



LJMU Research Online

Bouhafs, F, MacKay, M, Raschella, A, Shi, Q, den Hartog, F, Saldana, J, Munilla, R, Ruiz-Mas, J, Fernández-Navajas, J, Almodovar, J and van Adrichem, N

Wi-5: A Programming Architecture for Unlicensed Frequency Bands

<http://researchonline.ljmu.ac.uk/id/eprint/9608/>

Article

Citation (please note it is advisable to refer to the publisher's version if you intend to cite from this work)

Bouhafs, F, MacKay, M, Raschella, A, Shi, Q, den Hartog, F, Saldana, J, Munilla, R, Ruiz-Mas, J, Fernández-Navajas, J, Almodovar, J and van Adrichem, N (2018) Wi-5: A Programming Architecture for Unlicensed Frequency Bands. IEEE Communications Magazine. 56 (12). pp. 178-185.

LJMU has developed [LJMU Research Online](#) for users to access the research output of the University more effectively. Copyright © and Moral Rights for the papers on this site are retained by the individual authors and/or other copyright owners. Users may download and/or print one copy of any article(s) in LJMU Research Online to facilitate their private study or for non-commercial research. You may not engage in further distribution of the material or use it for any profit-making activities or any commercial gain.

The version presented here may differ from the published version or from the version of the record. Please see the repository URL above for details on accessing the published version and note that access may require a subscription.

For more information please contact researchonline@ljmu.ac.uk

<http://researchonline.ljmu.ac.uk/>

Wi-5: A Programming Architecture for Unlicensed Frequency Bands

Faycal Bouhafs, Michael Mackay, Alessandro Raschellà, Qi Shi, Frank den Hartog, Jose Saldana, Rubén Munilla, José Ruiz-Mas, Julián Fernández-Navajas, Jose Almodovar, and Niels van Adrichem

The authors advocate the need for a spectrum programming architecture that goes beyond managing networking resources and operation. Whereas SDN and SDWN architectures add programmability to the management of data traffic by defining a data plane, we here define the spectrum plane to introduce spectrum programmability.

ABSTRACT

We present Wi-5, a spectrum programming architecture for radio resource management in unlicensed frequency bands. It introduces a spectrum control plane that offers fine grained allocation of radio resources, flexible configuration of radio and wireless networking parameters, and continuous monitoring of the wireless network status. These features, along with the centralized nature of this architecture, can effectively address spectrum congestion which often occurs in unlicensed frequency bands. To demonstrate Wi-5's capabilities, we show results obtained from emulating various use case scenarios on our open source proof-of-concept.

INTRODUCTION

In recent decades, we have witnessed the transformation of WiFi from a secondary enterprise communication technology to the most prevalent option for accessing the Internet worldwide. Today, WiFi networks represent a major communication infrastructure in public spaces and residential buildings. Moreover, other wireless technologies, such as ZigBee and Bluetooth, are increasingly used to connect devices with each other and the Internet. These technologies commonly use so-called unlicensed spectrum: in most cases, individuals do not need to obtain a special licence to operate their wireless devices. Unfortunately, managing these wireless networks is often complex, as they need to operate in an increasingly hostile and uncontrolled spectral environment, and the number of devices and the amount of data traffic continues to grow.

Staying with WiFi as an example, conventional WiFi network management solutions focus on providing robust and fine-grained functionality to operators responsible for large networks such as those deployed in university campuses and airports. These products are usually vendor proprietary and require the operators to deploy only equipment from the same solution provider. With these management platforms, operators can monitor the number of devices connected to the network, the data traffic flowing through them, and radio-related performance parameters such as signal-to-noise ratio. Typically, they also enable the remote configuration of the communication

channel and the transmission power of each WiFi access point (AP).

More recently, new products developed by less established companies such as Open Mesh and Tanaza have emerged, offering an open southbound management interface (between management platform and devices) without requiring vendor-specific equipment. They target large and small networks deployed in shops, small offices, and large houses, and offer easy-to-use functionality to non-professional network managers trying to optimize the performance of their networks. Although some of these solutions include functionality for cloud management, secure access, and Internet of Things applications, they typically provide fewer operational management capabilities than the professional vendor-specific solutions, and do not provide an open northbound interface to, for instance, third-party management applications. As WiFi networks are relatively dynamic, and all have their own characteristics, these solutions therefore lack the flexibility to manage such networks optimally.

