

User-AP Association Management in Software-Defined WLANs

Suzan Bayhan*, Estefanía Coronado[‡], Roberto Riggio[‡], Anatolij Zubow[†]

*University of Twente, The Netherlands, s.bayhan@utwente.nl

[†] Smart Networks and Services, FBK, Trento, Italy, {e.coronado, rriggio}@fbk.eu

[‡]TU Berlin, Germany, anatolij.zubow@tu-berlin.de

Abstract—Despite the planned operation of enterprise wireless local area networks (WLANs), they still experience unsatisfactory performance due to several inefficiencies. One of the major issues is the so-called sticky user problem, in which users remain connected to an access point (AP) until the signal quality becomes too weak. In this paper, we leverage software-defined networking (SDN) to propose a user association solution for WLANs aiming to mitigate such inefficiencies, thus improving resource utilization. As it is a computationally hard problem, we also design various low-complexity user-AP association schemes that consider not only signal quality but also AP loads and minimum quality requirements for user traffic. Moreover, to provide simultaneous content distribution in a sustainable mode, we propose exploiting link-layer multicasting to decide on user-AP associations. Our analysis via simulations and experimentation on an open-source testbed shows that considering user-AP association jointly with multicast delivery leads to a significant performance increase over the default client-driven approach: the median throughput is 11x higher when all users request the same content and the achieved improvement decreases to 68% for 100 contents. Moreover, due to more efficient use of the airtime, unicast users achieve higher throughput if multicast delivery is exploited.

Index Terms—WLANs, multicast, SDN, Wi-Fi, AP association.

I. INTRODUCTION

Wi-Fi is the de-facto standard in radio access networks for residential and enterprise applications [1]. Despite being planned deployments continuously monitored by expert administrators, enterprise Wi-Fi networks still suffer from sub-optimal operation, especially on high user densities [2], [3]. This sub-optimality can be attributed to several reasons, including the fact that default Wi-Fi operation lets clients select the access point (AP) to connect to based on the received signal strength. But, this approach, known as *client-driven user-AP association*, has many shortcomings. First, a client has local knowledge of the network state and its association scheme considers only signal quality from the APs. Second, user fairness and load balancing among APs cannot be ensured. Third, even if a client is optimally associated, it may not remain optimal if the client moves or the network state changes. Moreover, several enterprise and home applications require to deliver the same content to multiple users using a process known as multicast. In this case, the sub-optimal user association problem is aggravated due to the absence of acknowledgements and error control in IEEE 802.11 [4].

Software-defined networking (SDN) has changed the classical architecture of wireless local area networks (WLANs)

towards a more flexible approach in which a global view of the network is made available at a software-defined controller [5]. In SDN, the virtual access point (VAP) concept enables active user steering and seamless handover [6], [7], [8], opening the door to more sophisticated user association solutions. However, most of the prior work target unicast traffic [9], [10], [11], and not many solutions leverage multicasting [12], fail to consider applications requirements or assume a single multicast service [13]. A user-AP association exploiting link-layer multicast delivery can be a suitable approach in dense WLANs requiring simultaneous content distribution to a group of users, e.g., in scenarios such as live events, in-flight video entertainment, and education courses, where most users tend to connect to Wi-Fi networks through mobile devices. It remains still an open issue to deliver multicast content efficiently in WLANs, which we investigate in this paper jointly with user-AP association. Our solution aims to be a suitable and effective deployment option to practical use cases involving a large number of users, services and mobility patterns, while being directly compatible with existing Wi-Fi devices in the market without sacrificing innovation and performance.

In this context, the contribution of this paper is three-fold. First, we extend our previous work in [14], where we proposed several centralized user-AP association heuristics considering various system parameters such as handover latency and AP load, by introducing software-defined link layer multicast delivery supporting adaptive user transmission rate. Second, we exploit network softwarization to propose a proactive solution that jointly performs efficient multi-content transmission, load balancing, and seamless user-AP association by considering not only AP loads, but also traffic requirements and user signal-to-noise-ratio (SNR). Finally, we extensively analyse the feasibility of our solution via simulation as well as in a practical setting leveraging an open-source SDN-based platform [7] fully compliant with the IEEE 802.11 standard.

The remainder of the paper is organized as follows. Section II reviews the related work, while Section III describes the system model. Next, Section IV introduces the user-AP association problem formulation maximizing the network-wide proportionally-fair throughput. Section V details various low-complexity heuristics for the problem while Section VI discusses user association in multicast scenarios. Section VII and VIII report on conclusive simulated and experimental results. Finally, Section IX draws the conclusions.

II. RELATED WORK

We can categorise the related work into two main strands. The former comprises research on user-AP association in enterprise WLANs, while the latter considers multicast delivery for deciding on user-AP association. Table I summarizes the characteristics of the most relevant related work.

User-AP association in enterprise WLANs. Depending on the entity responsible for the association decision, schemes can be centralized or distributed. Centralized solutions, e.g., [2], [10], [14], [23], [27], [33], have gained popularity with the increasing adoption of SDN in WLANs [5]. Moreover, centralization brings ease-of-management and improved security and control [34]. A network controller collects network state information and optimizes association considering various goals, e.g., aggregated throughput, user fairness [16], AP load balancing [17], [28], [35]. Such an SDN-based solution is Wi-Balance [18], which jointly manages user association and channel assignment. AP selection can be proactive [27] or reactive [8], [19], [20], [36]. Machine learning techniques are of great help for proactive approaches such as ABRAHAM [27] where a supervised learning model takes various performance indicators to find an efficient user-AP mapping. Several key considerations in AP selection can be listed as: signal quality, AP bandwidth [11], [37], traffic requirements, load balance [38], heterogeneity of users [21], and handover cost [14], [20]. Although the signal quality could be high, the AP might have many other users to serve, resulting in low user throughput. Hence, considering AP load is important in AP selection [23]. While mobility is hard to model, and simplified in most studies, analysis of network traces in [38] shows that users move in unison as a result of their social relations and ignoring this fact might lead to overloaded APs. Authors propose to distribute users with a tight social relation to different APs so that traffic load is fairly balanced.

Distributed solutions, e.g., [29], [6], [30], [31] are easier to implement and may not require changes to the standard. However, as Ercetin [29] shows, users acting on their self-interest might lead to an undesirable operation point, e.g., with high price of anarchy for load-based association. Conversely, authors in [6] propose adaptive probe scheduling for Wi-Fi clients to collect information on the neighbouring AP loads so that each station can decide on the best AP to associate with. Since switching might result in excessive number of re-associations with only marginal performance improvement, most of the solutions perform the change only if the promised rate exceeds a certain margin. Hence, different degrees of cooperation—between users, between APs, or even between network operators [32]—are possible. Notice that given the diversity of Wi-Fi equipment, centralized solutions might be easier to implement and more scalable compared to the changes at the client side. In general, solutions should not require changes to the standard or the client hardware.

Comparing our paper to these earlier works, our solution is a centralized, proactive method that considers not only AP loads or signal strength of each user-AP link but also users traffic requirements. This later aspect is mostly missing in the prior work. Most of the schemes aim at providing the highest

throughput regardless of such requirements. In addition to the analysis via simulation, we show the feasibility of our proposals by implementing them on an open-source testbed.

Multicast in WLANs for user-AP association. In scenarios where traffic is simultaneously delivered to multiple users, multicast on link layer is worthy to be explored in user-AP association and load balancing solutions to improve resource efficiency. Nevertheless, several problems of this mode need to be addressed to ensure that the content is transmitted at a rate that can be decoded by all receivers. However, due to user heterogeneity, ensuring resource efficiency and user satisfaction becomes challenging as the multicast rate is based on the weakest link. Moreover, automatic repeat-request (ARQ) mechanism is disabled, hence suppressing ACKs, as it would result in collisions at the sender [39]. To mitigate these issues, earlier research proposes pseudo-broadcast [39], leader-based solutions [40], [41], forward-error-correction (FEC) [42], [43] and improvements to the IEEE 802.11aa amendment [44]. Experimental analyses also confirm that user feedback, e.g., direct multicast service (DMS) in 802.11aa, increases multicast efficiency by facilitating data rate adaptation [45].

Coronado et al. [13] exploit an alternating DMS-legacy mode to collect link statistics through an SDN controller, which selects the modulation and coding scheme (MCS) for legacy multicast, thus reducing resource utilization. A later work of the same authors [24] extends [13] to improve scalability. In contrast to our paper, multicast group formation for user association is out of the scope of these works. Authors in [25] exploit smart antennas in a link-layer solution for multicast transmission to multiple groups concurrently. JurCast [12] focuses on the application layer and proposes a joint user and rate allocation scheme for video multicast. However, the approach only applies to video content and does not enforce any load balancing policy. Conversely, authors in [26] use multicast to reduce channel utilization and improve load balancing. Nevertheless, unicast delivery is not considered and it is assumed that the data rate from each AP is known.

Comparing our paper to these earlier works, our solution jointly solves user-AP association and multicast group formation considering the proportionally-fair network aggregate throughput. While we acknowledge that our multicast approach does not ensure reliable delivery of the multicast content, we believe that the experimental analysis in our test bed shows yet the feasibility of our solution as well as its higher performance compared to other schemes.

III. SYSTEM MODEL

We consider an enterprise WLAN with a central controller that implements SDN functionalities and is connected to all APs through a high speed Ethernet link. To represent the changes in the network state, e.g., user locations or channel states, we assume a time-slotted setting as in Fig. 1, where the time slot¹ duration reflects the dynamics of a WLAN. In this setting, the controller is in charge of user-AP association decisions and airtime allocation for every user associated to

¹Please note that the time slot represents the time granularity of the actions at the controller and APs, and it does not refer to the WiFi time slots.

TABLE I

OVERVIEW OF THE RELATED WORK. SDN: SDN SUPPORT, UA: USER ASSOCIATION, LB: LOAD BALANCING, MT: MULTICAST TRANSMISSION, CA: CHANNEL ASSIGNMENT, MS: MOBILITY SUPPORT, SH: SEAMLESS HANDOVER.

