

Dynamic Frequency Management in 802.11-based Multi-Radio Wireless Networks

George Athanasiou^a, Leandros Tassioulas^b

^aAutomatic Control Lab, School of Electrical Engineering,
KTH Royal Institute of Technology, Sweden. E-mail: georgioa@kth.se

^bDepartment of Computers and Communications Engineering,
University of Thessaly, Greece. E-mail: leandros@uth.gr

Abstract

Efficient channel selection is essential in 802.11 mesh deployments, for minimizing contention and interference among co-channel devices and thereby supporting a plurality of QoS-sensitive applications. A few protocols have been proposed for frequency allocation in such networks, however they do not address the problem *end-to-end*. In this paper, we present a general formulation of the channel selection problem taking into account the performance of both mesh-access and mesh-backhaul. Moreover, we propose ARACHNE, a routing-aware channel selection protocol for wireless mesh networks. ARACHNE is distributed in nature, and motivated by our measurements on a wireless testbed. The main novelty of our protocol comes from adopting a metric that captures the end-to-end link loads across different routes in the network. ARACHNE prioritizes the assignment of low-interference channels to links that (a) need to serve high-load aggregate traffic and/or (b) already suffer significant levels of contention and interference. Our protocol takes into account the number of potential interfaces (radios) per device, and allocates these interfaces in a manner that efficiently utilizes the available channel capacity. We evaluate ARACHNE through extensive, trace-driven simulations and we show that approaches the optimal channel selection. We observe that our protocol improves the total network throughput, as compared to three other representative channel allocation approaches in literature.

Keywords: Frequency management, Wireless networks, IEEE 802.11, Cross-layer

1. Introduction

Wireless mesh networking has been touted as the new technology that can support ubiquitous end-to-end connectivity. In wireless mesh networks, information has to be routed via multiple wireless hops before it can reach the destination [1, 2]. A critical requirement for the efficient routing of packets is the identification and use of interference-limited wireless links. Therefore, intermediate mesh hops along a route need to operate in frequencies, where contention and interference are as low as possible, especially in highly-dense mesh deployments. We ask the question: *How can we allocate frequencies in a mesh network, in order to maximize the total network throughput, in a distributed manner?*

In order to efficiently allocate the set of available channels to nodes, the load at each individual link needs to be taken into account. Here, the load at the mesh access and backhaul levels could be represented in various ways [3, 4], potentially involving the number of

neighbor transmitters, the traffic demand, the amount of traffic flowing through each node, etc., as we discuss later. Previous studies on frequency selection, however, do not consider the *end-to-end load distribution* across entire routes; they consider the sub-problems of either the access level, between clients and access points (APs) [4, 5], or the backbone portion of the network [6, 7, 8]. We argue that frequency selection algorithms should prioritize the assignment of low-interference channels at highly-loaded mesh links, both at the access and the backhaul levels. As a simple example, consider the connectivity graph of Fig. 1, where nodes *A*, *B* and *C* generate equal amounts of traffic towards node *E*, while the same channel is initially used by all links.

In this scenario, the *DE* link facilitates the aggregated traffic, destined to node *E*. Thus, *DE* should be assigned a frequency such that the aggregate traffic towards *E* is efficiently forwarded, i.e., the bottleneck situation with *DE*, due to potentially low SINR

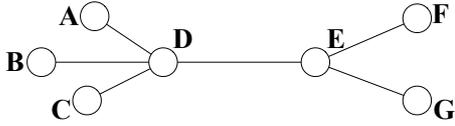


Figure 1: Load-aware channel selection.

or high contention levels, is avoided to the extent possible. Note here that the frequency selection outcome will likely affect the decision of load-aware routing protocols, such as RM-AODV (Radio Metric Ad Hoc On-Demand Distance Vector Routing) [9]. Hence, both the frequency selection and load-aware routing functionalities are inter-dependent and must be considered in conjunction.

In this paper we present a general formulation of the channel allocation problem, taking into account the important parameters (i.e., interference level, packet dropping and transmission rate) that affect the end-to-end throughput in the network. Moreover, we propose ARACHNE, a load and routing aware channel selection protocol for wireless mesh networks. ARACHNE performs end-to-end channel selection along a route, by adopting a variation of a load characterization metric [3]. Given that the load-aware routing choices are affected by the frequency selection policy, ARACHNE combines frequency selection and route selection under the same unified framework and approaches the maximum throughput in the network. To the best of our knowledge, this is one of the first works that present a frequency selection protocol across entire routes, involving both the access and the backhaul levels in conjunction with the load-aware routing of information between end-hosts. We evaluate ARACHNE through trace-driven simulations, using the OPNET [10] simulation platform. We observe that ARACHNE outperforms other channel allocation mechanisms for mesh networks in terms of overall network throughput, average delay and dropped data.

The rest of this paper is structured as follows. In section 2 we present our preliminary experiments and discuss the previous approaches related to our work, which motivate our problem formulation and the design of ARACHNE. Section 3 presents the modelling of channel allocation in wireless mesh networks and describes the airtime metric that we employ, based on which ARACHNE discovers the channel with the maximum long-term throughput. In section 4, we describe in detail the design of ARACHNE for both the access and the backhaul levels. We evaluate our protocol through

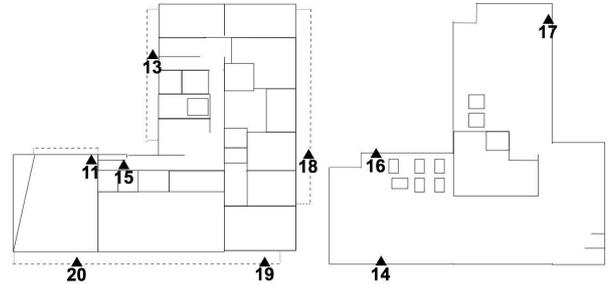


Figure 2: The testbed deployment in the 4th (left) and the 5th (right) floor of our building.

simulations, in section 5. Finally, section 6 concludes this paper.

2. Motivating our Channel Allocation Policy

In this section, we describe a set of preliminary experiments on our wireless testbed, which motivate the design of ARACHNE. We first provide a description of our testbed platform, and subsequently we discuss our experiments and their interpretations. We also discuss relevant previous work.

Our wireless testbed deployment (Fig. 2) consists of 9 nodes that are based on the ORBIT hardware configuration, and run a Debian Linux distribution with kernel v2.6 over NFS. Each node is equipped with 1 GHz CPU, 512 Mbytes of memory, and a WN-CM9 wireless card, which carries the AR5212 Atheros chipset. We set the cards to 802.11g mode and we use channels 1, 6 and 11. For the purposes of these preliminary experiments we consider fixed routes between source and destination. We inject UDP traffic with various constant bit rates (CBR) and with packet size equal to 1500 bytes. For each end-to-end traffic session we set different application data rates; we utilize the *iperf* measurement tool.

