



Adaptive ML-Based Frame Length Optimisation in Enterprise SD-WLANs

Estefanía Coronado¹ · Abin Thomas¹ · Roberto Riggio¹

Received: 10 January 2020 / Revised: 2 March 2020 / Accepted: 10 March 2020 /
Published online: 17 March 2020
© The Author(s) 2020

Abstract

Software-Defined Networking (SDN) is gaining a lot of traction in wireless systems with several practical implementations and numerous proposals being made. Despite instigating a shift from monolithic network architectures towards more modulated operations, automated network management requires the ability to extract, utilise and improve knowledge over time. Beyond simply scrutinizing data, Machine Learning (ML) is evolving from a simple tool applied in networking to an active component in what is known as Knowledge-Defined Networking (KDN). This work discusses the inclusion of ML techniques in the specific case of Software-Defined Wireless Local Area Networks (SD-WLANs), paying particular attention to the frame length optimization problem. With this in mind, we propose an adaptive ML-based approach for frame size selection on a per-user basis by taking into account both specific channel conditions and global performance indicators. By relying on standard frame aggregation mechanisms, the model can be seamlessly embedded into any Enterprise SD-WLAN by obtaining the data needed from the control plane, and then returning the output back to this in order to efficiently adapt the frame size to the needs of each user. Our approach has been gauged by analysing a multitude of scenarios, with the results showing an average improvement of 18.36% in goodput over standard aggregation mechanisms.

Keywords Machine learning · Supervised learning · WLANs · IEEE 802.11 · Frame length optimization · Aggregation · SDN · Network management

✉ Estefanía Coronado
e.coronado@fbk.eu

Abin Thomas
athomas@fbk.eu

Roberto Riggio
rriggio@fbk.eu

¹ Smart Networks and Services (SENSE), Fondazione Bruno Kessler, Trento, Italy

1 Introduction

Progress in the communications industry has generally been marked by hardware and computation-centric innovations. However, networking systems have gradually begun to evolve dynamically towards service-oriented architectures that break the chains of an outmoded dependence on monolithic network stacks and conventional hardware advancements. This change has been the objective pursued by Software-Defined Networking (SDN), which introduced a management architecture characterized by the decoupling of the control and data planes across various degrees of centralisation [1, 2], thus demonstrating that traditional approaches to network management are no longer adequate. This is especially true when it comes to wireless networks, where the level of complexity, which is due to their error-prone nature, longer delays, and the inefficient and inflexible use of resources, requires the use of advanced network management policies.

SDN marks a turning point in networking, giving rise to a new generation of programmable and service-oriented networks. However, the inherent behaviour of the Physical (PHY) layer is still defined by mathematical models and complex algorithms, whose accuracy and tractability can be compromised by the increase in data processing, and in the number and heterogeneity of the services provided. As a result, it is becoming increasingly evident that there is a need to enhance the flexibility of not only the resource provisioning, but also of the physical infrastructure. This involves a change from *reactive* to *proactive and automated* network management, in which the analytics made available at the SDN controller lay the basis for *knowledge-based* and *self-driven* networks in what is known as Knowledge-Defined Networking (KDN) [3]. This paradigm aims at building smarter networks able to autonomously optimize operation and management by extending the SDN architecture with a Knowledge Plane (KP), a new component characterized by the active inclusion of Machine Learning (ML) techniques.

ML has been effectively applied to several domains, including computer vision and medical diagnosis, as an accurate approach to manage large and heterogeneous volumes of data. Its contribution to many fields has proved its ability to overcome the drawbacks of traditional mathematical formulations and complex data analysis algorithms [4]. This is especially relevant in networking, where the growing diversification of services and the constantly changing channel dynamics require more effective ways to speed up the decision-making process, a task for which ML constitutes a very promising solution [5–7]. Far from being realistic and technically feasible, ML solutions for networking have often been a controversial matter. One of the major reasons for this is the inherently distributed control, where nodes have only a partial view of the system. A good example can be found in IEEE 802.11 wireless networks [8], in which Access Points (APs) handle traffic flows and radio resources independently of each other without considering specific requirements and channel fluctuations, thus leading to performance inefficiencies [9]. Moreover, dataset collection for training and validation represents a complex task given the diversity of parameters involved. Network softwarization

and programmability can serve as a uniting bond between networking and ML. SDN provides the perfect ecosystem for ML by bringing together multiple performance metrics at the controller as well as current and historical network information [10].

This trend has given rise to ambitious SDN-based approaches for Wireless Local Area Networks (WLANs) as a way not only to simplify network management but also to alleviate the complexity of algorithms aiming to optimize different aspects of the network [11–13]. As a result, a huge body of literature can be found on improvements on top of the Medium Access Control (MAC) layer, including the Modulation and Coding Scheme (MCS) [14, 15] and resource-based slicing [16, 17], among others. However, such improvements are limited by channel access and encapsulation [18], which makes frame size adaptation the focus of research, in an effort to reduce this transmission overhead. This is the idea behind the 802.11 standard, which introduces two methods of fixed frame aggregation, namely Aggregated MAC Service Data Unit (A-MSDU) and Aggregated MAC Protocol Data Unit (A-MPDU) [19]. Although it may seem logical that longer frames require fewer channel accesses (thus lowering the overhead), transmission errors affect a greater amount of information, which has to be retransmitted later. Consequently, this tradeoff between network status and frame length has attracted considerable interest [20–23]. Nevertheless, the optimal size may vary over time and is dependent on a wide range of factors such as bitrate, channel quality, mobility and MCS, which may also be different for each user [24].¹

Considering the above, the dynamic frame length selection problem requires more advanced techniques able to handle such amounts of information on practical settings in real-time. To deal with these limitations, in this paper we leverage ML techniques, and in particular Supervised Learning (SL), to introduce a low-complexity ML model for adaptive frame length optimization in enterprise SD-WLANs. This solution is transparently deployed on the management plane of the SDN architecture, and is periodically fed with the network *knowledge*, which is made available at the control plane, about channel conditions and user state. The main contribution of this work is the ability to provide a per-user specific frame size selection while attending to both global and individual performance indicators, and without making any assumptions about the channel model. This fact is of great importance since a common frame length may not be optimal for all the users, especially with changing and heterogeneous channel conditions, such as in mobility scenarios. An extensive performance assessment via simulation has shown that our approach can improve network goodput and user fairness with respect to the standard A-MSDU aggregation with constant frame length by up to 45%.

The rest of the paper is structured as follows. Section 2 describes the frame aggregation mechanisms introduced in the IEEE 802.11 standard, and discusses related work. The principles behind the main ML and SL techniques applied in networking are outlined in Sect. 3. A detailed description of the ML model for frame length optimization is presented in Sect. 4. Section 5 reports on the performance evaluation

¹ Note that the terms stations, clients and users are used interchangeably in this work.




Fig. 1 Schematic diagram of A-MSDU and A-MPDU aggregation formats

under various channel conditions. Finally, Sect. 6 draws the conclusions and outlines areas for future work.

2 Related Work

In this section, we first provide a background covering the standard frame aggregation mechanisms in 802.11 networks. Then, we discuss the main works aiming at overcoming the issues in the standard concerning frame length optimization.

2.1 Aggregation Mechanisms in IEEE 802.11

The IEEE 802.11 standard defines three levels of frame aggregation, namely A-MSDU, A-MPDU and a combination of the two, as shown in Fig. 1 [8]. The main goal of these mechanisms is to reduce the overhead generated by headers, preambles and channel access [25, 26].

A-MSDU aggregation improves efficiency by collating multiple MSDUs with a single PHY and MAC header. The MSDU subframe header must have as Destination Address (DA) and Sender Address (SA) the same as Receiver Address (RA) and Transmitter Address (TA) in the MAC header. The aggregated frame, i.e., the frame comprising multiple MSDUs, is complete when the length equals the maximum aggregation size. The standard limits the maximum aggregation size to 3839 or 7935 bytes, but the selection of any of these values is vendor-dependent. Nevertheless, A-MSDUs suffer in error-prone channel conditions as a result of generating a single Frame Check Sequence (FCS) for the whole packet. By contrast, A-MPDU aggregation involves a single PHY header and combines multiple A-MSDUs or MPDUs with their respective MAC headers. A-MPDU, unlike A-MSDU, generates the FCS for each subframe, which helps to reduce retransmissions in case of error and performs better in error-prone channel conditions. The A-MPDU maximum aggregation length has an upper bound of 65536 bytes.

2.2 Solutions for Frame Length Optimization

Frame length optimization has long been studied within IEEE 802.11 networks by relying on several frame aggregation mechanisms. The main literature on this topic

is discussed below in relation to the specific procedures employed for dynamically adjusting the frame size.

A-MSDU aggregation related solutions, which are the technique leveraged in this work, are principally focused on problems related to Quality of Service (QoS), real-time traffic and small frames. An example of A-MSDU aggregation is proposed by Maqhat et al. in [27], in which a scheduler for delay-sensitive traffic is presented. Control bits are modified for each sub-frame to facilitate swifter retransmissions. This work was later implemented using NS-2 in [28]. Similarly, the authors of [29] seek improvements for error-prone channels by adding control bits to every sub-frame with the aim of enabling per-subframe retransmissions. Conversely, Kim et al. [22] propose a method that deals with frame size estimation based on the Extended Kalman Filter for saturated networks. Attending to QoS requirements, in [30] a dynamic scheme is introduced to calculate the optimal size according to the packet traffic class that is based on a tradeoff between throughput and delay to form an aggregated frame. Similarly to [30], Saldana et al. [31] examine the tradeoff between throughput and latency of A-MSDU aggregation when considering mobile users. Moreover, Kriara et al. [9] study the impact of PHY rate and frame aggregation on the performance of an 802.11 network with respect to other factors for determining the optimal frame length. Finally, the authors in [21] propose a joint PHY-MAC link adaptation strategy with theoretical link quality analysis together with A-MSDU aggregation in error-prone channel conditions.

Several works also look to incorporate optimizations to the A-MPDU mechanism. The policy proposed in [23] is aimed at dynamically adapting the A-MPDU length by discerning the effects of user mobility on channel conditions. Furthermore, in [32] the authors use a Proportional Integral Derivative (PID) controller to select an appropriate A-MPDU aggregation size based on QoS indicators. Unlike [32], the objective in [33] is to find the optimal number of MPDUs with respect to the delay requirements for 802.11ac-based WLANs. In this work, the authors look for improved performance while handling delay requirements utilizing RTS/CTS via simulation. Seytnazarov et al. [34] introduce a QoS-aware A-MPDU aggregation scheduler for voice traffic. This approach, however, is not standard-compliant. In line with this, in [35] the authors deal with rate and frame size adaptation by modeling the network conditions in NS-2 using different Bit Error Rate (BER) values. Similarly, in [36] the authors perform rate adaptation, A-MPDU aggregation and Multiple-Input Multiple-Output (MIMO) mode selection based on Channel State Information (CSI).

