

# LAC: Load-Aware Channel Selection in 802.11 WLANs

(Invited Paper)

George Athanasiou\*, Ioannis Broustis\*, Thanasis Korakis<sup>†</sup>, Leandros Tassioulas\*

\* Department of Computer and Communications Engineering, University of Thessaly

Email: {*gathanas, broustis, leandros*}@uth.gr

<sup>†</sup> Department of Electrical and Computer Engineering, Polytechnic University

Email: *korakis@poly.edu*

**Abstract**—Dense deployments of hybrid WLANs result in high levels of interference and low end-user throughput. Many frequency allocation mechanisms for WLANs have been proposed by a large body of previous studies. However, none of these mechanisms considers the load that is carried by APs in terms of channel conditions, number of affiliated users as well as traffic-load, in conjunction. In this paper, we propose LAC, a load-aware channel allocation scheme for WLANs, which considers all the above performance determinant factors. LAC incorporates an *airtime cost* metric into its channel scanning process, in order to capture the effects of these factors and select the channel with the maximum long-term throughput. We evaluate LAC through extensive OPNET simulations, for many different traffic scenarios. Our simulations demonstrate that LAC outperforms other frequency allocation policies for WLANs in terms of total network throughput by up to 135%.

**Index Terms**—IEEE 802.11 Wireless Local Area Networks (WLANs), Channel Selection, Interference.

## I. INTRODUCTION

The growing demand for high-throughput wireless Internet connectivity has enabled the deployment of thousands of WLANs in urban areas, during the last decade. This has resulted to increased amounts of interference and contention among co-channel access points (APs) [1]. As a consequence, the end-users (affiliated clients with those APs) end up enjoying very low throughputs in the long term. Towards addressing this problem, various studies have proposed mechanisms for allocating the available number of channels to APs [2], [3], [4], [5], [6]. However, each of these studies considers only a subset of the parameters that affect the performance of the network, namely the number of associated clients, their traffic load, or the channel conditions in terms of projected interference (SINR) and thereby achievable transmission rates. In other words, there are no studies that embrace all these parameters under a common frequency selection framework; therefore the previously proposed solutions are efficient only when certain conditions are met (as we discuss in detail, in section II).

As our contribution in this paper, we propose LAC, a Load-Aware Channel assignment scheme for 802.11 WLANs, which finds the channel that approximates the maximum throughput of the AP, i.e., the sum of throughputs achieved by all its affiliated clients. LAC employs the *airtime cost* metric [7]<sup>1</sup>. Through this metric LAC discovers the most appropriate channel for every AP in a distributed fashion, by measuring: (a) both the downlink and uplink channel conditions in terms of supported transmission rates and packet error probability, and (b) the number of affiliated clients with every AP. As a result, the AP has a unified knowledge with regards to the quality of all its downlink and uplink connections. Note here that through these measurements the *airtime cost* directly reflects the environmental conditions in and around each device (AP or client), in terms of interference and contention experienced due to concurrent transmissions [2]. In addition, accounting for the number of clients per AP helps construct a more accurate view of which APs are idle or underutilized, so as for LAC to exclude these APs from the set of interfering devices.

In a nutshell, with LAC every AP and its clients perform a sequential scanning on all the available channels and collect measurements with regards to the above metrics. Furthermore, they exchange these measurements to compute the average *airtime cost* for the AP cell that they belong to, and for every channel. After a set of iterations they select the channel wherein the average *airtime cost* is minimal; this is the channel that will provide the estimated maximum long-term cell throughput, since this throughput can be considered as an inverse function of the average *airtime cost* (as we show in [8]). We implement LAC in OPNET and we evaluate its performance against 3 other channel selection mechanisms. We perform extensive simulation experiments, and we demonstrate the outperformance of LAC in terms of throughput, delay and number of dropped packets, for different traffic patterns

<sup>1</sup>The *airtime cost metric* was proposed in [7] for optimal user association in 802.11 networks.

and topological settings.

The rest of the paper is structured as follows. Section II provides the relevant previous work on channel allocation for WLANs. Section III describes the airtime metric that we employ, based on which LAC discovers the channel with the maximum long-term throughput. In section IV, we present the design of LAC. We evaluate our protocol through simulations, in section V. Finally, section VI concludes this paper.