We provide a representative experiment in what follows. Consider the following two simultaneously active routes (see Fig. 2): **(a)** 16→15→20→19, and **(b)** 13→15→20→14. These routes have one link in common, 15→20, while all links are of similar quality in terms of achievable throughput in isolation. We apply two different channel selection policies; for both policies we make sure that connectivity is maintained between the end-hosts, i.e., *two consecutive relays use the same channel on one of their interfaces*. We first consider the channel policy, A1, which assigns channels according to the interference experienced (allocates the

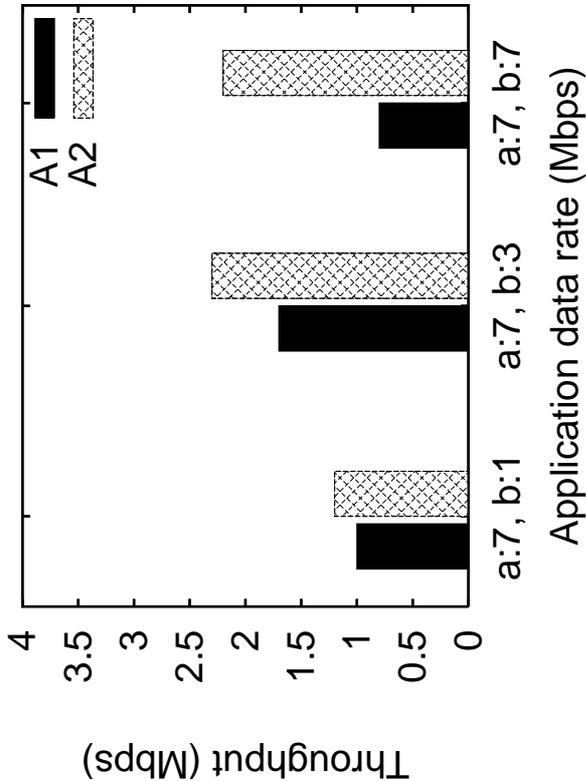


Figure 3: Policy A2 outperforms policy A1, especially at high loads.

channel with the minimum aggregate interference, observed through RSSI measurements). Subsequently we assign channels to nodes, additionally taking into account the link loads, in terms of both link quality and aggregate traffic service; we call this policy A2. In other words, we manually prioritize the channel selection on the link 15→20; the rest of the links choose their channels as previously, *given the selection on link 15→20*. For both approaches, we measure the total end-to-end throughput for the two routes. Fig. 3 depicts these measurements for different source data rates.

First, we observe that the total end-to-end throughput is always higher with A2. Second, we see that the throughput improvement with A2 is much higher than with A1, when the data rates of the routes are of similar magnitude. This is because A2 assigns a “better” channel to the 15→20 link, which has to serve more traffic than the other links. With A1, however, this acts as a bottleneck to the performance of the individual routes: this bottleneck situation cannot be captured

by simply measuring the interference at each channel. On the other hand, when one of the routes, (say (a)) has significantly lower traffic demands, the aggregate traffic that traverses link 15→20 comes mainly from the route (b) over a period of time; in such a case, A1 works quite well, too. This observation constitutes a key element for the design of ARACHNE, in terms of selecting channels and allocating the available interfaces to specific routes, as we explain in section 4.

The LCCS (Least Congested Channel Search) method [11] 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 [12].

Similarly, Leith and Clifford [13] propose a self-managed distributed channel selection scheme, wherein each AP passively measures the received power from the packets transmitted by neighbor APs.

Kauffmann et al. in [4] propose a distributed frequency selection algorithm that minimizes the global interference in the network. In this work the authors minimize the total interference which can result in improved user throughput. However, their algorithm does not consider the number of clients in the network; it assumes purely downlink saturated traffic and that all APs have affiliated clients.

Similarly, Rozner et al. in [5], [14] propose a channel assignment scheme for WLANs, taking into account traffic demands. They show that taking into consideration the current traffic demands at APs and clients, the quality of the channel assignment can be greatly improved.

The MaxChop mechanism [6] provides high levels of fairness among users using channel hopping. However, it requires tight synchronization between AP and clients, while it is difficult to implement efficiently with off-the-shelf hardware.

In [8] the authors study the joint channel allocation and routing problem, assuming that traffic demands and network topology are known. They propose a centralized algorithm that maximizes the aggregate throughput.

Raniwala et al. [7] propose a tree-based mesh architecture, Hyacinth, where the local channel load information exchange facilitates the channel selection. Hyacinth tries to address the joint problem of channel assignment and routing in wireless mesh networks. The latter two approaches, [8, 7], however, assume the avail-

ability of a global network view.

Our protocol differentiates from these approaches by providing efficient end-to-end channel selection in a distributed manner (access level and mesh backhaul). In ARACHNE there is no need of synchronized channel access. Moreover our protocol does not employ any tree-based architecture, which in some cases cannot represent the actual network topology and its dynamics. Finally, our work is fully compliant with 802.11s [9] wireless mesh networks and it can be implemented on top of the existing IEEE 802.11 standards. In what follows, we present the metric that our protocol uses during channel assignment.

3. Modelling Channel Allocation

It is clear that the selfish objective of the APs is to select the channel that provides the maximum average throughput. Therefore, the APs need to first estimate the average throughput it can obtain from when operating on a specific channel. In this section, we investigate important parameters that need to be taken into account for allocating the available channels to the APs in an efficient manner, through a sophisticated metric and a problem formulation that tries to maximize the throughput in the network.

3.1. Metric for channel selection

The metric that we consider in our work is called *airtime metric*. The airtime metric was first discussed in the 802.11s [9] 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*.

Formally, the *airtime cost* of a link l that supports the communication of two routers, using channel or frequency f , is given as:

$$C_{l,f} = \left[O_{ca} + O_p + \frac{B_t}{R_{l,f}} \right] \frac{1}{1 - e_{pt}^f}. \quad (1)$$

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 frame¹. Some representative values for these constants,

¹The transmission of test frames is necessary, in order to derive values for the computation of the airtime cost.

for 802.11g, are: $O_{ca} + O_p = 1.25\text{ms}$ and $B_t = 8224\text{bits}$. Furthermore, $R_{l,f}$ and e_{pt}^f are the current transmission rate and frame-error rate, respectively, in Mbps, for the test frame size B_t in channel f . In other words, the estimation of e_{pt}^f corresponds to transmissions of standard-size frames B_t at the current transmit bit rate $R_{l,f}$.

In order to understand the reasoning behind the airtime metric, we should refer to the recent work investigating the calculation of the average throughput of the 802.11 based WLANs [15], [16], [17], [18]. It has been observed that when there are several flows with different physical transmission rates, then the throughput of all flows is bounded by the slowest transmission rate [19]. In order to explain this anomaly, further studies have been conducted, and the average uplink throughput in a single cell environment is calculated under saturation and decoupling approximations in [15] and [18]. The saturation approximation states that there are always packets backlogged on every user. Meanwhile, with the decoupling approximation it is assumed that when there are n users, the aggregate attempt process of $(n - 1)$ nodes is independent of the back-off process of any given node.

