4

Wireless Software Defined Networking

Claude Chaudet[1] and Yoram Haddad[2]
[1] *Telecom ParisTech, Institut Telecom, Paris, France*
[2] *Jerusalem College of Technology, Jerusalem, Israel*

4.1 Introduction

The recent years have witnessed the advent of a digital, mobile, and connected society in which digital devices are used for countless applications. People use wireless devices such as computers, smartphones, or tablets for communicating, producing, accessing, and sharing knowledge; shopping; interacting with public services; taking and storing photographs: finding directions: listening to music: watching videos: studying; or playing. More usages are expected to appear in the next years as the terminal size and cost decrease. The availability and the performance of wireless communications have obviously played a major role in this evolution.

Compared to the digital world of the beginning of the year 2000, people now often possess multiple terminals (smartphone, tablet, computer) and select one of the terminals based on the situation. A professional can work on a presentation from his office computer, make a few modifications on his tablet on the plane, and rehearse it on his smartphone on the metro while getting to a meeting where he will use his laptop to display it. Someone may start watching a movie at home on his TV and finish it in the bus while traveling to work. These scenarios, which are a reality today, have an important impact on the networks.

First, data is more and more often synchronized among terminals, thanks to cloud services. However, this permanent synchronization process consumes wireless channel bandwidth and creates important data flows that remind peer-to-peer applications. If the average size of the exchanged files is smaller, the number of exchanges, compared to the wireless link capacity, is important. Besides, the deployment of self-hosted cloud solutions induces similar user-to-user traffic patterns, except that few peer-to-peer applications were hosted by moving devices.

Software Defined Mobile Networks (SDMN): Beyond LTE Network Architecture, First Edition.
Edited by Madhusanka Liyanage, Andrei Gurtov, and Mika Ylianttila.

Handovers are therefore more and more frequent, and concern multimedia data that can tolerate losses better than delays, as well as file transfer that can tolerate delays as long as the TCP connections remain. Finally, when looking at the scenario involving terminal switching while watching a streaming movie, content adaptation may be required.

To fully satisfy the user, wireless infrastructure networks and wireless LANs (WLANs) should therefore find efficient ways to manage horizontal and vertical handover content adaptation while preserving their scarce bandwidth and reducing as much as possible their exploitation costs. Numerous research initiatives are driving the technology toward these objectives. Femtocells and power control improve spatial efficiency; traffic off-loading and multihoming allow to share bandwidth between technologies; in-network data replication and information-centric networks (ICN) help content distribution; power savings and green networking improve the operational costs; etc. Some key challenges remain, though: how to implement these techniques and make sure they interact properly.

Operators therefore need a way to manage their wireless network as a whole efficiently and in an evolutionary manner. Besides, several optimization possibilities would also require coordination and cooperation between operators, users, and infrastructures, which is difficult to achieve in practice. This advocates for a paradigm shift in the way wireless networks are managed. A network should be able to increase or decrease its capacity on demand, providing the desired quality of service (QoS) to users when possible and avoiding disturbances to close networks when possible. This requires a bandwidth management entity capable of examining not only the demand and the user service-level agreements (SLAs) but also the state of the different access points (APs) or the wireless channel status and to take appropriate decisions. This management entity shall build its vision of the network by gathering data from various network devices and perform global optimization to ultimately suggest connection options to each user and for each application. This paradigm is referred to as *network virtualization* or *software defined networking* (SDN).

As exposed in the other chapters of this book, the SDN concept emerged as a way to foster innovation, allowing experimenters to use a production network without noticeable impact, and as an elegant way to provide better QoS or to improve network reliability. OpenFlow [1] was introduced in 2011 following these concepts and has since then been pushed by the industry, leading to the creation of the Open Networking Foundation (ONF), a nonprofit organization that manages what has now become a standard. OpenFlow basically consists in separating the control and forwarding planes by letting the interconnection equipment match packets based on a cross-layer set of 12 fields (the 12-tuple) and by deporting all the intelligence to central entities called controllers. It is these controllers that decide the policy that applies in case of successful/unsuccessful matching. Multiple controllers can coexist and manage multiple and independent visions of the network called *slices* in the OpenFlow terminology. SDN in the OpenFlow vision therefore consists in three functions: flow-based forwarding that requires packet matching against a flow table, status reporting from each interconnection device, and slicing that requires the capability for each interconnection element to isolate traffic. An OpenFlow architecture is composed of the interconnection devices that only perform matching and forwarding, the controllers that decide and publish policies regarding flows handling, and an entity, called *FlowVisor*, that presents a sliced vision of the physical infrastructure to the controllers.

We refer the reader to other chapters of this book for details on this aspect and will focus, in this present chapter, on the wireless extension of SDN that can be called *wireless SDN* or *software defined wireless networks* (SDWN). Section 4.2 elaborates around the concept of wireless SDN and the extension to OpenFlow for wireless. We then present related works and some existing projects in Section 4.3. Section 4.4 exposes the opportunities that a successful wireless SDN implementation could bring, while Section 4.5 lists some of the key challenges to overcome.