In the following section we illustrate the issue with two of the most challenging use cases. We then describe our spectrum programming approach, how it relates to Software-Defined Networking (SDN), and how it is applied in the European Horizon 2020 "Wi-5" project (www.wi5.eu).

USE CASES

SPECTRUM CONGESTION

In 2014, 54 percent of the world's population was urban, a number that is expected to rise to 66 percent by 2050 [1]. In addition, most cities are executing policies of urban consolidation. As a result, more and more people live and work in increasingly crowded built-up areas. In such areas, the lack of appropriate network management has resulted in the dense and generally uncoordinated deployment of WiFi networks that we see today. This is leading to various unwanted effects, including heavy traffic congestion and over-congestion [2, 3].

Congestion happens when an AP perceives all channels as being occupied with traffic virtually all the time. WiFi's MAC protocol still allows for such APs to transmit, due to its random-back-off mechanism, but ultimately the total capacity will be shared with the other APs, and the achievable throughput of those APs will go down.

Faycal Bouhafs, Michael Mackay, Alessandro Raschellà, and Qi Shi are with Liverpool John Moores University; Frank den Hartog is with the University of New South Wales; Jose Saldana is with the University of Zaragoza; Rubén Munilla, José Ruiz-Mas, and Julián Fernández-Navajas are with the University of Zaragoza; Jose Almodovar and Niels van Adrichem are with TNO.




Figure 1. Three approaches to WiFi network management and control: a) AP-based approach; b) centralized proprietary approach; c) spectrum programming approach.

Ozyagci *et al.* [2] showed that a system consisting of a continuously increasing number of APs in a confined space will ultimately end up in a so-called over-congested state: an increasingly larger portion of the total available capacity is used for control traffic trying to mitigate the congestion and resulting packet collisions. The result is actual depletion of the common spectral resource for all users. In [3] we showed that over-congestion is already happening in apartment blocks today.

CONNECTIVITY AND COVERAGE IN LARGE HOUSES

Indoor WiFi networks are often found to be a bottleneck for bandwidth-hungry applications in large houses that have a fast Internet connection. In such environments, WiFi network operators (i.e., anyone responsible for operating one or more APs) deploy multiple APs to guarantee full coverage as well as the best possible connectivity to the wireless users. Traffic within these networks is often characterized by differentiated Quality of Service (QoS) demands, as each wireless user might be running different types of applications. Therefore, operators need new ways to connect each wireless user to the best AP to satisfy their

Traffic within these networks is often characterized by differentiated QoS demands, as each wireless user might be running different types of applications. Therefore, operators need new ways to connect each wireless user to the best AP to satisfy their QoS demands. This usually requires a soft handover in which the user's device is steered from its current AP to another AP.

QoS demands. This usually requires a soft handover in which the user's device is steered from its current AP to another AP.

SPECTRUM PROGRAMMING

REQUIREMENTS

Efficient management of wireless networks using unlicensed frequency bands typically requires:

- *The ability to infer and keep track of the status of the wireless network:* For a network performance optimization algorithm to be effective, it needs to obtain periodic updates on the status of the network. Such updates will allow the algorithm to detect anomalies (i.e., quality degradation) and react to them. This includes the ability to measure the level of interference among APs, which allows the algorithm to determine the best configuration of these APs' radio and networking parameters.
- *The ability to configure radio and networking parameters of all APs:* To mitigate spectrum congestion, the management system must be able to configure the APs' radio and networking parameters such as the transmis-

Wi-5 introduces a spectrum control plane that offers fine grained allocation of radio resources, flexible configuration of radio and wireless networking parameters, and continuous monitoring of the wireless network status.




Figure 2. Wi-5 architecture and comparison with Software-Defined Wireless Networking: a) a high level description of Wi-5 Architecture; b) comparison between SDWN and Wi-5 with their respective control paths.

sion channel, the transmit power, the frame aggregation, and the access control list.

In theory, a system operator may choose to run a network performance optimization algorithm off-line (i.e., not in real-time), and the outcome of the algorithm can then be configured manually. However, the status of WiFi networks frequently changes, as clients (i.e., devices that can wirelessly connect with APs) often join and leave randomly, and their QoS requirements also vary over time. This requires the operator to intervene regularly to reduce spectrum congestion and guarantee connectivity. Thus, having the algorithm running in real-time and the spectrum utilization being dynamically programmable is desirable.