We consider the simple case, when all nodes have the same back-off parameters, each node is the transmitter for a single flow, and all packets have lengths are equal to L . As derived in [18], the total average network throughput $\theta(\beta)$ is given as:

$$\theta(\beta) = \frac{n\beta(1 - \beta)^{n-1}L}{1 + \sum_{i=1}^n \beta(1 - \beta)^{n-1}(\frac{L}{C_i} + T_0) + (1 - (1 - \beta)^n - n\beta(1 - \beta)^{n-1})T_c} \quad (2)$$

where β is the attempt rate (probability) in the equilibrium, C_i is the physical transmission rate of node i , T_0 is the fixed overhead with packet transmission and T_c is the fixed overhead for an RTS collision. Due to the exponential back-off behavior of the nodes and the decoupling approximation, it can be shown that the attempt probability of a node accessing the channel can be determined in terms of a given collision probability γ as:

$$G(\gamma) = \frac{\sum_{k=0}^K \gamma^k}{\sum_{k=0}^K \gamma^k b_k}, \quad (3)$$

where K is the maximum number of attempts allowed under the protocol, and b_k is the mean back-off at the k^{th} attempt. Meanwhile, the probability of collision of an attempt by a node is given by $\Gamma(\beta) = 1 - (1 - \beta)^{n-1}$ due to the decoupling approximation. The equilibrium behavior of the system is governed by the solution of

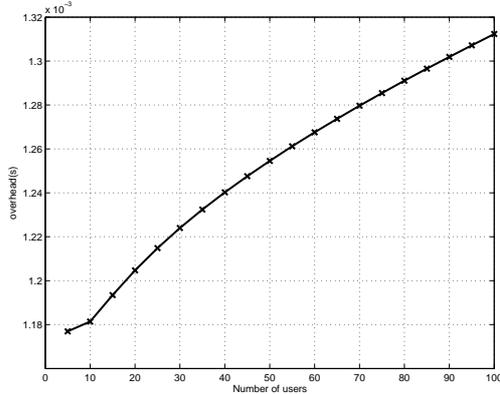


Figure 4: The contention and protocol overhead calculated according to (4).

the fixed point equation $\gamma = \Gamma(G(\gamma))$. If this equation is solved it yields the collision probability from which the attempt rate in the equilibrium β can be determined from (3).

Note that $\frac{L}{\theta(\beta)}$, is the average delay per packet in the equilibrium, which is given as:

$$\rho(n) = \frac{L}{\theta(\beta)} = \frac{1}{n\beta(1-\beta)^{n-1}} + (T_0 - T_c) + \frac{1 - (1-\beta)^n}{n\beta(1-\beta)^{n-1}} T_c + \frac{1}{n} \sum_{i=1}^n \frac{L}{C_i} \quad (4)$$

The first three terms in (4) represent the delay due to channel contention and protocol overheads, and the last term represents the average transmission time of an L length packet by a node in the cell.

Comparing the definition of the airtime metric $C_{l,f}$ and $\rho(n)$, we can see that the sum of O_{ca} and O_p correspond to the approximation of the first three terms in (4) with a constant value. Clearly, the transmission of RTS, CTS, DATA and ACK frames may also get corrupted not only due to collisions but also due to channel errors. The authors in [16] extended the analysis in [15] to account for wireless bit errors. The goodput expression in [16] assumes that the wireless channel can be modeled by an appropriate Gilbert model with known transition probabilities. In wireless networks with dynamically changing conditions, such an assumption is not practical. Therefore, in our method, the current frame error rate, e_{pt}^f is measured by the users and AP. For each packet attempted to be transmitted, it would be in error due to channel errors with probability e_{pt}^f . Clearly, the average number of attempts until successful trans-

mission would be $1/(1 - e_{pt}^f)$. For each attempt on the average $\rho(n)$ amount of time is experienced. The product of $1/(1 - e_{pt}^f)$ and $\rho(n)$ gives the airtime metric used in our algorithm.

When the overhead is calculated as described in (4), we observe that it varies with the number of nodes in the cell. The overhead with respect to the number of nodes is given in Figure 4, when $K = 7$, $b_0 = 16$, $b_k = 2^k b_0$, the slot length is $20\mu s$, $T_0 = 52$ and $T_c = 17$ slots. The overhead varies between 59 and 66 slots, and thus, it is clear that there is weak dependency between the number of users n and the observed overhead. Therefore, in our work the overhead is taken as a constant to avoid additional computational complexity.

3.2. Optimal channel allocation problem formulation

In this section we give a general formulation of the channel allocation in 802.11-based multi-radio wireless networks. We formulate the problem of optimal allocating the available channels, assuring that the maximum path end-to-end delay (total airtime cost of the links that are part of the path) in the network is minimized. Based on our previous analysis this will provide throughput maximization and balanced/stable network operation is assured [20]:

$$\min_{a_{l,f}} \max_{p \in P} \sum_{l \in L_p} \sum_{f \in F} a_{l,f} C_{l,f} \quad (5a)$$

$$\text{s.t. } a_{l,f} = [0, 1], \quad \forall l \in L_p, \quad \forall p \in P, \quad \forall f \in F, \quad (5b)$$

$$\sum_{f \in F} a_{l,f} = 1, \quad \forall l \in L_p, \quad \forall p \in P, \quad (5c)$$

$$C_{l,f} = \left[O_{ca} + O_p + \frac{B_t}{R_{l,f}} \right] \frac{1}{1 - e_{pt}^f}, \quad \forall l \in L_p, \quad \forall p \in P, \quad \forall f \in F \quad (5d)$$

Where P is the set of active paths in the network, F is the set of available channels/frequencies that could be used by the APs, L_p is the links that belong to path p , $a_{l,f}$ is the probability that frequency f is used in the communication through link l , and $C_{l,f}$ is the airtime cost of link l that operates in frequency f .

4. Dynamic Channel Selection Protocol

Inspired by the previous problem formulation and the *airtime metric* association with the average throughput in the network, we present ARACHNE, a load

and routing aware channel selection protocol. The design of ARACHNE makes the following assumptions: a) Nodes of the mesh backhaul are equipped with at least 3 radios (2 for the mesh backhaul communication and 1 for the access level), one of which is set to a pre-arranged channel, c_p , which is the same for all nodes in the network, b) A client associates to the AP with the strongest signal (RSSI) among all neighbor APs (802.11-based association procedure), c) The set of channels used at the access level (AP-client interfaces), is different from the channel-set used at the mesh backhaul. For simplicity in this paper, the AP-client interfaces operate in the 2.4 Ghz band; the 5 GHz band is used exclusively for the interfaces in the mesh backhaul, and d) We assume that traffic is exchanged among end hosts belonging to the same mesh network, i.e., traffic is not crossing different networks.

