Efficient Real–time Content Distribution for Multiple Multicast Groups in SDN–based WLANs

Estefanía Coronado*, Roberto Riggio‡, José Villalón*, Antonio Garrido*
*High-Performance Networks and Architectures (RAAP), University of Castilla–La Mancha, Albacete, Spain
Email: {Estefania.Coronado, JoseMiguel.Villalon, Antonio.Garrido}@uclm.es
‡FBK CREATE-NET, Trento, Italy
Email: rriggio@fbk.eu

***Abstract*—Wireless networks research and development efforts are largely driven by the increasing interest in multimedia applications. Video streaming services, which often involve strict Quality of Service (QoS) requirements and are very sensitive to delays, represent a significant proportion of these applications. In IEEE 802.11–based WLANs, these services raise several challenges in terms of robustness, reliability and scalability, specially when supporting multiple multicast streams at the same time. Nevertheless, traditional network architectures make it difficult to address these problems. In this context, the Software Defined Networking (SDN) paradigm opens new research possibilities by decoupling the control decisions from the data–plane and by improving network management and programmability. In this paper, we present *SM-SDN@Play*, an SDN–based solution for joint multicast rate selection and group formation in 802.11–based networks. Experimental results show the high performance and reliability capabilities of the scheme, regardless of the application bitrate, the number of clients, and the number of concurrent multicast streams. Furthermore, the channel utilization is greatly reduced with regard to the standard multicast schemes, which allows other applications to be supported without experiencing a performance degradation. We release the entire software implementation under a permissive APACHE 2.0 license for academic use.**

***Keywords*—Software Defined Networking, WLANs, IEEE 802.11, multicast, rate adaptation, multimedia, video distribution.**

I. Introduction

The emergence of platforms such as Netflix and Youtube has made multimedia content distribution a popular service in the recent years. Furthermore, it is becoming more common that real–time events such as conferences, social events or educational courses are simultaneously transmitted to a wide range of users. In view of this scenario, multicast communications represent an efficient way of delivering the same information to multiple destinations in a scalable fashion.

The IEEE 802.11 [1] standard is one of the most widespread technologies for the deployment of Wireless Local Area Networks (WLANs) and is found in domestic and professional settings such as enterprises, campuses and hotels. Nevertheless, multicast communications over 802.11–based WLANs incur in severe reliability issues. In fact, due to the lack of acknowledgements and retransmissions, multicast transmissions in 802.11 are performed using the basic Modulation and Coding Scheme (MCS) which results in a high channel occupation. This issue is exacerbated as the applications bitrate increases, and becomes worse in the cases of multiple simultaneous multicast streams. To address these reliability concerns, the IEEE 802.11aa [2] amendment introduces a set of multicast retransmission policies. Nevertheless, no mechanisms for the delivery rate adaptation and multicast group management are specified by the standard. Moreover, due to the widespread use of Wi–Fi compatible devices, IEEE 802.11 amendments aim at maximizing backward compatibility at the expense of innovation. In view of this, Software Defined Networking (SDN) changes the traditional network architecture by effectively decoupling the data–plane from the control–plane and by providing network developers with powerful programming abstractions to affect the state of the network.

The contribution of this paper is twofold. First, we take advantage of our previous work *SDN@Play* [3], a multicast MCS selection algorithm, in order to propose a novel multicast group management scheme. This new scheme, named *Scalable Multigroup SDN@Play* (*SM-SDN@Play*), jointly drives the multicast MCS selection and the multicast group formation in order to minimize the network–wide airtime utilization and maximize the multicast services reliability. Second, we implement and test *SM-SDN@Play* over a real world 802.11–based WLAN and we release the entire implementation under a permissive APACHE 2.0 license for academic use[1].

This work extends [3] in three ways. First, as opposed to *SDN@Play*, the solution presented in this paper independently selects the optimal MCS for each multicast group. Second, *SDN@Play* introduced a two–phase algorithm alternating unicast and multicast periods, however the duration of such periods was static. *SM-SDN@Play*, on the other hand, dynamically adapts the duration of the unicast and multicast periods according to the number of active multicast groups. Third, *SM-SDN@Play* distributes the unicast periods of each multicast group in such a way to minimize the chances that multiple multicast group will operate in unicast mode at the same time. Experimental results show that *SM-SDN@Play* outperforms the standard IEEE 802.11 multicast schemes in terms of both throughput and channel utilization without requiring any change to the wireless clients.

The remainder of this paper is structured as follows. Section II introduces the technical background on multicast communications in 802.11 WLANs. Section III provides an

[1] http://empower.create-net.org/

overview of most relevant related work. In Section IV we introduce the design of *SM-SDN@Play*, while in Section V the implementation details are presented. The results of the measurements campaign are discussed in Section VI. Finally, Section VII concludes the paper pointing out the future work.

II. Technical Background

Multicast transmissions are an efficient way to send the same data to many wireless clients. However, in IEEE 802.11, multicast services are specified as a simple broadcasting mechanism that does not make use of Acknowledgment (ACK) frames. As a result, multicast transmissions are usually performed at the lowest MCS (in order to increase both the range and the reliability of the transmission) and do not use any form of transmission feedback mechanism.

This problem is partially addressed by the IEEE 802.11v amendment [4], where the Direct Multicast Service (DMS) is introduced. DMS replicates each multicast frame into as many unicast frames as the number of receptors in a multicast group. In this way, each frame is retransmitted as many times as required until the Access Point (AP) receives the ACK or the retransmission counter reaches the limit. Although this approach ensures the same reliability level of a unicast transmission, it also presents serious scalability issues as the number of stations in a multicast group increases.

To partially address this scalability limitation, the IEEE 802.11aa amendment [2] introduces the Group Addressed Transmission Service. An in–depth analysis of the performance of this service is carried out by Daldoul *et al.* [5]. The Group Addressed Transmission Service is composed of two mechanisms: DMS and Groupcast with Retries. The latter defines three retransmission methods: Legacy multicast, Unsolicited Retries (UR) and Block ACK (BACK). Legacy multicast is the multicast mode defined in the original IEEE 802.11 standard. Unsolicited Retries specifies a number of retry attempts, N, so that a frame is transmitted $N+1$ times. In spite of increasing the frame delivery probability, this method reduces the network performance due to the retransmission of unnecessary frames. Furthermore, although the stations do not require acknowledgments, this mechanisms still suffers from scalability issues. In the Block ACK method, the AP agrees with the stations the number of consecutive unacknowledged frames. After that, the AP sends a burst of multicast frames up to that number, and requests the Block ACK to each station. Both the request and the ACKs are sent in unicast mode. Although the control traffic is reduced with regard to DMS, the scalability degree of this scheme is also limited.

III. Related Work

The lack of ACKs makes multicast frames in 802.11 to be transmitted at the basic MCS. In this regard, the channel congestion and the QoS restrictions mainly determine the data rate that is selected in most of the related proposals. In spite of achieving higher transmission rates, the performance improvement of all these works usually depends on the size of the multicast group and may suffer from scalability issues. Moreover, many of these works require significant modifications to the wireless clients stack, making them incompatible with the IEEE 802.11 standard. In this section we will summarize such works pointing out in which way our solution improves multicast communications with multiple multicast groups.

Feedback gathering from the stations can be carried out through leader–based schemes. J. Kuri *et al.* [6] and D. Dujovne *et al.* [7] seek to improve the transmission reliability by enabling ACKs for the group leader, which is selected as the receptor exhibiting the worst signal quality. However, a procedure for the leader selection is not provided. Signal–to–Noise Ratio (SNR) in combination with leader–based works have been widely used in the literature. The Auto Rate Selection Multicast mechanism [8] selects the multicast group leader during the first part of the algorithm, while in the second one the SNR obtained from the leader ACKs is considered to adapt the data rate. The SNR–based Auto Rate for Multicast algorithm [9] makes the AP periodically send beacons frames to the multicast stations with the aim of figuring out from their responses the perceived SNR level. Based on this information, the transmission rate is adapted according to the selected leader, which corresponds with the client exhibiting the worst SNR value. Lastly, the Hierarchical Auto Rate Selection Multicast mechanism [10] ensures that the clients under the worst channel conditions receive, at least, the base layer of the video, while the remaining ones also receive some enhancement layers. However, most of these approaches either require to make changes in the 802.11 standard, or they do not specify a procedure for the leader election or need to reach a trade–off between reliability and scalability.

Mathematical and analytical models have been also taken as reference to improve the performance of the multicast transmissions. M. Sun *et al.* [11] propose an analytical model to perform a multicast scheduling by gathering the channel state information and the quality of each Scalable Video Coding layer. The Batch Mode Multicast MAC scheme [12] enhances the network reliability by polling the receptors to obtain individual ACKs, which makes it not scalable to large multicast groups. The Enhanced Leader Based Protocol [13] relies on the use of multiple leaders for the ACKs handling and the Block ACK techniques. However, analytical models are usually applied on a saturated network and make assumptions that are not always met on a real–world scenario.

Research efforts on SDN–based multicasting in OpenFlow [14] networks can also be found. L. Bondan *et al.* [15] introduced a solution for multimedia multicasting based on OpenFlow which aims at calculating the best route between the multicast source and the destinations. Similarly, H. Egilmez *et al.* [16] also aim at enhancing multicast video transmissions by enabling the QoS support at the OpenFlow control layer. The reliability of the multicast traffic is also improved in ECast [17] by means of a novel packet retransmission scheme for packet loss mitigation. OpenFlow is also used by Y. Nakagawa *et al.* [18] to introduce a method for the multicast group management problem. However, these works are targeted at the wired segment of the network and are thus not applicable to the radio access segment.

SDN concepts have been also applied to a few solutions on multicast in WLANs. The work presented by

N. Soetens *et al.* [19] demonstrates how the SDN–based management improves the performance of WLANs. H. Kumar *et al.* [20] and P. Gallo *et al.* [21] provide the users with a set of controls to manage the quality of their services. Lastly, S. Tajik *et al.* [22] present a numerical analysis to improve multicast communications using SDN principles. Nevertheless, the channel occupation could be greater than the Legacy multicast one when the multicast group size increases.