Looking at a combined A-MSDU/A-MPDU aggregation strategy, the authors in [37] present an exhaustive performance evaluation of the A-MSDU and A-MPDU mechanisms with respect to different data rates and transport protocols. Based on these findings, Kim et al. [38] look to combine A-MSDU and A-MPDU aggregation to achieve airtime fairness and improve overall network throughput. The authors in [39] study the performance of A-MSDU and A-MPDU mechanisms in NS-2 under error-prone channel conditions. Then, the authors propose a frame length selection scheme leveraging the results obtained from the analysis of the relationship between frame length and BER. Similar approaches are also implemented in vehicular networks [40, 41]. Unlike the rest, Li et al. [18] propose a scheme with fragment

retransmission where multiple packets are aggregated and transmitted as a single frame. Instead of using MSDU or MPDU aggregation, this work proposes an algorithm in which a fragmentation threshold is set in such a manner that any packet longer than this value is fragmented before the aggregation process begins. This model is evaluated using NS-2 for TCP, HDTV and VoIP.

Despite the amount of literature on frame length optimization, to the best of our knowledge, there is little research exploring the use of ML techniques to tackle this problem. A representative example can be found in [18], in which Lin et al. introduce an ML approach to optimize the frame size selection when considering error-prone channel conditions. The algorithm is tested by extending Bianchi's model [42] with parameters such as retry limits and different data rates. Nevertheless, this work does not deal with frame aggregation and deaggregation but instead limits itself to fragmentation/defragmentation in order to maintain backward compatibility with 802.11a/b/g. Moreover, none of the existing solutions provides a specific frame length for each user. The use of a common frame length makes it difficult to adapt the transmission to all the users' conditions, which is especially relevant in the case of heterogeneous channel states. In contrast to this approach, our work proposes an SL-based solution that is able to dynamically calculate a per-user frame length based on diverse network conditions and to use such a length to compose specific aggregated A-MSDUs for each user at any time.

3 Machine Learning for Networking

Machine learning has been central to the creation of autonomous systems. This term can be defined as the algorithmic techniques with the ability to process data and learn from it, instead of those algorithms that are explicitly designed and written in a fixed way to perform a specific task. The application of ML in networking aims at projecting a new vision in which human interaction is reduced to creating *self-driven* networks that are able to configure and optimize themselves. In this section we first review the most relevant ML techniques for communication systems, paying particular attention to the one used in this work, i.e., supervised ML methods. Then, we discuss the initial attempts at deploying intelligent network management policies, focusing on ML solutions expressly used in the frame length selection problem in WLANs.

3.1 Machine Learning Principles and Techniques

The ML taxonomy can be broadly classified into three main categories according to principles and applications, namely Supervised Learning, Unsupervised Learning (UL) and Reinforcement Learning (RL), as shown in Fig. 2. As complements to these main categories, (Deep) Federated Learning (DFL) and Transfer Learning (TL) can also be distinguished [43].

According to the information available in the dataset, we can refer to supervised or unsupervised learning. In supervised learning, the training process is

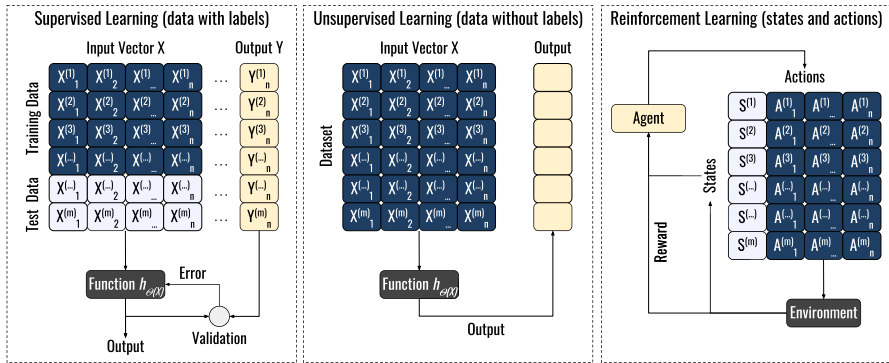


Fig. 2 Main techniques in the machine learning taxonomy

carried out on a set of examples establishing a map between an specific input and the corresponding output. The resulting algorithms are task-oriented and can be categorized into regression and classification techniques. Examples of this type of learning are linear/logistic regression, neural networks, and decision trees, among others. On the other hand, in unsupervised learning, the training data contains only input features and is used to extrapolate a statistical structure. Dimensionality reduction and clustering are the main groups of methods in this category. Some examples of these techniques are K-means and principal component analysis [44]. By contrast, reinforcement learning cannot be exactly defined according to the dataset structure since, in a sense, it lies somewhere between supervised and unsupervised learning. These algorithms enable a constant evolution of the dataset, i.e., the form of supervision comes from the feedback from the environment in which the model is executed after selecting an output for a given input feature. This feedback is then used to adjust the next action. The main learning algorithms in this category are Q-Learning and Multi-Armed Banding Learning (MAB), among others. Conversely, federated learning (especially suitable for mobile networks) aims to distribute data and computation tasks among federated devices that are coordinated by a central server. The server is in charge of combining the local models into a common neural network, which is based on, and updated according to, the local datasets [45]. Finally, the goal of (deep) transfer learning is to apply the knowledge from one task to another in a related context to reduce the amount of data required for training and validating new models [46].

Most of the current ML applications in networking fall into the supervised learning category due to the observable capabilities and the form of explicit feedback [5, 47]. For this reason, this paper focuses on supervised learning, leaving the investigation of unsupervised and reinforcement learning in the specific problem of frame length optimization in WLANs as future work.

3.2 A Focus on Supervised Learning

3.2.1 Main Characteristics

Supervised learning refers to the process of building a model, $h_{\theta}(x)$, from supervised data, which is characterized by n input features, $X = (X_1, \dots, X_n)$, and an output variable, Y . Therefore, data must be represented as a pair, (x, y) , in a dataset, $S := (x^{(1)}, y^{(1)}), \dots, (x^{(m)}, y^{(m)})$, which contains the information regarding m independent entries whose current outputs are already known. Based on this, the models must predict the output of other unlabeled data, y , from its input features, x . Depending on the output class, two types of supervised learning can be distinguished: classification and regression. Classification problems are those whose goal is to determine the class of a certain instance, i.e., suitable for binary/categorical classes, in such a manner that $Y \in c_1, \dots, c_w$. In regression, however, the target is a numerical class, and therefore, $Y \in \mathbb{R}$ and $h_{\theta}(x) \in \mathbb{R}$.

Many SL techniques can be found in the literature, such as linear regression [48], Support Vector Machines (SVM) [49] and decision trees [50]. The selection depends mainly on the requirements of the problem. Some of these techniques are much more powerful than others, i.e., they achieve greater precision and can identify a wider set of patterns in data. However, the ease of interpretation is also an issue in some scenarios. For example, models such as neural networks are considered a *black box* since the underlying patterns cannot be extracted. By contrast, decision trees can be easily interpreted and allow the identification of relationships between input and output features. Another important issue concerns computational complexity. For instance, processing a decision tree might only require a few comparisons, which becomes particularly important in real-time contexts.

Therefore, decision trees are a popular SL method that is commonly used for data exploration, classification and regression problems. This technique is a greedy, top-down binary and recursive partitioning algorithm that divides the feature vector, X , into sets of disjoint regions that are as pure as possible until a leaf node is reached. For each split, an input feature is selected according to the quality of the information provided with respect to the output class, Y , leveraging criteria such as Gini impurity, information gain and standard deviation reduction. Despite the numerous benefits, decision trees are prone to overfitting, especially as the depth increases. As a result, a non-representative model that generalizes poorly is obtained. Examples of ways to alleviate this problem are using postpruning techniques and/or stopping the growth of the tree when splits are not statistically significant. Nevertheless, it is still likely to produce noisy and inaccurate models.

3.2.2 Random Forest Models

Random Forest (RF), which has emerged as one of the most versatile ML approaches for classification and regression problems, is based on ensemble learning techniques to overcome the weaknesses of prediction trees [51]. The general principle of ensemble learning methods is to construct a linear prediction combination given by multiple ML algorithms to jointly provide more accurate

Table 1 Summary of ensemble learning techniques

	Bagging	Boosting	Stacking
Objective	Variance minimization	Bias minimization	Accuracy maximization
Data partitioning	Bootstrap sampling	Weighted sampling	Original dataset
Training stage	Parallel	Sequential	Parallel (heterogeneous models)
Combination function	Majority vote (class.), simple average (regr.)	Weighted majority vote (class.), weighted sum (regr.)	Meta model construction from first-level model output

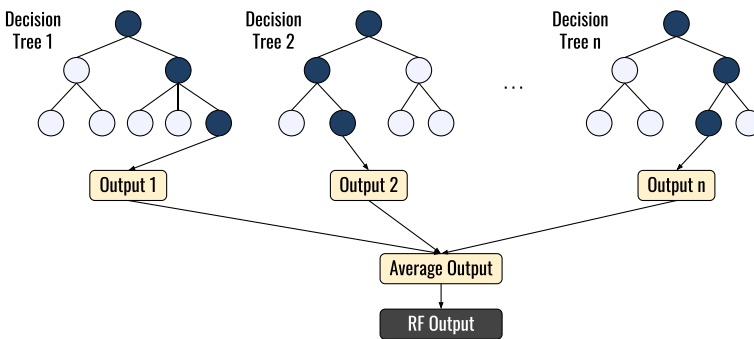


Fig. 3 Layout of a Random Forest structure

models. As shown in Table 1, three main ensemble learning techniques can be distinguished according to the objective and the learning process, namely bagging, boosting and stacking. In particular, the statistical principles of random forests are based on *bagging* techniques, which combine multiple decision trees that are constructed in parallel to average the noise approximated by each of them, with the goal of reducing variance and improving accuracy. A high level layout of the structure of an RF model is depicted in Fig. 3.