As previously mentioned, ARACHNE is using *airtime metric* that is an approximation of the per packet latency. The airtime metric was first discussed in the 802.11s [9] 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 ARACHNE for the purposes of our proposed channel selection functionality.

Several studies have shown that the number of erroneously received packets increases and the transmission rate decreases when the cells in the network interfere [21], [22], [4]. The *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. Besides, our analysis shows that the average *airtime cost* is an approximation of the average per-packet delay. Therefore, the average *airtime cost* is a representative metric that reflects the channel performance and also approximates the maximum throughput in the network.

ARACHNE captures the performance of a bidirectional link in terms of estimated throughput at a particular channel, by measuring the average *airtime cost* for both uplink and downlink, $C_{l,f} = \overline{C_{l,f}^{up}} + \overline{C_{l,f}^{down}}$ (*airtime cost* for a particular link l that operates on channel f), and applies a channel selection methodology where the channel with the minimum $C_{l,f}$ is chosen. This channel selection policy determines the frequency with the minimum average per-packet delay in both up-

link and downlink, thereby approximating the maximum throughput in the communication. The goal of ARACHNE is to assign channels to mesh nodes (their associated links), such that the average airtime metric is minimized, both at the access and the backhaul levels. Hence, ARACHNE involves two procedures, P1 and P2, one for each level. ARACHNE is executed exclusively by the nodes that belong to the mesh backhaul, i.e., the relay nodes as well as the APs that connect the end-hosts with the mesh network. The channel discovery is initiated by the AP of the source host. In what follows, we describe the operations of the protocol, P1 and P2, for each of the two levels of operation.

4.1. Channel Selection at the Access Level (P1)

The access level involves the communication of the end hosts (clients) with mesh nodes at the edge of the backhaul (APs). Let us assume that host A wishes to send traffic to another host B . This traffic will flow through A 's AP, M_A , to the mesh backhaul and it will eventually reach M_B ; the latter will finally forward the traffic to B . With ARACHNE, the frequency selection at the access level involves a channel discovery process at the two aforementioned APs, in order to find the channel with the minimum airtime cost value (we provide the steps of this process below). The calculation of this value is performed for every scanned channel, and involves all the downward and upward links of M_A , M_B . Let us again assume that host A wants to send traffic to host B . This part of the protocol, P1, is executed by the APs at the access level, and consists of the following steps.

1. **[P1a]. Deriving interference information for the current channel.** At the nominal start of ARACHNE, M_A is informed about the operational frequencies of its neighbor *co-channel* APs. This can be performed through passive scanning of periodically transmitted beacons [23], [24]. By the end of this step, M_A is aware of the total received power from all co-channel APs.
2. **[P1b]. Computing the downlink airtime cost.** M_A calculates the aggregated downlink airtime cost with its clients, through a link performance-measurement procedure, described in detail in [3] where airtime metric is used for user association.
3. **[P1c]. Computing the uplink airtime cost.** The clients of M_A calculate their individual uplink costs. They piggy-back this information through their data frame transmissions towards M_A . By the end of this step, M_A has received information regarding the uplink channel qualities from all its clients.

4. **[P1d]. Deciding if the current channel is appropriate.** M_A receives the information from the client and computes the average airtime cost (for both uplink and downlink) for the access level. If this is higher than a pre-defined threshold T , M_A remains in the same channel; otherwise it initiates a channel discovery process. The value of threshold T is decided and controlled by the system designer.
5. **[P1e]. Computing the cumulative airtime cost at the next available channel.** M_A and its clients switch to the next channel and repeat steps $P1a - P1e$. channels have been visited, M_A finally selects the channel with the minimum average airtime cost (for both uplink and downlink).

Note here that process P1 (of selecting a channel at the access level) is actually independent of P2; although they are executed in parallel, they do not affect each other to a large extent, in terms of convergence time, since they utilize different sets of channels.

4.2. Channel Selection at the Mesh Backhaul (P2)

The role of the mesh backhaul is to serve the forwarding of packets towards their final destinations. Undoubtedly, the network connectivity is affected by a potential channel selection policy that may be applied in the mesh backhaul. ARACHNE's channel selection framework ensures connectivity at all times in the network, while the selection of a frequency at a particular link takes into account the load of the routing sessions that are traversing the link.

For the purposes of this study, we assume that the RM-AODV routing protocol [9] is applied on the mesh backhaul. RM-AODV is the default routing protocol proposed in the context of 802.11s mesh networks, where the airtime cost is used as a routing decision metric in the mesh backhaul. In particular, the airtime metric is appended to the RREQ and RREP messages, during the route discovery process; finally, the end-to-end path with the minimum total airtime cost (the route that has the minimum load) is selected. Our choice of RM-AODV is motivated by the fact that this routing protocol is based on the airtime metric. *Note, however, that ARACHNE can operate in conjunction with any load-aware routing protocol!*

ARACHNE employs the Local Association Base Advertisement (LABA) mechanism introduced in 802.11s [9], in order to disseminate information with regards to inform the entire mesh network about the clients (end-hosts) that are associated with each mesh AP (MAP). MAPs periodically broadcast LABA frames, which consist of the MAC addresses of the hosts that are associ-

ated with. We have enhanced the LABA frames to include load related information, as we explain later in the description of the main parts of our protocol.

Our protocol prioritizes the channel assignment on the most loaded mesh APs, as well as on the mesh APs that are expected to be highly loaded in the near future. In particular, ARACHNE observes the per-link load, for the links that serve one or more routes in the mesh; the load is captured in the airtime cost metric. The main problem in designing distributed channel selection policies is the channel dependencies that arise between the nodes that are part of the mesh network. For example, we assume that in the simple network in Fig. 1 nodes are equipped with 2 wireless interfaces. Node D uses channel a in order to communicate with node E and node E uses channel b to communicate with nodes F and G. In case that in the link E-F the current channel is heavily loaded, node E uses a new channel c that operates better at the link E-F. As E has only two interfaces, channel c must be used at the link E-G too. Unfortunately, a strong dependency between links E-F and E-G is established. Channel c may not operate effectively at the link E-G and therefore the performance of the network is decreased. The aforementioned *ripple effect* could be further propagated in dense/huge mesh networks. In order to avoid *ripple effects* in the channel selection process [7], we categorize the wireless interfaces at each mesh AP that serves as a relay as:

- **IN network interfaces** that are used for data reception,
- **OUT network interfaces** that are used for data forwarding.

In addition, each node can assign channels only to its **OUT** interfaces. The **IN** interfaces follow the channels that are assigned by the nodes that communicate with the current node. In our protocol we assume that each node is equipped with at least one **IN** and one **OUT** interface.

In what follows, we describe the steps that are executed by ARACHNE, for assigning channels at the mesh backhaul (P2). To begin with, we consider a network state, in which: (a) All nodes have already dedicated one of their **OUT** interfaces (I_p), to control channel c_p ; (b) The RM-AODV protocol has converged to a set of routes between end-hosts; (c) All of the other wireless interfaces of a mesh AP (besides I_p) have been randomly assigned a channel in the 5 GHz band. The channel assignment in P2 is comprised of the following steps.