Ref.	Approach	Input Factors	SDN	UA	LB	MT	CA	MS	SH
[2]	Experimental	Channel util. per AP Expected user-AP Th. RSSI of probe requests	✗	✓	✓	✗	Different	✓	✗
[10]	Simulation Experimental	RSSI of probe requests	✗	✓	✗	✗	Different	✓	✓
[13]	Experimental	RSSI of beacons Channel util. per AP	✓	✗	✓	✓	Different	✓	✓
[14]	ILP formulation	Network utility Users-AP airtime SNR of probe requests	✗	✓	✓	✗	Shared or different	✗	✗
[15]	NLP formulation	Users per AP AP traffic load Expected user-AP th.	✗	✓	✗	✗	Shared or different	✗	✗
[16]	NLP formulation	BW fairness index SINR of beacons User (traffic) priority	✗	✓	✓	✗	Shared	✓	✗
[17]	Min-Max load alg.	TX power of beacons AP traffic load Users in range per AP	✗	✓	✓	✗	Different	✗	✗
[18]	Experimental	RSSI of probe requests Channel util. per AP	✓	✓	✓	✗	Different	✓	✓
[19]	Experimental	RSSI of probe requests Distance user-AP AP traffic load	✓	✓	✓	✗	Different	✓	✓
[20]	MILP formulation	Users in range per AP Users-APs MCS Expected user-AP th.	✗	✓	✗	✗	Different	✓	✓
[8]	Experimental	SNR of beacons SNR of probe requests Users activity time	✓	✓	✓	✗	Different	✓	✓
[21]	Simulation	RSSI of beacons Users-APs MCS Users in range per AP	✗	✓	✗	✗	Different	✗	✗
[22]	Simulation	RSSI level Users location	✗	✓	✓	✗	Different	✗	✗
[23]	Experimental	SNR of probe requests Channel util. per AP	✓	✓	✗	✗	Different	✗	✗
[24]	Experimental	RSSI of beacons Channel util. per AP	✓	✗	✓	✓	Different	✓	✓
[25]	ILP formulation Experimental	Users per AP SNR of beacons Packet reception ratio	✗	✓	✗	✓	Different	✗	✗
[12]	ILP formulation Experimental	PSNR of the video User BW demand RSSI of beacons	✓	✓	✗	✓	Different	✓	✗
[26]	Simulation	Users in range per AP Users MCS Channel util. per AP	✗	✓	✓	✓	Different	✗	✗
[27]	Experimental	RSSI of probe requests AP traffic load Location	✓	✓	✓	✗	Shared	✓	✓
[6]	Markov model ILP formulation	User BW demand Time between RTS RSSI of probe requests	✗	✓	✗	✗	Different	✗	✗
[28]	Simulation	Users per AP Channel util. per AP	✓	✓	✗	✗	Different	✗	✗
[29]	Congestion games Min-Max airtime alg.	Users MCS Users per AP Channel util. per AP	✗	✓	✗	✗	Different	✗	✗
[30]	QP formulation Simulation	RSSI of probe requests Users in range per AP	✗	✓	✓	✗	Different	✓	✗
[31]	ILP formulation	Channel energy level SINR of beacons	✗	✓	✓	✗	Shared or different	✓	✗
[32]	ILP formulation	Users per AP Users MCS Users-AP airtime	✗	✓	✗	✗	Shared or different	✓	✓

Centralized

Distributed

each AP. While the former decision concerns which AP should serve a user, for the latter, the controller determines the fraction of time an AP should serve each of its users. We refer to this service time per user as the user's *airtime*. Please refer to Table II for the key notations.

To control the overhead associated with the operation of the controller, as the latter needs certain information from the users in the network, we assume that the controller is active only at certain times, e.g., time slots 1, 6, 11 in Fig. 1. The controller collects statistics from the APs at the beginning

TABLE II
SUMMARY OF KEY PARAMETERS

Symbol	Description	Symbol	Description
$\mathcal{A} = \{AP_j\}, K$	Set of access points, number of APs	$d_{i,j}$	Distance of u_i from AP_j
$\mathcal{U} = \{u_i\}, n$	Set of users, number of users	$r_{i,j}$	Capacity of the channel between u_i and AP_j
\mathcal{U}_j, n_j	Set and number of users associated to AP_j	$R_{i,j}$	Throughput of u_i if associated to AP_j
B	Bandwidth available at each AP	$\mathcal{R}_{i,j}$	Utility of u_i if associated to AP_j
a_j^f	Binary variable yielding 1 if AP_j uses channel f	r_i^{min}	Min. throughput required for u_i 's application
P_j	Transmission power of AP_j	$\alpha_{i,j}^{min}$	Needed min. airtime for u_i if served by AP_j
C, F	Set and number of channels	α_j	Airtime AP_j allocates to each of its users
$x_{i,j}$	Binary decision variable showing if user i is assigned to AP_j	$\alpha_{i,j}$	Airtime for u_i who is already associated to AP_j
$v_{i,j}$	Binary state variable showing whether u_i is in AP_j 's service region	T	Period of information exchange between the controller and APs
σ	Path loss coefficient of the environment	$\gamma_{i,j}$	Received signal strength of AP_j at u_i

of each *controller period* after which it may change user-AP associations if they improve network-wide utility. We denote the controller period by T time slots, which represents the number of time slots elapsed between the two consecutive operations of the controller, e.g., $T = 5$ in Fig. 1. When the controller is not active, e.g., in time slots 2, 3, 4, 5, 7 in Fig. 1, each user performs handover and association according to its own local decision scheme, which we refer to as *client-driven approach* (CD). In CD, a user associates to the AP with the highest SNR regardless of the AP load and performs handover only if the signal level is below the handover SNR threshold. A network administrator can set the controller period based on the expected mobility of the users, e.g., for highly mobile settings, shorter periods should be preferred for the controller to react promptly to changes in user locations [14].

Let $\mathcal{A} = \{AP_j, \dots, AP_K\}$ be the set of APs and $C = \{1, \dots, F\}$ the set of channels these APs can operate on. As enterprise networks are planned to ensure high service continuity and high capacity for mobile users, AP coverage areas are mostly overlapping [46]. We assume that each AP covers a circular region of radius r meters. We can represent the network topology by a connectivity graph $G = (\mathcal{A}, \mathcal{E})$ where APs are abstracted as the vertices and an edge, $e \in \mathcal{E}$, between two APs, e.g., AP_j and AP_k , means that the two APs have some overlapping coverage region. We represent the channel allocation with $\mathcal{F} = [a_j^f]$, where a_j^f yields 1 if AP_j is assigned channel f for its operation. We denote by \mathcal{A}_f the set of APs operating on channel f , to which we refer as *co-channel APs*. Note that an efficient channel assignment should guarantee minimal interference among co-channel APs [47]. We assume that two APs in the interference range are assigned orthogonal channels in the network planning phase.²

Let $\mathcal{U} = \{u_i, \dots, u_n\}$ denote the set of n users in this network. As users might have different traffic types, e.g., voice vs. video, their throughput requirements may also vary. User traffic can be represented as a vector of tuples $\Gamma = [\langle c_i, r_i^{min} \rangle]$, where c_i and r_i^{min} are the requested content by user i and the minimum rate required to deliver this content, respectively. We assume that each user informs the AP about its minimum rate requirements. Moreover, we

²For dense settings where the APs cannot be assigned orthogonal channels, our model can be modified by reflecting the airtime available to each AP.

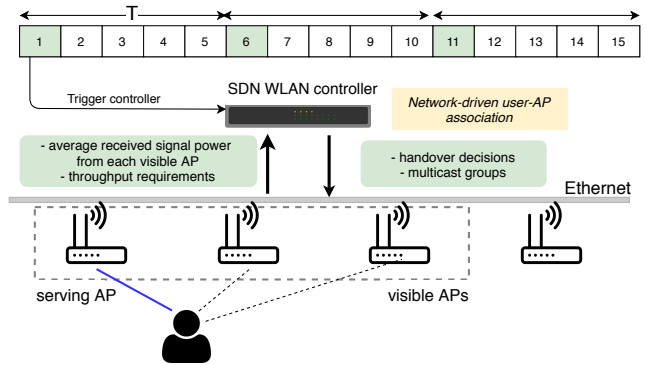


Fig. 1. Dense WLAN setting, where the controller collects channel state statistics and traffic requirements for optimal user-AP association.

assume that uplink traffic is negligible and focus only on the saturated downlink traffic.

As a result of overlapping cells, a user u_i may be in the coverage range of a number of APs. Let $d_{i,j}$ denote the distance between u_i and AP_j . We call the set of APs a user overhears as *visible APs*, i.e., received signal power from the AP is above the receiver sensitivity of the station. We denote the visibility of AP_j at u_i by a binary variable $v_{i,j}$, which yields value 1 if AP_j is a visible AP for u_i . Some users might have mobility with low speed and may change their location at the beginning of a time slot. The SDN controller collects user's statistics such as SNR and traffic requirements by building a traffic matrix for each user-AP pair. This matrix stores for each pair the received and transmitted bytes and packets, and maintains real-time and historical information, e.g., exponentially weighted moving average (EWMA), which are updated over time. This information is then used as the basis for triggering users switching from one AP to another.

We call the users whose visible AP set has more than one AP as *handover candidates*, \mathcal{U}^{ho} . The rest of the users, i.e., $u_k \in \mathcal{U} - \mathcal{U}^{ho}$, are either under outage or have only one visible AP. For the former, there is nothing a controller could do, and in fact, this case should rarely occur under a careful coverage planning, i.e., sufficiently dense AP deployment. For the latter, the controller assigns the user to its only option.

In the next section, we present several ways the controller can decide on these handover events, i.e., user-AP association. The controller reports the user-AP association decision to the APs, after which the users requiring a re-association are switched to the new APs. As prior research [48] has shown, SDN controllers ensure a smooth handover preventing service disruptions due to association changes.

IV. OPTIMAL USER-AP ASSOCIATION

In this section, we first model the throughput of a user under a particular AP association setting. Next, we formulate optimal user-AP association as a network-wide throughput maximization problem.

