

AI-Empowered Software-Defined WLANs

Estefanía Coronado, Suzan Bayhan, Abin Thomas, and Roberto Riggio

The authors introduce aiOS, which is an AI-based platform that builds toward the autonomous management of SD-WLANs. Their proposal is aligned with the most recent trends in in-network AI promoted by the ITU Telecommunication Standardization Sector (ITU-T) and with the architecture for disaggregated radio access networks promoted by the Open Radio Access Network Alliance.

ABSTRACT

The complexity of wireless and mobile networks is growing at an unprecedented pace. This trend is proving current network control and management techniques based on analytical models and simulations to be impractical, especially if combined with the data deluge expected from future applications such as augmented reality. This is particularly true for software-defined wireless local area networks (SD-WLANs). It is our belief that to battle this growing complexity, future SD-WLANs must follow an artificial intelligence (AI)-native approach. In this article, we introduce aiOS, which is an AI-based platform that builds toward the autonomous management of SD-WLANs. Our proposal is aligned with the most recent trends in in-network AI promoted by the ITU Telecommunication Standardization Sector (ITU-T) and with the architecture for disaggregated radio access networks promoted by the Open Radio Access Network Alliance. We validate aiOS in a practical use case, namely frame size optimization in SD-WLANs, and we consider the long-term evolution, challenges, and scenarios for AI-assisted network automation in the wireless and mobile networking domain.

INTRODUCTION

Network softwarization refers to the decoupling of the software implementing a network function from the hardware running it. The arguments in favor of network softwarization are manifold. First, it reduces the deployment cycles for new network functions. Second, by standardizing the underlying hardware and accessing it using well-defined abstractions, network control and management are made easier. Third, the physical network becomes an open arena for innovation, which leads to advances in the existing services and the creation of new ones. Software-defined networking (SDN) [1] is a key enabler of network softwarization.

SDN is essential for the control and management of wireless networks and, in particular, for wireless LANs (WLANs), which today constitute the most popular form of wireless access connectivity due to their performance and low deployment cost. In fact, a four-fold increase in the number of hotspots is foreseen by 2023, resulting in a total of 628 million public hotspots. In addition, the average Wi-Fi speed will exceed 91.6 Mb/s by 2023 [2]. Therefore, it is crucial to maximize resource utilization and efficiency in software-defined WLANs (SD-WLANs), thus ensuring

that future applications and services can be efficiently consumed by mobile users.

The literature on SD-WLANs is ample, and is excellently reviewed in [3]. Nevertheless, the promise of SDN to deliver a more manageable network, whose behavior could be specified through high-level applications running on top of a logically centralized controller, led to the proliferation of complex approaches to solve highly specific problems, and to the creation of a multitude of network configurations. Although the global network view at the SDN controller enables data-driven network management based on artificial intelligence (AI) and machine learning (ML) approaches [4], how these solutions should be integrated into the existing networks without increasing their complexity remains an open question.

By learning from the success of the Internet obtained through layering and standardized interfaces, we argue that the design of the next-generation AI-enabled SD-WLANs should follow a similar path, namely high-level abstractions, unified data models, and well-defined interfaces. This trend is also promoted by the International Telecommunication Union Telecommunication Standardization Sector (ITU-T) [5] and the Open Radio Access Network (O-RAN) Alliance [6]. In this work, we introduce aiOS, the first open source O-RAN near-real-time RAN intelligent controller (RIC).¹ We extend our previous work [7] as follows:

- First, we provide an overview of the state of the art of ML-based network management schemes for WLANs.
- Second, we describe the challenges, requirements, and architecture of the aiOS network automation platform.
- Third, by focusing on a practical use case, namely frame aggregation in SD-WLANs, we show how aiOS improves network goodput by up to 55 percent.
- Fourth, we discuss the future challenges and applications for automation in wireless and mobile networks.

ARTIFICIAL INTELLIGENCE IN SOFTWARE-DEFINED WIRELESS NETWORKS

By reinterpreting the concept of control and user plane separation, and by introducing a logically centralized controller and the associated control applications, SDN has played a key role in taming the complexity of current networks. However, full control and user plane separation in SD-WLANs is not trivial. This is because in SD-WLANs it is

¹ aiOS is released under APACHE 2.0 License; <http://5g-empower.github.io/>

essential to draw a line between network control, which deals with fast timescale operations that cannot be offloaded to the logically centralized controller, and network management, which deals with monitoring and reconfiguration operations. In this section we first introduce the concept of AI for networks (including the standardization aspects). Then we discuss the challenges it raises and describe a few deployment scenarios in which an AI-empowered SD-WLAN could provide benefits.

