Cloudlet- and NFV-based carrier Wi-Fi architecture for a wider range of services

Fatma Ben Jemaa, Guy Pujolle, Michel Pariente

Cloudlet- and NFV-based Carrier Wi-Fi Architecture for a Wider Range of Services

Fatma Ben Jemaa *†, Guy Pujolle *, and Michel Pariente †

* Sorbonne Universités, UPMC Univ Paris 06, CNRS, LIP6 UMR 7606, Paris, France. Email: Guy.pujolle@lip6.fr
† Meteor Network, Vitry-sur-Seine, France. Email: {fbenjemaa, mpariente}@meteornetworks.com

Abstract—Over the past few years, Wireless Local Area Networks (WLANs) have been extensively deployed and have significantly evolved. However, the deployment of large-scale WLAN still presents management issues. Moreover, while newer WLAN technologies and services have been emerging at a prolific rate, the architecture of WLAN networks has been quite static and has seen difficulties to evolve. In this paper, we present a novel architecture for carrier-managed WLAN networks which leverages Network Function Virtualization concepts and virtualization technology in general. It is based on a WLAN Cloudlet which offloads MAC layer processing from access points and consolidates network functions and value-added services. All these functions and services are based on software instances. This brings more flexibility and adaptability and allows operators to easily implement new services while reducing CAPEX/OPEX and network equipment costs (e.g., access points).

Keywords—Carrier-grade WLAN; MAC layer; Network Function Virtualization; WLAN emerging services; Software-based WLAN architecture; Cloudlet.

I. INTRODUCTION

Over the last decade, WLAN networks have experienced an incredible evolution with the emergence of new technologies and services which mainly enhance user experience and improve quality of service. Nevertheless, today's WLAN networks are unable to rapidly adapt to such evolution due to their rigid architectural design. In fact, this typically requires time-consuming and costly upgrades of existing infrastructures generally composed of proprietary hardware appliances. Consequently, the issue of implementing flexible architecture while boosting innovation and reducing the cost of network upgrades becomes a major concern to support evolving contexts and service needs.

A new industry trend, Network Function Virtualization (NFV) [1], has gained a lot of attention over the past few years. Through its basic idea, decoupling software from hardware, it promises cost efficient realization of network functions in software deployed over commodity hardware. Moreover, it encourages openness and innovation to quickly bring new services and new revenue streams at much lower risk.

Leveraging the same concepts, we introduce a Cloudlet-based WLAN architecture in order to obtain similar benefits in WLANs. More specifically, the Cloudlet is a server located in the end user premises that not only consolidates network functions but also some MAC functions and value-added services.

Using software-based functions on the three levels (i.e., MAC, network, and service levels) brings more flexibility and adaptability to the whole system. Thus, it allows for reducing the complexity and cost of introducing new functionalities and services in the different levels.

Moreover, the WLAN Cloudlet introduces proximate virtualization infrastructure in WLAN architecture which provides two major benefits. First, this plays a crucial role in reducing i) Access Point (AP) costs by offloading MAC layer processing to virtual machines provided by the Cloudlet and ii) other network equipment costs through consolidating multiple instances of network functions in the same hardware. Secondly, it decreases access latency by placing network functions and certain services close to end-users.

The remainder of the paper is structured as follows. Section 2 discusses the evolution of WLAN architectures while highlighting the major advantages and drawbacks of each one. In section 3, we describe our proposed Cloudlet-based WLAN architecture. Benefits and possible applications of this architecture are presented in section 4. Section 5 describes feasibility and implementation aspects and provides performance evaluation results. Finally, we respectively present deployment challenges and conclude this paper in section 6 and 7.

II. EVOLUTION OF WLAN ARCHITECTURE

In the first generation of WLAN, access points (APs) were individually managed and independent. Moreover, they contained all the intelligence to manage Wi-Fi traffic [2, 3]. As a result, duplicating this intelligence induces high cost especially for medium and large sized networks. In addition, this solution has limited capacities and is relatively static.

As the need for centralized monitoring and dynamic configurability grew, vendors introduced controller-based systems with "thin" low cost APs. In this architecture, the controller is responsible for controlling, configuring, and managing the entire WLAN access network. Furthermore, all the traffic is routed from the APs to the controller [4]. According to RFC 4118 [2], centralized WLAN architectures are categorized into three main variants: i) the Local MAC in which the MAC functions stay intact and local to APs, ii) the Remote MAC in which the MAC has moved away from the AP to a remote Access Controller (AC) in the network, and iii) the Split MAC in which the MAC is split between the APs and ACs.