A. Throughput for a User-AP Assignment

We can calculate the capacity of the downlink channel between AP_j at u_i , denoted by $r_{i,j}$, as a function of the signal-to-noise-plus-interference ratio (SINR) of AP_j 's signal received by u_i . More formally, SINR of an AP signal operating at channel f with bandwidth B equals to:

$$\gamma_{i,j} = \frac{P_j d_{i,j}^{-\sigma}}{B\eta_0 + \sum_{k \in \mathcal{A}_f} P_k d_{i,k}^{-\sigma}}, \quad (1)$$

where P_j denotes the transmission power of AP_j , σ is the path loss coefficient, and η_0 is noise power per unitary bandwidth. A link's SINR is a function of the channel assignment decision, which is reflected in \mathcal{A}_f —the APs assigned to channel f . The capacity of the channel with bandwidth B units is:

$$r_{i,j} = B \log(1 + \gamma_{i,j}) \text{ bits per second.} \quad (2)$$

If u_i is associated to AP_j , its downlink throughput is a function of $\gamma_{i,j}$ and the airtime it will get from AP_j .³ As we assumed orthogonal channel assignment, there is only one AP operating at a particular Wi-Fi channel in a collision domain. Hence, each AP uses all the airtime itself without sharing it with other APs. In case there are multiple APs operating at the same channel and within sensing range of each other, the airtime of an AP will be $1/K_{coch}$, where K_{coch} is the number of co-channel APs. Moreover, due to Wi-Fi's medium access inefficiency, the airtime that can be used for data transmissions is lower than 1. While in the following we will assume all airtime is used for data transmission, one can simply adjust the airtime for data transmission by accounting for MAC layer overhead, i.e., $\alpha = \alpha\Delta$ where Δ is the MAC layer efficiency accounting for the time lost due to backoff in the downlink and number of co-channel APs.

Let us assume that an AP allocates the airtime equally among its users. So, each user gets from its AP $\alpha_j = \frac{1}{n_j}$ fraction of the airtime, if there are total n_j clients connected to this AP. Then, we can calculate the expected throughput of u_i from AP_j , denoted by $R_{i,j}$, as:

$$R_{i,j} = r_{i,j} \alpha_j \text{ bits per second.} \quad (3)$$

³We use Shannon's capacity formula to calculate the rate of this user-AP link. However, the actual rate also depends on the selected MCS.

Although an AP allocates its airtime equally among its users, the actual airtime a user needs may differ across users. For instance, a user browsing the web would need less airtime compared to another having video conferencing. We denote the airtime need of a user as $\alpha_{i,j}^{min}$ and calculate it for u_i from AP_j with rate requirement r_i^{min} as:

$$\alpha_{i,j}^{min} = \frac{r_i^{min}}{r_{i,j}}. \quad (4)$$

Note that satisfying $\alpha_j \geq \alpha_{i,j}^{min}$ inequality is essential for some applications such as video communications. For best-effort traffic, we set $r_i^{min} = 0$. If the throughput a user gets from its AP is at least equal to the requested minimum throughput, we call this user a *satisfied user* and define its utility as a function of its throughput. For unsatisfied users, the utility is zero as the user cannot get the bare minimum for a pleasant user experience, e.g., a user might have low quality of experience due to buffer stalls for video streaming. To reflect the two goals of our controller, i.e., high throughput efficiency and user fairness, we define the utility of a user as its logarithmic throughput [49]. More formally, utility of a user u_i connected to AP_j is defined as:

$$\mathcal{R}_{i,j} = \begin{cases} \log(1 + R_{i,j}), & \text{if } \alpha_j \geq \alpha_{i,j}^{min} \\ 0, & \text{otherwise.} \end{cases} \quad (5)$$

In (5), a user is satisfied if its minimum airtime requirement is allocated and unsatisfied otherwise.

B. Problem Formulation

Let $x_{i,j}$ denote the binary decision variable yielding value 1 if the controller assigns user u_i to AP_j . We formulate the centralized optimal user-AP assignment problem as follows:

$$\mathbf{P1} : \max_{\mathbf{X}=[x_{i,j}]} \sum_{AP_j \in \mathcal{A}} \sum_{u_i \in \mathcal{U}} \log(1 + x_{i,j} r_{i,j} \alpha_j) \quad (6)$$

$$\sum_{AP_j \in \mathcal{A}} x_{i,j} \leq 1 \quad \forall u_i \in \mathcal{U} \quad (7)$$

$$x_{i,j} \leq v_{i,j} \quad \forall u_i \in \mathcal{U}, \forall AP_j \in \mathcal{A} \quad (8)$$

$$\sum_{u_i \in \mathcal{U}} x_{i,j} \alpha_{i,j}^{min} \leq 1 \quad \forall AP_j \in \mathcal{A} \quad (9)$$

$$x_{i,j} \alpha_{i,j}^{min} \leq \alpha_j, \forall u_i \in \mathcal{U}, \forall AP_j \in \mathcal{A} \quad (10)$$

$$\alpha_j = \frac{1}{\sum_{u_i \in \mathcal{U}} x_{i,j}}, \forall AP_j \in \mathcal{A} \quad (11)$$

$$x_{i,j} \in \{0, 1\} \quad \forall u_i \in \mathcal{U}, \forall AP_j \in \mathcal{A}. \quad (12)$$

The objective function in (6) states that users must be associated to the APs that result in the highest network utility which is a function of logarithmic throughput maintained by each user. Const. (7) signifies that each user can be associated to at most one AP, whereas Const. (8) is necessary for a feasible assignment, i.e., a user can only be associated to an AP that is within its receive range. Const. (9) states that the minimum airtime demand of associated users cannot exceed the capacity of an AP, i.e., 100% airtime, whereas Const. (10) ensures that the airtime share this user will get from an AP is higher than the minimum airtime needed to satisfy the

application requirements of this user. Finally, (11) formally states the airtime a user gets from its AP and Const. (12) denotes that each assignment variable is a binary variable.

The controller solves **P1** every T time units and sends the user-AP association vector, denoted by $X = [x_{i,j}]$ where $x_{i,j} = 1$, to all APs. During other times, the CD approach is effective. That is, if a client moves from its current location for which the controller has solved the association problem, the client might not receive sufficient signal from its associated AP. Then, the AP with the highest signal quality is selected regardless of the AP's load. Complexity of **P1** depends on the number of users in the cell edges with overlapping AP coverage and the number of APs visible to each user. In particular, it increases exponentially with the number of users: since in the worst case the number of users in these regions is n and the number of APs each user can get service from is K , complexity is $\mathcal{O}(K^n)$. The high complexity of **P1** renders it infeasible for practical operation. Hence, in the next section we present the design of lower complexity heuristics.

V. LOW-COMPLEXITY USER-AP ASSOCIATION SCHEMES

In this section, we present a set of low-complexity heuristics based on two design goals: (i) minimum throughput requirements must be satisfied, and if not, the utility of an unsatisfied user is zero; and (ii) a heuristic must have polynomial complexity in the number of users and APs. For all schemes, the controller first identifies handover candidates (\mathcal{U}^{ho}) and assigns the rest to their only visible AP, if any.

A. Highest-SNR AP Association (h-SNR)

A simple association scheme a controller can implement is to assign users to the AP with the highest signal strength. Note that the conventional scheme, namely CD, follows the same approach. However, CD triggers a handover only after the AP cannot provide the minimum signal level or after the user has completely disconnected from the AP. In h-SNR, the user does not stick to its AP but instead switches to the AP with the highest signal level, which allows us to compare the benefit of periodic handover management to CD handovers. While h-SNR mitigates the sticky user problem, it does not consider minimum rate requirements. The computational complexity of h-SNR is $\mathcal{O}(K \cdot n)$ since for each user the algorithm must iterate over the APs in range, which is independent across users and executed K times in the worst case.

B. Airtime-aware AP Association (AIR)

While a high SNR value ensures high link rate, it overlooks the time-sharing nature of Wi-Fi at the medium-access control layer. The throughput of a user is a product of its link rate and the airtime it receives. To account for both parameters, we design airtime-aware AP association (AIR). To have some notion of fairness, AIR starts with a randomly-picked user and assigns to it the AP that promises highest (airtime \times rate). After each assignment, the airtime a user can get from each AP is updated by considering the new number of associated users for each AP. Note that AIR does not consider minimum rate requirements. The procedure is repeated for each user, which incurs a complexity of $\mathcal{O}(K \cdot n)$.

C. Demand-aware AP Association (DAW)

While AIR calculates expected throughput, it does not consider how a new association affects the performance of existing users of the AP. Given that some users require minimum rate to have a satisfactory quality of experience, demand-aware AP association (DAW) avoids violating the minimum rate requirements. Similar to airtime-aware scheme, DAW first calculates (air-time \times rate) for each user-AP pair as in (3). However, DAW also checks if an AP has *spare* airtime. Given that an AP allocates its airtime equally among its users, the number of users it can serve is limited by the minimum bandwidth requirements. Let \mathcal{U}_j denote the set of users served by AP_j and total airtime demand from users of this AP as $\alpha_j^{min} = \sum_{u_i \in \mathcal{U}_j} \alpha_{i,j}^{min}$. While an AP with $\alpha_j^{min} = 0$ can in theory serve unlimited number of users, an AP with $\alpha_j^{min} > 0$ can serve at most n_j^{max} users to continue satisfying the minimum rate requirements of its users. Here, n_j^{max} is determined by the most-demanding user. More formally, we define n_j^{max} as follows:

$$n_j^{max} = \lfloor \frac{1}{\max_{u_i \in \mathcal{U}_j} \{\alpha_{i,j}^{min}\}} \rfloor. \quad (13)$$

Then, spare airtime s_j this AP can allocate to a new user is:

$$s_j = \begin{cases} \frac{1}{n_j+1}, & \text{if } n_j < n_j^{max} \\ 0, & \text{otherwise.} \end{cases} \quad (14)$$

Let \mathcal{A}^+ be the set of APs that can accommodate new users without violating the minimum rate requirements of the existing users, i.e., $\mathcal{A}^+ := \bigcup AP_j$ where $s_j > 0$ for $AP_j \in \mathcal{A}$. Our aim is to switch handover candidates to such APs in \mathcal{A}^+ .