AI FOR NETWORKS

With the success of AI in many domains (e.g., computer vision), the idea of an intelligent network that can observe its environment and adapt accordingly is gaining momentum. The broad consensus is that future wireless and mobile networks will need an increased level of intelligence. Nevertheless, it is still hard to see shared goals and methodologies when it comes to integrating AI/ML solutions into production networks. A key challenge is the availability of unified frameworks supporting the various steps in an ML pipeline, namely data collection, filtering, analysis, and decision making [8]. Moreover, interfaces between the data sources and the ML models as well as between the ML models and the sink nodes must be standardized for faster deployment of new solutions. Prior works [4] have reported various successful approaches involving the use of ML techniques to address different networking challenges, including resource scheduling at the medium access control (MAC) layer, mobility management at the network layer, and precise location at the application layer. Finally, there are plenty of examples highlighting the role that AI and ML will play in beyond 5G networks, including 6G [9].

STANDARDIZATION

A recent proposal by ITU-T [5] aims to define a framework meeting several requirements for integrating ML functions in wireless networks. These requirements are:

- Support for multiple data sources and communication networks
- Distributed ML functionalities at various network levels
- Flexible deployment of ML functionalities depending on the requirements and the available network resources
- Flexibility to change the data sources and sinks of the ML applications
- Standard syntax to define ML applications

ITU-T's proposal is technology-agnostic and can be adapted to Wi-Fi networks. However, no such effort has been initiated in IEEE, and there is no planned amendment aimed at introducing network automation into the 802.11 family of standards. Conversely, academia actively looks into the application of AI/ML to various aspects of (SD-)WLANs. Due to space constraints, we mention only two particular works [7, 10]. The first study [10] introduces deep learning into the low-level Wi-Fi stack, while the second work [7] focuses on the higher resource management layers, such as enhanced distributed channel access (EDCA) optimization and mobility management.

CHALLENGES

Several challenges must be tackled before ML techniques can be applied to SD-WLANs. First, collecting a sufficient volume of training data can take considerable time. Second, network operations can be negatively affected when the ML-based solutions are deployed on the production network. Network simulators can help cope with these challenges by generating training data from a wide range of scenarios. Moreover, analyzing the ML-based solutions on network simulators helps in assessing their performance and pitfalls prior to actual deployment on the production network.

Network simulators differ significantly from the tools that are typically used in AI communities. There is thus the need for interfaces and abstractions that hide the complexity of one domain from the other if in-network AI becomes a reality. Works aiming to provide well-defined ML libraries for network simulators can already be found. A notable example is ns3-gym [11], which integrates ns-3 with a widely used toolkit for reinforcement learning named OpenAI Gym. This allows the use of the AI functions implemented in OpenAI Gym by any ns-3 protocol. The authors in [12] go further and discuss the integration of simulators in ITU-T's architecture to provide the data for training and testing ML models.

In this work, we take a fundamental step forward in network intelligence by presenting aiOS, which is an AI-based platform for the control and management of SD-WLANs. aiOS embeds state-of-the-art ML toolboxes to provide a full intelligence platform, whose design is driven by AI and aims to drive future AI-powered networking applications and services. Like [11], the power of aiOS is that the AI/ML tools are hidden behind high-level programming abstractions that allow network experts to build novel and effective resource control policies with limited knowledge of the underlying AI/ML machinery. Nevertheless, a basic understanding of data filtering and the ML models used is needed.

DEPLOYMENT SCENARIOS

The range of services that Wi-Fi networks can deliver is extremely diverse and spans from broadband Internet access to Industry 4.0. AI-empowered network operations are the key to unlocking the full potential of SD-WLANs in such scenarios. In this section, we discuss three use cases (depicted in Fig. 1) in which aiOS can bring tangible benefits.

Residential Deployments: A residential Wi-Fi network typically comprises several access points (APs) and various stations consuming delay-sensitive services (e.g., online gaming) and delay-tolerant services (e.g., web browsing). In this scenario, aiOS could be deployed at the Internet service provider (ISP)'s premises, managing multiple residential deployments. Hierarchical solutions, in which a local aiOS instance runs at the customer site while a second-tier instance aggregates information at the ISP, can also be envisioned. In both cases federated learning approaches can be used to train the SD-WLAN control applications across a wide number of deployments.

Enterprise Deployments: As opposed to residential networks, enterprise deployments are

Network simulators differ significantly from the tools that are typically used in AI communities. There is thus the need for interfaces and abstractions that hide the complexity of one domain from the other if in-network AI becomes a reality. Works aiming to provide well-defined ML libraries for network simulators can already be found.

