

Performance Evaluation of the IEEE 802.11aa Multicast Mechanisms for Video Streaming

Antonio de la Oliva*, Pablo Serrano*, Pablo Salvador*[†], Albert Banchs*[†]

*Universidad Carlos III de Madrid, Avda. Universidad, 30, 28911 Leganés, Spain

Email: {aoliva,pablo}@it.uc3m.es

[†]Institute IMDEA Networks, Avda. del Mar Mediterráneo, 22, 28918 Leganés, Spain

Email: {josepablo.salvador,albert.banchs}@imdea.org

Abstract—Video traffic is foreseen to account for the majority of the Internet traffic in the near future. While the demand of video transmission keeps growing, the vast majority of wireless equipment deployed in the home environment, based on IEEE 802.11, cannot satisfy the amount of bandwidth that the video applications require. In order to cope with the increasing demand of multimedia traffic, the IEEE 802.11aa Task Group has recently standardized new mechanisms to allow efficient and robust transmission of multicast flows in Wireless LAN. However, the standard leaves open the choice of which one to use for a given scenario. In this paper, we explore the new mechanisms introduced by the 802.11aa Task Group, providing insights of the new choices for handling group addressed frames, by carrying out extensive simulations. Our results highlight the various trade-offs each mechanism has in terms of robustness, resource consumption and complexity, and provide a set of recommended guidelines for their use.

Index Terms—Groupcast; WLAN; 802.11aa

I. INTRODUCTION

The IEEE 802.11 standard for Wireless LAN (WLAN) [1] has become one of the most common technologies to provide broadband connectivity to the Internet at user's premises. The widespread deployment of high-rate access points (APs), with modulation schemes reaching up to 600 Mbps as defined by the amendment 802.11n [2], are enabling the introduction of applications with relatively large bandwidth demands, like e.g., YouTube, Skype videoconferencing or VideoLAN video-streaming.

The original 802.11 standard was poorly suited for the efficient support of multimedia flows, and in particular video streams, because of several reasons: (i) the originally supported physical transmission rates (1 and 2 Mbps) imposed a severe bottleneck on the maximum achievable rate, regardless of the efficiency of the MAC protocol; (ii) only *best-effort* service was supported, thus preventing any sort of traffic differentiation that could improve the performance of multimedia applications, which as compared to data traffic can relax their reliability requirements below 100 % to improve performance; (iii) multicast transmissions were very inefficient and unreliable (as widely reported in various works, see e.g. [3]), which eventually prevented its use on WLANs, as there is no proper support for video streaming to various receivers.

The subsequent amendments to the 802.11 standard have lessened the first two limitations. Indeed, as already mentioned,

the introduction of PHY-amendments have boosted the maximum achievable rates, starting with the 802.11b [4] that increased the maximum rate up to 11 Mbps, continuing with the 802.11a and 802.11g amendments that reach up to 54 Mbps, and finally with the 802.11n [2] which specifies modulation rates up to 600 Mbps. New amendments, i.e., 802.11ac and 802.11ad, both of them under development, propose to increase data rates beyond 1 Gb/s. 802.11ac [5] will work on the 5 GHz WLAN band and is built on 802.11n, while IEEE 802.11ad [6] proposes to use the unlicensed 60 GHz band. These amendments provide a large increase in terms of achievable bandwidth, however the new demands due to the arrival of high definition video impose several burden on the current MAC scheme, in particular for the delivery of multicast traffic. The 802.11e amendment [7] introduced traffic differentiation via the setting of the contention parameters, thus enabling both the ability to prioritize one type of traffic over other type and a more efficient operation of WLANs by proper tuning of the MAC parameters [8]. The remaining challenge, therefore, is to efficiently support multicast over 802.11 WLANs.

The IEEE 802.11aa Task Group has recently addressed this last limitation, with the definition of various mechanisms to support *Robust streaming of Audio Video Transport Streams*. Its focus is to extend the 802.11 standard [1] with mechanisms tailored to improve performance of multimedia streaming over WLAN. In particular, the new mechanisms target at significantly improving both the effectiveness in terms of reliability and the efficiency of the delivery of multicast traffic.

In this paper, we present a thorough simulation-based evaluation of the groupcast mechanisms defined in the current version of the 802.11aa standard [9], in order to assess their performance in a scenario with real video streams. We quantify both the reliability that each of the considered mechanisms provides to a video stream, for different number of receivers and radio channel conditions, as well as the efficiency in terms of the overhead introduced by each mechanism with respect to the scheme that consumes the least amount of resources. This metric will have a direct impact on the bandwidth left for data users. Our results show how the new defined mechanisms substantially enrich the set of services than can be offered to multimedia streams in 802.11 WLANs, and that there is

no *clear winner* as each of them offers different trade-offs between efficiency, reliability and complexity, depending on the considered scenario. In addition, we define a metric based on the ratio overhead over reliability, in order to compare the cost per service of each mechanism when a minimum delivery guarantee is required.

The rest of the paper is organized as follows. In Section II we present an overview of the mechanisms (legacy and novel) available for the transmission of video flows using IEEE 802.11. In Section III we detail the operation of a set of the mechanisms introduced by 802.11aa, namely, the Group Addressed Transmission Service, which constitutes the focus of this paper. In Section IV we present a performance assessment and comparison of the different schemes in terms of reliability, overhead and cost per service under a variety of scenarios, and for a given quality threshold. In Section V we summarize the related work, and in Section VI we present the conclusions from our study and our plans for future work.

II. VIDEO TRANSMISSION OVER IEEE 802.11 WLANS

In this section we summarize the main MAC mechanisms that can be used to support multimedia services over 802.11 WLANs, considering both the current standard and the new IEEE 802.11aa amendment.

A. Mechanisms with legacy 802.11

The choice of multicasting or unicasting audio and video content over IEEE 802.11 networks does not depend only on the number of receivers, since the multicast transmission in the IEEE 802.11 standard has some specificities that we consider in the following.

