

OpenSDWN: Programmatic Control over Home and Enterprise WiFi

Julius Schulz-Zander
TU Berlin
Berlin, Germany
julius@inet.tu-berlin.de

Carlos Mayer
TU Berlin
Berlin, Germany
carlosnmayer@gmail.com

Bogdan Ciobotaru
TU Berlin
Berlin, Germany
bogdan.ciobotaru1@gmail.com

Stefan Schmid
Deutsche Telekom Innovation
Laboratories / TU Berlin
Berlin, Germany
stefan@inet.tu-berlin.de

Anja Feldmann
TU Berlin
Berlin, Germany
anja@inet.tu-berlin.de

ABSTRACT

The quickly growing demand for wireless networks and the numerous application-specific requirements stand in stark contrast to today's inflexible management and operation of WiFi networks. In this paper, we present and evaluate OPENSNDWN, a novel WiFi architecture based on an SDN/NFV approach. OPENSNDWN exploits datapath programmability to enable service differentiation and fine-grained transmission control, facilitating the prioritization of critical applications. OPENSNDWN implements per-client virtual access points and per-client virtual middleboxes, to render network functions more flexible and support mobility and seamless migration. OPENSNDWN can also be used to out-source the control over the home network to a participatory interface or to an Internet Service Provider.

Categories and Subject Descriptors

C.2.3 [Network Operations]: Network management; C.2.1 [Network Architecture and Design]: Wireless Communication

Keywords

Software-Defined Wireless Networking, Software-Defined Networking, Network Function Virtualization, WLAN, Enterprise

1. INTRODUCTION

The popularity of WiFi networks is increasing at a fast pace, with more and more mobile end-devices becoming WiFi enabled. Today, many hotels and cafés—and sometimes also entire cities—offer free WiFi services. Several mobile operators also plan massive WiFi HotSpot as well as HotSpot 2.0 deployments for traffic offloading from cellular and future Internet-of-Things networks [11, 44].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or to publish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

SOSR2015 June 17 - 18, 2015, Santa Clara, CA, USA

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-3451-8/15/06 ...\$15.00.

DOI: <http://dx.doi.org/10.1145/2774993.2775002>.

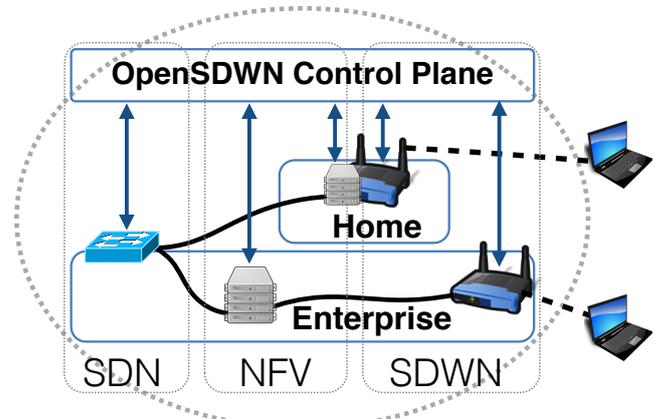


Figure 1: OPENSNDWN introduces programmability in home and enterprise WiFi networks using an SDN and NFV approach.

The increasing demand for WiFi networks imposes new requirements, e.g., on security, optimized medium utilization, and mobility support. The (last) wireless hop is often critical for network performance, as it can contribute a non-negligible delay and may constitute a bandwidth bottleneck.

These requirements stand in stark contrast to the state-of-the-art: The management and operation of off-the-shelf WiFi networks is often very inflexible, and today's networks largely ignore the specific needs of users and/or applications. Moreover, WiFi networks are often deployed in an unplanned and uncoordinated manner: different parties in a house or neighborhood typically deploy and run their own dedicated infrastructure; neighboring access points as well as public access points cannot be leveraged—but rather interfere with each other, introducing unnecessary transmission delays, and reducing network capacity. Also mobility support is often very limited, depriving users from essential services.

Software-Defined Networking is an interesting new paradigm which allows overcoming network ossification by introducing programmability. In a nutshell, Software-Defined Networks (SDNs) consolidate and outsource the control over a set of network devices to a logically centralized software controller. The decoupling of the data plane and control plane allows the control plane to evolve independently of the data plane, enabling faster innovations. Moreover, OpenFlow, the standard SDN protocol today, introduces interesting

generalizations. Openflow is based on a match-action paradigm, where switches can match not only the Layer-2 header fields of packets, but also Layer-3 and Layer-4 fields. These flexibilities can be used, e.g., to implement fine-grained traffic engineering [16], enforce complex network policies [20, 39], improve resource utilization in wide-area networks [21, 23], or enable network virtualization in datacenters [15].

SDN is also an enabler for a second paradigm shift in the Internet: *Network Functions Virtualization (NFV)*. Modern networks include many middleboxes to provide a wide range of network functions to improve performance as well as security. For example, middleboxes are used for caching and load-balancing, as well as for intrusion detection. NFV aims to *virtualize* these network functions, and replace dedicated network function hardware with software applications running on generic compute resources. The resulting orchestration flexibilities can be exploited for a faster and cheaper service deployment. SDN can be exploited to steer flows through the appropriate network functions [4, 18, 31, 39]. Thus, SDN and NFV together, recently also called *SDNv2* in the context of carrier WAN networks, support fine grained service level agreements, as well as an accurate monitoring and manipulation of network traffic.

In this paper, we argue that there is a major potential of introducing programmability and virtualization in *wireless networks*, i.e., following a Software-Defined Wireless Networking (SDWN) approach. Wireless networks are very different from wired networks—the domain where SDN/NFV has been studied most intensively so far. In wireless networks communication happens over a shared medium whose characteristics can change quickly over time and in an unpredictable manner, as users are often mobile and associations dynamic. WiFi networks offer several unique knobs to influence the probability of successful transmissions, such as transmission rate and power, as well as retry chains. This introduces opportunities for a fine-grained and application specific transmission control, e.g., for service differentiation.

Today’s OpenFlow protocol is not well suited for WiFi: it is restricted to programming flow table rules on Ethernet-based switches, and it is not possible to match on wireless frames, nor can measurements of the wireless medium be accommodated or per-frame receiver side statistics reported; it is also not possible to set per-frame or per-flow transmission settings for the WiFi datapath. In general, SDN+NFV have not received as much attention yet in the context of wireless

1.1 Our Contribution

This paper shows how to reap the benefits of SDN and NFV in home and enterprise WiFi networks. In particular, we present the design, implementation, and evaluation of OPENSDWN, a flexible WiFi architecture based on a unified, programmable control plane as illustrated in Figure 1. OPENSDWN allows to manage both the virtualized middleboxes as well as the wired and wireless datapath, e.g., to apply per-flow PHY and MAC layer transmission settings.

OPENSDWN comes with interesting use cases: (1) It enables service differentiation, and allows administrators or users to specify application and flow priorities on the wired and wireless portion of the network. These priorities are implemented using a fine-grained wireless transmission control. (2) Using its per-client virtual access points and virtual middleboxes, OPENSDWN supports seamless user mobility, as well as flexible function allocation (e.g., function collocation at night to save energy). (3) Network functions such as firewalls and NATs can be deployed flexibly, e.g., outside user premises. (4) OPENSDWN also introduces flexibilities in terms of network control: the system exposes a participatory interface à

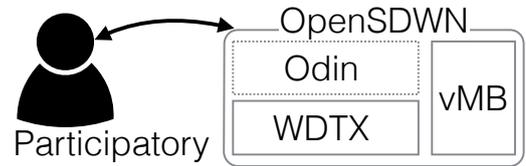


Figure 2: OPENSDWN extends Odin with WiFi datapath programmability (WDTX), a unified abstraction for virtualized middleboxes and access points, and a participatory interface.

la [17], through which users can indicate priorities for their applications. The control can also be outsourced to an Internet Service Provider (ISP), e.g., for troubleshooting.

OPENSDWN leverages the LVAP abstraction and extends Odin [42] by (see Figure 2): (1) WiFi datapath programmability, e.g., for fine-grained wireless datapath transmission control (WDTX): settings include transmission power, transmission rate as well as tailored retry chains. (2) A unified SDN and NFV abstraction through virtualized middleboxes and access points, e.g., to facilitate an easy handling and migration of per-client state. (3) A participatory interface which allows to share network control.

