

# Multi-level application-based traffic characterization in a large-scale wireless network

Manolis Ploumidis<sup>a,b \*</sup> Maria Papadopoulou<sup>a,b †</sup> Thomas Karagiannis<sup>c ‡</sup>

## Abstract

**With the increasing deployment of wireless networks, network management and configuration of wireless Access Points (APs) has become one of the main concerns of network operators. While statistics and measurements regarding the overall usage of individual APs are readily available, the limited knowledge of the wireless traffic demand, in terms of the type of application, hinders efficient network provisioning. This paper provides an extensive application-based characterization of a large-scale wireless network, going beyond the port-number limitation, across three levels, namely, network, clients, and APs. We found that the most popular application types, in terms of the number of flows, bytes, and clients, are Web and peer-to-peer, and while the majority of APs is dominated by them, APs of the same building type have large differences in their traffic mix. File transfer flows, such as FTP and P2P, are heavier in wired than in wireless networks. Finally, an interesting dichotomy among APs, in terms of their dominant application type and downloading and uploading behavior was observed.**

## 1 Introduction

Wireless local area networks (WLANs) are increasingly being deployed to address the growing demand for wireless access. As their user population increases, accurate traffic modeling and characterization of their workload are essential to facilitate efficient network management and better utilization of their scarce resources. While there have been several studies looking at the application cross-section at wired networks (e.g., [8]), such attempts are limited in the case of wireless networks [11]. Besides

the well-documented limitations of application identification [16], inherent additional complications in wireless networks, such as the increasing overheads of data collection due to the need of multiple monitoring points, cross-correlation of different type of traces, and transient phenomena due to the radio propagation and mobility, have led the community to assume that the expected workload of wireless networks follows the general trends of Internet applications.

This paper provides a multi-level application-based traffic characterization of a campus-wide wireless infrastructure from the perspective of the network, client, and AP. Such a characterization is crucial to our understanding of the application usage, build user profiles and develop better resource management and admission control mechanisms.

Using the port number to classify flows may lead to significant amounts of misclassified traffic due to the following reasons: Many modern applications, such as modern P2P protocols or streaming applications, use dynamic ports to communicate. Even worse, port ranges of different applications may overlap. Apart from this, numerous applications, such as malware or P2P, may try to masquerade their traffic under well-known “non-suspicious” ports, such as the port 80. To avoid this “known-port limitation” [19, 15], we employed the *BLINC* tool [16] which performs classification of flows into applications based on the transport-layer footprint of the various application types. To identify AP- and user-specific characteristics, flow-related statistics were coupled with client SNMP traces.

Our findings include: The most popular applications are Web browsing and P2P accounting approximately for 81% of the total traffic. Most users are also dominated by these two applications. Network management (NM) and scanning activity are responsible for 17% of the total flows in our trace. While building-aggregated traffic application usage patterns appear similar, the application cross-section varies within APs of the same building. Most wireless clients appear to use the wireless network for one specific application that dominates their traffic share. File transfer flows, such as FTP and P2P, are heavier in the wired network than in the wireless one. An interesting dichotomy among APs, in

<sup>\*</sup>a. Department of Computer Science, University of Crete. (ploumid@ics.forth.gr)

<sup>†</sup>b. Institute of Computer Science, Foundation for Research and Technology-Hellas. (mgp@ics.forth.gr)

<sup>‡</sup>c. Microsoft Research. (thomkar@microsoft.com)

terms of their dominant application type and downloading and uploading behavior, was observed.

The rest of the paper is organized as follows. Section 2 presents related work while Section 3 describes the analyzed traces and our methodology. Sections 4, 5 and 6 present our findings for the aggregate, AP and client traffic characterization. Finally, we discuss the implications of our findings in Section 7.