Quality of Experience (QoE) aspects play an important role in multimedia applications. K. Piamrat *et al.* [23] deploy a neural network to map QoE measurements into data rates, while G. Rubino *et al.* [24] present a new hybrid objective–subjective video quality metric. Finally, although some changes in the Linux kernel are required, S. Paris *et al.* [25] also explore this problem in a real–world environment.

When the size of the multicast group grows, some approaches present scalability problems due to the number of retransmissions, the control traffic overhead or the transformation of multicast frames into unicast ones. Scalability issues are partially solved by Y. Sangenya *et al.* [26] through a protocol that improves the delay and frame loss rate of the clients by dividing and scheduling the stations into several groups. AMuSe [27] is presented as an efficient leader–based algorithm to dynamically select a subset of feedback nodes and adjust the bitrate accordingly. Nevertheless, it is assumed that the location of the devices can be estimated. The concept of assisting stations is also presented by Y. Bokyung *et al.* [28]. The AP transmits the data in unicast mode to the client exhibiting the worst channel conditions. However, since the remaining stations need to sniff the ongoing transmission with this client in order receive the multicast stream, the proposed scheme is not compatible with the 802.11 standard.

Scalability issues are exacerbated when considering several multicast groups. Although there is not much research in 802.11, this problem has been studied in WiMAX [29], [30]. P. Sendn-Raa *et al.* [29] propose a group management by comprising in the same group the clients attached to each relay station. Nevertheless, a procedure to schedule the multicast groups is not provided. F. Han *et al.* present a mathematical model [30] where the stations are divided into two groups according to the distance to the Base Station. In this way, two time slots are needed and the data rate is adjusted based on the user with the worst channel conditions in each group. The Multi–View Group Management Protocol [31] analytically intends to facilitate the 3D video transmissions in Wi–Fi multicast. To this end, each view is associated with a multicast group, in a manner that a client may subscribe to a set of views by joining a set of multicast groups.

Despite the progresses made, most of the works are not validated in real–world environments or are not compatible with the IEEE 802.11 standard. Therefore, the lack of practical approaches to address the multicast data rate adaptation in Wi–Fi networks becomes highly noticeable. Moreover, the applications performance depends on several factors such as network congestion and distribution, QoS requirements and multicast group size. As a consequence, integration between rate adaptation features and multicast retransmissions policies while ensuring a high scalability level is still an open issue.

IV. System Design

Enterprise WLANs must support a wide spectrum of services. Nonetheless, the management of both these services and of the network itself becomes more difficult as the number of devices increases. This, along with the difficulty of adding new functionalities to the Wi–Fi MAC layer, has led to the concept of SDN–based WLANs. This new paradigm addresses such limitations by introducing a fully programmable and modular network, making it possible to implement control and management tasks on top of a (logically) centralized control plane instead of implementing them as distributed applications running across the various Wi–Fi APs in the network.

OpenFlow is one of the most widely adopted options to implement the link between the data–plane and the control–plane (the socalled southbound interface). Nevertheless, its features are targeted at wired packet switched networks and are not suited for controlling wireless networks [32]. As a result, in the last years several SDN solutions for wireless and mobile networks have emerged, examples include Odin [33], CloudMAC [34] and 5G–EmPOWER [32].

The work presented in this paper has been implemented taking as a reference the 5G–EmPOWER platform [32]. Nevertheless, it should be noted that, the multicast scheme presented in this work is absolutely general and can be in principle applied to any centrally controlled enterprise WLAN. In this section, we first describe the main features of the 5G–EmPOWER platform. Then we introduce the *Transmission Policy* abstraction designed to allow an SDN controller to reconfigure a Wi–Fi AP rate adaptation policy. Finally, we show how these two abstractions can be used to implement the *SM-SDN@Play* algorithm for multicast groups management.

A. 5G–EmPOWER

5G–EmPOWER is a network operating system for wireless and mobile networks. As shown in Fig. 1, it is composed of three layers: infrastructure, control, and application. The *Infrastructure Layer* consists of a programmable 802.11 data–path (*i.e.* the 802.11 APs). This layer is made up of four main modules, namely *Rate Control Statistics*, *Transmission Policies*, *IGMP Membership* and *Multicast Addresses Management*. The first two modules are used for MCS selection, while the last ones focus on multicast groups formation. These blocks are further described in the following sections. The data–plane network elements in the *Infrastructure Layer* are in constant communication with the (logically) centralized controller situated at the *Control Layer*. Notice that the communication between the data–path is implemented using a custom built protocol. The details of this protocol are out of the scope of the paper and are omitted due to space constrains. However, a full description can be found in [35]. Finally, applications, such as *SM-SDN@Play*, run at the *Application Layer* and leverage on the global network view exposed by the controller in order to implement the network intelligence.

B. The Transmission Policy Abstraction

The fundamentals of SDN call for a clear separation between control–plane and data–plane. This in time requires to

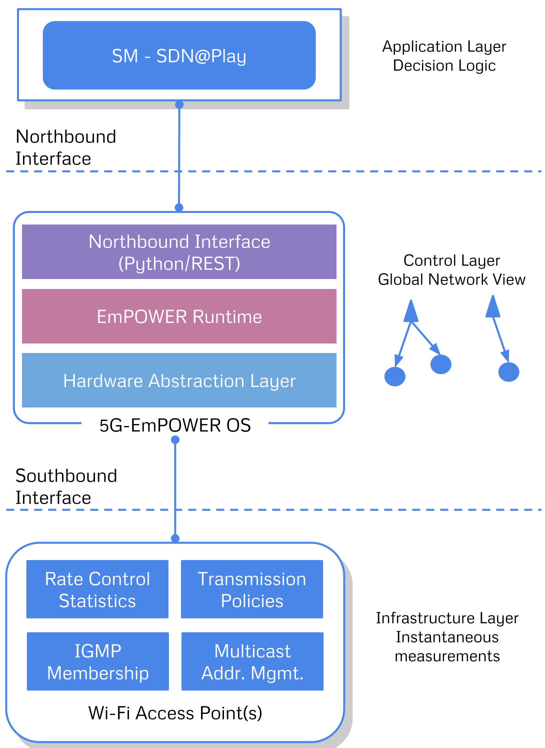

Fig. 1: 5G–EmPOWER System Architecture.

identify how network resources are exposed (and represented) to software modules written by developers and how those can affect the network state. Due to the stochastic nature of the wireless medium, the physical layer parameters that characterize the radio link between a Wi–Fi AP and a wireless client, such as transmission power, MCS, and Multiple Input Multiple Output (MIMO) configuration, must be adapted in real–time to the actual channel conditions. As a consequence, any programming abstraction for rate–adaptation in Wi–Fi networks must clearly separate fast–control operations that must happen very close to the air interface, such as rate adaptation, from operations with looser latency constrains, such as mobility management.

In this work we use the *Transmission Policy* abstraction [3]. A *Transmission Policy* is defined for each $\langle AP, client\rangle$ pair in the network and specifies the range of parameters the AP can use for its communication with that wireless client. Such parameters include:

- *MCSes*. The set of MCSes that can be used by the rate selection algorithm;
- *RTS/CTS Threshold*. The frame length above which the RTS/CTS handshake must be used;
- *No ACKs*. The AP shall not wait for ACKs if true;
- *Multicast policy*. Specifies the multicast policy, which can be Legacy, DMS, or Unsolicited Retries;
- *Unsolicited Retries Count*. Specifies the number of unsolicited retransmissions.

Table I presents three *Transmission Policy* configuration examples for unicast and multicast destination addresses. The first multicast entry (*01:00:5e:b4:21:90*) specifies the usage of Legacy as multicast mode, and 24 Mbps as transmission rate. By contrast, in the second multicast entry (*01:00:5e:40:a4:b4*), the DMS mode is selected. We remind the reader that DMS transmits each frame in unicast mode as many times as the number of receptors in the group. Therefore, the transmission rate is selected from the list of MCSes specified in the *Transmission Policy* of each receptor and the remaining parameters are not applicable.

The *Transmission Policy* configurations are manipulated by the controller via the southbound interface using a CRUD (Create, Retrieve, Update, Delete) model. Due to space constraints the details of the signaling protocol are omitted.

C. Multicast Rate Adaptation

The *SDN@Play* algorithm presented in our previous work [3] uses the unicast link delivery statistics computed at the Wi–Fi AP to calculate the MCS used for a multicast transmission. Notice how, link delivery statistics can only be computed for unicast transmissions. Therefore, *SDN@Play* alternates between the DMS and Legacy multicast modes in order to collect unicast link delivery statistics even when there are no ongoing unicast transmissions between the AP and the wireless clients. The ratio between the DMS and Legacy periods is fully configurable, hence allowing network programmers to trade reliability for channel. Fig. 2 depicts the high level operation of *SDN@Play*.

In the *first phase*, which extends over the shortest period of time, the controller sets DMS as multicast policy for a given multicast address A. This allows the APs to gather the statistical information of all the clients in a multicast group. In the *second phase*, the previous statistics are used to compute the MCS with the highest delivery probability, R_{tx}, for all the stations in the group. Then, the Legacy mode is set as *Transmission Policy* for the multicast address A, and R_{tx} is configured as single MCS for that destination.

Based on the current statistical information, the transmission rate for a certain multicast group is calculated as described below. Let us define M as the set of receptors in a multicast group and let $R(n')$ be the set of MCSes supported by the multicast receptor $n' \in M$. Moreover, let $P_r^{n'}$ be the delivery probability of the MCS index $r \in R(n')$ at the multicast receptor $n' \in M$. Accordingly, R_{valid}, the set of MCS indexes with a delivery probability higher than a given threshold r_{th} for all receptors, can be computed as follows:

$$R_{valid} = \bigcap_{n' \in M} \left\{ r \in R(n') | P_r^{n'} > r_{th} \right\} \quad (1)$$

Following from this result, the multicast transmission rate R_{tx} can be computed as follows:

$$R_{tx} = \begin{cases} max\left(R_{valid}\right) & if\ R_{valid} \neq \emptyset \\ min\left(\bigcup_{n' \in M} \left\{ \underset{r \in R(n')}{\operatorname{argmax}}(P_r^{n'}) \right\} \right) & otherwise \end{cases} \quad (2)$$

Notice how this approach ensures an appropriate data rate even for the clients with poor channel conditions. Furthermore,

TABLE I: *Transmission Policies* Configuration Examples.

