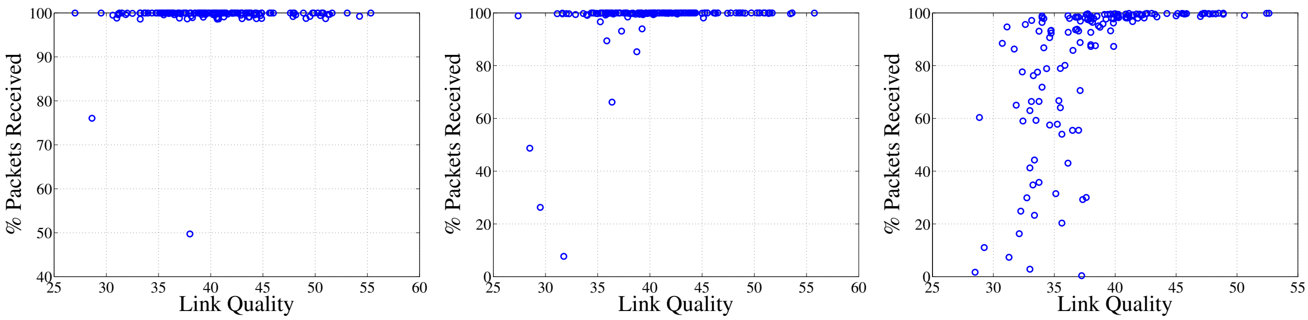

(a) PDR vs. LQ, TX_{AP} = 6 Mbps. (b) PDR vs. LQ, TX_{AP} = 36 Mbps. (c) PDR vs. LQ, TX_{AP} = 48 Mbps.

Fig. 3. Experimental results for evaluating the nodes PDR vs. LQ for different TX_{AP} rates.

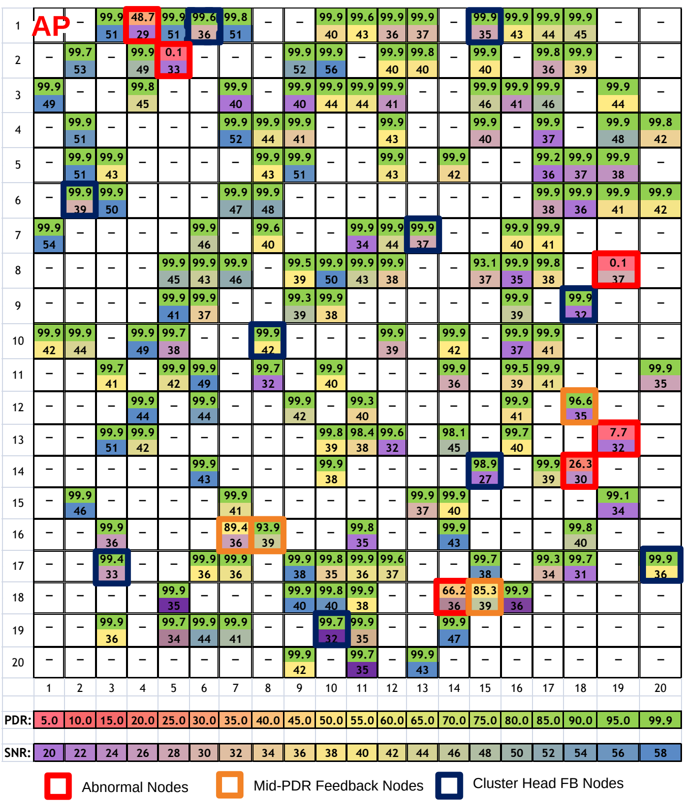

Fig. 4. Heatmap of the PDR and LQ values of all the nodes as well as the selected FB nodes in one experiment in which the node in location 1,1 operates as the AP with $TX_{AP} = 36$Mbps. Each square represents a single node, where the upper number is the nodes PDR while the lower number is its LQ.

Our Experiments: To avoid performance variability due to a mismatch of WiFi hardware and software, only nodes equipped with *Atheros 5212/5213* cards with *ath5k* driver were selected. For each experiment *we activated all the operational nodes* that meet these specifications (between 150 and 250 nodes). In all the experiments, one corner node served as a single multicast AP. The other nodes were multicast receivers. To imitate RTP flows of several video flows, the AP used 802.11a to send a multicast UDP flow, where each packet was 1400 bytes.