The centralized WLAN architecture is characterized by the ease of deployment especially for wide networks and it provides more security and control. However, it has two major drawbacks. Firstly, the WLAN controller is a single point of failure. Secondly, since all transmissions require passing through the controller, the latter became a major bottleneck eroding network performance.

To resolve these issues, the third generation of WLAN, called distributed WLAN architecture, introduced distributed data forwarding [5]. A controller still provides a central point of control for APs; however, all traffic is no longer backhauled to the controller. This solution eliminates network bottleneck but it is more much expensive than the other solutions.

Recently, a new trend has emerged in the world of WLAN by introducing virtualized WLAN architecture which has rapidly been gaining the attention of industry and academia. Some vendors offer the controller as a virtual appliance and even as a cloud-based hosted service [6-8]. The cloud-based controller centrally manages and monitors APs and the user data is not going through the controller. This solution has several advantages in terms of ease of deployment and availability but it exhibits a number of shortcomings. Note that even if this solution eliminates the need of deploying on premise hardware WLAN controllers, it does not involve cost reduction. In fact, controller functionality is distributed between the APs and the cloud. Thus, the cost of the controller is integrated into the price of the APs and the Cloud controller subscription that must be continuously renewed. Over the long haul, this solution is much more expensive than simply purchasing a WLAN controller from the start. Moreover, for security reasons, many organizations still require on premise WLAN controller to easily manage network traffic and there may be regulatory needs to tunnel traffic to the controller. Finally, this solution is only suitable for low- to medium-density locations as APs are not designed to support high-scale tunnel traffic. In contrast, centralized controllers have dedicated hardware that makes them extremely efficient at moving traffic through the network.

It's worth pointing out that virtualization did not only target WLAN controller but also access points. CloudMAC [9] attempts to partially offload the MAC layer processing to virtual machines in the cloud acting as virtual APs. This allows for reducing the software complexity of physical APs. Furthermore, adding support for new functionalities related to MAC processing requires updates only to the virtual APs. The deficiency of this solution is the additional latencies introduced due to the distant MAC frames processing.

To summarize, among the aforementioned WLAN architectures, there is no optimal and efficient solution in terms of cost, QoS (e.g., latency), flexibility by enabling dynamic configurability and adaptability, and simple deployment of new services.

III. CLOUDLET-BASED WLAN ARCHITECTURE

Cloudlet is not a new concept. It was first introduced by Satyanarayanan et al. [10] in the context of Mobile Computing to enable mobile users to rapidly instantiate customized service software on a nearby cloudlet and then use that service over a wireless LAN. This allows the resource-constrained mobile device to function as a thin
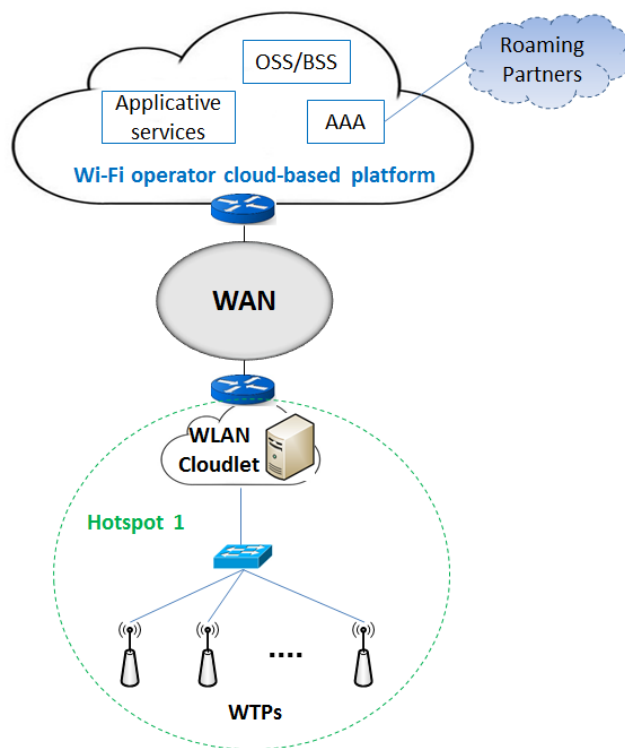

Fig. 1. Cloudlet-based WLAN architecture

client and to exploit virtual machine technology while avoiding long WAN latencies of Cloud Computing. It was proposed to integrate Cloudlet with access points [10], with wireless mesh networks [11], and with base stations [12]. Most of works related to Cloudlet concept focused on Mobile Computing for resource-demanding and delay-sensitive applications.