## II. RELATED STUDIES

In this section, we discuss the most relevant previous work on efficiently assigning channels to APs in 802.11 wireless LANs. A large set of studies on channel selection in WLANs exists; to the best of our knowledge, however, none of these mechanisms captures the total actual throughput at the AP cell for every particular channel.

The LCCS (Least Congested Channel Search) method [9] was the first effort towards allocating a set of available channels to wireless devices. With LCCS, devices (e.g. APs) periodically scan the set of available channels and select the one with the lowest *levels of contention* (as the name suggests). However, there are many topological scenarios where LCCS is unable to capture the total interference in the channel, as explained in [4]. Similarly, Leith and Clifford [10] propose a self-managed distributed channel selection scheme, wherein each AP passively measures the received power from the packets transmitted by neighbor APs.

Along similar lines, Mhatre et al. in [2] propose a distributed frequency selection algorithm, which is proved to minimize the global interference in the network. Simply put, minimizing the total interference can result in improved user throughput. Towards addressing this objective, each AP measures the total received power from all neighbor APs for every channel and selects the channel with the minimum total power. This is performed at each AP by measuring the RSSI of the received beacon frames from all the neighbor APs at every channel. The authors show that their proposed algorithm manages to converge to the global optimum of the optimization criterion, i.e., the minimization of interference across the entire network. However, this algorithm does not consider the number of clients in the network; it assumes purely downlink saturated traffic and that all APs have affiliated clients. Moreover, the channel with the minimum total interference does not guarantee maximal throughput at the cell. We compare LAC against the protocol in [2], in section V.

Moreover, the work in [6], by Mishra et al., belongs to a set of studies that propose a distributed channel hopping mechanism. The mechanism in [6], MaxChop, provides higher levels of fairness among users. Channel hopping, however, requires tight synchronization between AP and clients, while it is difficult to implement efficiently with off-the-shelf hardware. Note that the channel switching and

the subsequent restoration of traffic at the new channel may take from 700 to 1000 msec [11]; this is prohibitive in terms of incurred overhead.

Lee et al. [12] take into account the *expected* traffic demand points in the network. Their channel allocation strategy seeks to assign frequencies in such a way that the signal strength at these demand points is maximized. As a further step, Rozner et al. in [3] also consider the current traffic demands at the WLAN. In particular, they show that, taking into consideration the current traffic demands at APs and clients, the quality of the channel assignment can be greatly improved.

Furthermore, centralized channel allocation algorithms have been proposed in [4], [5]. Mishra et al. [4] propose a frequency allocation scheme, wherein clients play a large role in the decision for the best channel. Their proposed approach opts to perform joint load balancing and frequency allocation. However, the approach is based on conflict graph coloring and cannot be directly implemented in a distributed setting.

## III. A METRIC FOR CHANNEL ALLOCATION

In this section, we present the metric that we adopt in LAC, in order to effectively select the most appropriate channel in every cell.

The metric that is used in our load-aware channel allocation scheme is called *airtime metric* and it is an approximation of the per packet latency (as shown in [13], [14] and in our extended analysis [8]). The airtime metric was first discussed in the 802.11s [15] standard, for the purposes of load-aware routing (RM-AODV routing protocol). This metric reflects the load on a wireless router (AP) in terms of the average delay a transmission of a unit size packet experiences. RM-AODV which is the default routing protocol in 802.11s-based wireless mesh networks, employs the airtime metric in order to provide end-to-end paths with the minimum total *airtime cost*. We adopt this metric in LAC for the purposes of our proposed channel selection functionality.

Formally, the airtime metric of station  $i \in U_a$ , where  $U_a$  is the set of stations associated with AP  $a$  that communicate using channel  $c$ , is given as:

$$C_{a,c}^i = \left[ O_{ca} + O_p + \frac{B_t}{R_i^{a,c}} \right] \frac{1}{1 - e_{pt}^c}. \quad (1)$$

In (1),  $O_{ca}$  is the channel access overhead,  $O_p$  is the protocol overhead and  $B_t$  is the number of bits in the test frame<sup>2</sup>. Some representative values for these constants, for 802.11g, are:  $O_{ca} + O_p = 1.25\text{ms}$  and  $B_t = 8224\text{bits}$ . Furthermore,  $R_i^{a,c}$  and  $e_{pt}^c$  are the current transmission rate and frame-error rate, respectively, in Mbps, for the test frame size  $B_t$  in channel  $c$ . In other words, the estimation

<sup>2</sup>The transmission of test frames is necessary, in order to derive values for the computation of the airtime cost.

of  $e_{pt}^c$  corresponds to transmissions of standard-size frames  $B_t$  at the current transmit bit rate  $R_i^{a,c}$ .