We turn to describe some key observations from the ORBIT Testbed and the corresponding design challenges.

Low Transmission Power: All the ORBIT nodes are concentrated in a small room. When using typical WiFi transmission power, i.e., $30-100$mW, all nodes experience excellent channel condition and can decode messages at the maximal supported bit-rate of 54Mbps. Therefore, *we set the AP to use the lowest supported transmission power of* $1mW = 0dBm$ *to ensure that the channel conditions of some nodes are marginal.* As a result, the provided service become very sensitive to external interference and the optimal transmission rate has changed between experiments, in the range of 24-48Mbps.

Service Quality Evaluation Metrics: The WiFi nodes provide us two metrics for evaluating the channel condition. The first is the *Link Quality* (LQ) which is lineally related to the received signal strength. We expected the LQ to be a good measure of the SNR if there is no external noise. The second is the Packet Delivery Ratio (PDR). We evaluated the relation between nodes PDR vs. their LQ for different multicast transmission rate, denoted as TX_{AP}. Typical results for $TX_{AP} = 6, 36, 48$Mbps are given in Fig. 3. The figure (mainly Fig 3(c)) shows that the correlation between the PDR and LQ is not strong, suggesting that nodes with the same LQ value may have significantly different PDR. This observation indicates that *LQ is not an appropriate metric for evaluating the service quality.*

Abnormal Nodes: We refer to a node with PDR below the PDR Threshold $L = 85\%$ as *abnormal*, otherwise it is called *normal*. Fig. 3(a) demonstrates that even in the extreme case of very low TX_{AP} without any interference some of the nodes (two in this case) are abnormal and suffer from low PDR. The set of abnormal nodes remained small when we increase TX_{AP} to higher bitrates until 36 Mbps, as shown in Fig. 3(b). The number of abnormal nodes increases significantly once TX_{AP} reaches 48 Mbps. Surprisingly, the set of abnormal nodes is not the same in all experiments. These experiments demonstrate that *it is impossible to provide high service quality to all users, without scarifying the network utilization.*

Distance from the AP: We evaluate the impact of the distance from the AP on the experienced service quality. Fig. 4 shows a heatmap of the PDR and LQ values of the nodes for a single experiment with a $TX_{AP} = 36$Mbps. We checked the locations of the abnormal nodes and found out that one of the nodes at a distance of 3 meters from the AP (at location 1,4) suffers from PDR of 48%. Overall, we observed very weak correlation between the distance from the AP and the channel

condition. This observation results from the following aspects: (1) *Reflection and fast fading effects* - The testbed is located in a small room, therefore reflection from the walls and fast fading effects significantly impact the channel condition of the nodes.
(2) *Static Testbed* - The testbed is static, the WiFi nodes are mounted to the walls and the room typically does not contain people. As a result, the channel condition of a node does not change during an experiment, even if it is severely affected by fast fading effects.
Sporadic Interferences: In all of the experiments we observed that even at a low rate, the channel may suffer from sporadic interference events, which cause a sharp increase in the number of abnormal nodes. These interference spikes caused by unknown sources are beyond our control and their duration varies in time. Further, the location of the nodes affected by the spikes varies with time and does not follow a known pattern. These experiments show that even in a seemingly controlled environment, *nodes may suffer from sporadic continuous interference, which may cause service quality fluctuations*.
WiFi Beacon Messages: While analyzing the performance, we noticed that clients disconnect from the AP at high bit-rates, thereby causing performance degradation. This results from the fact that increasing the bit-rate also increases the WiFi beacon bit-rate which may not be decoded at some nodes. A sustained loss of beacons leads to node disconnection. To counter this, we modified the ath5k driver to send beacons at the minimum bit-rate.

VI. THE AMUSE SYSTEM OVERVIEW

The AMuSe system is composed of four main modules, as illustrated in Fig. 1. We provide a high level description of each module in the following subsections. Detailed descriptions of the first two modules are given in [10], [11], respectively. For our description we assume appropriate frequency planning, such that adjacent APs don't interfere with each other.

A. The Feedback Mechanism

One of the core operations of AMuSe is collecting limited yet sufficient feedback from the users for optimizing the network performance. This feedback can be used for variety of tasks, including: (i) multicast rate adaptation, (ii) FEC tuning and (iii) detection of interference sources. The feedback mechanism is based on the following hypothesis, which was reported in [12] and was also validated in our studies [8], [11].
Hypothesis: A cluster of adjacent nodes experience similar channel quality and suffer from similar interference levels. Hence, a node v with a worse channel condition than its adjacent neighbors can represent the service quality observed by the nodes in the cluster.