In our work, we extend this approach to not only support application-level services but also network services and some specific-802.11 functions. More specifically, we propose a novel carrier grade WLAN architecture that leverages network virtualization and cloud computing free of WAN delays, jitter, congestion, and failures. This allows greater flexibility, ease of management, and rapid creation and deployment of new services (L2-L7). Fig. 1 depicts the Cloudlet-based WLAN architecture composed of WLAN Cloudlet, Wireless Termination Points (WTPs), and a centralized Cloud-based platform.

A. WLAN Cloudlet

WLAN Cloudlet is located in the end-user premises and consolidates some MAC layer functions, network functions, and value-added services provisioned by software that run on an industry standard server hardware. In practice, these functions are bundled in virtual machines installed over hypervisors. As described in Fig. 2, WLAN Cloudlet incorporates the following types of virtualized functions and services:

- Virtualized MAC Functions: are handled by a Software AP (SoftAP) which partially offloads WLAN MAC processing from physical access points. In fact, it runs "non-real-time" MAC functions such as Distribution and Integration services or responding to Association/Authentication
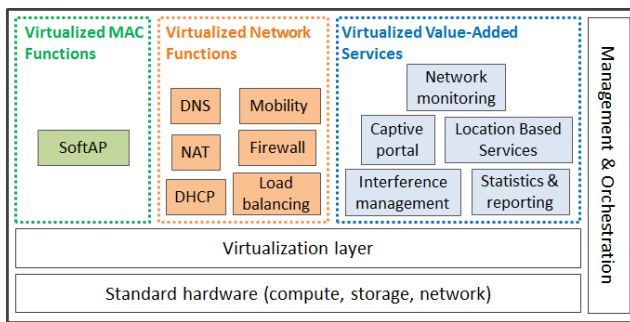
Fig. 2. WLAN Cloudlet architecture

MAC frames. A software AP can be connected to many WTPs. Consequently, a WLAN network (a hotspot) is only one Software AP.

This simplifies the management of WLAN and allows a simple update and rapid deployment of new functionalities using software modifications. Moreover, this allows the WTP to act as a thin AP by reducing software complexity.

- Virtualized Network Functions: include all network functions and services that are locally deployed in WLAN network as dedicated hardware infrastructures such as routing, DHCP, NAT, firewall, and DNS. Moving network functions from purpose-built appliances to equivalent functionalities implemented in Cloudlet, a COTS (Commercial Off The Shelf) hardware environment providing cloud computing capabilities, increases the flexibility to deploy new features and decreases setup and management cost.

- Virtualized Value-Added Services: include network and applicative services which improve QoS and user experience.

B. Wireless Termination Points

Wireless Termination Points (WTPs) are the physical entities that contain an RF antenna and implements 802.11 physical layer functions to transmit and receive station traffic over the air. They operate as a pass-through, forwarding MAC management frames between the WLAN clients and the SoftAP within the WLAN Cloudlet. The WTPs also implement the medium access using the Distributed Coordination Function [13]. In addition, WTPs process MAC services with real-time constraints such as synchronization, retransmissions, generating Control frames (e.g., RTS, CTS and ACK), and beacon and probe response frames.

C. Cloud-based platform

Cloud-based platform includes functions that need to be executed in a centralized way (e.g., OSS/BSS, AAA) or requires a lot of resources (e.g., analytics). Thus, the WLAN Cloudlet provides interfaces to access these functions.

Furthermore, this Cloud platform can offer more scalability to our system. Indeed, processing capabilities of the WLAN cloudlet server can be extended by cloud resources especially for non-real-time functionalities and the virtualized function can be scaled by creating additional instances of the function in the Cloud platform. This ability of WLAN cloudlet to scale dynamically leads to support a very large number of flows.

IV. CLOUDLET-BASED WLAN BENEFITS AND POSSIBLE APPLICATIONS

A. Benefits

Cloudlet-based WLAN brings many benefits to the network operator and potentially to end-users. These benefits include:

- Reduced equipment costs by avoiding, especially in large scale WLANs, to have a large number of expensive fat APs to ensure network coverage since the WTPs are kept the lightest as possible.

- Reducing CAPEX and OPEX by getting rid of expensive purpose-built middle-boxes (e.g., access controller used in many organizations for security and regulatory requirements, WLAN controller, firewall, etc.) and avoiding their management complexities due to NFV technology.

- Decreasing access latency by placing network functions and certain applications close to end-users.

- More flexible and simplified network management. In fact, adding new MAC or network functionality requires only a software upgrade.

- Increased speed of Time to Market and faster configuration of new services.

- The possibility of introducing targeted services based on geography or venue type. Furthermore, provisioning could be made remotely in software without any site visits required to install new hardware.