4.2 SDN for Wireless

The software defined network concept has initially been imagined with data centers and fixed networks in mind. However, the wireless world could greatly benefit from such a framework. Indeed, most wireless technologies have to face limited resources. Cellular networks have to deal with an ever-increasing mobile traffic that is handled with difficulty by the radio access networks. Mobile operators try to imagine solutions to off-load data to WLANs in urban areas, which can theoretically accommodate more elastic flows thanks to random access. However, wireless LANs also become congested, as the unlicensed frequency bands are utilized by multiple technologies whose traffic also increases. Moreover, the deployment of WLAN APs in urban areas is denser and denser and uncoordinated, making interference mitigation difficult.

Besides, the mobility of users also has an impact on the available resources. Users should be able to keep their communications and connections open while traveling, which requires the operator to predict user mobility and to reserve a part of the resources for users passing from cell to cell. Several handover procedures are possible, but this process often requires dedicating a part of the resources to mobility management.

These problems are all related to the scarce capacity offered by wireless channels compared to the demand of the applications, which is often scaled on the wired connection performance. Classical strategy to solve this issue, in the past, consisted in improving the spectral efficiency by working on modulation or coding. However, this process has a limit, and it is now common to read that we are reaching Shannon capacity limit. The most optimistic predictions place today's technology within 20% of this limit, and the ultimate efforts to improve signal over interference plus noise ratio are expected to be very difficult.

However, recent reports show, using frequency band scanning, that only 2% of the wireless spectrum between 30 MHz and 3 GHz is effectively used in some areas of the world.[1] However, the frequency bands are allocated by national regulators and are granted with caution to applications in order to keep spectrum free for the future and to preserve space for potential tactical communications. If software defined radio (SDR) (aka cognitive radio) provides ways to make different classes of users cohabit on a single frequency band, hardware is not mature enough to make this procedure seamless. Moreover, changing regularly channels and consequently the offered bandwidth requires a strong service adaptation layer.

The fact is that this large frequency band already hosts dozens of wireless transmission technologies. Most of them are unidirectional (FM radio, broadcast television, etc.) and are not suited for data communications. However, WLAN technologies (e.g., IEEE 802.11), cellular broadband networks (e.g., UMTS, LTE, WiMAX), or even satellite networks, for example, can all be utilized by connected data services, even though their performance levels are heterogeneous. However, using these technologies conjointly or successively requires to be able to perform handovers between technologies and between operators,

---

[1] http://www.sharedspectrum.com/papers/spectrum-reports/

which both pose serious technical issues. Operators already respond to the lack of bandwidth on certain networks by off-loading traffic to other networks. For instance, in the access networks, several mobile operators have a solution to off-load traffic from the 3G network to a partner Wi-Fi network when a user is in range [2], which nevertheless causes disconnection and performance issues for mobile users, as handover is not properly handled especially regarding security. Authentication and encryption keys often need to be reestablished, which causes delays that are often incompatible with high mobility. Besides, off-loading needs to be carefully managed, as it could easily cause overload. The most emblematic examples of such situations come from other domains. Electrical networks use off-loading extensively, and the lack of coordination between operators has been one of the causes of the escalade in electrical failure of November 4, 2006, that led to a blackout in occidental Europe [3]. If the consequences in telecommunication networks will never reach this level, undesired saturation can appear.

A third solution to the scarcity of the wireless spectrum consists in enhancing the spatial efficiency of wireless communications. Bringing the users closer to their serving base station would allow to reduce both mobile terminal and base station transmission powers, therefore generating less interference on close cells. Femtocells in cellular networks work this way, even though their goal is more to improve coverage than to reduce interferences. Nevertheless, a generalization of this principle brings its load of issues. If each user deploys a miniature plug-and-play base station, there will be no possibility for the operator to control its position, and global planning strategies are therefore impossible. There is a slight chance that wireless spectrum-related problems will appear in some areas instead of being solved, especially when different operators work on the same band of frequencies (e.g., Wi-Fi, UMTS).

These examples show that if solutions exist to alleviate issues related to the spectrum scarcity, the coordination is necessary to manage the radio resource globally, even across operators, and to adapt the network behavior to the user traffic. And that is precisely where software defined networks can help. Operators could use the feedback features from the APs or base stations as well as the slicing facilities. Moreover, as the free bands of the wireless spectrum become overcrowded, decoupling the traffic operators that provide users with connectivity from the infrastructure operators that could run different or heterogeneous technologies would certainly alleviate the problem.

In September 2013, the ONF published a brief paper entitled “OpenFlow-Enabled Mobile and Wireless Networks” [4]. This document describes several outcomes of wireless SDN, focusing on radio access network performance optimization. They list several key challenges for implementing wireless SDN before focusing on two major issues: wireless channel resource management through intercell interference reduction and mobile traffic management through roaming and off-loading. If this shows a part of the potential of wireless SDN, this document also reminds that the wireless channel has some specificity that makes the SDN concept tough to implement. That’s why this marriage has not been extensively studied yet. However, the set of problems it brings is also what makes it interesting to study.