The average *airtime cost* (in one direction: uplink or downlink) of AP  $a$  with  $N_a$  users, that operates on channel  $c$  is:

$$\begin{aligned} \overline{C_{a,c}} &= \frac{1}{N_a} \sum_{i=1}^{N_a} \left[ O_{ca} + O_p + \frac{B_t}{R_i^{a,c}} \right] \times \frac{1}{(1-e_{pt}^c)} = \\ & \left[ \frac{1}{N_a} \sum_{i=1}^{N_a} (O_{ca} + O_p) + \frac{1}{N_a} \sum_{i=1}^{N_a} \frac{B_t}{R_i^{a,c}} \right] \times \frac{1}{(1-e_{pt}^c)} \end{aligned} \quad (2)$$

Our analysis in [8] shows that the average *airtime cost* in the uplink and the downlink is an approximation of the average per-packet delay in both directions. Therefore, the average *airtime cost* is a representative metric that reflects the uplink and downlink channel performance and also approximates the maximum throughput in the cell.

The main advantage of our work is that it captures the performance of an AP in terms of estimated throughput at a particular channel, by measuring the average *airtime cost* for both uplink and downlink,  $C_a^c = \overline{C_{a,c}^{up}} + \overline{C_{a,c}^{down}}$  (*airtime cost* for AP  $a$ , in channel  $c$ ), and applies a channel allocation methodology where the channel with the minimum  $C_a^c$  is chosen. This channel selection policy determines the frequency with the minimum average per-packet delay in both uplink and downlink, thereby approximating the maximum throughput in the cell [8].

Various studies have shown that the number of erroneously received packets increases and the transmission rate decreases, in the presence of interfering cells in the network [16], [2]. Our proposed *airtime metric* takes into account the packet error rate as well as the transmission rate; hence, it reflects the performance at a particular communication channel. In the next section, we describe our channel allocation protocol; we explain how the *airtime metric* is used in order to optimize the allocation of the available channels and improve the network throughput.

#### IV. LAC: OUR CHANNEL ALLOCATION SCHEME

In this section, we describe LAC, our load aware channel selection mechanism. Our analysis from the previous section clearly shows that the *airtime metric* reflects the performance of the WLAN in terms of AP throughput. Hence, determining the channel with the lowest *airtime cost* will provide the maximum long-term throughput to the clients within a cell. This is the target of LAC. LAC is AP-centric, i.e., the channel choice is made by the APs. Note that clients also play a very important role in the channel decision, by informing their affiliated AP with regards to the uplink conditions. This way the AP has knowledge with regards to both the downlink and uplink. The AP and its clients perform scanning procedure for the set of available channels. For every scanned channel, the AP and its clients measure the downlink and uplink channel properties and exchange them through their control and data transmissions. This information is subsequently

used by the AP to select the channel with the minimum cumulative airtime cost. More specifically, the scanning procedure is comprised of the following 4 steps.

**[Step1]. Computing the downlink airtime cost.** At the nominal start of LAC, AP  $A_i$  of cell  $i$  initiates the downlink airtime cost calculation.  $A_i$  sets a special bit in its beacon template, thereby informing its clients that the airtime calculation process has been initiated.  $A_i$  calculates the average downlink airtime cost for the links with its clients by following the link performance-measurement procedure that is described in detail in [7].

**[Step2]. Computing the uplink airtime cost.** The clients of  $A_i$  read the airtime cost bit from  $A_i$ 's beacon frame transmissions, and they further calculate their individual uplink costs [7]. They may include these costs into their measurement report messages towards  $M_A$ , as per [17]. Alternatively, they can also piggy-back this information through their data frame transmissions towards  $A_i$ . In this work we adopt the latter approach (we discuss this choice later). By the end of this step,  $A_i$  has received information regarding the uplink channel qualities from all its clients.