Destination	Type	MCS	RTS/CTS	No ACK	Multicast	Unsolicited Retries Count
5c:e0:c5:ac:b4:a3	unicast	6, 9, 12, 18, 24, 36, 48, 54	2436	False	n.a.	n.a.
01:00:5e:b4:21:90	multicast	24	n.a.	n.a.	Legacy	n.a.
01:00:5e:40:a4:b4	multicast	n.a.	n.a.	n.a.	DMS	n.a.

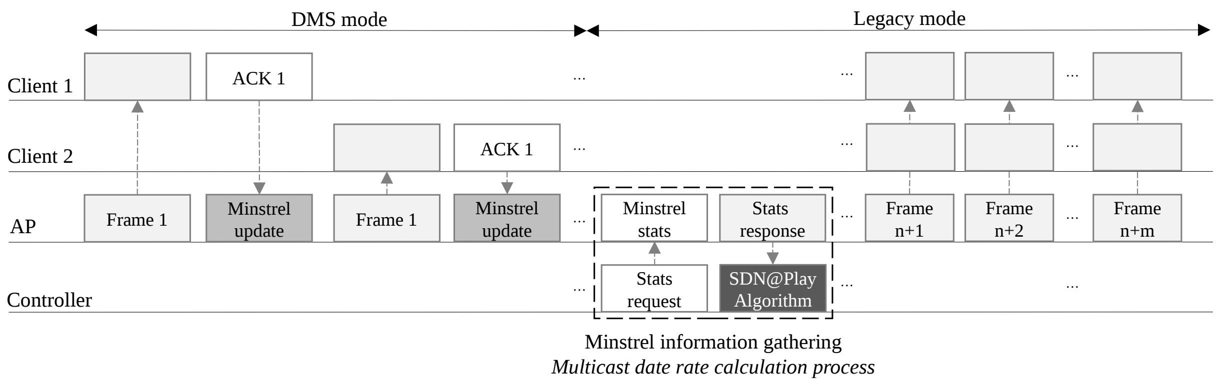

Fig. 2: *SDN@Play*'s scheme. In the *first phase DMS* is used as multicast policy allowing the link delivery statistics gathering. In the *second phase* the policy is switch to *Legacy* and the collected statistics are used to compute the optimal multicast MCS.

if the delivery probability of all MCS indexes is lower than the minimum required reliability level, r_{th}, the algorithm picks for each receptor the data rate that achieves the highest delivery probability. Then, it chooses as multicast rate the most robust one (*i.e.* the lowest) among those rates. Although gathering the link delivery statistics needs some signaling between the APs and the network controller, only a few small changes must be done at the APs and no modifications are required on the wireless clients. Consequently, *SDN@Play* is fully backward compatible with the standard IEEE 802.11 protocol.

D. *Multiple Multicast Group Management*

Building on the *SDN@Play* algorithm described in the previous section, the *SM-SDN@Play* multicast group management algorithm is introduced in this work. When several transmissions are targeted for multiple multicast groups, an instance of *SDN@Play* must be run for each of them to separately adapt their data rate. Consequently, the lack of coordination among the working phases of *SDN@Play* of each group with regard to the others may make the algorithm very inefficient. Especially, depending on the size of each group, the overlapping of several DMS phases may arise collisions, retranssmissions and performance issues. In other words, if the group management is not performed properly, the DMS phase of some of them may take place at the same time, which would result in a high number of simultaneous unicast transmissions (one for each receptor in each multicast group).

In order to show the importance of an efficient scheduler for the multicast groups, we will use as an example the scenario described below. Let us take 500 ms and 2500 ms for the duration of the DMS and Legacy periods of *SDN@Play*, respectively. In other words, during the first 500 ms the algorithm uses the DMS policy, while the Legacy one is used for the next 2500 ms. This is applied to all the multicast groups in the network. Consequently, in the first phase of the algorithm, the number of simultaneous unicast transmissions will increase with the number of active multicast groups. As described in Section II, DMS has serious scalability problems, which would also affect the performance of *SM-SDN@Play*.

In order to overcome the problem described above, the total length of the two phases, L, is divided into small parts, whose duration corresponds with the duration of the DMS period, dms_d. Let also leg_d be the duration of the Legacy period. Thus, we can define n as the total number of subphases and d_i as the length of each subphase $i \in L$ as follows:

$$d_i = \frac{L}{dms_d} \tag{3}$$

Accordingly, the Legacy period, leg_d, would be composed of $n-1$ consecutive subphases, and can be derived as follows:

$$leg_d = (n-1) \cdot d_i \tag{4}$$

In order to prevent the unicast transmissions (in the DMS phase) of all the multicast groups from taking place simultaneously, the DMS period of each group is set in a different subphase i. When an new multicast group is created, the controller assigns the DMS period of that group to one of the available subphases. This operation is sketched in Fig. 3. As can be observed in the situation displayed in *Step 3.1*, up to 6 multicast groups can be accommodated without overlaps in the DMS phase. This is achieved by using 500 ms and 2500 ms for the duration of the DMS and the Legacy phases, respectively, as stated in the example above.

However, as can be seen in the situation shown in *Step 3.2* in Fig. 3, it may be the case that all the slots are occupied when a new multicast group is created. In view of this, if possible, the duration of each subphase d_i for the DMS period must be recomputed. Let dms_{min} be the minimum amount of time needed to compute the link delivery statistics, and dms_{max}

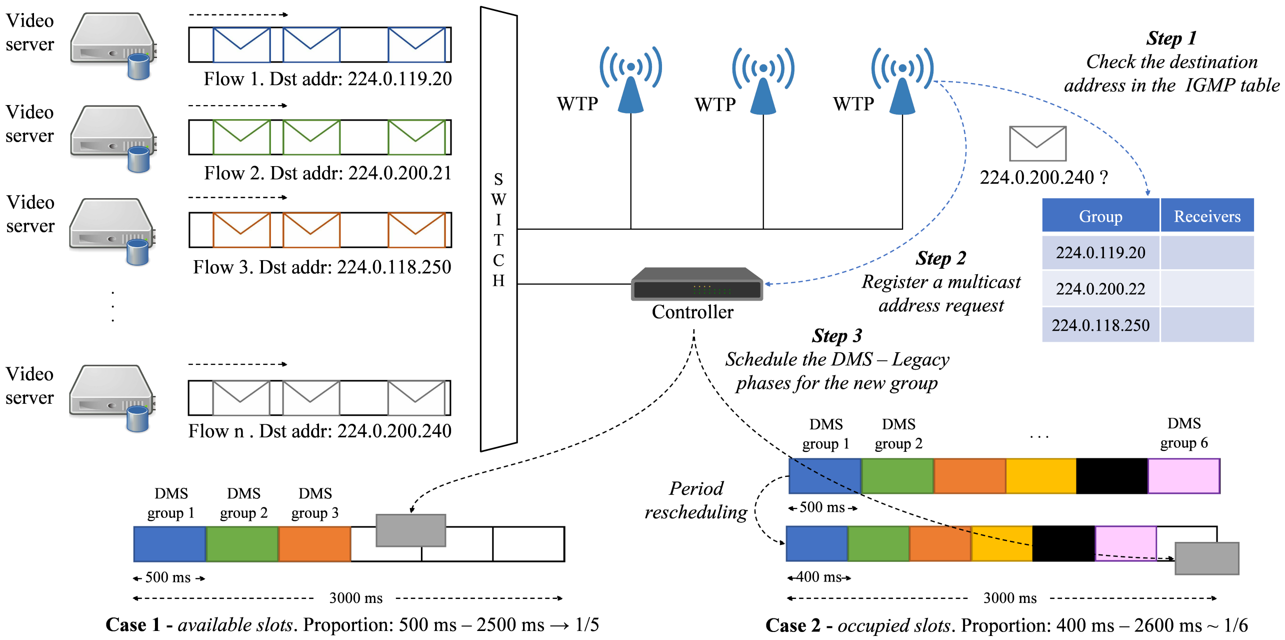

Fig. 3: *SM-SDN@Play* periods management and transmission policies coordination.

the maximum length for the DMS period to avoid causing performance degradation. Moreover, let S be the set of $s = |S|$ multicast groups in the network. Therefore, the new length for the DMS and Legacy periods can be expressed as follows:

$$dms_d = max(dms_{min}, min(dms_{max}, \lceil \frac{L}{S} \rceil)) \tag{5}$$

$$leg_d = max(L - dms_d, dms_{min}) \tag{6}$$

We would like to emphasize that, in some cases, the proportion between the policies ratio and the entire duration, L, may not be exact. In that case, the algorithm will approximate the duration d_i with the aim of not modifying the defined ratio. This phenomenon is also sketched in *Step 3.2* in Fig. 3.

Although it would be an extreme case, in the specific situation of simultaneously managing a huge number of multicast groups, and depending on the duration dms_{min}, it could happen that dms_{min} is equal to dms_d, and hence to d_i. In other words, the protocol subphases cannot be split again. In view of this, the DMS period of a new multicast group would coincide with one of the already scheduled groups. We consider this as an unlikely scenario which in any case would only result in the overlap of a few subphases with a negligible impact on the network performances.

SM-SDN@Play allows to dynamically schedule the DMS periods of the different multicast groups with the aim of avoiding collisions between the unicast transmissions of each group. This approach makes *SM-SDN@Play* suitable for managing huge multicast groups, increases the scalability level with regard to *SDN@Play* and makes it possible to maintain the network throughput.

E. Complexity Analysis

In this section, we would like to analyse the computational complexity of the *SM-SDN@Play* algorithm.

For each multicast frame to be transmitted by an AP the list of active multicast groups must be traversed. As a result, if the number of multicast groups is s, the complexity of this operation is $O(s)$. Notice however that the number of multicast groups is expected to be very small. As a result, this operation complexity can be considered constant.

The complexity of scheduling a new multicast group is also essentially constant. In fact, if a free DMS slot is available, then *SM-SDN@Play* simply assigns the new multicast group to a free slot. Conversely, if a free DMS slot is not available then *SM-SDN@Play* must recompute a new length for both the DMS and the Legacy periods. However, since this operation does not depend on the number of multicast groups nor on the number of active multicast receptors, the complexity of scheduling a new multicast group can be considered constant. Finally, if the periods have been recomputed, the algorithm must iterate through the list of multicast groups to assign the new periods to each of them. Consequently, in the worst case, the computational complexity of recomputing the groups periods is $O(s)$.