First of all, the wireless medium is a fundamentally shared medium. The few free frequency bands such as the ISM bands are shared by multiple clients and multiple technologies. Concerning slicing, if some solutions have been proposed to reduce the interference between technologies (e.g., Bluetooth and IEEE 802.11), the narrowness of the available bandwidth makes the number of independent channels too low to efficiently implement slicing when considering that close APs may interfere. In reserved frequency bands such as cellular networks, similar problems arise when the number of mobile virtual network operator (MVNO) increases.

Concerning status report, the problem is even more difficult. First, the wireless channel state changes very frequently, especially in an indoor scenario. Fading and shadowing, for example, can easily make a link disappear suddenly, provoking frequent updates on link state that need to be considered in the routing protocols. A controller therefore needs to evaluate more than the simple channel or device load; it also needs to acquire information on the link stability, for example. This variation is partially due to the variations in the physical environment (doors closing, people passing, etc.), but also to the presence of close APs that have their own traffic pattern and that do not necessarily belong to the same operator. Discovering these potential interferers is challenging.

Finally, given the variability of the channel conditions, the status report from the AP is likely to generate a lot of control traffic that could pass over wireless links whose bandwidth is limited.

4.2.1 Implementations: OpenRoads and OpenRadio

OpenRoads is the adaptation of OpenFlow to wireless networks. It relies on OpenFlow to separate control path and data path through the FlowVisor open API. As for the wired case with OpenFlow, in OpenRoads, the network OS constitutes the interface between the infrastructure and the applications that observe and control the network.

The OpenRoads project built a demonstration platform composed of Wi-Fi and WiMAX APs that has been used for academic courses. Successful student projects demonstrate the potential of the approach through the development of applications or of an n-casting mobility manager. Basically, this consist in multiplexing at the receiver the same traffic coming from multiple base stations at the same time and from different networks such as Wi-Fi and WiMAX. Packet duplication facilitates the handover between technologies and increases the QoS. Some other mobility manager implementations succeed to improve the handover process with a reported reduced packet loss rate [5] and represent no more than a dozen lines of code [6]. It is worth to mention that there exist lots of wireless platforms over the academic world. The specificity of wireless SDN labs is that they provide wireless platforms but with a wired backbone control [7].

If these projects were only early developments, full utilization of all the features of OpenRoads is expected in the midterm, when the infrastructure includes programmable radio hardware platform. These flexible radio interfaces will be controlled through an API that allows external selection of various physical parameters, modulation, and coding scheme, for example. This control over the entire protocol stack is difficult to achieve, but it is necessary to allow full flexibility at the data plane level. Even with such capabilities, it is still not trivial to implement the rule-action abstraction in wireless world.

In the current version, slicing in OpenRoads is implemented by creating virtual interfaces on the same AP and assigning different service set identifier (SSID) to each interface. Each SSID can be considered as a separated slice and may be managed by a different controller. Even though the controller can apply different policies to different users, it is limited by the physical constraints and the hardware capabilities of today’s APs. All slices should use the same channel and power settings, for example, which limits isolation. This demonstrates that it is fundamentally different to run SDN on regular wired switches than on wireless APs.

More recently, Bansal et al. [8] describe OpenRadio, an implementation of the separation between a decision and a processing plane that are similar to the control and forwarding plane of OpenFlow. They demonstrate the effectiveness of the concept by implementing Wi-Fi and LTE protocols over generic DSPs. OpenRadio, which is one of the SDN projects of the Open Networking Research Center (ONRC) at Stanford University, targets dense wireless networks and the development of a wireless controller, as well as the provision of slicing functionalities.

4.2.2 SDR versus SDN

SDR, also called cognitive radio, designates a set of techniques designed to cope with the shortage of available wireless channels. SDR makes a limited use of hardware components (for digital/analog conversion) and runs most of the tasks in software (filtering, coding, modulation, etc.). This software-based architecture is exactly in the line of the software defined network philosophy, as it allows remote definition of the channel parameters. Hence, SDR technologies appear at least as a fundamental building block for wireless SDN.

However, if SDR can clearly help in implementing wireless SDN, it is not a complete solution. First of all, SDR technologies generally try to fill in the so-called white space, that is, the frequency bands that are not used. However, as soon as a channel is used, even partly, it is considered as busy by the SDR and classified out of the usable bands. In addition, the focus of SDR is on the physical layer, but SDN require to consider the full protocol stack.

When looking from a higher point of view, SDN should allow to balance load between different operators, which not only requires to adapt the physical parameters of the base stations to facilitate the user handoff but also to exchange information between the operators to make the service providers aware of the resource blocks (in the LTE terminology) they were granted. SDR can therefore provide the capability to define dynamically and remotely the physical layer parameters, but does not represent a complete solution.

4.3 Related Works

Although the concept of SDN is recent, the developments around the various platforms and implementation issues already lead to a few publications. Some articles investigate the benefits and challenges behind the deployment of SDN in wireless networks of embedded devices like personal area networks or body area networks [9], which require device energy management and in-network aggregation. We stand here more from the infrastructure's point of view, even though terminal cooperation should be highly beneficial.

Our intent in this section is not to provide here a complete state of the art, but to select some works that illustrate some of the challenges and opportunities mentioned in Sections 4.4 and 4.5.