While this hypothesis is generally valid, as explained in Section V, some users, i.e., abnormal nodes, may suffer from atypical low service quality. For monitoring both the number of abnormal nodes as well as the overall user satisfaction from the service, AMuSe uses two types of FB nodes.

- *Singleton* – An FB node that represents just itself. Singletons are used for monitoring nodes with low-PDR, such as Abnormal Nodes.
- *Cluster-Head* – The node with the lowest channel quality in its vicinity (its *cluster*), excluding singletons. Such FB node represents itself and the other non-FB-nodes in its cluster.

Evaluation Metric: As described in Section V, the service quality that the nodes experience can be evaluated by both the PDR and LQ metrics. Fig. 3 shows that LQ provides inaccurate estimation of the signal quality. Therefore, we prefer to use the node PDR to estimate a user experience. However, for nodes with very high PDR values, i.e. above 98%, this metric cannot be used to identify the node with the weakest channel condition. To overcome this problem, AMuSe uses a *mixed metric* that considers both the PDR and the LQ. For nodes with PDR $\leq 98\%$, the ordering is based on PDR, while for nodes with PDR $> 98\%$, the comparison is based on their LQ. Thus, the channel quality is defined by the following tuple in lexicographic order: $(\min(PDR, 98), LQ)$.
Clustering Requirements: For a given distance D, two nodes are termed *D-adjacent* if they are separated by a distance of at most D. In order to find a small set of FB nodes that can provide accurate reports, the selected cluster-head FB nodes should satisfy the following requirements.

(i) Each node should be D-adjacent to an FB node.
(ii) An FB node must have similar or weaker channel quality than the other nodes in its cluster, excluding singletons.
(iii) Any two FB nodes cannot be D-adjacent.

Figure 4 illustrates the selected set of FB nodes according to the PDR and LQ values. The FB nodes marked with red and orange squares suffer from low PDR (below 98%) and are singletons, while the nodes marked with blue squares are cluster heads with maximal cluster radius of $D = 6$ meters. Observe that all cluster heads have PDR above 98% and their selection is based on their LQ values. Although only a few cluster heads are selected, they are well distributed throughout the transmission area and the maximal distance of any non-FB node to a cluster head is at most 6 meters.
FB Node Selection: This is a quasi-distributed process in which singleton and cluster head FB nodes are selected according to their PDP and LQ values as well as their locations. The process works as follows: At the beginning of each *reporting interval* the AP sends a message with the FB node list, which contains the selected singletons and cluster-heads, their locations and their latest channel quality measurements. Upon receiving this message, each FB node, either cluster-head or singleton, waits a short random time for avoiding collisions and then reports its measured PDR and LQ to the AP. Every other node checks whether it is adjacent (with maximal distance of D) to a cluster-head with lower channel quality[1]. If it cannot detect such a cluster-head within 2 reporting intervals, it volunteers to serve as a cluster-head.

To avoid a swarm of volunteer messages, a node postpones its message a small random number of reporting intervals proportional to its channel quality (short delay for bad channel

[1]To avoid oscillations, the scheme allows some margins.

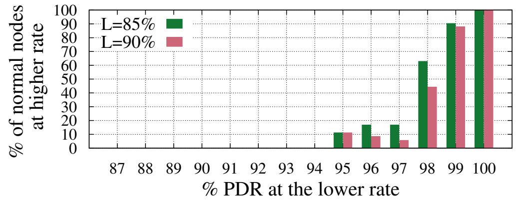

Fig. 5. The percentage of nodes that remain normal after increasing the TX_{AP} from 36Mbps to 48Mbps vs. their PDR values at the 36Mbps for different PDR-thresholds (L).

quality and long delay for good one). After each round, the AP refreshes the set of FB nodes and selects the ones with lowest channel quality at each cluster. Volunteers with exceptional low PDR are selected as singletons.

The FB selection process ensures an upper bound on the number of FB nodes, regardless of the receiver group size, by proper selection of configuration parameters (e.g., the cluster radius). This upper bound is required for limiting the interference due to FB reports.