CURRENT APPROACHES VS SPECTRUM PROGRAMMING

Current approaches to WiFi network management and control can be divided into two different categories: distributed and centralized. They are schematically depicted in Figs. 1a and 1b. The distributed approach, or AP-based approach, is the basis of most IEEE 802.11 and WiFi Alliance standards today. APs provide a best effort service to coordinate spectrum use among each other by exchanging coordination messages. Optimization algorithms are included in the MAC layer (covered by IEEE 802.11) or in additional embedded software (WiFi Alliance). However, the APs' lack of processing capabilities and limited view of the total network makes it difficult to implement an efficient management approach.

Current platforms for centralized management (Fig. 1b) are either proprietary, or unable to provide the capabilities identified above. They lack the flexibility to deal with newly added APs from different vendors and/or to deploy novel management routines as needed. These limitations call for a novel management architecture that offers open programming interfaces: a southbound interface that makes it possible to manipulate radio management capabilities of APs flexibly and optimally,

and a northbound interface that exposes primitives allowing the administrator to develop their own RRM solutions.

Interestingly, this paradigm is similar to what SDN offers to wired network engineers today, as SDN makes another network resource programmable, namely routing. WiFi networks have already benefited from the features that some SDN-based applications offer. These applications are, however, limited to traffic engineering and orchestration of networking resources, as current SDN architectures do not support fundamental IEEE 802.11 functions such as managing Service Set Identifiers (SSIDs), network associations, and so on [4].

Software-Defined Wireless Networking (SDWN) provides an extension of SDN to support fast and flexible large-scale management of wireless networks [5]. The Odin framework [6] offers such an extension, as it enables orchestration of large enterprise WiFi networks. With Odin, each wireless client is associated to a light virtual access point (LVAP) that runs on the physical AP to facilitate the management of wireless connections and the clients' mobility. The OpenSDWN architecture [7] extends Odin, enabling per-client virtual APs and per-client virtual middleboxes. Other architectures such as Empower [8] apply SDN to the management of heterogeneous wireless networks with different radio access technologies, a key feature in 5G networks. The authors in [9] proposed an architecture based on SDWN whereby the control plane is enhanced to support basic radio parameter configuration. However, the authors did not acknowledge the need for making these parameters programmable to include the desired management flexibility, and the proposed SDN model was justified with simulations only.

We advocate the need for a spectrum programming architecture that goes beyond managing networking resources and operation, as is currently the case with SDN and SDWN. Besides

an open southbound application programming interface (API) to the APs radio primitives, such as spectrum programming architecture should also offer an open northbound API that allows network managers to deploy radio resource management (RRM) algorithms and other management tools, possibly provided by third party developers (Fig. 1c).

WI-5 ARCHITECTURE

Our architecture implements spectrum programming for networks operating on unlicensed frequency bands. A high-level overview of the architecture is illustrated in Fig. 2a. APs expose the main radio primitives to the Wi-5 central controller. The Wi-5 controller in turn exposes a northbound API, enabling various applications to program the wireless network as desired. Applications can be added and removed as desired, and algorithms may be added that orchestrate various applications.

Whereas SDN and SDWN architectures add programmability to the management of data traffic by defining a data plane, we here define the *spectrum plane* to introduce spectrum programmability (Fig. 2b). Therefore, the four planes in our architecture are as follows.

Infrastructure Plane: This plane is shared by Wi-5 and all SD(W)N solutions. Although our architecture focuses on WiFi, it also considers the use of other techniques in the unlicensed and licensed frequency bands, such as LTE, to alleviate spectrum congestion.

Spectrum Plane: This new plane enhances the operational capabilities of APs by defining new monitoring and configuration primitives, and making the APs programmable, thus enabling fine-grained spectrum allocation and management. This plane sits alongside the data plane that is part of traditional SDN and SDWN architectures, and where data traffic management policies reside. The spectrum plane provides an additional interface to the control plane.

Control Plane: This plane holds the SD(W)N and Wi-5 controllers. The SD(W)N controllers control the data plane (the blue dashed line in Fig. 2b), and the Wi-5 controllers control the spectrum plane (the red dashed line in Fig. 2b). The control plane provides an API to the application plane.

Application Plane: The application plane (not shown in Fig. 2b) is also common among Wi-5 and the SD(W)N architectures. In Wi-5, this plane will implement and run applications that execute management policies to achieve better connectivity and spectrum efficiency while remaining aware of potential issues with the QoS of the wireless users' services.

In a multi-domain context, addressing spectrum congestion and connectivity issues optimally requires cooperation among different AP operators. For this, the proposed architecture may be enhanced by an additional plane above the application plane through which operators interact and negotiate, and execute cooperation agreements and policies.