Indeed, middleboxes are an integral part of OPENSDWN. First, to abstract and decouple user-specific state, OPENSDWN introduces the notion of per-client virtual middleboxes (*MBs*). Second, to identify and classify flows, and hence enable service-differentiation, OPENSDWN relies on a *Bro Intrusion Detection System (IDS)* [30]. Once flows have been detected, per-flow transmission rules are installed according to specific requirements such as policies specified by the users. Bro may also be used to tag packets, e.g., for a live streaming application where key frames should be transmitted in a prioritized way, as these frames are more critical for service quality.

We demonstrate the feasibility and usefulness of our system by reporting on different case studies and experiments conducted using two deployments: one at our university and one in a large home network.

Paper Scope. We understand our system as an *enabler* of more flexible WiFi networks. How to optimally *exploit* the resulting flexibilities (e.g., in order to provide QoS guarantees) or how to fine-tune performance (e.g., of function migration), are orthogonal questions, and are left for future research.

1.2 Paper Organization

The remainder of this paper is organized as follows. Section 2 gives an overview of the goals and benefits of OPENSDWN. Section 3 presents the architecture of OPENSDWN, and Section 4 reports on our deployments and experiments. Section 5 discusses the prototype implementation. After reviewing related literature in Section 6, we conclude our work in Section 7.

2. USE CASES AND OVERVIEW

OPENSDWN is based on programmable network devices in the spirit of SDN and NFV. Before we give an overview of the architecture, we discuss some use cases for the envisioned system. See also Figure 3 for some illustrations.

2.1 Use Cases

1. **Service differentiation:** OPENSDWN offers visibility into the network’s state and supports a fine-grained transmission control, by allowing administrators and users to set per-flow



(a) **Mobility support:** Virtual middleboxes (e.g., encapsulating firewall connection state) can be migrated in the presence of mobility and cloned for redundancy.

(b) **Transmission control:** The controller sets specific wireless transmission and OpenFlow rules for per-flow wireless transmission control.

(c) **Participatory interface:** A participatory application provides an interface to the user. Service detection is achieved through DPI.

Figure 3: Three basic operations supported by OPENSDown.

and per-packet specific transmission settings (such as transmission rate, power, retransmission and RTC/CTS strategy). For instance, as we will demonstrate, OPENSDown can protect latency-sensitive flows (e.g., live media streams) from competing with background traffic (e.g., Dropbox synchronization).

2. **Mobility and migration:** By virtualizing not only the per-client access points, but also the middleboxes, OPENSDown supports both seamless user mobility and dynamic resource allocation. The more dynamic resource management introduced by OPENSDown enables the adjustment and migration of resources and functionality with the user, e.g., for flexibly scaling up or down resources depending on the demand. By collocating network functions, e.g., at night, also energy may be saved.
3. **Flexible deployment:** Network functions (firewalls, NATs, functionality for service differentiation) can be allocated and deployed flexibly. For instance, the different users of a house may use a shared box, outside their individual user premises, to run a middlebox or controller. The specific deployment requirements will depend on the scenario (fiber-to-the-home, endpoints of encryption tunnels, etc.).
4. **Flexible control and participatory networking:** OPENSDown provides unified programmability and control over the network devices and middleboxes. It also offers customization flexibilities through a participatory interface [17]: the interface can be used by the users to specify priorities over different applications (e.g., Youtube over Dropbox), and the control may also be handed over to an Internet Service Provider (ISP) for troubleshooting or for defining requirements and updating transmission rules. A local controller can also maintain connectivity between users in a neighborhood during network failures on the uplink.

2.2 Overview

OPENSDown is based on an SDN+NFV (a.k.a. *SDNv2*) approach and consists of the following components:

1. **Unified Programmability and Abstractions:** The logically centralized control plane unifies SDN and NFV through programmatic abstractions. That is, OPENSDown virtualizes both access points and virtualized middleboxes (see Figure 3(a)), which facilitates an easy handling and migration of per-client state, also beyond CPE boundaries. The OPENSDown abstractions can be seen as an extension of

Odin [42] to NFV: Odin's *LVAP* concept abstracts the complexities of the IEEE 802.11 protocol stack (client associations, authentication, and handovers), and enables the unified slicing of both the wired and wireless portions of the network. The former is achieved by encapsulating the client's Openflow state. OPENSDown additionally introduces per-client virtual middleboxes, short *vMBs*, which can be transferred seamlessly across the network. Specifically, a *vMB* encapsulates the client's *MB* state as a virtual *MB* object. Thus, OPENSDown achieves control logic isolation as SDN/NFV applications running on top the controller can only operate on their respective *LVAPs* and *vMBs*.

2. **Programmable Datapath:** The programmable datapath gives the possibility to set per-flow specific transmission settings as shown in Figure 3(b). The settings include transmission power, transmission rate as well as tailored retry chains. It is even possible to differentiate between different packets of the *same flow* (5-tuple): for instance, key frames of a live stream may be given higher priority. This is achieved by using an Intrusion Detection System (*IDS*, in our case: *Bro*) for packet classification and *tagging*: transmission settings are chosen depending on the tag.
3. **Participatory Interface:** OPENSDown's participatory interface allows us to define flow priorities as well as priorities over customers. The chosen priorities are translated by the controller into meaningful network policies. Priorities can be adjusted anytime. Figure 3(c) depicts the participatory interface.

3. THE OPENSDown SYSTEM

We first describe the wireless SDN component of OPENSDown, then the virtual middlebox, and finally the participatory interface.

3.1 Wireless SDN

WiFi networks have several unique properties which do not exist in wired networks. For instance, WiFi networks offer several knobs to influence the probability of successful transmissions, such as transmission power or rate. This introduces opportunities for a fine-grained and application specific transmission control.

The wireless subcomponent of OPENSDown builds upon Odin [42], from which OPENSDown also inherits: (1) The Light Virtual Access Point (*LVAP*) abstraction: essentially the client's association state (the BSSID, SSIDs, client IP address, and OpenFlow rules). (2) Mobility support: by migrating a client's *LVAP*

between physical APs, the infrastructure can control the client's attachment point to the network, without triggering a re-association at the client. (3) Slicing: the accommodation of multiple logical networks on top of the same physical infrastructure with different policies and control applications. A network slice is a virtual network with a specific set of SSIDs, where for example, the traffic may be VLAN tagged or directed to a specific destination port.

OPENDSDWN introduces service differentiation through per-flow WiFi datapath transmission rules, organized into per-flow transmission rule tables. Rules are bound to one or more OpenFlow rules and assign meta or direct transmission properties to one or more OpenFlow entries. Specifically, fine-grained wireless transmission control is achieved by combining Openflow match-action rules with *wireless transmission rules* (WDTX) within the wireless access points. Regarding actions, assigning fixed and/or meta transmission settings is possible. Meta transmission settings include: *best probability rate*, *best throughput rate*, *second best throughput rate*, *common maximum rate* or *fixed rates* (e.g., a basic rate or a specific modulation and coding scheme rate). Based on the capabilities of the WiFi NIC, the transmission settings can be set for the device multirate retry chains.

Furthermore, in order to account for the dynamic nature of the wireless network and in order to support client mobility, agents in OPENDSDWN implement a publish/subscribe interface, allowing the controller to subscribe to network events (see Section 5 for more details).

3.2 Virtual Middleboxes

Middleboxes are an integral part of OPENDSDWN. First, our service differentiation mechanism relies on a deep-packet inspection middlebox, to identify and classify flows. Moreover, OPENDSDWN integrates *MBs* in the virtual network, and allows us to set and migrate state to support client mobility and to scale dynamically.

At the core of our system lies the concept of virtual *MBs*, short *vMBs*. *vMBs* are used to fully reap the virtualization benefits: the handling of *vMBs* is important to guarantee the decoupling of the per-client middlebox state and the inner workings of the middlebox from the physical instance.

The *vMB* keeps user-specific state information and can be transferred from one *MB* instance to another. On top of a physical *MB* runs a *MB agent* which needs to accomplish three primary tasks: (i) interface with the physical resources of the *MB*, (ii) handle *vMBs* and (iii) expose the control of the *MB* to a remote entity (the controller). The middlebox agent also provides the necessary hooks for the controller (and thus applications) to instantiate, destroy, monitor and manage its functionality.

In OPENDSDWN, a stateful *vMB* is characterized by a configuration file (a *MB*-specific list of tunable parameters), the state of the active connections, the statistics (counters) and a list of subscribed events in order to completely define its behavior. When a *vMB* is moved from one *MB* Agent to another, the new *MB* is able to handle the user's traffic in the exactly same way the old one. *vMBs* were designed to give applications the possibility to manage user related *MB* state across physical *MBs*, without any awareness of the user's traffic.