1) *Multicast service*: IEEE 802.11 defines layer-2 multicast as the transmission of data frames with a multicast address as the Destination Address. This multicast address accounts for a set of stations. The basic access procedure corresponds to the DCF mechanism, where no RTS/CTS exchange is used. In addition, no ACK is transmitted by any of the recipients of the frame. The lack of MAC-level recovery on multicast frames entails a reduced reliability of this kind of traffic, due to the increased probability of losing frames from interference or collisions. Moreover, the multicast frames must be transmitted at one of the rates included in the Basic Rate Set. This set is defined at the Access Point (AP) and includes the minimum set of rates that a station must support in order to join the BSS. Although it is not a requirement, usually the Basic Rate Set includes only rates with lower order modulations, so the transmission of multicast frames is performed at a reduced bit rate, hence decreasing the overall performance of the BSS but increasing reliability.

2) *Unicast service*: The second choice for transmitting audio/video frames is the use of unicast traffic. Unlike multicast, unicast traffic can be transmitted at any rate and it is acknowledged and retried, so its reliability is higher than standard multicast traffic. The counterpart is the amount of bandwidth required to transmit video to multiple receivers as the number of flows grows with the number of stations

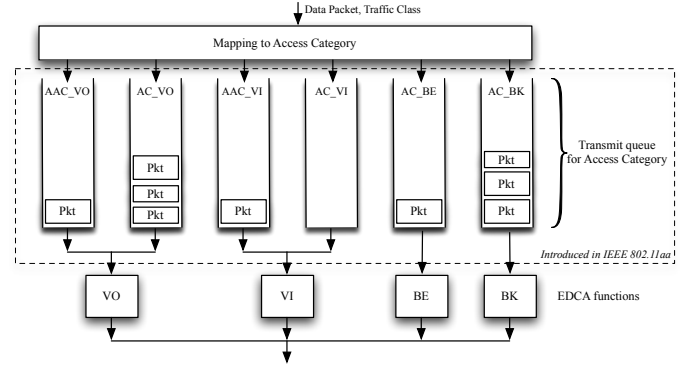


Fig. 1: Stream Classification Service over the standard EDCA mechanism.

receiving it, hence it is only suited for low bit rate flows and a reduced number of receiving stations.

3) *Prioritization*: With the introduction of 802.11e, multicast traffic can be differentiated, so the network can be configured to prioritize the multimedia flows while accessing the medium compared with the rest of traffic. Although this is an advantage, one of the problems with this approach is the definition of only one queue for video traffic, as current video codecs generate video flows as a sequence of frames that relate to one another, with different levels of relevance at the decoding stage, hence multimedia traffic may benefit from intra-traffic category differentiation.

B. Enhancements provided by 802.11aa

As explained in the previous section, the availability of mechanisms to perform video streaming over 802.11 WLANs is rather limited. The IEEE 802.11aa amendment has been designed to specifically address the challenge of multimedia transmission, implementing a set of new functionalities over the base specification. In this section, we describe the main novelties introduced by 802.11aa [9]. The objective of these extensions is to efficiently improve the reliability of audio and video streaming, while maintaining and even improving the service as perceived by the other streams. Following this, we consider the next mechanisms¹:

- Interworking with IEEE 802.1 Audio Video Bridging (AVB).
- Stream Classification Service (SCS).
- Group Addressed Transmission Service (GATS).

In the following we present the first two functionalities, while the third one, which is the focus of our work, is detailed in Section III.

1) *Interworking with IEEE 802.1AVB*: Within the 802.11aa standard the use of IEEE 802.1Qat [10], and specifically the Stream Reservation Protocol (SRP), is specified in order to enable end-to-end reservation of resources across IEEE 802

¹The IEEE 802.11aa amendment also specifies mechanisms to manage overlapping BSSs (OBSS), but for the purpose of this work we only consider the case of a single WLAN.

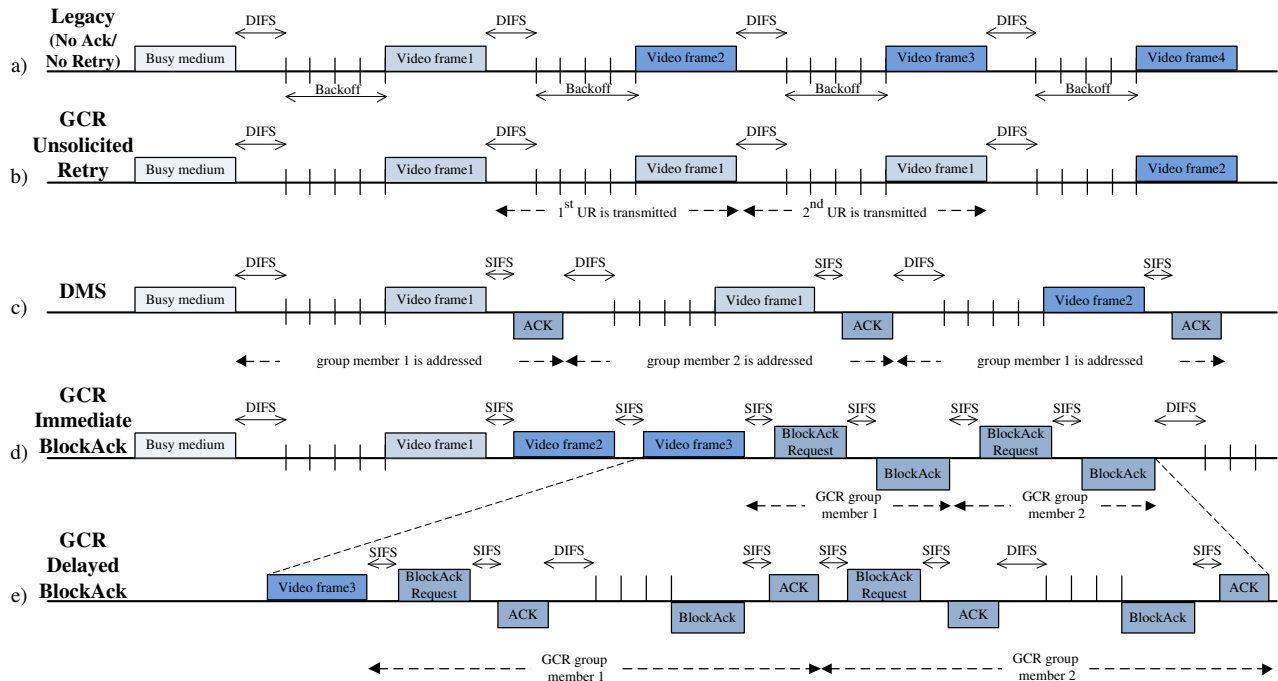


Fig. 2: Available mechanisms with the Group Addressed Transmission Service (GAST).

networks. Multimedia flows are the clear example of traffic that may take benefit of end-to-end resource reservation and as such it is envisioned to support and integrate the SRP.