## 2 Related work

While there have been several measurement attempts in deployed wireless networks over the past few years [12, 21], the focus has mostly been in identifying roaming patterns [18, 4, 3, 24, 7, 22, 23, 20], traffic load characterization [24, 7, 17, 12], packet losses and delays [13, 3, 2, 5] and forecasting [21], with the goal of provisioning the wireless APs. These studies reflect a variety of wireless environments, such as campuses, auditoriums, and enterprise networks. While our work touches upon the workload of APs, we are interested in studying the application cross-section across APs rather than the overall traffic.

Closer to this research, there have been studies on port-based application classification [11, 3, 24] which were shown to be highly inaccurate [19, 15]. Furthermore, our work goes beyond simple application breakdown by examining the variation of the application cross-section across APs and clients. Tutschku investigated the difference of the uploading from the downloading traffic of a popular P2P application in a wired network and reported a significant amount of uploaded P2P traffic in [25]. Guo *et al.* [10] studied asymmetry in bit torrent traffic and concluded that it is highly affected by high speed downloading. Chambers *et al.* [6] characterized the online games usage, in terms of user sessions and periodicities of the workload. Our client profiling is based on their traffic mix across applications and the existence of an application through which they transfer the majority of their traffic. Finally, Guha *et al.* [9] studied the structure of a VoIP based network and contrasted it to P2P systems. Our method distinguishes chat and streaming applications from P2P by focusing on the graph that is created by flows among clients rather than using network topology information.

## 3 Data description

We analyzed traces from one of the largest campus wireless networks [1]. Out of the total 488 APs, 382 were monitored that were located in 231 different buildings. Our traces consist of packet traces collected at one of the access routers at UNC and client SNMP data from all APs. Client SNMP data were used to add MAC address and AP-based

information to the flow-level data. Our traces correspond to approximately 9,125 distinct internal IPs which were mapped to approximately 6,593 unique MAC addresses. Using BLINC, we were able to classify 86% of our flows into application types. Some cases of misclassification were due to outlying user behavior. Nearly 5% of the users were responsible for 98% of misclassified Web traffic and thus excluded all misclassified Web flows. We were able to correlate with SNMP data 74.6% of the flows corresponding to 77.6% of the total traffic measured in our trace. Note that not all APs in the campus were monitored, hence the unmapped flows.

Table 1 presents general characteristics of our 7.5-day packet traces, such as the number of source and destination IPs observed, the total number of packets and bytes as well as the average utilization observed for both the wired and the wireless sources found in our traces. We will occasionally use the later data to compare application and client behavior across the wired and wireless network.

## 4 Aggregate traffic characterization

First using BLINC [16], we classified flows into applications and then identified the dominant application types, in terms of the total number of bytes, flows, and popularity among clients. Finally, we compared our findings with previous studies both in wired and wireless environments.

Although Web traffic has been traditionally prevalent, recent studies have highlighted an increasing interest in P2P applications. Table 2 shows that although less than half of the total flows are Web flows, it still accounts for nearly 57% of the traffic. P2P applications have a similar share, in terms of flows, that corresponds to roughly 24% of the total traffic. While some applications have a relatively small share of the overall traffic, they contribute with a large number of flows having thus, a significant impact on APs. Similarly, although the network management and scanning activity are both responsible for 0.42% of the bytes transferred, they account collectively for approximately 17% of the total flows! Thus, their effect on the overall network should not be ignored.

Although byte and flow statistics reveal the dominant applications, in terms of network traffic, they only indirectly hint on the popularity of each application. We define as *popularity* of an application the number of clients that had at least one flow of that specific application. Our findings are summarized at the last column of Table 2. Note that, these percentages do not sum up to 100%, since each client engages in more than one application. Interestingly, all clients have at least one Web flow, while a significant percentage of them (43%) appears to be P2P users. Further, 4% of them use online games in our wireless network, while most of

UNC set	Date	Start	Duration	Src IP	Dst IP	Number of packets	Bytes	Average Util.	Average number of flows
Wireless	2005-04-13	12:00	178.2 h	2650 K	2454 K	1046 M	583 GB	7.3 Mbps	25.8 K
Wired	2005-04-13	12:00	178.2 h	1998 K	1933 K	1022 M	521 GB	6.5 Mbps	19.5 K

