SDN@Play: Software-Defined Multicasting in Enterprise WLANs

Estefanía Coronado*, Elisenda Temprado Garriga†, José Villalón‡,
Antonio Garrido‡, Leonardo Goratti†, Roberto Riggio*
*Wireless and Networked Systems, FBK CREATE-NET, Trento, Italy
Email: {e.coronado, rriggio}@fbk.eu
†Zodiac Inflight Innovations, Munich, Germany
Email: {Elisenda.Temprado, Leonardo.Goratti}@zii.aero
‡High Performance Networks and Architectures, University of Castilla-La Mancha, Albacete, Spain
Email: {JoseMiguel.Villalon, Antonio.Garrido}@uclm.es

***Abstract*—Concurrent distribution of multimedia streams to a high number of end users through a wireless network is extremely popular nowadays. Nevertheless, multicast traffic orchestration over traditional Wi-Fi networks is still an important problem for network operators. Low multicast data rates and poor coexistence with other services are just some of the challenges that must be faced when deploying multicast solutions over Wi-Fi. In this article we present *SDN@Play*, a centralized approach to manage multicast flows based on Software-Defined Networking (SDN) principles. We also report on a proof-of-concept implementation of *SDN@Play* that can be used in several scenarios ranging from vehicular networks to in-flight entertainment systems, and from enterprise group to machine type communications applications. Finally, we make the prototype fully available to researchers and practitioners under a permissive APACHE license.**

***Index Terms*—WLANs, IEEE 802.11, software-defined networking, multicast, mobility, multimedia.**

I. Introduction

Multicast provides an efficient way to deliver the same content to multiple users. Unfortunately, its use has so far been limited owing to the fact that the most popular Internet content delivery platforms, like Youtube and Netflix, use unicast streams to transmit their services. Recently, an increasing number of scenarios, ranging from enterprise networking to in-flight entertainment systems, are demanding improved support for multicast. For example, recent trends in commercial aviation call for a significant weight reduction by diminishing the cables for non-critical functions such as the entertainment system, where multicast would allow airlines to provide a broader range of services and a wider offer to customers. Likewise, large sport events, in which replaying a video and providing different views of the event can be required, is another situation where multicast delivery could reduce the network load due to the number of users watching the same content simultaneously.

Wi-Fi has proved the ability to support these heterogeneous use cases [1], where unicast frames would greatly increase the network load. This is particularly important in wireless networks where radio resources are scarce and must be used efficiently. Multicast communications are suitable for business and entertainment purposes, and have made content distribution one of the most profitable applications for the future 5G networks. Mobile network operators are accommodating this spike in capacity demand by deploying denser heterogeneous Radio Access Networks (RANs), in which Wi-Fi is becoming the traffic offloading technology for 5G systems. The significant number of practitioners-oriented [2], as well as scientific papers [3] are good evidence of this aspect, which is also supported by the standardization of the evolved Multimedia Broadcast Multicast Services (eMBMS) in 3GPP [4].

Multicast incurs severe reliability issues due to the lack of acknowledgements and retransmissions. Consequently, in Wireless Local Area Networks (WLANs) frames are distributed using the lowest Modulation and Coding Scheme (MCS). Although this maximises the delivery probability, it also results in a higher channel occupation. The IEEE 802.11aa amendment [5] intends to address these concerns through a set of retransmission policies. However, it does not consider multicast rate adaptation and groups management. Moreover, given the widespread use of Wi-Fi compatible devices, IEEE 802.11 amendments aim at maximizing backward compatibility at the expense of innovation.

This article examines *SDN@Play*, a Software-Defined multicast solution for WLANs. *SDN@Play* builds on the fundamental pillars of Software-Defined Networking (SDN), namely control-user plane separation, programming abstractions, and open Application Programming Interfaces (APIs). The proposed solution aims at reducing the overall spectrum utilization while ensuring reliable service provisioning. To this end, we rely on the *Transmission Policy* abstraction to reconfigure the multicast MCS selection policy used by Access Points (APs). This abstraction is then used to implement a set of multicast policies for mobility management and services orchestration.

In contrast to our previous work [6], [7], this paper shows the effectiveness of an integrated multicast solution through practical use cases with clear business opportunities (see Sec. V). This approach intends to prove that the low complexity and the standard compliant features of our solution can make multicast an attractive proposition to be applied directly to commercial services on the market and in the industry. Experimental results demonstrate how *SDN@Play* can distribute multicast content in an stable manner over time with high efficiency without requiring changes to the wireless clients. Finally, we make available the implementation under a permissive APACHE 2.0 license for academic use[1].