Mathematically speaking, RF can be described as an ensemble of unpruned tree-structured models, $h(X, \Theta_k)$. Instead of examining all possible feature-splits, X , each tree is built from a subset of features, which is represented by a random vector, Θ_k , which is independent from the previous vectors $\Theta_1, \dots, \Theta_{k-1}$, but with the same distribution. Moreover, each tree selects a sample from the dataset to carry out the training. The nature and dimensionality of Θ depends on its use in tree construction. In other words, the most computationally expensive aspect of the tree building is the feature split decision. Therefore, a narrower set of features reduces the learning time of each submodel. The goal of these models is to find a function, $h_\theta(X)$, that is able to accurately predict the output feature, Y . The prediction function is given by a loss function, $L(Y, h_\theta(X))$, which determines the difference between the prediction and the real value, Y . In classification problems, the loss function is defined by a zero-one loss equation (Eq. 1), while choices for regression problems consider the squared error loss (Eq. 2).

$$L(Y, h(X, \Theta_k)) = \begin{cases} 0 & \text{if } Y = h(X, \Theta_k) \\ 1 & \text{otherwise} \end{cases} \quad (1)$$

$$L(Y, h_\Theta(X)) = (Y - h(X, \Theta_k))^2 \quad (2)$$

The resulting submodels are combined by taking the majority vote in the case of classification, as shown in Eq. 3, or by averaging the output in the case of regression, as shown in Eq. 4. In such equations $h_\Theta(X)$ represents a combination of models, while $h(X, \Theta_k)$ refers to a single decision model, Y is the output variable and $I(\cdot)$ is the indicator function [52].

$$h_\Theta(X) = \arg \max_Y \sum_{i=1}^k I(h(X, \Theta_i) = Y) \quad (3)$$

$$h_\Theta(X) = \frac{1}{k} \sum_{i=1}^k h(X, \Theta_i) \quad (4)$$

3.3 From Software-Defined Networking to Knowledge-Defined Networking

SDN is expected to deliver a much simpler network whose behaviour can be easily modified and adapted. However, this simplicity can be compromised by the increase in the number of services, as well as in the variety of bitrate and latency requirements. In view of this complexity component, high-level policies and network software abstractions are not sufficient to efficiently manage wireless networks while accommodating rigorous service requirements. Future research efforts in network softwarization should enforce the deployment of management policies that are intelligent enough to automatically adapt to changes in the constantly varying wireless scenario. This aim of automatizing network management led to the introduction of Self-Organizing Networks (SONs) [53–55] by the Next Generation Mobile Networks (NGMN) alliance. Although these works take into account network conditions and QoS requirements, SONs failed to achieve dynamic data acquisition (which may be spread across various network nodes) and the corresponding processing in order to allow all segments to rapidly react to varying traffic or even infrastructure changes due to connectivity or hardware failures.

This change of perspective leads us to believe that future wireless networks, including WLANs, must follow an ML-native approach and become smart, agile, and able to *learn* from, and *adapt* to, the changing environment. If this transition from network *softwarization* to network *brainitisation* is to take place, ML cannot be treated as an afterthought but as a central element that must be taken into account from the requirements phase, and work together with mathematical models and standards to make this vision a reality. This statement is set out in [3, 56], in which the SDN paradigm is extended with a knowledge plane in the so-called KDN paradigm. As shown in Fig. 4, the network architecture is composed of four

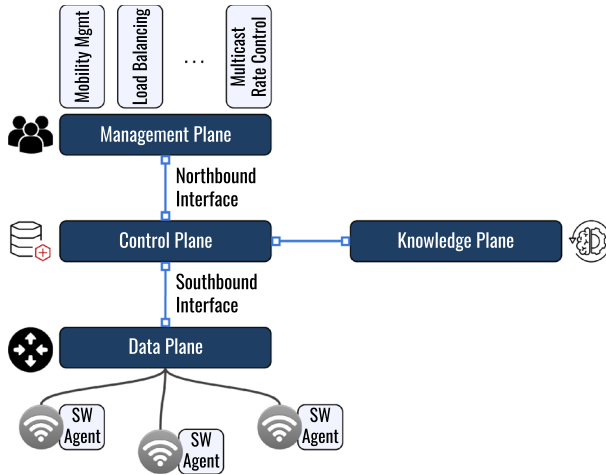


Fig. 4 Knowledge-defined networking structure

planes: in addition to the data, control and management planes, it incorporates the KP to provide the network with the ability to gain knowledge, generalize, learn from past experience and even *reason*. If decoupling data from the control plane offers a full network view in real-time and simplifies network management, the inclusion of ML techniques as the main enabler of the KP will allow, in time, the building of smarter networks that are able to optimize operation and management.

In the KDN paradigm, the data plane covers the infrastructure devices, which are relieved from control decisions and follow the operations issued by the control plane. Conversely, the control plane takes the raw data coming from the data plane and feeds it to the knowledge plane for further analysis. This interface is also used as a way for the control plane to receive the *intelligent* directives from the knowledge plane. Finally, the management plane takes advantage of the privileged position of the control plane to obtain information about the network state and about the analytics provided by the knowledge plane. The instructions provided by the knowledge plane can be specifically stored and enforced on a service base in the management plane in the form of network applications, e.g., mobility management or load balancing. Note how this structure is used in a control loop, i.e., the knowledge plane must be able to specify and adapt the required management policies based on historical and current network information despite channel or performance issues. This control loop requires the integration of data acquisition and analysis modules, which could, however, be implementation-dependent.

KDN can bring benefits to many networking problems. The following provides a review of representative topics in SDN-based wireless networks in which ML techniques have achieved an impressive performance [57], paying special attention to the focus of this work: frame length optimization in WLANs.

- *User Association and Handoff Management* User association is a problem of high computational complexity due to the number of features involved, therefore attracting the attention of ML applications. One of the most widely studied features is the Received Signal Strength Indicator (RSSI). In [58] this is used by a Recurrent Neural Network (RNN) to perform seamless user-AP association. Similarly, in [59] the authors rely on SL (SVMs and Naive Bayes) to estimate the number of active nodes in a Wi-Fi network. Examples of similar works are [60, 61].
- *Resource Management* The varying nature of wireless channels calls for efficient resource allocation schemes that consider the network as a whole instead of focusing on isolated performance indicators, e.g., RSSI and available bandwidth. This increase in complexity makes the problem worthy of being tackled using ML solutions. With this in mind, in [62] the authors propose a method for reducing the amount of interference and congestion based on Linear Programming (LP) and regression that is updated periodically upon changes in the channel. Similar ML tools are explored in [63] for active channel selection and channel extension (via channel hopping) when resources are exhausted. This topic is discussed in [64].
- *Localization* This topic has received increasing attention in ML-empowered networks given the device heterogeneity and user mobility, especially in noisy environments. In this respect, the authors in [65] use deep learning to achieve accurate fingerprinting localization across different APs. This solution is enhanced in [66], in which a deep learning model is fed with the CSI to obtain more fine-grained information on the wireless channel. Further research on this matter is carried out in [67].
- *Traffic Prediction and Analysis for QoS/QoE* Works falling into this category seek to ensure QoS and Quality of Experience (QoE) by optimizing network resources based on traffic requirements. The application of ML techniques has already enabled optimisations of this problem. For example, in [68] a neural network is trained for traffic prediction based on flow level statistics for network balancing. Conversely, the authors in [69] leverage reinforcement learning to control an SDN-based 802.11 network by taking the Mean Opinion Score (MOS) QoE metric as input for the model. Another work with the same objective is [70].
- *Frame Length Selection* The tradeoff between payload length and channel access cost presents a challenging problem in wireless channels (characterized by changing conditions) with regard to providing reliability and spectrum efficiency. As mentioned in the discussion in Sect. 2, several factors such as MCS, bitrate and interference can alter the channel state, a situation which calls for advanced ML tools. Similarly to our work, in [71] the authors introduce a solution based on Multilayer Perceptron (MLP) as the main ML component for frame size optimization by extending Bianchi's analytical model with parameters such as error-prone channel conditions, retry limits and different data rates. However, unlike our scheme, this work does not proceed with aggregation policies but limits its contribution to dynamic frame fragmentation and defragmentation. An analogous aim is pursued in [72], in which a DNN based on CSI data, FER and effective throughput is proposed for vehicle-to-vehicle communications based on

the IEEE 802.11 standard. To the best of our knowledge, this is the first time SL techniques have been used to tackle user-specific frame size adaptation in WLANs following standard frame aggregation mechanisms.

4 ML-Based Model for Optimal Frame Length Optimisation

In this section we first discuss the design decisions made with respect to the network configuration and present the system model considered in this work. Then, we describe the learning process followed to build the model and its interaction with the rest of the system.

4.1 Design Decisions

As discussed in Sect. 2, adapting the frame length to channel conditions has been widely discussed in the literature. Among the factors determining such channel conditions, many works have proved the greater role of the PHY rate with respect to others [20, 21, 73]. In other words, it has been demonstrated how stations with different channel conditions have a different frame length that maximizes the goodput at the receiver, and that could change according to the selected MCS (which is in turn partially determined by the distance). It is worth highlighting that in this work the selected frame length is achieved at the transmitter by leveraging A-MSDU aggregation.

This work targets the 802.11n version of the standard, given its wide spread use in the market. Nevertheless, our solution can be easily extended to work with other PHY layers. The 802.11n release defines four basic modulation schemes (BPSK, QPSK, 16-QAM, and 64-QAM), each of them with different coding schemes, which results in eight basic MCS values (from MCS 0 to MCS 7). Although up to four MIMO streams are supported, the majority of commercial APs support only two MIMO streams. MCS values higher than 7 are essentially the same as the lower ones but with an increasing number of streams, e.g., MCS 15 has two streams, MCS 23 has 3 streams and MCS 31 has four streams. For this work, we have focused on the effect of MCS on the optimal frame length, leaving the inclusion of MIMO features for a future work. Therefore, the training process has considered only MCS values from 0 to 7. Note also that the standard defines two maximum aggregation lengths for an A-MSDU: 3839 and 7935 bytes. Wi-Fi clients can support either of the two values. However, only 3839 bytes is used in our solution as the maximum length since it is the most common value in Wi-Fi clients for 802.11n interfaces.

4.2 System Model

The aggregation size specified by the IEEE 802.11 standard is static, regardless of traffic classes and network state. Despite reducing the channel access overhead, utilizing the maximum aggregation size may not be optimal in all situations since it may lead to an increase in delivery errors and retransmissions. Therefore,

to effectively serve multiple mobile stations with an increasing resource demands and different channel qualities, configuring the frame length at a finer granularity becomes essential. This is precisely the motivation for our work and the reference point for defining the system model.

The system model is based on a Wi-Fi network composed of D stations, M APs in the data plane and an SDN controller in the control plane. The D stations receive traffic in downlink from the AP, while U additional stations transmit traffic to the AP to simulate more realistic scenarios. The controller periodically collects state information from the APs, including network-wide statistics such as channel utilization, and per-user statistics such as data from the rate control algorithm, including successfully received data, packet loss, etc. This data is then fed to the KP plane, which enables the analysis of the information from the network and allows the proposed ML model to compute the frame length for each AP/client pair. After each computation period, such values are communicated to the SDN controller, which applies the new directives on the nodes in the data plane. To ensure compatibility with the IEEE 802.11 standard, the frame length optimization is only implemented in the downlink direction as it does not require changes to the client hardware when performing frame aggregation. By contrast, uplink aggregation length optimization would result in the solution not being compliant with the standard.