It is important to note that the Stream Reservation Protocol is a *Higher Layer Protocol*. As such, the IEEE 802.11 system of a non-AP station is not able to interpret the SRP messages and its integration consists on the station sending the requests to the AP. In order to support end-to-end reservations on the wireless link, the AP has co-located a higher layer entity called Designated MSRP² Node (DMN) that is in charge of processing the signaling required by SRP and invoking the required 802.11 reservation primitives.

2) *Stream Classification Service*: This service is built upon the EDCA access mechanism, which is based on different access categories (ACs), and supports traffic differentiation by means of priorities between queues and heterogeneous configuration of the contention parameters (namely, *AIFS*, *CW* and *TXOP*). The SCS, along with the intra-Access Category (AC) Traffic Stream (TS) prioritization, enables classification using layer 2 and/or layer 3 signaling and increases the granularity of the service differentiation already provided by the EDCA mechanism.

More specifically, as illustrated in Fig. 1, this service introduces two additional queues within the existing EDCA access categories to support prioritization for audio and video streams, respectively. Furthermore, in addition to this intra-AC prioritization, packets are tagged with their drop eligibility

(DE), which defines a different maximum number of (short and long) retries.

With this availability of additional access categories and the drop eligibility bit, *graceful degradation* of video quality is supported in case of bandwidth shortage. However, it should be noted that the SCS does not alter the standard MAC behavior, as it builds on the existing EDCA mechanism. Therefore, it inherits the same limitations of the EDCA mechanism with respect to video streaming, being either: (i) extremely unreliable, in case of legacy multicast, or (ii) extremely inefficient, in case of unicast transmissions to a moderate number of receivers (as we will see in Section IV).

III. GROUP ADDRESSED TRANSMISSION SERVICE

Two of the main issues of multicast schemes defined by legacy 802.11, as described in Section II-A, are: (i) the poor reliability of the service, since multicast frames are neither acknowledged nor retried, and (ii) the high inefficiency due to the use of a low modulation and coding scheme, as it uses the largest rate available in the Basic Rate Set. In order to overcome these limitations, the IEEE 802.11aa standard extends the set of mechanisms available for multicast transmissions. These mechanisms include the Direct Multicast Service (DMS), which was first introduced in 802.11v [11] and is extended to target group addressed frames, and the Groupcast with Retries (GCR) service. The GCR service defines new retransmission schemes for *group* addressed frames, a group being a set of stations listening to the same (non-multicast) address, called *group concealment address*.

²MSRP: Multiple Stream Reservation Protocol

In the following, we describe the different mechanisms included in the 802.11aa standard for multicast service. We illustrate the operation of these mechanisms in Fig. 2, which also includes the *legacy* multicast service described in the previous section as reference.

A. GCR unsolicited retry (UR)

This delivery method *preemptively* retransmits a frame one or more times, to increase the delivery probability to the STAs (see Fig. 2b). The number of retransmissions to use is not specified in the 802.11aa standard, but stated as implementation-dependent. In this way, the mechanism aims at improving the reliability of the legacy multicast without introducing the associated overhead and complexity of a per-station acknowledgment mechanism.

B. Direct Multicast Service (DMS)

This mechanism performs, for every group addressed frame and destination, a standard unicast transmission (as illustrated in Fig. 2c for two groupcast members). The transmission follows the legacy procedure and the frames are retransmitted until an ACK is received, or the maximum retransmission limit is reached (and the frame is discarded). Although this mechanism provides very high reliability, it also poses scalability issues as the consumed resources linearly grow with the number of group members.

C. GCR Block Ack (BA)

This approach extends the Block Ack mechanism defined in the version of the standard [1] to account for group addressed frames. In this case the AP transmits no more than GCR buffer size frames to the GCR group address, and polls each destination STA with a *BlockAckRequest* frame. After receiving the *BlockAck* frame from one STA, the AP sends another *BlockAckRequest* frame to the next STA and so on, and then proceeds to send another burst of multicast frames.

GCR Block Ack comprises two mechanisms depending on the control exchange procedure: GCR Immediate Block Ack and GCR Delayed Block Ack. Within GCR Immediate Block Ack (BA-I for short), the recipient of a *BlockAckRequest* frame replies immediately (after a SIFS time) with a *BlockAck* frame, see Fig. 2d. In contrast, in the GCR Delayed Block Ack (BA-D for short), after receiving a *BlockAckRequest* frame the recipient starts a backoff process before sending the *BlockAck* frame (Fig. 2e). This backoff procedure was introduced to enable stations with low CPU capabilities to use this mechanism, as it gives them more time to compute CRC values and generate the ACK bitmap. With the GCR Delayed Block Ack, both *BlockAckRequest* and *BlockAck* frames are acknowledged with an ACK frame.

D. Parameters of the mechanisms

The mechanisms available with 802.11aa offer a variety of parameters to set, but the standard provides no guidelines towards their optimal configuration for any scenario (described in terms of, e.g., radio conditions, number of receivers, video

BW). For simplicity, we set the Modulation and Coding Scheme (MCS) of Data frames to $MCS_D=54$ Mbps, and for the case of Control and Management frames we set it to $MCS_C=24$ Mbps. We represent with \mathcal{M} the set of mechanism under study, with $\mathcal{M} \in \{L, UR_i, DMS, BA-I_M, BA-D_M\}$, which therefore consists on the following five schemes:

- L: Legacy multicast, transmitted at MCS_C^3 .
- UR_R : UR with R retries ($R \geq 0$) at MCS_D .
- DMS: at MCS_D .
- $BA-I_M$: with a maximum burst of M frames, at MCS_D .
- $BA-D_M$: with a maximum burst of M frames, at MCS_D .

IV. PERFORMANCE EVALUATION

To perform the evaluation of the new multicast mechanisms, we use the OMNeT++ simulation framework,⁴ configured with the 802.11e EDCA module, which we extend to support the new schemes. For all results we present the average of 10 simulation runs, and we confirmed that 95% confidence intervals were below 1% of the average.