In this article we first provide a short background on multicast challenges in Wi-Fi networks. Then, we report on a mature proof-of-concept of *SDN@Play* and on its evaluation. Finally, we introduce the use cases that benefit from the design of our platform: Aviation, Internet of Things (IoT) and Automotive, and discuss the improvement given by *SDN@Play*.

II. Multicasting in Wi-Fi Networks

IEEE 802.11's unicast transmission is based on a two-way handshake where each frame is acknowledged by the receiver. However, in multicast, acknowledgements would inevitably collide at the transmitter. This situation raises two issues:

- *Reliability.* Given the impossibility to know if a multicast frame has arrived properly, incorrect frames cannot be retransmitted. Therefore, the reliability of the multicast service is more limited than the unicast one.
- *Adaptability.* In the absence of feedback the sender cannot adapt the data rate to the channel conditions. For this reason, the lowest MCS is used to maximize delivery probability, especially for the receivers under severe channel conditions. As a result, the channel is occupied for longer periods, which puts a strain on radio resources and demeans the capacity available for other services.

The IEEE 802.11aa amendment [5] aims to support robust multicast transport and to address the limitations caused by the lack of feedback. To this end, it introduces the Directed Multicast Service (DMS) and the Groupcast with Retries (GCR) service. Moreover, GCR consists of three retransmission methods: Legacy Multicast, Unsolicited Retries (UR) and Block ACK (BACK). This outlook allows transmitting each stream in a different manner fulfilling their requirements. Following this, Table I summarises the strengths and weaknesses of each scheme according to the specific scenario.

III. Programmable System Architecture

Typically, multicast does not consider the particular needs of the users or other simultaneous applications. To cope with this situation, *SDN@Play* is an SDN-based solution that makes multicast services user-oriented and encompasses aspects such as rate adaptation, user mobility, and services orchestration.

SDN@Play takes as a reference 5G-EmPOWER [8], a programmable Mobile Network Operating System supporting heterogeneous radio access technologies such as Wi-Fi and LTE. 5G-EmPOWER builds upon a hardware abstraction layer converging control and management protocols into a unified set of programming primitives that are then exposed to the application layer. Such primitives include support for slicing and virtualization, and for accessing the global network view maintained by 5G-EmPOWER. Primitives are available to control applications through a Python-based Software Development Kit (SDK) or through a REST interface. Notice that such primitives can be used by control/management applications, as well as by network monitoring and analytic platforms. The

[1]Online resources available at: http://5g-empower.io

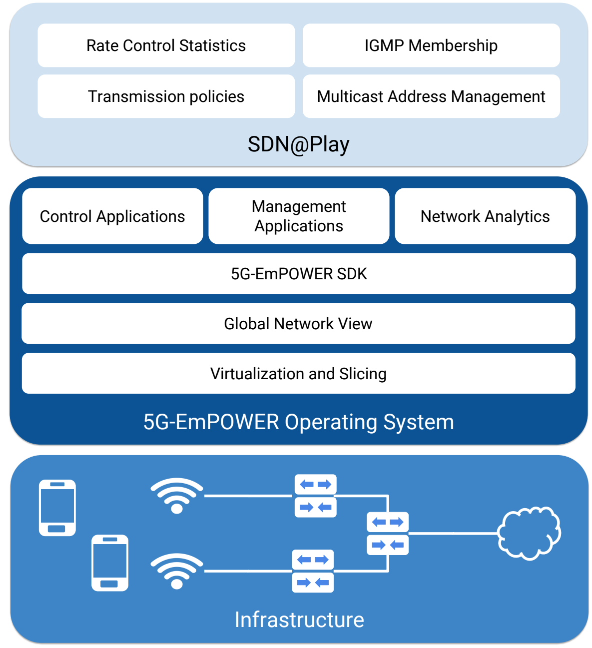

Fig. 1: High level architecture of *SDN@Play*.

high level architecture of the system is depicted in Fig. 1. In particular, it can be seen that *SDN@Play* is located in the application layer, and composed of four main modules that become the cornerstones for deploying the multicast features described in the following sections.