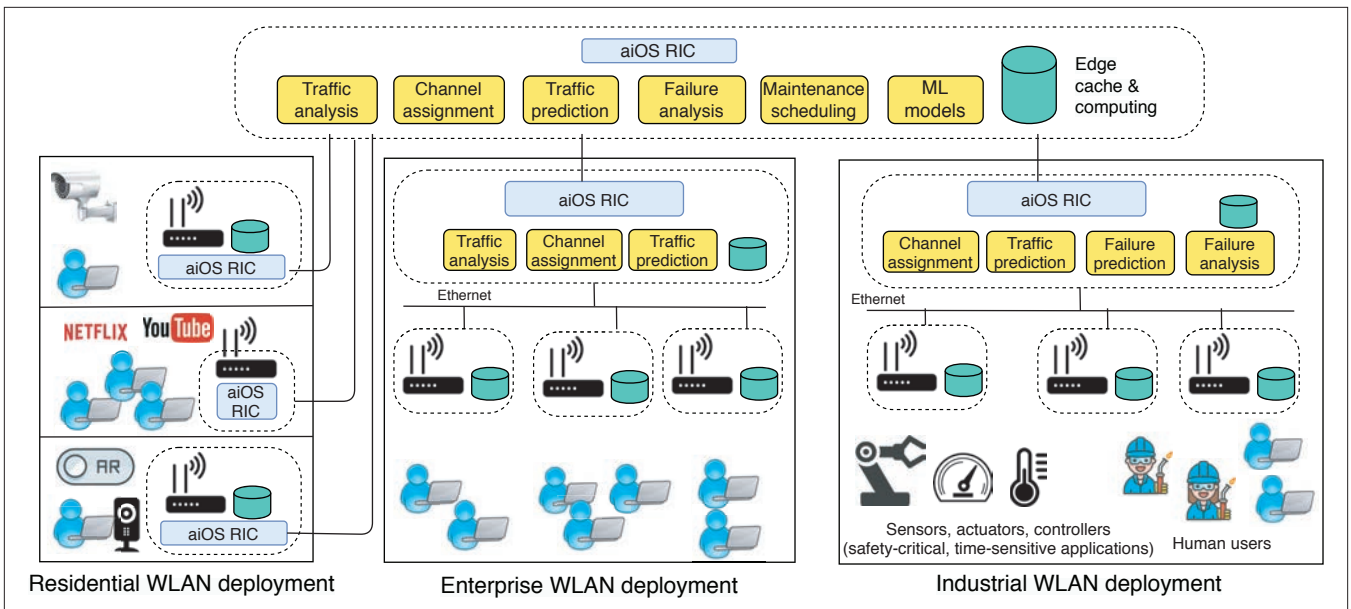


FIGURE 1. Residential, enterprise, and industrial aiOS deployment scenarios.

typically planned. In this scenario, we can envision aiOS running at the customer's site, adapting the network configuration to the changing environment. Nevertheless, a second-tier aiOS instance can still be used as a way to collect (in a privacy preserving fashion) training data from a more diverse set of scenarios, thus improving the accuracy of the ML models (privacy and security aspects need to be carefully considered in this case).

Industrial Deployments: Industrial Wi-Fi deployments may consist of cyber-physical systems with numerous sensors, controllers, and actuators coexisting with human users [13]. Industrial applications are often time-sensitive and require ultra-reliability. Moreover, the amount of data exchanged can vary depending on the application (control system or failure prediction). While such challenges may seem more suitable for 5G networks, in May 2019 the IEEE 802.11be Extremely High Throughput Task Group started work on extending Wi-Fi into the time-sensitive network domain [14]. In this context, aiOS can be the element on top of which data-driven optimization tasks such as traffic prediction, failure analysis, and demand-attentive resource management can be implemented.

THE AIOS SYSTEM

Early works on SDN attempted to tame network complexity by putting into the hands of *network programmers* powerful abstractions and languages to control SDN networks [15]). Such attempts resulted in solutions that either hid too many aspects of the network, up to the point of preventing the implementation of any meaningful task, or exposed too many low-level details to the network programmer. While SDN, instead of introducing new concepts, proposed a different way of arranging network capabilities, the aim of aiOS is to provide a coherent, practical, and data-driven AI platform for SD-WLANs. It is our belief that this approach is pivotal in enabling reutilization of the best practices in AI within the networking

domain and in leveraging the huge amount of data that will be generated by applications. While our work focuses on SD-WLANs, since they are the most popular wireless access technology, the same principles can also be extended to 5G networks and beyond.

The aiOS design leverages the ML pipeline concept proposed by ITU-T [5] and the intelligent and disaggregated RAN principles put forward by O-RAN [6]. In this section, we review the key challenges, both ML-related and network-related, limiting the deployment of ML in wireless networks, and then we discuss the aiOS architecture.

ML-RELATED CHALLENGES

Dataset Availability and Labeling (ML-1): The availability of datasets is an integral part of the success of ML and, in particular, of supervised and semi-supervised learning in SD-WLANs. However, labeled datasets are usually difficult to obtain for two main reasons:

- High-quality datasets usually come from operational networks, and operators can be very reluctant to share them as they do not wish to provide competitors with valuable insights into their business strategy.
- Significant amounts of time and resources are needed to collect and label datasets, which are typically collected using field measurements.

