WiFiMon: Combining Crowdsourced and Probe Measurements for Wi-Fi Performance Evaluation

Nikos Kostopoulos
School of Electrical & Computer Engineering
National Technical University of Athens
Athens, Greece
nkostopoulos@netmode.ntua.gr

Sokol Gjeçi
IT Department
RASH
Tirana, Albania
sgjeci@rash.al

Kurt Baumann
Network Department
SWITCH
Zurich, Switzerland
kurt.baumann@switch.ch

Pavle V. Vuletić
Department of Computer Science & Information Technology
University of Belgrade, School of Electrical Engineering
Belgrade, Serbia
pavle.vuletic@etf.bg.ac.rs

Kostas Stamos
Computer Engineering & Informatics Department
University of Patras
Patras, Greece
kostas.stamos@gmail.com

***Abstract*—We present WiFiMon, a lightweight active monitoring open source system for Wi-Fi network performance evaluation and verification. Using an objective Quality of Experience (QoE) assessment approach, WiFiMon is capable of capturing Wi-Fi throughput as experienced by end users without requiring their intervention. This is possible via crowdsourced measurements that are automatically triggered when the end users visit websites preconfigured with WiFiMon technology. Crowdsourced results are combined with probe measurements collected from small form factor devices located at fixed points within the monitored network; this enables Wi-Fi administrators to reach stronger conclusions about the performance of their networks. When deployed in IEEE 802.1X networks, WiFiMon provides additional benefits, e.g. fine-grained performance analysis per access point, by integrating data from available RADIUS and DHCP logs, which are correlated with the crowdsourced and probe measurements. Our system capabilities are demonstrated based on WiFiMon pilots, which took place in recent events with several hundreds of participants (TNC19 Conference, GÉANT Symposium 2020). We conclude that WiFiMon can effectively detect Wi-Fi performance fluctuations.**

***Index Terms*—Wi-Fi performance monitoring and verification, crowdsourced and probe measurements, mobile crowd-sensing, Quality of Experience (QoE), lightweight active monitoring**

I. Introduction

Wi-Fi networks have become an essential part of our lives and one of the most common Internet access methods. The constantly increasing popularity of Wi-Fi becomes evident in various aspects of everyday life, e.g. working environments, education facilities and recreational activities. Therefore, development of systems that provide accurate and efficient Wi-Fi network performance monitoring is of vital importance.

Although there are many tools that provide wireless network monitoring services, they fail to effectively address Quality of Experience (QoE) in performance monitoring and verification. Network Management Systems (NMS's) that monitor Wi-Fi infrastructure, i.e. Access Points (APs) and/or Wi-Fi controllers, mainly focus on determining if a Wi-Fi network is operational on its whole. Therefore, they can provide an insight view into the current status of the Wi-Fi network, number of associated devices, channel utilization and so on. However, the way end users perceive network quality is often neglected and, as it will be shown in this paper, objective Wi-Fi infrastructure metrics do not always accurately represent the quality of the network service as experienced by the user. There are some tools that monitor end user devices in an attempt to determine how the users experience Wi-Fi as they roam the network. However, these tests often rely on dedicated applications and are triggered manually by users, which is not always a solution that attracts users to widely adopt such tools.

The system described in this paper, WiFiMon, monitors the performance of Wi-Fi networks as experienced by their end users. Estimating QoE primarily assumes gathering subjective marks from users regarding how they perceive service quality. Objective QoE assessment is also possible as a time-efficient alternative using monitoring tools [1]. In that sense WiFiMon can be seen as an objective Wi-Fi network QoE estimation tool because the monitoring process includes the complete end-to-end system effects that may influence user experience, i.e. end user equipment, as well as network and service infrastructure [2]. As it will be described below, WiFiMon is based on measurements of the HTTP service which constitutes the majority of the average user's traffic today. However, since the QoE is strongly dependent on the application, WiFiMon cannot be used as a QoE tool for all the services (especially interactive voice or video). WiFiMon operates in passive objective testing mode [3], which means that the measurements are taken without the user intervention or any dedicated application.

Our system supports administrators by detecting Wi-Fi throughput degradation, hence allowing them to detect underperforming areas in their networks and, accordingly, enhance Wi-Fi performance, e.g. by installing more APs. WiFiMon

The work on WiFiMon receives funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement No. 856726 (GN4-3).

ISBN 978-3-903176-35-5 © IFIP

can be used in any Wi-Fi environment, but offers additional benefits in the eduroam-enabled [4], or generally IEEE 802.1X, campus networks. Crowdsourced measurements are triggered by end users, when they visit WiFiMon-enabled websites and/or use WiFiMon-enabled mobile applications. Although triggering measurements does not require any user intervention or additional tools apart from the web browser, in the WiFiMon terminology, the user's device is called WiFiMon Software Probe (WSP). WiFiMon combines crowdsourced measurements with deterministic ones derived from appropriately placed WiFiMon Hardware Probes (WHPs) at fixed points within the monitored networks to provide complete insight into network performance. WSP and WHP measurements are correlated with the identity of end users and their location, i.e. their associated AP. This information is exported from RADIUS and DHCP logs associated with the monitored end devices. This approach allows different forms of Wi-Fi network performance analyses (per user, per AP, per browser, etc.).

The basic principles of WiFiMon have been introduced in [5]. In this paper, we elaborate on new features incorporated in WiFiMon in the last two years. In particular, we (i) introduce WHPs and describe their operation, (ii) detail privacy enhancements in RADIUS/DHCP logs processing that safeguard end user sensitive data and (iii) evaluate the performance of our system, specifically WSPs and WHPs, in monitoring the Wi-Fi network of two recent conference venues, i.e. the TNC19 Conference and the GÉANT Symposium 2020.