**Table 1. General workload characteristics of our packet trace.**

Application type	Flows (%)	Bytes (%)	Packets (%)	Clients (%)
Network management	9.95	0.42	1.54	17
Chat	2.05	0.48	1.47	73
Web	35.6	57.59	46.80	99
P2P	30.04	24.85	34.46	43
On-line games	1.11	0.01	0.075	4
FTP	0.91	1.57	1.72	7
Mail	0.07	0.33	0.21	1.5
Address scan	6.40	0.12	0.58	73
Port Scan	0.39	0.32	0.28	1.4
Streaming	0.10	0.177	0.196	0.5
Unknown	13.2	14.09	12.64	84

**Table 2. Application cross-section with respect to flows, bytes, packets and clients.**

them had at least one unclassified flow. While mail traffic is only observed for a small number of our clients, this is the effect of our monitoring point at the edge of the network which does not capture any traffic for which both the source and the destination reside inside the UNC WLAN. Finally, note that, the majority of the clients are scanned at least once by an address scan or port scan flow!

We compared our results with three other application-based characterization studies in wireless and wired networks, namely, two wired campus networks (residential Campus, and UNC), and two wireless campus networks (UNC, and Dartmouth). Note that, although a direct comparison is not straightforward due to the differences in the monitored networks, time of collection, and lack of an identical application classification algorithm across studies, we can still observe general trends.

To this end, we first compared the traffic share of the most dominant and popular applications, of the UNC wireless network with the share of the same applications captured at the wired component of the UNC network (i.e., traffic originating from wired clients) within the same time interval. Similarly, we contrasted our findings to the “residential campus” trace studied in the original BLINC work [16], and the ones from the Dartmouth wireless network [11]. Table 3 summarizes the percentages of Web and P2P traffic for each one of these networks. Note that, the residential campus trace, and the two UNC traces were all classified by BLINC, hence the findings can be directly comparable; port numbers were used in the case of the Dartmouth trace. The results are remarkably similar for the two “wired” traces (residential and UNC campus networks), especially for P2P. On the contrary, the share of P2P traffic is significantly lower in both wireless traces amounting to approximately one fifth for Dartmouth and one fourth for the UNC of the total traffic. The most prominent difference

across the wireless traces is the share of Web traffic which is significantly higher in our traces. While such a difference may simply reflect different usage patterns across the two wireless networks, we speculate that the port-based classification of traffic may have missed all Web traffic that was not accessed through one of the well-known ports for Web (e.g, port 80).


## 5 AP traffic characterization

As observed also in previous studies, the overall distribution of traffic across APs is not uniform (Figure 1, left) and only few APs are responsible for the largest amount of traffic. For example, 3.5% of the APs are responsible for nearly 30.5% of the traffic while three APs are responsible for 10% of the total network traffic. The reason for such a skewed distribution of traffic among APs is the varying popularity among APs. To validate this assumption, we plot the CCDF of the number of clients across APs in Figure 1 (right). A few APs are significantly more popular than others, having been visited by more than 400 clients, and correspond to the ones with the highest traffic aggregation.

We found that the distribution of traffic across APs varies with the application. Figure 2 (left) presents the CCDF of the traffic share of three distinct application types relative to the total traffic across APs. While high percentages of Web traffic appear in most of the APs, a small portion of them, roughly 10%, access a significant amount of P2P data. Other application types rarely dominate any of the APs with the exception of three APs being dominated by mail traffic and one by FTP. This unevenness is especially pronounced in the case of P2P. Figure 2 (right) shows the combined percentage of Web and P2P traffic across APs sorted by their Web traffic percentage. There are two modes namely, either most APs are dominated by P2P traffic, or their share

	Residential campus	UNC Wired	UNC Wireless	Dartmouth Wireless
Web	37.5%	48.68%	57.59%	28.6%
P2P	31.9%	34.85%	24.85%	19.3%

**Table 3. Percentage of Web and P2P traffic (bytes) across four different networks.**



**Figure 1. Left: CDF of percentage of total network traffic across APs. Right: CCDF of the number of distinct clients across APs.**

of P2P traffic is minimal. Although Web is the most popular and dominant application type at the infrastructure-wide level, this is not the case at the AP level.