In order to support e.g., scale-out upon certain network events, or to monitor the middlebox, OPENDSDWN implements a publish/subscribe interface (see Section 5).

3.3 Participatory Interface

OPENDSDWN's participatory interface allows the WiFi users, the network provider or even the content provider, to express their preferences in terms of flow differentiation. Specifically, we allow external entities to rank—by assigning priorities—their transmissions. The rationale behind this prioritization approach is simplicity: the participatory interface hides network complexity from end-users. Concretely, a user could express his or her preference to prioritize Netflix over Dropbox, by assigning a higher priority to the former. This preference will then be taken into account by the controller, which installs transmission rules which favor flows tagged as Netflix over flow tagged as Dropbox. This could be done, for example, by assigning different AC Queues or setting distinct rate chains.

As a static service mapping based on, e.g., content server IPs is cumbersome and unreliable, OPENDSDWN uses a signature-based Intrusion Detection System (IDS) which also considers packet payload. Once the IDS detects a service of interest, it immediately informs the OPENDSDWN controller, which applies the necessary policies accordingly.

In order to keep the system evolvable, and to account for the advent of new services, our participatory API also supports the installation of new signatures by external applications. This for example also enables content providers to install their own signatures, ensuring a better probability of correctness.

Technically, the participatory interface can be implemented based on a *URI* included in a HTTP GET request, or a domain name within a certificate.

4. EVALUATION

The key benefit of OPENDSDWN is its flexibility and the potential use cases it *enables*. How to optimally *exploit* the resulting flexibilities (e.g., in order to provide QoS guarantees) or how to fine-tune performance (e.g., of function migration), are orthogonal questions, and also depend on the context.

Nevertheless, in order to show the potential of OPENDSDWN, we implemented and evaluated different applications using our proof-of-concept prototype. The first case study focuses on the system's service differentiation capabilities, and in particular, we consider the optimization of a video-on-demand application. In the second case study, we consider an optimized multicast service based on direct multicasting. The third focuses on the middlebox virtualization, and we discuss the migration of a personalized stateful firewall.

4.1 Deployments and Methodology

Our proof-of-concept implementation of OPENDSDWN has been deployed in two real networks:

- Our research group's indoor WiFi network. This deployment consists of more than 25 IEEE 802.11n enabled APs, distributed across one floor of an office building.
- A centrally administrated home network which covers an entire building of ~21500 square feet. It provides internet connectivity for roughly 30 households with more than 70 active devices per day, using Ethernet and 10 WiFi APs (indoor and outdoor).

All APs run OpenWrt release Chaos Calmer with the ath9k Linux driver, user-level Click modular router [12], and Open vSwitch (OvS) version 2.3.90 supporting OpenFlow (OF) version 1.3 and *conntrack* table management. The off-the-shelf WiFi access points are either based on ARM, MIPS or x86. The variety

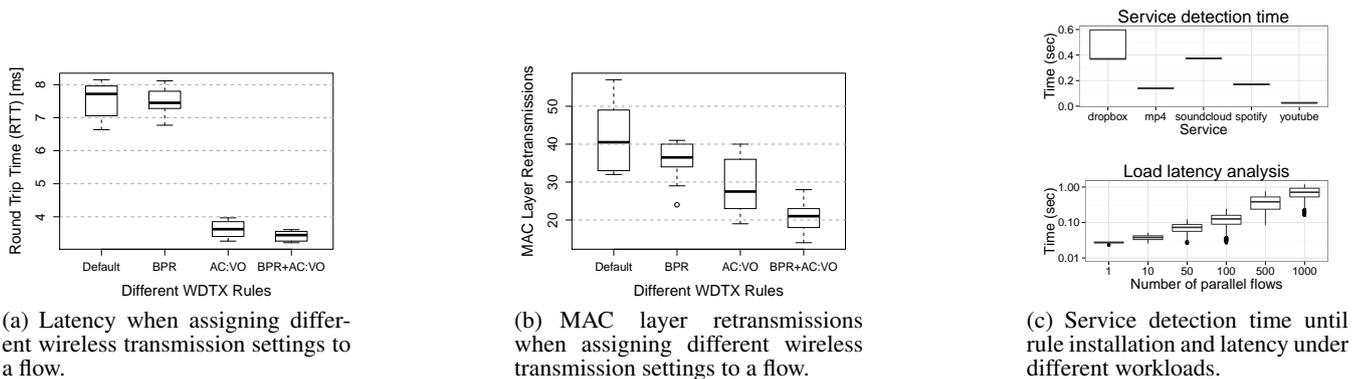


Figure 4: Evaluation of OPENSDown’s fine grained transmission control through *WDTX* rules and *vMB* handling.

of WiFi AP hardware ranges from IEEE 802.11g only to IEEE 802.11abgn boards equipped with one or more WiFi NICs based on Atheros chipsets.

Our controller and *MBs* are evaluated on non-virtualized servers with 4 CPU cores supporting hyper-threading and at least 8 GB RAM. All servers run a Debian-based OS with OvS 2.0.2 or 2.3.90. We monitor data through a dedicated port for the IDS at the core switch. We did not hit CPU or memory limitations in any of our experiments. Furthermore, for the performance evaluation of the controller and middleboxes, we use three dedicated servers: 1) an OpenFlow controller, 2) a middlebox, and 3) a traffic generator.

4.2 User-Defined Service Differentiation

The first case study concerns OPENSDown’s service differentiation capabilities. Before presenting our video-on-demand optimizer in more detail, we will discuss some more general aspects of our system.

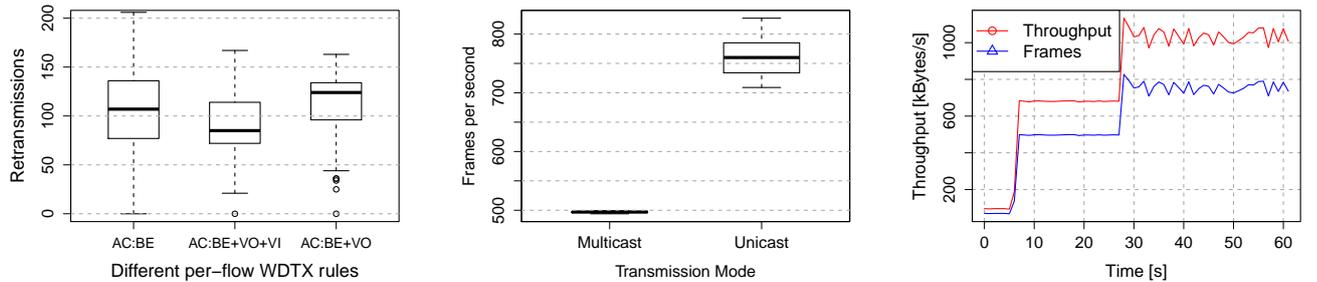
Today, most public internet downlink traffic is sent as *best effort*, also due to *network neutrality* requirements. But also in small offices, home offices or home networks, traffic is often treated equally, although this is legally not required. We believe that there is a high potential benefit of differentiating services in home networks, e.g., by prioritizing voice traffic over regular web traffic. Especially given today’s trend to deploy more and more wireless devices in the user’s premises, traffic can significantly interfere, e.g., an unimportant system update for a device can easily interfere with requested *on demand* services such as Spotify or Netflix, resulting in poor performance.

Benchmarking the Transmission Rule Extension: There are several ways to prioritize traffic through specific *WDTX* rules, bound to a particular flow entry. We investigate, as a benchmark, the effect of assigning a meta transmission rate and a medium access priority, on the latency and MAC layer retransmissions of a single flow. To this end, we first study the effect on MAC layer retransmissions (*cf.* Figure 4(b)) when assigning a per-flow transmission rule to a latency sensitive UDP flow. In our experiment, we use two OPENSDown APs and two clients. Each client is connected to one of the APs in our indoor testbed. We start generating best effort TCP traffic on the link between one AP and client, and start a latency sensitive flow on the link between the other client and the AP. In the beginning, the latency sensitive flow and the background traffic are treated equally, which results in a round trip time (RTT) of roughly 8 *ms*. Next we assign the *best probability rate* (BPR) to the flow; this leaves the RTT unchanged. When changing the medium access to the highest priority (AC:VO), i.e., the voice access category, the RTT drops by half to less than 4 *ms* as depicted

in Figure 4(a). This is as expected since a higher medium access probability constitutes a change in the RTT. Note, in today’s home network traffic is typically sent as *best effort* and rarely differentiated as in OPENSDown. However, only looking at the RTT of an UDP flow is not sufficient as it does not take the MAC-layer (L2) packet loss into account; this however has a significant effect on the jitter and performance of transport protocols (L4) such as TCP. Thus, we next study the effect of the meta transmission rates on the packet loss. We assign a *WDTX* entry to the *OF* flow rule that matches the flow, and assign the *best probability rate* and highest medium access priority (AC:VO) which increases the transmission probability on the L2. Figure 4(b) shows that this significantly reduces the MAC layer retransmissions compared to default flow properties. We conclude that combining the medium access strategy by a meta transmission rate within OPENSDown significantly reduces the number of MAC layer retransmissions, and the the RTT by roughly 50%. That said, OPENSDown can achieve a per-flow resource utilization which is better suited for the diversity of traffic requirements in today’s home networks.