4.3 ML-Based Model Deployment

4.3.1 Model Selection

Due to the inherent variability that characterizes a wireless network and the need to execute the model in real-time, the ML technique to be applied in this problem must offer low computational complexity, and the ability to handle large datasets and adapt itself to changes over time. In this work we have leveraged SL techniques for two main reasons. Firstly, the algorithm must be able to predict a specific output, i.e., the expected goodput when selecting a specific MCS and frame length. Given the key role of MCS in the frame length, this choice has the goal of selecting the (*MCS, frame length*) combination giving the best performance. Secondly, in addition to accuracy, the interpretability of the model is an essential requirement since after being built, it must be possible to analyse and modify it to be easily adjusted to the network outputs.

With these requirements in mind, this decision has involved the analysis of several observable SL algorithms including not only decision trees, but also many others such as Naive Bayes. After this analysis, we selected a Random Forest Regressor (RFR) model given its low computational complexity, low variance, low overfitting level and self-explanatory capacity. The reason for focusing on regression models is that the output class of this problem is a numerical value, i.e., the expected goodput, Y , for a specific frame length given by the configuration X for a particular network and user state. The process for building and deploying the ML model is depicted in Fig. 5, and is described in detail in the following subsections.

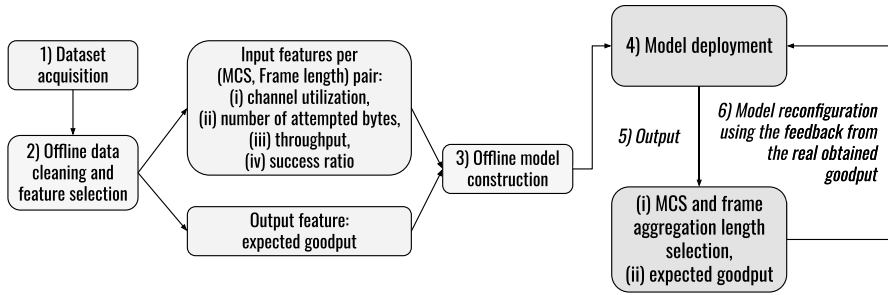


Fig. 5 Flowchart of the construction of the ML model for frame size optimisation

Table 2 Parameters used in the dataset acquisition process

Parameter	Value
MCS	[0–7]
Frame aggregation length	[512, 1024, 2048, 3839] bytes
Payload size	[200, 500, 1000, 1470] bytes
Aggregated bitrate	[5, 10, 15, 20, 25, 30, 35, 40] Mbps
Distance from AP	[15, 35, 50, 70] m
Number of stations	[1, 2, 4, 10]
Background traffic	2 uplink transmissions (1 Mbps)

4.3.2 Dataset Acquisition

The dataset acquisition process involves collecting the reference data containing the training and test sets for building the ML model. Complexity and privacy aspects are an important issue for collecting data from operational networks, and for this reason, we chose a synthetic dataset generation approach using an experimental WiFi testbed (*Step 1* in Fig. 5). The network layout deployed is similar to the one described in Sect. 4.2 but comprising a single AP. The AP was deployed using a PCEngines ALIX 2D (x86) processing board, equipped with a Wi-Fi card based on the Atheros AR9220 chipset and running OpenWRT 18.06.4. The AP was set to channel 36, isolated from other external noise. The SDN controller was built using the 5G-EmPOWER Software-Defined Radio Access Network (SD-RAN) controller [11]. The stations and the controller were deployed on Dell laptops powered by an Intel i7 CPU and running Ubuntu 18.04.02.

Using the aforementioned system we collected measurements from a wide set of test scenarios. The configuration used for each iteration was a unique combination of the parameters shown in Table 2, these being strictly related to the stations receiving downlink traffic from the AP. The background uplink traffic maintained the same configuration throughout the tests. The most significant factor is the frame aggregation length for which, in addition to the standard value, i.e., 3839 bytes, we

have proposed an extended range of values. The D stations were randomly placed in a static position within the coverage area, which has a radius of R , equal to the distance shown in Table 2. Conversely, 2 stations transmit uplink traffic to the AP, and these were set in a fixed position 10 m from the AP and used the basic MCS. We remind the reader that the uplink traffic is beyond the scope of this paper and is included in the setup to simulate more realistic environments.

The time series for each test case has a duration of 30 s and is repeated 10 times. For each combination, we collected the statistics of the rate control algorithm for each station. In this particular case, Minstrel [74] was implemented as rate control algorithm. The statistics provided have been extended to account not only in terms of packets but also in terms of bytes. Moreover, we have measured additional indicators such as goodput, throughput, success ratio, delivery probability and channel utilization, among others. As a result of this process, approximately 60k observations have been generated. Then, the dataset is processed offline using the *Sklearn* library, which has been deployed on an *a1.medium* instance on the Amazon EC2 platform. This process includes various subtasks, namely data cleaning, feature selection and model construction (*Step 2* and *Step 3* in Fig. 5).

4.3.3 Data Cleaning

First of all, we performed a data cleaning process. This tedious task requires great attention since the format of the dataset can considerably affect the accuracy of the predictions provided by the resulting model. One of the main issues is data imbalance, which usually reflects an unequal distribution of features. ML techniques such as Random Forests fail to cope with imbalanced training datasets since they are sensitive to the proportions of the different features. In this particular case, we observe this situation in the *aggregated length* variable since, for instance, given the payload sizes selected, the number of occurrences of the *512 bytes* value is lower than the one of *3839 bytes*. To address this problem, we carried out an *undersampling* process, in which some of the observations from the majority value were randomly deleted to match the number of the minority one. In addition, we verified that the data set does not contain missing values or duplicates.

4.3.4 Feature Selection

Despite the various features collected for each scenario, not all of them have a clear impact on prediction. As a matter of fact, the use of unrelated features can cause the opposite effect. Bearing this in mind, we carried out a variable selection process leveraging Random Forest techniques, which are able to rank input features in such a manner that the purity of the nodes is maximized [75]. As a result, the input features selected are: (i) channel utilization, (ii) number of attempted bytes in the last window of the rate control algorithm, (iii) throughput, and (iv) success ratio of the selected MCS. It should be noted that from this process, we generated a structure composed of K models, where K is the product of the number of MCS values and the number of frame aggregation lengths. Since 8 MCS and 4 frame aggregation

length values have been taken into consideration, this results in a total of 32 RFR models. Nevertheless, note how only 4 of them are evaluated per execution, since a single MCS is chosen each time per user by the rate selection algorithm. Then, each model provides the expected goodput on a per (*MCS, frame length*) pair basis.

4.3.5 Model Construction

By means of the aforementioned input features we have built an *RFR* model for each MCS, limiting the depth of the trees to 3 levels to reduce overfitting. The models undergo a tenfold cross validation to guarantee that the training and the test datasets are independent. The output of this process reports a mean absolute error of 9.80%, which shows the accuracy of the model and the relationship between the parameters involved. After these steps, the ML model is deployed on the management plane. It is run once per second for every station, producing as output the frame length for a specific MCS, and the predicted goodput for that combination of values (*Step 5* in Fig. 5). However, the model is not static: in the next run, the real goodput obtained can be compared with the predicted one, thus correcting the next predictions with a factor, f , that represents the prediction error (*Step 6* in Fig. 5).

5 Performance Evaluation

This section first provides an overview of the methodology used to analyse the performance of the proposed ML model. Then, we discuss the results obtained from the measurements campaign carried out following this methodology.

5.1 Experiments Methodology

To assess the solution proposed in this work, we considered an IEEE 802.11 enterprise WLAN, which was modelled via simulation with the aim of generating diverse channel conditions in a controlled environment. For this purpose, we used Network Simulator (NS-3.29), a tool widely used in the community for research purposes. NS-3 is open-source software that is composed of several modules and libraries, and that encompasses a variety of wireless and wired technologies. This feature allows recreating different experimental scenarios that otherwise would be difficult to realize with a limited budget and resources. The simulations were performed on Dell laptops powered by an Intel i7 CPU and running Ubuntu 18.04.02.

Figure 6 provides a schematic view of the setup used in the evaluation. Following the KDP layout discussed in Sect. 3.3, all scenarios comprise an AP on the data plane and an SDN controller on the control plane. Conversely, the proposed ML model for frame length selection is deployed on the management plane. We remind the reader that our solution applies only to the downlink direction since frames are aggregated at the AP. Note that for the sake of simplicity, the control and the knowledge plane are combined in the figure into the entity of the controller. Moreover, we employed a propagation loss model based on the log normal distribution

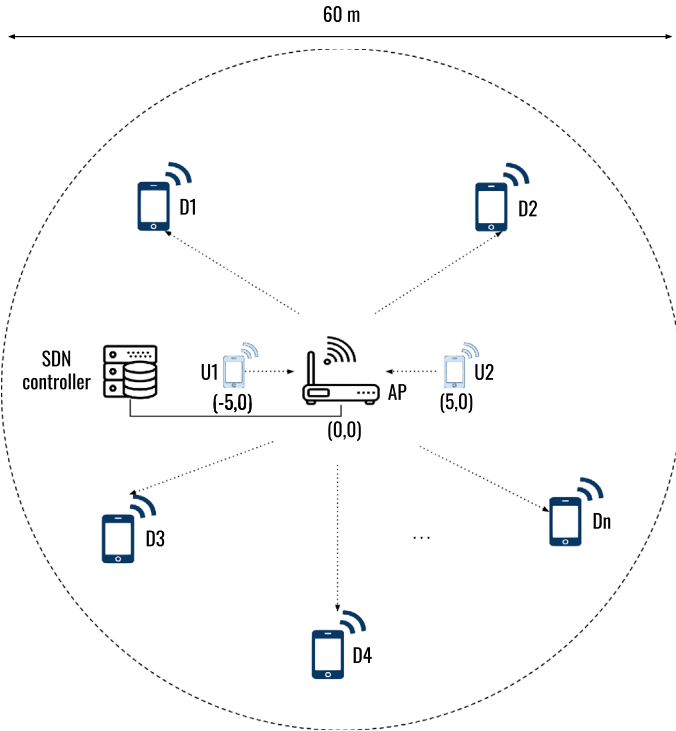


Fig. 6 Schematic diagram of scenarios used in the evaluation including the fusion of the control and knowledge planes, as well as the data plane, which comprises 2 ubiquitous stationary uplink stations and N stations consuming downlink traffic

Table 3 NS-3 Simulation parameters used for the performance evaluation

Parameter	Value	
	Downlink	Uplink
Stations	[5, 20]	2
MCS	[0, 2, 7]	0
Payload size	200 bytes	200 bytes
Max. aggr. size	3839 bytes	–
Aggregated bitrate	[1, 5, 10, 20, 25] Mbps	[2 Mbps, 18% of aggr. DL bitrate]
Position	Variable (30 m radius from AP)	Fixed (at 5 m from AP)
Velocity	Uniform random variable [2–50] m/s	–

(*LogDistancePropagationLossModel* in NS-3), with the exponent equal to 3 and the reference loss set to 46.6777 dB.