A. Scenario and channel model

Our WLAN scenario consists on one AP sending a video flow towards K destinations, using the 802.11g PHY layer. Our video flow is based on a trace-file from a movie clip of 2-minute duration encoded with the H.264 standard [12], one of the most popular video codecs [13]. Its spatial resolution is fixed to 720x480, using a variable q parameter with constant bit rate of 3.91 Mbps. We refer to the direction from the AP (stations) to the stations (AP) as downlink (uplink). Stations are 40 m away from the AP and configured with a transmission power of 20 dBm. Our channel model closely follows the OFDM model described in [14].

In order to properly analyze the robustness and cost of the use of the 802.11aa mechanisms under a wide variety of scenarios, in addition to the above *channel losses* we also introduce, following [15], *interference losses*, to account for the impact of, e.g., hidden nodes. Based on this, in our simulations we will provide results based on the per-frame loss probability p , which accounts for both radio losses (with probability p_R) and interference losses (with probability p_I), as follows:

$$p = 1 - (1 - p_R)(1 - p_I).$$

B. Open-loop schemes

Here we analyze the performance of those mechanisms that do not require traffic in the uplink direction, which we refer to as *open-loop* schemes (i.e., $\mathcal{M} \in \{L, UR_i\}$). We first analyze their *reliability*, i.e., how good they are at successfully delivering frames. More specifically, with K as the set of stations, if N frames were sent in the WLAN and out of them the subset N_i were successfully received by station i , the reliability of the scheme $\rho_{\mathcal{M}}$ can be computed as:

$$\rho = \frac{\frac{1}{K} \sum_{i=1}^K N_i}{N}, \quad (1)$$

³This is the maximum MCS available in the Basic Rate Set.

⁴<http://www.omnetpp.org>

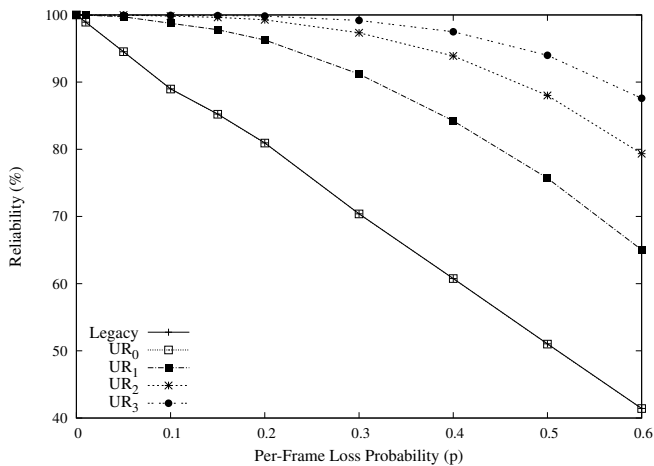


Fig. 3: Reliability of the open-loop schemes.

which can be interpreted as the probability that a given video frame is successfully delivered to all stations. We plot the reliability of the considered open-loop schemes for a 5-receiver scenario (i.e., $K = 5$) and different values of p in Fig. 3.

We use solid lines to represent the legacy mechanism, and dotted lines for the GCR unsolicited retry mechanism.⁵ The results confirm that the legacy mechanism has the worst reliability, which, as expected, goes with $1 - p$ (e.g., a 60% reliability for $p = 0.4$). In contrast, the UR_i mechanisms improve reliability as i increases. Indeed, for the same value of $p = 0.4$, with UR_3 reliability is practically 100%. Even for the case of only one retransmission (UR_1), the reliability is improved by approximately 25% on average.

It should be noted that the above discussion does not take into account the *price to pay* to obtain a given reliability figure. In order to perform a *fair* comparison between the schemes, we need to take into account their corresponding *cost* in terms of resource consumption, as it is evident that, e.g., for good channel conditions, the legacy scheme will incur into a huge resource wastage due to the overly robustness of the MCS, or with UR_3 most of retransmissions will be unnecessary.

To quantify resource wastage, we define the *overhead* η as the ratio between the channel time required by the mechanism \mathcal{M} to send the frames over the minimum channel time that would be required to send them (regardless of the ρ achieved). Let $\mathcal{T}_{\mathcal{M}}$ denote the channel time occupied by \mathcal{M} to deliver the frames. Given that the GCR unsolicited retry with $R = 0$ (UR_0) is the scheme that consumes the least amount of resources, we can compute the overhead as:

$$\eta = \frac{\mathcal{T}_{\mathcal{M}}}{\mathcal{T}_{UR_0}}, \quad (2)$$

where we compute \mathcal{T} by adding the time that the medium is busy (re)transmitting data and control frames.⁶

⁵Note that for every p value the case UR_0 matches the legacy mechanism in terms of reliability.

⁶Note that, the overhead cost is not the only cost that can be associated with a given scheme, as there are other costs, e.g., complexity/CPU, memory. We postpone the discussion on those until Section IV-E.

TABLE I: Overhead of the open-loop schemes.

	UR_0	UR_1	UR_2	UR_3	L
η	1.00	2.00	3.00	4.00	2.25

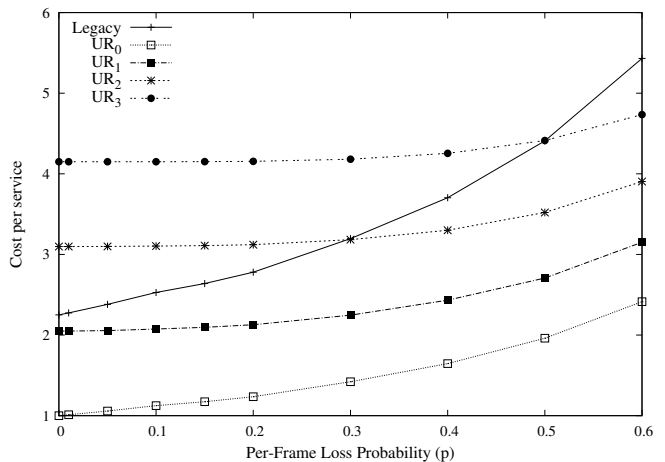


Fig. 4: Cost per service (η/ρ) for the open-loop schemes.

We report in Table I the overhead values of the open-loop schemes, which do not depend on the number of receivers (in contrast to the mechanisms of the next section) and, as expected, grow linearly with the number of retries. Given the MCS considered (54 Mbps and 24 Mbps for data and legacy, respectively), in this case we have that with two retries (i.e., UR_3) the scheme already consumes more wireless resources than the legacy case, despite its low MCS. We confirm that the overhead of the GCR unsolicited retry for an $R = 1, 2, 3$, are two, three and four times more than case without retries (UR_0).