WiFiMon is offered as a service by GÉANT, the academic backbone interconnecting the European National Research and Education Network (NREN) organizations. The documentation and open-sourced code of WiFiMon are available in [6].

The remainder of this paper is structured as follows: Section II describes related works; Section III includes the baseline design of WiFiMon; Section IV provides implementation details; Section V involves the system evaluation. Finally, Section VI summarizes our work and discusses future directions.

II. Related Work

The WiFiMon approach was inspired by [7], namely to take advantage of opportunistic measurements. WiFiMon builds upon this approach and uses frequently visited websites to monitor how the users experience the quality of a wireless network. This method differs from other approaches that monitor and verify performance using controllers to monitor standalone APs and are based on ping or scripts and application plugins that attempt to detect network events [8].

AP status can be monitored by processing SNMP objects from the Wi-Fi infrastructure. One such system that performs network monitoring based on data from controllers, base stations and the fixed network is Aruba AirWave [9]. However, it does not take into account the end user perspective. Another similar tool is Paessler PRTG network monitor [10] which uses SNMP, NetFlow and packet sniffing to verify device status including information from APs. Tools that focus on the end user side include NetSpot [11] and Tarlogic Acrylic [12] suite of Wi-Fi tools allowing for Wi-Fi signal quality monitoring. Similarly, NetAlly AirMagnet WiFi Analyzer PRO [13] provides details about Wi-Fi network performance as well as issues with coverage, interference and roaming. Comparable products are the Microsoft Wifi Analyzer and Scanner [14], Riverbed Wi-Fi Inspector [15], LizardSystems Wi-Fi Scanner [16], MetaGeek inSSIDer [17] and Netradar [18]. These approaches require an installation from the end user, while the performance results are not automatically communicated to the network administrators, responsible for setting up the wireless network and optimizing its performance. Such approaches are suitable for checking Wi-Fi network quality in the premises of users, but not for sites with a large number of APs where a different approach is needed.

There are several online tools/pages where users may manually initiate measurements and retrieve information about the performance of their connection, including Wi-Fi network, e.g. Speedtest by Ookla [19]. Performance metrics usually include information about the download/upload throughput and Round-Trip Time (RTT) [20], [21] towards a server installed in a generally unknown place in the Internet. This means that test results include the impact of the cross traffic along all the links between the user and the test server. Such online tools require that users initiate the tests and only themselves are informed of the results and not the network administrators.

Cnlab Speedtest [22] gathers test results executed by users in a central database and makes them available to licensees. As several thousand private customers use the cnlab Speedtest daily, it is possible to extract statistically relevant conclusions.

Research in wireless network traffic monitoring within university campuses also includes [23], where Pazl is proposed. Pazl is an indoor Wi-Fi monitoring system, which is based on mobile crowd-sensing to deliver low cost and automated wireless performance monitoring. Contrary to WiFiMon, Pazl requires the installation of a mobile application by end users.

In a recent study [24], Javed et al. created a system for crowdsourced wireless network monitoring. However, in this case, monitoring results are not gathered by the users and their devices, but from the APs as a channel utilization metric.

NetBeez [25] is a network monitoring platform that verifies end-to-end connectivity and network performance via dedicated hardware and software agents. The agents can be installed at multiple locations to monitor the status of the network and applications in real time from the end user perspective. Network engineers can then review real-time and historical performance data on the NetBeez dashboard. The NetBeez platform resembles WiFiMon performance measuring using soft and hard probes, but differs regarding metrics and granularity of data analysis capabilities, while users monitored by WiFiMon are not required to install additional software.

WiFiMon differs from related works in that: (i) it enables "non-invasive" performance measurements without the intervention of users, (ii) there is a correlation among performance data, AP identifiers and end user information (RADIUS/DHCP Logs), (iii) correlated data are posted to the same database, permitting a centralized view of Wi-Fi performance (complete

ISBN 978-3-903176-35-5 © IFIP

history) and end user behavior within the network (heat maps).

III. Baseline Design

In this section we describe the overall architecture of our system and discuss its main design features.

A. Overall Architecture

Fig. 1 depicts an overview of the WiFiMon overall architecture. It follows the well-known monitoring system architecture [26], which consists of monitoring agents and sources of measurement data (in our case WSP, WHP and WiFiMon Test Server - WTS) and the monitoring collector (WiFiMon Analysis Server - WAS). The WAS is responsible for data processing, correlation, storage and visualization. It also acts as a monitoring system controller.

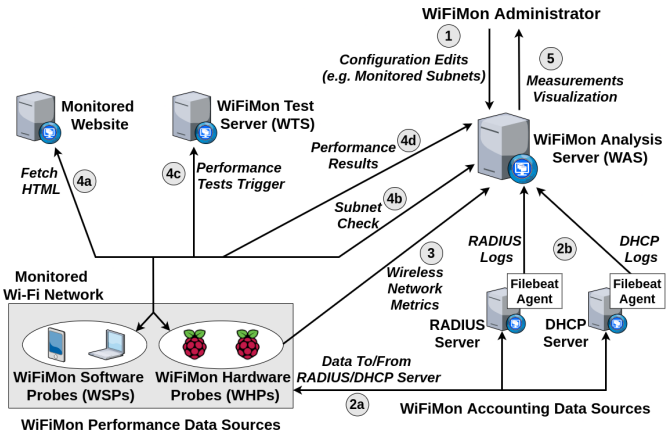

Fig. 1: The overall architecture of WiFiMon

B. WiFiMon Operation

Initially, the WiFiMon Administrator provides the configuration required for the WAS (step 1, Fig. 1). Essential configuration required by the WiFiMon Administrator includes specifying the monitored Wi-Fi subnets. WiFiMon relies on end users visiting a popular website monitored by WiFiMon, e.g. the university site for campus networks. Since the monitored website may be visited by various Internet users, this step is required to discard performance tests originating from end user devices outside the monitored Wi-Fi subnets. Thus, measurements from external networks are excluded, while data from wired network segments in the same organization are not considered as Wi-Fi monitoring results. This practice prevents WTS and WAS from being overloaded by processing irrelevant requests, but also enables focusing on the network segments that are of interest for performance evaluation.