In order to compute the list of valid rates R_{valid}, the *SM-SDN@Play* algorithm must first traverse the list of receptors M and for each of them it must then traverse the list of supported transmission rates R. In the worst case the length of this list is mr where m is the number of receptors in the group and r is the number of transmission rates. Such list must then be traversed again in order to find the actual multicast transmission rate R_{tx}. As a result, the overall computational complexity for this operation is $O((mr)^2)$. Notice how this operation is performed once for each multicast group at the end of the group DMS period.

TABLE II: Minstrel Retry Chain Configuration.

Rate	Look–around		Normal transmission
	Random < Best	Random > Best	
$r0$	Best rate	Random rate	Best rate
$r1$	Random rate	Best rate	Second best rate
$r2$	Best probability	Best probability	Best probability
$r3$	Base rate	Base rate	Base rate

V. Implementation details

A. Statistics gathering

The 5G–EmPOWER platform provides a full set of programming primitives for the network management trough a Python–based SDK [32]. These primitives can be used in *polling* or *trigger* mode. The *polling* mode allows the controller to periodically poll the APs for specific information, while in the *trigger* mode this information is sent by the APs to the controller when the firing condition is verified.

In this work, the *polling*–based primitives presented by E. Coronado *et al.* [3] are used to collect the rate adaptation algorithm statistics for a given multicast receptor. This information includes, for each supported MCS, the Exponentially Weighted Moving Average (EWMA) of the frame delivery probability, the expected throughput, and the number of successful and failed transmissions in the last observation window. This primitive is used by *SM-SDN@Play* to periodically gather and update the link delivery statistics of all the receptors in a multicast group. This information is updated by the MCS selection algorithm implemented by the AP. Therefore, no extra computation complexity is added to the APs.

B. Data–path Implementation

APs are composed of one OpenvSwitch [36] instance for the wired backhaul and one Click modular router [37] instance for the 802.11 data–path implementation. In this work, Click is used to handle the clients/APs frame exchange, while the remaining network intelligence is managed by the controller. The controller communicates with Click via the southbound interface through a persistent TCP connection.

The MCS selection mechanism is implemented in Click using the Minstrel algorithm [38]. Minstrel follows a multi–rate retry chain model where four rate–count pairs, $r0/c0$, $r1/c1$, $r2/c2$ and $r3/c3$ are defined, as shown in Table II. They specify the rate that must be used to transmit a given number of retry attempts. If a frame is successfully transmitted, the remaining part of the retry chain is ignored. Otherwise, the next pair is used until the frame is properly transmitted or is finally dropped. To adapt to channel conditions, the statistics are recomputed every 500 ms. Minstrel spends the 90% of the time using the collected link delivery statistics to configure the retry chain, while in the remaining 10% of the time, other MCSes are randomly selected to gather new statistics.

For a multicast address, Minstrel will use the first MCS in the list if the retransmission mode is set to Legacy. If the policy is set to DMS, the entry is ignored and the policy associated to each receptor is used instead. Finally, if the Unsolicited Retries mechanism is selected, the frame is sent N times at the specified rate.

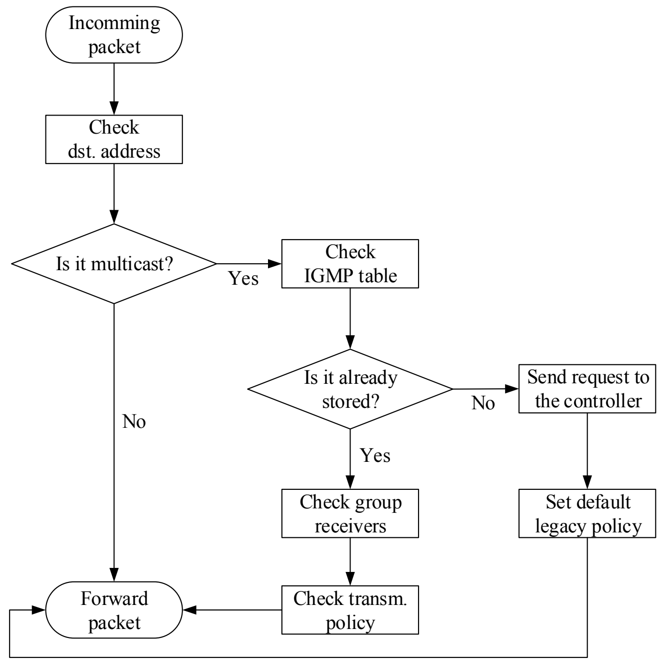

Fig. 4: *SM-SDN@Play* AP flowchart.

C. Transmission Policy Abstraction

The *Transmission Policy* abstraction is exposed to the programmers to configure the delivery features of a destination address through the `tx_policy` property of a *Resource Block* object. A *Resource Block* is the minimum allocation block in the network and is defined as a 2–tuple $\langle f, b\rangle$, where f and b are, respectively, the center frequency and the band type. Therefore, the each AP has as many *Resource Block*s as the number of installed Wi–Fi interfaces.

The *Transmission Policy* configuration only requires to specify the information for the MCS and multicast policy. The following example shows the configuration needed to set the DMS retransmission policy for the *01:00:5e:00:00:fb* address:

```
txp = block.tx_policies['01:00:5e:00:00:fb']
txp.mcast = TX_MCAST_DMS
```

In a similar manner, the `tx_policy` can be reset to the Legacy mode, for which the new multicast rate is also defined:

```
txp = block.tx_policies['01:00:5e:00:00:fb']
txp.mcast = TX_MCAST_LEGACY
txp.mcs = [24]
```

This solution is directly extensible to *SM-SDN@Play* given that it easily allows the specification of a different policy for each multicast group without introducing extra complexity.

D. Multicast Groups Management Abstraction

In this work, a *Multicast Group Management* abstraction is introduced to properly handle the stations requests to join or leave a certain multicast group. To achieve this goal, both the APs and the network controller are involved. However, APs merely forward the information to the controller, which is in charge of making the corresponding decisions.

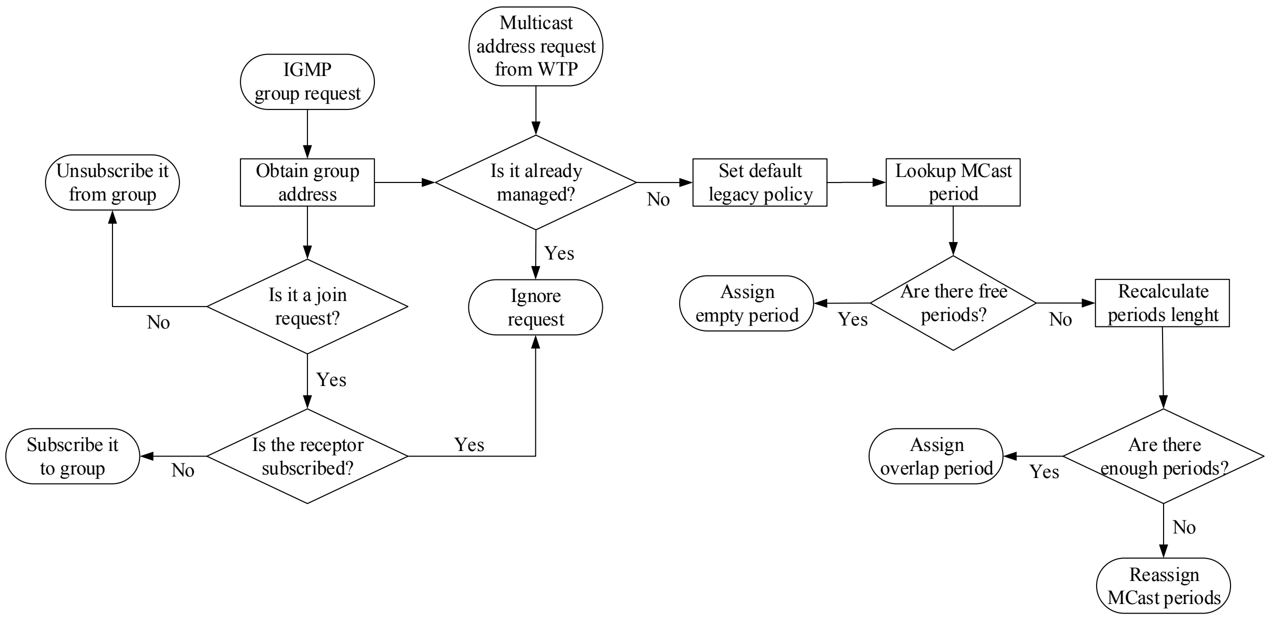

Fig. 5: *SM-SDN@Play* controller flowchart.

When an AP receives a multicast frame, it must check if there is already defined a forwarding rule for that multicast address. The flowchart followed by an AP is shown in Fig. 4. When the frame comes from an station that wants to join a multicast group, an Internet Group Management Protocol (IGMP) frame is also received. This management frame stores the multicast address and the IGMP request type, which mainly corresponds with *join* and *leave* requests. To this end, the *IGMP table* object is defined, as depicted in *Step 1* in Fig. 3. This structure includes the multicast addresses in use and the receptors of each group. On the one hand, if the group is already registered in the table, it means that it is already managed by the controller. Therefore, the receptors subscribed to the group are directly obtained and the frame is forwarded using the *Transmission Policy* defined for that address. On the other hand, if none of the entries corresponds with the group address, the request is sent to the controller, as depicted in *Step 2* in Fig. 3. While this request is being processed by the controller in *Step 3* to schedule the DMS phase of the new group, the frame is transmitted using the Legacy policy.

At the controller side two types of inputs can be distinguished, as can be seen in Fig. 5. The controller may detect a multicast transmission to which there are no clients subscribed yet or receive IGMP requests from the AP for a group inclusion or exclusion of a certain client.

When a new multicast address request is received, the Legacy multicast *Transmission Policy* is temporarily specified. Then, the controller must look for an available DMS period for this group, as described in Subsection IV-D. However, if all the slots are occupied, the protocol periods must be recalculated. If after this procedure, there are still free slots, the multicast group is scheduled in the first available one. Otherwise, it means that there is a huge number of multicast groups and this one must be scheduled in conjunction with another one.