A. Traffic Forwarding Rules

The idea of *SDN@Play* is to define precise forwarding rules to provide user-oriented multicast services. This is done through the *Transmission Policies* abstraction [9] shown in Fig. 1, which specifies the parameters that each AP must use to deliver a certain multicast stream. These parameters include:

- *MCSes*. It specifies the set of MCSes that can be used.
- *Retransmission policy*. It indicates the frame retransmission method, which can be Legacy, DMS or UR.
- *Unsolicited Retries Counter*. It specifies the number of retries for the UR policy, if selected.
- *RTS/CTS Threshold*. It determines the frame length above which the RTS/CTS handshake must be used.

This approach enables the dynamic reconfiguration of the transmission mode in order to meet specific application requirements, e.g. one multicast stream can choose the Legacy Multicast mode at 18 Mbps, while another stream can select the DMS mode.

B. Group Identification and Management

To ensure user-oriented transmissions, identifying the users attached to each AP becomes a major point. Equally important is to know the receivers subscribed to each multicast service. To this end, the *IGMP Membership* abstraction [7] presented in Fig. 1 extends the Internet Group Management Protocol (IGMP) in a way that when a user joins or leaves a multicast group, an IGMP frame is generated. These frames are received

TABLE I: Capabilities Analysis of the Group Addressed Transmission Services introduced in IEEE 802.11aa.

Policy	**Description**	**Strengths**	**Weaknesses**
GCR Legacy Multicast	This is the original IEEE 802.11 delivery method, specified as a simple broadcast mechanism without error control or recovery mechanisms.	High scalability, easy deployment and absence of traffic overhead.	Fixed PHY rate and high channel utilization.
GCR Unsolicited Retries	Multicast frames are transmitted several times to increase the probability that all the users receive them successfully. The number of retransmissions is implementation dependent. However, frames are transmitted without confirmation or recovery mechanisms.	High scalability, easy deployment and higher reliability than original standard.	Fixed PHY rate and increase in the traffic overhead and channel utilization due to the retransmission of correct frames.
GCR Block ACK	It consists on acknowledging several multicast packets using a single ACK. To this end, the AP first establishes a block ACK agreement with each receiver. From that moment, the AP regularly sends Block ACK requests to these stations in order to discover the missing frames that must be retransmitted.	High reliability, moderate overhead, adaptive PHY rate and high efficiency.	Very complex deployment and implementation, and moderate scalability and overhead.
Directed Multicast Service	It converts a multicast frame into as many unicast frames as the number of receivers in the multicast group. Hence, it is not valid for large groups.	High reliability, adaptive PHY data rate and easy deployment.	Very low scalability.

by the APs and contain the type of IGMP request (e.g. *join* or *leave*). However, the APs merely forward the data to the SDN controller, which makes the management decisions.

The APs store in a table the *Transmission Policy* for each multicast application and the group of users subscribed. When an AP receives a join IGMP request it checks if the multicast address is already registered. In that case, the user is added to the group of receivers. Otherwise, the AP creates an entry in the table, and asks to the SDN controller the *Transmission Policy* to be used. Conversely, when receiving a leave IGMP request, the AP just removes the receiver from the multicast group, and informs the controller if it becomes empty.

C. User-oriented Multicast Transmissions

The MCSes configured in each *Transmission Policy* are calculated at the controller for each AP according to the channel conditions of the users. In unicast, the status of the receiver determines the data rate used for the transmission. Following this idea, *SDN@Play* intends to compute the data rate for a multicast service based on the rate control statistics of the receivers in that multicast group. Nevertheless, multicast traffic lacks of feedback information from the users.

Given the relevance of feedback, the design leverages the IEEE 802.11aa policies, which allows us to retrieve the data required while ensuring minimal implementation complexity. Based on the analysis shown in Table I, the operation of *SDN@Play* relies on the alternation of short unicast (to avoid excessive overhead) and long multicast periods (to ensure high scalability), as can be observed in Fig. 2. In the unicast period, DMS is set as retransmission policy, while in the multicast one, Legacy Multicast is chosen. Notice that UR and BACK have been left aside because they would imply higher computational and implementation complexity in both the APs and the stations [10].

In the unicast phase, DMS allows gathering from the APs the data rates with the highest delivery probability for each user. This process leverages the *Rate Control Statistics* module sketched in Fig. 1. Then, the multicast rate is calculated as the intersection of these data rates set. If the intersection is empty, the lowest rate from the ones with the highest delivery probability is selected to ensure that all the users receive the information properly. Finally, the SDN controller sets the Legacy Multicast policy and instructs the AP to use the MCS calculated for the specific *Transmission Policy*. This is repeated for all the APs with a configurable ratio for the DMS-Legacy periods [9].