There are three types of data processed by WiFiMon; (i) the WiFiMon accounting data, (ii) the wireless network metrics and (iii) the WiFiMon performance data. Accounting data are obtained from RADIUS and DHCP servers (step 2b, Fig. 1). These servers hold the log entries which are created when the WiFiMon performance data sources, i.e. WSPs/WHPs, connect to the monitored Wi-Fi network. Leveraging on Filebeat Agents installed within the RADIUS and DHCP servers, logs are exported to the WAS for further processing. The information from the logs is necessary when WiFiMon Administrators wish to carry out performance analysis that considers the location of end devices within the network, e.g. throughput measurements per AP in the network. Moreover, WHPs periodically collect wireless network metrics and stream them to the WAS (step 3). Collected metrics include statistics extracted from the wireless network interface of the WHPs, e.g. link quality and signal strength, along with information regarding the surroundings of the WHP, such as the nearby ESSIDs and statistics related to them.

Performance data are gathered from active performance tests executed by WSPs/WHPs and they provide estimations of the download/upload throughput and HTTP ping round trip time. Performance tests are initiated when Wi-Fi enabled end devices, i.e. WSPs and WHPs, visit a monitored website (step 4a). Upon fetching the HTML code from the monitored website, end devices check whether they are located within one of the monitored subnets by inspecting the configuration available from the WAS (step 4b). WSPs and WHPs monitored by the WAS trigger performance tests against the WiFiMon Test Server (WTS) (step 4c) that holds the test data and JavaScript code necessary for the operation of the WiFiMon performance monitoring test tools. Test data are then exchanged between WSPs/WHPs and the WTS, performance results are calculated and they are, subsequently, streamed to the WAS for further processing (step 4d). There, the results of the measurements are correlated with the accounting data to provide more accurate wireless network analysis.

C. Design Features

The main design features of the WiFiMon system are:

- Combination of crowdsourced and deterministic measurements: WiFiMon uses measurements from both WSPs and WHPs to reach stronger conclusions about the wireless network performance. While crowdsourced measurements report on how Wi-Fi is experienced by end users (WSPs), deterministic ones from probes (WHPs) monitor Wi-Fi performance from fixed points in the network.
- Correlation with RADIUS and DHCP logs respecting end user privacy: The WAS correlates performance measurements received from WSPs/WHPs with data available from the RADIUS and DHCP servers that hold the entries of devices connecting to the Wi-Fi network. Correlation enables more accurate performance analysis by delivering Wi-Fi analysis reports per AP within the network. Moreover, sensitive information originating from RADIUS and DHCP logs, i.e. end user IP and MAC addresses, are anonymized before they are stored in the WAS in order to safeguard end user privacy.
- Independence of Wi-Fi technology and hardware vendor: WiFiMon does not impose particular restrictions on the type of monitored Wi-Fi network or specific hardware requirements. However, additional benefits are provided for

IEEE 802.1X networks, e.g. eduroam, where correlation with RADIUS/DHCP logs may be performed.
- Lightweight, active monitoring that does not impact end user browsing experience: WiFiMon utilizes JavaScript-based active measurement test tools. These are triggered automatically when end devices visit a monitored website, including the HTML script tags that redirect to the WTS. Test tools are based on downloading/uploading small chunks of test data or images, hence they do not consume significant network bandwidth. Moreover, tests are initiated after the monitored website has loaded in the device of the end user. Therefore, test tools used by WiFiMon do not interfere with the user Wi-Fi experience.

IV. Implementation Details

In the following, we elaborate on the details pertaining to the WiFiMon components.

A. WiFiMon Test Server (WTS)

The purpose of the WTS is to hold the code and test data required for WiFiMon performance measurements. The WTS is implemented using an Apache web server and is based on JavaScript technology to deliver performance analysis. The test tools utilized by WiFiMon are: (i) NetTest [27], (ii) Akamai Boomerang [28] and (iii) LibreSpeed Speedtest [29].

Performance tests are initiated when end devices, i.e. WSPs and WHPs, within the monitored Wi-Fi networks visit a website. This website contains HTML script tags pointing users to initiate the performance tests with the WTS using the available JavaScript-based WiFiMon test tools. End devices exchange test data with the WTS, calculate performance metrics and stream them to the WAS for further processing. Test data include either images of diverse sizes (up to 5 MB) for NetTest and Akamai Boomerang or streams of data chunks for LibreSpeed Speetest. Test tools provide WiFiMon with reports on the download/upload throughput and HTTP ping RTT as measured by end devices.