We say that an AP has a certain type of application as *home application*, if more than  $p\%$  of its total traffic in bytes is of that application type. Table 4 presents the breakdown of APs that have such a home application when more than half of the traffic of the AP belongs to an application type ( $p=50$ ). Approximately 4% of the APs are dominated by unknown traffic, while for roughly another 4%, a home application could not be defined, since no application type is dominating this AP’s traffic. Overall, there is an application preference towards specific APs in the wireless network. This observation can direct traffic engineering decisions about admission control or P2P traffic filtering at certain locations of the network.

We grouped APs based on their building category. Analysis at building level circumvents several problems that occur at AP-level, such as non amenability to statistical processing, higher sensitivity of monitored traffic variables to the short-term propagation conditions, and lack of scalability [17]. Thus, we grouped APs based on their category that reflects their functionality. We would expect for example to notice insignificant amounts of online games and P2P traffic at APs residing in academic buildings or near lecture halls.


To this end, a similar analysis is performed using the notion of the home application for an AP as defined in previ-

ous section. Table 5 presents the number of APs for nine building categories and shows the percentage of APs for which a home application existed. There is a weak correlation of the building category with the number of APs that have a home application (e.g., mail exists only in academic buildings as a home application, while FTP is present only in the business category). This reinforces our intuition that distinct APs may require different configuration settings depending on the application or type of building functionality.

The uneven traffic distribution in the application cross-section that was observed across buildings exists also across APs of the same building. Table 6 presents the percentage of Web and P2P traffic for all APs located in building 22. This building was chosen randomly among the buildings with the largest number of APs. While in most cases Web traffic dominates the overall traffic share, there are distinct APs (highlighted in the table) in which P2P contributes with the largest amount of traffic. Note that, these are not transient traffic phenomena, since our tracing period corresponds to several days.

### 5.1 Download and upload traffic asymmetry

Previous studies [14] have observed that certain APs are dominated by uploaders. In this section, we first explore if this phenomenon persists over time, and then, examine which application is responsible for the asymmetry between the downloaded and uploaded traffic. For that, we define



**Figure 2.** Left: CCDF of the percentage of traffic across the various APs. Half of the APs are dominated by Web traffic (>70%), while approximately 10% of APs by P2P traffic. Right: the share of P2P and Web across APs.

$p$	Web (%)	P2P (%)	FTP (%)	Mail (%)	Unknown (%)
50	85.9	6.17	0.28	0	4.2
75	55.8	0	0	0	0.84
90	25.52	0.28	0	0	0

**Table 4.** Percentage of APs with a home application.

the asymmetry index of an AP to be the fraction of the total downloaded to total uploaded traffic (bytes).

Approximately 5% of the APs show asymmetry indexes less than one, that is more bytes are uploaded than downloaded by clients associated to these APs. Figure 3 (left) is a scatter plot presenting the asymmetry index for each AP versus the total uploaded bytes at the same AP. It is interesting to note that, especially APs with significant amounts of uploaded traffic, have asymmetry indexes less or very close to one, which implies symmetric usage of these APs. The asymmetry trend observed in this study using packet header traces from 2005 is very similar to the one with traces from 2004 [14]. In 2005, there were slightly fewer APs with asymmetry indexes less than one, suggesting that overall downloading activity has increased.

