

Contention and Traffic Load-aware Association in IEEE 802.11 WLANs: Algorithms and Implementation

Stratos Keranidis

Thanasis Korakis

Iordanis Koutsopoulos

Leandros Tassiulas

Department of Computer and Communication Engineering, University of Thessaly, Greece
Centre for Research and Technology Hellas, CERTH, Greece
E-mail: {efkerani, korakis, jordan, leandros}@uth.gr

Abstract—Efficient association of a station with the appropriate access point has always been a challenging problem. The standard approach of considering only the Received Signal Strength, has recently been substituted by more efficient schemes that consider channel conditions, cell population etc. However, in spite of the large variety of approaches, several factors that determine to a large extent user throughput after association with an access point have been overlooked. In this work, we propose innovative metrics on which association should be based. First, we capture the contention from one-hop and interference from two-hop neighbors that is inherent in IEEE 802.11 WLAN environments. Second we include the PHY transmission rate and show preference to higher rates that reduce the above effects. Third, unlike most relevant approaches, we define an activity factor that reveals the anticipated activity due to backlogged traffic. We devise an association protocol suite, through which messages containing the information above are passed between the AP and the user to support association decisions for the uplink and downlink. We implement the proposed mechanism using the MAD-WiFi open source driver and moreover show through experiments in a wireless testbed that it significantly improves user performance in real conditions.

Index Terms—Wireless communications, Association, Handoff, MAC, IEEE 802.11

I. INTRODUCTION

In IEEE 802.11 WLANs, each station (STA) has to first associate with an access point (AP), before it can start transmitting data to other nodes in the network. This association procedure consists of four phases. During the first phase, a STA has to discover the networks in its vicinity before it can join a Basic Service Set (BSS). This process is called *scanning* and can be either passive or active. In passive scanning, a STA scans all available channels and listens to information periodically broadcasted by the APs in their beacon frames. In active scanning, a STA tries to find the BSSs in its vicinity by transmitting a *Probe Request* frame on each channel of the channel list. APs respond by sending *Probe Response* frames. Having collected these frames, the STA *decides* which AP it will associate with, in the second phase. According to the standard [1], AP selection is based on the Received Signal Strength Indication (RSSI). A STA simply selects the AP from which it has received the strongest signal during the scanning process. In the third phase, the STA has to follow the *authentication* process if the selected AP follows some authentication mode. Finally, the STA sends an *Association Request* frame to the selected AP and sequentially the AP responds with an *Association Response* frame. If the Association Response

frame is received with a "successful" status value, the STA is now associated with the AP.

The rest of the paper is organized as follows. In the remaining of this section the state of the art related work is presented and a summarization of our contribution follows. A detailed analysis of our metric definition follows in section II. Details about the proposed algorithms and their implementation are provided in section III. The configuration of our experiments, concerning the testbed and the methodology used is then described in section IV. In section V, we experimentally evaluate the performance of our implementation. Finally, in section VI, we present the conclusion and discussion of future work.

A. Related Work

The performance of the standard AP selection policy has been extensively studied and it is well known that it leads to inefficient use of the network resources [2],[3]. In addition, due to the asymmetric nature of the wireless medium, this policy becomes unsuitable, as RSSI is an indicator just for the downlink channel and not for the uplink. An association mechanism considering signal to interference and noise (SINR) per connection, as well as asymmetric traffic was proposed in our previous work [4]. Although our approach considered uplink channel conditions as well, thus offering a significant improvement, it was not able to lead to the best available throughput performance.

One of the major issues studied among relevant works has been the proper definition of AP load. The authors in [5], proposed an AP selection policy that estimates AP load based on instantaneous measurements of the transmission rate and the fraction of time an AP acquires the channel for its transmissions. However, this model faces the disadvantage of considering only downlink traffic and therefore assumes that channel contention is only among APs. Another common assumption of works on the field has been to denote AP load as a factor reflecting the AP's inability to satisfy the requirements of its associated users [2],[6]. Another approach followed in [7], bases association decisions on a metric denoted as *airtime cost*, which considers both uplink and downlink traffic as well as AP load. The above approaches, have the common characteristic of considering the effect induced by transmissions only of associated users in the AP load estimation.

However, since the IEEE 802.11 MAC layer is based on contention, the efficiency of an AP is not only dependent on

the STAs associated with it, but also on other neighboring stations and their activity.

Trying to address this issue the authors in [8] consider AP load over all neighboring nodes. The new scheme incorporated both the effects of associated and contending nodes in its throughput estimation algorithm. However, this approach was restricted in considering only downlink transmissions and setting fixed transmission rates, neglecting the importance of rate adaptation mechanisms.