B. The Multicast Rate Adaptation Algorithm

The rate adaptation mechanism is composed of two components:

(a) Rate Decision: Utilizes the limited and infrequent FB reports to determine the highest possible rate, termed the *target-rate*, while meeting the requirements in Section IV-B.

At any given time, the FB reports are available only for the current rate. Hence, a key challenge is to *determine if the AP operates at the target-rate, without having FB reports from higher rates*[2]. We refer to this assessment as the *target condition*. Unfortunately, the target-rate cannot be detected from RF measurements, such as LQ. As shown in [38], [39] as a well as our studies, different nodes may have different receiver sensitivities, which may result in substantial PDR gaps between nodes with similar RF measurements. However, large scale multicast environments enable us to efficiently predict the target condition as described next.

Observation I: When operating below the target-rate, almost all the nodes have PDR close to 100%. However, when operating at the target-rate, a noticeable number of receivers experience PDR below 97%.

For instance, the system in Fig. 4 contains 161 nodes so the number of permitted abnormal nodes $A_{max} = \lceil 161 \cdot 5\% \rceil = 8$. At 36Mbps, 10 nodes had PDR below 97%. We derive the next observation from Fig. 5, which shows the average percentage of nodes that remain normal vs. their initial PDR when increasing TX_{AP} from 36Mbps to 48Mbps averaged for 3 different sets of experiments. The total number of nodes in these experiments were between 160 to 170.

Observation II: There is a PDR threshold, $H = 97\%$, such that every node with PDR between L and H becomes abnormal after the rate increase with very high probability. We refer to these nodes as *mid-PDR nodes*.

Observation II is not surprising. As reported in [38], [40], each receiver has an SNR band of $2-5$dB, in which its PDR drops from almost 100% to almost 0%. The SNR of mid-PDR nodes lies in this band. Increasing the rate requires $2-3$dB higher SNR at the nodes. Hence, mid-PDR nodes with SNR in the transition band before the rate increase will be below or at the lower end of the transition band after the increase, and therefore, become abnormal nodes.

Observations I and II imply that it is possible to assess the target condition by monitoring the nodes close to transitioning from normal to abnormal.

(b) Stability Control: Users are very sensitive to changes in video quality [41]. Therefore unnecessary rate changes should be avoided. AMuSe uses a window based method for maintaining rate stability in the event of sporadic interference and after a rate adaptation decision.

C. Loss Recovery

Packet loss may significantly hinder the video quality. AMuSe can handle mild amount of losses (below 15%) by adding application level FEC [42], [43] to the multicast streams. Our PDR-Threshold $L = 85\%$ was selected to allow receivers to handle losses in the event of short simultaneous transmission of another node. In such a situation, the collision probability is below $2/W_{min}$, where W_{min} is the minimal contention window. For 802.11a/g/n $W_{min} = 16$, which implies collision probability is below 12.5%. Therefore receivers with high PDR (near 100%) should be able to compensate for lost packets due to sporadic interference events or temporary simultaneous flows.

D. Content Control

This module contains content selection and video coding components to ensure high QoE to the users. As we show in Section IX, AMuSe quickly converges to the target rate, however this rate may change due to user mobility or interference[3]. Therefore at any given time the Video Control module is required to adjust the video quality to the available bandwidth.

If the system suffers from strong interference, other means should be used. For instance, the multicast content can be divided into high and low priority flows, augmenting the high priority flow with stronger FEC during the interference time, while postponing low priority flows.

Authentication and Copy Right Management (CRM): Content distribution systems are also required to authenticate the users and handle CRM aspects. Since AMuSe is designed for wireless content distribution in a specific venue, we assume that all the users in the venue are entitled to receive the content. If this is not the case, the CRM aspects are beyond the scope of this paper.

VII. Large Scale Deployment

AMuSe can leverage existing WiFi infrastructure to provide high quality video to thousands of users by using WiFi multicast technology. The AMuSe system is composed of three types of components as seen in Fig. 6: a single AMuSe server,

[2]In [10] we show that when acceding the target rate numerous users may suffer from poor service for a few seconds.

[3]Although we assume appropriate frequency planning sporadic interferences may be generated by the users or rogue APs.