Table 3 presents the simulation parameters with respect to the downlink and uplink traffic configuration. On the one hand, the stations receiving downlink traffic from the AP are allowed to move within a circle centered at the origin (where the AP is located) with a radius of 30 m using the *RandomWalk2dMobilityModel* available in NS-3. This distance was chosen while keeping in mind that the maximum range of reception deteriorates as the MCS improves. Then, the position of the stations were altered every 100 ms at a velocity determined by a uniform random variable with a minimum of 2 m/s and a maximum of 50 m/s. Also with regard to the stations, we selected a set of MCS values (0, 2, and 7) instead of letting the default rate control algorithm perform this selection, with the aim of evaluating the performance in a systematic and controlled manner given the close relationship between MCS, channel quality and frame size. The minimum payload size is fixed at 200 bytes, while the maximum A-MSDU aggregation size is set to 3839 bytes. The aggregated downlink bitrate is varied in the range [1, 5, 10, 15, 20, 25] Mbps for two network sizes comprising either 5 or 20 stations. On the other hand, the two uplink stations are set in static positions at $(-5,0)$ and $(5,0)$ m on the X-Y axes, as shown in Fig. 6. The payload size is also 200 bytes while, as opposed to the downlink case, the basic MCS, i.e., MCS 0, is set for the uplink transmissions. Note that no aggregation mechanisms are used for this traffic. To analyse the impact of the background traffic, we distinguish two cases in terms of delivered bitrate to the AP. In the first case, the aggregated uplink traffic is fixed at 2 Mbps while, in the second case, it is adjusted in such a manner that it represents 18% of the aggregated downlink bitrate in each scenario.

We have compared the performance of the proposed scheme, i.e., the RFR model, with the existing baselines: (i) standard delivery without using aggregation mechanisms, and (ii) maximum length defined by the 802.11 standard for A-MSDU aggregation, i.e., 3839 bytes. Each test case is a unique combination of the parameter values shown in Table 3. The simulations were repeated 10 times with a duration of 14 s. However, 2 s were trimmed from the beginning and the end of each simulation in order to collect the measurements when the system has reached a steady state. The results discussed in the next section were obtained by considering several performance metrics, namely goodput, channel utilization, packet loss and retransmission ratio.

5.2 Results and Discussion

As mentioned in the section above, each scenario is evaluated 10 times to ensure the reliability of the results. Then, the output obtained for each scenario is plotted with a 95% confidence interval. The following provides a discussion of the results in relation to two network sizes.

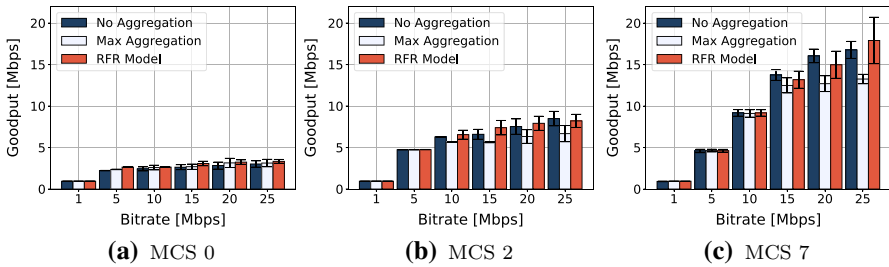


Fig. 7 Goodput comparison in a small-sized network for MCS 0, MCS 2 and MCS 7 in the presence of fixed uplink bitrate




Fig. 8 Retransmission ratio comparison in a small-sized network for MCS 0, MCS 2 and MCS 7 in the presence of fixed uplink bitrate

5.2.1 Small-Sized Network

In this section we consider a small-sized network comprising 5 stations receiving data from the AP. Then, we separate the discussion into scenarios with fixed uplink bitrate and scenarios with proportional uplink bitrate.

5.2.1.1 Fixed uplink bitrate Considering a fixed uplink bitrate, i.e., 2 Mbps regardless of the downlink load, Fig. 7 shows the aggregated downlink goodput achieved by the different schemes with respect to the MCS. As might be expected, the goodput for various downlink bitrates increases with the MCS for all the schemes and saturates as the bitrate gets closer to the PHY rate. In particular, we can observe a greater performance improvement for our solution (RFR model) as the bitrate increases given that in scenarios with good channel qualities and low traffic load the standard schemes are still capable of transmitting the information without needing further optimizations. Furthermore, it can be seen how the RFR model improves upon or equals the results for the standard baselines in nearly all cases.

The most significant exception is found in the case of MCS 7 for bitrates equals to 15 and 20 Mbps (see Fig. 7c), where disabling frame aggregation provides a better performance. This is because the PHY rate is still conducive to efficient transmission regardless of the channel access overhead. However, delivery errors cause all frames in an A-MSDU to be retransmitted, a situation that worsens as the network

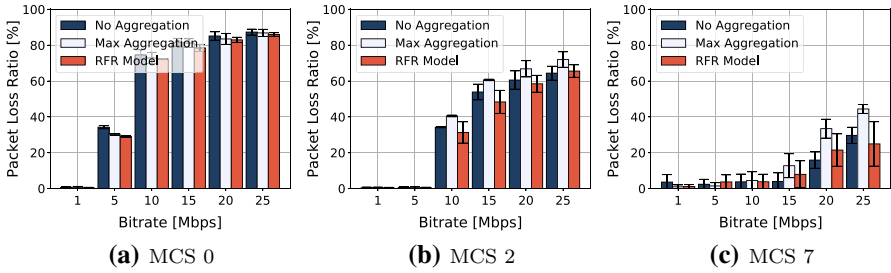


Fig. 9 Packet loss ratio comparison in a small-sized network for MCS 0, MCS 2 and MCS 7 in the presence of fixed uplink bitrate




Fig. 10 Channel utilization comparison in a small-sized network for MCS 0, MCS 2 and MCS 7 in the presence of fixed uplink bitrate

load rises. Consequently, the forceful fixed aggregation scheme performs worse than all the others. This can also be confirmed from the results shown in Fig. 8, where we can see how the retransmission ratio of the delivery without frame aggregation is almost negligible except for low bitrates. This effect is caused by the greater number of frames generated by this scheme which, in case of saturation, are dropped at the MAC layer before even reaching the channel. We also notice that the RFR model requires the lowest retransmission ratio for high MCS values. This fact demonstrates that the maximum aggregation length is not always appropriate and that an adaptive solution is able to provide a better performance.

Figure 9 shows the packet loss ratio for the downlink traffic. First, it can be seen that, as can be expected, the packet loss falls considerably from MCS 0 to MCS 7 due to a higher physical capacity. In fact, high traffic congestion is shown in Fig. 9a when using the basic MCS. Moreover, these measurements demonstrate again how the RFR model offers superior performance with respect to the baselines, as could be deduced from the goodput results presented above in Fig. 7.

Figure 10 illustrates channel utilization, where it is no surprise to find higher values for the basic MCS 0 because the airtime needed to transmit a frame is much higher, therefore causing greater channel contention. It is also important to note that although this ratio could seem low with respect to the steep saturation found, for example, in Fig. 10a, this plot refers purely to the airtime used by the frames that reach the wireless medium. Since the contention causes a high proportion of frames




Fig. 11 Goodput vs. time comparison in a small-sized network for MCS 2 at 15 Mbps in the presence of a fixed uplink bitrate (stations 1, 2 and 3)

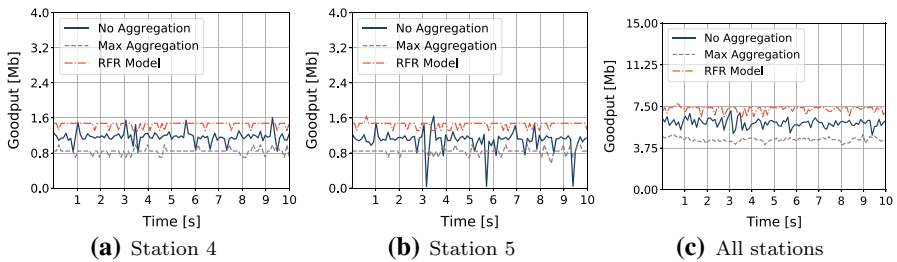


Fig. 12 Goodput vs. time comparison in a small-sized network for MCS 2 at 15 Mbps in the presence of a fixed uplink bitrate (stations 4 and 5, and overall aggregated value)




Fig. 13 Goodput comparison in a small-sized network for MCS 0, MCS 2 and MCS 7 in the presence of proportional uplink bitrate

to be dropped early, the effective channel utilization is lower than what might be expected for such bitrates.

The above results basically reflect the average performance of the stations in each scenario. However, it must be remembered that such stations are placed at random locations and follow different mobility patterns during the simulations. Therefore, considering that the frame length selection is performed on a per-user basis, the average results do not allow the drawing of conclusions on individual user states. In order to have a more accurate vision, Figs. 11 and 12 present the user goodput over time in an example scenario taken from the above ones characterized by the use of




Fig. 14 Retransmission ratio comparison in a small-sized network for MCS 0, MCS 2 and MCS 7 in the presence of proportional uplink bitrate

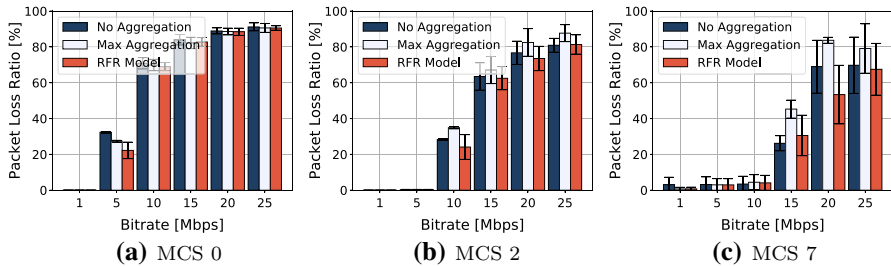


Fig. 15 Packet loss ratio comparison in a small-sized network for MCS 0, MCS 2 and MCS 7 in the presence of proportional uplink bitrate

MCS 2 and a 15 Mbps bitrate. In particular, we can see that except for a few isolated cases (for instance at time equals 2 s in Fig. 11a) the RFR model presents more stable results and a better performance than the standard counterparts. Furthermore, our model is able to provide a fair frame length selection for all the users regardless of their channel conditions. The same conclusions can be drawn from Fig. 12c, in which the evolution of the aggregated goodput is shown.