1. **[P2a]. Constructing a priority list.** With ARACHNE, an AP starts the channel-scanning process at a time-instant dictated by its priority ranking. This priority is highly-related to the load of each mesh AP; the higher the load (the data that must forward), the higher the priority. In addition, we believe that in this prioritization the estimated load of a mesh AP in near future must be taken into account. In this way we guarantee that our protocol will converge to a long-lasting and a stable channel allocation. Hence, it is imperative that this list is constructed, before APs start scanning each channel. Each mesh AP a calculates its priority rank according to its current load and its expected load, as:

$$P_r^a = L_{cmt}^a w_1 + L_{est}^a w_2,$$

where L_{cmt}^a is the current load, L_{est}^a is the estimated load and w_1, w_2 are the weights that are used in the calculation (we will give more details about the selection of w_1, w_2 in the evaluation of our protocol). As far as the calculation of the estimated load L_{est}^a is concerned, we adopt the estimation method (based on historical data), proposed in [25]. In this approach the authors design a trace-based traffic model in order to predict the aggregated traffic demands of an AP in near future. Time-series analysis is the basis of this traffic estimation model. The accuracy of this model is high and in combination with our load-aware channel selection protocol we achieve long-lasting and stable channel allocation in the mesh network.

Besides, the priority ranks are incorporated into the LABA frames and are broadcasted in the network. Consequently, a sorted list (in terms of channel selection priority rank) is disseminated and maintained by all APs. Hence, by the end of this step each AP knows the priority of all APs in the network.

2. **[P2b]. Performing channel-scanning.** Let us assume that each AP can scan a set of K channels. The first mesh AP in the priority list measures the cumulative airtime cost for each of the K channels, and constructs a local, *temporary*, sorted list with the cumulative airtime cost, per channel.
3. **[P2c]. Assigning route sessions to available *OUT* interfaces.** Each mesh AP typically prefers to assign low-airtime channels to links that serve high-load end-to-end traffic sessions. ARACHNE manages to efficiently utilize the available spectrum, by providing a balanced channel and interface assignment in the network. A mesh AP a calculates the

maximum load L_{sh}^a that can be assigned at each of its *OUT* interfaces, as (clearly, when the number of available interfaces is higher than the number of served sessions, ARACHNE allocates an interface to a particular session): $L_{sh}^a = L_{cmt}^a / N_{OUT}^a$, where N_{OUT}^a is the number of the available *OUT* interfaces of mesh AP a , and L_{cmt}^a is the current load that must be served by the *OUT* interfaces. The main constrains of the interface-assignment strategy are: (a) The load assigned to an *OUT* interface is less than L_{sh}^a , and b) The load is proportionally allocated to the available *OUT* interfaces in terms of the load of the flows that pass through the current mesh AP. In other words, the load is balanced to the available *OUT* interfaces.

4. **[P2d]. Assigning channels to *OUT* interfaces.** In order to assign a channel to an *OUT* interface, a coordination with the *IN* interface of the mesh AP that receives the traffic is required. In particular, a mesh AP A that selected the channel with the minimum airtime cost in the communication with a mesh AP B , sends an RTC frame to B announcing in this way that a new channel must be used in their communication. Then, B responds with a CTC frame and sets one of its *IN* interfaces to the selected channel. In case that B hasn't responded in a time window (timeoff), A retransmits its RTC frame. We must note here that there are special situations where B has recently assigned channels to its *IN* interfaces (during the same protocol iteration) and there are no available interfaces to assign the channel proposed by A . In other words, higher priority APs (compared to the priority of A) have previously assigned those channels to B 's *IN* interfaces and there are no available *IN* interfaces to use. In this case B responds with an XTC frame announcing this situation to A and the available channels that are assigned to its *IN* interfaces. Lastly, A is restricted to use one of those channels in the communication with B (the channel with the minimum airtime cost in the link A-B). After the completion of the aforementioned process, the AP sends a "good-to-go" unicast² message to the next AP in the list, to initiate its channel scanning and assignment process.
5. **[P2e]. Selecting channels iteratively.** When all APs have completed a round of channel scanning, RM-AODV updates the routing tables at the

²This unicast message speeds up the process of completing a full iteration of channel scanning for all participating mesh APs.

mesh backhaul, taking into account the new state of the network. Note that RM-AODV discovers routes based on the computation of the airtime cost, which will have likely changed after an iteration of the channel selection process. Steps P2a-P2e are repeated until the channel selection has converged.

5. Evaluating ARACHNE

In this section we present the evaluation study of our approach. During the first part of our validation methodology we use the optimal solutions (resulted from solution of the channel allocation problem described in section 3) as a benchmark in the direction of evaluating ARACHNE. We are using IBM ILOG CPLEX Optimizer [26] to find the optimal solutions. Moreover, we evaluate ARACHNE through extensive simulations, which import both synthetic [10] and real traces [27], [28]. We compare our protocol against other channel selection schemes, and we present ARACHNE’s predominance in terms of the total network throughput, average packet dropping and average transmission delay.

We have implemented ARACHNE in OPNET [10]. We have also implemented the RM-AODV protocol and we consider this routing protocol in our simulations. The clients are uniformly distributed (at random) in the 1000m×1000m simulation area. All nodes use a default transmit power of 20 dBm and the source-destination pairs are randomly chosen in the network. We experiment with: (a) fully-saturated, end-to-end UDP traffic, (b) VoIP traffic and (c) real traffic traces from Dartmouth College [27] and IBM [28]. The weight values that are used in our simulation experiments are: $w_1 = 0.6$ and $w_2 = 0.4$ (the current load affects the channel selection procedure more than the estimated load in near future). As far as the duration of the measurement period is concerned, in our experiments the APs remain 100ms in each scanned channel in order to gather the necessary information. The convergence of our mechanisms is reached rather quickly in our network topologies, as will be presented later in the current section (close to 20sec in average and after a small number of iterations). Moreover, the traffic keeps flowing during the execution of ARACHNE and the network operation is not affected. We chose accordingly the values of the performance threshold T in order to avoid unnecessary and frequent protocol executions. Throughout our simulation experiments we compare the network performance with ARACHNE, against the single-channel assignment, a random-channel allocation strategy, as well as the Hyacinth protocol [7].