Notably, throughput results of the WiFiMon performance measurements are expected to be lower than the maximum achievable throughput at the moment of monitoring. The reasons for this are the following:

- The WTS location. Increasing the distance, and hence the RTT, between the WTS and WSPs/WHPs reduces the maximum achievable TCP session throughput between them (tests use HTTP, therefore TCP).
- The TCP slow start and a gradually increasing congestion window mechanism. This prevents connections from quickly achieving their maximum possible throughput especially for smaller transfer units.
- The need to exchange relatively small files by the tests. Smaller sizes reduce the network overhead created by the measurements, but impact test accuracy as the files are transferred while the window size still increases.

Accuracy loss of WiFiMon measurements can be reduced by placing the WTS as close as possible to the monitored networks so that the distance between the end devices and the WTS is minimized. This way, the most realistic results can be achieved. However, even if the placement of the WTS close to the monitored network is impossible, WiFiMon will still provide useful insight regarding Wi-Fi performance by capturing the relative changes among the received measurements.

B. WiFiMon Software Probes (WSPs)

WSPs are user devices: smartphones, laptops or any other Wi-Fi enabled device in the monitored network. Devices participate in the measurements when users visit the websites monitored by WiFiMon. Since the end user device's protocol and software stacks (including operating system and browser) participate in the monitoring, the results present an objective measure of the network quality as experienced by users who surf the web [1]. WiFiMon approach does not require any additional software to be installed on user devices and measurements do not impact browsing experience. Additional traffic that is injected into the network from active probing is comparable with one average web page download. Unlike WHPs which are installed at fixed locations, WSPs can capture the perceived quality in any network section the user may be.

As end users may repetitively visit the same web page and/or frequently refresh their browser, WiFiMon components, i.e. the WTS and the WAS, could become overloaded with excessive processing. This problem is alleviated by introducing a cookie with a customizable blocking time parameter (e.g. 5 minutes) that is stored within WSPs. New tests by the same WSP are accepted by the WTS if the cookie parameter has expired or has not been set yet.

C. WiFiMon Hardware Probes (WHPs)

WHPs are small form factor single-board computers (e.g. Raspberry Pi's) with Wi-Fi capability. The purpose of the WHPs is to measure Wi-Fi performance from fixed locations within the monitored wireless network where they are installed. These measurements are useful for network administrators as they complement WiFiMon crowdsourced measurements performed by WSPs.

WHPs initiate tests to collect monitoring data about the Wi-Fi network and stream them to the WAS. Unlike WSPs, WHPs monitor the network periodically using the period defined by the WiFiMon Administrator. WHPs support two types of tests which are independent to one another. Specifically:

- Similarly to WSPs, WHPs trigger performance measurements with the WTS utilizing the same set of WiFiMon test tools. WHPs continuously monitor the Wi-Fi network from fixed points. Therefore, measurements of WHPs are regarded as deterministic since the distance between them and their associated AP remains relatively constant. Thus, a baseline comparison of the WHP measurements with the crowdsourced ones is possible. Apart from the period set by the WiFiMon Administrator, WHP test frequency is also regulated by a cookie parameter, like WSPs.
- WHPs collect data about the overall quality of the monitored Wi-Fi ESSID as well as the surrounding ones. Metrics are collected directly from the wireless network

ISBN 978-3-903176-35-5 © IFIP

interface; these involve the signal level, link quality, bit rate and transmission (TX) power reported by the WLAN NIC. Moreover, WHPs track nearby Wi-Fi ESSIDs and record their associated APs and signal strengths.

D. WiFiMon Analysis Server (WAS)

The WAS collects data from the monitored wireless networks, processes them, correlates them with the available accounting data and provides visualizations to the WiFiMon Administrator. This component receives (i) log records from the RADIUS and DHCP servers associated with the end users of the monitored network (accounting data), (ii) wireless network metrics from WHPs and (iii) performance test results from WSPs and WHPs (performance data). The location of the WAS does not affect measurement result accuracy. The overall functionality of the WAS is depicted in Fig. 2.

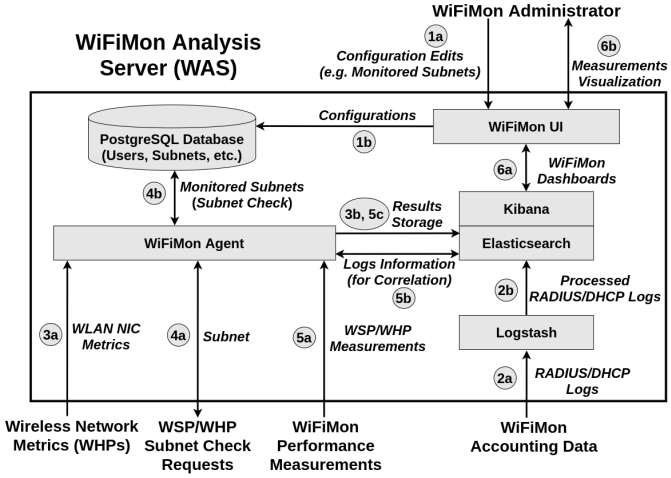

Fig. 2: The operation of the WiFiMon Analysis Server (WAS)

There are two modules that carry out the operations of the WAS. These are (i) the WiFiMon Agent, which is responsible for collecting and processing the received monitoring data, and (ii) the WiFiMon User Interface (WiFiMon UI), which depicts the results of data processing. Both of these modules are developed in Java and utilize the Spring Boot framework.

The WiFiMon Administrator initially specifies parameters essential for the operation of the WAS through the WiFiMon UI; these include (i) the subnets monitored by WiFiMon, (ii) details related to the monitored APs, (iii) WiFiMon UI users and their administrative privileges and (iv) options regarding data correlation. The aforementioned choices of the WiFiMon Administrator are stored within a PostgreSQL database.