Concerning wireless SDN, Dely et al. [10] propose an interesting implementation of a wireless mesh network using SDN. In this work, each radio interface is split into two virtual interfaces, one for data transmission and one for the control packets. These virtual interfaces use two different SSIDs. At the controller level, a monitoring and control server is used in parallel of the controller, whose role is to maintain a topology database used by the controller to compute the optimal data paths. The implementation provides performance results that demonstrate that some serious issues exist regarding rule activation time (i.e., the time required to set up a new rule for a new flow by a remote controller), rule processing time (which can be important if lots of rules have to be parsed before the matching to succeed for an incoming flow), and the famous scalability issue related to the volume of traffic generated by the control plane. Yang et al. [11] describe briefly the architecture of a radio access network based on SDN that separates the network in four hierarchical levels: one that represents each operator, one that virtualizes infrastructure devices, one that manages the wireless spectrum, and the last one that constitutes the wireless SDN.

Some works deal with the performance of OpenFlow in the wired case. For instance, see Ref. [12], which focuses on the interaction between a router and the controller and evaluates the mean sojourn time of a packet in the system, taking into account the probability that no rule matches the flow the packet belongs to. It is shown that the sojourn time depends mainly on the processing speed of the controller used where the measurements show that it lies between 220 and 245 μs. It also evaluates the probability that a packet is dropped due to a limited buffer space at the controller. Bianco et al. [13] focus on the data plane and compare network efficiency with and without OpenFlow regarding throughput and packet latency. For instance, it is shown that OpenFlow experiences a performance (latency and throughput) drop of 11% of the packets compared to regular layer 3 routing, when small packets (64 bytes) are considered. However, almost equivalent performance is measured when packet size is slightly increased (96 bytes or more). As for the OpenRoads wireless extension of OpenFlow, Yap et al. [7] showed that the increased load generated by the communications between the devices (switches, etc.) and their controller is not important and stands for <0.05% of all traffic volume transferred in the reported experiences.

4.4 Wireless SDN Opportunities

The potential of SDN and the OpenFlow approaches has already been demonstrated into numerous examples concerning the network infrastructure. This paragraph lists a few technical opportunities that are specific to wireless, either by addressing wireless-specific problems (interferences, fast channel quality evolutions, etc.) or by using the specificities of the wireless medium (e.g., the broadcast channel) to improve generic tasks.

4.4.1 Multinetwork Planning

WLAN standards, as well as most unlicensed technologies, have a limited number of independent channels. Wi-Fi, for example, has only three nonoverlapping channels in the 2.4 GHz band and eight in the 5.2 GHz band. New technologies at the physical layer, such as ultrawide band, could increase this number, but there will always be a limit to the number of terminals that can operate without any interference in a given geographic area. Today, such interference problems appear at the building level, between APs of neighbor users, and tomorrow, close personal or body area networks could also interfere in a similar manner. Figure 4.1 represents the channels attributed to different Wi-Fi APs by their owners. Terminals located at the intersection of the two red or green circles are likely to experience poor network performance as the emissions of the two corresponding APs will interfere and cause collisions. Figure 4.2 represents a more efficient allocation in which close cells operate on different frequencies.

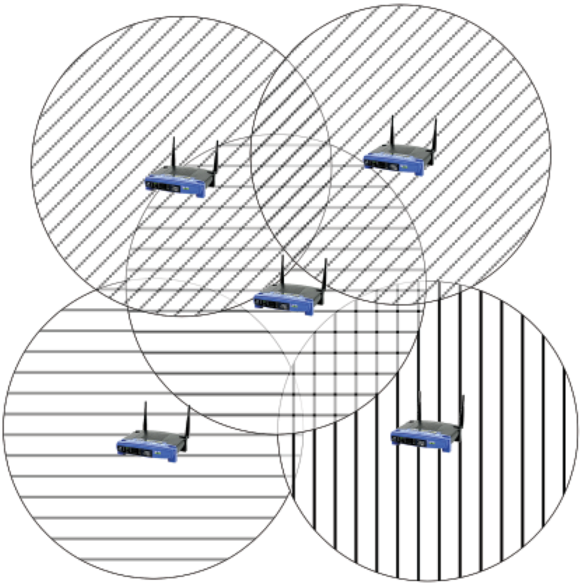

Figure 4.1 Uncoordinated Wi-Fi channel allocation.

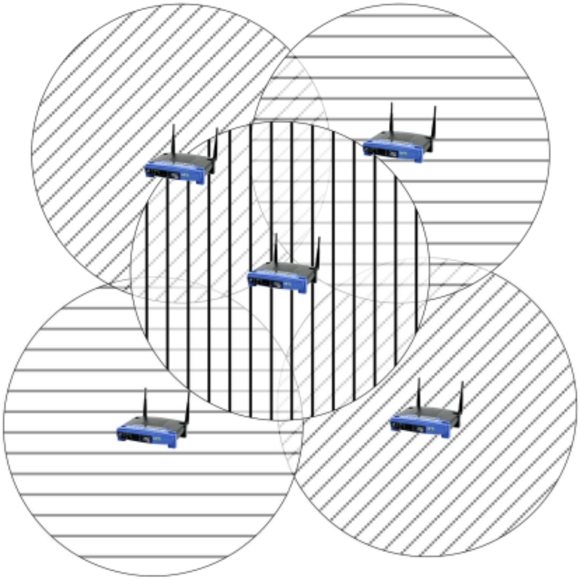

Figure 4.2 Interference-free Wi-Fi channel allocation.