Nevertheless, some open datasets are publicly available, and synthetic ones can be generated using network simulators (prior to validation on a real deployment).

Support for Heterogeneous Data Sources (ML-2): To be successful, ML-based network management solutions need to pull data not only from the network stack but also from the applications and services running on the network so that the impact of such applications (e.g., tele-medicine) on network configuration or optimization can be assessed. Future IEEE 802.11ax-based WLANs will by themselves be a huge source of highly distributed monitoring data. Given the expected num-

ber of hotspots in the coming decades, this could easily outpace the data generated by 5G and beyond networks. It is also worth mentioning that different data gathering systems may deal with different constraints in terms of data timestamping and retention, which are two key aspects to consider for training/updating ML models.

Support for Current and Future ML Toolboxes (ML-3): The success of ML in many fields has led to a rich set of libraries that simplify the application of ML solutions to a variety of problems. Such libraries also shield programmers from the complexities of the hardware acceleration tools, such as graphics processing units (GPUs) or field programmable gate arrays (FPGAs). On the downside, due to their success, these libraries evolve at a very high pace. Hence, it is crucial for an AI-based resource management platform for SD-WLANs to shelter network experts from the low-level details of the ML toolkits. Conversely, great care is required in abstracting the expected outputs (i.e., the knobs of the network configuration that the ML solutions must turn and tune).

Training Time (ML-4): Training time refers to the time needed to build an ML model from the available datasets. This includes offline training and re-training while new data are gathered. This process can be slow, especially if the initial training set is big. Moreover, due to the stochastic nature of the channel, the behavior of wireless networks is highly volatile, which can cause offline trained models to fail to generalize when deployed in the field. One way to address this problem is by using federated learning to train different models at distributed sites, consolidate them at a central site, and then redistribute the consolidated model. Another option is to leverage reinforcement learning, which, instead of retraining, adapts its behavior on the basis of past decisions following the punishment/reward model.

NETWORK-RELATED CHALLENGES

Ease of Use (N-1): In order for AI-empowered SD-WLANs to be adopted in a wide range of use cases and scenarios, popular AI concepts and tools need to become an integral part of the development pipeline of SDN experts. If we draw a parallel with the everything-as-a-service concept that reshaped cloud computing by encapsulating the management of complex infrastructures behind a service-based model, we need an ML-as-a-service model providing SDN developers with easier access to the most relevant ML toolkits for the problem they need to tackle. Such toolkits should also be integrated with the current network management workflows and best practices.

Interpretability (N-2): One of the key challenges limiting the inclusion of ML in wireless and mobile networks is the interpretability of the results. Operators are reluctant to deploy black-box solutions that make the network even harder to debug. State-of-the-art ML solutions impose a trade-off between highly accurate but non-interpretable models (suitable for nonlinear relationships and requiring long computation times) vs. interpretable but not very accurate models (suitable for linear relationships and more computationally tractable). The former category comprises approaches such as deep learning, while the latter covers regression and classification techniques. A

particularly promising solution is random forest (and derivatives), which combines interpretability, fast training, and a relatively good ability to generalize results [7].

Computational Complexity (N-3): It refers to the complexity of using a trained model. Some techniques, such as deep learning, are very demanding in terms of computational and storage resources, and this can be a challenge in certain deployments. For example, the expected deluge of Wi-Fi hotspots and small 5G cells will not be able to use, due to cost and power consumption requirements, specialized hardware to run complex ML models. Large macrocell deployments already have a notable power consumption footprint coupled with generally larger deployment sites. In such cases, deploying specialized acceleration units might be a better option. Moreover, it is reasonable to assume that in time the hardware required to support ML solutions will be embedded in the FPGAs already used by vendors in their products.

Support for Multiple Radio Access Technologies (N-4): Something that is closely related to the *heterogeneous data sources* challenge is support for a diverse ecosystem of radio access technologies (RATs). In fact, it is expected that 2G (3G could be the first technology to be withdrawn), 4G, 5G, Wi-Fi, and beyond 5G networks will coexist for decades to come. This requires the AI-based management platform to interface with highly heterogeneous technologies, each characterized by very different design choices. If we just consider the MAC layer, networks use time-division multiple access (TDMA) in 2G, carrier sense multiple access with collision avoidance (CSMA/CA) in Wi-Fi, orthogonal frequency-division multiple access (OFDMA) in 4G, and mixed numerology in 5G.

ML-BASED ARCHITECTURE

The aiOS architecture is based on the open next-generation RAN envisaged by O-RAN [6], which seeks to extend the control/user plane decoupling with AI-empowered radio control, hardware and software openness, and virtualization. With this in mind, below we describe the main components of aiOS, as depicted in Fig. 2, and how the challenges discussed in the previous section are tackled.