RADIUS and DHCP logs, exported from the corresponding servers using Filebeat Agents, are preprocessed by the WAS Logstash instance. Logstash detains the data required for correlations and stores them in the WAS Elasticsearch. Information extracted from the RADIUS logs includes the IP and MAC address of end devices, i.e. the Framed-IP-Address and the Calling-Station-Id RADIUS attributes [30], the IP and MAC address of the associated AP, i.e. the NAS-IP-Address and Called-Station-Id RADIUS attributes [30] and the RADIUS event timestamp. From the DHCP logs Logstash extracts the IP and MAC addresses of end users as well as the DHCP event timestamp. Personally Identifiable Information (PII), i.e. user IP and MAC addresses, are hashed using the HMAC-SHA-512 implementation of the Logstash fingerprint plugin [31] and a secret password known to the WiFiMon Administrator. Thus, PII is stored anonymized in the WAS. Additionally, the logs are secured in transit using a TLS-encrypted channel based on the X-Pack Elastic Stack extension [32]. The overall procedure of obtaining and processing the logs is illustrated in Fig. 3.

The WiFiMon Agent receives wireless network metrics and performance measurements from the monitored networks. Wireless metrics from WHPs are simply converted into the appropriate storage format and they are fed in Elasticsearch. Performance measurements received from WSPs/WHPs are correlated with the exported logs. Thus, identifying performance per network AP becomes possible. WiFiMon Administrators may opt between two types of correlation: based on (i) Framed-IP-Address, i.e. the end device IP, thus depending solely on RADIUS logs, or (ii) Calling-Station-Id, i.e. the end device MAC address, therefore utilizing both RADIUS and DHCP logs. Concerning the first option, when a performance measurement is received, the available RADIUS logs are searched to determine the most recent record containing a Framed-IP-Address that matches the observed end device IP. As the Framed-IP-Address attribute may not always be present in RADIUS logs, there is a second option that leverages on DHCP logs. When a performance measurement is received, DHCP logs are searched to find the MAC address of the end device based on the observed IP. Then, RADIUS logs are utilized to locate the most recent record including a Calling-Station-Id that equals the MAC address of the end device.

Since Logstash stores sensitive data extracted from logs anonymized, correlation comparisons in the WiFiMon Agent are performed directly on hashed strings. Thus, the IP addresses obtained from performance measurements are first hashed with HMAC-SHA-512 algorithm utilizing the same secret password used by Logstash fingerprint plugin. Data from RADIUS logs are combined with the user measurements and the correlation results are stored in Elasticsearch.

Finally, the WiFiMon UI retrieves the WiFiMon dashboards from Kibana and offers visualizations of the processed data to the WiFiMon Administrator. Such visualizations are: throughput per WHP, crowdsourced measurements per monitored subnet and measurements per operating system and browser.

V. Evaluation

In this section, we demonstrate the capabilities of WiFiMon in detecting Wi-Fi performance degradation. The evaluation is based on two WiFiMon pilots that took place in major GÉANT networking events with several hundreds of participants; the TNC19 Conference and the GÉANT Symposium 2020.

In the sequel, we elaborate on details regarding the pilots, present their results and discuss their outcomes. Notably, the results depicted hereafter are averaged over all WiFiMon test tools, i.e. NetTest, Boomerang and Speedtest.

ISBN 978-3-903176-35-5 © IFIP

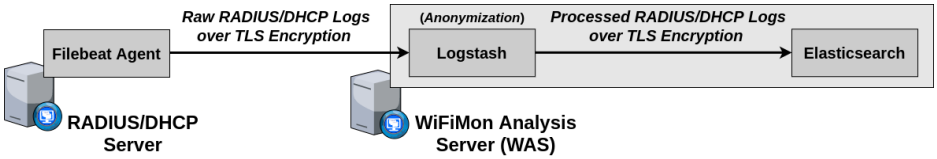

Fig. 3: Exporting and processing RADIUS and DHCP logs

A. WiFiMon Pilot at TNC19

The first WiFiMon pilot was held in Tallinn, Estonia, during TNC19 [33], June 17-19 2019. The conference had more than 800 participants. The pilot monitored the conference Wi-Fi network that was setup for the conference days using only WHPs. Wi-Fi at the venue was based on Cisco Aironet 3700 with the controller installed locally.

During the pilot, five WHPs (Raspberry Pi 3 model B devices), have been distributed across the rooms of the venue. These rooms included (i) the main hall which hosted the opening/closing plenary sessions and other major sessions, e.g. lightning talks which gathered the majority of the conference participants, (ii) the area where the coffee/lunch breaks and an opening ceremony took place as well as (iii) other rooms of the venue that hosted smaller sessions. Each WHP triggered performance measurements every 20 minutes. The WTS component was installed in the Tallinn University of Technology, while the WAS was hosted at SWITCH, the Swiss NREN. This setup met the requirements of WiFiMon as the WTS was placed close to the conference venue, hence the RTT between the WHPs and the WTS was small (less than 4 msec).