SDN could provide help when it comes to such wireless network planning. It is indeed possible to create zone-specific and operator-independent controllers, which can even be distributed processes, that would be capable of aggregating statistics coming from the APs and could decide of channel allocations and transmission power of the access points in order to minimize interdependencies and interferences, similarly to the Control and Provisioning of Wireless Access Points (CAPWAP) protocol [14]. CAPWAP, which has been implemented in some Wi-Fi AP models, could integrate smoothly in global multitechnology SDN, with CAPWAP serving as a monitoring and control interface for Wi-Fi and compatible APs.

Nevertheless, realizing such a controller is not easy for several reasons. First, the algorithmic problem to solve can be solvable or not. Frequency allocation is a graph vertex k-coloring problem, which may or may not be feasible, depending on the number of independent channels. Brooks' theorem proves that at most $\Delta + 1$ channels are required to provide independent access in a graph whose maximum degree is Δ, but the density of the terminals within a single transmission range will most likely reach a value superior to the number of available channels. Power control could help, as it reduces the transmission range and hence the density, but it requires solving a plane or 3D packing problem, and the disconnection probability increases as the transmission ranges decrease.

In any case, such a power and frequency allocation problem requires full collaborations from the APs, which are the only ones capable of monitoring the different channel occupancy levels. A single nonparticipating access point could break the solution, and game-theoretic algorithms should be imagined to solve these issues.

4.4.2 Handovers and Off-Loading

Today's devices are highly portable, ergonomic enough to allow users to interact with online services while moving, and they are also equipped with multiple wireless interfaces. A state-of-the-art smartphone can potentially be connected to the Internet through a cellular access (LTE, UMTS, etc.), Wi-Fi, and even Bluetooth LE if a compatible gateway is around. A device could therefore select the best network AP(s) to access remote services based on a combination of criteria that include link quality (throughput, SINR, etc.), stability, billing, QoS capabilities, as well as the operators' preferences regarding, for example, off-loading. As the operators providing each technology access are not necessarily the same, the terminal is likely to be multihomed and to realize conjointly three types of handovers.

Classical "horizontal" handovers across APs are common in cellular networks and are well handled by operators through mobility prediction and resource soft reservation. As Wi-Fi localization is relatively precise today in urban areas, a cellular operator may expect a slight improvement in mobility prediction. One can imagine to implement local mobility management-dedicated controllers that aggregate statistic on signal strength or disconnection events coming from various technologies and send back prereservation orders to the appropriate neighbor cells.

Vertical handover across technologies happens when a cellular operator favors off-loading to Wi-Fi networks. Off-loading today suffers from the short range of Wi-Fi APs, which makes the connection very episodic, and from the long authentication delay. Even if EAP-SIM 802.1X method allows sharing the authentication token across technologies, the process still takes too much time, and a user traveling by car in a city often experiences connection and disconnection events that are not properly handled by the terminal. In this case, SDN could help by providing cross-technology soft handover, transmitting simultaneously the same data across both networks until the connection stability is confirmed. SDN could also realize the authentication procedure without involving the terminal.

Cross-operator handovers may happen when a mobile user temporarily or permanently leaves the coverage area of an operator he used to start a connection, for example, leaving the home Wi-Fi network when going outside, leaving the cellular operator when taking an elevator, or even crossing a frontier and roaming to a new cellular operator. Today, the user expects its TCP connection to break; however, SDN could help, in this case, to preserve it by automating the implementation of a Mobile IP-like scenario. In these situations, using SDN could help to reduce the handoff latency by facilitating the initial Mobile IP handshake. Local user location and travel monitoring can help predict handoffs and trigger prereservation of resources in the upcoming operator network, while a roaming-dedicated controller can prepare the control packets necessary to change Care-of Address and register to the new foreign agent.

Going a bit further, a multihomed terminal could easily decide dynamically how to use its multiple connections based on QoS considerations. Specific wireless technologies can be bound for a specific application. For instance, Wi-Fi is suitable for applications that require high throughput but can tolerate a variable delay. Cellular 3G technology provides a more consistent bandwidth and a good coverage. If SDN are able to have a good vision of which networks and technologies are available in different areas, they could "intelligently" allocate specific resource of a specific wireless technology for a specific user QoS requirement [15]. Traffic could also, in particular cases, be split among technologies, as illustrated in Figure 4.3. Let us take the example of a videoconferencing chat. The images of the participants require a high throughput, but the speech is the most important flow, as its clarity has a real impact on the quality of experience. One could imagine transferring the video data over a WLAN, while the speech arrives through a link that supports resource reservation such as a cellular network.