The controller can also receive IGMP requests. On the one hand, if the multicast group specified in the request is already managed, the controller checks only the request type. The request could come from a client already subscribed to a group or from a new one. Depending on this fact, the controller will register it as a group member or ignore the request. On the other hand, a client could send a request for a multicast transmission that has not started yet. Then, in addition to the previous procedure, the operations described above for a new multicast address scheduling must be also performed.

VI. Performance evaluation

The performance evaluation presented in this section has been carried out from two points of view to show the scalability level of *SM-SDN@Play* and how efficient it is in managing multiple simultaneous multicast applications. This evaluation is performed in a real–world scenario, and establishes a comparison between our proposal, and the Legacy multicast and DMS schemes defined in the IEEE 802.11 standard. In the next subsection we will describe the characteristics of the scenarios. Then, an in–depth analysis of the results obtained during the measurements campaign will be shown.

A. Evaluation Methodology

The testbed used for the evaluation is displayed in Fig. 6 and is composed of an AP, the 5G–EmPOWER controller, a video server and a set of multicast receptors (MRx). All these devices, apart from the APs, are Dell-branded laptops powered by an Intel i7 CPU, equipped with 8GB RAM memory modules and running Ubuntu 16.04.01.

The AP is built upon a PCEngines ALIX 2D (x86) board, to which a Wi–Fi card based on the Atheros AR9220 chipset is connected. This AP uses the OpenWRT Operating System (15.05.01 version) and runs a Click instance for the 802.11

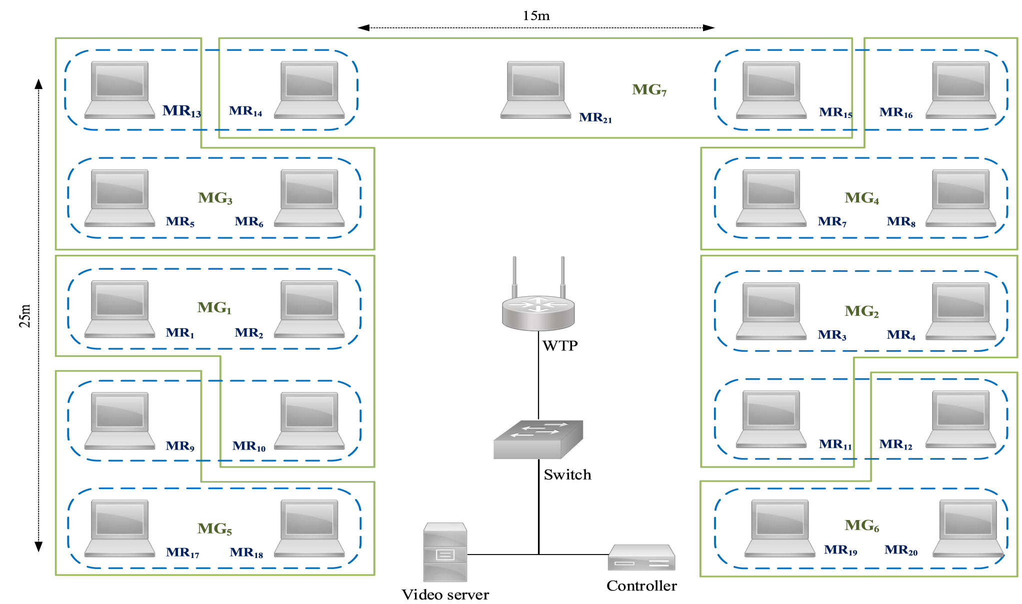

Fig. 6: Testbed deployment layout. Groups marked in blue color corresponds with the *Scalability Analysis*, whereas green markings are related to the *Multiple Groups Analysis*.

data–path. The multicast clients are widely distributed over the network coverage and are divided into groups according to the analysis type. Notice that in Fig. 6 the clients distribution corresponds to the *Scalability Analysis* when they are surrounded by blue marks. By contrast, the clients distribution in green corresponds to the *Multiple Groups Analysis*.

In the *Scalability Analysis*, a variable number of multicast receptors, ranging from 2 to 20 in steps of two stations, has been considered. The server generates and transmits a video streaming application that is delivered to the multicast receptors. This application consists on one minute video sequence with a resolution of 1920×1080 encoded using the High Efficiency Video Coding Standard (HEVC) [39] and transmitted using FFmpeg [40]. This video has been encoded making use of two different compression levels, resulting in 1.2 Mbps and 6.2 Mbps bitrate transmissions. This allows us to test how the network performance is determined by different traffic loads. Due to space limitation it was not possible to report on the *SM-SDN@Play* performance using different videos and/or compression schemes.

In the *Multiple Groups Analysis* a variable number of multicast groups is considered. This number ranges from 1 to 7 groups, each of them being made up of three receptors. The same one minute video sequence encoded for the *Scalability Analysis* is used. However, in this case, the video server transmits this video at 1.2 Mbps as many times as the number of multicast groups. Moreover, since the effect of using different bitrates has been already shown in the previous analysis and those results can be equally applied to this one, it is omitted in this test due to space constraints. Notice how we decided not to change the number of receptors involved in the experiment given that the goal of this section is to demonstrate the scalability of *SM-SDN@Play* for an increasing number of multicast groups. The scalability of the scheme for an increasing number of receptors was already studied in [3].

These scenarios have been considered for both analyses: Legacy multicast, DMS and *SM-SDN@Play*. The tests are performed in the 5.2 GHz band using the 802.11a physical layer. We remind the reader that Legacy multicast transmissions are carried out at the basic rate. Hence, due to the selected physical layer, Legacy transmissions will be sent at 6 Mbps. For *SM-SDN@Play* the duration ratio between the DMS and the Legacy phases has been set to $(500, 2500)$ ms, respectively. In order to show the analysis outcomes, we have selected as metrics the normalized throughput, the channel occupancy ratio and the percentage of retransmitted frames. The link delivery statistics have been cleared after each test. The multicast application is the only transmission that takes place in the network. This ensures that, in *SM-SDN@Play*, statistical data from the Minstrel algorithm can be only gathered during the DMS period. Finally, it should be noted that the experiments have been conducted within a 95% confidence interval and repeated 10 times to avoid possible fluctuations.

B. Experimental Results

1) Scalability Analysis: To ensure a proper QoS level, a high throughput must be achieved in video applications. Fig. 7 shows the average normalized throughput for the multicast schemes with an increasing number of receptors transmitting a video application at 1.2 Mbps. The performance of the Legacy multicast and *SM-SDN@Play* schemes remains practically constant in the $96 - 100\%$ interval. Conversely, the performance of DMS is highly damaged when increasing the number of receptors. Although at the beginning its performance is similar to the one achieved by the other schemes, it is slightly below 90% from 8 to 12 receptors. Moreover, it is highly degraded from the point in which 14 clients are considered, which shows the serious scalability issues of DMS.

Although Legacy multicast provides good delivery throughput with low bitrates, using a basic rate for all the transmis-

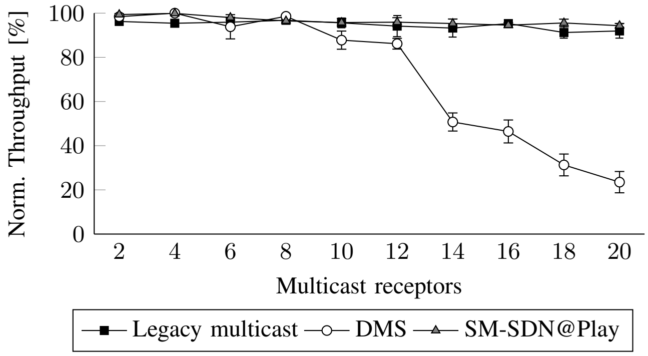

Fig. 7: Normalized throughput for an increasing number of multicast receptors using a video transmission at 1.2 Mbps.

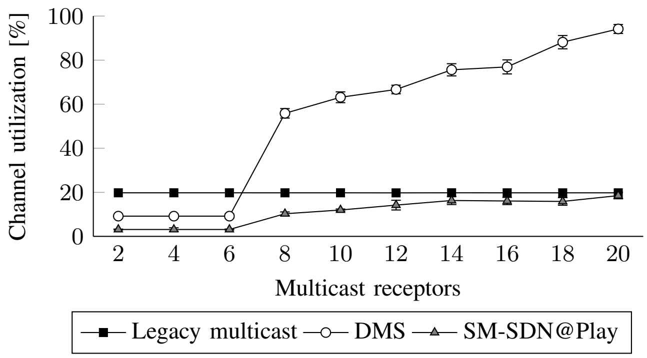

Fig. 8: Channel utilization for an increasing number of multicast receptors using a transmission at 1.2 Mbps.

sions results in a high channel utilization. In Fig. 8 it can be observed how this ratio is around 20% for a video streaming at 1.2 Mbps. In the DMS case, the channel utilization increases with the number of receptors. This issue is due not only to the increasing number of simultaneous unicast transmissions but also to the growing percentage of retransmissions. In the case of the application at 1.2 Mbps, the channel utilization becomes higher from the moment in which the network is made up of 8 receptors until the end of the measurements, when this ratio reaches 90%. By contrast, the channel occupancy ratio of *SM-SDN@Play* remains below the one achieved by the standard schemes in all the cases. The use of higher MCS indexes with regard to Legacy makes it possible to reduce the period of time that the channel is busy. Moreover, given that it only uses DMS in the smallest phase of the algorithm, the channel occupancy ratio is also lower than the DMS one. Specifically, this ratio is under 10% until the half of the test and is below 20% even for 20 receptors. These results show that, although the channel utilization of *SDN@Play Mobile* increases with the number of receptors, this growth is far lower than the DMS one and it does not raise scalability problems.

