Characterising Eduroam WLANs Usage Trends: A Case Study

Marangaze Munhepe Mulhanga, Solange Rito Lima, Paulo Carvalho
University of Minho, Department of Informatics, 4710–057 Braga, Portugal
e-mail:{solange,pmc}@di.uminho.pt

***Abstract*—The recent commitment of European Universities to join a global wireless access and roaming infrastructure through the *eduroam initiative* has brought many advantages to the academic community. In a new environment of increasing user mobility, it is useful to study and understand the users' behaviour when accessing the wireless network infrastructure. This paper is focused on characterising the real network usage within university campuses, based on WLAN traffic traces collected at the University of Minho (Portugal). The obtained results, allowing to identify user trends and quantify the influence of access location on the number of associated users, the number and duration of sessions, the volume of generated traffic, and user mobility indicators, provide important directions for planning future improvements of university WLANs.**

I. Introduction

The need for easy and ubiquitous access to information and the low to moderate cost of personal devices are factors influencing the steady growth of wireless technology usage. Network Access Points (APs) tend to be critical points in WLANs due to multiple aspects such as user mobility, traffic dynamics, location and density of devices, which impact on network performance. Understanding traffic characteristics and the usage of network resources is an essential step to assure quality of service (QoS) and improve quality of experience (QoE) of end users when accessing network services.

Within university context, Tang and Baker [1], based on data collected during twelve weeks in one building of Stanford University, studied the behaviour of WLAN users, namely the benefits of mobility, the volume and characteristics of traffic involved. QoS metrics such as delay and bandwidth were also measured. Kotz and Essien [2] work reports the analysis of data traffic collected during eleven weeks from all buildings at the Dartmouth University. Later on, the work reported in [3] states that the applications used over the already mature WLAN changed dramatically. Schwab and Bunt [4] studied WLAN usage patterns at the University of Saskatchewan, comprising a small number of APs (18) strategically placed, in order to plan network expansion. Papadopouli et al. [5] investigated roaming activity at aggregate level in the University of North Carolina network infrastructure. Based on syslog data from three monitoring periods (between 2004 and 2005), the authors identify the regions with high roaming activity and derive topological models of the university infrastructure, involving 488 APs. Kumar et al. [6] study classifies users into social groups and investigates the WLAN usage behaviour of these groups in the USC campus (MobiLib). Based on a month long WLAN trace, the authors analysed the differences on the average session duration for male and female users across the campus. Kim and Helmy [7], based on traffic traces collected from Dartmouth University WLAN between 2001 to 2006, studied how changes in wireless devices and network affect WLAN users, and influence location prediction.

Facing the constant evolution of wireless devices, number of users, variety of applications and services, and human behaviour regarding wireless technologies, it is relevant to keep an up-to-date analysis of today's university WLANs usage, in particular, within the *eduroam* context. The present study is a further step in this direction, presenting a traffic analysis and characterisation study involving WLANs of University of Minho (Portugal), within the *eduroam (Education Roaming) initiative* [8]. The analysis was carried out based on real traffic traces gathered between April and June of 2010, which corresponds to a typical academic term. The study assesses several network usage metrics related to AP utilisation, session characterisation, access location influence, and user mobility patterns. This analysis, when compared with other case studies, aims to provide guidelines for planning future university WLANs deployment.

II. Case Study: Eduroam at University of Minho

The University of Minho (UMinho), located in the north of Portugal, was founded in 1973 and started its academic activity in 1975. Currently, with a population of nearly 15,000 students, 1,200 teachers, and 600 technical and administrative staff, it is one of the biggest Portuguese universities. The academic and scientific activities at UMinho are developed in two campuses: the campus of Gualtar in Braga and the campus of Azurém in Guimarães. The students' accommodation buildings have capacity to accommodate 1400 students, around 60% in Braga and 40% in Guimarães.

Both campuses participate in the eduroam initiative, which provides wireless network access for research and education to the university population and visitors. The WLAN infrastructure comprises a total of 429 APs, 310 located in Braga and 119 in Guimarães. The WLAN technology used in both campuses is based on IEEE 802.11b, 802.11g and, more recently, 802.11n in strategic locations, such as libraries.

The core of network operation is located in Gualtar, where the main network services are assured to users inside and outside the campus, providing a 10Gbps access to the Internet. A 768Mbps link interconnects Gualtar and Azurém campuses.