**[Step3]. Deciding if the current channel is appropriate.**  $A_i$  receives the client reports and computes the average airtime cost for both uplink and downlink [8]. If this is higher than a pre-defined threshold,  $T$ , then  $A_i$  remains in the same channel; otherwise it initiates a channel discovery process.

**[Step4]. Computing the cumulative airtime cost at the next available channel.**  $A_i$  and its clients switch to the next channel and repeat steps 1 to 3. If all available channels have been visited,  $A_i$  finally selects the channel with the minimum average airtime cost (for both uplink and downlink).

In what follows, we discuss some properties of our LAC mechanism.

**Band dwell duration.** Calculating  $C_{A_i}^c$  in cell  $i$  for channel  $c$ , requires that  $A_i$  and its clients conduct link measurements in order to derive values for the transmission rate and frame error rate, as explained in the previous section. Towards collecting accurate values for these parameters, a sufficient amount of time of residing at a particular channel is to be consumed by all devices of cell  $i$ . Note, however, that: **(a)** The measurement collection process for every channel is performed through data exchange between AP and clients, and therefore traffic keeps flowing in the network even during channel scanning. In other words, with LAC traffic does not stop flowing between AP and clients, as with other channel allocation protocols (e.g. [10], [2], etc.) **(b)** In typical today's WLANs, where clients are statically located most of the time, the levels of interference are not fluctuating to a large extent. Thus, LAC does not need to be executed frequently, as we discuss below. Hence, one can expect that due to (a) and (b) above, the projected overheads because of the channel

dwelling duration are minimal.

**Convergence and frequency of invocation.** The set of steps 1-4 belong to an iterative process, where the set of available channels is re-scanned by the APs, until convergence has been reached. This is because, whenever AP  $A_i$  decides upon a certain channel  $c_i$ , the cumulative airtime cost for the neighbor cells of  $A_i$  will likely change for channel  $c_i$ ; this will enable the APs of those cells to also perform channel scanning [18], for a potentially better channel than  $c_i$ . Our simulations in section V have shown that in a static deployment with 20 APs and 40 clients, convergence is reached rather quickly (approximately 2 iterations, on average). We expect that in a more dynamic topology, where clients join or de-associate from the network, the invocation of LAC would be more frequent, since the violation of threshold  $T$  would occur more often. This problem is not so prominent with the schemes proposed in [2], [10], since those mechanisms do not consider the potential transmissions of clients. Due to this, however, those approaches yield a poorer performance than LAC, as we demonstrate in the following section.

**Embedding the airtime cost value into control and data frames.** With LAC, clients piggy-back their individual uplink airtime costs into their data transmissions towards the AP. As we mentioned earlier, another way of sending this information to their AP is through probe response frames, or special management report frames [17]. Note however that, as reported in [18], such control frames are not always received intact by their destinations, since they are not acknowledged. Therefore, the converge of mechanisms that depend on these frames may be delayed, since the AP does not manage to collect accurate information within the channel dwell duration [18]. Thus, we consider that the airtime cost value is repeatedly piggy-backed into every data packet transmission, for reliability.

## V. EVALUATING OUR LOAD-AWARE CHANNEL ALLOCATION SCHEME

In this section, we evaluate our load-aware channel allocation scheme through extensive OPNET [19] simulations. We compare our scheme against the frequency selection approach, proposed in [2], and we present LAC's predominance in terms of the total network throughput, average packet dropping and average transmission delay.

### A. Simulation set-up details

We have implemented LAC in OPNET [19], taking into account the 802.11 protocol operations. We have modified the beacon and data frames of 802.11 to facilitate the information exchange process that our protocol design requires. The clients and the APs are uniformly (at random) distributed in the  $1000\text{m} \times 1000\text{m}$  simulation area. All nodes use a default transmit power of 20 dBm. We simulate the network behavior with three types of traffic: (a) fully-saturated, downlink UDP traffic, where




Fig. 1: Total network throughput with saturated downlink UDP traffic.

the APs send traffic to their associated clients, (b) fully-saturated, bi-directional UDP traffic (where the source-destination pairs are chosen randomly) and (c) VoIP traffic. We repeat our simulations with both 802.11a and 802.11g modes of operation. We compare the network performance of LAC against the single-channel assignment (SC), a random-channel allocation strategy<sup>3</sup> (RC), as well as the scheme in [2], which we call *GFS* (for “Gibbs-based Frequency Selection”). We evaluate the efficiency of LAC in selecting the most appropriate channel by measuring the total achieved network throughput, the average end-to-end delay and the average dropped data packets in the network.