- Encouraging openness and more innovation to bring new services and new revenue streams quickly at much lower risk.

B. Possible application scenarios

Cloudlet-based WLAN enables a range of applications and use cases. In this section, we present some examples.

- Simple upgrade to new technologies such as Hotspot 2.0 [14]. Actually, to adopt this technology, this requires a time-consuming and costly upgrade of existing infrastructure or even change of legacy equipment (i.e., AP and controllers). With our proposed solution, the deployment of this new feature will be rapid and cost-effective as it only requires a software upgrade of MAC functions to support this standard which can be performed centrally and there is no need to upgrade WTPs.

- Providing network gateway functions in an on-demand and customized fashion. Particularly, this enables customers to insert new network gateway functions as security measures according to their requirements and in order to face new security threats. For example, a virtual firewall could be implemented for each group of users with specific security policies.
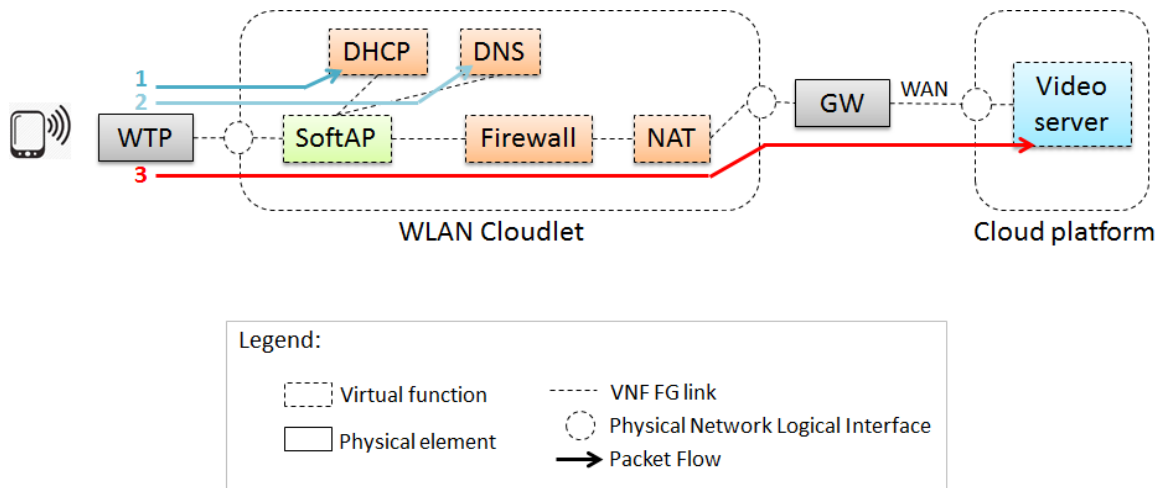
Fig. 3.   Use Case: VNF forwarding graph

- Introducing new services which enable to improve quality of service and network throughput. As an example of these services, we mention interference management service [15, 16] which allocates the optimum frequency channel to each WTP based on information related to the radio environment gathered from WTPs.

- Simple integration of venue based services (e.g., indoor mapping, special offers) that could leverage Wi-Fi Location Based Services. This could enhance user experience and offer monetization opportunities to Wi-Fi operators.

V.   Feasibility and Implementation

This section explores the feasibility of the Cloudlet-based WLAN architecture through a practical use case. Based on this reference case, we have built a proof-of-concept prototype and evaluated its performance.

A. Case Study

We assume the following basic scenario: a user desires to associate to a Wi-Fi network through his device to watch a video available on a video server. To describe the functional aspect of this scenario, we define a service chaining graph. To do this, we use the Virtual Network Function Forwarding Graph (VNF FG) [17], which is defined by the ETSI as one of the NFV components. This graph defines the sequence of network functions that packets go through and provides the logical connectivity between virtual appliances. A VNF FG can also include physical network functions to provide an end-to-end network service. As illustrated in Fig. 3, in our scenario, the VNF FG is composed of the WTP, a set of virtual functions (SoftAP, DHCP, DNS, firewall and NAT), a gateway, and a video server.

B. Implementation aspects

In our testbed for the proof-of-concept, the WTP is implemented using a machine equipped with Intel Pentium dual CPU 2.00 GHz processor with 2GB of RAM, Linux Ubuntu OS (release 14.04.2 LTS) and a TL-WN722N wireless card using Atheros 9271 chipset which uses ATH9K_HTC as a driver. This card supports the monitor mode, which allows for receiving and transmitting raw MAC frames.