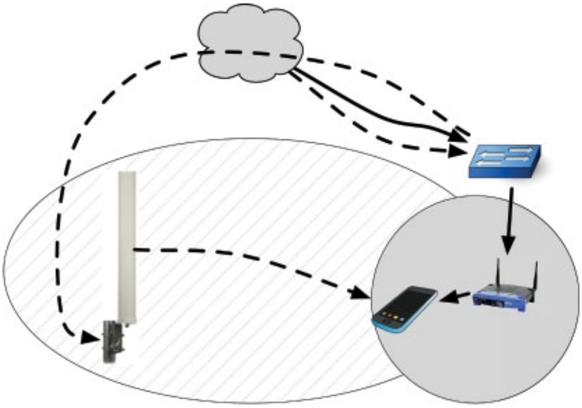

Figure 4.3 Example of a terminal receiving two parts of a data flow from heterogeneous networks: the switch decides to forward a part of the traffic through a more reliable cellular network rather than using the congested Wi-Fi channel.

Nevertheless, the vision from the infrastructure is imprecise and may not be accurate enough to take a clever decision for a given terminal. In this case, SDN could bring to the user information on the available networks, their stability, their coverage, and so on in order to feed a local decision process.

4.4.3 Dead Zone Coverage

One interesting case in cellular coverage is the dead zone. This is critical in some rural zones or large forest or also occurs when there is a disaster. Hasan et al. [16], for example, examine the scenario in which multiple operators share the same network infrastructure, and not only the antennas, to cover rural areas.

In some cases, a user located in a dead zone is not covered by an AP but is still within the range of another user. Both users can communicate via ad hoc networks, but the “uncovered” user will not be able to be reached by the AP since the AP does not know that the user is in fact reachable through ad hoc. Thanks to the concept of wireless SDN (with respect to what we mention in Section 4.5.2), we can use the topology discovery capabilities at the controller to help backhauling users in some dead zone through ad hoc network via neighboring users. Of course, this implies some reports by the end user to its covering AP to help it build an accurate view of the topology of the users covered. But since the controller gathers in the same place the network topology information, it won’t be difficult to run an algorithm that identifies the possible holes and in consequence establishes “virtual links” to the devices physically disconnected from the AP. Therefore, merging ad hoc network and mesh network with wireless SDN exhibits very large amount of opportunities that needs to be further investigated.

4.4.4 Security

The monitoring capacity of OpenFlow and OpenRoads can provide a clear vision of a network status to an entity in charge of detecting intrusions or abnormal behavior. The network load and the packet distributions per protocol can be compared to statistics, and a process could decide if the current traffic matches the expected values for this date and time. Suspicious situations may indicate an intrusion or the presence of inside computers participating to a botnet, for example, control traffic can also be examined. A high ARP activity may raise a warning when the network topology and traffic did not change. Besides generic intrusion detection based on signatures, more specific situations can be addressed:

- The status reports can be used to detect misbehaving users or routers in a collaborative network (e.g., mesh network or public Wi-Fi network) in which users are expected to route traffic for mobile clients in exchange for a similar service. The owner of a participating AP could insert a traffic shaper behind its shared AP to preserve its uplink bandwidth. Comparing the different statistics coming from close APs and for upstream and downstream interconnection devices, this kind of behavior could be detected.
- Similarly, in a QoS-enabled wireless network, an emitter always classifying its frames into the highest priority class to gain access to the medium more often than its share could be detected by close APs and terminals (provided that terminals also participate into statistics collection; see following text).

- User locations and movements can be tracked by looking at the MAC addresses uploaded by the APs of a LAN [7]. These statistics could be correlated with larger-scale movements obtained from a 3G system (which may cause privacy issues, as mentioned in the following), which could help detect physical intrusion attempts and orient a sensor network.

4.4.5 CDN and Caching

Wireless SDN could use content-centric paradigm where the base station could gather and store some data so that it can be delivered in a timely manner to delay sensitive application. The controller could identify the users' needs and according to this associate the users with the closest base stations that hold the required data. An interesting use case could be people working in finance in the Wall Street district in New York. It is likely that a large number of the users require the same data from some marketplace all the time and with very short delays. Instead of forwarding the query of each user to the marketplace server, it could instead gather and store the data in the base stations and deliver it in downlink as soon as requested.

4.5 Wireless SDN Challenges

The previous section detailed some potential benefits of SDN for the general wireless scenario. However, the wireless medium specificities also bring its load of challenges that need to be addressed.

4.5.1 Slice Isolation

As mentioned earlier in Section 4.2, defining slices in a wireless network is not easy as link or channel isolation cannot be guaranteed. Among a single infrastructure, it is possible to plan frequency usage and to control on which frequency each AP operates to minimize or avoid interferences. Depending on the number of available channels, it can be possible to create independent slices if the operator has a full control over the channel space. However, when it comes to open technologies, Wi-Fi, for example, a given operator does not control the AP locations, which are placed by the customers in their homes, and does not control the close AP frequencies either, especially when multiple operators are present in the area. Therefore, the creation of a wireless link that is isolated and that does not disrupt close networks cannot be guaranteed unless the channel space is under control.