5.2.1.2 Proportional uplink bitrate These scenarios include an uplink load that is proportional to the downlink bitrate (being equal to 18% of the latter), which makes it range from 0.18 to 4.5 Mbps. Figure 13 shows the goodput achieved in each scenario. In comparison with the fixed uplink load in Fig. 7, we can observe an upward trend up to 10 Mbps since the channel is less frequently accessed by the clients delivering traffic to the AP. By contrast, the opposite effect is noted for bitrates from 15 to 25 Mbps. The channel activity has a greater negative impact on the mechanism without aggregation capabilities (due to channel overhead) and the maximum aggregation scheme (due to the need to retransmit a greater number of MSDUs, as can be seen in Fig. 14). Consequently, the ability to adapt to changes in the channel state and the network load allows the RFR model to obtain better results regardless of the bitrate and the MCS.

Figure 14 shows a decrease in the retransmission ratio in comparison with the previous scenario (Fig. 8). On the one hand, for a low uplink load, i.e. up to 10 Mbps

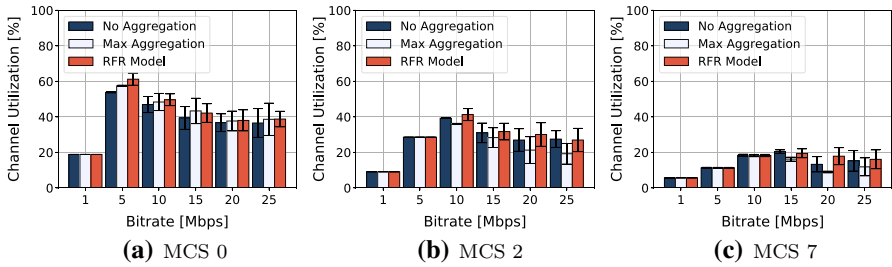


Fig. 16 Channel utilization comparison in a small-sized network for MCS 0, MCS 2 and MCS 7 in the presence of proportional uplink bitrate




Fig. 17 Goodput vs. time comparison in a small-sized network for MCS 2 at 15 Mbps in the presence of a proportional uplink bitrate (stations 1, 2 and 3)

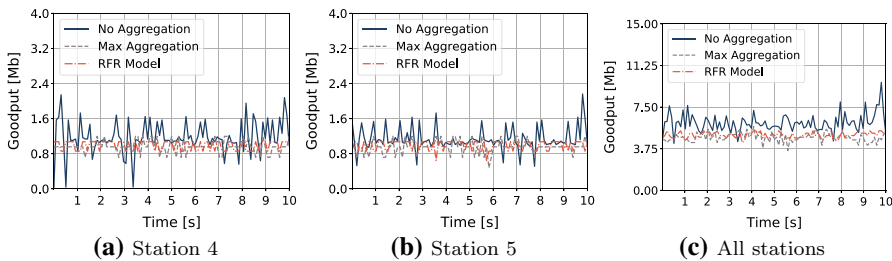


Fig. 18 Goodput vs. time comparison in a small-sized network for MCS 2 at 15 Mbps in the presence of a proportional uplink bitrate (stations 4, and 5, and overall aggregated value)

in downlink, retransmissions fall due to the sharp reduction in the overall network load. On the other hand, for a greater uplink load, the ratio is also lower because of network saturation. Although the RFR model requires the retransmission of more frames due to prediction failures in certain cases, e.g., for MCS 0 in Fig. 14a, this ratio goes down for higher PHY rates.

Figure 15 summarizes the aforementioned findings. Although the network starts to saturate at a certain point, e.g., after exceeding 10 Mbps downlink bitrate (plus the proportional uplink load), on average the RFR model delivers up to 10% more information than the standard baselines. Furthermore, Fig. 16 shows how the performance improvement is achieved at the expense of a negligible increase in channel




Fig. 19 Goodput comparison in a medium-sized network for MCS 0, MCS 2 and MCS 7 in the presence of fixed uplink bitrate

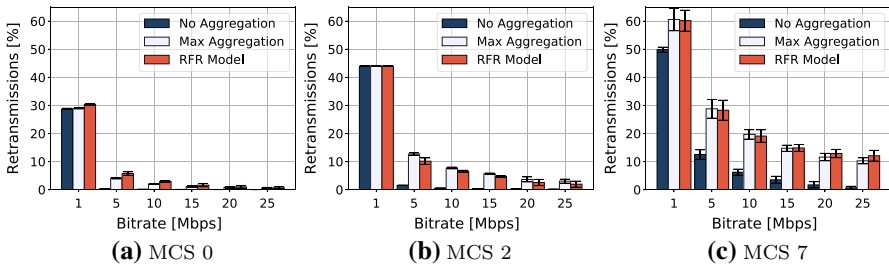


Fig. 20 Retransmission ratio comparison in a medium-sized network for MCS 0, MCS 2 and MCS 7 in the presence of fixed uplink bitrate

utilization. These conclusions are confirmed by the analysis of the user goodput over time. To draw a fair comparison with the previous analysis (in Figs. 11 and 12), Figs. 17 and 18 preserve the same network configuration, namely MCS 2 and a 15 Mbps downlink bitrate. In this case, the per-user improvement provided by the RFR model is less noteworthy than in the previous scenario due to the saturation added by the uplink traffic (almost 3 Mbps). It is worth noting that while the RFR model delivers constant goodput over time, the standard baselines present spikes that significantly affect the user experience. Moreover, as was the case for a fixed uplink bitrate, fairness among stations is also ensured.

5.2.2 Medium-Sized Network

In this section we consider a medium-sized network composed of 20 stations, differentiating, as in Sect. 5.2.1, between scenarios with a fixed and with a proportional uplink bitrate. Note that for the sake of clarity and ease of comprehension, individual measurements over time for each user are omitted.

5.2.2.1 Fixed uplink bitrate Figure 19 shows the goodput comparison in the same scenarios as those discussed in Sect. 5.2.1. Although the trend is similar to the one in the smaller setup, i.e., the goodput rises with the bitrate and the MCS, the performance in this case is substantially lower. The reason for this issue is that, despite maintaining the same aggregated bitrate, the increase in the number of users (4 times




Fig. 21 Packet loss ratio comparison in a medium-sized network for MCS 0, MCS 2, MCS 7 in the presence of fixed uplink bitrate




Fig. 22 Channel utilization comparison in a medium-sized network for MCS 0, MCS 2 and MCS 7 in the presence of fixed uplink bitrate

higher) leads to greater channel contention, therefore reducing the information that is effectively transmitted. Similarly, the retransmission measurements in Fig. 20 present a pattern akin to the experiment in Fig. 8, i.e., they drop gradually as the downlink bitrate increases. Note that the slight rise is simply due to the higher number of stations. It should be remembered that the lower retransmission ratio experienced by the delivery without frame aggregation is not related to a higher performance. In fact, this effect is due to the massive number of frames dropped at the MAC layer due to the transmission of very small packets.

Figure 21 illustrates the difference in terms of packet loss between the RFR model and the standard baselines. The larger number of users causes individual channel conditions to be very different from each other and to change more sharply with mobility. This issue has a substantial impact on the mechanism disabling frame aggregation. Neither is this an optimal scenario for the maximum frame length given the heterogeneity of the user conditions, for which aggregating MSDUs up to 3835 bytes leads to transmission errors and delays. By contrast, our solution improves upon these results due to its capacity to adapt the frame length to the specific user status, by up to 20.83% (with regard to disabling frame aggregation) and 16.28% (with regard to maximum frame length). Finally, a slight increase in channel utilization can be observed in Fig. 22 in comparison with Fig. 10. Although the downlink bitrate is maintained for both network sizes, channel access contention grows in this case with the number of users. Note also that, as was the case in the smaller setup, the increment in channel utilization for the RFR model is almost negligible.

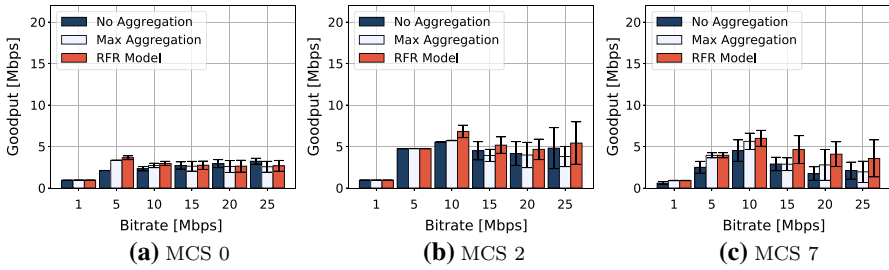


Fig. 23 Goodput comparison in a medium-sized network for MCS 0, MCS 2 and MCS 7 in the presence of proportional uplink bitrate




Fig. 24 Retransmission ratio comparison in a medium-sized network for MCS 0, MCS 2 and MCS 7 in the presence of proportional uplink bitrate

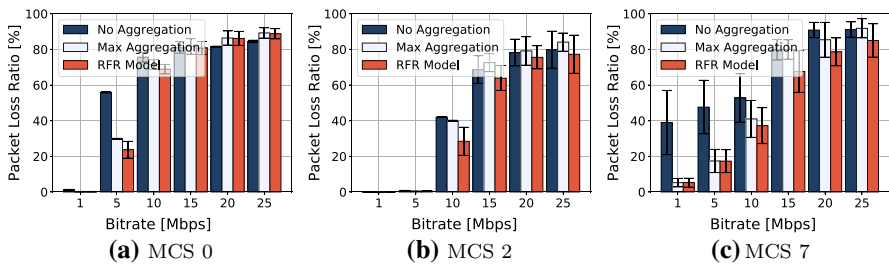


Fig. 25 Packet loss ratio comparison in a medium-sized network for MCS 0, MCS 2 and MCS 7 in the presence of proportional uplink bitrate

5.2.2.2 Proportional uplink bitrate Figure 23 shows the average goodput in the medium-sized setup when setting the uplink load to 18% of the downlink bitrate. By looking at the trends we can see that, as opposed to the above scenarios, the goodput is substantially impaired by both the greater channel contention and the increasing network-wide load. Comparing these results with the ones in Fig. 13, it can be seen that the goodput is lower than in the former case because there are 4 times more stations vying for the channel. In this scenario the standard solutions are only able to deliver a high performance when the traffic load is low. By contrast, the RFR model

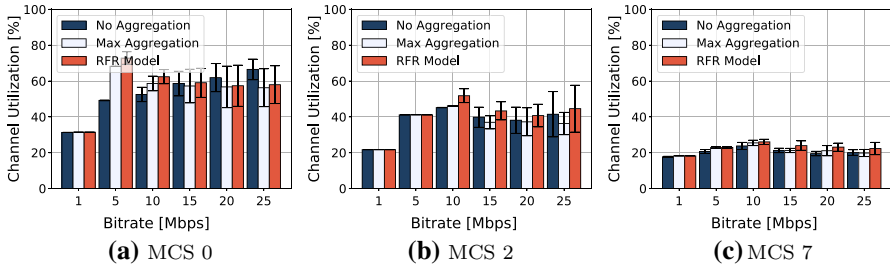


Fig. 26 Channel utilization comparison in a medium-sized network for MCS 0, MCS 2 and MCS 7 in the presence of proportional uplink bitrate

improves upon the goodput of such solutions by up to 69.2% and 83.2% with regard to disabling aggregation and using the maximum frame length, respectively. Minor differences are observed in terms of retransmissions in Fig. 24 in comparison with the use of a fixed uplink bitrate in Fig. 20. As in the previous cases, this ratio generally increases across the downlink bitrates as the MCS rises. Furthermore, the values shown by the RFR model present only slight differences with respect to the mechanism applying the maximum frame length.