Benchmarking the DPI Interface: In order to understand latency and induced load of service discovery, we replay traces of typical streaming services collected at a university campus network in our testbed. Concretely, we replay the traces 100 times per service at first and then vary the number of simultaneous *youtube* flows: 1, 10, 50, 100, 500, 1000 to identify eventual bottlenecks on the service detection engine. The traffic is injected on one server and tapped on a second server running a Bro *MB* instance handled by our agent. The controller is hosted on a third server with a dedicated out-of-band control channel running a service discovery SDN/NFV application. Figure 4(c) depicts the measured service detection time and the load latency analysis, i.e., the latency added during high workload pattern.

In order to estimate possible performance bottlenecks of the service detection chain, we measure the delay added by the different components involved in the detection, in bursty scenarios. We are interested in how our system reacts to different rates of events. We mock the detection of a service by triggering an event from Bro at different intervals. Specifically, we schedule events sequentially, from a Bro script, adding a determined delay between 2 consecutive events. We send 300 events in total over multiple runs for each delay, starting from 2 μ s up to 1 second. We run this procedure in two different scenarios: First, we keep both the controller and the MB Agent on the same host to eliminate the network delay. In the second scenario, we run the MB Agent and controller on a different host. Table 1 presents the results. They include, for



(a) Flows sent as BE do not experience significantly more retransmissions in the presence of higher prioritized traffic.

(b) HD IPTV stream exceeds capacity of multicast bandwidth. Switching to unicast mitigates this problem.

(c) OPENSDown’s multicast Application switches to unicast for a single subscriber.

Figure 5: Evaluation of the service differentiation and smart multicast OPENSDown applications

Table 1: Service Detection Delay

Frequency (ms)		2us	5us	10us	50us	100us	500us	1ms	5ms	10ms	50ms	100ms	1s
Same host	Bro - MB Agent	0.060651	0.058026	0.05992	0.058391	0.064939	0.049412	0.030326	0.009835	0.003022	0.000381	0.000371	0.000421
	Bro - Rule Installation	0.06167	0.0589	0.060837	0.059249	0.063867	0.050346	0.031277	0.011736	0.005432	0.002918	0.002871	0.003129
Different hosts	Bro - MB Agent	0.057146	0.063181	0.062309	0.064495	0.055805	0.031668	0.021651	0.001527	0.000388	0.000389	0.000394	0.000425
	Bro - Rule Installation	0.058662	0.064439	0.069839	0.06587	0.058102	0.033013	0.023731	0.003773	0.003733	0.003183	0.00362	0.003861

each scenario, two distinct measurements: First, we measure the mean of the delay between the instant Bro sends the event and the MB Agent processes it (basically, the delay caused by the queue of events). Second, we measure the time between the instant Bro fires the event and the moment that our controller installs the required rules for this flow. This delay includes therefore the MB Agent processing time, the delay caused by the MB Protocol and the controller handling time of this event. As expected, our system presents a lower response time for smaller event rates (bigger delays). The worst performance, for the highest rate, indicates a total delay of around 60 ms.

Case Study: Medium Access Optimizer. Given these benchmarks, we now consider a simple case study: the optimization of a video-on-demand transmission. Our setup consists of a single AP and three clients. One client performs a system update, one requests a Video-on-Demand (VoD) stream and the third client does a UDP-based VoIP call. In the beginning, all flows are treated equally as best effort traffic. Next we put the voice flow into the highest priority queue. As expected, the prioritized traffic now achieves a slightly higher throughput than the best effort traffic. However, in mac80211, the voice queue does not perform aggregation and hence, can easily suffer from too many competing stations. Specifically, even if the medium access probability is high, the performance without 802.11 frame aggregation is significantly lower. That said, if a flow suffers from background traffic, e.g., caused by a neighboring WiFi network, switching to the highest queue with aggregation can significantly increase the throughput. Due to the bursty nature of DASH based VoD traffic, the BE traffic is just slightly decreased while VoD services benefit from a more aggressive medium access, which in turn leads to a faster switching of the video quality. BE flows do not experience significantly more retransmissions in the presence of higher prioritized traffic. In other words, using prioritization reduces the achievable throughput of BE flows without a big impact on the the MAC layer retransmissions (see Figure 5(a)).

4.3 Smart Direct Multicast Service

OPENSDown can also be used in conjunction with group communication abstractions such as multicast. Especially with the ad-

vent of IPv6, the fraction of multicast traffic is likely to grow in the future: IPv6 realizes broadcast over multicast, and mDNS to broadcast features to neighboring stations.

In IEEE 802.11, multicast packets are typically sent at basic rate. However, wireless networks may benefit from a Direct Multicast Service (DMS): DMS has the potential of reducing the transmission time over regular multicast, by sending 802.11 packets as unicast. Unfortunately, DMS requires a client to signal its DMS capabilities to the AP, which is the reason why DMS is rarely used in 802.11 networks today.

With OPENSDown, a controller can detect the number of subscriptions for a particular multicast service, and control the transmission accordingly. Specifically, a controller can install an OpenFlow rule to switch from multicast to unicast for the transmission. Moreover, OPENSDown allows to assign a WDTX transmission rule to a particular stream of multicast data, to send the data at the *maximum common transmission rate* for a group of wireless devices.

We evaluate OPENSDown’s smart multicast application with a single access point and a IPTV set-top-box from a major European ISP. First, we transmit a IPTV continuous stream of multicast data to the box. With a single station, our application installs a rule to send the multicast stream as unicast on the wireless medium. With multiple stations, the application switches back to multicast at the maximum common rate for the transmission. Figure 5(c) shows that the throughput and frame count increase after 28 seconds, when the application switches from multicast to unicast. Figure 5(b) indicates that an HD IPTV stream easily exceeds the basic rate of IEEE 802.11g networks. Note, switching to unicast or to a higher datarate mitigates this issue.

4.4 User Mobility

As a second case study, we consider OPENSDown’s support for user mobility, where also middlebox functionality is migrated. Supporting client mobility is a crucial feature in WiFi deployments with multiple physical APs. The application migrates a stateful firewall vMB object between MBs, i.e., installs the client’s flow state at the new AP before or during the handoff.

Benchmarking the Stateful Firewall vMB Interface: We study the performance of the vMB stateful firewall module of

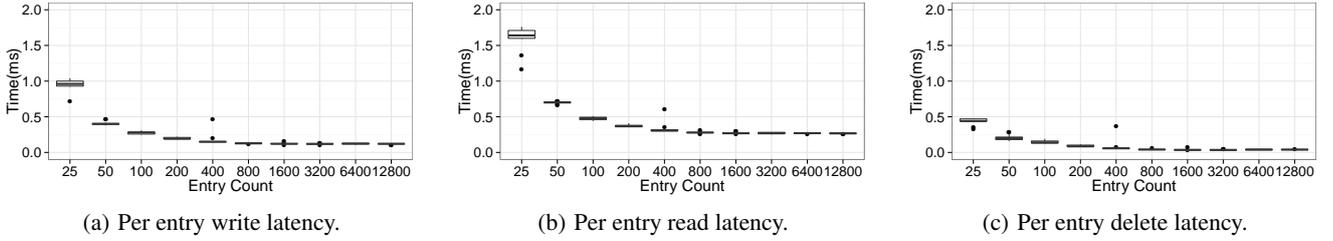


Figure 6: Latency for a stateful firewall vMB object read, write and delete operation. Latency in milliseconds (time) is normalized to a per-entry time. vMB object size is increased from 25 entries to 12,800 entries.