4.5.2 Topology Discovery and Topology-Related Problems

One of the required functions of SDN concerns the upward communication of network status. However, the status of a wireless network is not easy to measure, especially when multiple interferences can impact the quality of transmissions. Ideally, an AP could want to identify close APs and determine various parameters such as their channel(s), their output power, or their traffic patterns. This process relates to topology discovery in multihop networks and could benefit from the classical techniques used in these networks (periodically transmit neighbor discovery packets that include list of 1-hop neighbors to automatically discover 2-hop neighborhood). However, it has to be performed carefully:

- First, there may exist extended versions of the classical hidden/exposed terminal situations in which a mobile user is in range of two APs that do not see each other. As the two APs are not mutually aware of their existence, their transmissions are not correlated, and the middle user suffers from collisions on its downstream traffic, getting a poor QoS. This means that these two APs are not as independent as they are able to detect, and the only possible solution to this kind of issues is to implicate the terminal using uplink feedback or explicit control packets.
- Second, discovering topology may require more than simply identifying the APs and their operation frequencies. Two close access point operating on the same channel could indicate either a problematic situation if the two APs interfere or could correspond to a desired situation, in which one extends the range of the second one. If channel usage reports are used to derive policies, these situations should be identified.

4.5.3 Resource Evaluation and Reporting

Orthogonally to topology discovery, an infrastructure, for status report purposes, should be able to evaluate its available resources (e.g., channel capacity). However, this is far from obvious. The problem of identifying potential interferers has already been mentioned when discussing topology discovery. It is even more serious when it comes to resource evaluation because even if it is not necessary to identify the different interferers, their timely pattern needs to be imagined. If it is possible to measure and report the current network status, predicting its evolution even in the next few seconds is almost impossible. It means that the usage of these statistics is limited and that this information shall not be trusted to implement strong guarantees.

4.5.4 User and Operator Preferences

Slicing and status report are necessary to implement SDN and, as seen earlier, are difficult to implement in the wireless domain. However, there are also other issues to solve that are more related to the user experience than to the proper network controller operation. The network can define access point physical parameters and behavior in order to minimize interferences or to enhance the global network performance. Nevertheless, these objectives may be blind to the individual user preferences or can be counterproductive.

The question of how to specify and take into account user preferences is also expected to arise in the wireless context, as wireless technologies enable mobility and mobility implies that a given user will be successively or simultaneously connected to various networks with different pricing policies and different levels of QoS and trust. Today's solution has some logical yet simple policies (prefer Ethernet over Wi-Fi as it has a better performance, prefer Wi-Fi over UMTS because it off-loads the operator's network, etc.). The questions of how to implement a personal mobility manager that makes the user able to specify its own preferences in a simple way and how to mix these preferences with global- or operator-level objectives are important.

4.5.5 Nontechnical Aspects (Governance, Regulation, Etc.)

4.5.5.1 Interactions between Operators

A desired efficient network would be one where users could migrate freely between the infrastructures of different providers. The service provider to whom the users subscribe could realize the payment to the infrastructure holders. For this purpose, we need a clear distinction between the network infrastructure and the service delivered. But this openness gives rise to nonobvious economic and regulatory issues.

Let us bring the example of the electrical market that, in Europe, has seen the separation from the electrical produced companies and the electrical transporter companies. There are numerous advantages over this model, the first one being allowing a more open concurrency without requiring new providers to deploy their own power transport and distribution lines. Even though associations of transport providers exist (UCTE in Europe, ETRANS in Switzerland), the existence of multiple such associations and their limited communication impact the vision operators have of other infrastructures, even though they are interdependent. This lack of a shared real-time vision of the network status has been one of the major causes of the 2003 Italian–Swiss electrical blackout [17]. Obviously, the consequences of a separate vision are not that critical in IP networks. To avoid similar situations in IP networks (which would indeed have less severe consequences), there is a need for redefining peering and transit contracts, as the operators connecting the users shall not own the infrastructure and consequently not have a real-time vision of the network status.

4.5.5.2 Service and Forwarding Provides Interactions

As for every service where there is a decoupling between the service provided and the infrastructure needed to provide this service, we face also in our case the major question which is: Can the service provider be also an infrastructure provider, and vice versa? Historically, when the service is only at its beginning, the same company that built the infrastructure also provides the service. This is understandable since the deployment of infrastructure is very expensive. But with time, governments seek more competitiveness, and the market is open to new competitors that are generally service providers that lease infrastructure from original operators. To some extent, we can cite the MVNO case in the cellular field, but this should be generalized to all wireless and wired technologies. One of the issues that arise from this challenge is the question of fairness. For instance, assume that we allow a service provider to be also an infrastructure provider, and then we will have some problem of objectivity of the measurements that are required for the users to be able to compare the service quality of different providers and in different networks. This is only one out of the many questions that show how complex are the regulatory decisions regarding the separation between authorities.

4.5.5.3 Privacy-Related Issues

As mentioned before, SDN relies a lot on status reports from the access points and even from the users. These status reports can easily include user identities or MAC addresses, which can be used to predict mobility but also to finely track users, causing some serious privacy issues.

Such tracking issues already exist in cellular networks; nevertheless, the cellular operators are clearly identified, and the tracking resolution is far lower than what Wi-Fi could achieve.