The WLAN Cloudlet is a machine equipped with an Intel Core i5-3337U 1.8 GHz processor with 4GB of RAM. It is connected to the WTP through a Gigabit Ethernet switch and it uses VMware Player that runs two VMs.

The first VM is the Software AP which uses Xubuntu OS and communicates with the WTP using OpenCAPWAP protocol [18-20], an open source version of CAPWAP (RFC 5416) [21]. This protocol supports the Split MAC approach in which only the real-time IEEE 802.11 services, such as the beacon and probe response frames, are handled on the WTP and all remaining MAC management frames and distribution and integration services reside on the SoftAP. We have chosen this approach to make the WTP as light as possible while satisfying 802.11 timing constraints. Moreover, most of the MAC services can be updated centrally.

The second VM is a FreeBSD guest machine running pfSense [22], an open source firewall which is able to run a set of services such as DHCP, DNS, routing and NAT.

C. Performance evaluation

In this section, we report experimental results concerning the performance of our Cloudlet-based WLAN prototype implementation. Using the testbed described above, we have measured the delay and throughput metrics between different elements in our architecture. We describe below the different sets of measurement as well as the obtained results for each metric.

1)   Delay

To evaluate delay performance, we measure the Round Trip Time (RTT) metric given by the ping tests. In these tests, we compare our proposal to a reference architecture which includes a standalone AP based on Hostapd [23] (using the same hardware as the WTP) and is connected to the virtual firewall.

The delay measurements are performed in three different segments of the network referred to in the following as case 'a', case 'b' and case 'c'.

- Case 'a': We consider the delay between the user device and the video server.

(a) RTT between user device and video server (case 'a')

(b) RTT between user device and the gateway (case 'b')

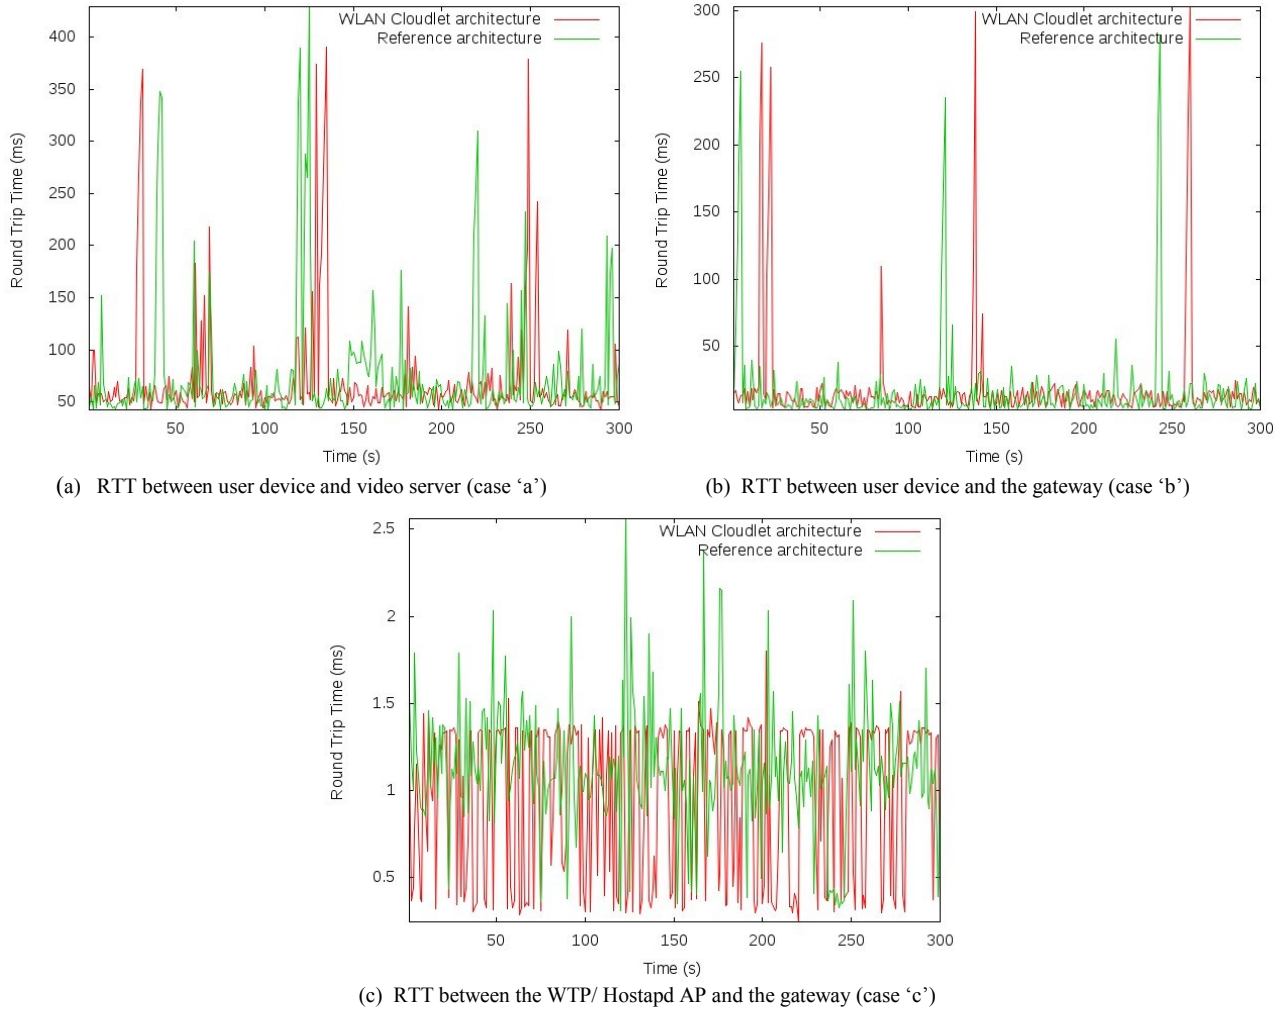(c) RTT between the WTP/ Hostapd AP and the gateway (case 'c')