As expected, uploading behavior is more intense in P2P than Web, highlighting the need for application profiling of the various APs. Heavy uploaders that are mostly associated to a specific AP could significantly affect the usage of that AP. For roughly 40% of the APs, the asymmetry index for P2P is less than one, while this is true for only 1.4% of the APs for Web. Network operators should dimension APs according to their application usage characteristics to exploit such phenomena. There has been also significant uploading behavior for chat and streaming applications. However, P2P is the main reason for almost all small asymmetry indexes

observed in Figure 3 (left).

## 6 Client traffic characterization

The AP- and building- level characterization is now extended to client-level to gain a better understanding of the user behavior, which is critical in the design of admission control and AP-selection mechanisms. We define the *home application* of a client as the application that is responsible for more than  $x\%$  of that client's traffic. Wireless clients have strong application preferences, both in terms of number of flows, and bytes. For example, for  $x$  equal to 90%, nearly half of the wireless clients have a home application, that is, approximately half of the clients transfer nearly all of their traffic through a specific application type. As shown in Table 7, the most prominent home application is the Web, while P2P appears to be the home application for only a minority of clients. The variation of traffic mix among clients without a home application is also interesting. Even in this case, most clients are still dominated by Web. The second largest share of their data is accessed either through P2P or an undefined application.

While wireless user application preferences appear to have similar trends as the wired ones [8, 15] (i.e., Web and P2P popularity), it is unclear whether client behavior is implicitly affected by the application performance over

Building Type	APs	Web (%)	P2P(%)	FTP(%)	Mail(%)	Unknown(%)
ACADEMIC	165	79	6	-	1.8	3
ADMINISTRATIVE	36	66	8.3	-	-	8
CLINICAL	16	62	6	-	-	6
ATHLETIC	15	53	20	-	-	6
RESIDENTIAL	42	83	5	-	-	2
BUSINESS	15	73	-	5	-	-
LIBRARY	12	86	-	-	-	-
CONFERENCE	9	55	-	-	-	11
THEATER	4	90	5	-	-	-

**Table 5. APs per building category and the percentage of APs with a home application.**

AP ID	44	45	46	47	48	49	50	51	52
Web (%)	89	81	87	85	78	<b>61</b>	95	<b>45</b>	74
P2P (%)	0.01	0.09	0.1	0.1	0.2	<b>35</b>	0	<b>51</b>	<b>11</b>

**Table 6. Web vs. P2P traffic share for building ID 22 across APs.**

the wireless channel. To shed some light on client behavior over the wireless network, we compare the characteristics of the dominant applications over the wired and wireless networks. As observed in our earlier study [13], the distribution of flow sizes for wireless users is very similar for both the wired and wireless component when looking at the aggregate traffic. Similarly, uploading and downloading flows also follow very similar distributions in both the wired and wireless components. However, flow sizes appear quite different over wireless and wired network per specific application type. Flow sizes are smaller for bulk file-transfer applications, such as FTP and P2P, over the wireless network. Figure 4 (left) presents the CCDF of flow sizes in bytes for the FTP, Web and P2P applications as seen by wired and wireless clients. In the case of Web, the two curves look very similar, with the wired Web flows being slightly heavier than the wireless ones. On the other hand, P2P and FTP wireless flows are lighter than the wired ones, especially the ones accessing large amounts of traffic. We speculate two possible explanations for this, one application-dependent and the other user-driven. First, especially in the case of P2P applications, loss or disconnected TCP flows severely affect performance; broken TCP connections will result in disconnecting from existing peers, which will further trigger peer discovery mechanisms and increase queue waiting times, hence ultimately decreasing overall flow sizes. The large number of retransmissions at the 802.11 MAC layer increases both the packet delay and number of retransmitted or failed packets at the transport layer [13]. This is consistent with Figure 4 (right) which compares the number of flows per client between Web and P2P in the wireless network. Note that, while overall the Web flows are heavier, in terms of bytes and number of transferred bytes per client<sup>1</sup>, the number of flows per client is larger in P2P than in Web. Essentially, each P2P client appears to have a large overhead in the wireless network with numerous small flows,

<sup>1</sup>The figure is not presented due to space limitations.

corresponding to control traffic. Second, users may avoid transferring large files over the wireless network because of the limited throughput, indicated also in Table 3 that compares the P2P traffic in wireless and wired. This observation holds also for FTP and P2P transfers.