The enormous number of simultaneous transmissions sent in DMS causes an increase in the retransmission ratio. This effect is shown in Fig. 9, where it is appreciable how this ratio is over 50% when using a wide range of receptors. In contrast, using the DMS policy in only a small part of its protocol makes the retransmission ratio of *SM-SDN@Play* not exceed 20%. In Fig. 10 the behavior of this ratio over time for a 1.2 Mbps transmission is reported. Notice how when using *SM-SDN@Play*, the retransmitted packets ratio remains constant over the test. Nevertheless, a great amount of information is retransmitted by DMS, which becomes even higher when the network is saturated. Specifically, in Fig. 10 can be seen that this figure arises from the half of the test when holding 10 receptors and it practically reaches 50% from the beginning of the transmission for 20 receptors. We would like to stress that since in Legacy multicast packets are neither acknowledged nor retransmitted, the retransmission ratio analysis is omitted for this scheme in all the experiments.

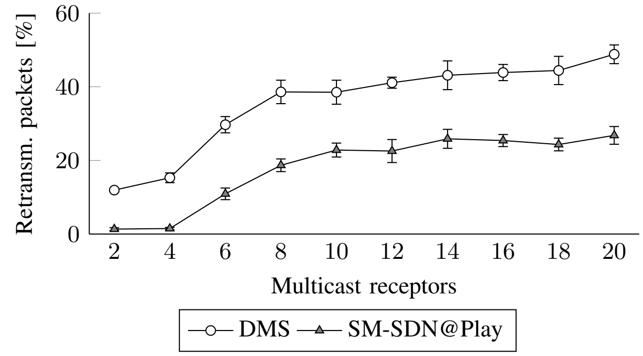

Fig. 9: Retransmitted packets for an increasing number of multicast receptors using a video application at 1.2 Mbps.

Fig. 11 displays the rates distribution used by DMS and *SM-SDN@Play* for a 1.2 Mbps service. Notice that the information of the Legacy multicast scheme is omitted as it always uses the 6 Mbps basic rate. DMS uses higher MCS indexes than *SM-SDN@Play* until the number of retransmissions arises from 14 receptors. The use of acknowledgments and the constant update of the link delivery statistics make this possible for DMS. Furthermore, in order to achieve a high reliability level, the MCS index is decreased in *SM-SDN@Play* if the success delivery probability does not exceed a given threshold. In the case of *SM-SDN@Play*, the percentage of transmissions in which high data rates, such as 54 and 48 Mbps, are used is as high as 66% with a 95% CI [61.55% - 72.04%].

This analysis has shown how *SM-SDN@Play* improves the efficiency of the multicast video transmissions. Due to the lower channel utilization, our proposal may allow a greater amount of information to be simultaneously transmitted. In order to show this effect, the multicast schemes are also evaluated using a bitrate of 6.2 Mbps with the aim of checking how it compromises their performance and efficiency.

The outcomes for a 6.2 Mbps multicast application present some differences with regard to the previous analysis. First, DMS impairs the network throughput in an earlier stage. As can be seen in Fig. 12, this value is below 30% with only 6 receptors. Meanwhile, although the performance of Legacy multicast keeps constant, it is not able to achieve a normalized throughput higher than 70%. By contrast, *SM-SDN@Play* is only slightly degraded with respect to the previous test and

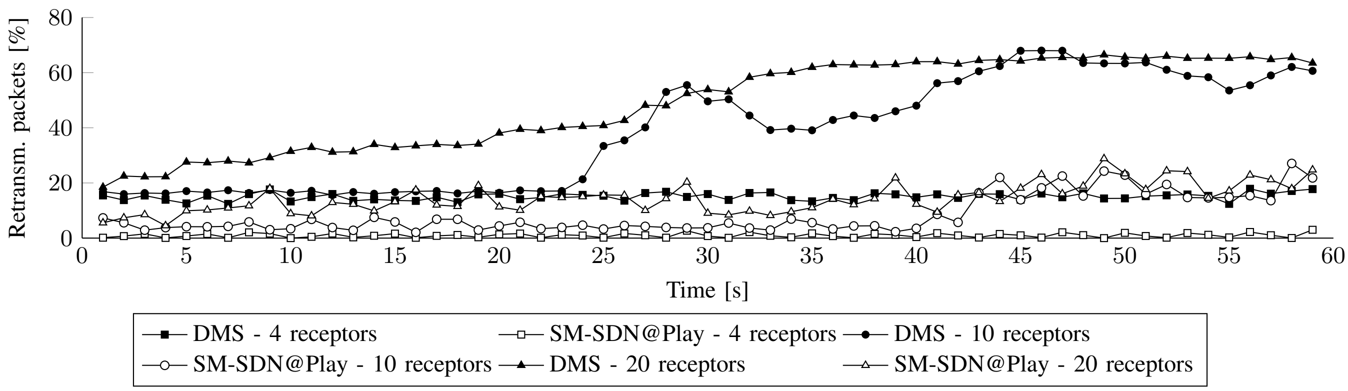

Fig. 10: Retransmitted packets over time using a multicast video application at 1.2 Mbps.

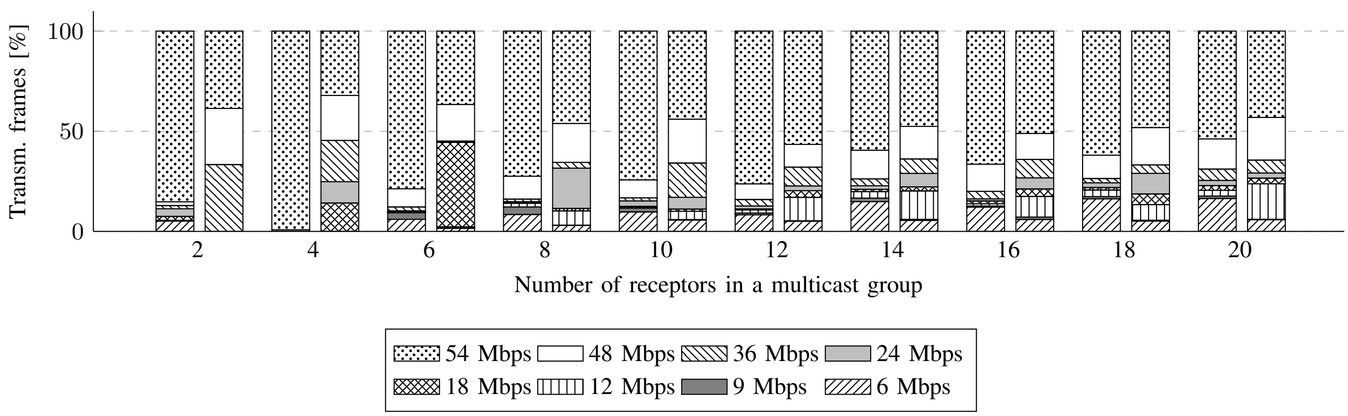

Fig. 11: Rates distribution corresponding to the DMS and *SM-SDN@Play* schemes in a multicast transmission at 1.2 mbps with an increasing number of receptors.

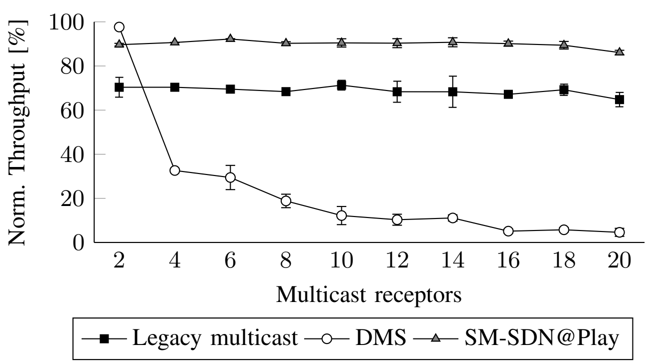

Fig. 12: Average normalized throughput for an increasing number of receptors using a video transmission at 6.2 Mbps.

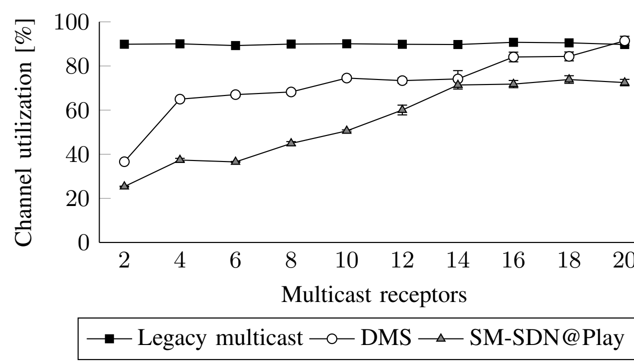

Fig. 13: Channel utilization for an increasing number of multicast receptors using a transmission at 6.2 Mbps.

outperforms the other schemes with a normalized throughput higher than 90%, regardless of the number of receptors.

As mentioned above, the basic rate used by Legacy multicast makes the channel be busy for long periods of time. The channel utilization shown when analyzing a 1.2 Mbps transmission significantly arises until reaching a value close to 90%, as plotted in Fig. 13. This proves that despite its performance, this scheme is not suitable for applications with a high bitrate and impairs the performance of other transmissions in the network. Similarly, the DMS problem is exacerbated in this scenario, in which it makes the network become saturated in an earlier point. When the video bitrate is increased, the channel utilization of *SM-SDN@Play* also arises with regard to the first scenario. However, this ratio allows the video to be delivered without loosing a significant part of the information.

The retransmissions issue becomes even worse when DMS needs to handle a 6.2 Mbps multicast service. In Fig. 14 can be seen that this ratio arises from the half of the test when

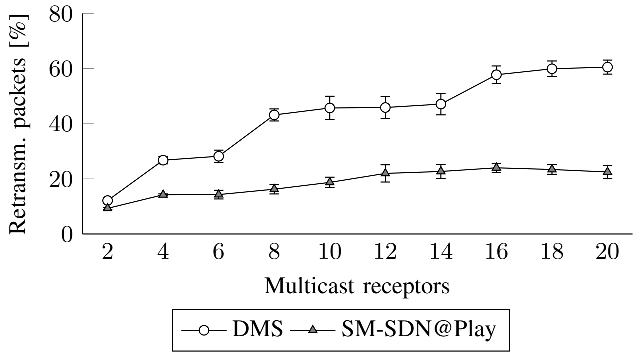

Fig. 14: Retransmitted packets for an increasing number of multicast receptors using a video application at 6.2 Mbps.