Table 1: Performance with UDP traffic and 12 orthogonal channels available

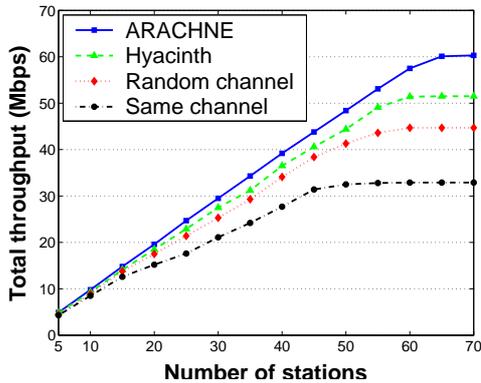
Number of APs	Optimal Throughput (Mbps)	Throughput achieved by ARACHNE (Mbps)	Average ARACHNE iterations
10	48.2	45.1	3
20	79.4	76.4	5
30	131.2	121.3	8
40	164.8	151.1	11
50	208.6	192.7	15
60	249.1	230.2	19

Table 2: Performance with UDP traffic and 3 orthogonal channels available

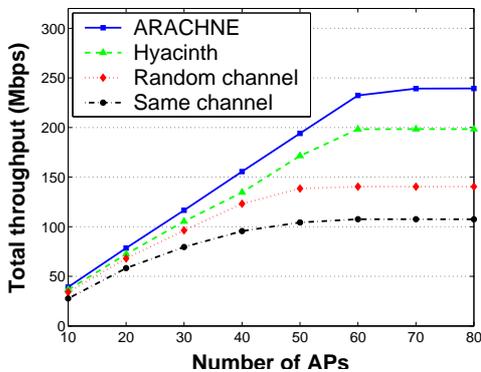
Number of APs	Optimal Throughput (Mbps)	Throughput achieved by ARACHNE (Mbps)	Average ARACHNE iterations
10	42.7	38.4	2
20	64.8	59.2	3
30	112.2	102.2	6
40	139.1	121.7	9
50	194.2	172.4	11
60	227.5	211.8	15

5.1. Simulations with saturated UDP traffic and VoIP

To begin with, we opt to observe the variation in the total network throughput, as we increase the number of source-destination pairs. For this, we progressively increase the number of associated clients from 5 to 70. Here the network consists of 10 mesh APs, each of which is equipped with 2 wireless interfaces for the backhaul communication. We observe in Fig. 5(a) that the performance with ARACHNE is similar to the other 3 policies, when the number of clients (and therefore the communication load) is low. However, with increased load ARACHNE manages to provide much higher end-to-end throughputs (up to **85%** difference!), due to its efficient end-to-end channel selection strategy. Hyacinth is designed especially for wireless internet traffic which is directed to/from the wireless gateways. In our simulation scenario where saturated UDP traffic is supported, the tree-based architecture is incapable to support dynamic traffic variations in the network (especially in high communication load). In low load conditions Hyacinth performs close to ARACHNE. Contrarily, when the number of clients in the networks increases the performance drops and the throughput saturation point is reached quite quickly.



(a) Total throughput Vs. # clients.



(b) Total throughput Vs. # APs.

Figure 5: ARACHNE provides very high total end-to-end throughput with saturated UDP traffic.

Next, we examine how close ARACHNE is performing to optimal and its scalability. We measure the total end-to-end network throughput as we increase the number of APs (i.e., the mesh density) from 10 to 80 (we increase the number of clients along with the mesh APs: 10 APs - 40 clients, 20 APs - 80 clients, etc.). Note that the interference is dynamically changing while the number of the APs in the network increases; the airtime cost metric manages to effectively capture the varying co-channel interference.

The first objective in our study is to examine how efficient ARACHNE approaches the optimal network performance. Our methodology includes the incorporation of the operational parameters in the IBM ILOG CPLEX solver while the number of the existing APs in the network increases. The optimal channel allocation then is applied in our simulation scenarios (OPNET). We apply ARACHNE under the same operational parameters of the network and compare the achieved throughput to the optimal throughput. In table 1, we observe the total net-

Table 3: Performance with VoIP traffic and 12 orthogonal channels available

Number of VoIP sessions	Optimal end-to-end delay (s)	End-to-end delay achieved by ARACHNE (s)	Average ARACHNE iterations
4	0.0001	0.0002	2
8	0.0008	0.001	4
12	0.001	0.004	7
16	0.006	0.009	11
20	0.011	0.016	14
24	0.014	0.019	18

Table 4: Performance with VoIP traffic and 3 orthogonal channels available

Number of VoIP sessions	Optimal end-to-end delay (s)	End-to-end delay achieved by ARACHNE (s)	Average ARACHNE iterations
4	0.0009	0.0012	2
8	0.0014	0.0028	3
12	0.0044	0.0069	4
16	0.0092	0.014	8
20	0.016	0.02	9
24	0.019	0.024	12

work throughput that is achieved by ARACHNE after a small number of iterations till convergence, compared to the optimal network throughput (when 12 orthogonal channels are available). Moreover, table 2 contains similar results when there are just 3 available orthogonal channels available to allocate in the network. A general outcome is that ARACHNE is able to approach the optimal network performance and scales efficiently while the network topology and the load vary. The best performance is achieved within the case of 12 orthogonal channels, where ARACHNE gets as input more frequencies to allocate and converges to channel allocation solutions close to optimal (by applying its sophisticated policy).

Fig. 5(b) depicts the network performance in terms of total network throughput. We observe that ARACHNE is able to adapt to topology and load variations, and therefore it manages to provide the most beneficial selection of the available channels at the mesh backhaul. The tree-based architecture that is introduced by Hyacinth does not scale well and therefore the achieved performance is worst.

In order to observe the performance of our protocol with delay-sensitive data exchange, we utilize varying, parallel, end-to-end VoIP traffic sessions. Tables 3 and

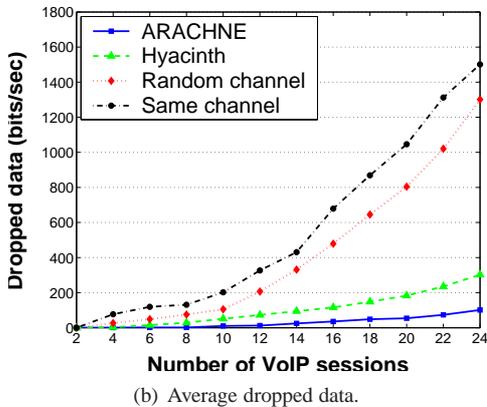
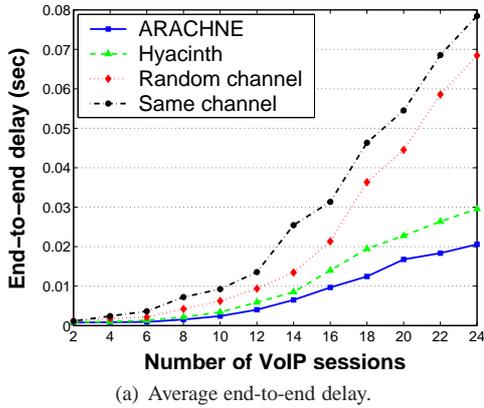


Figure 6: VoIP simulation results.