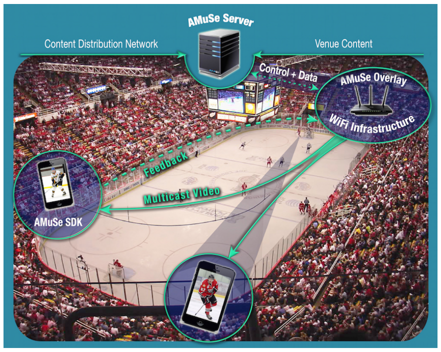

Fig. 6. Deployment of AMuSe System at a venue. The AP multicasts video to mobile devices equipped with AMuSe mobile application. A server exchanges data and control information with the AP.

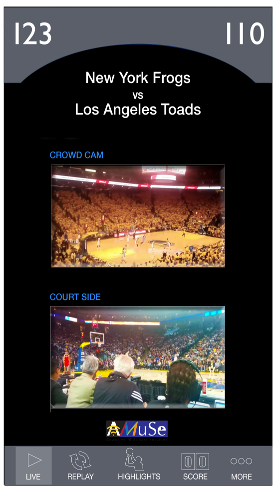

Fig. 7. The AMuSe Mobile Application where users can select different camera angles or replays, and watch high quality video content.

wireless APs, and mobile devices. The AMuSe server controls the multicast parameters, such as the AP muticst rate and the video coding. It also sends the different video streams to the APs. The APs multicast the video content to the mobile phones, which are equipped with a mobile application for watching the video channels and interacting with AMuSe servers.

The number of APs required depends on the size of the venue and the WiFi technology employed. In the case of IEEE 802.11a/g deployment, an AP can typically support a cell radius of 20 to 30 meters in which the users can benefit from high bit-rate of 54Mbps. For instance, for a stadium with 20,000 seats, this equates to serving 2,000-3,000 seats per AP. Using a conservative estimate of 1,000 seats per AP, only 20 APs are required to deliver the desired content. This shows that AMuSe offers a more cost effective solution than unicast based systems, which require hundreds of APs to provide the same service.

VIII. AMuSe Mobile Application

At an event the audience can benefit from the AMuSe system by using the venue mobile application, which allows users to obtain desired content while supporting AMuSe user-side feedback mechanism. Through an Electronic Program Guide (EPG) the venue can deliver unique multimedia content to the end user. An example of what the mobile application would look like for a sports game can be seen in Fig. VIII. The application will allow the fans to watch different live views, replays, statistics, and other behind the scenes content.

During the experimental tests using IEEE 802.11a technology and marginal channel quality, the evaluated AMuSe system usually transmited at a bit rate ranging from 18 Mbps to 36 Mbps, giving a throughput of 10 Mbps to 20 Mbps. This allows for transmission of 5 to 10 parallel video streams. This number can be scaled for newer WiFi standards, i.e., IEEE 802.11n/ac is an appropriate AP replacement. The video flows can be curated and updated during the event. Additional content can be downloaded to the application prior to the event to increase the information that can be consumed at any given time. The combination of these options along with anything else the venue wants to offer to the customers will greatly improve the watching experience.

IX. Experimental Evaluation

For evaluating the performance of AMuSe on the ORBIT testbed, we use the performance metrics described below:
(i) *Multicast rate and throughput*: The time instants when the target condition is satisfied are marked separately.
(ii) *PDR at nodes*: Measured at each node.
(iii) *Number of abnormal and mid-PDR nodes*: We monitored *all* the abnormal and mid-PDR nodes (not just the FB nodes).
(iv) *Control traffic*: The feedback overhead (this overhead is very low and is measured in Kbps).

We compared AMuSe to the following schemes:
(i) *Fixed rate scheme*: Transmit at a fixed rate of 36Mbps, since it is expected to be the target rate.
(ii) *Pseudo-multicast*: Unicast transmissions to the node with the lowest SNR/RSS. The unicast RA is the driver specific RA algorithm *Minstrel* [44]. The remaining nodes are configured in promiscuous mode.
(iii) *Simple Rate Adaptation (SRA) algorithm [8]*: This scheme also relies on measuring the number of abnormal nodes for making rate adaptation decisions. Yet, it is not designed to achieve the target rate, maintain stability, or respond to interference.