Algorithm 1: Demand-aware Association (DAW)

- 1: Input: user-AP link qualities $\gamma_{i,j}$, traffic profiles Γ
 - 2: Output: user-AP association \mathcal{U}_j for each AP_j
 - 3: $\mathcal{U}^j \leftarrow \emptyset$, $n_j = 0$, $\forall AP_j \in \mathcal{A}$
 - 4: **while** $\mathcal{U}^{ho} \neq \emptyset$ and $\mathcal{A}^+ \neq \emptyset$ **do**
 - 5: **for** $u_i \in \mathcal{U}^{ho}$ **do**
 - 6: **for** $AP_j \in \mathcal{A}$ and $v_{i,j} == 1$ **do**
 - 7: Get $\alpha_j = 1/(n_j + 1)$
 - 8: Calculate $\mathcal{R}_{i,j}$ using α_j as in (3)
 - 9: $\Delta \mathcal{R}_{i,j} = \mathcal{R}_{i,j}$
 - 10: **for** $u_k \in \mathcal{U}_j$ **do**
 - 11: $\Delta \mathcal{R}_{i,j} = \Delta \mathcal{R}_{i,j} + \log(\frac{1+r_{k,j}\bar{\alpha}_j}{1+r_{k,j}\alpha_j})$
 - 12: $(i^*, j^*) = \arg \max \Delta \mathcal{R}_{i,j}$
 - 13: $\mathcal{U}_{j^*} = \mathcal{U}_{j^*} \cup u_{i^*}$ and $n_{j^*} = n_{j^*} + 1$
 - 14: $\mathcal{U}^{ho} \leftarrow \mathcal{U}^{ho} - u_{i^*}$
 - 15: Update s_{j^*} using (14) and update \mathcal{A}^+
 - 16: **return** \mathcal{U}_j
-

For a candidate AP_j we also need to consider the decrease in aggregated throughput of users in \mathcal{U}_j after a new user joins. As the number of users increases by 1 in AP_j 's service region, existing users of AP_j will have less airtime and therefore will sustain lower throughput. Since our aim is to maximize the utility in (6), we calculate the change in the utilities of the users due to the change in the airtime values. Let α_j represent




Fig. 2. A toy example showing the impact of u_3 's association to AP-1 in unicast/multicast mode or to AP-2 in unicast mode. In Scenario-1 and Scenario-2, u_3 is closer to AP 1. In Scenario-3 and Scenario-4, u_3 is closer to AP 2.

the airtime if n_j users are connected to this AP and $\bar{\alpha}_j$ for $n_j + 1$ connected clients. Moreover, let $\Delta\mathcal{R}_{i,j}^-$ represent the decrease in total utility of the already-connected users when u_i joins AP_j . We calculate it as follows:

$$\Delta\mathcal{R}_{i,j}^- = \sum_{u_k \in \mathcal{U}_j} \log\left(\frac{1+r_{k,j}\bar{\alpha}_j}{1+r_{k,j}\alpha_j}\right). \quad (15)$$

Then, considering user i 's expected utility, we calculate the net utility of assigning user i to AP_j as:

$$\Delta\mathcal{R}_{i,j} = \log(1+r_{i,j}\bar{\alpha}_j) + \Delta\mathcal{R}_{i,j}^-. \quad (16)$$

Then, we take the pair user i^* and AP j^* achieving the highest net utility: $i^*, j^* = \arg \max \Delta\mathcal{R}_{i,j}$. We update \mathcal{U}^{ho} by removing the assigned user u_{i^*} from the set of handover candidates. To avoid violating minimum rate requirements, we update the spare capacity of each AP in \mathcal{A}^+ using (13) and (14), and remove those APs with zero spare capacity from \mathcal{A}^+ . The operation of DAW, sketched in Alg. 1, terminates when \mathcal{A}^+ or \mathcal{U}^{ho} equals to empty set. The computational complexity of DAW is $\mathcal{O}(K \cdot n^2)$.

VI. MULTICAST-AWARE USER-AP ASSOCIATION (MAA)

When multiple users are interested in the same content, multicasting is desirable to save from redundant transmission. For video applications, e.g., a common use case in corporate settings is video conferencing, saving by multicasting has a huge potential. Treating such users as a single bundle may have two consequences. On the positive side, airtime allocated per user increases. For example, if two users are in the same multicast group in a network of n users, the airtime is divided into $n - 1$ transmissions instead of n unicast transmissions. Hence, the airtime increase per user is $\frac{1}{n(n-1)}$. On the negative side, the transmission parameters must be adapted to the weakest link in the multicast group to guarantee successful decoding at each user. As a result, users with good links may maintain lower data rate compared to unicast. Considering these pros and cons, an SDN controller decides on whether delivering the traffic in multicast mode or unicast mode.

Fig. 2 illustrates a toy example showing different association options for user 3 denoted by u_3 . Scenario-1 and Scenario-2 depict a case where u_3 is located closer to AP₁ whereas u_3

TABLE III
DATA RATES AND UTILITY OF THE EXAMPLE SCENARIOS IN FIG. 2.

Sc.	u_3 's connection	u_1	u_2	u_3	u_4	Total	Utility
1	AP ₁ unicast	4	7	5	30	46	3.87
2	AP ₁ multicast	6	10.5	6	30	52.5	4.24
3	AP ₁ multicast	3	10.5	3	30	46.5	3.76
4	AP ₂ unicast	6	10.5	9	15	40.5	4.11

is closer to AP₂ in Scenario-3 and Scenario-4. Table III lists the resulting performance under different AP association and delivery options. In this example, if u_3 gets its content via a unicast connection to AP₁ as in Scenario-1, then its rate is $15/3 = 5$ Mbps due to equal airtime sharing with two other clients whereas the rate of (u_1, u_2, u_4) are (4, 7, 30) Mbps, respectively. As a result, total utility is $\log(1+5) + \log(1+4) + \log(1+7) + \log(1+30) = 3.87$. If u_3 and u_1 are put in a multicast session in AP₁ as in Scenario-2, then the rates for (u_1, u_2, u_3, u_4) are as follows: (6, 10.5, 6, 30) Mbps resulting in a higher utility of 4.24. Here, multicast delivery improves the utility compared to unicast option. However, there might be cases (as depicted in Scenario-3 and 4) where the airtime gain may not justify the loss in a user's rate as a result of rate adaptation according to the weakest link in the multicast group. Therefore, unicast delivery might be preferred mode of service in such scenarios as reflected with a higher utility in Scenario-4 in comparison to the utility of Scenario-3.

Let us define a multicast group as a set of users that request the same content in the coverage region of the same AP and will be served in a single session by this AP. More formally, m^{th} multicast group in the coverage of AP_j is defined as: $G_m^j = \{u_i, \dots, u_k\}$ where $c_i = c_k, \forall u_i, u_k \in G_m^j$. A single user is also trivially a multicast group with only one user. Note that two users with the same content request can be in different multicast groups even if they are connected to the same AP. We will denote the set of all multicast groups of AP_j by $G^j = \{G_1^j, \dots, G_m^j\}$ and the number of groups in AP_j by g_j . As an AP serves the users in a multicast group via a single downlink session, we calculate airtime simply as $\alpha_{m,j} = \frac{1}{g_j}$, assuming equal airtime for each multicast group.

The downlink rate that AP_j will transmit to G_m^j is given by the SNR of the weakest link in this group. We calculate the downlink rate $r_{m,j}$ as in (2), where we replace $\gamma_{i,j}$ by

Algorithm 2: Multicast-aware association (MAA)

```

1: Input: user-AP link qualities  $\gamma_{i,j}$ , traffic profiles  $\Gamma$ 
2: Output: multicast groups  $G^j$  for each  $AP_j$ 
3:  $G^j \leftarrow \emptyset$  and  $g_j = 0 \quad \forall AP_j \in \mathcal{A}$ 
4: while  $\mathcal{U}^{ho} \neq \emptyset$  do
5:   for  $u_i \in \mathcal{U}^{ho}$  do
6:     for  $AP_j \in \mathcal{A}$  and  $v_{i,j} == 1$  do
7:       for  $\forall G_m^j$  and  $c_i == c_m$  do
8:         Calculate rate as if  $u_i$  joins  $G_m^j$  as in (17)
9:         Calculate  $\Delta \mathcal{R}_{i,j,m}$  as in (19)
10:        Calculate  $\Delta \mathcal{R}_{i,j,g_j+1}$  as in (21)
11:         $(i^*, j^*, m^*) = \arg \max \Delta \mathcal{R}_{i,j,m}$ 
12:        if  $m^* == g_j + 1$  then
13:          Create a new multicast group,  $G_{g_j+1}^j \leftarrow \{u_i^*\}$ 
14:           $G^j \leftarrow G^j \cup G_{g_j+1}^j$ 
15:          Update airtime and utility for all entries
16:        else
17:           $u_i^*$  joins group  $m^*$  of  $AP_{j^*}$ :  $G_{m^*}^{j^*} \leftarrow G_{m^*}^{j^*} \cup u_i$ 
18:          Update multicast rate of  $G_{m^*}^{j^*}$  as in (17)
19:         $\mathcal{U}^{ho} \leftarrow \mathcal{U}^{ho} - u_i^*$ 
20: return  $G^j \quad \forall AP_j \in \mathcal{A}$ 

```

$\gamma_{m,j} = \min_{u_i \in G_m^j} \gamma_{i,j}$. Consequently, the throughput for $u_i \in G_m^j$ is straightforward as: $R_{i,j} = r_{m,j} \alpha_{m,j}$ bits per second. Utility calculation remains the same as in (5) and all members of a group have the same utility. Next, we introduce our heuristic for user-AP association which exploits the possibility of multicast delivery.

Multicast-aware association (MAA) in Alg. 2 is a greedy approach that considers both multicast and unicast delivery options. When no users are associated, MAA calculates for each u_i and AP_j pair the unicast delivery utility $\mathcal{R}_{i,j,1}$ where 1 stands for the first multicast group belonging to AP_j . Then, MAA picks the user-AP pair (u_{i^*}, AP_{j^*}) with the highest $\mathcal{R}_{i,j,1}$ value, meaning that user u_{i^*} will associate to AP_{j^*} and construct the first multicast group $G_1^{j^*} = \{u_{i^*}\}$. Later, for a user u_i , there are two possibilities for getting service from AP_j , namely in multicast and unicast mode.