With the above, we have on the one hand ρ to quantify how good a mechanism is at delivering data, and on the other hand, η to quantify how much it costs to use it. In order to better compare the schemes, we introduce the ratio η/ρ as the *cost per service* of a scheme, i.e., the amount of resources consumed per each unit of reliability obtained. The results are depicted in Fig. 4.

The results show that the legacy scheme (solid line) provides a cost per service that notably increases as p increases. The reason is that although the overhead η is fixed, the reliability ρ decreases with p , and therefore η/ρ grows. The same qualitative behaviour is obtained for the UR_i schemes, although in these cases the decrease of ρ is less steep, therefore the cost per service is flatter. The figure also serves to illustrate that depending on the number of retries configured, the cost per service could be higher or smaller than the one with the legacy scheme, which illustrates the flexibility introduced by 802.11aa and motivates the design of mechanisms to adequately tune this parameter. Finally, it should be noted that this (η/ρ) variable provides an adequate metric on the efficiency of a given scheme, but it does not take into account the actual performance experienced by a video flow. In Section IV-E we will take this into account when comparing open and closed-

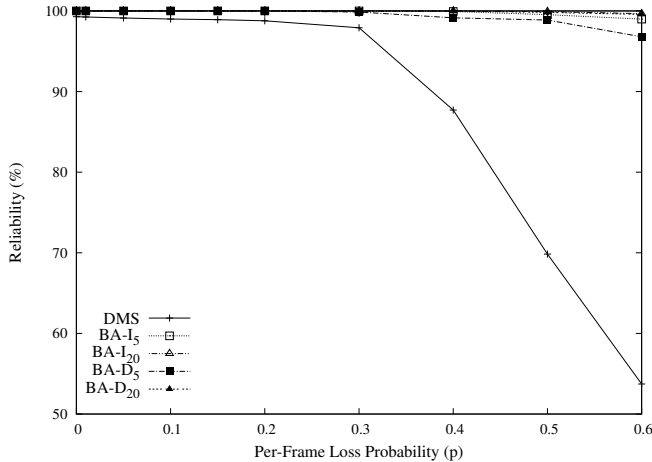


Fig. 5: Reliability of the closed-loop schemes.

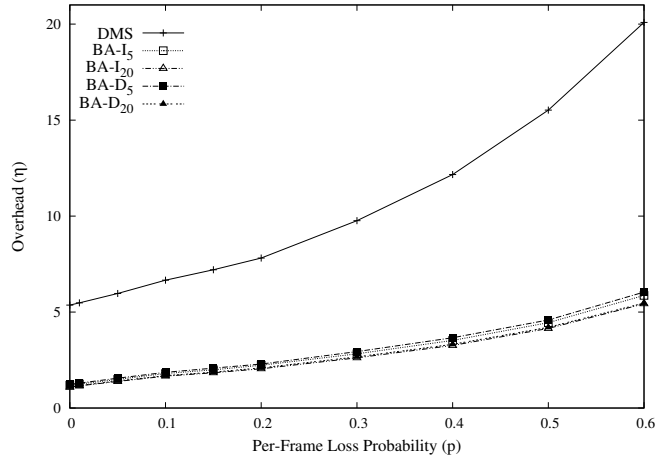


Fig. 6: Efficiency of the closed-loop schemes.

loop schemes.

C. Closed-loop schemes

In this section we analyze the performance of those mechanisms based on acknowledgments, i.e., $\mathcal{M} \in \{\text{DMS}, \text{BA-I}_M, \text{BA-D}_M\}$, whose performance therefore depends on the number of receivers K . We start our evaluation with the computation of the reliability ρ for the same WLAN conditions as in the previous section ($K = 5$). It should be noted that, with this schemes, and in particular with DMS, the need to retransmit a frame until it is acknowledge can severely choke the available bandwidth, up to a point where losses are also caused by frame drops at the output queue (we limit this queue to 200 frames).

We depict in Fig. 5 the results for the reliability ρ of the closed loop schemes. We use solid lines to represent the DMS mechanism, and dotted lines for all GCR Block Ack mechanisms, whose performance in terms of ρ is very similar for the considered cases (note that for both schemes we use $M = \{5, 20\}$, to understand the impact of the maximum frame burst). Compared to the open-loop results in Fig. 3, in this case ρ figures are much higher in general, thanks to the feedback from the stations. For the case of DMS, reliability is well above 95% until the frame loss probability reaches 0.3 –from this point on, drops at the output queue dominate the loss process. For the case of Block ACK schemes, reliability is practically 100% for all p values below 0.5, and it decreases for the case of the delayed scheme with $M = 5$ to approximately 97% when $p = 0.6$.

We next compute the overhead η of each mechanism for the same scenario, with the results illustrated in Fig. 6. As expected, the DMS mechanism presents a significant overhead, which is at least five times larger than the UR_0 (note that in this scenario $K = 5$). The GCR Block Ack mechanisms, in contrast, presents a significantly smaller overhead and very similar for all configurations (a closer look to the figure reveals that $\text{BA} - \text{I}_{20}$ shows the smallest overhead).

We next compute the cost per service η/ρ for the same

scenario ($K = 5$) and present the results in Fig. 7a. The figure presents the same qualitative results as the ones obtained for η , which is caused by the relatively constant η values.⁷ Furthermore, it also confirms that for $K = 5$, there are little differences between the considered Block Ack schemes.

In order to understand the differences between Block Ack schemes, we provide the (η/ρ) values for $K = 10$ and $K = 20$ receivers in Fig. 7b and Fig. 7c, respectively (note that we only provide the DMS values for $K = 10$, due to its scalability issues). The figures show that, as the number of receivers increases, differences arises between BA schemes: the delayed version with a maximum burst of $M = 5$ frames provides the largest cost per service, while the immediate version with $M = 20$ provides the smallest. Furthermore, from the results and configurations considered, it seems that the choice of the maximum burst length M has a more significant impact that the difference between the immediate and the delayed versions of the Block Ack schemes.