D. Multicast Services Orchestration

When an AP holds several multicast transmissions, an instance of *SDN@Play* must be run for each of them. However, this might cause the unicast periods of each instance to take place at the same time, hence increasing the traffic overhead.

To reduce this overhead, *SDN@Play* schedules the unicast period of each multicast service in different time slots. The duration of the two periods (Legacy and DMS) is divided into n parts: one used for the DMS period, and the remaining $n-1$ parts dedicated for the Legacy one. In this way, when a multicast group is created, our solution schedules its DMS period in a slot in which no other multicast application transmits unicast frames. The module responsible for calculating the effective periods is named *Multicast Address Management* in Fig. 1. Notice that if all the slots are busy, the DMS periods of two groups would coincide. However, this is considered an unlikely event with negligible impact on the network performance [7].

E. User Mobility Management

The channel conditions of the mobile users vary along the time. Therefore, maintaining the transmission quality may involve their handover to other APs. Tackling this problem requires the controller to know the quality perceived by each user with respect to all the APs.

To do this, the receivers periodically inform their AP about the signal quality through beacon reports as defined in IEEE 802.11k [11]. This data is forwarded to the controller to detect signal drops and examine if a handover to another AP would increase the performance. However, with the goal of serving homogeneous multicast groups, the controller considers as candidates for a user's handover only the APs whose receivers experience similar channel quality. Then, from this group, the AP with the highest signal strength is selected. Finally, the *Rate Control Statistics* are updated and the *Transmission Policies* are reconfigured for the APs involved [6].

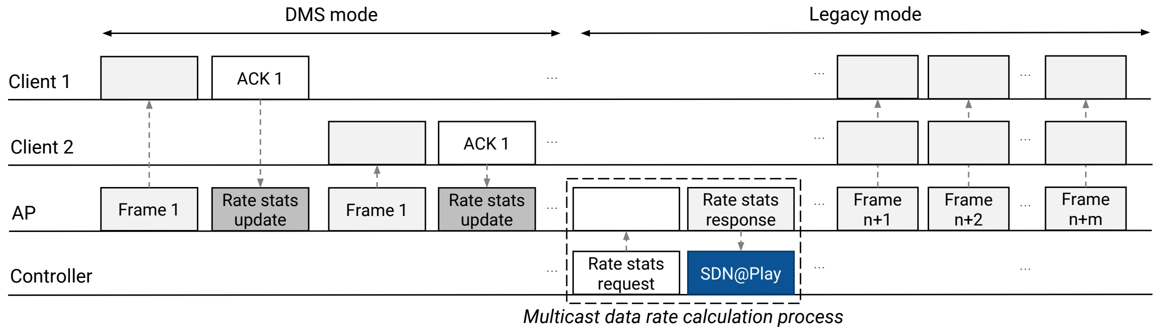

Fig. 2: Overview of the working mode of *SDN@Play*.

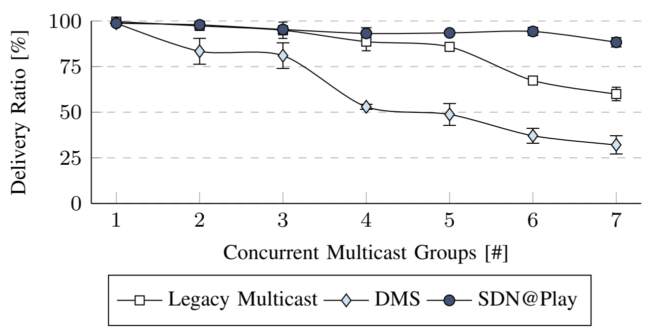

Fig. 3: Delivery ratio of a multicast application for an increasing number of multicast groups.

IV. Proof-of-Concept and Results

SDN@Play has been evaluated in an experimental testbed assessing the capacity to: (i) handle simultaneous multicast services; and (ii) satisfy the requirements of mobile and static users. In all the experiments a comparison is drawn between *SDN@Play* and the standard multicast policies DMS and Legacy Multicast. The evaluation focuses on these specific policies since they are taken as a reference to implement *SDN@Play*. Furthermore, as described in Sec. II, the multicast schemes defined in IEEE 802.11aa are not able to meet the requirements of these use cases or involve changes at the client side [10]. In the experiments, the time relationship between the DMS and the Legacy phases in *SDN@Play* is $100-900$ms. This period is chosen following the conclusions reached in [9], where the reader can find an extensive evaluation on different time configurations. Notice that each experiment has been repeated 10 times to confirm the accuracy of the results.