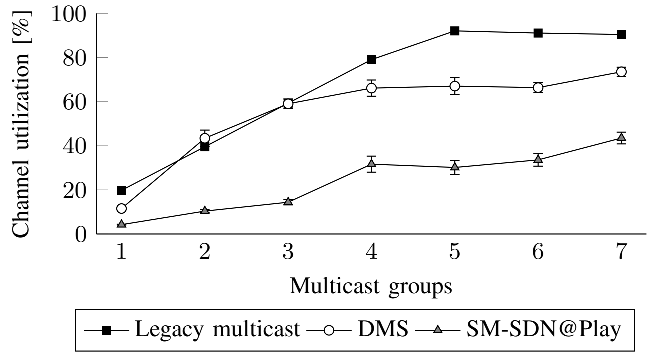

Fig. 16: Channel utilization for a transmission at 1.2 Mbps targeted at multiple multicast groups.

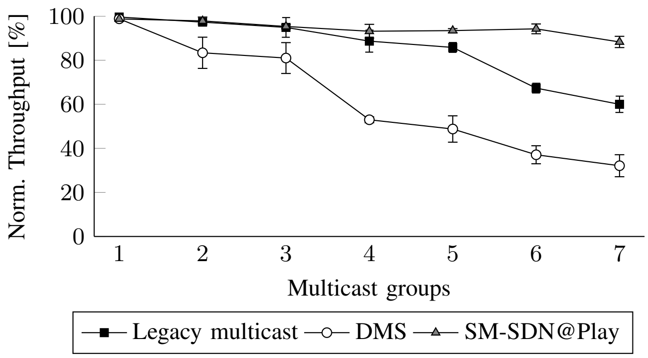

Fig. 15: Average normalized throughput for a video transmission at 1.2 Mbps targeted at multiple multicast groups.

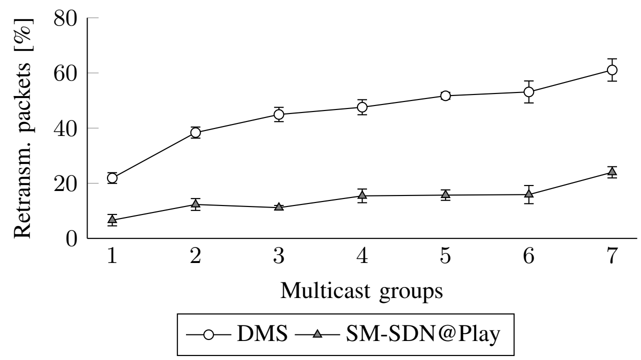

Fig. 17: Retransmitted packets for an increasing number of multicast groups using a video application at 1.2 Mbps.

holding 10 receptors and it practically reaches 70% from the beginning of the transmission for 20 receptors. By contrast, small differences can be found for *SM-SDN@Play*, whose retransmission ratio stands at 20%, irrespective of the bitrate.

2) Multiple Groups Analysis: After studying the scalability level of the proposal, a similar analysis is performed to evaluate its efficiency when managing multiple multicast groups.

Fig. 15 reports the average normalized throughput of the evaluated schemes upon an increasing number of multicast groups. This shows how the performance of Legacy multicast is highly degraded with regard to the case of a single multicast group. We remind the reader that, for a 1.2 Mbps application, the channel utilization of this scheme is around 20%, which makes it be practically saturated when 4 groups are managed. Hence, a throughput fall can be appreciated from this point for the Legacy mechanism. The performance of DMS is similar to the one provided in the previous analysis. However, it is also slightly reduced due to the increase in the simultaneous transmissions and the need to forward them accordingly. Conversely, the normalized throughput of *SM-SDN@Play* remains practically constant and similar to the figure obtained in the single group analysis ($96-100\%$) until practically the end of the measurements, where the performance is slightly impaired by the amount of traffic in the network.

Closely connected to the previous metric, Fig. 16 plots the channel utilization of each scheme. In contrast with the first analysis, where the channel occupancy ratio remained constant for the Legacy multicast scheme, in this case it proportionally arises with the number of multicast groups. The ratio achieved by DMS is similar to the one obtained in the single group analysis, until the AP is completely saturated and is not able to forward on time all the frames. At this point, the period of time that the channel is busy by *SM-SDN@Play* falls far short of the remaining schemes. In fact, the channel occupancy ratio is only slightly risen with regard to the *Scalability Analysis* and it only reaches a 40% utilization ratio when managing 7 simultaneous multicast transmissions.

The retransmissions distribution of DMS and *SM-SDN@Play* is depicted in Fig. 17. This view is almost equal to the one observed in the first analysis of the evaluation for both schemes. The only small difference can be seen for *SM-SDN@Play* when the network holds 7 multicast groups, when the percentage of retransmitted packets lies minimally above 20%.

Finally, Fig. 18 displays the distribution of the MCSes used by DMS and *SM-SDN@Play*. We remind the reader that, once again, Legacy multicast rates distribution is omitted. Figures obtained for DMS are practically the same as presented above given that, regardless of the multicast group a receptor belongs to, the unicast transmission for each of them is forwarded independently from the remainder. However, it is worthy to highlight that in *SM-SDN@Play* the percentage of the frames

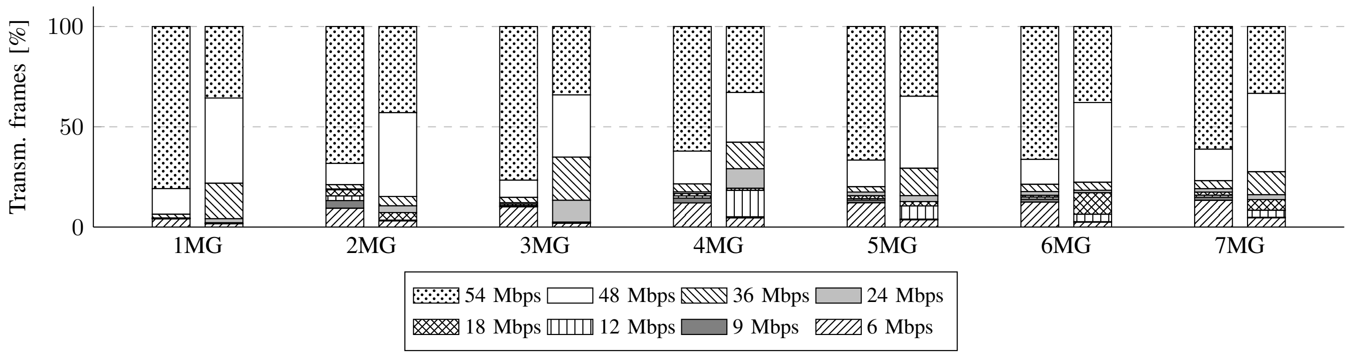

Fig. 18: Rates distribution corresponding to the DMS and *SM-SDN@Play* schemes based on a multicast video transmission at 1.2 Mbps for an increasing number of multicast groups.

that is transmitted using high data rates (both 48 and 54 Mbps) is, in some cases, higher than in the *Scalability Analysis*, reaching it up 83% with a 95% CI [80.22% - 86.29%]. This is due the fact that the data rate of each group is independently calculated only considering the receptors in that group, which allows it to provide more accurate results.

VII. Conclusions

In this work we have proposed *SM-SDN@Play* as a novel solution for multicast group management in SDN–based WLANs. *SM-SDN@Play* is fully backward compatible with the 802.11 standard and does not require any change to the wireless clients. Only minimal changes to the APs are needed.

The performance of *SM-SDN@Play* has been evaluated in a real–world scenario implemented over the 5G–EmPOWER platform and compared with the one achieved by the standard DMS and Legacy multicast schemes. The results prove that, in contrast with the standard mechanisms, our proposal scales properly with respect to both the number of receptors in a multicast group and the number of multicast groups.

A particularly important open issue regards the security and performance isolation properties of the *SM-SDN@Play* scheme. Further studies are needed here in order to properly assess the impact of a misbehaving multicast stream on the other groups. Moreover, we also plan to extend *SM-SDN@Play* in order to make use of Scalable Video Coding and investigate the video quality layers prioritization according to the channel status of the multicast groups.

Acknowledgements

This work has been supported by the Spanish Ministry of Economy under Grant Agreement BES-2013-065457, by the European Union's MINECO/FEDER funds under project TIN2015-66972-C5-2-R, and by the H2020 Research and Innovation Action under Grant Agreement H2020-ICT-671639 (COHERENT).

References

[1] *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, ANSI/IEEE Std 802.11, LAN/MAN Standards Committee of the IEEE Computer Society Std., 2012.

[2] *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications. Amendment 3: MAC Enchancements for Robust Audio Video Streaming*, ANSI/IEEE Std 802.11aa, LAN/MAN Standards Committee of the IEEE Computer Society Std., 2011.

[3] E. Coronado, R. Riggio, J. Villalón, and A. Garrido, "SDN@Play: A Multicast Rate Adaptation Mechanism for IEEE 802.11 WLANs," in *Proc. of IEEE CCNC*, Las Vegas, NV, USA, 2017.

[4] *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications. Amendment 8: Local and Metropolitan Area Networks*, ANSI/IEEE Std 802.11v, LAN/MAN Standards Committee of the IEEE Computer Society Std., 2011.

[5] Y. Daldoul, D.-E. Meddour, T. Ahmed, and R. Boutaba, "Performance and scalability evaluation of IEEE 802.11v/aa multicast transport," *Wireless Communications and Mobile Computing*, vol. 16, no. 14, pp. 1987–2000, 2016.

[6] K. Joy and K. Sneha Kumar, "Reliable Multicast in Multi-Access Wireless LANs," in *Proc. of IEEE INFOCOM*, New York, USA, 1999.

[7] D. Dujovne and T. Turletti, "Multicast in 802.11 WLANs: An Experimental Study," in *Proc. of ACM MSWiM*, Terromolinos, Spain, 2006.

[8] J. Villalón, P. Cuenca, L. Orozco-Barbosa, Y. Seok, and T. Turletti, "ARSM: a cross-layer auto rate selection multicast mechanism for multi-rate wireless LANs," *IET Communications*, vol. 1, no. 5, pp. 893–902, 2007.

[9] Y. Park, Y. Seok, N. Choi, Y. Choi, and J.-M. Bonnin, "Rate-adaptive multimedia multicasting over IEEE 802.11 wireless LAN," in *Proc. of IEEE CCNC*, Las Vegas, NV, USA, 2006.