D. Comparing open and closed-loop mechanisms

Here we evaluate the efficiency of all considered schemes when a minimum QoS is guaranteed. More specifically, we set a threshold on the reliability ρ , and compare the cost per service of each mechanism as long as their performance is above this threshold. Following [16], [17], which report that video quality is severely degraded when loss rates are above 5-10%, we fix 10% as the minimum requirement on reliability. Based on this, we plot in Fig. 8 the cost per service of the considered schemes. For the sake of clarity we do not depict in the figure the case of DMS, whose cost is significantly higher than any other scheme, nor the case of Delayed Block Ack, given that its performance is very similar to the one obtained with the Immediate version.⁸

⁷It is worth noting that with open-loop schemes we have a fixed overhead, while in this case we have a relatively fixed reliability.

⁸As explained before, the choice between the Delayed or the Immediate version will mostly depend on hardware constraints.

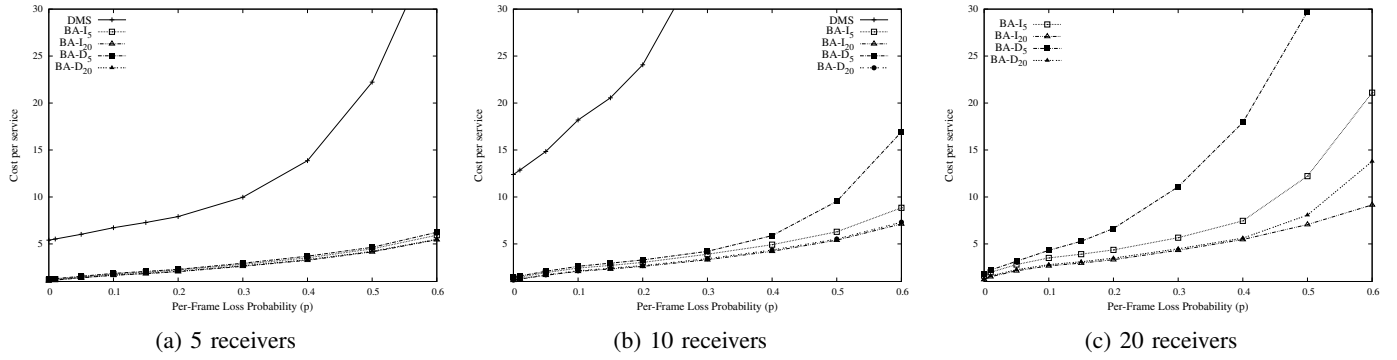


Fig. 7: Cost per service (η/ρ) for the closed-loop schemes for $K=5, 10$ and 20 .

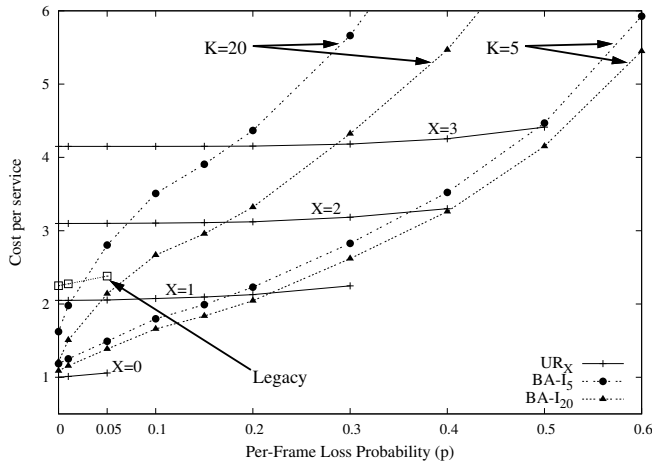


Fig. 8: Cost per service (η/ρ) for multicast schemes for $\rho \geq 90\%$.

We can split the analysis of the results in Fig. 8 in two parts, namely, when losses are very small ($p \leq 0.05$) and when losses are significant ($p > 0.05$). For the first case, UR_0 achieves the minimum cost per service, and in contrast to the Block Ack schemes, its performance is independent of the number of receivers K . Indeed, BA schemes are *second best* when $K = 5$, but for $p \leq 0.05$ it is interesting to observe that UR_1 achieves a smaller cost per service than BA when the maximum burst size $M = 5$ and $K = 20$, and that even the legacy scheme is more efficient than BA in that case.

As losses become significant, the UR scheme has to increase the number of retries to provide the required reliability, up to a certain point where it cannot achieve $\rho > 90\%$. For instance, UR_1 provides the best performance for $p = 0.3$, but from this point on, it cannot guarantee enough reliability. Similarly, for $p = 0.4$ and $p = 0.5$ the cost per service of UR_2 and UR_3 is very similar the minimum one, which confirms its good efficiency in terms of cost per service, but beyond these probabilities the video flow suffers from severe losses. These results confirm that a key issue when using UR is to properly tune the number of retransmissions, as depending on WLAN conditions this could result in a very inefficient scheme (e.g., UR_3 when $p < 0.1$) or one of the most efficient ones (e.g.,

UR_2 for $p \approx 0.1$), which in addition has a performance that does not depend on the number of receivers.

Considering the Block Ack schemes, the figure shows that for the case of a small number of receivers ($K = 5$) their cost per service is relatively small, being able to provide enough reliability for all the considered values of p . When the number of receivers increases, the overhead associated with the polling scheme becomes more evident, leading to the largest values of (η/ρ) for $p \geq 0.2$.⁹ Furthermore, in these circumstances (i.e., when the control traffic is significant) the maximum burst length M has a notable impact on performance, as opposed to the case of $K = 5$ receivers, where there are little differences between the two M configurations.

The above results quantify the differences in terms of performance between the multicast schemes. We next summarize qualitatively their differences, discussing the pro's and con's of each scheme.

E. Summary and discussion

Here we summarize the main results from our simulation-based assessment while taking into account the qualitative differences between the mechanisms, both in terms of their ease of use (in case, e.g., they require a proper tuning of a parameter to achieve good performance) and the complexity introduced.

- The *legacy* mechanism, as expected, is the approach that provides in general the poorest reliability, and given its low MCS, it also provides very low efficiency values. However, it should be noted that in terms of cost per service, it can outperform very aggressive configurations of UR_i , and that for good WLAN conditions¹⁰ it can outperform Block Ack schemes when the number of receivers is high. Given its non-existent implementation complexity, we believe that still it cannot be discarded as a valuable scheme to transmit multicast traffic, although its use is limited to very specific conditions: large number of receivers, some with weak radio links as they require a low MCS, and low traffic activity in the WLAN.