## 7 Conclusions and future work

This work provides a detailed three-level characterization of an operational campus-wide wireless network across APs, clusters of APs and clients. Our results can be employed to support better admission control and AP selection mechanisms, indicate usage trends, and guide per-application traffic modeling efforts.

We found that while Web traffic appears to dominate both the client and AP traffic share, P2P applications bear a significant impact on the wireless network. First of all, most P2P protocols usually transfer data during all of their active time (i.e, seeding in bit torrent). Secondly, the large number of flows initiated by P2P applications, increase the contention at APs, which further increases the overall TCP RTT. Clients have strong application preferences and mostly access the wireless network for one specific application.

Extending our earlier work [13], we will focus on applications with real-time constraints and explore the impact of mobility and wireless network on their performance. We will use heuristics and statistical clustering techniques to profile clients based on their application characteristics and roaming patterns, and contrast the results using data from different wireless environments. To encourage further experimentation we have made our datasets and tools available to the research community [1].

## 8 Acknowledgment

This work was partially supported by the General Secretariat for Research and Technology (05NON-EU-238) and

$x(\%)$	NM	Chat	Web	P2P	Games	FTP	Mail	Streaming	Address scan	Port scan	Unknown
50	0.40%	1%	88.8%	2.36%	0%	0.13%	0.21%	0.1%	0.03%	0.03%	5.64%
75	0.13%	0.77%	78.16%	1.5%	0%	0.07%	0.16%	0.06%	0.01%	0.0115%	2.5%
90	0.06%	0.56%	59.26%	0.8%	0%	0.06%	0.06%	0.015%	0%	0%	1.1%

Table 7. Percentage of clients per home application for various thresholds ( $x\%$ ).





Figure 3. Left: CCDF of the asymmetry index for Web, P2P and total traffic. Right: asymmetry index vs. uploaded traffic APs.

by European Commission with a Marie Curie IRG (MIRG-CT-2005-029186 ForthRepository) grant.

## References

- [1] UNC/FORTH archive of wireless traces, models, and tools. <http://www.ics.forth.gr/mobile/>.
- [2] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris. Link-level measurements from an 802.11b Mesh network. In *Symposium on Communications Architectures and Protocols (SIGCOMM)*, Portland, OR, August 2004.
- [3] A. Balachandran, G. Voelker, P. Bahl, and V. Rangan. Characterizing user behavior and network performance in a public wireless lan. In *ACM Sigmetrics Conference on Measurement and Modeling of Computer Systems*, California, USA, June 2002.
- [4] M. Balazinska and P. Castro. Characterizing mobility and network usage in a corporate wireless local-area network. In *First International Conference on Mobile Systems, Applications, and Services (MobiSys)*, San Francisco, USA, May 2003.
- [5] J. Bicket, D. Aguayo, S. Biswas, and R. Morris. Architecture and evaluation of an unplanned 802.11b mesh network. In *ACM International Conference on Mobile Computing and Networking (MobiCom)*, Cologne, Germany, August 2005.
- [6] C. Chambers, W. Feng, S. Sahu, and D. Saha. Measurement-based characterization of a collection of on-line games. In *Symposium on Communications Architectures and Protocols (SIGCOMM)*, Berkley, USA, October 2005.
- [7] F. Chinchilla, M. Lindsey, and M. Papadopouli. Analysis of wireless information locality and association patterns in a campus. In *IEEE Conference on Computer Communications (Infocom)*, Hong Kong, March 2004.
- [8] S. A. R. Group. <http://ipmon.sprintlabs.com/packstat/packetoverview.php>.
- [9] S. Guha, N. Daswani, and R. Jain. An experimental study of the skype peer-to-peer voip system. In *IPTPS, 2006*, Santa Barbara, CA, USA, February 2006.
- [10] L. Guo, S. Chen, Z. Xiao, E. Tan, X. Ding, and X. Zhang. Measurements, analysis, and modeling of bittorrent-like systems. In *Symposium on Communications Architectures and Protocols (SIGCOMM)*, Berkley, USA, October 2005.
- [11] T. Henderson, D. Kotz, and I. Abyzov. The changing usage of a mature campuswide wireless network. In *ACM International Conference on Mobile Computing and Networking (MobiCom)*, Philadelphia, PA, USA, September 2004.
- [12] F. Hernandez-Campos, M. Karaliopoulos, M. Papadopouli, and H. Shen. Spatio-temporal modeling of traffic workload in a campus WLAN. In *Second Annual International Wireless Internet Conference*, Boston, USA, August 2006.
- [13] F. Hernandez-Campos and M. Papadopouli. Assessing the real impact of 802.11 wlans: A large-scale comparison of wired and wireless traffic. In *14th IEEE Workshop on Local and Metropolitan Area Networks*, Chania, Greece, September 2005.