4 present the average end-to-end delay that ARACHNE achieves (when 3 and 12 orthogonal channels are supported in the network). The optimal end-to-end delay values are used (again) as a benchmark. We observe that ARACHNE can perform close to optimal after a lightweight execution in the network (i.e. converges after a small number of iterations).

Fig. 6(a) depicts the average end-to-end delay of VoIP packet transmissions. We observe that ARACHNE achieves low end-to-end delays, since it considers the load of the individual links across a route in the process of assigning frequencies to wireless interfaces. Recall here that ARACHNE is expected to provide high benefits when operating in conjunction with load-aware routing protocols, such as RM-AODV (this is case with our simulations), since they both take into consideration the load that is experienced by individual links. Furthermore,

Fig. 6(b) shows the average number of dropped data packets due to channel errors and contention. The performance of ARACHNE is impressive, since packet

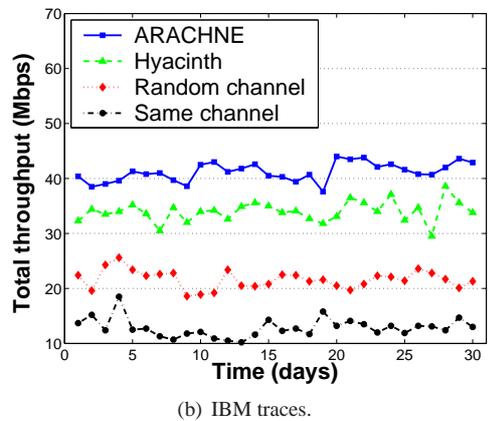
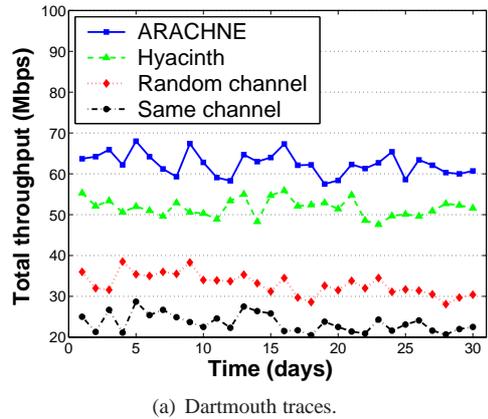
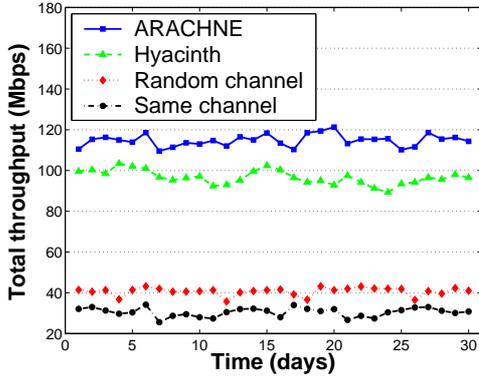


Figure 7: Simulation results with real traces: 2 interfaces per node.

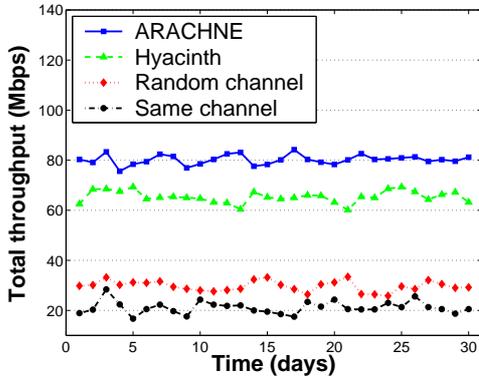
dropping is kept in very low levels, as compared to Hyacinth and to the other strategies. Hyacinth does not consider the channel allocation at the access level and therefore, it can support less VoIP sessions than ARACHNE can (which provides end-to-end channel selection and therefore the provided QoS to the VoIP clients is higher).

5.2. Simulation results with real traces

Next, we present our simulation experiments with real traces from Dartmouth College and IBM. This will reveal the behavior of ARACHNE in more realistic scenarios. By analyzing the SNMP logs from each AP, we derive the dynamic behavior of the aggregated traffic demand. Note that these traces have been gathered from real wireless networks settings. Therefore, they represent actual traffic demands and behaviors of the APs and clients in a real network. The traces from Dartmouth College [27] (3/2001) involve the buildings labeled as



(a) Dartmouth traces.

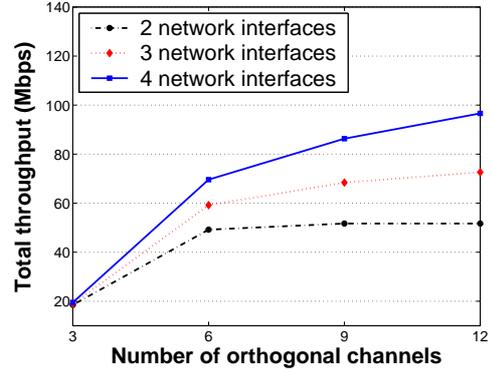


(b) IBM traces.

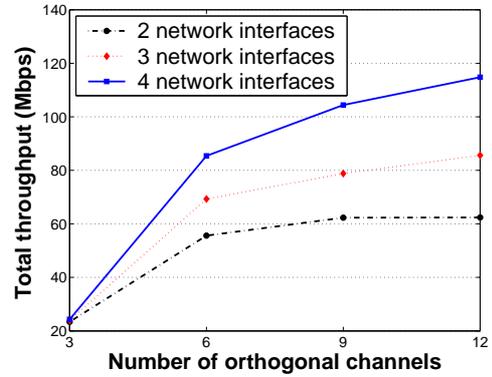
Figure 8: Simulation results with real traces: 4 interfaces per node.

AcadBldg10, *SocBldg4*, and include approximately 40 APs (the traces contain the AP location; we have placed them accordingly in our simulation space). In addition, the traces from IBM [28] (8/2002) have been gathered from the buildings labeled as *MBldg*, and *SBldg*. The number of the APs that are placed in these buildings varies (the traces don't contain their locations and therefore we have simulated various random topologies with them). Note that the number of active clients (end-hosts) varies in both trace sets and they are uniformly distributed in our experiments (the traces do not contain information about client locations). In order to import the traces in our simulation we periodically compute the average rate for each client, once every 3 minutes, and we correspond the computed average values to the traffic demands of each client.

To begin with, we consider 2 wireless interfaces at each AP for the backhaul communication. Fig. 7(a) and 7(b) show that ARACHNE achieves the highest total network throughput in all cases. Our simulation



(a) Hyacinth performance.

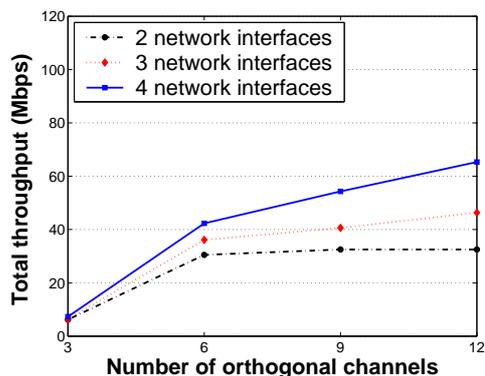


(b) ARACHNE performance.