The user joins the multicast group m . If there is a multicast group G_m^j whose requested content is the same as c_i , u_i can be served in this multicast group. Consequently, we update the multicast rate of G_m^j as follows:

$$\bar{r}_{m,j} = \min(r_{m,j}, r_{i,j}), \quad (17)$$

where $r_{i,j}$ is the rate of the link between u_i and AP_j . Next, we calculate the decrease in total utility of G_m^j as:

$$\Delta \mathcal{R}_{i,j,m}^- = \begin{cases} |G_m^j| \log\left(\frac{1+r_{m,j}\alpha_{m,j}}{1+\bar{r}_{m,j}\alpha_{m,j}}\right) & \text{if } \bar{r}_{m,j} \neq r_{m,j} \\ 0 & \text{otherwise.} \end{cases} \quad (18)$$

For other multicast groups m' of AP_j , there will be no change in their utility and airtime. Hence, change in network utility if u_i joins AP_j 's multicast group m is:

$$\Delta \mathcal{R}_{i,j,m} = \log(1 + \bar{r}_{m,j}\alpha_{m,j}) + \Delta \mathcal{R}_{i,j,m}^- \quad (19)$$

The user gets service in unicast mode. This is equivalent to a new multicast group being added to AP_j . As a result, the airtime of the existing multicast groups will decrease to

$\bar{\alpha}_{m,j} = \frac{1}{g_j+1}$ which consequently leads to a decrease in utility of the existing multicast groups. The utility decrease for G_m^j denoted by $\Delta \mathcal{R}_{j,m}^-$ is as follows:

$$\Delta \mathcal{R}_{j,m}^- = (|G_m^j|) \log\left(\frac{1+r_{m,j}\bar{\alpha}_{m,j}}{1+r_{m,j}\alpha_{m,j}}\right), \quad (20)$$

while utility of the new multicast group is $\log(1 + r_{i,j}\bar{\alpha}_{m,j})$. Then, similar to (19), utility change is:

$$\Delta \mathcal{R}_{i,j,g_j+1} = \log(1 + r_{i,j}\bar{\alpha}_{m,j}) + \sum_{\forall m, m \neq g_j+1} \Delta \mathcal{R}_{j,m}^- \quad (21)$$

Considering the utility in both cases, i.e., in (19) and (21), the controller selects the association achieving the highest sum utility at each step. This increase in utility represents the main performance gain of MAA which is, at the same time, given by the reduction of airtime utilization. Note that the airtime use is reduced by lowering the number of simultaneously delivered separate streams with respect to n unicast transmissions.

VII. PERFORMANCE ANALYSIS VIA SIMULATIONS

In this section, we analyse the performance of our solutions in a simulated environment using our in-house system level simulator developed in Python. In the following, we present the evaluation setting and discuss the results obtained.

A. Scenario and Setting

Network deployment: Our scenario follows the indoor scenario guidelines defined in IMT [50], e.g., indoor environments isolated from external interference and consisting of stationary or low-mobility pedestrians. We consider a conference-like setting with a large number of people in a meeting room equipped with 3 APs. For the initial placement, 90% of the users are located in this room. Outside the conference hall, there are 7 APs deployed in a grid-like topology to serve the remaining 10% of the users. The total area covers a region of [150 m, 100 m]. The conference hall is located at the centre of the area and its size is [50 m, 30 m].

The total bandwidth is 100 MHz. After solving a graph-colouring problem, we find the bandwidth per AP based on the total number of frequencies needed for this network to avoid inter-AP interference [14]. Since the locations of the users affect the user-AP association decisions, e.g., via the visibility parameter, we introduce a metric called *density balance* to differentiate among different user and AP deployment settings. To compute it, for each AP, we count the number of users closer to this AP compared to other APs. Then, we compute the Jain's fairness index considering user distribution over all APs. For instance, if all APs have equal number of users in their proximity, the system has a density balance of 1. Note that despite different settings, having the same number of users and APs, may have different density balance as it depends (only) on the locations of users and APs. We define *load balance* of APs which represents how uniformly the users are associated to the APs. Load balance is calculated as Jain's fairness index considering the distribution of number of users associated to each AP. In the described setting, the density balance is low as users are concentrated in a small area and




Fig. 3. Impact of increasing number of Wi-Fi clients. 10 content items. Only 0.3 fraction of users are mobile. 50% of users have some traffic with a minimum bandwidth requirement while others are best effort traffic users.

the APs in the conference room have many users to serve while other APs are idle. We set the controller period to 1.⁴ We model the user-AP links using the Keenan-Motley model [51]. The channel SNRs are updated at the beginning of each time slot based on the user’s location and the channel model.

Content model: Since our aim is to analyze multicast scenarios, we consider a small catalogue comprising a number of contents, N_c , varying from 1 to 100. For each user, we assign a content from the content catalogue based on Zipf content popularity distribution with Zipf exponent 1. Unless otherwise stated, we set $N_c = 10$. As for content’s minimum rate requirements, we derive a rate between [5, 15] Mbps for high-bandwidth demanding users.

Mobility model: We consider users moving with speed values uniformly distributed in interval [1, 5] m/s. We assume a Random-Waypoint mobility with 0.3 probability of pausing. After a user pauses or hits the borders, it changes the movement direction with an angle $\sim U(0, 2\pi)$. We set the fraction of mobile users to θ of the users. While we calculate fairness of users’ throughput, we consider the throughput accumulated in a time window of 5 slots. We report the average results of 40 repetitions along with 95% confidence intervals for each scenario and a simulation time of 300 time slots.

Performance metrics: In addition to utility, APs load balance, and aggregated/median throughput, we report the following metrics. We define the *fraction of satisfied users* as the ratio of the number of satisfied users to the number of users. The *improvement ratio* is calculated as: $\frac{\mathcal{M}_H - \mathcal{M}_{CD}}{\mathcal{M}_{CD}} \cdot 100$, where \mathcal{M}_H is the performance metric of interest, e.g., median throughput, under our heuristic $H = \{\text{AIR, DAW, h-SNR, MAA}\}$.

B. Performance Comparison via Simulations

Fig. 3 compares the performance of all schemes under increasing number of users where only 0.3 fraction of users are mobile and 50% of users have some traffic with a minimum bandwidth requirement while the rest are best-effort traffic users. As observed in Fig. 3a, the median throughput is significantly higher if MAA is enabled. Here, the improvement ratio is around 145% for low number of Wi-Fi users as compared to CD. The performance of the remaining schemes in descending order is as follows: DAW, AIR, h-SNR, and CD.

⁴Please see our earlier paper [14] for other settings with users being more uniformly distributed or the controller period being longer than one time slot.

The improvement of h-SNR over CD varies between 5-10% considering the median throughput, while it is between 145-330% for MAA, 32-56% for DAW, and 29-48% for AIR. With increasing number of Wi-Fi users, the performance improvement decreases for unicast schemes. For example, DAW and AIR provides around 30% improvement over CD while it is only 5% for h-SNR for $n = 200$. In contrast, improvement ratio of MAA over CD increases with increasing n , reaching to 330% for $n = 200$. This trend is due to the capability of MAA to deliver the content in multicast mode and to use the airtime efficiently. As Fig. 3b shows, MAA sustains high user satisfaction close to 100%, while fraction of satisfied users degrades markedly for all others. Finally, Fig. 3c shows that MAA’s utility outperforms all others significantly while its running time is also higher than that of h-SNR, CD, and AIR. The running time increases with higher user density for all schemes, but the higher complexity of DAW, i.e., quadratic vs. linear, is confirmed in Fig. 3c. But, MAA takes between 5-70% longer to execute compared to CD while h-SNR and AIR need maximum 7% longer time.

Next, we analyze the impact of increasing fraction of mobile users, as shown in Fig. 4. Intuitively, for a static scenario, we expect the performance gain of periodic association schemes over CD to be lower as the change in user-AP link qualities is minimal compared to a mobile scenario where users’ distance from their APs changes. As Fig. 4a shows, all schemes except MAA perform better with increasing number of mobile users. We attribute this trend to the increase in density balance (not plotted). Initially, all clients are in the conference room that has three APs and the resulting density balance is around 0.3. With mobility, the users might move outside this room and go into footprint of other APs, resulting in a gradual increase in the density balance. So, the AP loads become more balanced and the density when all users are mobile is around 0.55. For CD, as users are sticky, CD has significantly less handovers compared to others as in Fig. 4b. Note that although our controller-based schemes result in slight increase in handover probability, the resulting performance improvement is significant, justifying this increase in handovers. For MAA, the possibility of forming multicast groups becomes lower with users being distributed in a larger area where the channel qualities among users might differ significantly. Consequently, the desirability of the multicast delivery decreases due to the

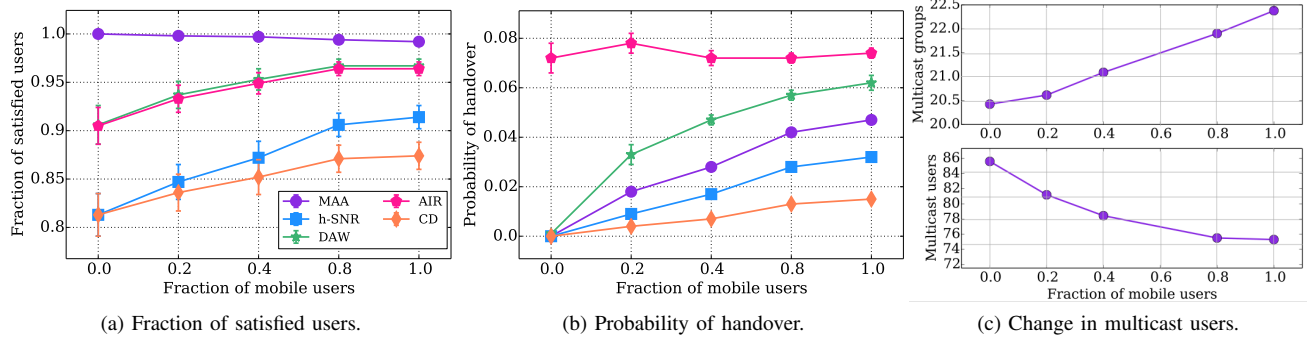


Fig. 4. Impact of fraction of mobile Wi-Fi clients for 100 Wi-Fi clients, 10 content items, Zipf popularity with exponent 1.