The first test assesses the performance when handling various multicast services. To this end, a new multicast transmission starts, i.e. a new multicast group is added, every 25s until reaching 7 groups after 180s. Services are delivered at 1.2 Mb/s by an AP to the multicast groups composed of 3 users. Notice that a constant number of users was considered in order to focus on the ability of the schemes to manage concurrent transmissions. An study on how the number of users determines the network's throughput is presented in [7].

Figure 3 depicts the performance of the experiment, showing that despite the number of transmissions, *SDN@Play* achieves remarkably higher delivery ratio than the standard schemes. This scalability improvement is given by the reduction in the channel utilization. On the one hand, as opposed to Legacy Multicast that uses the basic MCS, *SDN@Play* adapts and increases the data rate, hence using less radio resources. On the other hand, since DMS transforms multicast into unicast frames, it saturates the channel when the traffic grows. By contrast, our solution transmits unicast frames in just a short period, thus leaving the channel available for other services for more time. It must be noted that by default IEEE 802.11 does not include error control and recovery mechanisms for multicast traffic. Consequently, the delivery ratio is slightly lower than 100% when the number of transmissions increases. Conversely, Fig. 4 shows how the channel utilization changes over time when the traffic load rises. Specifically, it can be seen that the airtime used by the standard solutions increases until exhausting network resources when enlarging the number of transmissions. By contrast, the channel utilization of *SDN@Play* is maintained below 50% for the entire experiment, therefore ensuring greater scalability.

The second experiment examines the user mobility management. It is important to note that neither Legacy Multicast nor DMS support seamless handover. Nevertheless, this feature has been introduced in *SDN@Play*, allowing users to be migrated between APs without disrupting their connection. In this scenario a single multicast stream is delivered by 3 APs to 4 users (one mobile and three static) in order to evaluate the network-wide effect of user mobility in multicast environments. The reason for having only one global mobile user and one static user per AP is that the multicast rate of each AP participating in the handover depends only on the clients attached to it. However, it should be noted that the re-computation of the data rate would have the same impact on the throughput of the users already attached to the APs, regardless if they are static or mobile. The bitrate in this experiment is the same used in the first test.

Figure 5 illustrates the performance achieved in this scenario by the mobile user. At the beginning, this user is attached to an AP and gradually moves away over time until forcing a handover to a closer AP. This situation entails for the multicast standard schemes, a disassociation and an association process (at 155s and 165s for DMS and Legacy Multicast, respectively) due to the lack of seamless handover mechanisms. However, *SDN@Play* is able to move the user to a more appropriate AP in a proactive and transparent manner (at 148s), which

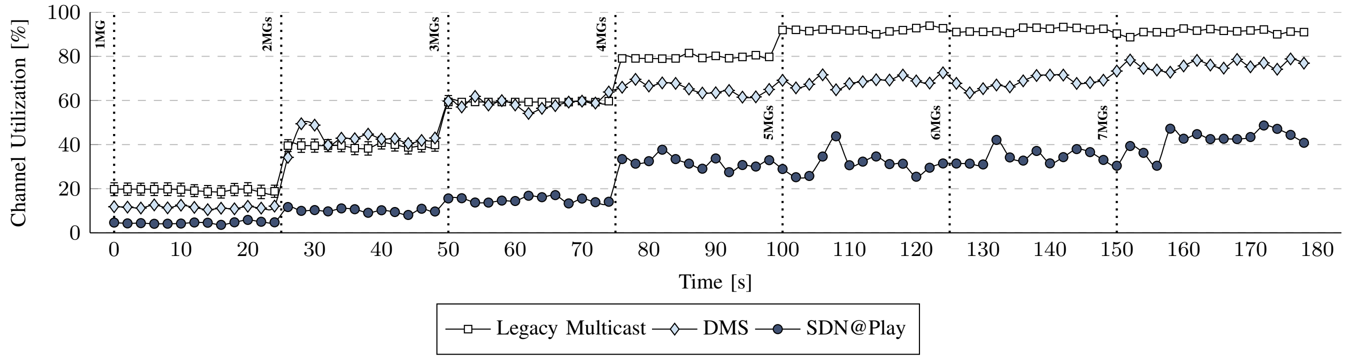

Fig. 4: Channel utilization over time for an increasing number of multicast groups (one group is added every 25s).