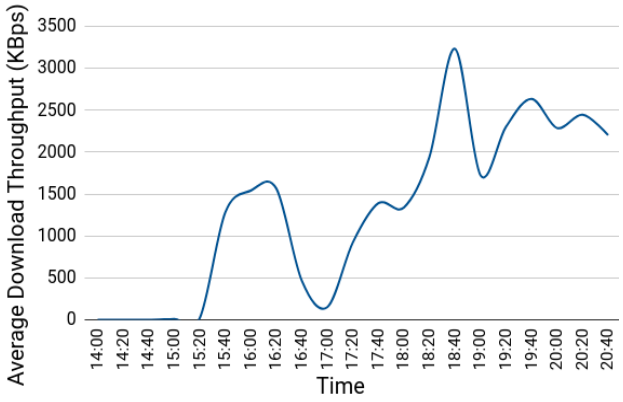

Fig. 4: Average download throughput reported by the WHP placed within the main hall of the TNC19 venue

Fig. 4 depicts the average download throughput measured by the WHP placed in the main hall during the first conference day between 14:00 and 20:40 hours. Notably, APs were not installed directly in the main hall and Wi-Fi connectivity there was provided by those present in the adjacent spaces where coffee/lunch breaks and the opening ceremony occurred.

Fig. 4 shows the expected throughput metrics behavior during an event where a lot of participants gather in one place at the same time: throughput fell down in those places where participants gathered and used Wi-Fi heavily, often with multiple devices. Between 14:00 and 15:20, during the lightning talks session in the main hall which attracted nearly all conference participants, the WHP reported almost zero throughput. Notably, apart from measurements reporting low throughput, we observed that many measurements were lost during this time interval, presumably due to lost connectivity of the WHP. The installed APs were not able to connect all the users' devices in the main hall. Between 15:20 and 16:30, we noticed that the venue was less crowded since lots of participants had left the conference. Consequently, we observe that the Wi-Fi performance was restored during this period. Then, the average download throughput dropped again at 17:00 when the opening reception, which was organized in the same space, started. Finally, the Wi-Fi performance was restored after 18:00 when most of the participants had left the venue.

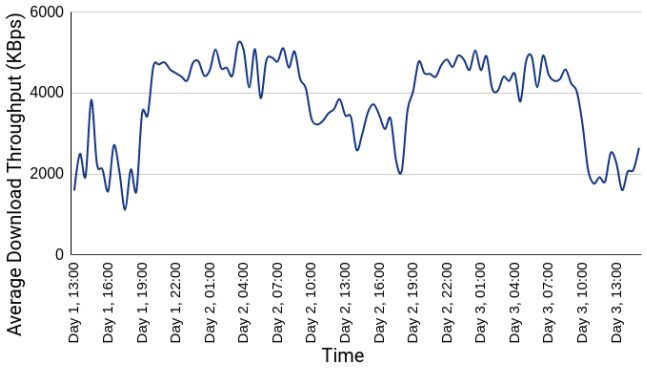

Fig. 5: Average download throughput reported by the WHP placed within the room of the TNC19 venue where coffee/lunch breaks and the opening ceremony took place

Fig. 5 depicts the average download throughput of the WHP placed in the room where the coffee/lunch breaks and the opening ceremony happened for the larger part of the conference duration (Day 1 13:00 up to Day 3 15:00). Similarly, it can be observed that the Wi-Fi performance degraded when people were at the venue, while the throughput was higher and more stable when participants were absent. In the periods when the venue was empty, download throughput fluctuated between 4 and 5 MBit/s, with an average value of 4.61MBit/s and standard deviation of 0.35 MBit/s. This value can be considered as a baseline value for the achievable download throughput of our setup which should be used to estimate later the performance degradation. This throughput is by no means the maximum achievable throughput in the Wi-Fi network. It is the metric obtained in this specific situation in which the WTS was not at the venue and with relatively small transfer

ISBN 978-3-903176-35-5 © IFIP

files selected in order not to create a too large traffic overhead.

B. WiFiMon Pilot at GÉANT Symposium 2020

The second WiFiMon pilot took place in Ljubljana, Slovenia, during the GÉANT Symposium 2020, February 4-5 2020. During the pilot, WiFiMon monitored the performance of the eduroam ESSID. The Wi-Fi setup was the same as in Tallinn, however on a smaller scale as there were around 250 participants at the symposium and the venue was a hotel.

During the pilot, we placed seven WHPs (Raspberry Pi 3 model B devices), across the rooms of the symposium venue. Similarly to our first pilot, these rooms included the main hall that hosted the major sessions, the room in which coffee and lunch were served and rooms with smaller sessions. WHPs triggered a measurement every 5 minutes. In addition to TNC19, the pilot added measurements performed by WSPs, thus combining crowdsourced results with deterministic ones obtained from WHPs. The necessary HTML lines to perform the WSP measurements were injected in the symposium agenda which was often visited by the symposium participants; their consent was obtained via the online registration platform. Both the WTS and WAS were installed in ARNES (Slovenian NREN) premises, just a couple of kilometers from the venue, and were connected with an uncongested 1Gbps link.

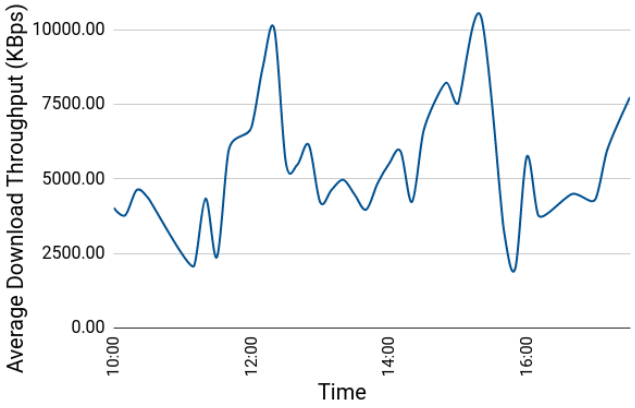

Fig. 6: Average download throughput of the crowdsourced measurements in the GÉANT Symposium 2020 pilot

Fig. 6 depicts the average download throughput reported by the crowdsourced measurements (WSPs) during the first symposium day between 10:00 and 17:30. The average download throughput dropped significantly in two time periods, between 11:00 and 11:40 as well as between 15:30 and 16:00. As both periods followed a coffee break, we may assume that more participants were visiting the symposium agenda at approximately the same time and position within the venue to decide the next session to attend. A notable drop also exists between 12:30 and 14:00 hours, i.e. during and right after lunch time when most participants gathered in less space. Around 12:20 and 15:20, higher levels of download throughput were reported by crowdsourced measurements since participants were distributed across many different sessions within the symposium venue and, presumably, specific APs were not overloaded. After 17:00, most people left the venue as the symposium day had ended. This is apparent in the crowdsourced measurement results since Wi-Fi performance raised to greater levels.