All the above approaches follow the assumption of fully saturated traffic, which considers that all users transmit and require service at all times. In [9], the authors suggest that APs should assign an activity level estimator to their associated STAs based on observations of their traffic intensity. Nevertheless, this approach does not manage to characterize the traffic intensity of neighboring nodes that belong to adjacent cells, although these contend for channel usage or even interfere with transmissions in the cell under consideration. Moreover, they suggest that an Inter-AP protocol is required, that is used to collect activity estimations about all STAs in the WLAN and feed this information to a central entity that calculates the optimal association scheme, considering aggregate WLAN throughput. However, such centralized approaches can only apply to centrally managed deployments, which is not the general case.

One more issue that has not received much attention in the association process, is the effect of hidden node terminals, which appears very often in dense WLANs. In a later work in [10], a metric is proposed, that comprises contention and interference as well. The authors trying to estimate the effect of interfering nodes, use a factor that captures the error probability due to collisions, considering it as a value proportional only to the number of STAs associated to each AP and STAs that belong to neighboring cells and operate on the same channel. The disadvantage of this assumption is that it does not consider APs transmitting on downlink as potential interference. In addition, their approach is not able to distinguish between nodes that just contend for channel usage and nodes that appear hidden.

B. Our Contribution

In this paper, we propose a novel approach that resolves the issues mentioned above not as individual parameters but in a joint manner. We contribute by developing a comprehensive metric, that is based on estimation of end user throughput in 802.11 infrastructure networks. In order to capture the asymmetric nature of the wireless medium, we estimate performance both on uplink and downlink channel.

As a first contribution we encapsulate in our throughput estimation formula the effect of contention. In contradiction to the aforementioned approaches, we state that AP load should be considered over all neighboring nodes, due to the shared nature of the medium. The IEEE 802.11 medium access is performed by the distributed coordination function (DCF), that is based on the CSMA/CA protocol. This medium access control (MAC) protocol provides all compliant nodes with the same chance to access the medium and transmit

frames in the long term. As a result every node in a WLAN shares the medium with its neighboring nodes. Moreover, due to the multi-rate capability at the physical layer (PHY), supported by rate adaptation mechanisms, the transmission duration of a frame depends on the transmission rate selected by the transmitter. As a result, transmitters that use low PHY rates, capture the medium for longer duration, muting their surrounding nodes during their transmissions. The combined effect of shared medium in accordance with the multiple PHY rates used, can cause the well known 802.11 "anomaly phenomenon", where low transmission rate STAs negatively affect high bit rate ones [11]. The result is that all STAs finally get throughput of the same order of magnitude. Consequently, we have to take a step further than the previous approaches and take into account transmissions of all active nodes in *STA's* neighborhood, in accordance with their transmission rates, in order to estimate the levels of contention and extend the definition of load.

Another key contribution of the proposed scheme is its ability to adapt to the varying traffic patterns that each corresponding node follows. Thus it manages to adapt to realistic traffic conditions. We state that activity estimation should be performed by each individual node and this information can be exchanged through neighboring WLANs, by using specially generated for this purpose management frames. Through this approach, all nodes are able to detect transmissions in adjacent cells in a distributed way.

As a third contribution we investigate how simultaneous transmissions of hidden nodes affect user performance. Due to the shared medium, transmissions of interfering hidden nodes can cause collisions and erroneous receptions, that lead to decreased packet delivery ratios (PDR) at the receivers. Counter to relevant approaches, we incorporate in our proposed metric the effects of contending and interfering nodes separately.

Our mechanism, integrating all the above features, results in algorithms proposed both for the association and the handoff procedure. One more important contribution is that we move one step further than simulation and implement the proposed algorithms using open source drivers and also validate their performance in a wireless testbed, to evaluate the performance in real world settings.

II. SYSTEM MODEL AND METRIC DEFINITION

We consider an IEEE 802.11 based WLAN that consists of a large number of APs and STAs. We use \mathcal{M} to denote the set of APs that define a network coverage area. We assume that there is a set of available channels, denoted by \mathcal{C} . Each $AP_j \in \mathcal{M}$ operates on a single predefined channel $c_j \in \mathcal{C}$, where \mathcal{C} denotes the set of non-overlapping channels that the operating band offers. The coverage areas of multiple APs may be overlapping. Within the network coverage area resides a set of mobile STAs, denoted by \mathcal{N} , which tend to stay in the same physical locations for long time periods. At any time instant, a $STA_i \in \mathcal{N}$ chooses to associate with a single $AP_j \in M_i$, where M_i denotes the set of APs that operate in the vicinity of STA_i . We use N_j to denote the set of $STAs$ that are associated with AP_j .