A. Performance Comparison

We evaluated the performance of AMuSe in several experiments on different day with $160 - 170$ nodes. Fig. 8 shows one instance of such an experiment over 300s with 162 nodes. Fig. 8(a) shows the mid-PDR and abnormal nodes for the duration of one experiment run. Fig. 8(b) shows the rate determined by AMuSe . The AP converges to the target rate after the initial interference spike in abnormal nodes at 15s. The AP successfully ignored the interference spikes at time instants of 210, 240, and 280s to maintain a stable rate. The target-condition is satisfied except during the spikes. The overall control overhead measured was approximately 40Kbps. The population of abnormal nodes stays around $2 - 3$ for most of the time which implies that more than 160 nodes ($> 98\%$) have a PDR $> 85\%$. The actual throughput is

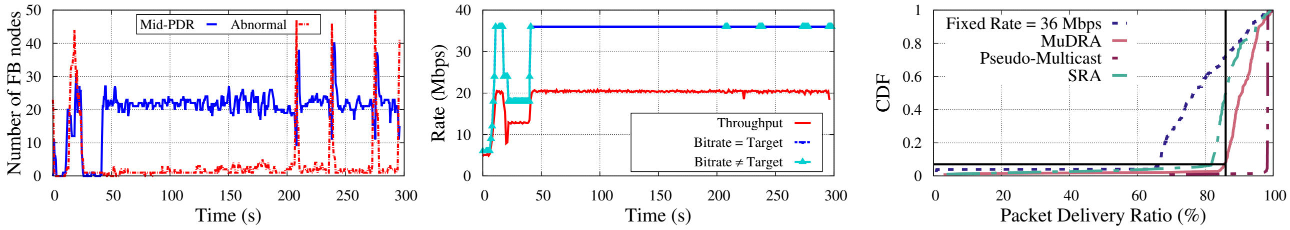

(a) (b) (c)

Fig. 8. (a) Mid-PDR and abnormal nodes for a typical sample of AMuSe's operation over 300s with 162 nodes:, (b) Multicast rate and throughput measured at the AP for a typical sample of AMuSe's operation over 300s with 162 nodes, and (c) CDF of PDR distributions of 162 nodes over several experiments for fixed rate, AMuSe , Pseudo-Multicast, and SRA schemes.

TABLE II
AVERAGE THROUGHPUT (MBPS) OF PSEUDO-MULTICAST, AMUSE , AND SRA SCHEMES.

	Throughput
Fixed rate = 36Mbps	20.42
Pseudo-Multicast	9.13
AMuSe	18.75
SRA	19.30

stable at around 20Mbps which after accounting for 15% FEC correction implies a goodput of 17Mbps.

The average throughput for different schemes over 3 experiments of 300s each (conducted on different days) with 162 nodes is shown in Table II. AMuSe achieves 2x throughput than pseudo-multicast scheme. The fixed rate scheme yields approximately 10% higher throughput than AMuSe. SRA has similar throughput as AMuSe .

Fig. 8(c) shows the distribution of average PDR of 162 nodes for the same 3 experiments. In the pseudo-multicast scheme, more than 95% of nodes obtain a PDR close to 100% (we did not consider any retransmissions to nodes listening in promiscuous mode). AMuSe meets the QoS requirements of 95% nodes with at least 85% PDR. On the other hand, in SRA and the fixed rate schemes 45% and 70% of the nodes have PDR less than 85%, respectively.

In pseudo-multicast, more reliable transmissions take place at the cost of reduced throughput, since the AP communicates with the node with the poorest channel quality in unicast. The significant difference in QoS performance of the fixed rate and SRA schemes is because the target rate can change due to interference, etc. In such a situation, AMuSe can achieve the new target rate while the fixed rate and SRA schemes lead to significant losses (we observed that exceeding the target rate even 10% of time may cause up to 20% losses and less than 5% throughput gain).

X. CONCLUSION

In this paper, we present the design and large-scale experimental evaluation of the AMuSe system for providing scalable and efficient WiFi multicast services to a large number of users. AMuSe only needs access to the channel quality measurements such as Link Quality (LQ) and Packet Delivery Ratio (PDR) on WiFi devices and can be implemented as an application layer protocol on existing devices. Our extensive experiments on the ORBIT testbed with hundreds of nodes show that AMuSe can reliably support applications such as large scale multimedia content delivery while optimizing the system performance.

Future Work: In the future we intend to refine our feedback mechanism to distinguish between losses due to channel conditions and collisions. While rate adaptation is the right approach for dealing with degradation of channel conditions, it may not be appropriate for handling collisions where more aggressive loss recovery methods are required.