Fig. 7 depicts the average download throughput reported by WHPs #2 and #5 during the same day as on the measurements of Fig. 6. We observe that both WHPs follow similar trends and they were able to conceive the throughput drops reported by the WSP measurements of Fig. 6 around 11:00 and 16:00 hours. WHPs were placed near the available power plugs of the venue which were typically farther from the APs than the audience. Thus, WHPs reported less throughput than WSPs.

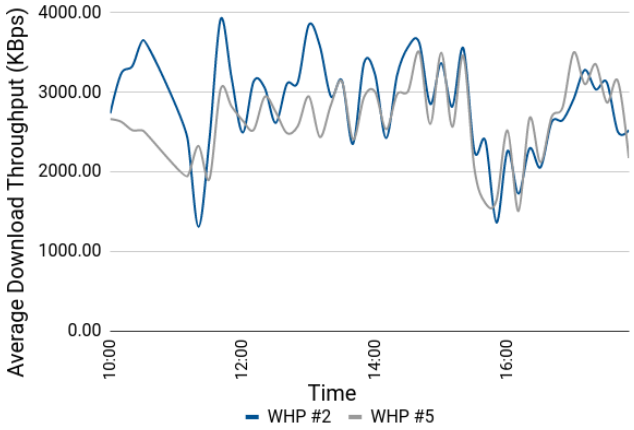

Fig. 7: Average download throughput of WHPs #2 and #5 placed within the main hall of the GÉANT symposium 2020 venue

Finally, Table I provides the average values of the wireless network metrics and the averages of the performance measurements as reported by each WHP of our setup. Wireless metrics included signal level, bit rate, link quality and TX power, while performance measurements involved download/upload throughput and ping latency. The considered time interval is the first symposium day between 10:00 and 18:00 hours. We observe that the trends of the wireless network metrics did not necessarily follow those of the performance measurements. Indicatively, while WHP #1 reported the best average link quality, its throughput results were among the worst. Moreover, while WHP #5 had the worst link quality, the measured throughput results were among the best. Correlation coefficients between the throughput and the wireless interface metrics time series for the measurements given in Fig. 7 were in the range between -0.12 and 0.1, meaning almost no correlation at all. This brings us to a conclusion that for the proper evaluation of Wi-Fi performance it is necessary to combine multiple sources of performance information: Crowdsourced and deterministic measurements from WSPs and WHPs respectively are essential to reach solid conclusions about the Wi-Fi performance and perceived network quality as high values of signal strength and link quality do not necessarily guarantee high Wi-Fi throughputs [34].

VI. Conclusion & Future Work

We described WiFiMon, an open source system for monitoring Wi-Fi performance without requiring the intervention

ISBN 978-3-903176-35-5 © IFIP

TABLE I: Wireless network metrics and performance results reported by WHPs in the GÉANT Symposium 2020 pilot

WHP No	Average Signal Level (dBm)	Average Bit Rate (Mbps)	Average Link Quality	Average TX Power (dBm)	Average Download Throughput (KBps)	Average Upload Throughput (KBps)	Average Ping Latency (msec)
1	-43	71	67/70	31	1588	763	48
2	-52	49	58/70	31	2883	1500	30
3	-59	78	51/70	31	2644	1429	44
4	-59	59	51/70	31	1431	650	41
5	-66	75	44/70	31	2678	1514	23
6	-62	65	48/70	31	1758	890	41
7	-55	66	55/70	31	2730	1562	32

of end users. WiFiMon combines crowdsourced measurements collected from end users with deterministic ones gathered from hardware probes placed at fixed locations within the network. Moreover, WiFiMon correlates performance measurements with data available from the RADIUS and DHCP logs associated with the end devices within the monitored networks, therefore delivering fine-grained analytics. We validated the capabilities of our system based on data obtained from recent WiFiMon pilots in highly attended events, i.e. TNC19 and GÉANT Symposium 2020. Based on the outcomes of our pilots that monitored a large number of diverse users and devices, we conclude that our system is capable of detecting performance fluctuations in Wi-Fi networks.

In the future, we plan to enrich the capabilities of WiFiMon. Specifically, we will investigate the utilization of time series analysis and/or machine learning methods to promptly predict Wi-Fi outages. Furthermore, we will research additional options suitable for Wi-Fi performance monitoring, thus enriching the current WiFiMon toolset. Finally, we will install WiFiMon in campus/enterprise networks and, thus, test our system for longer time periods, while including information from RADIUS and DHCP logs, omitted from our pilots.