Fig. 4. RTT measurements

- Case 'b': We measure the RTT for both architectures between the user device and the gateway in order to have more accurate values without the impact of dynamic nature of Internet network.

- Case 'c': We focus on the delay between the WTP (Hostapd AP in the reference architecture) and the gateway, thus excluding the radio link delay.

Fig. 4 shows the RTT values measured each second during two minutes for the three cases. The minimum, average, maximum and standard deviation values of RTT are presented in Table 1.

TABLE I.   RTT VALUES (ms)

		Min	Avg	Max	Std Dev
Case 'a'	WLAN Cloudlet architecture	43.310	71.069	390.764	52.655
	Reference architecture	42.672	75.609	429.529	57.865
Case 'b'	WLAN Cloudlet architecture	4.180	18.632	303.343	37.777
	Reference architecture	2.897	16.334	282.601	32.785
Case 'c'	WLAN Cloudlet architecture	0.247	0.964	1.808	0.458
	Reference architecture	0.309	1.149	2.560	0.484

In case 'a' and 'c', the curves for both architectures are almost indistinguishable and RTT values are very close. However, in case 'b', the results reveal more visibly that the delay in our architecture is slightly higher than the reference architecture. This is mainly due to the presence of separate physical and software access points in our architecture, while, in the reference architecture, we have only a standalone access point. Thus, user traffic is passing through one more hop in the proposed architecture than the reference one. Moreover, this delay is also due to the virtualization layer.

The results, in general, show that WLAN Cloudlet solution achieves delay performance metrics comparable to the reference architecture.

2) Throughput

In the second set of experiments, we evaluate the throughput performance in the proposed WLAN Cloudlet architecture between the user device and the SoftAP (Fig. 5(a)), on the one hand, and between the WTP and the SoftAP (Fig. 5(b)), on the other hand. To do this, we use Iperf tool which enables to measure the maximum TCP bandwidth. We perform tests during 10 minutes when results are reported every two seconds in the first experiment and during 5