SPECTRUM PLANE FUNCTIONALITY

The spectrum plane consists of the following set of functions making APs programmable via the controller.

Monitoring: This function measures the performance of the wireless network, including the interference level and the load in each channel. It also keeps track of the number of clients associated with each AP, the downstream traffic, and the nature and QoS requirements of each flow.

DFS and TPC: This function implements Dynamic Frequency Selection (DFS) and Transmit Power Control (TPC) as defined in IEEE 802.11h. It is worth noting that the IEEE 802.11h standard does not specify a particular way of implementing DFS and TPC. Including these functions in a spectrum plane makes them accessible as primitives by which applications can optimize spectrum utilization over a complete network.

Frame Aggregation: This function enables APs to group small frames into larger ones, as defined in IEEE 802.11n, to reduce communication overhead, resulting in spectrum saving.

APPLICATIONS

To validate the proposed architecture, we have so far designed the following applications as part of the application plane.

Horizontal Handover: This application enables clients to perform a fast soft handover to another AP if required, apparently seamless to the user. It always connects the client to the AP that provides the best connectivity, without interruption of service.

Vertical Handover: This application moves wireless clients to networks operating in other frequencies (e.g. 4G) to reduce spectrum congestion, while maintaining the QoS requirements of the users' services as much as possible. This can only be applied to clients with dual networking capabilities, and is the opposite of what has been standardized in 3GPP as ANDSF (Access Network Discovery and Selection Function). ANDSF is a mobile core network functionality that enables a (4G) network operator to offload a device from the 4G network to a WiFi network.

Radio Parameter Configuration: Based on the algorithm presented in [10], this application finds an optimal radio parameter configuration to achieve the dual objective of minimizing the interference level in congested environments while simultaneously satisfying the QoS requirements of the wireless users' services. This application is based on a flow-centric radio management approach where the spectrum allocation is based on the QoS requirements of each traffic flow.

Smart AP Selection: This application can dynamically connect clients to other APs when spectrum congestion is detected or the initial APs' performance drops. The aim is to maintain the QoS requirements of the users' services as much as possible.

Frame Aggregation Adjustment: This application dynamically adjusts the aggregation parameters in APs to save air time, especially during periods of congestion. The decision to modify frame aggregation is based on the nature of the flows present on each AP [11], such as their real-time requirements.

It needs to be stressed that any of the functions provided by the applications above can also be implemented statically and in hardware or embedded software, and many of them already have (e.g. ANDSF), often with limited success. In

In a multi-domain context, addressing spectrum congestion and connectivity issues optimally requires cooperation among different AP operators. For this, the proposed architecture may be enhanced by an additional plane above the application plane through which operators interact and negotiate, and execute cooperation agreements and policies.




Figure 3. Our implementation of the Wi-5 architecture.




Figure 4. Client throughput before and after manually introducing co-channel interference at $t = 115$ s. Wi-5’s Radio Parameter Configuration application automatically restores the network at $t = 138$ s.

Wi-5, however, these functions are implemented as mere software applications that run on a platform, and can be tailored and orchestrated at any time to any need.

IMPLEMENTATION AND EXPERIMENTAL RESULTS

An open source proof-of-concept (<https://github.com/Wi5>) of our architecture has been implemented using the Odin framework, described in detail in [12]. It consists of an Odin controller, which is an extension of Floodlight OpenFlow, and an Odin Agent, built using Click Modular Router, that runs on each AP in parallel with the OpenFlow Agent. The Odin Agent interacts with the controller via the Odin Master which presents an extension of the Floodlight southbound API to support network management operations. The controller also extends the Floodlight northbound API to allow the implementation of WiFi network management applications. The controller runs on a Raspberry Pi, and the agent runs on commodity APs typically used in small private networks.

Our proof-of-concept extends Odin by further enhancing the controller’s southbound API to support radio configuration and monitoring. This extension is represented by the Wi-5 Master, located at the Wi-5 controller, and the Wi-5 Agent, located at the AP, as illustrated in Fig. 3. We also extended the controller’s northbound API to expose the primitives that allow the programming of the spectrum plane. In addition, our proof of concept improves SSID and LVAP management by taking radio configuration parameters into account. No changes or additions are required in the client, which runs standard WiFi.

To validate our architecture, we implemented the top three applications described earlier. We have deliberately limited our experiments to networks of relatively small scale, as they provide easier to understand proofs that our concepts work. Besides, for these applications we can show that, even in relatively small-scale networks, our solution already provides significant improvements. Large-scale field pilots, including various other applications, are currently underway.