Algorithm 1: Mobility Service

```

begin
  if handoverEvent = True then
    oldMbid ← AP2MMap.get(oldAPid);
    newMbid ← AP2MMap.get(newAPid);
    vMB ← createvMB(clientIP, oldMbid);
    if vMB.migrate(newMbid) = True then
      signalOdin(migrationComplete);

```

OPENSDown for different workloads in more detail. Specifically, we measure the read, write and delete performance of a stateful FW vMB extension that utilizes the *netlink* interface of the Linux Kernel *conntrack* module for connection tracking, which is typically part of a stateful firewall. We repeat each experiment 12 times for each workload. The vMB object workloads vary from 25 up to 12,800 entries. As shown in Figure 6(a), we first measure the performance of the per-entry execution time of the write (*setState*). The write duration for a single entry decreases constantly with the workload, and stabilizes at around $130 \mu s$ for a single entry in a vMB object. Next, we evaluate the read time (*getState*) which decreases constantly. The average value stabilizes at around $270 \mu s$ (see Figure 6(b)). Finally, we evaluate the delete operation in order to fully understand the required time for the migrate operation, which requires a read, write and a delete of the old vMB object. The average value of a *delState* stabilizes at around $40 \mu s$ (see Figure 6(c)). That said, a migrate operation takes at least the time of a combined read and write, times the number of entries. Thus, the time can be estimated by the measured results. Specifically, the delete of the old vMB state can be called after the object was correctly fetched and while it is installed into the new MB.

Case Study: Firewall State Migration The firewall state migration service is a reactive application triggered through external events to move state between MB instances. The algorithm in form of pseudo-code is shown in Algorithm 1. For example, when Odin detects a client with a higher RSSI at a new AP, a handover event is generated and the client’s firewall state migrated to the AP before the handover. The firewall state migration service then decides whether the state associated with the mobile user needs to be migrated and executes the operation. The application keeps a mapping between APs and firewalls. If the client is moving over to an AP that corresponds to a different stateful firewall than the current, a migration of the client’s connection tracking state is performed.

During the state migration operation, the controller uses the three operations that were evaluated previously. The *getState* call on the serving middlebox is followed by a *setState* operation with the target MB identifier as argument. Finally, the state is removed through a *delState* call. The last two operations are vir-

Entry count	Mean execution time (ms)			
	Write	Read	Delete	Migrate
1	11.6	38.4	6.4	45.0
10	12.3	48.6	6.8	60.9
100	20.3	121.6	10.7	141.9
1000	115.9	778.0	43.0	893.9
10000	1119.3	5201.2	385.3	6320.5

Table 2: Average execution time of the *setState*, *getState* and *delState* operations for different workloads.

tually simultaneous because RPC method calls are asynchronous, and called at different agents. Table 2 shows the measured average migration time for different vMB object sizes. The total time of a *migrate()* call on a vMB object with 100 entries averages at around $140 ms$. This underlines the potential power that the simplicity of the vMB abstraction exposes to a network programmer. Note, the agent to kernel communication for a single rule is below one millisecond. The RPC interface and entry processing from the Linux Kernel *netlink* interface contribute the most to the processing time.

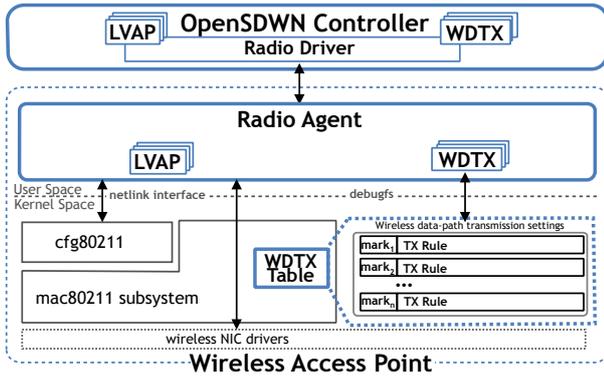
5. PROTOTYPE IMPLEMENTATION

This section presents more details about our prototype implementation. We first describe the different radio and middlebox interfaces implemented by OPENSDown, then present the control plane, and finally discuss the support for reactive and proactive applications. The Radio Agent is implemented in C/C++ while the controller is based on the Java-based Floodlight *OF* controller. The MB agent is realized in python and implements a newly defined MB protocol.

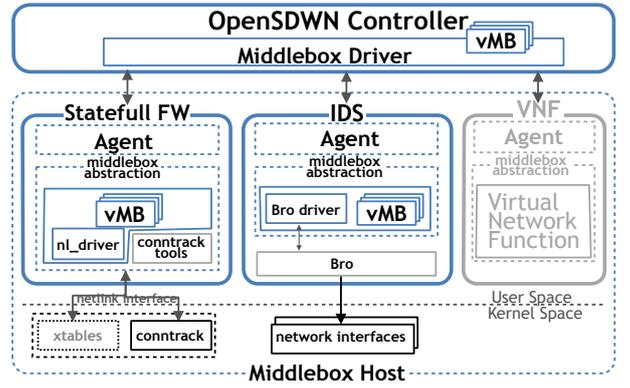
5.1 Interfaces

Interfaces to the physical WiFi and middlebox resources are provided by agents. We describe the radio and middlebox interfaces in turn. Moreover, Table 3 depicts the south-bound interface between the agents and the controller.

Radio Interface: OPENSDown’s wireless APs run a radio agent which exposes the necessary hooks for the controller (and thus applications) to orchestrate the WiFi network and report measurements. All time-critical aspects of the WiFi MAC protocol (such as IEEE 802.11 acknowledgments) continue to be performed by the WiFi NIC’s hardware. On the other hand, non time-critical functionality including management of client associations, is implemented in software on the controller and the agents. Specifically, this realizes a distributed WiFi split-MAC architecture. In addition, matching on incoming frames is performed to support a publish-subscribe system wherein network applications can subscribe to per-frame events.



(a) Packets matching an OpenFlow rule are annotated with a *mark* and then matched by a *WDTX* rule to control wireless transmission settings on a per-flow level.



(b) vMB agent structure: vMB protocol interpreter and state machine; MB interface with specific handlers for different types of middleboxes.

Figure 7: Architecture of the Wireless Radio Interface and Middlebox Interface components.

Table 3: Subset of South-bound APIs provided by the framework

Radio API: (Controller to agent)	Description
{add/remove/set}-lvap	Add/remove/update an LVAP on an agent
read-lvap-table	Obtain the list of LVAPs on an agent
read-rx-stats	Query per-station rx-stats at the agent
{read/set}-wdtx	Query/set per-flow transmission rules on an agent
{read/set}-subscriptions	Query/set the list of subscriptions at the agent
{read/set}-channel	Query/set the channel the agent listens and transmits on
{read/set}-beacon-interval	Query/set the beacon interval on the agent
Middlebox API: (Controller to Agent)	Description
{get/set}-config	Get/Set configuration of parameters a virtual middlebox
{get/set}-state	Get/Set the state of a virtual middlebox
{get/set}-stats	Get/Set statistics (e.g. packet counter) of a vMBs
getAvailableEvents	Get a list of available events supported by a middlebox.
subscribe/unsubscribe	Un-/subscribe from receiving notifications

In order to realize the fine grained wireless transmission rule interface, we have extended the `mac80211` subsystem of the Linux Kernel. Thus, OPENSNDWN benefits from its driver abstraction and the *minstrel* rate control algorithm of the `mac80211`. *WDTX* rules control per-flow physical layer settings. Assigning fixed and/or meta transmission settings is possible, e.g., assigning fixed MCS transmission rates or *best throughput rate*. Based on the capabilities of the WiFi NIC, the transmission settings can be set for the device’s multirate retry chains. With Atheros cards such as the AR9280, there are four segments for the transmission rate, power and retry count. We are currently investigating the possibility to assign functions such as a maximum common transmission rate for a given set of LVAPs or maximum transmission time to *WDTX* rules. *WDTX* rules are bound to *OF* rules through a newly defined action that attaches a tag to all packets that match an *OF* flow entry at the ingress port. The defined tags are passed through the Linux kernel down to the WiFi driver. Figure 7(a) depicts OPENSNDWN’s *WDTX* interface.

Moreover, for effective control decisions, wireless network applications need access to statistics not only at a per-frame granularity, but also measurements of the medium itself (for instance, to infer interference from non-WiFi devices operating in the same spectrum). Thus, applications can access measurements (e.g., RSSI, *OF* statistics or spectral measurements) from multiple layers, and work either reactively (e.g., trigger-driven) or proactively (e.g., timer-driven).