Multi-RAT Protocol Stack: This includes various network elements based on Wi-Fi, 4G, and 5G that are able of distributing the load across several (possibly disaggregated) nodes.

Challenges addressed: This diversified deployment tackles the N-4 (*support for multiple RATs*) and ML-2 (*heterogeneous data sources*) challenges, allowing per-user and per-node RAN telemetry to be pushed to the near-real-time RIC through the standard E2 interface.

Near-Real-Time RIC: This implements control functions on the order of 10 ms–1 s, and comprises two main elements:

- **Applications Layer.** This is a software layer comprising network applications called *xApps* in O-RAN terminology. The *xApps* control one or several radio management operations by exchanging data with the infrastructure devices over the E2 interface. Examples of these applications are mobility management and frame size selection.

To be successful, ML-based network management solutions need to pull data not only from the network stack but also from the applications and services running on the network so that the impact of such applications (e.g., tele-medicine) on network configuration or optimization can be assessed.

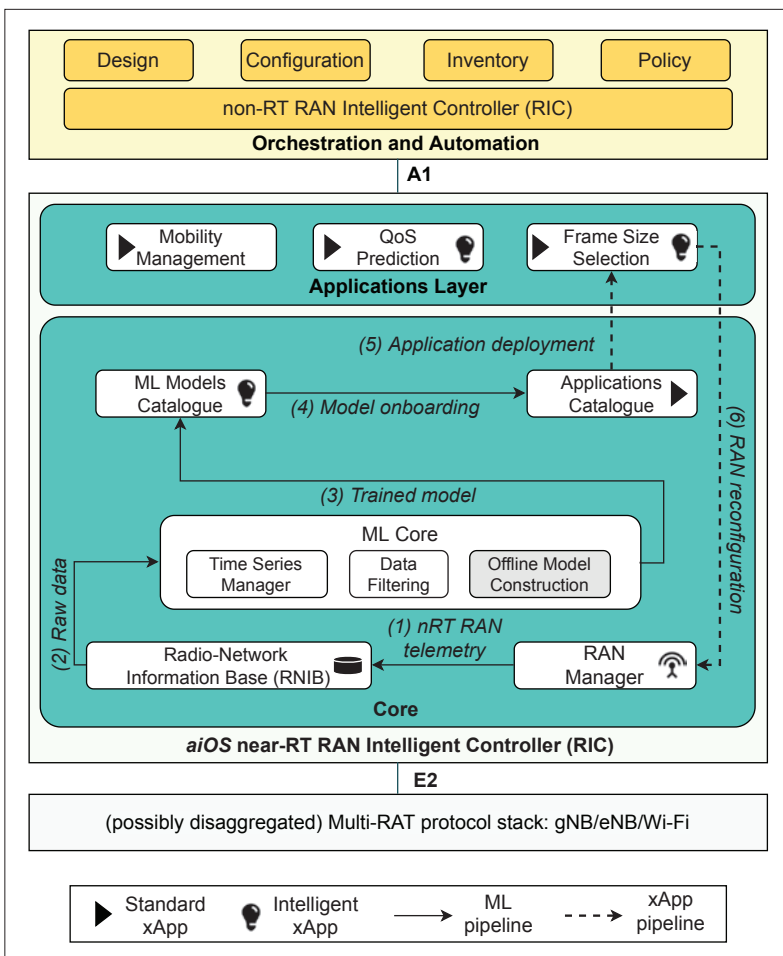


FIGURE 2. aiOS reference system architecture based on O-RAN standards.

- Core.** This embeds the network intelligence and per-user near-real-time functions. The information flow is divided into six steps. In *Step 1*, the RAN manager collects near-real-time RAN telemetry by capturing the state of the underlying infrastructure and storing it in the radio-network information base (RNIB). In *Step 2*, the ML core is fed with these new raw data. The time series manager is the first entry point, and is responsible for processing the data from the RNIB, and for providing the ML core with a merged and clean dataset. This processing combines possibly non-synchronized data sources (e.g., from various APs) and weights historical data and new data. After that, the ML core produces a filtered dataset that can be used offline by any ML framework to build a model. In *Step 3*, the ML core stores the new model in the ML models catalogue, while in *Step 4* it onboardes the model in the applications catalogue as an xApp. In *Step 5*, the xApp is deployed in the applications layer. Finally, in *Step 6* the output of the application is forwarded to the RAN manager for RAN reconfiguration. Note that although Steps 1 to 4 correspond to the ML pipeline, Steps 5 and 6 are identical for the deployment and reconfiguration of standard and intelligent xApps. This is highlighted in Fig. 2 by dashed lines.