RADIO PARAMETER CONFIGURATION

This application has been implemented to address the spectrum congestion problem in dense and uncoordinated deployments of WiFi networks, as evidenced in [3]. It relies on the monitoring function to measure the network-wide interference impact, which is a measure of the total interference caused by channel overlap [10].

For this test, we use two Wi-5 coordinated APs (AP1 and AP2), connecting one client to AP1 and another client to AP2. The APs are operating at different, non-overlapping channels. The results are shown in Fig. 4. At $t = 60$ s, the client connected to AP1 starts a file download with TCP, and the client connected to AP2 starts sending a UDP flow in the uplink (black line). We then create co-channel interference by forcing AP2 to switch to the same channel as AP1 (channel 11) at $t = 115$ s. This interference is detected by the monitoring func-

tion which measures the interference impact (yellow line), and coincides with a sharp drop in the throughput of the client connected to AP2.

At $t = 138$ s the client's throughput starts increasing while the value of the interference impact drops. This is due to the Radio Parameter Configuration application which, based on the channel assignment algorithm as described in [10], then switches AP2 to channel 6, thus relieving the interference. Note that the algorithm has been designed to address congestion in much more densely deployed environments than tested here. However, since the purpose of this experiment is to assess the performance of the proposed architecture rather than the algorithm, we opted for a small-scale spectrum congestion experiment that is easier to execute and understand. Also note that we only implemented DFS, and left TPC for future work. During the experiment we observed a burst transmission in the throughput, just after AP2 switches to channel 6, which lasted for one second (not shown in the figure). This can be addressed to packets stored in the wireless interface buffer during the period of low throughput.

HORIZONTAL HANDOVER

Fast and seamless handover is typically required in situations where the network consists of many APs covering a large area. Through the LVAP abstraction mechanism, our architecture can implement handover solutions that enable clients to move among APs faster and smoother than traditional handover solutions, and without interruption of service. To allow proper channel planning, the handovers must happen between APs operating in different channels. This is achieved by using beacon signals containing the Channel Switch Announcement element. In addition, the APs are extended with an auxiliary wireless interface that allows scanning in other channels than the one in which the data plane is operating. The ability to perform inter-channel handover goes beyond the original Odin implementation [12].

A horizontal handover is triggered when the monitoring functionality in the spectrum plane detects that the signal received by the client drops below a predefined threshold. The application first instructs all APs in the neighborhood to use their auxiliary interfaces to scan for the client MAC address during a short period of time, and each AP reports back the signal strength received from the client. Accordingly, the application associates the client with the AP from which it has received the strongest signal.

The main performance metric of a handover is the seamlessness with which clients can move from one AP to another. This is observed by measuring variations in throughput at the client. We therefore extended the testbed with a third AP. A client is moving along a corridor with varying proximities to the three APs in different channels, while receiving a 100 packets per second flow of 80-byte UDP packets. Figure 5 shows the received throughput (black line) with a) standard IEEE 802.11 connection and disconnection, and b) our Horizontal Handover application. The colored lines represent the Received Signal Strength Indicator (RSSI) as observed by the individual APs.

Figure 5a shows that the client first connects to the nearest AP, that is, AP1 (blue line), and stays




Figure 5. Client throughput (black) and the APs' signal strengths (RSSI) as observed with two different mechanisms for horizontal handover: a) as standardized by IEEE 802.11; b) using the Wi-5 Horizontal Handover application.

connected to it despite moving away from it, and being closer to AP2 (green) after 25s. A connection to AP3 (red) at $t = 62$ s, is initiated after the client lost its connection with AP1 (around $t = 62$ s). At this point the service has been seriously interrupted, as proven by the sharp drop in the measured throughput. This could have been avoided if AP2 (green) were to be used for some period of time. This issue is widely known as the "sticky client" problem [13], and our architecture is solving it [12]. In Fig. 5b, we can observe that at around $t = 20$ s the signal received at AP1 drops below the predefined threshold (amber line). This is notified to the controller, which then instructs AP2 and AP3 to scan for that client for 4 seconds. As the signal level reported by AP2 (green) now appears to be highest, the controller moves the LVAP from AP1 to AP2. In practice, the client hardly notices this, as the drop in throughput is only ~10 percent over a few seconds. A similar behavior is shown at $t = 40$ s, when the client hands over to AP3.