Moreover, besides the users' tracking issues, the published statistics could be exploited for malicious activities. Let us imagine, for example, that a user publishes statistics on his home network usage to a neighborhood-level controller to regulate the area wireless network. A malicious user gaining access to the controller could use these statistics to determine when the user is home and when he is away.

4.6 Conclusion

In this chapter, we examined the potential benefits and the challenges behind the adaptation of an OpenFlow-like software defined network paradigm to wireless networks. The presence of the controller that can be a mixture of a centralized entity and a collection of distributed small controllers brings multiple advantages when it comes to managing the radio resource and the user mobility. As this controller gathers data from various measurement points and potentially from various technologies, it is able to take informed decisions on all the radio parameters and to optimize, even using machine learning, the network operation. This controller can even exchange information with other operators' controller entities, which would allow collaboration without giving full access to confidential data.

However, implementing wireless SDN also poses some intrinsic challenges, and even though some implementations have already been demonstrated, the demonstration scenarios remain modest or controlled today. The task appears difficult but not impossible and the aim of this chapter is to point out hard points that need to be solved.

References

[1] McKeown N, Anderson T, Balakrishnan H, Parulkar G, Peterson L, Rexford J, Shenker S, Turner J. OpenFlow: Enabling innovation in campus networks. ACM SIGCOMM Computer Communication Review. 2008;38(2), 69–74.

[2] Lee K, Rhee I, Lee J, Chong S, Yi Y. Mobile data offloading: How much can WiFi deliver? In: Proceedings of ACM CoNEXT 2010. Philadelphia, USA; 2010.

[3] Union for the Co-ordination of Transmission of Electricity (UCTE). Final Report on the European System Disturbance on 4 November 2006. Union for the Co-ordination of Transmission of Electricity; 2006.

[4] Open Networking Foundation. OpenFlow-Enabled Mobile and Wireless Networks; 2013. ONF Solution Brief.

[5] Yap KK, Sherwood R, Kobayashi M, Huang TY, Chan M, Handigol N, McKeown N, Parulkar G. Blueprint for introducing innovation into wireless mobile networks. In: Proceedings of the Second ACM SIGCOMM Workshop on Virtualized Infrastructure Systems and Architectures (VISA'10). New Delhi, India; 2010, 20–32.

[6] Yap KK, Kobayashi M, Sherwood R, Huang TY, Chan M, Handigol N, McKeown N. OpenRoads: Empowering research in mobile networks. ACM SIGCOMM Computer Communication Review. 2010;40(1), 125–126.

[7] Yap KK, Kobayashi M, Underhill D, Seetharaman S, Kazemian P, McKeown N. The Stanford OpenRoads deployment. In: Proceedings of the 4th ACM International Workshop on Experimental Evaluation and Characterization (WINTECH'09). Beijing, China; 2009.

[8] Bansal M, Mehlman J, Katti S, Levis P. OpenRadio: A programmable wireless dataplane. In: Proceedings of the First Workshop on Hot Topics in Software Defined Networks (HotSDN'12). Helsinki, Finland; 2012.

[9] Costanzo S, Galluccio L, Morabito G, Palazzo S. Software defined wireless networks: Unbridling SDNs. In: Proceedings of the 2012 European Workshop on Software Defined Networking. Darmstadt, Germany; 2012.

[10] Dely P, Kassler A, Bayer N. OpenFlow for wireless mesh networks. In: Proceedings of 20th International Conference on Computer Communications and Networks (ICCCN). Maui, HI, USA; 2011.

[11] Yang M, Li Y, Jin D, Su L, Ma S, Zeng L. OpenRAN: A software-defined RAN architecture via virtualization. In: Proceedings of the ACM SIGCOMM 2013 Conference. Hong Kong, China; 2013.

[12] Jarschel M, Oechsner S, Schlosser D, Pries R, Goll S, Tran-Gia P. Modeling and performance evaluation of an OpenFlow architecture. In: Proceedings of the 23rd International Teletraffic Congress (ITC); 2011, 1–7.

[13] Bianco A, Birke R, Giraudo L, Palacin M. OpenFlow switching: Data plane performance. In: Proceedings of the IEEE International Conference on Communications (ICC). Cape Town, South Africa; 2010.

[14] Calhoun P, Montemurro M, Stanley D. Control And Provisioning of Wireless Access Points (CAPWAP) Protocol Specification; 2009. RFC 5415.

[15] Yap KK, Katti S, Parulkar G, McKeown N. Delivering capacity for the mobile internet by stitching together networks. In: Proceedings of the 2010 ACM Workshop on Wireless of the Students, by the Students, for the Students (S3'10). Chicago, IL, USA; 2010.

[16] Hasan S, Ben-David Y, Scott C, Brewer E, Shenker S. Enhancing rural connectivity with software defined networks. In: Proceedings of the 3rd ACM Symposium on Computing for Development (ACM DEV'13). Bangalore, India; 2013.

[17] Johnson CW. Analysing the causes of the Italian and Swiss blackout, 28th September 2003. In: Proceedings of the 12th Australian Workshop on Safety Critical Systems and Software and Safety-Related Programmable Systems. Adelaide, Australia; 2007.