Challenges addressed: The modular design eliminates the need to perform the entire ML pipeline on the network nodes, reduces *computational complexity* (challenge N-3), and facilitates the support of future ML toolboxes (challenge ML-3). Furthermore, it enables *ease of use* (challenge N-1) and *interpretability* for networking experts (challenge N-2), and allows them to rectify the decisions made by ML engines.

Orchestration and Automation: This builds on cloud components acting as a single distributed system (from edge to data centers). This layer is orchestrated by the non-RT RIC and implements above-1-s functions such as configuration, inventory, and policy management. The standard A1 interface allows the specification of individual control policies on the RAN and their conveyance for runtime execution.

Challenges addressed: Besides the deployment of network policies and reconfiguration, this layer allows ML models to be trained at the edge or in data centers. Consequently, *training time* (challenge ML-4) is greatly reduced by leveraging powerful cloud sites. Moreover, clean data from other sources (e.g., operational networks) can be inserted, enabling *dataset availability* (challenge ML-1).

ADAPTIVE FRAME AGGREGATION USING aiOS

In this section we describe an xApp designed to assess the viability of aiOS in a network management task of practical relevance, namely data-driven adaptive frame aggregation in IEEE 802.11-based SD-WLANs. The xApp has been implemented and validated on a Wi-Fi testbed.

DESCRIPTION

The channel access of 802.11 incurs a high overhead given its contention-based nature. Figure 3 depicts an example of this issue when transmitting three packets, where the overhead accounts for 60 percent of the airtime. To overcome this, the standard incorporates frame aggregation by employing two techniques: the aggregated MAC service data unit (A-MSDU) and the aggregated MAC protocol data unit (A-MPDU). Figure 3 shows the huge chunk of airtime saved by A-MSDU aggregation. However, both techniques are too rigid as they define a fixed frame size. Figure 4 shows how, under different channel conditions determined by different distances from the AP and modulation and coding schemes (MCSs), diverse frame sizes provide the best goodput for each user, which is where aiOS's data-driven capabilities prove their worth by adapting the frame length on a per-user basis.

On the basis of the in-network information made available by aiOS, we designed offline ML models to adapt the per-user frame length to maximize network goodput. We took as reference the Random Forest Regressor (RFR) and M5P ML models, which were deployed as xApps, as shown in Fig. 2. To ensure standard compatibility, the frame size was only adapted in the downlink direction. The deployment considered a Wi-Fi network comprising one AP, N static stations connected in downlink, and M static stations in uplink representing background traffic. The AP is based on a PCEngines ALIX 2D board mounting an Atheros AR9220 Wi-Fi interface and running

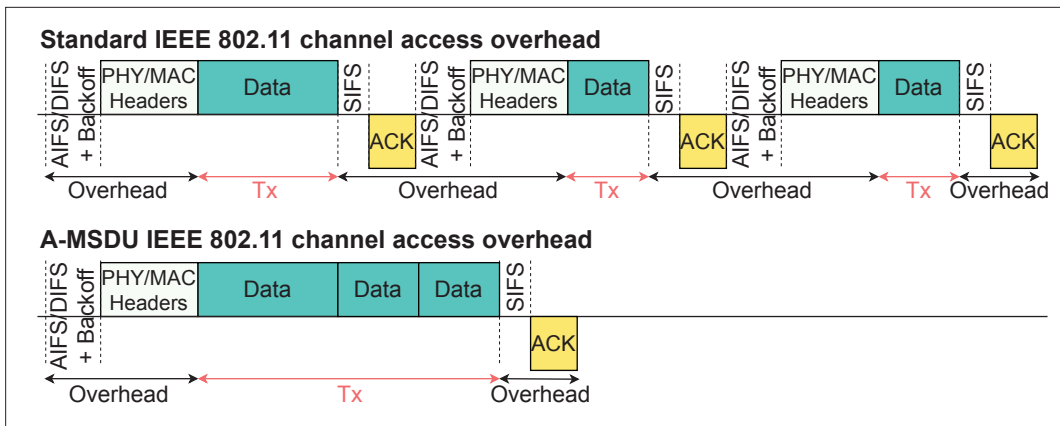


FIGURE 3. Overhead savings in IEEE 802.11 channel access enabled by frame aggregation.

OpenWRT 18.06.04. The AP, which was configured on channel 36 and free from external interference, had a fixed location, while the stations were placed randomly. For the training phase, the aggregated downlink bit rate was 20 Mb/s, and the uplink bit rate was 1 Mb/s. The frame size was set to 200 B, with a maximum aggregation length of 3839 B. From this dataset, the variables selected for the models' construction were the channel utilization, and specific information of the MCS, for example, success probability, attempted bytes in the last 100 ms, and expected throughput in perfect channel conditions. Please refer to [7] for more details about the training.