VERTICAL HANDOVER

This application executes a vertical handover by proactively de-authenticating clients connected to a given AP and the removal of the corresponding LVAP. An actual handover to a third party net-




Figure 6. Throughput of client 1 (black), client 2 (red), and client 3 (green) where client 1 is handed over to a non-WiFi network between $t = 90$ s and $t = 150$ s.

work such as 4G is then left to the client. To avoid the client to handover back to the WiFi network immediately after the initial de-authentication, its MAC address is added to a blacklist preventing re-authentication. When network resources permit, the controller can then reintroduce clients by removing them from the blacklist.

To test the effectiveness of this application with regard to spectrum congestion alleviation, we conducted the following experiment in our testbed. AP1 and AP2 are in each other's direct vicinity, and both set on channel 11 using IEEE 802.11g in the 2.4 GHz band. AP1 runs a Wi-5 Agent and is connected to an *iperf* server and to the Wi-5 controller running the Vertical Handover application. The Click Modular Router, which is part of the Wi-5 Agent, has an internal interface *ap* connected to the virtual switch (Open vSwitch) on AP1. Interface *ap* operates at a maximum of 10 Mb/s. AP2 is not part of the Wi-5 network.

The clients are Linux laptops supporting CUBIC TCP stacks and running an *iperf* traffic generator. MAC Address Randomization was disabled to guarantee quick reconnects. Two clients (1 and 2) are connected to AP1 and transmit TCP traffic. The third client is connected to AP2 and transmits UDP traffic at 4 Mb/s. The actual throughputs obtained with *Wireshark* and an *AirPcap* WiFi traffic analyzer are shown in Fig. 6, for client 1 (black), client 2 (red), and client 3 (green), for the following sequence of actions. At $t = 0$ s, client 1 is connected to AP1. At $t = 30$ s, client 2 is also connected to AP1, and TCP equally divides the available capacity over the two clients. At $t = 60$ s, client 3 is connected to AP2. The performance of both clients 1 and 2 now deteriorates because of spectral congestion. At $t = 90$ s, the Wi-5 controller de-authenticates and blacklists client 1, and the performance of client 2 improves accordingly. At $t = 120$ s, client 3 stops emitting, resulting in a further improvement of the performance of client 2. At $t = 150$ s, the Wi-5 controller allows client 1 back into the network, and clients 1 and 2 again share the available capacity.

The results show that the detrimental effect of client 3 on the throughputs of clients 1 and 2 can be effectively mitigated by handing over client 1 to

another network, of course assuming that client 1 can obtain a better performance on that other network than it is currently experiencing. In a real-life implementation of this application as a service, the Wi-5 controller needs to be instructed how to deal with this trade-off.

DISCUSSION

One strength of our architecture is its ability to deal with APs and clients in the physical presence of a Wi-5-enabled network, but not being part of it, either because their owners or the hardware do not allow it. The spectrum resources that these devices consume are beyond the control of the managed network, which can only optimize the use of the remaining resources. However, initial simulations applying game theory [14] show that due to the flexibility that our architecture provides, it is always beneficial for an incoming AP to participate in the network if it can.

Wi-5 also introduces an interesting new attack vector to all the security issues that already come with SDN-like architectures. As our architecture provides flexible controlled optimization of shared spectrum, so could a hacker use our system for massive destruction of shared spectrum resources ("jamming"). Wi-5 controllers should therefore include significant defence mechanisms.

Unlike with many other SDN-like architectures, scalability is not an issue here. The number of APs served by a single controller will be limited, as WiFi has a limited range, and only the APs that have an interference impact on each other need to collaborate.

A commercial deployment of our architecture demands a good understanding of who will be operating the various architectural components, the interactions between these actors, and what will be their responsibilities. The business model we propose, briefly described in [15], introduces two new actors: the Spectrum Usage Broker and the Wi-5 System Operator. The former devises and maintains sensible spectrum sharing strategies between AP operators, whereas the latter operates the Wi-5 controller. For the use case of spectrum congestion in an apartment block, the role of the Spectrum Usage Broker would typically be taken by the apartment owners' corporation and the caretaker of the building. The Wi-5 System Operator could be an independent IT subcontractor, or one of the broadband access providers servicing the apartments.

CONCLUSION

We presented Wi-5, a radio spectrum programming architecture that enables new possibilities for managing radio resources in unlicensed frequency bands. By focusing on WiFi networks, we identified two major issues with current management systems: the inability to implement and execute new management applications, and to extend the network with hardware and software from other vendors. In short, they lack open APIs. We therefore extended the SDN control plane with radio primitives, and introduced a new plane (the spectrum plane) with various functions that enable fine-grained and QoS-aware radio resource management.