### B. Simulation results

We present our simulation experiments and the interpretations thereof, in what follows.

**Applying downlink UDP traffic.** To begin with, we opt to compare the performance with LAC against the performance with the other approaches, with downlink traffic. We are mainly interested in comparing the performance against the case with GFS, which assumes downlink traffic. For this, we apply saturated downlink UDP traffic in a network with 20 APs and 40 STAs, all uniformly-randomly distributed in the area. In other words, we assume that the APs have always packets in their buffer, to send to their associated clients. We use the default UDP packet size (1500 bytes). Fig. 1 depicts the average total network throughput when there are 11 (802.11a) and 3 (802.11g) orthogonal channels available. The best performance is achieved by LAC, which uses the airtime metric to capture the cell performance at every channel. The number of the associated clients is not considered by GFS, since the latter considers that the imposed interference on an AP comes only from its neighbor APs, while the interference at the client is assumed to be approximately equal as that at the AP. Note however that in this scenario GFS does not

<sup>3</sup>This is the most frequent approach in practice today, as discussed in [18].

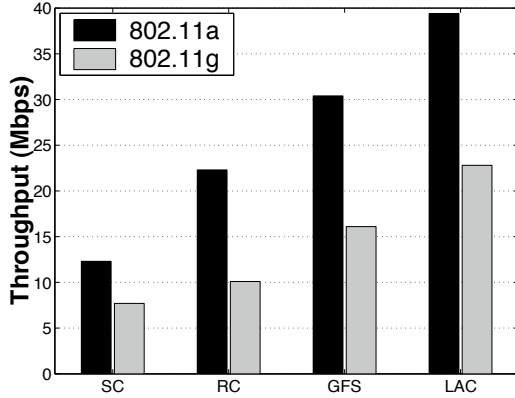


Fig. 2: Total network throughput with both-directions UDP traffic.

achieve much poorer throughput than LAC. In other words, the above assumption of GFS is weak and rather realistic in a purely downlink traffic scenario. Nevertheless, our load-aware scheme performs better than the other three channel selection approaches. In particular, LAC outperforms GFS by 8% and RC by 59% when 802.11a is used, and by 20% and 104% respectively when 802.11g is used.

**Applying bi-directional UDP traffic between random source-destination pairs.** We are also interested in the network performance with LAC, in more realistic (than downlink-only) traffic scenarios. In this case, each client sends saturated UDP traffic to another, randomly selected client in the network. Note that the client-destination may be associated to a different AP than the client-source<sup>4</sup>. Hence, both uplink and downlink UDP traffic takes place in every cell (uplink traffic in the client→AP (source) link and downlink traffic in the AP→client (destination) link). Fig. 2 depicts the average total network throughput, achieved with the different channel selection approaches. We observe that GFS underperforms LAC to a large extent (LAC improves network throughput achieved by GFS by 27% in 802.11a and by 47% in 802.11g). This is because GFS does not take into account the fact that also clients may be sending traffic towards their APs; hence, the uplink channel conditions and the load of each communication channel are not considered by GFS. Finally, we observe that LAC improves the total network throughput as compared to RC, by 70% in 802.11a and by 135% in 802.11g.

**Scalability of LAC:** Next, we seek to observe the throughput with LAC as the network density increases, in terms of number of clients. For this, we progressively increase the number of clients from 5 to 70 in the network, while we maintain the same number of APs (20). Fig.

<sup>4</sup>In our simulations, APs are connected through a wireline Ethernet network, and they use the Ethernet interface to exchange packets with other cells.