WiFi vs. LTE: The concepts presented in this paper are also applicable for improving the performance of LTE evolved Multimedia Broadcast Multicast Services (eMBMS). The use of WiFi or cellular eMBMS depends on the venue and the copyright issues of the video streams. Stadium owners may prefer to deploy the existing WiFi networks for multicast which is more cost effective and provides control over content. Cellular networks may provide more value to outdoor venues and other public events since cellular networks offer wider coverage than WiFi.

XI. ACKNOWLEDGMENTS

We would like to thank Dr. Jaime Ferragut, Dr. Katherine Guo, and Dr. Thyaga Nandagopal for their several key contributions to the project. We also thank Ivan Seskar for his critical support and several useful technical discussions. This work was supported in part by NSF CIAN ERC under grant EEC-0812072 and the People Programme (Marie Curie Actions) of the European Unions Seventh Framework Programme (FP7/20072013) under REA grant agreement no. [PIIF-GA-2013-629740].11.

REFERENCES

[1] Y. Tanigawa, K. Yasukawa, and K. Yamaoka, “Transparent unicast translation to improve quality of multicast over wireless LAN,” in *IEEE CCNC'10*, 2010.

[2] “Cisco, white-paper, Cisco connected stadium Wi-Fi solution,” 2011. [Online]. Available: http://www.cisco.com/web/strategy/docs/sports/c78-675063_dSheet.pdf

[3] “Yinzcam,” http://www.yinzcam.com/.

[4] S. Narayan, C. Jayawardena, J. Wang, W. Ma, and G. Geetu, “Performance test of ieee 802.11ac wireless devices,” in *ICCCI'15*, 2015.

[5] N. Hajlaoui and I. Jabri, “On the performance of IEEE 802.11n protocol,” in *ACM WiNTECH'12*, 2012.

[6] K. Pelechrinis, T. Salonidis, H. Lundgren, and N. Vaidya, “Experimental characterization of 802.11n link quality at high rates,” in *ACM WiNTECH'10*, 2010.

[7] C. Perkins and M. McBride, “Multicast WiFi problem statement,” Working Draft, IETF Internet-Draft, 2015, https://tools.ietf.org/html/draft-mcbride-mboned-wifi-mcast-problem-statement-00.

[8] Y. Bejerano, J. Ferragut, K. Guo, V. Gupta, C. Gutterman, T. Nandagopal, and G. Zussman, "Scalable WiFi multicast services for very large groups," in *IEEE ICNP'13*, 2013.

[9] ——, "Experimental evaluation of a scalable WiFi multicast scheme in the ORBIT testbed," in *3rd GENI Research and Educational Experiment Workshop (GREE)*, 2014.

[10] V. Gupta, C. Gutterman, Y. Bejerano, and G. Zussman, "Experimental evaluation of large scale WiFi multicast rate control," in *Proc. IEEE INFOCOM'16*, 2016.

[11] V. Gupta, Y. Bejerano, C. Gutterman, J. Ferragut, K. Guo, T. Nandagopal, and G. Zussman, "Experimental evaluation of large scale wifi multicast rate control," *IEEE Trans. Netw. (to appear)*, 2016.

[12] 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.

[13] J. Vella and S. Zammit, "A survey of multicasting over wireless access networks," *IEEE Commun. Surveys Tuts.*, vol. 15, no. 2, pp. 718–753, 2013.

[14] J. K. Kuri and S. Kumar, "Reliable multicast in multi-access wireless LANs," *ACM/Kluwer Wirel. Netw.*, vol. 7, pp. 359–369, 2001.

[15] M.-T. Sun, L. Huang, A. Arora, and T.-H. Lai, "Reliable MAC layer multicast in IEEE 802.11 wireless networks," in *IEEE ICPP'02*, 2002.

[16] J. Villalon, P. Cuenca, L. Orozco-Barbosa, Y. Seok, and T. Turletti, "Cross-layer architecture for adaptive video multicast streaming over multirate wireless LANs," *IEEE J. Sel. Areas Commun.*, vol. 25, no. 4, pp. 699 –711, 2007.

[17] 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 *IEEE ICNP'09*, 2009.

[18] N. Choi, Y. Seok, T. Kwon, and Y. Choi, "Leader-based multicast service in IEEE 802.11v networks," in *IEEE CCNC'10*, 2010.

[19] Z. Li and T. Herfet, "BLBP: a beacon-driven leader based protocol for MAC layer multicast error control in wireless LANs," in *IEEE WiCOM'08*, 2008.