II.1. Data collection strategies and tools

The data collection process involved the articulation of several technologies, namely: (i) *SQL database*: the information of eduroam authenticated users, directly accessed from a Radius server, is stored in an SQL database on an AP-basis. This data includes the MAC addresses from wireless devices of authenticated users, and corresponding events (such as association, disassociation, roaming, etc.), the number and duration of user associations, the initial and final APs involved in a session, and the traffic volume in packets and bytes per AP. The WLAN traffic history can be obtained through an SQL server; (ii) *DHCP logs*: private or public IP addresses are dynamically assigned to wireless devices on campus by a DHCP server. The DHCP logs contain IP addressing assignments to MAC addresses, including a timestamp of that occurrence: (iii) *SNMP-based tool*: a proprietary tool allows monitoring APs at regular intervals (configured for polling each five minutes). This tool performs traffic collection of associated users regardless its authentication status. (iv) *Port Mirroring*: for analysis at protocol level, traffic was captured using `tcpdump` via port mirroring at main router providing connectivity between Gualtar and Azurém campuses. Traffic analysis was carried out resorting to Colasoft's Capsa Network Analyzer; (v) *Data Confidentiality*: the use of anonymizing tools to pre-process network traces, e.g. replacing consistently MAC addresses, allowed to perform WLAN analysis without compromising data confidentiality.

III. Results of Data Analysis

III.1. Classification criteria

A first step toward WLANs characterisation is to group and classify APs. This work considers a classification criteria according to the location, activity sector and usage degree of APs. University buildings at UMinho are usually associated with a main activity, e.g. Department of Civil Engineering. Using this first classification criterion, 30 distinct locations were identified. APs were also grouped per activity sector, resulting in six distinct sectors, namely: (i) *Social* - bars, canteens and sports facilities; (ii) *Residential* - student accommodation premises; (iii) *Services* - administrative, technical and student support services; (iv) *Library* - libraries; (v) *Research* - research laboratories; and (vi) *Academic* - departments and schools. A third classification criterion was based on the number of distinct users registered in each AP, which led to the definition of five distinct groups:

```
Group 1: APs with number of users ≥ 500
Group 2: APs with 300 ≤ number of users < 500
Group 3: APs with 100 ≤ number of users < 300
Group 4: APs with 50 ≤ number of users < 100
Group 5: APs with 0 ≤ number of users < 50
```

The following sections detail the results of analysing and characterising the traffic traces from the Azurém Campus.

III.2. Analysis of associated users and location

The analysis of the number of different mobile users associated with each AP reveals an asymmetric behaviour, ranging from 928 different users on the busiest AP to 2 users on underutilised APs. Considering the usage groups defined above, the results show that the APs utilisation follows approximately a normal distribution centred on Group 3 with 38% of AP associations. A total of a 13.4% APs handle more than 500 users, whereas around 16% handle less that 50 users. This analysis also indicates that 56% of APs on campus support between 100 and 500 users. Figure 1 shows the values recorded for the different intervals.

In addition, for each group, the spatial distribution of APs on campus was assessed. Figure 1 illustrates the obtained results when considering 30 spatial areas (sites). As shown, the busiest APs cover about 30% of the total area, whereas less utilised APs (Group 4 and 5) around 20%.

III.3. Analysis of traffic volume per activity sector

The analysis of WLAN traffic volume (eduroam) on campus was carried out considering the APs grouped in the six classes defined according to their location. Figure 2 shows the distribution of the volume of inbound and outbound traffic by type of location, with a clear dominance of inbound traffic. As illustrated, the residential and academic sectors are responsible for most of inbound traffic, with nearly 80% of the total. Regarding the volume of outbound traffic, residential and academic sectors are again the major contributors to the overall traffic load.

III.4. Analysis of sessions

An important aspect of users' pattern behaviour is the time users remain associated with each AP. This time analysis, based on the duration of user sessions, resorts to aggregated data comprising the total number of sessions and their average duration. Crossing this information with the number of users registered per AP, it is possible to evaluate the average time of each user session. It was observed that the AP with the highest monthly associated time per user was approximately 58 hours, resulting from the sum of its individual sessions.

The average duration of sessions varies between 1 and 110 minutes. Figure 3 shows a histogram representing the number of APs for distinct session time intervals. As illustrated, five representative time intervals were identified, showing that, from a total of 119 APs, 34 (28%) handle sessions with average duration from 10 to 15 minutes. Based on this information it was found that: (i) APs with a low number of users (Groups

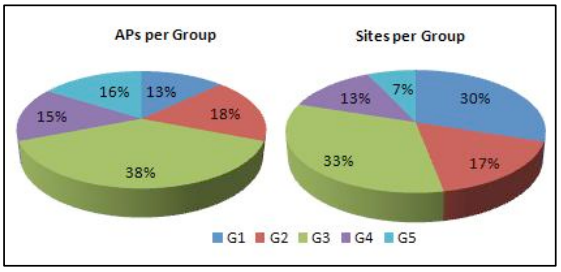

Fig. 1. Distribution of APs and sites per group (%)