⁹Note that the DMS mechanism is not considered due to its high costs.

¹⁰We remark here that p accounts for both radio conditions and collision from other traffic sources.

TABLE II: Overview of the schemes available with 802.11aa.

Scheme	Complexity	Reliability	Overhead	Cost per service	Scalability	Target scenario
Legacy	Low	Low	Medium	Medium	High	Legacy stations, good WLAN conditions
UR _R	Medium	Medium ^a	Low ^a	Low	High	Large number of receivers, relatively good WLAN conditions
DMS	Medium	High	High	High	Low	Good WLAN conditions, small number of receivers
BA-I _M	High	High	Medium ^b	Low	Medium	Moderate number of receivers, average WLAN conditions
BA-D _M	High	High	Medium ^b	Medium	Medium	Same as BA-I, interfaces with limited capabilities

^a Depending on the number of retries

^b Depending on the number of receivers

- The *GCR unsolicited retry* mechanism offers a relatively large variety in its performance, which obviously depends on the number of retries configured but not in the number of receivers. The latter is a major advantage of the scheme, in contrast to closed-loop mechanisms, and we have seen that even for moderate frame loss probabilities its performance can be as good as the one obtained with BA (for 5 receivers) or significantly better (for 20 receivers). Although the mechanism does not introduce significant overhead in terms of implementation, its main weakness is that it requires a proper tuning of the number of retries, as otherwise the above good features are difficult to achieve.
- The *Directed Multicast Service* mechanism poses severe scalability issues, as its overhead is at least the number of receivers, and greatly increases with frame losses due to retransmissions, causing frame drops at the output queue. Its implementation complexity is, on the other hand, moderate, the unicast mechanism is already available with existing interfaces and the only modification is basically to replicate each frame towards each intended destination. Based on this, the DMS is only suited for scenarios with small number of receivers and losses, and when the required reliability is relatively high.
- The *GCR Block Ack* mechanisms achieve the high reliability of DMS without its enormous overhead, at the expense of a high implementation complexity (both at the AP and the stations). Indeed, the AP has to keep a copy of all frames as long as they are not acknowledged, and the stations have to compute the CRC for each frame and send the feedback in the Block ACK, which requires not only a new control exchange but also relatively high computation capabilities –if these are not available, the delayed version lessens this requirement. This complexity supports a good performance in terms of cost per service, although when the number of receivers is large, it can be outperformed by simpler mechanisms (properly configured).

The main highlights of the above discussion are presented in Table II, in which we compare for every scheme, its main features in terms of complexity, reliability, overhead and cost per service, and also we provide the target scenario. In this way, one of the main conclusions of our work is that there is no *clear winner* among the mechanisms available with 802.11aa, as each of them offers a different trade-off in terms of performance and complexity, and therefore their

use depends on the target scenario and its parameters.

V. RELATED WORK

Most of the related work present in the literature focuses on finding new mechanisms or enhancements for the multicast transmission in IEEE 802.11 WLANs, motivated by the low tolerance to losses and lack of retransmissions of the legacy scheme. Some of these works [18], [19] propose the use of forward error correction (FEC) codes, or multi-rate adaptation based on the feedback obtained from the stations.

A significant amount of work builds on the Leader Based Protocols (LBP). With LBP [3], one station is selected as the leader of the multicast group. This leader will send feedback to the AP with a positive acknowledgement (ACK) on behalf of the stations in the group, although they can send negative acknowledgements (NACK) to notify the reception of incorrect frames. An enhancement by [20] proposes the Enhanced Leader Based Protocol (ELBP), which is a hybrid between the Batch Mode Multicast MAC protocol (BMMM) [21] and LBP protocols, similar to the GCR solicited mechanism specified in 802.11aa, by using the Block Ack mechanism from 802.11e. In a follow-up work [22], the authors extend the previous work by proposing two mechanisms for leader selection. Authors in [23] advocate for the use of propagation-related metrics for leader selection (i.e., the station with worst link quality), supporting some results with experiments, although there are some practical issues due to synchronization and the capture effect. In [24], the authors propose a combination of LBP with adaptive transmission rate, relying on the feedback from the stations and combined with a CTS-to-Self so that the multicast receivers will be aware of any pending transmission. Authors in [25] define a hybrid model that combines an LBP approach, adaptive rate transmission and tuned power control based on the positive and negative feedback jamming ratio.

Apart from LBP-based proposals, in [26] authors propose to use *tones* to transmit the feedback information, i.e., NACKs or Negative CTS (NCTS), so the number of collisions gets reduced. However, this solution requires a dedicated signaling channel that is implemented with two wireless interfaces. Other proposed works mix schemes along the lines of DMS: in [27], authors propose to transmit the video I-frames using unicast, while the rest of video frames are transmitted using legacy multicast. Similarly to DMS, this proposal poses scalability issues.

Based on the above, most of the existing work has proposed non-standard schemes to improve video performance. Apart

from the qualitative description in [28], only [29] presents some numerical figures on the performance of the recent 802.11aa standard. However, in contrast to our work, in that paper authors assume an 802.11b scenario with a low-bandwidth video, they only consider a subset of the existing 802.11aa mechanisms, and the performance evaluation does not address the efficiency of the schemes (in particular, for a minimum guarantee required).

VI. CONCLUSIONS AND FUTURE WORK

In this paper, we have performed an evaluation of the novel mechanisms available with IEEE 802.11aa in terms of reliability, overhead and cost per service, and compared them to one another and against the legacy multicast service. We have confirmed that the new mechanisms significantly enrich the available choices to deliver video over WLANs, as they provide different trade-offs in terms of complexity, efficiency and reliability. We have identified the main limiting factors of each mechanism, which paves the way for the derivation of guidelines to optimally configure the parameters of each mechanism and to select the best suited for a given scenario.

We are currently developing an analytical model for the performance of each mechanisms to derive these guidelines, and we are also implementing them in a prototype using commercial off-the-shelf devices (COTS) devices. We plan to perform a similar evaluation to the one carried out in this paper in the testbed, and to validate the results from the analytical model against experimental figures.

ACKNOWLEDGMENTS

This work has been supported by the European Community's Seventh Framework Programme (FP7-ICT-2009-5) under grant agreement n. 258053 (MEDIEVAL project) as well as by the Spanish Government, MICINN, under research grant agreement TIN2010-20136-C03.