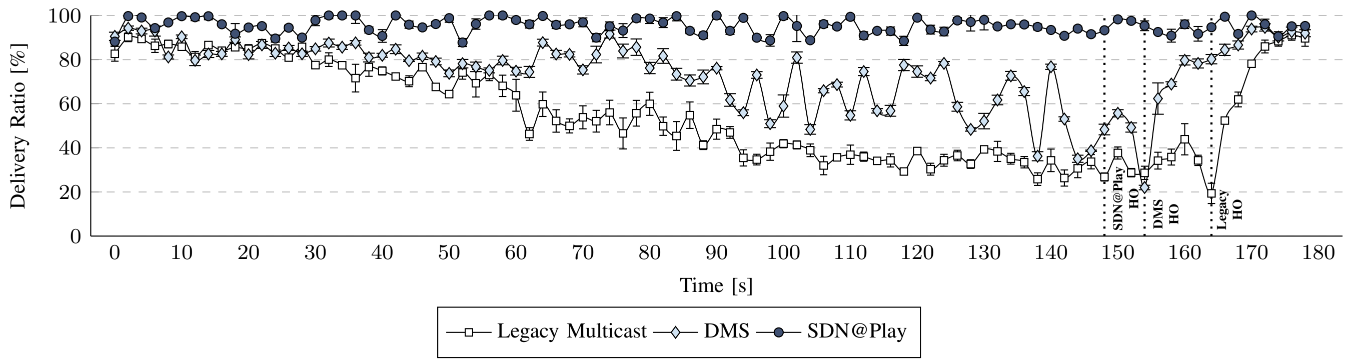

Fig. 5: Delivery ratio over time of a multicast application for a mobile receptor undergoing a handover between APs. The handover occurs at time 148s, 155s and 165s for SDN@Play, DMS and Legacy Multicast, respectively.

makes the performance of the mobile user be comparable with the one in the absence of mobility. By contrast, DMS and Legacy Multicast work reactively when experiencing connection issues. Notice that the throughput of the remaining three static users is the same than the one depicted for a single multicast group in Fig. 3. Further details concerning this aspect can be found in [6]. Finally, it is also worth mentioning that despite only 2 APs are involved in the handover, the inclusion of 3 APs intends to allow the multicast schemes deciding the destination AP of the handover process, e.g. based on the Received Signal Strength Indicator (RSSI) of such APs.

V. Discussion

The range of services delivered over wireless is extremely diverse and spans from broadband Internet access to ultra-reliable communications. Moreover, as soon as new markets will emerge, techniques such as multicast will be mandatory to enable a more efficient delivery of traffic with different patterns (machine-type, HD video stream and so on). The addition of SDN in this context is seen as the turnkey to unlock all potentials of multicast in Machine to Machine (M2M) as well [12]. In this section we will shed light onto three advanced scenarios where *SDN@Play* can bring tangible benefits.

A. Aviation

In the civil aviation market, wireless In-Flight Entertainment and Communications (IFEC) denotes the ensemble of entertainment content and wireless communication technologies. This is an appealing cost-effective solution making the content accessible directly to both Personal Electronic Devices (PEDs), relying on the Bring-Your-Own-Devices (BYOD) model, and seat screens. In this context, multicast groups over Wi-Fi (through the combination of IEEE 802.11n and IEEE 802.11ac [5]), which is the current dominant wireless communication technology inside aircrafts, can enable bandwidth efficient delivery of new services such as safety advertisements. An advanced wireless IFEC could exploit the opportunity of transmitting selected content like videoclips directly over the passengers' PED. Referring to the layout of an Airbus A320, which is the typical single body aircraft model for mid-range travels where built-in seat screens are not always available, messages can be directly transmitted to the PEDs by means of the wireless system onboard and downloading an app offered by the airline. To this extent, multicast can be used to distribute different versions of the safety video to groups of passengers depending on their location, for instance, the closest emergency exits. *SDN@Play* allows supporting several multicast groups with better performance than the IEEE 802.11aa amendment both in terms of packet delivery rate and channel utilization, without requiring any modification on the PED side. Finally, improving the channel utilization is particularly important since it allows the wireless medium to be free for other services, such as in-cabin communication.