Middlebox Interface: A middlebox agent (MB Agent) runs either on a server or WiFi AP and accomplishes three primary tasks: interface the physical resources of the middlebox, handle virtual middleboxes (vMB) and expose the control of the middlebox to the

control plane. In OPENSNDWN, each agent handles exactly one MB functionality through the middlebox interface. Figure 7(b) depicts the agent’s structure with its interfaces and abstractions.

We have implemented two interfaces for different types of middleboxes in OPENSNDWN: 1) a stateful firewall and 2) an interface for deep packet inspection. The former targets firewall handling within the Linux Kernel. Moreover, we have implemented two versions of the stateful firewall vMB: 1) the first one uses wrappers of the `iptables` and `contrack` user-space tools and 2) the other one uses the `python-iptables` and `pynetfilter_contrack` libraries to communicate with the Linux kernel netfilter modules. For the latter, we had to extend the libraries to support insertion of new entries to the connection tracking table, and to monitor changes inside the connection tracking table for event generation. Specifically, the latter brings a significant performance improvement: e.g., a state insertion call of 10000 entries is almost 70 times faster over the former interface. However, the former brings advantages for simpler extensibility for non-time critical parts of the firewall handling. The connection tracking table inside the kernel space keeps track of all traffic passing through the firewall in both directions, and represents the internal *traffic-dependent* state. For each connection or flow, the number of bytes and packets sent in each direction is recorded. This serves as the *statistics* state of the middlebox.

Moreover, the SDN control plane needs to react to events such as threats like DoS attacks or load changes within the network. To this end, the MB agent implements a publish/subscribe system together with the controller. Our Bro IDS and stateful firewall abstraction implement an interface to receive events at the controller, e.g., if someone scans the network. In the case of the stateful firewall, events must be generated whenever something changes in the connection tracking table. The agent leverages the `pynetfilter_contrack` API to filter events that match a subscription from the controller. Specifically, the agent offers a list of parameters that can be used to create an event mask. The controller can request this information through the `Event_List_Req` message. For each event mask, the agent creates a filter and an event ID. In this way, the controller can deactivate notifications it is no longer interested in, according to the ID. An `Event` message is sent every time a change occurs in the internal state of the MB.

Table 4: Subset of APIs provided by the framework

North-bound API for Radio	Description
getClients ()	Get slice-specific-view of associated clients
getAgents ()	Get a view of agents in the application's slice
handoffClientToAp ()	Perform an LVAP migration of a client to an AP
getRxStatsFromAgent ()	Query agent for per-station rx-statistics
{register/unregister}Subscription ()	Subscribe to a per-frame event of interest at agents
{add/remove}Network ()	Add or remove an SSID to the application's slice
Northbound API specific for DPIs	Description
{start/stop}DPI ()	Start or stop the DPI daemon running on the agent
{get/set}InterfaceToMonitor ()	Get/Set the network interface the DPI should monitor
{unsub/sub}scribeForService ()	Un/Subscribe for services
{uninstall/install}Service ()	Un/Install the capability of detecting a service
availableInterfaces ()	Get the network interfaces available at MB
getServicesInstalled ()	Get the list of services installed on a DPI instance
isRunning ()	Check whether the service is currently running
getEventTypes ()	Get event types that DPI supports
getFieldsToSearch ()	Get a list of header fields that DPI is able to inspect
Virtual Middlebox Northbound API	Description
migrate_vMB ()	Move a vMB from one physical MB to another
add_vMB ()	Add a vMB to a physical MB
remove_vMB ()	Remove a vMB from a physical MB
clone_vMB ()	Clones a vMB from a physical MB to another

5.2 Control Plane

The OPENSOWN controller exposes a set of interfaces to the applications (the northbound API shown in Table 4) and then translates these calls into a set of commands on the network devices (the southbound API). The controller also maintains a view of the network including clients, APs, MBs and OF switches, which the applications can then control.

Reactive applications can leverage a publish-subscribe system of the radio and MB agent which invokes a handler at the application, whenever an event of interest occurs at the agents. Our current implementation supports applications to register thresholds for events to reduce the amount of events, e.g., receive link-based (PHY and MAC layer) rx-statistics like the receiver signal strength indicator (RSSI), the bit-rate, and the timestamp of the last received packet, only when necessary. That said, an application can ask to be notified whenever a frame is received at a radio agent at an RSSI greater than -70dBm. Moreover, applications can access data from multiple measurement sources outside the framework, too.

Participatory Interface:

The participatory interface is implemented as a RESTful API exposed by an SDN application, that we call *Service Ranking*. The *Service Ranking* is implemented as a simple Web service, written in NodeJS. It receives priorities as input and feeds the controller with requests through the Northbound API. Figure 8 depicts the components associated with the participatory interface.

A dedicated *Traffic Manager* module within the OPENSOWN Controller, is responsible for compiling requests into meaningful transmission rules, namely assigning a QoS class or WDTX rule on a matched flow.

Concretely, the Traffic Manager is responsible to apply network policies to flows, taking into consideration the service being carried by the flow. We define flows as a group of packets that share a 5 header tuple, composed by source and destination IPs, source and destination ports and the transport protocol, following Bro's connection concept. The algorithm is shown in Algorithm 2. During its initialization (not shown by the algorithm), the Traffic Manager subscribes to all DPI events (by sending subscriptions to MB Agents associated to middleboxes whose type is DPI), passing as parameter the callback to handle the upcoming events. If an event occurs, the message is parsed to a flow and its *tag*, which is determined by Bro and indicates the service being carried in the flow's payload. It then builds the Openflow matches that will be used later to set the rules, and checks the storage for the pre-installed policies for the client IP and the specific service (tag). The retrieved policies, which are defined in terms of Openflow actions, are installed

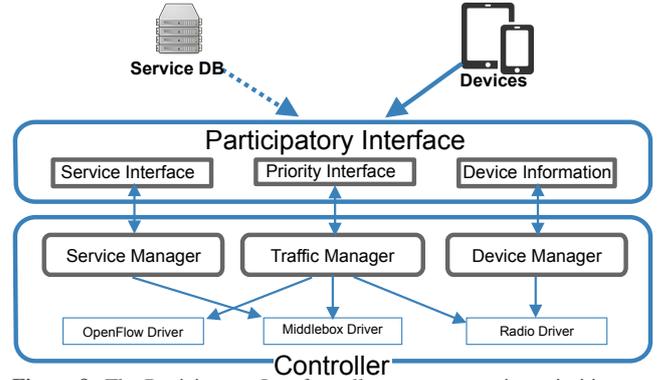


Figure 8: The Participatory Interface allows users to assign priorities to a particular service on a per-device basis.

Algorithm 2: Service Differentiation through DPI

```

begin
  if ServiceDetectEvent = True then
    (5 tuple, service) ← parsevMB(vMBMessage);
    match ← toOFMatch(5 tuple);
    policy ← getNetworkPolicy(ipClient,service);
    action ← buildOFAction(ipClient,markPkt,setTOS);
    lvap ← getLVAP(ipClient);
    if lvap exists then
      physicalAP ← getPhysicalAP(lvap);
      switchAP ← getOFSwitch(physicalAP);
      if physicalAP and switchAP then
        # Install rules at last hop;
        addOpenFlowRule(switchAP, match, action);
        addWDTXRule(lvap, policy);
        # Install Selective Tap at DPI;
        addOpenFlowRule(switchDPI, match, action:drop);

```

in the switch currently responsible for the traffic of the given client. The exact switch can be determined, in case Odin is integrated, by querying which physical access point is hosting the LVAP associated to this client. After the proper Openflow actions are installed, the Traffic Manager installs a rule on the switch attached to the host running Bro, to drop all traffic regarding this flow, in order to make sure Bro does not spend any resource analyzing already tagged flows, hence avoiding unnecessary load on Bro.

An application developer can interact with the *Service Ranking* interface by defining: 1) a Service Name, 2) a Device ID or IP address and 3) a Priority. The former identifies a content provider (e.g., Youtube, Spotify) or a generic application layer service which the OPENSOWN framework system detects through deep packet inspection, e.g., by looking at the SSL certificate, URI or IP address space.

New services can be added by the network operator or public database through the *Service Ranking* interface. With the user participatory interface, the controller exposes a list of detectable services through the northbound API along with a list of IPs of devices connected to a particular network slice. Moreover, the *Service Ranking* interface can also be configured to only expose the list of services to a specific (e.g., connecting) client.

Reactive and Proactive Applications

Network applications written on top of OPENSOWN can function both reactively and/or proactively. Proactive applications are timer-driven whereas reactive applications use triggers and callbacks to handle events. The latter mode of operation is particularly interesting in the context of WiFi networks due where channel qual-

ity can change quickly. To this end, in our current implementation, an application can utilize multiple measurement sources.