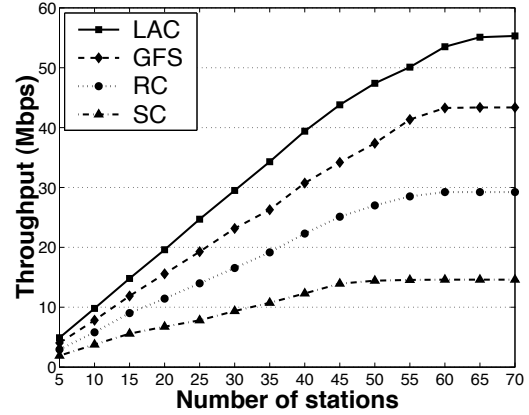


Fig. 3: Total network throughput Vs. Number of clients.

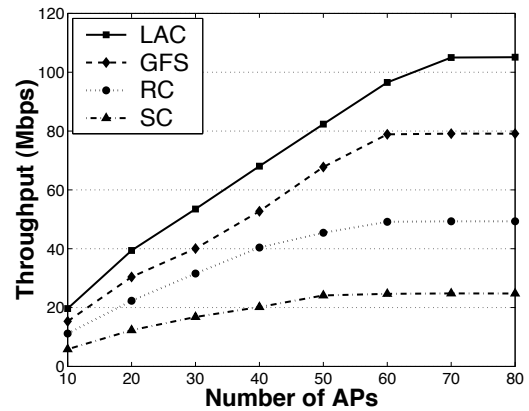


Fig. 4: Total network throughput Vs. Number of APs.

3 depicts the throughput results for the four channel selection strategies. We observe that the performance with LAC is similar to the other 3 policies, when the number of clients (and therefore the load) is low. However, with increased load LAC manages to provide much higher throughputs due to its load-aware channel allocation strategy.

Moreover, we measure the average total network throughput as we increase the number of APs, from 10 to 80 (meanwhile, we deploy twice the number of clients: 10 APs - 20 clients, 20 APs - 40 clients, etc.). Fig. 4 depicts the throughput gains with LAC. We observe that our scheme manages to scale much better than all other 3 approaches. In other words, the maximum achievable total network throughput is reached when the network includes 10 more APs for the case of LAC than in the case of GFS.

**Simulating LAC with VoIP traffic.** In order to observe the performance of our protocol with delay-sensitive applications, we utilize varying, parallel, end-to-end VoIP traffic sessions. The simulation set-up is the same as for the previous bi-directional experiments, i.e., VoIP traffic is exchanged among clients in our network – hence we have

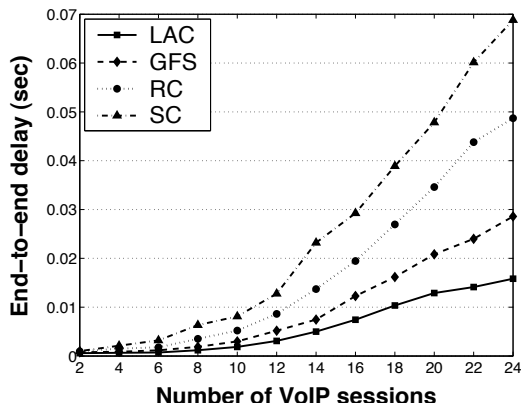


Fig. 5: Average end-to-end delay with VoIP traffic.




Fig. 6: Average dropped data with VoIP traffic.

again both uplink and downlink traffic at every cell. Fig. 5 and 6 present the network performance with VoIP.

In particular, Fig. 5 depicts the average end-to-end delay of VoIP packet transmissions. We observe that LAC achieves low end-to-end delays, due to its sophisticated channel allocation strategy. GFS achieves quite good performance in low load communication conditions. However, as the number of the supported VoIP sessions increases, the channel setting with GFS is unable to efficiently support them. Finally, Fig. 6 shows the average number of dropped data packets due to channel errors and contention. The performance of LAC is impressive, since packet dropping is kept in very low levels, as compared to the other strategies.