Our open source implementation is used to realize three different management applications to

reduce spectrum congestion and improve connectivity, and we assessed their performances in WiFi networks. We observed that Wi-5 indeed provides the tools to reduce spectrum congestion effectively, and to solve the “sticky client” problem.

ACKNOWLEDGMENT

We thank Ronald in 't Velt for his contribution to the vertical handover experiments. This work is partly funded by the H2020 Wi-5 Project (G.A. no. 644262); TISFIBE (TIN2015-64770-R) Government of Spain, European Social Fund and Government of Aragon (Research Group T31_17R).

REFERENCES

- [1] United Nations, “World Urbanization Prospects: The 2014 Revision,” Department of Economic and Social Affairs, Population Division, 2015.
- [2] A. Ozyagci, K. Sung, and J. Zander, “Effect of Propagation Environment on Area Throughput of Dense WLAN Deployments” *Proc. Globecom Workshops (GC Wkshps)*, Atlanta, GA, USA, 2013.
- [3] F. den Hartog *et al.*, “A Pathway to Solving the WiFi Tragedy of the Commons in Apartment Blocks,” *Proc. Int'l. Telecommunication Networks and Applications Conf. (ITNAC 2017)*, Melbourne, Australia, 2017.
- [4] V. Sivaraman *et al.*, “Virtualizing the Access Network via Open APIs,” *Proc. 9th ACM Conf. Emerging Networking Experiments and Technologies*, Santa Barbara, CA, USA, 2013.
- [5] I. T. Haque and N. Abu-Ghazaleh, “Wireless Software Defined Networking: A Survey and Taxonomy” *IEEE Commun. Surveys & Tutorials*, vol. 4, no. 18, 2016, pp. 2713–37.
- [6] J. Schulz-Zander *et al.*, “Programmatic Orchestration of WiFi Networks,” *Proc. USENIX Annual Technical Conf. (USENIX ATC 14)*, 2014.
- [7] J. Schulz-Zander *et al.*, “OpenSDWN: Programmatic Control over Home and Enterprise WiFi,” *Proc. 1st ACM SIGCOMM Symposium on Software Defined Networking Research*, 2015.
- [8] R. Riggio, M. K. Marina, and J. Schulz-Zander, “Programming Abstractions for Software-Defined Wireless Networks,” *IEEE Trans. Network and Service Management*, vol. 12, no. 2, 2015, pp. 146–62.
- [9] P. Gallo *et al.*, “CADWAN: A Control Architecture for Dense WiFi Access Networks,” *IEEE Commun. Mag.*, vol. 56, no. 1, 2018, pp. 194–201.
- [10] M. Seyedebrahimi *et al.*, “Fine-Grained Radio Resource Management to Control Interference in Dense WiFi Networks,” *Proc. IEEE Wireless Commun. Networking Conf. (WCNC)*, San Francisco, CA, USA, 2017.
- [11] J. Saldana, J. Ruiz-Mas, and J. Almodovar, “Frame Aggregation in Central Controlled 802.11 WLANs: The Latency vs. Throughput Trade-off,” *IEEE Commun. Lett.*, vol. 11, no. 21, 2017, pp. 2500–03.
- [12] J. Saldana *et al.*, “Unsticking the WiFi Client: Smarter Decisions Using a Software Defined Wireless Solution,” *IEEE Access*, 2018.
- [13] J. Shi *et al.*, “A Walk on the Client Side: Monitoring Enterprise WiFi Networks using Smartphone Channel Scans,” *Proc. 35th Annual IEEE Int'l. Conf. Computer Communications (INFOCOM)*, San Francisco, CA, USA, 2016.
- [14] M.P.P. van Heesch, “Combined Cooperation and Non-Cooperation for Channel Allocation and Transmission Power Control,” Tilburg University, 2016.
- [15] F. den Hartog and J. de Nijs, “The Role of Regulation in Preventing WiFi Over-Congestion in Densely Populated Areas,” *Australian J. Telecommunications and the Digital Economy*, vol. 5, no. 2, 2017, pp. 1–16.

BIOGRAPHIES

FAYCAL BOUHAFS is an assistant professor of Wireless Communications and Networking and holds a Ph.D. degree in computer science from LJMU, UK. Since 2015, he has been the technical lead for the H2020 Project Wi-5. His research interests revolve around spectrum congestion, radio resource management, and Beyond 5G communications and architectures.