Radio agent interface: Reactive applications can use of a publish-subscribe system of the radio agent. The former can register a handler to receive notifications on a per-frame granularity. In our current implementation, applications register thresholds for link-based (PHY and MAC layer) rx-statistics like receiver signal strength indicator (RSSI), bit-rate, and timestamp of the last received packet. For instance, an application can ask to be notified whenever a frame is received at an agent at an RSSI greater than -70dBm . In addition, applications can make use of measurements such as spectral scans or the channel busy time which can be collected by the agents.

Middlebox agent interface Communication over the south-bound interface is realized through the exchange of messages according to the vMB protocol. It functions based on two mechanisms:

- **Request-response model:** A controller interested in the contents of a remote middlebox, can send a request message to the MB's agent and the corresponding action is performed: part of the existing internal state is read and sent back, modified or deleted, or new data is added. This feature is useful for the remote control over the behavior of the machine and represents the *proactive* behavior of the controller.
- **Publish-subscribe model:** The role of the publisher is taken by the agent, where the controller acts as the subscriber. The agent offers a set of events and event parameters to which the controller can register. Because a controller is usually not interested in all event messages that can be sent by the publishing agent, a filter is used for selecting the content or the type of event messages.

OpenFlow statistics: OpenFlow provides per flow and port-based statistics of entries through the switch flow tables. Applications can query these statistics through the controller to make traffic-aware routing decisions.

6. RELATED WORK

While software-defined networking and network virtualization principles have been studied intensively for wired environments, not much is known today about how to reap the corresponding benefits in the wireless and home network context. In general, it is difficult to port systems such as FlowVisor [37] to WiFi networks, and provide, e.g., bandwidth and CPU isolation on the access point.

While there exist a plethora of commercial enterprise WiFi solutions, which typically manage APs centrally via a controller (hosted either in the local network [2], or remotely in the cloud [1]), these solutions do not extend into the purview of cheap low-cost commodity AP hardware that is used by provider networks, nor do they support common, open and programmable interfaces.

OPENSOWN exploits the LVAP abstraction and builds upon Odin [33, 42] and Aeorflux [32], by introducing datapath programmability, network function virtualization and participatory interface. Over the last years, several interesting architectures have been proposed towards a more programmable WiFi, for example Dyson [29], an architecture for extensible wireless LANs which also defines a set of APIs for clients and APs to be managed by a controller. Flashback [10] proposes a control channel technique for WiFi networks, by allowing stations to send short control messages concurrently with data transmissions, without affecting throughput. This ensures a low overhead control plane for WiFi networks that is decoupled from the data plane. BeHop [46] is a programmable wireless testbed for dense WiFi networks as they occur in residential and enterprise settings. Atomix [6] is a modular software framework for building applications on wireless infrastructure

which achieves hardware-like performance by building an 802.11a receiver that operates at high bandwidth and low latency. FlexRadio [8] aims to unify RF chain techniques (MIMO, full-duplex and interference alignment), into a single wireless node, and enables a flexible RF resource allocation. DIRAC [49] proposes a split-architecture wherein link-layer information is relayed by agents running on the APs to a central controller to improve network management decisions. However, the requirement for special software or hardware on the client, and violates the design requirements for OPENSOWN. There are also systems that do not modify the client in order to deliver services. In DenseAP [28], channel assignment and association related decisions are made centrally by taking advantage of a global view of the network. However, slicing is not supported and also client association management is limited. Also Centaur [38] seeks to improve the datapath in enterprise WiFi networks by using centralization to mitigate hidden terminals and to exploit exposed terminals.

Picasso [22] enables virtualization across the MAC/PHY and uses spectrum slicing. It allows a single radio to receive and transmit on different frequencies simultaneously. MAClets [7] allows multiple MAC/PHY protocols to share a single RF frontend. These advances can be used by OPENSOWN (and already Odin) to operate multiple LVAPs with different characteristics on top of the same AP. Alternative approaches, such as [43] and [45], are incompatible with today's WiFi MAC/PHY and thus do not fit our design requirements. FICA [43] introduces a new PHY layer, that splits the channel into separate subchannels which stations can simultaneously use according to their traffic demands. Jello [45], a MAC overlay where devices sense and occupy unused spectrum without central coordination or dedicated radio for control. Enabling per-flow transmission settings will allow applications to centrally implement rate and power control. With OpenRadio [5], our system could also benefit from a clean-slate programmable network dataplane.

There is also a number of interesting works in the context of programmable *cellular* networks. C-RAN [9] (i.e., *Cloud-RAN*), is a new cellular network architecture for the future mobile network infrastructure. It combines centralized processing, cooperative radio and cloud, to render the radio access network more flexible. SoftCell [24] simplifies the operation of cellular networks and supports high-level service policies to direct traffic through sequences of MBs. Fine-grained packet classifications are pushed to the access switches, and to ensure control-plane scalability, a local agent at the base station caches the service policy. Openflow-based SDN also offers a number of benefits for mobile networks, including wireless access segments, mobile backhaul networks, and core networks. SoftRAN [19] uses SDN principles to redesign the radio access network, and seeks to provide the "big-base station abstraction": it coordinates radio resource management through its logically centralized control plane, managing interference, load, QoS, etc. through plug and play algorithms.

OPENSOWN promotes a unified programmable control over network and middleboxes. Middleboxes are ubiquitous in today's computer networks [36]. Besides virtualization, middleboxes also play an important role in OPENSOWN for the fine-grained transmission control, which is based in deep-packet inspection [3, 14, 40]. Sekar et al. were one of the first to emphasize the importance of middleboxes, and in their middlebox manifesto [35], the authors argued for software-centric middlebox implementations running on general-purpose hardware platforms that are managed via open and extensible management APIs. Also Gember-Jacobson et al. [18] argue for a joint control of NFV and SDN components, and present the OpenNF architecture to coordinate the different control plane

tasks, and to enable an efficient reallocation of flows across network function instances. Concretely, the southbound interface of OpenNF deals with the network function state diversity and seeks to minimize modifications. The northbound interface allows control applications to flexibly move, copy, or share subsets of state between NF instances. Merlin [41] is a language to provision network functions and entire network function chains. An interesting NFV platform is *ClickOS* [26], a virtualized software middlebox platform, based on light virtual machines. OPENSNDWN is based on the Click modular router [12].

Home networks have received particular attention over the last years [13, 48]. Users are offered more flexibilities on how their network can be optimized [27, 47], sometimes even over participatory interfaces [13], helping home users to improve performance [34]. Programmable middleboxes can also be exploited to provide a faster ISP service delivery [25].

7. CONCLUSION

We have presented OPENSNDWN, a programmable and virtualized WiFi network which may be used as a prototype to experiment and demonstrate a more flexible and fine-grained network management environment for WiFi networks. Through prototype implementations of unified programmability and abstractions, a programmable datapath, and a user interface, we proposed use cases of service differentiation mobility and migration, fleible deployment and flexible control. We believe that OPENSNDWN demonstrated interesting and valuable capabilities particularly in the context of home and enterprise networks, which can benefit from the capabilities demonstrated in this research.

Acknowledgments. We are thankful for many discussions with Henry Owen, James Kempf, and Thomas Hühn. We also like to thank Sven Zehl and Tobias Steinicke for their efforts on the wireless datapath. Research supported by the Federal Ministry of Education and Research (BMBF) Software Campus "SDWN" Project Grant (Reference number 01IS12056).