REFERENCES

- [1] *IEEE 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*, IEEE Std. 802.11, 2012.
- [2] *Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications - Amendment 5: Enhancements for Higher Throughput*, IEEE Amendment 802.11n, 2009.
- [3] J. Kuri and S. Kaser, "Reliable Multicast in Multi-access Wireless LANs," in *INFOCOM '99. Eighteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, vol. 2, Mar. 1999, pp. 760–767 vol.2.
- [4] *Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications - Amendment 2: Higher-speed Physical Layer (PHY) extension in the 2.4 GHz band*, IEEE Amendment 802.11b, 2001.
- [5] *IEEE P802.11ac/D4.1 Draft Standard - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications - 4: Enhancements for Very High Throughput for Operation in Bands below 6 GHz*, IEEE Amendment 802.11ac, 2012.
- [6] *IEEE P802.11ad/D9.00 Draft Standard - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications - Amendment 3: Enhancements for Very High Throughput in the 60 GHz Band*, IEEE Amendment 802.11ad, 2012.
- [7] *Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications - Amendment 8: Medium Access Control (MAC) Quality of Service Enhancements*, IEEE Amendment 802.11e, 2005.
- [8] P. Serrano, A. Banchs, P. Patras and A. Azcorra, "Optimal Configuration of 802.11e EDCA for Real-Time and Data Traffic," *Vehicular Technology, IEEE Transactions on*, vol. 59, no. 5, pp. 2511–2528, Jun. 2010.
- [9] *Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications - Amendment 3: MAC Enhancements for Robust Audio Video Streaming*, IEEE Amendment 802.11aa, 2012.
- [10] *IEEE 802.1Qat/D6.1 Draft Standard for Information Technology-Telecommunications and information exchange between systems-Local and metropolitan area networks: Virtual Bridged Local Area Networks - Amendment 9: Stream Reservation Protocol (SRP)*, IEEE Std. 802.1Qat, 2010.
- [11] *Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications - Amendment 8: IEEE 802.11 Wireless Network Management*, IEEE Amendment 802.11v, 2011.
- [12] "ITU-T Recommendation H.264 : Advanced video coding for generic audiovisual services," International Telecommunications Union, Nov. 2007.
- [13] P. Seeling and M. Reisslein, "Video Transport Evaluation With H.264 Video Traces," *IEEE Communications Surveys and Tutorials*, in print, vol. 14, no. 4, pp. 1142–1165, 2012.
- [14] G. Pei and T.-R. Henderson., "Validation of the OFDM error rate model in ns-3," 2010.
- [15] M. Garetto, T. Salonidis and E.-W. Knightly, "Modeling Per-flow Throughput and Capturing Starvation in CSMA Multi-hop Wireless Networks," *IEEE/ACM Trans. Netw.*, vol. 16, no. 4, Aug. 2008.
- [16] T. Stockhammer, M.-M. Hannuksela and T. Wiegand, "H.264/AVC in Wireless Environments," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 13, pp. 657–673, 2003.
- [17] T. Schierl, H. Schwarz, D. Marpe, and T. Wiegand, "Wireless Broadcasting Using the Scalable Extension of H.264/AVC," in *ICME. IEEE*, 2005, pp. 884–887.
- [18] O. Alay, T. Korakis, Y. Wang, and S. Panwar, "Dynamic Rate and FEC Adaptation for Video Multicast in Multi-rate Wireless Networks," *Mob. Netw. Appl.*, vol. 15, no. 3, pp. 425–434, Jun. 2010.
- [19] C. Xi, W. YunHeng and L. JianHua, "Cross-layer Link Rate Adaptation for High Performance Multimedia Broadcast over WLANs," in *GLOBECOM Workshops (GC Wkshps), 2010 IEEE*, Dec. 2010, pp. 965–969.
- [20] A. Lyakhov, V. Vishnevsky and M. Yakimov, "Multicast QoS Support in IEEE 802.11 WLANs," in *IEEE 4th International Conference on Mobile Adhoc and Sensor Systems, MASS 2007, 8-11 Oct. 2007, Pisa, Italy. IEEE*, 2007.
- [21] M.-T. Sun, L. Huang, A. Arora, and T.-H. Lai, "Reliable MAC Layer Multicast in IEEE 802.11 Wireless Networks," in *Proceedings of the 2002 International Conference on Parallel Processing*, ser. ICPP '02. Washington, DC, USA: IEEE Computer Society, 2002, pp. 527–.
- [22] A. Lyakhov and M. Yakimov, "Analytical Study of QoS-Oriented Multicast in Wireless Networks," *EURASIP J. Wireless Comm. and Networking*, vol. 2011, 2011.
- [23] J. Miroll, Z. Li and T. Herfet, "Wireless Feedback Cancellation for Leader-Based MAC Layer Multicast Protocols," in *IEEE International Symposium on Consumer Electronics (ISCE)*, 2010.
- [24] S. Choi, N. Choi, Y. Seok, T. Kwon, and Y. Choi, "Leader-Based Rate Adaptive Multicasting for Wireless LANs," in *Global Telecommunications Conference, 2007. GLOBECOM '07. IEEE*, Nov. 2007.
- [25] J. Miroll and T. Herfet, "Efficient OFDM-based WLAN-multicast with feedback aggregation, power control and rate adaptation," *Computer Communications*, vol. 35, no. 18, pp. 2245–2253, 2012.
- [26] S. Gupta, V. Shankar, and S. Lalwani, "Reliable Multicast MAC Protocol for Wireless LANs," in *Communications, 2003. ICC '03. IEEE International Conference on*, vol. 1, May 2003, pp. 93 – 97 vol.1.
- [27] Zhengyong Feng and Guangjun Wen and Chunlei Yin and Hongsheng Liu, "Video Stream Groupcast Optimization in WLAN," in *International Conference on Internet Technology and Applications*, 2010.
- [28] K. Maraslis, P. Chatzimisios, and A. Boucouvalas, "802.11aa: Improvements on Video Transmission over Wireless LANs," in *Communications (ICC), 2012 IEEE International Conference on*, 2012.
- [29] M. A. Santos, J. Villalon, and L. Orozco-Barbosa, "Evaluation of the IEEE 802.11aa Group Addressed Service for Robust Audio-Video Streaming," in *Communications (ICC), 2012 IEEE International Conference on*, Jun. 2012, pp. 6879–6884.