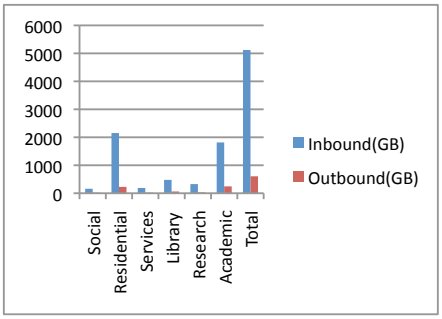

Fig. 2. Distribution of traffic volume per activity sector

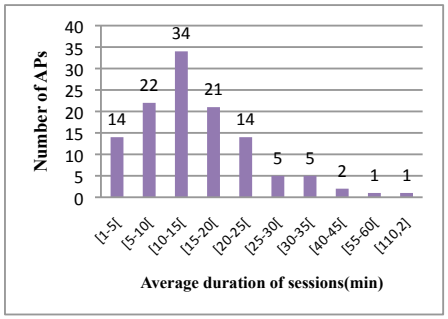

Fig. 3. Number of APs per average time duration of the sessions

4 and 5) support a large number and duration of sessions. These observations correspond mainly to APs associated with the Residential sector and other APs in locations of reduced mobility; (ii) the average length of sessions does not exhibit major variations, especially in Groups 1 and 2. The largest deviation occurs in all AP with few users. There are sessions of long duration depending on the AP location (e.g. home or research center), and sessions of short duration in APs with many users.

A fine grain analysis of WLAN users' behaviour was also carried out through the observation of user and session variations in small timescales. Taking a weekly traffic trace as an example, traffic data was analysed on a daily basis, hour-by-hour. The variation of the number of users during working days follows a pattern with a daily peak at 11/12 and 16pm, with about 400-550 wireless users. The low activity period occurs at night, with a minimum number of users at 4am (about 15-20 users), mainly from student residences. The variation on the number of sessions presents a similar behaviour. During busy hours, the average number of sessions per user is around 4, increasing to 7 overnight.

III.5. Mobility of Users

To assess to what extent users take advantage of mobility in the WLAN, we identified (by MAC address) the users that presented high mobility patterns during April and May. We analysed the number of associations of one user with two or more different APs. The results show that most of authenticated users (90% of a total of 3480 users) exhibited effective mobility and only a reduced percentage (9%) used a single AP to access the WLAN.

Attending to the high percentage of users associated with distinct APs, a more detailed analysis was carried out, identifying the number of users per mobility scenario, grouped in intervals of number of associated APs. As shown in Figure 4, a significant number of users (around 37,5% of the total) visited from 2 to 5 APs. Conversely, a single user may visit monthly a large number of different locations on campus, reaching up to 60 APs visited. For users with the highest mobility indicators (corresponding to the slice "other ranges"), approximately 37% were identified as the same users in April and May.

IV. Conclusions and Future Work

This paper reports ongoing work on the characterisation of University WLANs usage, within the eduroam initiative. The study, based on real traffic traces collected from UMinho WLANs, has covered and assessed several network usage metrics related to APs utilisation, session characterisation, access location influence, and user mobility within campuses. These metrics are relevant to better understand both the real utilisation of the WLAN infrastructure and the characteristics in terms of user behaviour, providing a useful feedback for planning, dimensioning and managing university WLANs. Future work includes expanding the characterisation of UMinho WLAN traffic, including the study of mobility patterns, protocols and applications per activity sector and user location.

References

[1] Diane Tang and Mary Baker. Analysis of a local-area wireless network. In *MOBICOM*, pages 1–10, 2000.
[2] David Kotz and Kobby Essien. Analysis of a Campus-wide Wireless Network. In *ACM Mobicom*, pages 107–118. ACM Press, 2002.
[3] Tristan Henderson, David Kotz, and Ilya Abyzov. The Changing Usage of a Mature Campus-Wide Wireless Network. In *MOBICOM'04*, 2004.
[4] David Schwab and Rick Bunt. Characterising the Use of a Campus Wireless Network. In *INFOCOM'04*, 2004.
[5] Maria Papadopouli, Michael Moudatsos, and Merkourios Karaliopoulos. Modeling Roaming in Large-scaleWireless Networks Using Real Measurements. In *Proceedings of WOWMOM'06*, 2006.
[6] Udayan Kumar, Nikhil Yadav, and Ahmed Helmy. Gender-based feature analysis in campus-wide WLANs. *SIGMOBILE Mob. Computer Communications Review*, 12:40–42, January 2008.
[7] Jeeyoung Kim and Ahmed Helmy. Analyzing Mobility Evolution in WLAN Users: How Predictable Are We? *SIGMOBILE Mob. Computer Communications Review*, 14:10–12, December 2010.
[8] eduroam. Eduroam Initiative, http://www.eduroam.org, 2003.

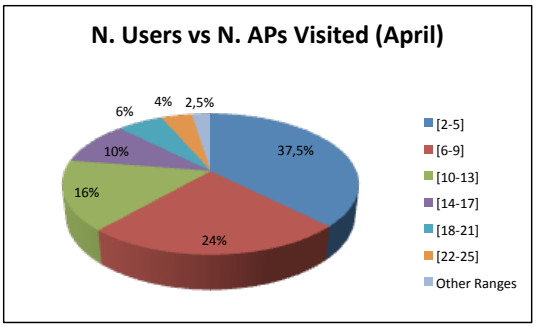

Fig. 4. Relative view of Users Mobility