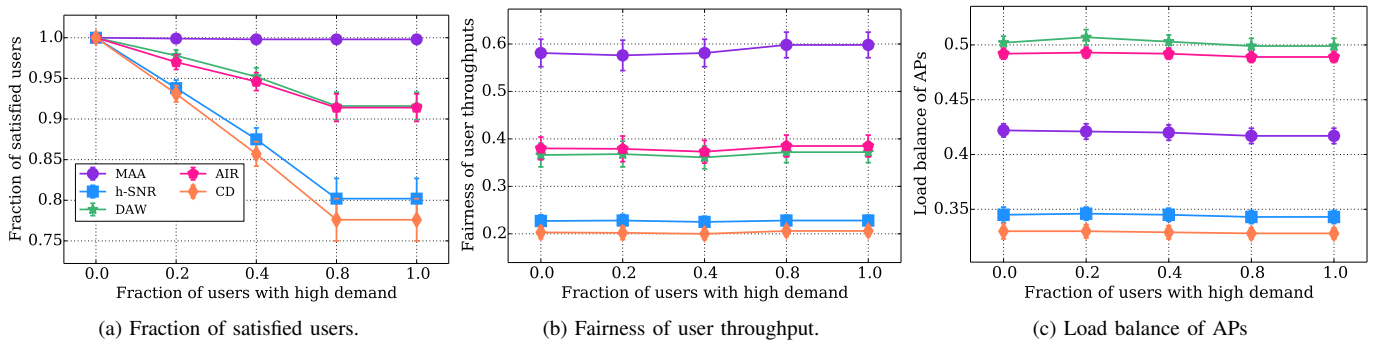


Fig. 5. Impact of fraction of high traffic demanding Wi-Fi clients for 100 Wi-Fi clients and 10 content items.




Fig. 6. Impact of increasing number of content items for Zipf popularity with exponent 1 and 100 Wi-Fi clients.

weakest link's quality determining the multicast rate. As seen in Fig. 4c, number of multicast users decreases while the number of groups increases with increasing mobility. Please note that other mobility patterns might affect performance differently, e.g., if users move together as a group and the density balance does not change significantly with mobility.

Fig. 5 shows the impact of increasing fraction of users with some bandwidth requirements. From Fig. 5a, we observe that only MAA can sustain the satisfaction of users as the congestion in the network is managed by airtime allocation to only a few multicast groups as compared to other unicast schemes which allocate airtime per user. Fig. 5b shows that MAA maintains relatively fair throughput distribution among its users compared to other schemes. The higher fairness is due to the users getting their content in the same multicast group with the same rate. But, resulting AP load balance is lower for MAA compared to DAW and AIR (Fig. 5c).

Finally, Fig. 6 shows the impact of increasing number of contents (N_c). Intuitively, diversity in consumed content de-

creases the chances of multicast delivery thereby its promises. As Fig. 6a shows, when all users request the same content, MAA can introduce 11x improvement in median throughput. With more diverse content, throughput difference between MAA and other schemes decreases. However, throughput gain is still significant, e.g., 1.8x when $N_c = 20$ and 68% for $N_c = 100$. MAA results in a lower AP load balance (not plotted) as multicasting option through an AP is favored over unicast delivery through two APs. With increasing N_c , multicast potential decreases which is reflected as increasing load balance among APs. Fig. 6b plots the distribution of multicast group size in terms of number of users in the group. As expected, under higher N_c , the group size becomes smaller. Lastly, Fig. 6c plots the throughput for multicast and unicast users. While multicast delivery increases the median throughput as apparent in Fig. 6a, this improvement is not at the expense of the unicast users. Interestingly, Fig. 6c shows that unicast users achieve higher performance compared to multicast users. A closer look to SNR values for unicast




Fig. 7. Testing scenario used in the experimental evaluation.

reveals that SNR levels are lower in case of unicast. Then, higher performance of unicast must be due to the increased airtime unicast users are allocated as a result of multicast users getting content all in a single downlink delivery. We argue that the lower SNR of unicast users might be due to our algorithm selecting the AP that might have higher airtime for this unicast user despite its signal being weaker. With higher content diversity, the number of multicast users decreases or the number of multicast groups increases, both resulting in a lower airtime per downlink delivery. Consequently, unicast throughput decreases and approaches to that of multicast.

VIII. PERFORMANCE ANALYSIS VIA EXPERIMENTATION

In this section, we validate our solutions using the open-source SDN-based platform named 5G-EmPOWER [7]. As before, we first introduce the methodology followed for the evaluation, and then we report on the results obtained.

A. Scenario and Setting

Network deployment: The scenario considers similar indoor settings as the ones mentioned in Sec. VII involving stationary and low-mobility users. We select an office building composed of 2 floors with a size of [20 m, 35 m], as sketched in Fig. 7, comprising 2 APs and 5 users. Each AP runs an instance of the 5G-EmPOWER Agent, a software agent used to communicate to the SDN controller (via TCP through an Ethernet link), who is in charge of issuing the management instructions to the APs, e.g., set an MCS, perform a handover, etc. The APs are deployed on PCEngines ALIX 2D (x86) processing boards running OpenWRT 18.05.01. However, any off-the-shelf AP supporting OpenWRT can be also used. The wireless cards are based on the Atheros AR9220 chipset using IEEE 802.11n. The APs are configured on non-overlapping channels. The controller is deployed on an Intel NUC with an i7 Intel processor and 16 GB of RAM running Ubuntu 18.04.1. In 5G-EmPOWER, each AP maintains for each connected client a light virtual access point [7] comprising its management state, e.g., association, authentication, and handover. This abstraction allows the controller to actively migrate a client by simultaneously asking (i) to the target AP, to instantiate a new virtual access point (with the same information); and (ii) to the old AP, to remove it. Consequently, a transparent handover is performed without changing the client’s configuration.

B. Performance Comparison in Testbed

Users setup and radio measurements gathering: Dell laptops equipped with two wireless cards, an Intel i7 CPU and running Ubuntu 18.04.1 are used as clients. The main wireless card is used in managed mode to connect to the AP, while the second one is configured in monitor mode for radio measurements collection. For the evaluation, signal values from 802.11 data frames are required, for which we leverage spectral snapshots based on Fast Fourier Transform (FFT) and time-lapse RF spectrogram available in Qualcomm/Atheros AR92xx and AR93xx chipsets. These NICs have the ability to report FFT information about the received signal [52]. When the spectral scan is triggered, the card performs an FFT every 4 μ s for all sub-carriers. The output of the FFT samples are provided in binary form, which is interpreted using the tool available in [53]. These small samples (less than 8KB), collected at the wireless users, are sent to the 5G-EmPOWER controller through an SDK REST server every 1 s. Following the description in [52], where the card assumes a noise floor level N at -96 dBm, the controller uses this data to report the SNR for a complete HT20 channel [54]. The FFT summary includes the In-phase (I_i) and Quadrature (Q_i) for each subcarrier (56 and 114 bins in 20 and 40 MHz, respectively) and the maximum signal magnitude. Based on that, the absolute magnitude (z_i) for each FFT bin is calculated as: $z_i = I_i + Q_i$. Hence, the signal strength above the noise floor is calculated as follows for each subcarrier: $SNR_i = N + RSSI + 10 \cdot \log(z_i^2) - 10 \cdot \log(\sum_{i=1}^{56} (b(i)^2))$.

Experiments details: The users are added incrementally in the scenario, where *user 1*, *user 2* and *user 3* are located in the 1st floor, and *user 4* and *user 5* in the ground floor. Note that all users are stationary except *user 2*, who follows the path in blue shown in Fig. 7 at an average speed of 1 m/s. This deployment model aims to provide reproducible results and to show the applicability of the solution in mobility scenarios. In fact, since the controller processes user information in order of request arrival, similar results would be obtained under higher mobility. In each experiment, a 40 Mb/s UDP stream is sent from a server at the backhaul to each user, hence doubling the overall load with each new user added. Apart from the downlink transmissions, no other uplink or downlink traffic exists between the APs and the stations. The duration of the experiments is 60 s and experiments are repeated 10 times. Similar results can be extracted for longer tests since it would not determine the behavior of the association policies.

Performance metrics: For the evaluation, we have considered aggregated throughput, throughput per user at the highest congestion level, retransmission ratio, network utility as well as Jain’s fairness index of such utility, and MCS distribution.

Multicast considerations: For MAA, the controller considers both airtime and SNR, similarly to AIR, and assigns several users to the same multicast stream to improve efficiency. However, as mentioned in Sec. VI, IEEE 802.11 sets the lowest data rate for multicast delivery, which highly hampers the performance. To improve this data rate, we have leveraged the multicast solution presented in [13], and configured it to periodically use DMS, i.e., unicast transmissions, in 10%




Fig. 8. Impact on the throughput and retransmission ratio for an increasing number of users (1 mobile user).