[20] V. Srinivas and L. Ruan, "An efficient reliable multicast protocol for 802.11-based wireless LANs," in *IEEE WoWMoM'09*, 2009.

[21] X. Wang, L. Wang, Y. Wang, Y. Zhang, and A. Yamada, "Supporting MAC layer multicast in IEEE 802.11n: Issues and solutions," in *IEEE WCNC'09*, 2009.

[22] E. Park, S. Han, H. Kim, K. Son, and L. Jing, "Efficient multicast video streaming for IPTV service over WLAN using CC-FEC," in *IEEE ICICSE'08*, 2008.

[23] O. Alay, T. Korakis, Y. Wang, and S. Panwar, "Dynamic rate and FEC adaptation for video multicast in multi-rate wireless networks," *ACM Mobile Netw. and Appl.*, vol. 15, no. 3, pp. 425–434, 2010.

[24] 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 *IEEE BMSB'12*, 2012.

[25] 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.

[26] Y. Seok and Y. Choi, "Efficient multicast supporting in multi-rate wireless local area networks," in *IEEE ICOIN'03*, 2003.

[27] A. Basalamah, H. Sugimoto, and T. Sato, "Rate adaptive reliable multicast MAC protocol for WLANs," in *IEEE VTC'06*, 2006.

[28] 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.

[29] X. Wang, L. Wang, and D. Wang, Y.and Gu, "Reliable multicast mechanism in WLAN with extended implicit MAC acknowledgment," in *IEEE VTC'08*, 2008.

[30] V. Srinivas and L. Ruan, "An efficient reliable multicast protocol for 802.11-based wireless LANs," in *IEEE WoWMoM'09*, 2009.

[31] Z. Feng, G. Wen, C. Yin, and H. Liu, "Video stream groupcast optimization in WLAN," in *IEEE ITA'10*, 2010.

[32] S. Sen, N. K. Madabhushi, and S. Banerjee, "Scalable WiFi media delivery through adaptive broadcasts," in *USENIX NSDI'10*, 2010.

[33] Y. Park, C. Jo, S. Yun, and H. Kim, "Multi-room IPTV delivery through pseudo-broadcast over IEEE 802.11 links," in *IEEE VTC'10*, 2010.

[34] Z. Li and T. Herfet, "HLBP: a hybrid leader based protocol for MAC layer multicast error control in wireless LANs," in *IEEE GLOBECOM'08*, 2008.

[35] V. Sgardoni, M. Sarafianou, P. Ferre, A. Nix, and D. Bull, "Robust video broadcasting over 802.11a/g in time-correlated fading channels," *IEEE Trans. Consum. Electron.*, vol. 55, no. 1, pp. 69–76, 2009.

[36] K. Lin, W. Shen, C. Hsu, and C. Chou, "Quality-differentiated video multicast in multi-rate wireless networks," *IEEE Trans. Mobile Comput.*, vol. 12, no. 1, pp. 21–34, January 2013.

[37] "ORBIT testbed," http://orbit-lab.org/.

[38] C. Reis, R. Mahajan, M. Rodrig, D. Wetherall, and J. Zahorjan, "Measurement-based models of delivery and interference in static wireless networks," in *ACM SIGCOMM'06*, 2006.

[39] D. Halperin, W. Hu, A. Sheth, and D. Wetherall, "Predictable 802.11 packet delivery from wireless channel measurements," in *ACM SIGCOMM'10*, 2010.

[40] M. R. Souryal, L. Klein-Berndt, L. E. Miller, and N. Moayeri, "Link assessment in an indoor 802.11 network," in *IEEE WCNC'06*, 2006.

[41] N. Cranley, P. Perry, and L. Murphy, "User perception of adapting video quality," *Int. J. Hum. Comput. St.*, vol. 64, no. 8, pp. 637–647, 2006.

[42] K. N. D. Vukobratovic, "A survey on application layer forward error correction codes for IP datacasting in DVB-H," in *3rd COST 2100 MCM*, 2007.

[43] J. Ababneh and O. Almomani, "Survey of error correction mechanisms for video streaming over the internet," in *IJACSA*, vol. 5, 2014, pp. 155–161.

[44] "Minstrel," https://wireless.wiki.kernel.org/en/developers/documentation/mac80211/ratecontrol/minstrel.