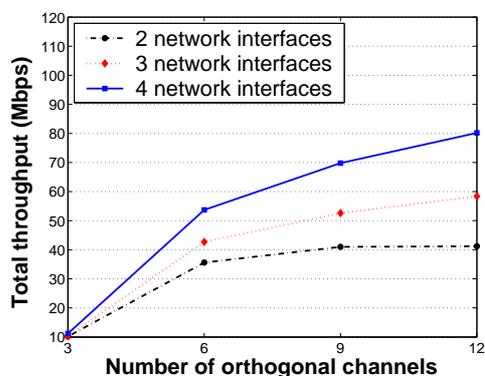
Figure 9: Simulation results with real traces: Varying the number of the supported interfaces and the available channels in the network with Dartmouth traces.

experiments reveal that channel selection at the access level plays an important role in the overall end-to-end performance. Being based on the airtime cost metric, ARACHNE takes into account the channel conditions as well as the communication load in both access level and the mesh backhaul (not considered by the other approaches, like Hyacinth) it manages to boost the network throughput. Furthermore, Fig. 8(a) and 8(b) depict the performance of the compared strategies when the mesh APs are equipped with 4 wireless interfaces. The results are similar as above.

Last but not least, we perform some simulations where we vary both the number of available channels and the number of the wireless interface supported by each AP. Fig. 9(a), 9(b), 10(a) and 10(b) show that ARACHNE achieves higher throughput values compared to Hyacinth [7]. Throughput improvement is significantly increasing while we the number of available channels and the wireless interfaces increase.



(a) Hyacinth performance.



(b) ARACHNE performance.

Figure 10: Simulation results with real traces: Varying the number of the supported interfaces and the available channels in the network with IBM traces.

6. Conclusions and Future Work

In this paper we presented a channel allocation methodology with main objective to maximize the end-to-end throughput in wireless multi-radio mesh networks. The modeling of the channel allocation problem inspired the design of a routing-aware channel selection mechanism (ARACHNE). ARACHNE adopts a metric that provides an estimation of the average packet transmission delay in a communication channel. In addition, ARACHNE captures the interference effects in the networks and applies an end-to-end channel selection policy that manages to provide approximately the maximum end-to-end network throughput. Our work was compared against 3 other representative channel selection approaches and we showed that it outperforms all of them, for different traffic scenarios and network densities. Moreover, the simulation results showed that ARACHNE converges after a small number of itera-

tions, proving in this way its' lightweight operation.

References

- [1] Microsoft Research. Self-organizing neighborhood wireless mesh networks project. <http://research.microsoft.com/mesh/>.
- [2] Meraki Networks Inc. <http://meraki.net>.
- [3] G. Athanasiou, T. Korakis, O. Ercetin, and L. Tassiulas. Dynamic Cross-Layer Association in 802.11-based Mesh Networks. In *IEEE INFOCOM*, 2007.
- [4] 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.
- [5] E. Rozner, Y. Mehta, A. Akella, and L. Qiu. Traffic-Aware Channel Assignment in Enterprise Wireless LANs. In *IEEE ICNP*, 2007.
- [6] A. Mishra, V. Shrivastava, D. Agarwal, and S. Banerjee. Distributed Channel Management in Uncoordinated Wireless Environments. In *ACM MOBICOM*, 2006.
- [7] A. Raniwala and T. Chiueh. Architecture and Algorithms for an IEEE 802.11-Based Multi-Channel Wireless Mesh Network. In *IEEE INFOCOM*, 2005.
- [8] M. Alicherry, R. Bhatia, and L. Li. Joint Channel Assignment and Routing for Throughput Optimization in Multi-radio Wireless Mesh Networks. In *ACM MOBICOM*, 2005.
- [9] IEEE 802.11s: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: Simple Efficient Extensible Mesh (SEE-Mesh) Proposal.
- [10] OPNET-Radio/Wireless Models. <http://www.opnet.com>.
- [11] J. Geier. Assigning 802.11b Access Point Channels. In *WiFi planet*, 2002.
- [12] 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.
- [13] D. J. Leith and P. Clifford. A Self-Managed Distributed Channel Selection Algorithm for WLANs. In *WiOPT*, April 2006.
- [14] P.H. Pathak. A Survey of Network Design Problems and Joint Design Approaches in Wireless Mesh Networks. In *IEEE Communications Surveys and Tutorials*, 2011.
- [15] G. Bianchi. Performance Analysis of the 802.11 DCF. In *IEEE JSAC*, vol. 18, pp. 535-547, March 2000.
- [16] N. Gupta and P. R. Kumar. A Performance Analysis of the 802.11 Wireless LAN Medium Access Control. In *Comm. in Inform. and Systems*, Vol. 3, No. 4, pp. 279-304, September 2004.
- [17] T-C Hou, L-F Tsao, and H-C liu. Analyzing the Throughput of 802.11 DCF Scheme With Hidden Nodes. In *IEEE VTC*, 2003.
- [18] A. Kumar, E. Altman, D. Miorandi, and M. Goyal. New Insights from a Fixed Point Analysis of Single Cell IEEE 802.11 WLANs". In *IEEE INFOCOM*, 2005.
- [19] M. Heusse, F. Rousseau, G. Berger-Sabbatel, and A. Duda. Performance Anomaly of 802.11b. In *IEEE INFOCOM*, 2003.
- [20] D. Bertsekas and R. Gallager. Data Networks. In *Englewood Cliffs, NJ: Prentice-Hall*, 1992.
- [21] D. Niculescu. Interference Map for 802.11 Networks. In *ACM IMC*, 2007.
- [22] D. Niculescu, S. Bhatnagar, and S. Ganguly. Channel Assignment for Wireless Meshes with Tree Topology. In *COMM*, 2010.
- [23] 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.
- [24] I. Broustis, K. Papagiannaki, S. V. Krishnamurthy, M. Faloutsos, and V. Mhatre. Measurement-Driven Guidelines for 802.11

- WLAN Design. In *IEEE/ACM Transactions on Networking*, June 2010.
- [25] L. Dai, Y. Xue, B. Chang, Y. Cao, and Y. Cui. Integrating Traffic Estimation and Routing Optimization for Multi-Radio Multi-Channel Wireless Mesh Networks. In *IEEE INFOCOM*, 2008.
- [26] IBM ILOG CPLEX Optimizer. <http://www-01.ibm.com/software/integration/optimization/cplex-optimizer/>.
- [27] Dartmouth Campus-Wide Wireless Traces. <http://crawdad.cs.dartmouth.edu/meta.php?name=dartmouth/campus>.
- [28] IBM Wireless Traces. <http://nms.lcs.mit.edu/mbalazin/wireless/>.