(a) Throughput between user device and the SoftAP

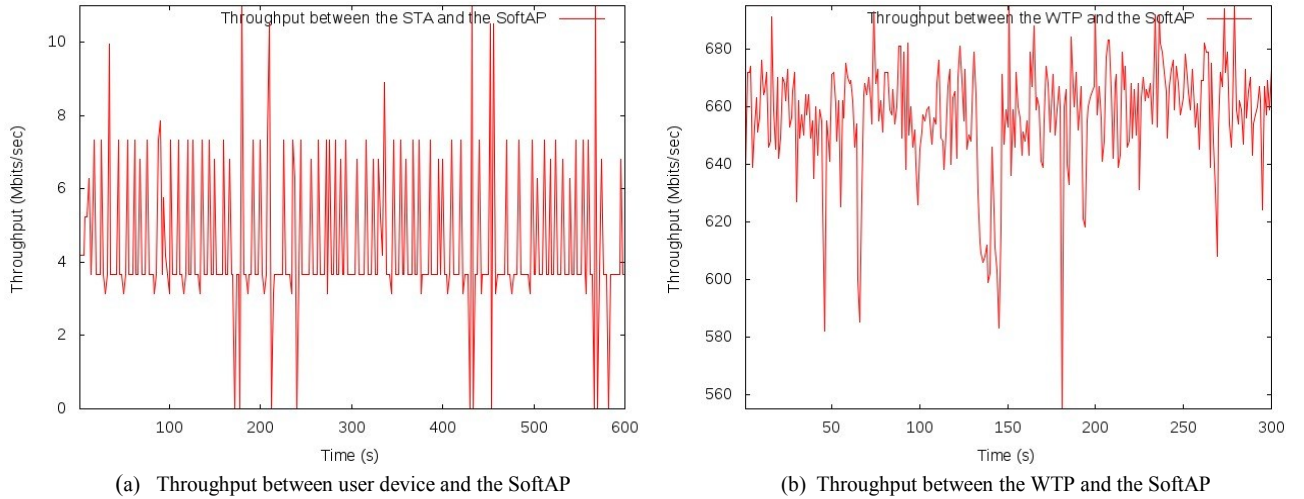

(b) Throughput between the WTP and the SoftAP

Fig. 5. TCP Throughput measurements

minutes when results are reported each second in the second experiment. So, we have in total 300 throughput measurement values for each experiment. The minimum, average, maximum and standard deviation of these values are depicted in Table 2 for both segments of network mentioned above.

We notice that the maximum reached TCP throughput value between the client and the SoftAP is 11 Mbps and the average value is 4.45 Mbps. These values are justified because the used WLAN standard in the tests is 802.11b which has as maximum theoretical throughput 11 Mbps. We also have to mention that experiments are performed in an operational wlan network where interferences with other networks cannot be avoided.

In order to eliminate the radio link effect and evaluate the maximum bandwidth of our internal system, we measure the throughput between the WTP and the SoftAP. Fig. 5(b) and Table 2 show that the maximum throughput is 695 Mbps and the average value is 656.64 Mbps. Consequently, in this current testbed, the SoftAP can handle about 130 users with a maximum of 5 Mbps for each one.

TABLE II. THROUGHPUT VALUES (Mbps)

	Min	Avg	Max	Std Dev
Between user device and SoftAP	3.15	4.45	11	1.92
Between the WTP and the SoftAP	555	656.64	695	19.72

According to the aforementioned results, WLAN Cloudlet solution does not affect network performance metrics and, at the same time, provides many advantages in terms of cost, flexibility and agility.

VI. CHALLENGES

In order to realize the concept of Cloudlet-based WLAN, several challenges need to be addressed. This section summarizes the main challenges.

- *Service description*. The characteristics of each virtualized service component should be accurately

described in an information model. Service description needs to cover the required resources such as compute, memory, storage, and bandwidth.

- *Management and orchestration*. The Cloudlet-based WLAN framework should incorporate mechanisms for the automated management and the dynamic orchestration of virtualized functions and services to enhance availability, flexibility, elasticity and to meet targeted performance constraints.

As particular management mechanisms, we highlight below monitoring and optimization functionalities and VM placement.

- *Monitoring and optimization*. Adaptive monitoring and optimization approaches are needed to constantly monitor networking and computing/storage assets and optimize resources between multiple service instances relative to usage and policy constraints under dynamic conditions.

- *VM placement*. As a major optimization mechanism, a resource mapping function tailored for Cloudlet-based WLAN deployment is needed to determine where to place the VM of each service, either in the WLAN Cloudlet or in the Cloud. This function requires a prior knowledge of physical machines' capacity and VM requirements. Additionally, it may be subject to different objectives such as optimizing system utilization and performance and reducing backbone network workload with respect to a set of criteria. The latter could be related to several requirements such as QoS (e.g., response time), location constraints (as some services may require to be in the proximity of end-users), and security aspects.

- *Elasticity and scalability*. Automatic scaling mechanisms should be supported to satisfy service level agreements and resolve eventual bottlenecks in the WLAN Cloudlet. For example, virtualized functions could be scaled by creating additional instances of the function in the Cloud platform or by migrating the VM to the Cloud.

VII. CONCLUSION AND FUTURE WORK

Operating large-scale WLAN deployments with heterogeneous hardware and software components, especially when technology and services innovation accelerates, is becoming very challenging for network operators. In this paper, we have presented a Cloudlet-based WLAN architecture in which functions on the MAC, network and service layers are virtualized onto an industry standard server located in the end user premises. Such architecture provides many benefits to network operators by reducing time and cost to integrate new services and afford flexibility and adaptability that would enhance user experience in next generation networks. It should be mentioned that the described solution is still the object of further research studies, especially related to the management and orchestration of virtualized functions.