B. Internet of Things

Internet of Things has become one of the prime priorities for both industry and academia. Significant research and standardization works were done to bring the IoT ecosystem to mass scale. To add on this, IoT has also become one of the most compelling use cases in the roadmap of 5G. In this regard, M2M includes the ensemble of communication technologies and protocols that enable things to be part of a ubiquitous network [13]. For instance, Non-Orthogonal Multiple Access (NOMA) is one candidate technology to unveil IoT over 5G in a way that, just to name a few, rising Industry 4.0 and remote healthcare use cases can be delivered by the next generation of mobile network technology [14]. In scenarios where the support of massive number of connections is a must (i.e. thousands of connected machines per km^2), multicast can provide a suitable manner to enable software and firmware upgrades to a large number of end devices with minimum human intervention. In this context, the main benefit introduced by *SDN@Play* consists of improved delivery ratio for an increasing number of groups. In particular, *SDN@Play* can facilitate a much faster upgrade by enabling grouping the devices based on their purpose and technical features, materializing the higher flexibility of software networks.

C. Automotive

The automotive industry is experiencing a major technological shift since the inception of unmanned vehicles that can autonomously cope with rising contingent situations that may suddenly occur over the roads. The challenge is nothing but fierce since the fully autonomous vehicle (i.e. Level 5) should react to traffic congestion, re-plan the route, and communicate with other vehicles (e.g. IEEE 802.11p [5]) and any other infrastructure in range. This topic falls traditionally within the concept of Vehicular Ad Hoc Networks in which road-side units are meant to provide the vehicle-to-infrastructure communication means and in which the opening to 5G has enriched them even more [15]. There is a consensus around the fact that the connected vehicles should avail data processing either performed directly onboard or remotely in the cloud. To make a practical example, an analytic engine located in the cloud processes the data collected from the vehicles, and then it feeds the results back to the vehicles via the communication infrastructure. In this case, *SDN@Play* can allow reliable delivery of such results to multiple multicast groups according to the message priority and the type of vehicle (e.g. cars, trucks, motorcycles). Moreover, *SDN@Play* can enable significant bandwidth savings leaving the channel available for other communications (e.g. emergency braking).

VI. Conclusions

Software-Defined Networking is a promising paradigm for both wireless and wired networks. In this article we introduce *SDN@Play*, an SDN-based multicasting orchestration scheme for Wi-Fi-based networks. The multicast approach presented in this article has high effectiveness in reducing the overall channel utilization and in improving the multicast frame delivery ratio for both static and nomadic users.

The approach we have pursued is meant to enable more programmability and customization in multicast applications. In particular, it is our standpoint that this approach is particularly well suited for different verticals, namely in-flight entertainment systems, connected and cooperative road mobility scenarios, and IoT applications. We are currently working toward the implementation of more complex multicast orchestration scenarios involving larger number of users per service and more diverse mobility patterns. Finally, we aim to further investigate the behaviour of the solution in other competing schemes such as 3GPP eMBMS.

Acknowledgements

This work has been supported by the European Union's H2020 Research and Innovation Action under Grant Agreement H2020-ICT-644843 (5G-ESSENCE Project) and by the European Union's REGIONAL/FEDER funds under Grant SBPLY/17/180501/000353.

References

[1] H. Wang, W. T. Ooi, and M. C. Chan, "JurCast: Joint user and rate allocation for video multicast over multiple APs," in *Proc. of IEEE INFOCOM*, San Francisco, California, USA, 2016.

[2] GSMA, "Network 2020: The 4G Broadcasting Opportunity," March 2017, Tech. Report.

[3] G. Araniti, M. Condoluci, P. Scopelliti, A. Molinaro, and A. Iera, "Multicasting over Emerging 5G Networks: Challenges and Perspectives," *IEEE Network*, vol. 31, no. 2, pp. 80–89, March 2017.

[4] D. Lecompte, and F. Gabin, "Evolved multimedia broadcast/multicast service (eMBMS) in LTE-advanced: overview and Rel-11 enhancements," *IEEE Communications Magazine*, vol. 50, no. 11, pp. 68–74, 2012.

[5] *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., 2016.

[6] E. Coronado, R. Riggio, J. Villalón, and A. Garrido, "Joint Mobility Management and Multicast Rate Adaptation in Software-Defined Enterprise WLANs," *IEEE Transactions on Network and Service Management*, vol. 15, no. 2, pp. 625–637, 2018.

[7] E. Coronado, R. Riggio, J. Villalón, and A. Garrido, "Efficient Real-Time Content Distribution for Multiple Multicast Groups in SDN-Based WLANs," *IEEE Transactions on Network and Service Management*, vol. 15, no. 1, pp. 430–443, 2018.

[8] 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.

[9] E. Coronado, R. Riggio, J. Villalón, and A. Garrido, "Programming Abstractions for Wireless Multicasting in Software-Defined Enterprise WLANs," in *Proc. of IEEE IM*, Lisbon, Portugal, 2017.