The most significant differences between the schemes are most noteworthy in terms of packet loss for a large number of stations. In this sense, Fig. 25 shows that the baseline mechanism avoiding frame aggregation performs worse compared with both maximum aggregation length and the RFR model because the transmission of frames at the original payload size is not effective enough. In addition to that, it can be seen that the inclusion of a variable uplink load imposes greater diversity regarding channel status, therefore making the maximum aggregation length an unsuitable option. This is especially clear with high MCS values, as illustrated in Fig. 25c. By contrast, adjusting the frame size for each user allows the RFR model to hugely reduce the packet loss in most of the cases. Finally, we can observe in Fig. 26 that the performance improvement does not involve an increase in channel utilization. As a matter of fact, these values remain practically constant with respect to the previous scenarios.

After an extensive performance evaluation, we have observed that factors like PHY rate/conditions, absolute downlink load, simultaneous uplink transmissions, and the sheer number of stations play a significant part in the interpretation of the results depicted above. In particular, we can conclude that for small-sized networks in case of fixed uplink rate our model performs better in almost all scenarios. Moreover, the goodput lowers as the uplink load is made proportional to the downlink one, while a consistent increase in absolute goodput can be observed when the uplink rate is fixed. Despite the good performance, this effect can be especially seen as the downlink traffic increases and the channel conditions improve. Specific examples of these event can be found, for instance, for 20/25 Mbps downlink delivery using MCS 7. We also notice this pattern in case of the medium sized network. The goodput improvement given by our model is proven to be better as the PHY rate increases.

6 Conclusions and Future Research

This paper makes the case for the inclusion of ML techniques in the SDN paradigm for automated network management, giving rise to *knowledge-based* and *self-driven* networks. In this work we have focused on assisting the network in making decisions regarding adaptive frame length in enterprise WLANs. The ML-based component introduced is able to perform proactive frame length selection on a per-user basis by considering individual channel conditions and global performance statistics. The performance of our solution has been proved through an extensive performance evaluation and compared with standard baseline mechanisms, showing considerable improvements especially in cases where there is substantial diversity in user channel conditions and the network size increases.

Acknowledgements This work has been performed in the framework of the European Union's Horizon 2020 project 5GZORRO under Grant Agreement No. 871533.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Singh, S., Jha, R.K.: A survey on software defined networking: architecture for next generation network. *J. Netw. Syst. Manag.* **25**(2), 321–374 (2017)
2. Feng, M., Mao, S., Jiang, T.: Enhancing the performance of future wireless networks with software-defined networking. *Front. Inf. Technol. Electron. Eng.* **17**(7), 606–619 (2016)
3. Mestres, A., Rodríguez-Natal, A., Carner, J., Barlet-Ros, P., Alarcón, E., Solé, M., Muntés-Mulero, V., Meyer, D., Barkai, S., Hibbett, M.J., Estrada, G., Ma'arif, K., Coras, F., Ermagan, V., Latapie, H., Cassar, C., Evans, J., Maino, F., Walrand, J., Cabellos, A.: Knowledge-defined networking. *ACM SIGCOMM Comput. Commun. Rev.* **47**(3), 2–10 (2017)
4. Carbonell, J.G., Michalski, R.S., Mitchell, T.M.: Machine learning: a historical and methodological analysis. *AI Mag.* **4**(3), 69–79 (1983)
5. Jiang, C., Zhang, H., Ren, Y., Han, Z., Chen, K., Hanzo, L.: Machine learning paradigms for next-generation wireless networks. *IEEE Wirel. Commun.* **24**(2), 98–105 (2017)
6. Keller, A.: Challenges and directions in service management automation. *J. Netw. Syst. Manag.* **25**(4), 884–901 (2017)
7. Nawrocki, P., Sniezynski, B.: Adaptive service management in mobile cloud computing by means of supervised and reinforcement learning. *J. Netw. Syst. Manag.* **26**(1), 1–22 (2018)
8. Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications. ANSI/IEEE Std 802.11 (2016)
9. Kriara, L., Marina, M.K., Farshad, A.: Characterization of 802.11n wireless LAN performance via testbed measurements and statistical analysis. In: *Proc. of IEEE SECON*. New Orleans, USA (2013)
10. Xie, J., Yu, F.R., Huang, T., Xie, R., Liu, J., Wang, C., Liu, Y.: A survey of machine learning techniques applied to software defined networking (SDN): research issues and challenges. *IEEE Commun. Surv. Tutor.* **21**(1), 393–430 (2019)

11. Riggio, R., Marina, M.K., Schulz-Zander, J., Kuklinski, S., Rasheed, T.: Programming abstractions for software-defined wireless networks. *IEEE Trans. Netw. Serv. Manag.* **12**(2), 146–162 (2015)
12. Li, L., Beach, M., Nejabati, R., Simeonidou, D.: Building SDN agent for wireless local area networks. In: Proc. of IEEE WCNC. Marrakesh, Morocco (2019)
13. Jang, R., Cho, D., Noh, Y., Nyang, D.: RFlow+: an SDN-based WLAN monitoring and management framework. In: Proc. of IEEE INFOCOM. Atlanta, GA, USA (2017)
14. Li, L., Oikonomou, G., Beach, M., Nejabati, R., Simeonidou, D.: An SDN agent-enabled rate adaptation framework for WLAN. In: Proc. of IEEE ICC. Shanghai, China (2019)
15. Coronado, E., Riggio, R., Villalon, J., Garrido, A.: Joint mobility management and multicast rate adaptation in software-defined enterprise WLANs. *IEEE Trans. Netw. Serv. Manag.* **15**(2), 625–637 (2018)
16. Richart, M., Baliosian, J., Serrat, J., Gorricho, J.L.: Resource slicing in virtual wireless networks: a survey. *IEEE Trans. Netw. Serv. Manag.* **13**(3), 462–476 (2016)
17. M. Richart, J. Baliosian, J. Serrati, J. L. Gorricho, R. Agiero, Agoulmine, N.: Resource allocation for network slicing in WiFi access points. In: Proc. of IEEE CNSM. Tokyo, Japan (2017)
18. Li, T., Ni, Q., Malone, D., Leith, D., Xiao, Y., Turletti, T.: Aggregation with fragment retransmission for very high-speed WLANs. *IEEE/ACM Trans. Netw.* **17**(2), 591–604 (2009)
19. Khorov, E., Kiryanov, A., Lyakhov, A., Bianchi, G.: A tutorial on IEEE 802.11 ax high efficiency WLANs. *IEEE Commun. Surv. Tutor.* **21**(1), 197–216 (2018)
20. Kim, D., An, S.: Throughput enhancement by dynamic frame aggregation in multi-rate WLANs. In: Proc. of IEEE SCVT. Eindhoven, Netherlands (2012)
21. Shin, C., Park, H., Kwon, H.M.: PHY-supported frame aggregation for wireless local area networks. *IEEE Trans. Mob. Comput.* **13**(10), 2369–2381 (2014)
22. Kim, G., Shin, C., Park, H.: Adaptive frame size estimation using extended Kalman filter for high-stressed WLANs. In: Proc. of IEEE PIMRC. Sydney, Australia (2012)
23. Byeon, S., Yoon, K., Lee, O., Cho, W., Oh, S., Choi, S.: MoFA: Mobility-aware frame aggregation in Wi-Fi. In: Proc. of ACM CoNEXT. Sydney, Australia (2014)
24. Sarret, M.G., Ashta, J.S., Mogensen, P., Catania, D., Cattoni, A.F.: A multi-QoS aggregation mechanism for improved fairness in WLAN. In: Proc. of IEEE VTC. Las Vegas, NV, USA (2013)
25. Gautam, B.: AMSDU vs. AMPDU: A Brief Tutorial on WiFi Aggregation Support (2017)
26. Skordoulis, D., Ni, Q., Chen, H., Stephens, A.P., Liu, C., Jamalipour, A.: IEEE 802.11n MAC frame aggregation mechanisms for next-generation high-throughput WLANs. *IEEE Wirel. Commun.* **15**(1), 40–47 (2008)
27. Maqhat, B., Baba, M.D., Rahman, R.A.: A-MSDU real time traffic scheduler for IEEE802.11n WLANs. In: Proc. of IEEE Symposium on Wireless Technology and Applications. Bandung, Indonesia (2012)
28. Maqhat, B., Dani Baba, M., Rahman, R.A., Saif, A.: Performance analysis of fair scheduler for A-MSDU aggregation in IEEE802.11n wireless networks. In: 2014 2nd International Conference on Electrical, Electronics and System Engineering (ICEESE). Kuala Lumpur, Malaysia (2014)
29. Saif, A., Othman, M.: A reliable A-MSDU frame aggregation scheme in 802.11n wireless networks. *Procedia Comput. Sci.* **21**, 191–198 (2013)
30. Emna, C., Gueguen, C., Fourati, L., Cousin, B., Kamoun, L.: Dynamic frame aggregation scheduler for multimedia applications in IEEE 802.11n networks. *Trans. Emerg. Telecommun. Technol.* **28**, 1–17 (2015)
31. Saldana, J., Ruiz-Mas, J., Almodóvar, J.: Frame aggregation in central controlled 802.11 WLANs: the latency versus throughput tradeoff. *IEEE Commun. Lett.* **21**(11), 2500–2503 (2017)
32. Azhari, S.V., Gurbuz, O., Ercetin, O.: QoS based aggregation in high speed IEEE802.11 wireless networks. In: Proc. of Mediterranean Ad Hoc Networking Workshop. Barcelona, Spain (2016)
33. Hyoung Lee, W., Young Hwang, H.: A-MPDU aggregation with optimal number of MPDUs for delay requirements in IEEE 802.11ac. *PLoS ONE* **14**, 1–17 (2019)
34. Seytazarov, S., Kim, Y.: QoS-aware adaptive A-MPDU aggregation scheduler for enhanced VoIP capacity over aggregation-enabled WLANs. In: Proc. of IEEE NOMS. Taipei, Taiwan (2018)
35. He, X., Li, F.Y., Lin, J.: Link adaptation with combined optimal frame size and rate selection in error-prone 802.11n networks. In: Proc. of IEEE ISWCS. Reykjavik, Iceland (2008)
36. Abdallah, S., Blostein, S.D.: Joint rate adaptation, frame aggregation and MIMO mode selection for IEEE 802.11ac. In: Proc. of IEEE Wireless Communications and Networking Conference. Doha, Qatar (2016)