REFERENCES

[1] ETSI, "Network Functions Virtualisation: An Introduction, Benefits, Enablers, Challenges & Call for Action", White paper, SDN and OpenFlow World Congress, October 2012. [Online]. Available: http://portal.etsi.org/nfv/nfv_white_paper.pdf.

[2] L. Yang, P. Zerfos, and E. Sadot, "Architecture Taxonomy for Control and Provisioning of Wireless Access Points (CAPWAP)", RFC 4118, June 2005.

[3] Infonetics Research, "The Evolution of the Enterprise-Class Wireless LAN Access Point", January 2004. [Online]. Available: http://www.nsaservices.com/pdf/airespace/WP_Enterprise_Class_WLAN_AP.pdf.

[4] 3e Technologies International, "802.11 WLAN architecture – Best practices", White paper, February 2005. [Online]. Available: http://www.ultra-3eti.com/assets/1/7/WirelessArchitectureBestPractices.pdf

[5] Motorola, "Distributed Intelligence: The Future of Wireless Networking Architecture", White paper, April 2011. [Online]. Available: http://www.csc.villanova.edu/~nadi/csc8580/S13/DistributedIntelligence.pdf. Accessed : 2015-11-02.

[6] Meraki, "Meraki hosted architecture", White Paper, February 2011. [Online]. Available: http://www.voyager.net.uk/wp-content/uploads/downloads/2014/02/WP-Meraki-Hosted-Architecture.pdf. Accessed : 2015-11-02.

[7] Motorola, "The next logical evolution in WLAN architecture", Technical brief, 2014.

[8] Meru, "Virtualized architecture enables choice, efficiency, and agility for enterprise mobility", White paper, 2012.

[9] J. Vestin, P. Dely, A. Kassler, N. Bayer, H. Einsiedler, and C. Peylo, "CloudMAC: towards software defined WLANs," ACM SIGMOBILE Mobile Computing and Communications Review, vol. 16, no. 4, pp. 42–45, 2013.

[10] M. Satyanarayanan, P. Bahl, R. Caceres, and N. Davies, "The case for vm-based cloudlets in mobile computing," Pervasive Computing, IEEE, vol. 8, no. 4, pp. 14–23, 2009.

[11] K. A. Khan, Q. Wang, C. Grecos, C. Luo, and X. Wang, "MeshCloud: Integrated cloudlet and wireless mesh network for real-time applications," IEEE 20th International Conference on Electronics, Circuits, and Systems (ICECS), pp. 317–320, 2013.

[12] M. Felemban, S. Basalamah, and A. Ghafoor, "A distributed cloud architecture for mobile multimedia services," IEEE Network Magazine, vol. 27, no. 5, pp. 20–27, 2013.

[13] "IEEE standard for Information technology— Telecommunications and information exchange between systems Local and metropolitan area networks—Specific requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications", IEEE Std 802.11™-2012.

[14] Wi-Fi Alliance. Hotspot 2.0, Release 1, Technical Specification, Version 1.0.0, 2012.

[15] ETSI, "Network functions virtualisation (NFV): Use cases", Group Specification, version 1.1.1, October 2013. [Online]. Available: http://docbox.etsi.org/ISG/NFV/Open/Published/gs_NFV001v010101p%20-%20Use%20Cases.pdf. Accessed: 2015-04-11.

[16] W. El-Hajj and H. Alazemi, "Optimal frequency assignment for IEEE 802.11 wireless networks," Wireless communications and mobile computing, vol. 9, no. 1, pp. 131–141, 2009.

[17] A. Farsi, N. Achir, and K. Boussetta, "WLAN planning: Separate and joint optimization of both access point placement and channel assignment," Ann. Telecommun., vol. 70, no. 5–6, pp. 263–274, Sep. 2014.

[18] M. Bernaschi, F. Cacace, G. Iannello, M. Vellucci, and L. Vollero, "OpenCAPWAP: An open source CAPWAP implementation for the management and configuration of WiFi hot-spots," Computer Networks, vol. 53, no. 2, pp. 217–230, 2009.

[19] E. Agostini, M. Bernaschi, M. Vellucci, and L. Vollero, "The new version of OpenCAPWAP", submitted to the International Journal of Network Management.

[20] OpenCAPWAP code source. [Online]. Available: https://github.com/vollero/openCAPWAP/tree/elena.ago.

[21] P. Calhoun, M. Montemurro, and D. Stanley, "Control and provisioning of wireless access points (CAPWAP) protocol binding for IEEE 802.11", RFC 5416, March 2009.

[22] pfSense, https://www.pfsense.org/

[23] Hostapd, http://w1.fi/hostapd/