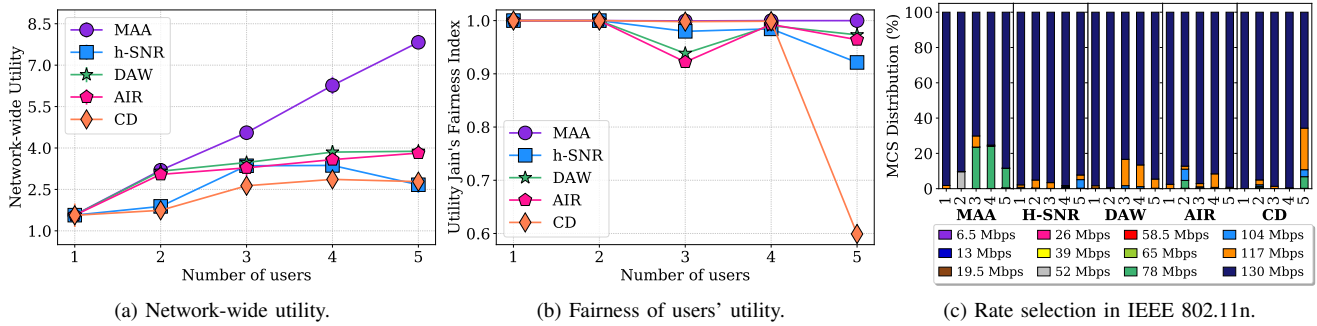


Fig. 9. Impact on the user utility as well as on the data rate selection for an increasing number of users (1 mobile user).

of the time for the rate control statistics collection, and legacy multicast in the remaining 90%. The data rate for each multicast group is estimated using such rate control statistics based on the conclusions of the same work. Note that to avoid excessive computational cost leading to network re-configuration providing minimal performance improvements, before executing MAA (Alg. 2), the SDN controller analyses one by one the SNR reported by the users. Hence, MAA is executed only if the signal strength above the noise floor varies beyond 5 dBm. This check only involves simple arithmetical comparisons, which leads to the computational complexity of $O(n)$ and would not impact the system performance even under high mobility and larger number of users.

The experimental campaign covers the schemes described in previous sections and analyzes the impact of increasing user density and mobility. We remind the reader that in the highest saturation state 20% of the users are mobile. User association for MAA works as follows: (i) for a single user, the content is transmitted in unicast mode; (ii) for two users, the algorithm makes a decision based on the utility achieved by joining them in a single multicast group, and by splitting them into two unicast streams, attending to channel conditions; (iii) in the presence of one multicast group and one unicast transmission, i.e., 3 users, the algorithm considers 2 transmissions for calculating the utility since the multicast stream is seen as a single user from a resource utilization viewpoint; and (iv) for 4 or more users, the reasoning is same as in (iii).

Figure 8 shows the network-wide and the individual throughput for an increasing number of users. In particular, Fig. 8a depicts the aggregated throughput for each policy, where we observe that the performance of CD is the lowest

one in all the cases due to the stickiness of the users. Although h-SNR performs slightly better, not considering the AP loads also determines its low performance. Moreover, the introduction of the mobile user (from 2 users) also hampers the performance of CD and h-SNR due to the AP stickiness effect and the channel variability, respectively. AIR and DAW outperform, in average by 50%, the results of the previous association schemes. However, this figure is still low since the throughput tendency remains at 75 Mbps from 3 users due to the high network load and the resource divisions. By contrast, MAA is able to introduce 2.5x improvement in the network-wide throughput, especially when the number of users increases. From Fig. 8b, which plots the median user throughput in the presence of 5 users, we draw the same conclusions; CD offers the lowest individual throughput and a high dispersion in the results of a user due to the poor load distribution. DAW, h-SNR, and AIR are able to better allocate the users, however the throughput per user is 1.5x higher for the schemes considering airtime utilization. Conversely, MAA increases the throughput per user and shows higher performance equality. The effect of poor user association and load balancing is translated in Fig. 8c into an increase of the retransmission ratio, especially as the number of users increases. MAA presents the lowest retransmission ratio due to the efficient multicast mode [13], where DMS and legacy multicast are used in the 10% and 90% of the time, respectively.

Figure 9 studies the utility in several scenarios and the impact on the data rates selected for transmission. In Fig. 9a, we observe that MAA achieves the highest network-wide utility, which shows an increasing trend with the number of users. As a matter of fact, this figure is twice as high as

the one offered by the rest of the schemes when the several simultaneous transmissions saturate the channel given that MAA uses the resources more efficiently. This difference is even higher for h-SNR, and especially for CD. As sketched in Fig. 9b, the utility per user is considerably affected due to the inadequate user distribution caused by the fact of not considering the airtime utilization and by the number of simultaneous transmissions. In fact, the Jain's fairness index practically achieves the highest value for all the cases in MAA. Finally, Fig. 9c presents the MCS distribution with respect to the association policy and the number of users. Due to the use of ACKs and retransmissions, the unicast solutions, i.e., h-SNR, DAW, AIR and CD, show high data rates until the retransmissions increase due to the poor performance, as it is the case of CD for 5 users. Conversely, MAA shows a more prudent MCS selection based on the 10% of the time in which it collects rate control statistics from the users to ensure that all of them receive the information properly. Nevertheless, the increase of the MCS with respect to the standard allows achieving a tradeoff between reliability and channel utilization, leading to an improved performance given by the more efficient user-AP association and load balancing. Notice that for all policies, users follow a CD approach if the controller is not active. However, due to the client/AP stickiness, unless the signal of the associated AP was too weak, we observed active handovers from the users in a negligible number of occasions.

IX. CONCLUSIONS

In this work, we have leveraged an SDN controller to mitigate the inefficiencies in enterprise WLANs, e.g., suboptimal user-AP association caused by client-driven AP association. More specifically, we have proposed several polynomial-time complexity user-AP association schemes that an SDN controller executes periodically to improve resource utilization based on information collected from the network such as SNRs, APs load, traffic requirements, and requested content by each user. We have improved network utility, e.g., in terms of proportional-fair throughput, also by defining link-layer multicast groups when users request the same content. We have validated our solutions via simulation and in a practical setting leveraging an SDN-based WLAN testbed, showing an improvement of up to 11x in median throughput over client-driven approach when multicast delivery mode is enabled and all users request the same content. While the improvement decreases with increasing content diversity, we still observe higher throughput, e.g., 68% when there are 100 contents. As a future work, we plan to leverage the mobility information in IEEE 802.11mc as part of the user-AP association decision-making process.

REFERENCES

[1] E. Khorov, A. Kiryanov, A. Lyakhov, and G. Bianchi, "A Tutorial on IEEE 802.11ax High Efficiency WLANs," *IEEE Commun. Surveys Tuts.*, vol. 21, no. 1, pp. 197–216, 2019.
 [2] R. Murty, J. Padhye, R. Chandra, A. Wolman, and B. Zill, "Designing High Performance Enterprise WiFi Networks," in *Proc. of NSDI*, San Francisco, CA, USA, 2008.

[3] A. Bhartia, B. Chen, D. Pallas, and W. Stone, "Clientmarshal: Regaining control from wireless clients for better experience," in *Proc. of ACM Mobicom*, Los Cabos, Mexico, 2019.
 [4] *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, ANSI/IEEE Standard 802.11, LAN/MAN Standards Committee of the IEEE Computer Society Std., 2016.
 [5] B. Dezfouli, V. Esmaelzadeh, J. Sheth, and M. Radi, "A review of software-defined WLANs: Architectures and central control mechanisms," *IEEE Commun. Surveys Tuts.*, vol. 21, no. 1, pp. 431–463, 2018.
 [6] T. Sun, Y. Zhang, and W. Trappe, "Improving Access Point Association Protocols Through Channel Utilization and Adaptive Probing," *IEEE Trans. Mobile Comput.*, vol. 15, no. 5, pp. 1157–1167, 2016.
 [7] R. Riggio, M. K. Marina, J. Schulz-Zander, S. Kuklinski, and T. Rasheed, "Programming Abstractions for Software-Defined Wireless Networks," *IEEE Trans. Netw. Service Manag.*, vol. 12, no. 2, pp. 146–162, 2015.
 [8] A. Zubow, S. Zehl, and A. Wolisz, "BigAP-Seamless Handover in High Performance Enterprise IEEE 802.11 Networks," in *Proc. of IEEE NOMS*, Istanbul, Turkey, 2016.
 [9] D. Xenakis, N. Passas, L. Di Gregorio, and C. Verikoukis, "A Context-Aware Vertical Handover Framework Towards Energy-Efficiency," in *Proc. of IEEE VTC*, Yokohama, Japan, 2011.
 [10] K. Nakauchi and Y. Shoji, "WiFi Network Virtualization to Control the Connectivity of a Target Service," *IEEE Trans. Netw. Service Manag.*, vol. 12, no. 2, pp. 308–319, 2015.
 [11] E. Coronado, R. Riggio, J. Villalón, and A. Garrido, "Lasagna: Programming Abstractions for End-to-End Slicing in Software-Defined WLANs," in *Proc. of IEEE WoWMoM*, Chania, Greece, 2018.
 [12] 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, CA, USA, 2016.
 [13] E. Coronado, R. Riggio, J. Villalón, and A. Garrido, "Joint Mobility Management and Multicast Rate Adaptation in Software-Defined Enterprise WLANs," *IEEE Trans. Netw. Service Manag.*, vol. 15, no. 2, pp. 625–637, June 2018.
 [14] S. Bayhan and A. Zubow, "Optimal Mapping of Stations to Access Points in Enterprise Wireless Local Area Networks," in *Proc. of ACM MSWiM*, Miami, Florida, USA, 2017.
 [15] M. Amer, A. Busson, and I. Guérin Lassous, "Association Optimization in WiFi Networks: Use of an Access-based Fairness," in *Proc. of ACM MSWiM*, Malta, 2016.
 [16] W. Li, S. Wang, Y. Cui, X. Cheng, R. Xin, M. A. Al-Rodhaan, and A. Al-Dhelaan, "AP association for proportional fairness in multirate WLANs," *IEEE/ACM Trans. Netw.*, vol. 22, no. 1, pp. 191–202, 2014.
 [17] Y. Bejerano and S.-J. Han, "Cell breathing techniques for load balancing in WLANs," *IEEE Trans. Mobile Comput.*, vol. 8, no. 6, pp. 735–49, 2009.
 [18] E. Coronado, R. Riggio, J. Villalón, and A. Garrido, "Wi-balance: Channel-aware user association in software-defined Wi-Fi networks," in *Proc. of IEEE/IFIP NOMS*, Taipei, Taiwan, 2018.
 [19] E. Zeljković, J. M. Marquez-Barja, A. Kessler, R. Riggio, and S. Latré, "Proactive Access Point Driven Handovers in IEEE 802.11 Networks," in *Proc. of IEEE CNSM*, Rome, Italy, 2018.
 [20] P. Dely, A. Kessler, N. Bayer, H. Einsiedler, and C. Peylo, "Optimization of WLAN associations considering handover costs," *EURASIP J. Wirel. Comm.*, vol. 2012, no. 1, p. 255, 2012.
 [21] D. Gong and Y. Yang, "On-line AP association algorithms for 802.11n WLANs with heterogeneous clients," *IEEE Trans. Comput.*, vol. 63, no. 11, pp. 2772–86, 2014.
 [22] F. Cao, Z. Zhong, Z. Fan, M. Sooriyabandara, S. Armour, and A. Ganesh, "User association for load balancing with uneven user distribution in IEEE 802.11ax networks," in *Proc. of IEEE CCNC*, 2016.
 [23] F. R. d. Sa, A. M. d. Cunha, and C. d. A. C. Cesar, "Effective AP Association in SDWN Based on Signal Strength and Occupancy Rate," in *Proc. of IEEE LCN*, Singapore, Singapore, 2017.
 [24] E. Coronado, R. Riggio, J. Villalón, and A. Garrido, "Efficient Real-Time Content Distribution for Multiple Multicast Groups in SDN-Based WLANs," *IEEE Trans. Netw. Service Manag.*, vol. 15, no. 1, pp. 430–443, 2018.
 [25] D. Gong and Y. Yang, "Link-layer multicast in large-scale 802.11n WLANs with smart antennas," *IEEE Trans. Comput.*, vol. 65, no. 7, pp. 2118–2133, 2016.
 [26] J. Song and K. Chin, "A Novel Association Control Strategy for Supporting Multiple Multicast Sessions in WLANs," *IEEE Commun. Lett.*, vol. 16, no. 12, pp. 1933–1936, 2012.