PERFORMANCE EVALUATION

The evaluation considered two scenarios covering homogeneous (*Scenario 1*) and heterogeneous (*Scenario 2*) channel conditions. In Scenario 1, the stations were placed at 20/30 m from the AP, while in Scenario 2 the distance was 20/50 m. The evaluation took the basic frame transmission as baseline and compared the goodput improvement achieved by the standard A-MSDU aggregation, the M5P model, and the RFR model for an increasing number of stations in downlink.

Figure 5 shows that the basic transmission always offers the lowest performance because of the high overhead induced by small payloads. As depicted in Fig. 5a, in Scenario 1 the ML models outperform A-MSDU aggregation due to their continuous size adaptation to the network status. In contrast, the use of a fixed size results in a greater number of errors in the transmissions. Moreover, the improvement achieved by the ML models increases with the number of stations due to their greater efficiency in dealing with channel congestion. Figure 5b, which corresponds to Scenario 2, compares the same mechanisms under heterogeneous channel conditions. Similarly to the previous scenario, the M5P and RFR models outperform the standard mechanisms. In this case, the goodput improvement is comparatively lower for three and four stations because multiple stations are using a lower MCS given the greater distance, thus increasing the channel utilization. In general, RFR provides the best results as it is more suitable for problems with high variance and high bias such as those found in wireless networks, in which channel conditions greatly vary.

The results demonstrate that the relative goodput improvement with respect to the stan-

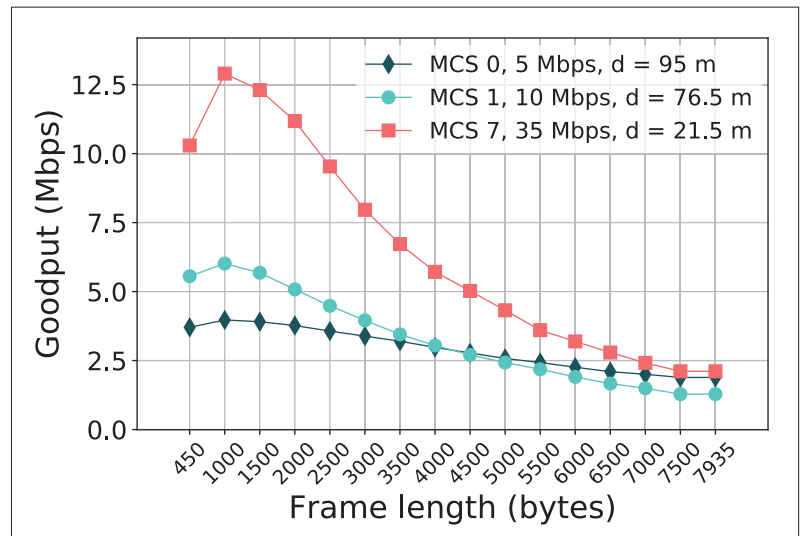


FIGURE 4. Goodput vs. frame length for various distances and MCSs.

dard mechanisms is higher in heterogeneous conditions. Note that although the improvement is higher in the homogeneous scenario, the heterogeneous case is much more significant as it more closely resembles a real-life environment. This fact proves the relevance of AI-based network management when taking as input real-time in-network information such as that provided by aiOS.

DISCUSSION AND FUTURE CHALLENGES

AI-enabled SDN is a promising paradigm for future wireless and mobile networks. In this article, we introduce aiOS, which is an AI-native platform for control and management policies in SD-WLANs. aiOS, whose design pillars are high-level abstractions, unified data models, and well-defined interfaces, is well suited to serve the needs of different verticals, including residential, enterprise, and industrial deployments. In particular, we have shown how aiOS can be used to dynamically adapt the frame aggregation length in IEEE 802-11-based SD-WLANs.

Currently, aiOS supports only offline and centralized model construction. However, we are already working on larger volumes of data collected from multiple sites to extend aiOS with online and federated learning capabilities. We expect

Currently, aiOS supports only offline and centralized model construction. However, we are already working on larger volumes of data collected from multiple sites to extend aiOS with online and federated learning capabilities. We expect this to allow aiOS to react more promptly to changes and provide more accurate decisions in a timely manner, as the models will be updated constantly with new batches of data.




FIGURE 5. Goodput improvement for an increasing number of stations: a) homogeneous channel conditions; b) heterogeneous channel conditions.

this to allow aiOS to react more promptly to changes and provide more accurate decisions in a timely manner, as the models will be updated constantly with new batches of data.

ACKNOWLEDGMENTS

This work has been performed in the framework of the European Union's Horizon 2020 project 5GZORRO co-funded by the EU under grant agreement No. 871533.