MICHAEL MACKAY is a senior lecturer in Computing at LJMU. He received a B.Sc. (hons) in computer science from Lancaster University in 2000 and a Ph.D. from the same institution in 2005. He is widely published in a range of research areas focussed around networking technologies including wireless communi-

cations, Beyond 5G networks and the Future Internet, QoS and mobility, and cloud computing and the Internet of Things.

ALESSANDRO RASCHELLÀ received the M.Sc. in telecommunications engineering from the University Mediterranea of Reggio Calabria (UNIRC), Italy in 2007, and the Ph.D. degree in wireless communications from the Universitat Politècnica de Catalunya (UPC), Barcelona, Spain in 2015. From 2007 to 2009, he was a research assistant with UNIRC. He is currently a research fellow with Liverpool John Moores University (LJMU), UK. His research interests include wireless networks optimization, cognitive radio networks and heterogeneous networks.

QI SHI received his Ph.D. in computing from the Dalian University of Technology, P.R. China. He is now a professor in computer security and the Director of PROTECT Research Centre in the Department of Computer Science at Liverpool John Moores University in the UK. He has many years research experience in several areas, including computer networks and security. He has published over 200 papers and has also played a key role in many research projects.

FRANK DEN HARTOG is an associate professor at the University of New South Wales, specializing in complex systems security. Until 2016 he was a senior scientist with TNO in the Netherlands, and acquired and managed various large collaborative projects. From 2012-2016 he was Chair of the Technical Working Group of the worldwide Home Gateway Initiative (HGI) industry consortium. He has co-authored 64 peer-reviewed articles, 23 patent applications, 67 contributions to standardization, and seven international standards.

JOSE SALDANA is a senior postdoc researcher at the University of Zaragoza. His research interests focus on quality of service in real-time multimedia services, VoIP and networked online games, traffic optimization and resource management in wireless LANs. He serves as an editor of *IEEE Access*, and as an area editor of KSII TIS. He has served on the Organization Committee of the IEEE Consumer Communications and Networking Conference (CCNC), and in other Technical Program Committees.

RUBÉN MUNILLA worked as a support technician of Computer and Telecommunication Service of University of Zaragoza (SICUZ) from 2011 to 2014. He worked as a software developer and support personnel at the Higher Engineering and Architecture School of UZ (EINA) until 2016. In 2017, he moved to Araelec, focused on telecommunications infrastructure. He is currently with the Department of Electronics Engineering and Communications in UNIZAR, working in wireless networks and SDN development and testing.

JOSÉ RUIZ-MAS is an associate professor in the Department of Electronics Engineering and Communications, University of Zaragoza. He is a member of the Aragn Institute of Engineering Research (ISA). He is a co-investigator of several research projects supported by companies including Telefonica and Teltronic. At present his research activity lies in the area of communication networks with special emphasis on wireless networks, distributed multimedia system, Quality of Service (QoS) and Quality of Experience (QoE).

JULIÁN FERNÁNDEZ-NAVAJAS is an associate professor in the Department of Electronics Engineering and Communications, University of Zaragoza. He is a member of the Aragn Institute of Engineering Research (ISA). He is a co-investigator of several research projects, funded by the EU, the Ministry of Science and Technology, the Sanitary Research Funds and the Government of Aragn, in the areas of distributed multimedia systems and wireless networks, Quality of Service (QoS) and Quality of Experience (QoE).

JOSE ALMODOVAR obtained his B.Sc. degree from the University Carlos III (Madrid) in 2007 and his M.Sc. degree in telecommunications at Delft University in 2012. Since 2013, he has worked as a network control scientist at TNO, Netherlands. He has participated in several European research projects and is also an active contributor to 3GPP.

NIELS VAN ADRICHEM is a communication networks scientist at TNO and holds a doctorate in computer science from Delft University of Technology. His research interests revolve around Future Internet Architectures with a specialization in SDN, NFV and cloud infrastructures. In his work, he addresses computer networking problems that show both practical relevance and theoretical complexity, the results of which, besides academic recognition through peer-reviewed publications, also show in a significant number of open-source code contributions.

Our open source implementation is used to realize three different management applications to reduce spectrum congestion and improve connectivity, and we assessed their performances in WiFi networks. We observed that Wi-5 indeed provides the tools to reduce spectrum congestion effectively, and to solve the “sticky client” problem.