[27] E. Zeljković, N. Slamnik-Kriještorac, S. Latré, and J. M. Marquez-Barja, "ABRAHAM: Machine Learning Backed Proactive Handover Algorithm Using SDN," *IEEE Trans. Netw. Service Manag.*, vol. 16, no. 4, pp. 1522–1536, 2019.

[28] T. M. De Oliveira, M. W. Da Silva, K. V. Cardoso, and J. F. de Rezende, "Virtualization for load balancing on IEEE 802.11 networks," in *Proc. of Springer MobiQuitous*, Sydney, Australia, 2010.

[29] O. Ercetin, "Association games in IEEE 802.11 WLANs," *IEEE Trans. Wireless Comm.*, vol. 7, no. 12, pp. 5136–43, 2008.

[30] S. K. Dandapat, B. Mitra, R. R. Choudhury, and N. Ganguly, "Smart association control in wireless mobile environment using max-flow," *IEEE Trans. Netw. Service Manag.*, vol. 9, no. 1, pp. 73–86, 2012.

[31] P. B. Oni and S. D. Blostein, "Decentralized AP selection in large-scale wireless LANs considering multi-AP interference," in *Proc. of IEEE ICNC*, Silicon Valley, USA, 2017.

[32] A. Baid, M. Schapira, I. Seskar, J. Rexford, and D. Raychaudhuri, "Network cooperation for client-AP association optimization," in *Proc. of IEEE WiOpt*, Paderborn, Germany, 2012.

[33] L.-H. Yen, T.-T. Yeh, and K.-H. Chi, "Load balancing in IEEE 802.11 networks," *IEEE Internet Comput.*, vol. 13, no. 1, pp. 56–64, 2009.

[34] V. Shrivastava, N. Ahmed, S. Rayanchu, S. Banerjee, S. Keshav, K. Papagiannaki, and A. Mishra, "CENTAUR: realizing the full potential of centralized w lans through a hybrid data path," in *Proc. of ACM MobiCom*, Beijing, China, 2009.

[35] H. Velayos, V. Aleo, and G. Karlsson, "Load balancing in overlapping wireless WLAN cells," in *Proc. of IEEE ICC*, Paris, France, 2004.

[36] X. Fafoutis and V. A. Siris, "Handover incentives for WLANs with overlapping coverage," in *Proc. of Springer WWIC*, 2009.

[37] S. Keranidis, T. Korakis, I. Koutsopoulos, and L. Tassiulas, "Contention and traffic load-aware association in IEEE 802.11 WLANs: algorithms and implementation," in *Proc. of IEEE WiOpt*, Princeton, USA, 2011.

[38] G. Xue, Q. He, H. Zhu, T. He, and Y. Liu, "Sociality-aware access point selection in enterprise wireless LANs," *IEEE Trans. Parallel Distrib. Syst.*, vol. 24, no. 10, pp. 2069–78, 2013.

[39] J. Vella and S. Zammit, "A survey of multicasting over wireless access networks," *IEEE Comm. Surveys Tuts.*, vol. 15, no. 2, pp. 718–53, 2013.

[40] 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*, Göttingen, Germany, 2013.

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

[42] M. Choi, W. Sun, J. Koo, S. Choi, and K. G. Shin, "Reliable video multicast over Wi-Fi networks with coordinated multiple APs," in *Proc. of IEEE INFOCOM*, Toronto, Canada, 2014.

[43] H. S. Oh, J. G. Shin, W. S. Jeon, and D. G. Jeong, "Reliable video multicast based on AL-FEC and H.264/AVC video traffic characteristics," in *Proc. of IEEE CCNC*, Las Vegas, NV, USA, 2017.

[44] Y. Daldoul, D.-E. Meddour, T. Ahmed, and R. Boutaba, "Performance and Scalability Evaluation of IEEE 802.11v/aa Multicast Transport," *Wirel. Commun. Mob. Com.*, vol. 16, no. 14, pp. 1987–2000, 2016.

[45] V. Gupta, C. Gutterman, Y. Bejerano, and G. Zussman, "Experimental evaluation of large scale WiFi multicast rate control," *IEEE Trans. Wireless Commun.*, vol. 17, no. 4, pp. 2319–2332, 2018.

[46] Cisco, "Channel Planning Best Practices," https://documentation.meraki.com/MR/WiFi_Basics_and_Best_Practices/Channel_Planning_Best_Practices, 2017, Accessed on 23.03.2020.

[47] S. Chiochan, E. Hossain, and J. Diamond, "Channel assignment schemes for infrastructure-based 802.11 WLANs: A survey," *IEEE Commun. Surveys Tuts.*, vol. 12, no. 1, pp. 124–136, 2010.

[48] 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.

[49] S. Huaizhou, R. V. Prasad, E. Onur, and I. Niemegeers, "Fairness in wireless networks: Issues, measures and challenges," *IEEE Commun. Surveys Tuts.*, vol. 16, no. 1, pp. 5–24, 2013.

[50] R. ITUR, "Guidelines for evaluation of radio interface technologies for IMT-Advanced," 2009.

[51] J. M. Keenan and A. J. Motley, "Radio coverage in buildings," *British Telecom Technical Journal*, vol. 8, no. 1, pp. 19–24, 1990.

[52] L. Wireless, "Ath9k spectral scan," https://wireless.wiki.kernel.org/en/users/drivers/ath9k/spectral_scan, 2014, Accessed on 23.03.2020.

[53] S. Wunderlich, "FFT evaluation," https://github.com/simonwunderlich/FFT_eval, 2014, Accessed on 23.03.2020.

[54] M. Rademacher and M. Kessel, "An Empirical Evaluation of the Received Signal Strength Indicator for fixed outdoor 802.11 links," *VDE ITG-Fachbericht Mobilkommunikation*, vol. 20, pp. 62–66, 2015.

BIOGRAPHIES



wireless networks, radio resource management, and edge computing.

Suzan Bayhan is an Assistant Professor at the University of Twente, the Netherlands and an Adjunct Professor (Docent) at the University of Helsinki, Finland. Suzan earned her Ph.D. in Computer Engineering in 2012 from Bogazici University, Turkey and worked at the University of Helsinki, Aalto University, and TU Berlin between 2012-2019. She was a short-term visitor at Princeton University and Aalto University. She received with her co-authors the best paper award at ACM ICN 2015. Her current research interests include spectrum sharing, coexistence of



mobile communications, network slicing, SDN, NFV, AI-driven networks and automated network management.

Estefanía Coronado is an Expert Researcher in the Smart Networks and Services Unit at FBK. In 2018 she completed her PhD at the University of Castilla-La Mancha (Spain) on multimedia distribution over SD-WLANs using ML techniques. She received the M.Sc. degrees in Computer Engineering and Advanced Computer Technologies in 2014 and 2015 from the same university. She published around 20 papers in international journals and conferences and received the IEEE INFOCOM Best Demo Award 2019. Her research interests include wireless and



FOCOM Best Demo Award (2013 and 2019) and the IEEE CNSM Best Paper Award (2015). He serves on the TPC/OC of top conferences and is an associate editor for the Wiley International Journal of Network Management, the Springer Wireless Networks Journal, and the IEEE Transactions on Network and Service Management. He is ACM member and IEEE Senior member.

Roberto Riggio is Head of the Smart Networks and Services Unit at FBK. His research interests revolve around networked and distributed systems with a particular focus on 5G and beyond networks. His fields of applications are distributed and decentralized edge computing, intelligent networks, and human-driven networking. He has published more than 125 papers in internationally refereed journals and conferences and has generated more than 3.5 MEuro in competitive funding in the last 5 years.



tence of heterogeneous wireless technologies in unlicensed spectrum, high-performance IEEE 802.11 networks, software-defined wireless networking, and ultra-reliable low latency communication.

Anatolij Zubow received the M.Sc. degree in computer science and the Ph.D. degree from the Humboldt University of Berlin in 2004 and 2009, respectively. He has been an Interim (Gast-) Professor of electrical engineering and computer science with the Chair for Telecommunication Networks, Technische Universität Berlin, since October 2018. His research interests are in architectures and protocols of wireless communication networks as well as in protocol engineering with impact on performance and QoS aspects. He is currently focusing mainly on coexistence of heterogeneous wireless technologies in unlicensed spectrum, high-performance IEEE 802.11 networks, software-defined wireless networking, and ultra-reliable low latency communication.