[10] J. Villalón, P. Cuenca, L. Orozco-Barbosa, Y. Seok, and T. Turletti, "Crosslayer Architecture for Adaptive Video Multicast Streaming over Multirate Wireless LANs," *IEEE Journal on Selected Areas in Communications*, pp. 699–711, 2007.

[11] S. P. Chuah, Z. Chen, and Y. P. Tan, "Energy-Efficient Resource Allocation and Scheduling for Multicast of Scalable Video Over Wireless Networks," *IEEE Transactions on Multimedia*, vol. 14, no. 4, pp. 1324–1336, Aug 2012.

[12] M.-T. Sun, L. Huang, A. Arora, and T.-H. Lai, "Reliable MAC layer multicast in IEEE 802.11 wireless networks," in *Proc. of IEEE ICPP*, Vancouver, BC, Canada, 2002.

[13] Y. Song, J. Chung, W. Sung, B. Kim, D. Hyun, and J. Jang, "An Analysis of Reliable MAC Layer Multicast in Wireless Networks," in *Proc. of IEEE ICTC*, Gyeongju, South Korea, 2007.

[14] N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S. Shenker, and J. Turner, "OpenFlow: Enabling Innovation in Campus Networks," *ACM Computer Communication Review*, vol. 38, no. 2, pp. 69–74, 2008.

[15] L. Bondan, L. Mller, and M. Kist, "Multiflow: multicast clean-slate with anticipated route calculation on Open-Flow programmable networks," *Journal of Applied Computing Research*, vol. 2, no. 2, pp. 68–74, Feb 2012.

[16] H. E. Egilmez, S. Civanlar and A. M. Tekalp, "An Optimization Framework for QoS-Enabled Adaptive Video Streaming Over OpenFlow Networks," *IEEE Transactions on Multimedia*, vol. 15, no. 3, pp. 710–715, 2013.

[17] X. Zhang, M. Yang, L. Wang and M. Sun, "An OpenFlow-Enabled Elastic Loss Recovery Solution for Reliable Multicast," *IEEE Systems Journal*, doi:10.1109/JSYST.2016.2634325, 2017 [in press].

[18] Y. Nakagawa, K. Hyoudou and T. Shimizu, “A Management Method of IP Multicast in Overlay Networks Using Openflow,” in *Proc. of ACM HotSDN*, Helsinki, Finland, 2012.

[19] N. Soetens, J. Famaey, M. Verstappen and S. Latre, “SDN-based Management of Heterogeneous Home Networks,” in *Proc. of IEEE CNSM*, Barcelona, Spain, 2015.

[20] H. Kumar, H. Gharakheili and V. Sivaraman, “User Control of Quality of Experience in Home Networks using SDN,” in *Proc. of IEEE ANTS*, Kattankulathur, India, 2013.

[21] P. Gallo, K. Kosek-Szott, S. Szott and I. Tinnirello, “SDN@home: A method for controlling future wireless home networks,” *IEEE Communications Magazine*, vol. 54, no. 5, pp. 123–131, May 2016.

[22] S. Tajik and A. Rostami, “MultiFlow: Enhancing IP Multicast over IEEE 802.11 WLAN,” in *Proc. of IFIP WD*, Valencia, Spain, 2013.

[23] K. Piamrat, A. Ksentini, J.-M. Bonnin, and C. Viho, “Q-DRAM: QoE-Based Dynamic Rate Adaptation Mechanism for Multicast in Wireless Networks,” in *Proc. of IEEE Globecom*, Honolulu, Hawaii, USA, 2009.

[24] G. Rubino, “Quantifying the Quality of Audio and Video Transmissions over the Internet: the PSQA Approach,” *Imperial College Press*, London, 2005.

[25] S. Paris, N. Facchi, F. Gringoli, and A. Capone, “An Innovative Rate Adaptation Algorithm for Multicast Transmissions in Wireless LANs,” in *Proc. of IEEE VTC*, Dresden, Germany, 2013.

[26] Y. Sangenya, F. Inoue, M. Morikura, K. Yamamoto, F. Nuno and T. Sugiyama, “Grouping-Based Reliable Multicast Protocol for Large-Scale WLANs,” *IEICE Transactions on Communications*, vol. E97-A, no. 1, pp. 40–48, Jan 2014.

[27] Y. Bejerano, J. Ferragut, K. Guo, V. Gupta, C. Gutterman, T. Nandagopal and G. Zussman, “Scalable WiFi multicast services for very large groups,” in *Proc. of IEEE ICNP*, Goettingen, Germany, 2013.

[28] B. Yoon and H. S. Kim, “A Scheme Improving Performance of IEEE 802.11 Multicast Protocol,” in *Proc. of INNOV*, Venice, Italy, 2012.

[29] P. Sendín-Raa, F. J. González-Castao, F. Gil-Castieira and P. S. Rodríguez-Hernández, “Dynamic multicast groups with adaptive path selection in MMR WiMAX networks,” in *Proc. of IEEE MILCOM*, San Jose, CA, USA, 2010.

[30] H. Fang, D. Liang, W. Wang and Y. Liu, “A Time-division multicast transmission scheme using multiple multicast groups,” in *Proc. of IEEE ICCAE*, Singapore, Singapore, 2010.

[31] C. H. Lin, D. N. Yang, J. T. Lee and W. Liao, “Efficient Error-Resilient Multicasting for Multi-View 3D Videos in Wireless Network,” in *Proc. of IEEE Globecom*, Washington, DC, USA, 2016.

[32] R. Riggio, M. K. Marina, J. Schulz-Zander, S. Kuklinski, and T. Rasheed, “Programming Abstractions for Software-Defined Wireless Networks,” *IEEE Transactions on Network and Service Management*, vol. 12, no. 2, pp. 146–162, 2015.

[33] L. Suresh, J. Schulz-Zander, R. Merz, A. Feldmann, and T. Vazao, “Towards Programmable Enterprise WLANS with Odin,” in *Proc. of ACM HotNets*, New York, USA, 2012.

[34] P. Dely, J. Vestin, A. Kassler, N. Bayer, H. Einsiedler and C. Peylo, “CloudMAC – An OpenFlow based architecture for 802.11 MAC layer processing in the cloud,” in *Proc. of IEEE Globecom*, Anaheim, CA, USA, 2012.

[35] “EmPOWER Protocol,” 2017. [Online]. Available: https://github.com/5g-empower/5g-empower.github.io/wiki/EmPOWER-Protocol

[36] B. Pfaff, J. Pettit, K. Amidon, M. Casado, T. Koponen, and S. Shenker, “Extending Networking into the Virtualization Layer,” in *Proc. of ACM HotNets*, New York, USA, 2009.

[37] E. Kohler, R. Morris, B. Chen, J. Jannotti, and M. F. Kaashoek, “The Click Modular Router,” *ACM Transactions on Computer Systems*, vol. 18, no. 3, pp. 263–297, 2000.

[38] D. Xia, J. Hart, and Q. Fu, “Evaluation of the Minstrel rate adaptation algorithm in IEEE 802.11g WLANs,” in *Proc. of IEEE ICC*, Budapest, Hungary, 2013.

[39] ISO/IEC and ITU-T, “High Efficiency Video Coding (HEVC). ITU-T Recommendation H.265 and ISO/IEC 23008-2 (version 4),” Dec. 2016.

[40] “FFmpeg project.” [Online]. Available: https://ffmpeg.org/, Accessed on 15.06.2017.

Estefanía Coronado Estefanía Coronado is a Ph.D. student in Advanced Computer Technologies at the University of Castilla–La Mancha (Spain). She received her M.Sc. degrees in Computer Engineering and in Advanced Computer Technologies in 2014 and 2015, respectively, from the aforementioned university. Her M.Sc. Theses versed on the evaluation and optimization of multicast and multimedia transmissions over IEEE 802.11 WLANs. In 2013, she joined the High Performance Networks and Architectures (RAAP) research group as a research assistant. Her current research interests and areas of publication include wireless networks, QoS and QoE over IEEE 802.11 networks, Software Defined Networking and multimedia communications.

Roberto Riggio Roberto Riggio (SM’16) is currently Chief Scientist at CREATE-NET where he is leading the Future Networks Research Unit efforts on 5G Systems. His research interests include Mobile Network Operating Systems for 5G Networks, Performance Isolation in Multi-tenant 5G Networks, 5G Radio Access Network Slicing, and Multi-domain Network Service Orchestration. He is the creator of 5G–EmPOWER the first Network Operating System for Mobile Networks which is now used in several 5G–PPP Projects. He has 2 granted patents, 70 papers published in internationally refereed journals and conferences, and has generated more than 1 M€ in competitive funding. He received several awards including: the IEEE INFOCOM 2013 Best Demo Award, the IEEE ManFI 2015 Best Paper Award, and the IEEE CNSM 2015 Best Paper Award. He serves in the TPC and in the OC of leading conferences in the networking field and he is associate editor for the Wiley International Journal of Network Management and Editor for the Spring Wireless Networks journal. He is the co-founder of the IEEE Soft5G and of the IEEE 5GMan workshops. He is a member of the ACM and a Senior Member of the IEEE.

José Villalón José Villalón received his M.Sc. degree in Computer Science from the University of Castilla-La Mancha, Albacete, Spain in 2003. He received a Ph.D. degree in Computer Engineering in 2007 from the University of Castilla–La Mancha. In 2007 he joined the Department of Computer Systems at the University of Castilla–La Mancha where he is currently an Assistant Professor. He has also been a visiting researcher at INRIA, Sophia Antipolis (France). His research interests include high-performance networks, wireless networks, Quality of Service and Quality of Experience over IEEE 802.11, multicast transmission, Software Defined Networking (SDN), Quality of Service over WiMAX, video transmission and error–resilient protocol architectures.

Antonio Garrido Antonio Garrido is full professor of computer architecture and technology in the University of Castilla–la Mancha (Spain). His research interests include high-performance networks, and wireless networks and video compression. He has published over 50 papers in international journals. He has led several research projects in telemedicine, computer networks and advanced computer system architectures. Antonio Garrido collaborates with the Spanish Ministry of Education in the National plan for quality evaluation of Spanish Universities from the year 2000. He has been dean of the of the School of Computer Engineering from 2000 to 2008.