37. Ginzburg, B., Kesselman, A.: Performance analysis of A-MPDU and A-MSDU aggregation in IEEE 802.11n. In: Proc. of IEEE Sarnoff Symposium. Princeton, NJ, USA (2007)
38. Kim, M., Park, E.C., Choi, C.H.: Adaptive two-level frame aggregation for fairness and efficiency in IEEE 802.11n wireless LANs. *Mob. Inf. Syst.* **2015**, 1–14 (2015)
39. Lin, Y., Wong, V.W.S.: Frame aggregation and optimal frame size adaptation for IEEE 802.11n WLANs. In: Proc. of IEEE Globecom. San Francisco, CA, USA (2006)
40. Liu, K., Li, C.: An efficient adaptive frame aggregation scheme in vehicular Ad Hoc networks. In: Proc. of IEEE WCSP. Nanjing, China (2017)
41. Zhou, X., Boukerche, A.: AFLAS: an adaptive frame length aggregation scheme for vehicular networks. *IEEE Trans. Veh. Technol.* **66**(1), 855–867 (2017)
42. Bianchi, G.: Performance analysis of the IEEE 802.11 distributed coordination function. *IEEE J. Sel. Areas Commun.* **18**(3), 535–547 (2000)
43. Simeone, O.: A very brief introduction to machine learning with applications to communication systems. *IEEE Trans. Cogn. Commun. Netw.* **4**(4), 648–664 (2018)
44. Usama, M., Qadir, J., Raza, A., Arif, H., Yau, K.A., Elkhatib, Y., Hussain, A., Al-Fuqaha, A.: Unsupervised machine learning for networking: techniques, applications and research challenges. *IEEE Access* **7**, 65579–65615 (2019)
45. Tran, N.H., Bao, W., Zomaya, A., Nguyen, M.N.H., Hong, C.S.: Federated learning over wireless networks: optimization model design and analysis. In: Proc. of IEEE INFOCOM. Paris, France (2019)
46. Zappone, A., Renzo, M.D., Debbah, M.: Wireless networks design in the era of deep learning: model-based, AI-based, or both? *IEEE Trans. Commun.* **67**, 7331–7376 (2019)
47. Shalev-Shwartz, S., Ben-David, S.: *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press, Cambridge (2014)
48. Weisberg, S.: *Applied Linear Regression*, third edn. Wiley, Hoboken (2005)
49. Boser, B.E., Guyon, I.M., Vapnik, V.N.: A training algorithm for optimal margin classifiers. In: Proc. of Workshop on Computational Learning Theory. Pittsburgh, USA (1992)
50. Quinlan, R.J.: Learning with continuous classes. In: Proc. of Conference on Artificial Intelligence. San Jose, CA, USA (1992)
51. Breiman, L.: Random forests. *Mach. Learn.* **45**(1), 5–32 (2001)
52. Cutler, A., Cutler, D., Stevens, J.: *Random Forests*, vol. 45, pp. 157–176. Springer, Berlin (2011)
53. Ramirez-Perez, C., Ramos, V.: SDN meets SDR in self-organizing networks: Fitting the pieces of network management. *IEEE Commun. Mag.* **54**, 48–57 (2016)
54. Zehl, S., Zubow, A., Döring, M., Wolisz, A.: ResFi: a secure framework for self organized Radio Resource Management in residential WiFi networks. In: Proc. of IEEE WoWMoM. Coimbra, Portugal (2016)
55. Nunes, B.A.A., Santos, M.A.S., de Oliveira, B.T., Margi, C.B., Obraczka, K., Turletti, T.: Software-defined-networking-enabled capacity sharing in user-centric networks. *IEEE Commun. Mag.* **52**(9), 28–36 (2014)
56. Clark, D.D., Partridge, C., Ramming, J.C., Wroclawski, J.T.: A knowledge plane for the internet. In: Proc. of ACM SIGCOMM. Karlsruhe, Germany (2003)
57. van der Hoof, J., Claeys, M., Bouten, N., Wauters, T., Schönwälder, J., Pras, A., Stiller, B., Charalambides, M., Badonnel, R., Serrat, J., dos Santos, C.R.P., De Turck, F.: Updated taxonomy for the network and service management research field. *J. Netw. Syst. Manag.* **26**(3), 790–808 (2018)
58. Zeljković, E., Slamnik-Kriještorac, N., Latré, S., Marquez-Barja, J.M.: ABRAHAM: machine learning backed proactive handover algorithm using SDN. *IEEE Trans. Netw. Serv. Manag.* **16**(4), 1–14 (2019)
59. Del Testa, D., Danieletto, M., Di Nunzio, G.M., Zorzi, M.: Estimating the number of receiving nodes in 802.11 networks via machine learning techniques. In: Proc. of IEEE GLOBECOM. Washington, DC, USA (2016)
60. Han, Z., Lei, T., Lu, Z., Wen, X., Zheng, W., Guo, L.: Artificial intelligence-based handoff management for dense WLANs: a deep reinforcement learning approach. *IEEE Access* **7**, 31688–31701 (2019)
61. Cheng, Y., Chou, H., Chang, R.Y.: Machine-learning indoor localization with access point selection and signal strength reconstruction. In: Proc. of IEEE VTC Spring. Nanjing, China (2016)
62. Jeunen, O., Bosch, P., Herwegen, M.V., Doorselaer, K.V., Godman, N., Latré, S.: A machine learning approach for IEEE 802.11 channel allocation. In: Proc. of IEEE CNSM. Rome, Italy (2018)

63. Wang, F., Gong, W., Liu, J., Wu, K.: Channel selective activity recognition with wifi: a deep learning approach exploring wideband information. *IEEE Trans. Netw. Sci. Eng.* 1–14 (2018)
64. Ferreira, C., Guardalben, L., Gomes, T., Sargento, S., Salvador, P., Robalo, D., Velez, F.J.: Supporting unified distributed management and autonomic decisions: design, implementation and deployment. *J. Netw. Syst. Manag.* **25**, 416–456 (2016)
65. Abbas, M., Elhamshary, M., Rizk, H., Torki, M., Youssef, M.: WiDeep: wiFi-based accurate and robust indoor localization system using deep learning. In: *Proc. of IEEE PerCom. Kyoto, Japan* (2019)
66. Wang, X., Gao, L., Mao, S., Pandey, S.: CSI-based fingerprinting for indoor localization: a deep learning approach. *IEEE Trans. Veh. Technol.* **66**(1), 763–776 (2017)
67. Salamah, A.H., Tamazin, M., Sharkas, M.A., Khedr, M.: An enhanced WiFi indoor localization system based on machine learning. In: *Proc. of IEEE IPIN. Alcalá de Henares, Spain* (2016)
68. Azzouni, A., Pujolle, G.: NeuTM: A neural network-based framework for traffic matrix prediction in SDN. In: *Proc. of IEEE NOMS. Taipei, Taiwan* (2018)
69. Moura, H.D., Fernandes Macedo, D., Vieira, M.A.M.: Automatic quality of experience management for WLAN networks using multi-armed bandit. In: *Proc. of IEEE IM. Arlington, VA, USA* (2019)
70. Allahdadi, A., Morla, R.: Anomaly detection and modeling in 802.11 wireless networks. *J. Netw. Syst. Manag.* **27**(1), 3–38 (2019)
71. Lin, P., Lin, T.: Machine-learning-based adaptive approach for frame-size optimization in wireless LAN environments. *IEEE Trans. Veh. Technol.* **58**(9), 5060–5073 (2009)
72. Elwekeil, M., Wang, T., Zhang, S.: Deep learning for joint adaptations of transmission rate and payload length in vehicular networks. *Sensors* **19**, 1113–1134 (2019)
73. Song, L., Striegel, A.: Leveraging frame aggregation for estimating wifi available bandwidth. In: *Proc. of IEEE SECON. San Diego, CA, USA* (2017)
74. Xia, D., Hart, J., Fu, Q.: Evaluation of the Minstrel rate adaptation algorithm in IEEE 802.11g WLANs. In: *Proc. of IEEE ICC. Budapest, Hungary* (2013)
75. Genuer, R., Poggi, J.M., Tuleau-Malot, C.: Variable selection using random forests. *Pattern Recognit. Lett.* **31**(14), 2225–2236 (2010)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

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. She received the M.Sc. degrees in Computer Engineering and Advanced Computer Technologies in 2014 and 2015. 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 mobile communications, network slicing, SDN, NFV, AI-driven network management.

Abin Thomas is currently pursuing the Ph.D. degree with FBK (Smart Networks and Services Unit) and University of Trento. He received the Bachelor Degree in Electronics and Telecommunications Engineering from Goa University, India in 2014. He received the M.Sc degree from EIT Digital (KTH and University of Trento) in 2019. His research interests include SDN, NFV, and application of machine learning in wireless and mobile networks.

Roberto Riggio is Head of the Smart Networks and Services Unit at FBK. His research interests include smart networks, network slicing and distributed network service orchestration. He has published more than 120 papers in journals and conferences, and has generated more than 3M Euro in competitive funding. He received several awards including the IEEE INFOCOM Best Demo Award (2013, 2017) and the IEEE CNSM Best Paper Award (2015). He serves on the TPC/OC of leading conferences and is associate editor for the Wiley International Journal of Network Management, the Springer Wireless Networks, and the IEEE Transactions on Network and Service Management. He is ACM member and IEEE Senior member.