[10] A. de la Oliva, P. Serrano, P. Salvador, and A. Banchs, "Performance evaluation of the IEEE 802.11aa multicast mechanisms for video streaming," in *Proc. of IEEE WoWMoM*, Madrid, Spain, 2013.

[11] *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications. Amendment 1: Radio Resource Measurement of Wireless LANs*, ANSI/IEEE Std 802.11k, LAN/MAN Standards Committee of the IEEE Computer Society Std., 2008.

[12] Y.-J. Chen, L.-C. Wang, M.-C. Chen, P.-M. Huang, and P.-J. Chung, "SDN-Enabled Traffic-Aware Load Balancing for M2M," *IEEE Journal on Internet of Things*, vol. 5, no. 3, pp. 1797–1806, 2018.

[13] Y. Mehmood, N. Haider, M. Imran, A. Timm-Giel, and M. Guizani, "M2M Communications in 5G: State-of-the-Art Architecture, Recent Advances, and research Challenges," *IEEE Communications Magazine*, vol. 55, no. 9, pp. 194–201, 2017.

[14] Z. Ding, X. Lei, G. Karagiannidis, R. Schober, J. Yuan, and V. Bhargava, "A Survey on Non-Orthogonal Multiple Access for 5G Networks: Research Challenges and Future Trends," *IEEE Journal on Selected Areas in Communications* , vol. 35, no. 10, pp. 2181–2195, 2017.

[15] B. Di, L. Song, Y. Li, and Z. Han, "V2X Meets NOMA: Non-Orthogonal Multiple Access for 5G-Enabled Vehicular Networks," *IEEE Wireless Communications* , vol. 24, no. 6, pp. 14 – 21, 2017.

BIOGRAPHIES

Estefanía Coronado is an Expert Researcher in the Wireless and Networked Systems Unit at FBK CREATE-NET (Italy). In 2018 she completed her PhD at the University of Castilla-La Mancha (Spain) on multimedia distribution over SD-WLANs. She received the M.Sc. degrees in Computer Engineering and Advanced Computer Technologies in 2014 and 2015 from the same university. She published 10 papers in international journals and conferences. Her research interests include wireless and mobile communications, SDN/NFV and network management.

Elisenda Temprado Garriga is a R&D Systems Engineer at Zodiac Inflight Innovations. She received her Bachelor and M.Sc. degrees in Telecommunications from the Polytechnic University of Catalonia, Barcelona, in 2015. Ms. Temprado has worked on emulation platforms for ground and aircrafts networks. She has also been involved on future product development, focusing on different wireless telecommunications technologies such as Wi-Fi, WiGig, LiFi or 5G. She has submitted a patent for improving wireless communication onboard aircrafts.

José Villalón received the M.Sc. degree in computer science and the PhD in computer engineering from the University of Castilla-La Mancha, Spain, in 2003 and 2007, respectively. In 2007, he joined the University of Castilla-La Mancha as Assistant Professor. He has been a Visiting Researcher at INRIA, France. His research interests include high-performance networks, wireless networks, QoS and QoE over IEEE 802.11 and WiMAX, multicast transmission, software defined networking, video distribution, and error-resilient protocol architectures.

Antonio Garrido is a Full Professor at the University of Castilla-La Mancha, Spain. He has led several research projects in telemedicine, computer networks, and advanced computer architectures. He collaborates with the Spanish Education Ministry on the quality evaluation of Spanish universities since 2000. He was the Dean of the School of Computer Engineering from 2000 to 2008. His research interests include high-performance networks, wireless networks and video compression. He has published over 50 papers in international journals.

Leonardo Goratti received his PhD degree from the University of Oulu (Finland) and his MSc from the University of Firenze (Italy). He is currently a Senior System Engineer at Zodiac Inflight Innovations, Germany. His interests span across SDN and NFV for 5G, LTE-Advanced and millimeter wave communications. He has authored more than 70 research papers in journals and conferences. He is TPC member in international conferences, and serves as reviewer in journals such as IEEE JSAC and IEEE Communications Magazine.

Roberto Riggio is Head of the Wireless and Networked System Unit at FBK CREATE-NET. His research interests include software-defined mobile networks, network slicing, and distributed management and orchestration of network services. He has published more than 100 papers and has generated more than 2.7M Euro in competitive funding. He received several awards including the IEEE CNSM 2015 Best Paper Award. He serves in the TPC/OC of leading conferences in networking and is associate editor of several journals including the IEEE Transactions on Network and Service Management.