8. REFERENCES

- [1] Meraki. <http://www.meraki.com/>.
- [2] Meru Networks. <http://www.merunetworks.com>.
- [3] Re-examining the performance bottleneck in a {NIDS} with detailed profiling. *Journal of Network and Computer Applications*, 36(2):768 – 780, 2013.
- [4] B. Anwer, T. Benson, N. Feamster, D. Levin, and J. Rexford. A Slick Control Plane for Network Middleboxes. In *Proc. HotSDN '13*.
- [5] M. Bansal, J. Mehlman, S. Katti, and P. Levis. OpenRadio: a programmable wireless dataplane. In *HotSDN '12*.
- [6] M. Bansal, A. Schulman, and S. Katti. Atomix: A Framework for Deploying Signal Processing Applications on Wireless Infrastructure. In *Proc. NSDI*, 2015.
- [7] G. Bianchi, P. Gallo, D. Garlisi, F. Giuliano, F. Gringoli, and I. Tinnirello. MAClets: active MAC protocols over hard-coded devices. In *Proc. CoNEXT '12*.
- [8] B. Chen, V. Yenamandra, and K. Srinivasan. FlexRadio: Fully Flexible Radios and Networks. In *Proc. NSDI 15*, 2015.
- [9] China Mobile Research Institute. C-RAN: The road toward green RAN. In *White Paper*, 2011.
- [10] A. Cidon, K. Nagaraj, S. Katti, and P. Viswanath. Flashback: decoupled lightweight wireless control. In *ACM SIGCOMM '12*.
- [11] Cisco. Cisco Service Provider Wi-Fi: A Platform for Business Innovation and Revenue Generation. In *Cisco*, 2015.
- [12] Click modular router project. <http://read.cs.ucla.edu/click>.
- [13] C. Dixon, R. Mahajan, S. Agarwal, A. Brush, B. Lee, S. Saroiu, and P. Bahl. An operating system for the home. In *Proc. NSDI*, 2012.
- [14] H. Dreger, C. Kreibich, V. Paxson, and R. Sommer. Enhancing the accuracy of network-based intrusion detection with host-based context. In *Proc. DIMVA*, 2005.
- [15] D. Drutskey, E. Keller, and J. Rexford. Scalable network virtualization in software-defined networks. *Internet Computing, IEEE*, 17(2):20–27, March 2013.
- [16] N. Feamster, J. Rexford, and E. Zegura. The road to SDN. *Queue*, 11(12):20:20–20:40, Dec. 2013.
- [17] A. D. Ferguson, A. Guha, C. Liang, R. Fonseca, and S. Krishnamurthi. Participatory networking: An api for application control of SDNs. *SIGCOMM Comput. Commun. Rev.*, 2013.
- [18] A. Gember-Jacobson, R. Viswanathan, C. Prakash, R. Grandl, J. Khalid, S. Das, and A. Akella. OpenNF: Enabling Innovation in Network Function Control. In *Proc. ACM SIGCOMM*, 2014.
- [19] A. Gudipati, D. Perry, L. E. Li, and S. Katti. SoftRAN: Software Defined Radio Access Network. In *Proc. HotSDN '13*.
- [20] A. Gupta, L. Vanbever, M. Shahbaz, S. P. Donovan, B. Schlinker, N. Feamster, J. Rexford, S. Shenker, R. Clark, and E. Katz-Bassett. SDX: A Software Defined Internet Exchange. In *Proc. SIGCOMM '14*.
- [21] C.-Y. Hong, S. Kandula, R. Mahajan, M. Zhang, V. Gill, M. Nanduri, and R. Wattenhofer. Achieving High Utilization with Software-driven WAN. *SIGCOMM Comput. Commun. Rev.* 2013.
- [22] S. S. Hong, J. Mehlman, and S. Katti. Picasso: Flexible RF and Spectrum Slicing. In *ACM SIGCOMM 2012*.
- [23] S. Jain, A. Kumar, S. Mandal, J. Ong, L. Poutievski, A. Singh, S. Venkata, J. Wanderer, J. Zhou, M. Zhu, J. Zolla, U. Hölzle, S. Stuart, and A. Vahdat. B4: Experience with a globally-deployed software defined wan. *SIGCOMM Comput. Commun. Rev.*, 2013.
- [24] X. Jin, L. E. Li, L. Vanbever, and J. Rexford. SoftCell: Scalable and Flexible Cellular Core Network Architecture. In *CoNEXT '13*.
- [25] K. R. Khan, Z. Ahmed, S. Ahmed, A. Syed, and S. A. Khayam. Rapid and scalable isp service delivery through a programmable middlebox. *SIGCOMM Comput. Commun. Rev.* 2014.
- [26] J. Martins, M. Ahmed, C. Raiciu, V. Olteanu, M. Honda, R. Bifulco, and F. Huici. ClickOS and the Art of Network Function Virtualization. In *Proc. NSDI*, 2014.
- [27] R. Mortier, T. Rodden, T. Lodge, D. McAuley, C. Rotsos, A. W. Moore, A. Kolioussis, and J. Sventek. Control and understanding: Owning your home network. In *Proc. COMSNETS*, 2012.
- [28] R. Murty, J. Padhye, R. Chandra, A. Wolman, and B. Zill. Designing high performance enterprise Wi-Fi networks. In *Proc. NSDI '08*.
- [29] R. Murty, J. Padhye, A. Wolman, and M. Welsh. Dyson: an architecture for extensible wireless LANs. In *Proc. USENIX ATC '10*.

- [30] V. Paxson. Bro: A system for detecting network intruders in real-time. *Comput. Netw. Dec. 1999*.
- [31] Z. A. Qazi, C.-C. Tu, L. Chiang, R. Miao, V. Sekar, and M. Yu. SIMPLE-fying Middlebox Policy Enforcement Using SDN. In *ACM SIGCOMM '13*.
- [32] J. Schulz-Zander, N. Sarrar, and S. Schmid. AeroFlux: A Near-Sighted Controller Architecture for Software-Defined Wireless Networks. In *Proc. Open Networking Summit (ONS)*, 2014.
- [33] J. Schulz-Zander, L. Suresh, N. Sarrar, A. Feldmann, T. Hühn, and R. Merz. Programmatic Orchestration of WiFi Networks. In *Proc. USENIX ATC '14*.
- [34] M. S. Seddiki, M. Shahbaz, S. Donovan, S. Grover, M. Park, N. Feamster, and Y.-Q. Song. FlowQoS: QoS for the Rest of Us. In *Proc. ACM HotSDN*, 2014.
- [35] V. Sekar, S. Ratnasamy, M. K. Reiter, N. Egi, and G. Shi. The middlebox manifesto: Enabling innovation in middlebox deployment. In *Proc. ACM HotNets*, 2011.
- [36] J. Sherry, S. Hasan, C. Scott, A. Krishnamurthy, S. Ratnasamy, and V. Sekar. Making middleboxes someone else's problem: Network processing as a cloud service. In *Proc. ACM SIGCOMM 2012*.
- [37] R. Sherwood, G. Gibb, K.-K. Yap, G. Appenzeller, M. Casado, N. McKeown, and G. Parulkar. Can the production network be the testbed? In *OSDI '10*.
- [38] V. Shrivastava, N. Ahmed, S. Rayanchu, S. Banerjee, S. Keshav, K. Papagiannaki, and A. Mishra. CENTAUR: realizing the full potential of centralized wlans through a hybrid data path. In *Proc. MobiCom '09*.
- [39] S.K. Fayazbakhsh et al. Enforcing Network-Wide Policies in the Presence of Dynamic Middlebox Actions using FlowTags. In *Proc. USENIX NSDI*, 2014,.
- [40] R. Sommer, M. Vallentin, L. De Carli, and V. Paxson. HILTI: An Abstract Execution Environment for Deep, Stateful Network Traffic Analysis. In *Proc. IMC*, 2014.
- [41] R. Soulé, S. Basu, P. J. Marandi, F. Pedone, R. Kleinberg, E. G. Sirer, and N. Foster. Merlin: A language for provisioning network resources. In *Proc. CoNEXT*, 2014.
- [42] L. Suresh, J. Schulz-Zander, R. Merz, A. Feldmann, and T. Vazao. Towards programmable enterprise WLANS with Odin. In *HotSDN '12*.
- [43] K. Tan, J. Fang, Y. Zhang, S. Chen, L. Shi, J. Zhang, and Y. Zhang. Fine-grained channel access in wireless LAN. In *ACM SIGCOMM 2010*.
- [44] P. Valerio. Using carrier wifi to offload iot networks. In *InformationWeek: Network Computing*, 2014.
- [45] L. Yang, W. Hou, L. Cao, B. Y. Zhao, and H. Zheng. Supporting demanding wireless applications with frequency-agile radios. In *USENIX NSDI '10*.
- [46] Y. Yiakoumis, M. Bansal, A. Covington, J. van Reijndam, S. Katti, and N. McKeown. BeHop: A Testbed for Dense WiFi Networks. In *Proc. WiNTECH '14*.
- [47] Y. Yiakoumis, S. Katti, T.-Y. Huang, N. McKeown, K.-K. Yap, and R. Johari. Putting home users in charge of their network. In *Proc. UbiComp '12*.
- [48] Y. Yiakoumis, K.-K. Yap, S. Katti, G. Parulkar, and N. McKeown. Slicing home networks. In *Proc. HomeNets '11*.
- [49] P. Zerfos, G. Zhong, J. Cheng, H. Luo, S. Lu, and J. J. Li. DIRAC: a software-based wireless router system. In *MobiCom*, 2003.