**Figure 4. Left: CCDF of flow sizes for FTP, P2P and Web. FTP and P2P flows appear heavier over the wired network. Right: flows per client for P2P and Web in the wireless network.**

- [14] F. Hernandez-Campos and M. Papadopouli. A comparative measurement study of the workload of wireless access points in campus networks. In *16th Annual IEEE International Symposium on Personal Indoor and Mobile Radio Communications*, Berlin, Germany, September 2005.
- [15] T. Karagiannis, A. Broido, M. Faloutsos, and K. claffy. Transport layer identification of p2p traffic. In *Symposium on Communications Architectures and Protocols (SIGCOMM)*, Portland/Oregon, USA, August-September 2004.
- [16] T. Karagiannis, K. Papagiannaki, and M. Faloutsos. Blinc: Multi-level traffic classification in the dark. In *Symposium on Communications Architectures and Protocols (SIGCOMM)*, Philadelphia, PA, USA, August 2005.
- [17] M. Karaliopoulos, M. Papadopouli, E. Raftopoulos, and H. Shen. On scalable measurement-driven modeling of traffic demand in large wlans. Technical Report 383, ICS-FORTH, Heraklion, Crete, Greece, August 2006.
- [18] M. McNett and G. M. Voelker. Access and mobility of wireless pda users. In *Mobile Computing Communications Review*, April 2005.
- [19] A. Moore and K. Papagiannaki. Toward the accurate identification of network applications. In *Passive and Active Measurement Workshop*, Boston, MA, USA, March 2005.
- [20] M. Papadopouli, M. Moudatsos, and M. Karaliopoulos. Modeling the roaming in large-scale wireless networks using real measurements. In *International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM06)*, Niagara-Falls/Bufalo, New York, USA, June 2006.
- [21] M. Papadopouli, E. Raftopoulos, and H. Shen. Evaluation of short-term traffic forecasting algorithms in wireless networks. In *2nd Conference on Next Generation Internet Design and Engineering*, Valencia, Spain, April 2006.
- [22] M. Papadopouli, H. Shen, and M. Spanakis. Characterizing the duration and association patterns of wireless access in a campus. In *11th European Wireless Conference*, Nicosia, Cyprus, April 2005.
- [23] M. Papadopouli, H. Shen, and M. Spanakis. Modeling client arrivals at access points in wireless campus-wide networks. In *14th IEEE Workshop on Local and Metropolitan Area Networks*, Chania, Greece, September 2005.
- [24] D. Tang and M. Baker. Analysis of a local-area wireless network. In *ACM International Conference on Mobile Computing and Networking (MobiCom)*, pages 1–10, Boston, Massachusetts, USA, August 2000.
- [25] K. Tutschku. A measurement-based traffic profile of the edonkey filesharing service. In *Passive and Active Measurement Workshop*, Antibes Juan-les-Pins, France, April 2004.