REFERENCES

- [1] H. Kim and N. Feamster, "Improving Network Management with Software Defined Networking," *IEEE Commun. Mag.*, vol. 51, no. 2, Feb. 2013, pp. 114–19.
- [2] "Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2016–2021," tech. rep., 2017.
- [3] B. Dezfouli et al., "A Review of Software-Defined WLANs: Architectures and Central Control Mechanisms," *IEEE Commun. Surveys & Tutorials*, vol. 21, no. 1, 2019, pp. 431–63.
- [4] C. Zhang, P. Patras, and H. Haddadi, "Deep Learning in Mobile and Wireless Networking: A Survey," *IEEE Commun. Surveys & Tutorials*, vol. 21, no. 3, 2019, pp. 2224–87.
- [5] ITU-T Rec. Y.3172, "Architectural Framework for Machine Learning in Future Networks Including IMT-2020," June 2019.

- [6] O-RAN Alliance, "O-RAN Architecture Description v1.0," Feb. 2020.
- [7] E. Coronado et al., "aiOS: An Intelligence Layer for SD-WLANs," *Proc. IEEE NOMS*, Budapest, Hungary, 2020.
- [8] F. Wilhelmi et al., "A Flexible Machine-Learning-Aware Architecture for Future WLANs," *IEEE Commun. Mag.*, vol. 58, no. 3, Mar. 2020, pp. 25–31.
- [9] K. B. Letaief et al., "The Roadmap to 6G: AI Empowered Wireless Networks," *IEEE Commun. Mag.*, vol. 57, no. 8, Aug. 2019, pp. 84–90.
- [10] K. Davaslioglu et al., "DeepWiFi: Cognitive WiFi with Deep Learning," *IEEE Trans. Mobile Comp.*, 2019.
- [11] P. Gawłowicz and A. Zubow, "Ns-3 Meets OpenAI Gym: The Playground for Machine Learning in Networking Research," *Proc. IEEE MSWIM*, Miami Beach, FL, 2019.
- [12] F. Wilhelmi et al., "Usage of Network Simulators in Machine-Learning-Assisted 5G/6G Networks," arXiv preprint arXiv:2005.08281, 2020.
- [13] M. Wollschlaeger, T. Sauter, and J. Jasperneite, "The Future of Industrial Communication: Automation Networks in the Era of the Internet of Things and Industry 4.0," *IEEE Trans. Ind. Electron.*, vol. 11, no. 1, 2017, pp. 17–27.
- [14] D. López-Pérez et al., "IEEE 802.11be Extremely High Throughput: The Next Generation of Wi-Fi Technology beyond 802.11ax," *IEEE Commun. Mag.*, vol. 57, no. 9, Sep. 2019, pp. 113–19.
- [15] N. Foster et al., "Languages for Software-Defined Networks," *IEEE Commun. Mag.*, vol. 51, no. 2, Feb. 2013, pp. 128–34.

BIOGRAPHIES

ESTEFANÍA CORONADO (estefania.coronado@i2cat.net) is a senior researcher at Fundació i2CAT, Spain. She obtained her Ph.D. in SD-WLANs from the University of Castilla-La Mancha, Spain, where she also obtained her M.Sc. degrees in computer engineering and advanced computer technologies in 2014 and 2015. She has published around 25 papers in international journals and conferences. Her research interests include AI-driven network automation, wireless/mobile networks, network slicing, and SDN/NFV.

SUZAN BAYHAN (s.bayhan@utwente.nl) is an assistant professor at the University of Twente and an adjunct professor (docent) at the University of Helsinki. She earned her Ph.D. in computer engineering in 2012 from Bogazici University. She received the best paper awards at ACM ICN '15 and IEEE WoWMoM '20, and best demo award at IEEE INFOCOM '20. Her current research interests include spectrum sharing, coexistence of wireless networks, WiFi and LTE resource management, and edge computing.

ABIN THOMAS (athomas@fbk.eu) is currently pursuing a Ph.D. at the University of Trento. He received a Bachelor's degree in electronics and telecommunications engineering from Goa University, India, in 2014. He received an M.Sc. degree from EIT Digital (KTH and University of Trento) in 2019. His research interests include SDN, NFV, and ML for wireless and mobile networks.

ROBERTO RIGGIO [SM] (roberto.riggio@i2cat.net, roberto.riggio@ri.se) is a senior researcher in the Connected Intelligence Group at RISE AB, Stockholm, Sweden. He received his Ph.D. from the University of Trento, Italy, after which he was a postdoctoral researcher at the University of Florida, researcher/chief scientist at CREATE-NET, Trento, a head of unit at FBK in Trento, and a senior 5G researcher at the i2CAT Foundation in Barcelona, Spain. He has published more than 130 papers in internationally refereed journals and conferences. He has received several awards including the IEEE INFOCOM Best Demo Award (2013 and 2019) and the IEEE CNSM Best Paper Award (2015). He is a member of ACM.