REFERENCES

[1] ITU-T Recommendation G.1011, "Reference Guide to Quality of Experience Assessment Methodologies", July 2016, https://www.itu.int/rec/T-REC-G.1011-201607-I/en, accessed 2/2021

[2] ITU, "Quality of Service Regulation Manual", 2017, https://www.itu.int/pub/D-PREF-BB.QOS_REG01-2017, accessed 2/2021

[3] T. Hoßfeld and S. Wunderer, eds., "White Paper on Crowdsourced Network and QoE Measurements – Definitions, Use Cases and Challenges (2020)", March 2020, Würzburg, Germany, doi: 10.25972/OPUS-20232

[4] eduroam Homepage, https://www.eduroam.org/, accessed 2/2021

[5] C. Bouras, K. Baumann, V. Kokkinos, N. Papachristos and K. Stamos, "WiFiMon: A Tool for Wi-Fi Performance Monitoring and Verification", in the Int. Journal of Wireless Networks and Broadband Technologies (IJWNBT), IGI Global, Vol. 8, Iss. 1, pp. 1-18, 2019

[6] WiFiMon Homepage, https://wiki.geant.org/display/WIF

[7] H. Ma, D. Zhao and P. Yuan, "Opportunities in Mobile Crowd Sensing" in IEEE Communications Magazine, Vol. 52, Iss. 8, pp. 29-35, August 2014

[8] D. R. Choffnes, F. E. Bustamante and Z. Ge, "Crowdsourcing Service-Level Network Event Monitoring", in the ACM SIGCOMM 2010, pp. 387-398, New Delhi, India, August 2010

[9] Aruba AirWave, https://www.arubanetworks.com/products/network-management-operations/airwave/, accessed 2/2021

[10] Paessler PRTG Network Monitor, https://www.paessler.com/wifi-monitoring, accessed 2/2021

[11] NetSpot, https://www.netspotapp.com/, accessed 2/2021

[12] Acrylic Wi-Fi, https://www.acrylicwifi.com/en/, accessed 2/2021

[13] NetAlly AirMagnet WiFi Analyzer PRO, https://www.netally.com/products/airmagnet-wifi-analyzer/, accessed 2/2021

[14] Wifi Analyzer and Scanner, Microsoft, https://www.microsoft.com/en-us/p/wifi-analyzer-and-scanner/9nblggh5qk8q?activetab=pivot:overviewtab, accessed 2/2021

[15] Riverbed Wi-Fi Inspector, https://www.riverbed.com/au/archives/forms/trial-downloads/xirrus-wifi-inspector.html, accessed 2/2021

[16] LizardSystems Wi-Fi Scanner, https://lizardsystems.com/wi-fi-scanner/, accessed 2/2021

[17] MetaGeek inSSIDer, https://www.metageek.com/products/inssider/, accessed 2/2021

[18] S. Sonntag, J. Manner and L. Schulte, "Netradar - Measuring the Wireless World", in the 11th Int. Symposium and Workshops on Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks (WiOpt), IEEE, pp. 29-34, Tsukuba Science City, Japan, 2013

[19] Speedtest by Ookla, https://www.speedtest.net/, accessed 2/2021

[20] A. Saeed, S. F. Naseem and Z. R. Zaidi, "Mobility Estimation for Wireless Networks using Round Trip Time (RTT)", in the 6th International Conference on Information, Communications and Signal Processing (ICICS), IEEE, pp. 1-5, Singapore, Singapore, 2007

[21] S. Park, H. S. Ahn and W. Yu, "Round-Trip Time-based Wireless Positioning without Time Synchronization", in the 2007 International Conference on Control, Automation and Systems, IEEE, pp. 2323-2326, Seoul, Korea, October 2007

[22] cnlab Speedtest, https://speedtest.cnlab.ch/en/, accessed 2/2021

[23] V. Radu, L. Kriara and M. K. Marina, "Pazl: A Mobile Crowdsensing Based Indoor WiFi Monitoring System", in the 9th International Conference on Network and Service Management (CNSM 2013), IEEE, pp. 75-83, Zurich, Switzerland, October 2013

[24] Z. Javed, Z. Khan, J. J. Lehtomäki, H. Ahmadi and E. Hossain, "Eliciting Truthful Data From Crowdsourced Wireless Monitoring Modules in Cloud Managed Networks", in IEEE Access, Vol. 8, pp. 173641-173653, September 2020

[25] NetBeez Wi-Fi Monitoring, https://netbeez.net/wireless-monitoring/, accessed 2/2021

[26] P. Eardley et al., "A Framework for Large-Scale Measurement of Broadband Performance (LMAP)", RFC 7594, September 2015

[27] NetTest Code Repository, https://code.google.com/archive/p/nettest/, accessed 2/2021

[28] Akamai Boomerang GitHub Repository, https://github.com/akamai/boomerang, accessed 2/2021

[29] LibreSpeed Speedtest GitHub Repository, https://github.com/librespeed/speedtest, accessed 2/2021

[30] P. Congdon, B. Aboba, A. Smith, G. Zorn and J. Roese, "IEEE 802.1X Remote Authentication Dial In User Service (RADIUS) Usage Guidelines", RFC 3580, September 2003

[31] Logstash Fingerprint Filter Plugin, https://www.elastic.co/guide/en/logstash/current/plugins-filters-fingerprint.html, accessed 2/2021

[32] X-Pack Elastic Stack Extension, https://www.elastic.co/guide/en/elasticsearch/reference/current/setup-xpack.html, accessed 2/2021

[33] TNC19 Conference, https://tnc19.geant.org/, accessed 2/2021

[34] CommScope Blog, "Understanding Wi-Fi Signal Strength vs. Wi-Fi Speed", https://www.commscope.com/blog/2015/understanding-wi-fi-signal-strength-vs.-wi-fi-speed/, accessed 2/2021

ISBN 978-3-903176-35-5 © IFIP