## VI. CONCLUSIONS AND FUTURE WORK

In this paper, we propose LAC, a Load-Aware Channel selection mechanism for 802.11-based WLANs. LAC adopts the airtime cost metric, which was originally proposed in the 802.11s standard, and which provides an estimation of the average packet transmission delay. LAC employs a channel scanning procedure, which converges to the channel with the minimum average airtime cost (for both uplink and downlink), thus managing to provide

approximately the maximum cell throughput. We compare LAC against 3 other channel selection approaches and we show that it outperforms all of them, for different traffic scenarios and network densities. So far we have designed two performance improvement mechanisms for WLANs, one for user association [7] and one for channel allocation, both being based on the airtime cost. As our future work, we opt to design and implement a common unified framework, which will combine these two functionalities.

## VII. ACKNOWLEDGMENTS

The authors acknowledge support of European Commission STREP WIP (FP6-IST-027402) and support of Greek Secretariat for Research and Technology through a PENED grant.

## REFERENCES

- [1] A. Akella, G. Judd, S. Seshan, and P. Steenkiste. Self-Management in Chaotic Wireless Deployments. In *ACM MOBICOM*, 2005.
- [2] B. Kauffmann, F. Baccelli, A. Chainteau, V. Mhatre, K. Papagiannaki, and C. Diot. Measurement-Based Self Organization of Interfering 802.11 Wireless Access Networks. In *IEEE INFOCOM*, 2007.
- [3] E. Rozner, Y. Mehta, A. Akella, and L. Qiu. Traffic-Aware Channel Assignment in Enterprise Wireless LANs. In *IEEE ICNP*, 2007.
- [4] A. Mishra, V. Brik, S. Banerjee, A. Srinivasan, and W. Arbaugh. A Client-Driven Approach for Channel Management in Wireless LANs. In *IEEE INFOCOM*, 2006.
- [5] B.-J. Leung and K. K. Kim. Frequency Assignment for IEEE 802.11 Wireless Networks. In *IEEE VTC*, Vol. 3, pp. 1422-1426, 2003.
- [6] A. Mishra, V. Shrivastava, D. Agarwal, and S. Banerjee. Distributed Channel Management in Uncoordinated Wireless Environments. In *ACM MOBICOM*, 2006.
- [7] G. Athanasiou, T. Korakis, O. Ercetin, and L. Tassiulas. Dynamic Cross-Layer Association in 802.11-based Mesh Networks. In *IEEE INFOCOM*, 2007.
- [8] G. Athanasiou, I. Broustis, T. Korakis, and L. Tassiulas. Load-Aware Channel Selection Scheme for 802.11 WLANs, Technical report, University of Thessaly, 2008, <http://inf-server.inf.uth.gr/~gathanas/Channel2008.pdf>.
- [9] J. Geier. Assigning 802.11b Access Point Channels. In *WiFi planet*, 2002.
- [10] D. J. Leith and P. Clifford. A Self-Managed Distributed Channel Selection Algorithm for WLANs. In *WIOPT*, April 2006.
- [11] R. Vedantham, S. Kakumanu, S. Lakshmanan, and R. Sivakumar. Component Based Channel Assignment in Single Radio, Multi-channel Ad Hoc Networks. In *ACM MOBICOM*, 2006.
- [12] Y. Lee, K. Kim, and Y. Choi. Optimization of AP Placement and Channel Assignment in Wireless LANs. In *IEEE LCN*, 2002.
- [13] O. Ercetin. User Association Games in 802.11 Wireless Local Area Networks. In *preprint*, 2008.
- [14] G. Athanasiou, T. Korakis, O. Ercetin, and L. Tassiulas. A Cross-Layer Framework for Association Control in Wireless Mesh Networks. In *IEEE Trans. on Mobile Computing*, (to appear), 2009.
- [15] IEEE 802.11s: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: Simple Efficient Extensible Mesh (SEE-Mesh) Proposal.
- [16] D. Niculescu. Interference Map for 802.11 Networks. In *ACM IMC*, 2007.
- [17] IEEE 802.11 WG. Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications: Specification for Radio Resource Measurement, IEEE 802.11k/D3.0. New York, USA: The Institute of Electrical and Electronics Engineers, Inc., October 2005.
- [18] I. Broustis, K. Papagiannaki, S. V. Krishnamurthy, M. Faloutsos, and V. Mhatre. MDG: Measurement-Driven Guidelines for 802.11 WLAN Design. In *ACM MOBICOM*, 2007.
- [19] OPNET-Radio/Wireless Models. <http://www.opnet.com>.