Each node of the network $n \in M \cup N$, has a set of neighbors, that reside in its sensing area and operate on the same channel with n . This set of "1-hop" neighbors, that can be either APs or STAs, is denoted by A_n .

Based on the discussion of the previous section, we notice that the throughput, which is experienced by a node in an IEEE 802.11 network, depends apart from channel quality also on the transmission of frames by other nodes in the network and its selected PHY rate. The authors of [12], based on the well known analysis of Bianchi [13], have shown that when there are multiple transmitters with different PHY rates, that lie in the contention domain of node n , which uses a PHY rate of R_n , each node of the network enjoys an equally shared value of throughput, that is approximated as:

$$T_n = \frac{1}{\frac{1}{R_n} + \sum_{k=1}^{|A_n|} \frac{1}{R_k}} \quad (1)$$

where R_k denotes the PHY rate that each node $k \in A_n$ uses while transmitting. This equation ignores the overhead resulting from the 802.11 MAC mechanism. The deficiency mentioned here, is not important in our analysis, as we use this equation to decide about the association that can lead to the best available throughput performance and not in order to calculate the actual resulting end user throughput. Moreover, this equation considers saturated traffic conditions and requires that all traffic flows consist of equal packet lengths. Throughout this paper, we follow the assumption of equal packet lengths, but later in our analysis we transform the above equation, in order to capture realistic varying traffic conditions. The above expression is based on the estimation of channel usage time that each transmitter node gains access to the medium, given the existence of other 802.11 nodes operating on its channel and transmitting in its vicinity. We modify the above equation, which refers to the general case of a network with multiple flows generated between 802.11 compliant nodes, to fit our needs about the special case of infrastructure 802.11 networks.

A. Contention Effect

We start forming our formula by considering only saturated downlink traffic. Assume the usual case, where an $AP_j \in \mathcal{M}$ has $|N_j|$ associated STAs and serves them with downlink traffic. We also consider that in the vicinity of AP_j , there are $|A_j|$ "1-hop" neighboring nodes operating on the same channel and contending to capture the channel. An AP has an equal probability among its contending nodes to capture the medium for its own transmissions and in each such instant it uses the medium to transmit to only one of its associated STAs. The AP's service rate is equally shared between the associated STAs, if the number of frames destined to each STA is equally distributed among them. Supposing this equal distribution and the upcoming association of STA_i with AP_j , the mean PHY rate that AP_j uses when transmitting downlink traffic to its associated STAs can be approximated as follows:

$$\overline{R_j} = \left(\sum_{m=1}^{|N_j|} R_{jm} + R_{ji} \right) / (|N_j| + 1) \quad (2)$$

where R_{jm} and R_{ji} denote the PHY rate that AP_j uses when transmitting to each $STA_m \in N_j$ and STA_i accordingly. Concluding, in order to estimate the equally shared value of transmitted bits destined to each STA_i while associated to AP_j , we transform expression (1) as:

$$T_{ij}^{down} = \frac{1}{\left(\frac{1}{R_j} + \sum_{k=1}^{|A_j|} \frac{1}{R_k} \right) \cdot (|N_j| + 1)} \quad (3)$$

Thus, we capture the effect of equal sharing of the AP_j 's service rate among its $|N_j|$ associated users, while we assume that AP_j uses a mean value of PHY rate when transmitting to all its associated STAs.

B. Hidden-Node Effect

As previously stated, these transmissions are still subject to frame errors, due to interference at the receivers' side. The factor that plays the key role in interference, is the effect of hidden terminals, which appears very often in dense WLANs. In a simple downlink scenario with one AP_j and one associated STA_i , the set of nodes that appear hidden to transmissions of AP_j , consists of nodes existing in the "1-hop" neighborhood of STA_i , that do not belong in the "1-hop" neighborhood of AP_j . We call this set of nodes as the "2-hop" neighbors of the transmitter, denoted as B_j , equal to the relative complement of set A_j in A_i ($A_i \setminus A_j$). Since the "2-hop" neighbors of the transmitter AP_j do not sense its ongoing transmissions, collisions occur, leading to decreased packet delivery ratio (PDR) and consequently result in throughput decrease. The negative effect of the hidden-node problem is proportional to the number of "2-hop" neighbors ($|B_j|$) and their transmissions duration. This effect can be imported in expression (3) to model the decrease in performance as follows:

$$T_{ij}^{down} = \frac{1}{\left(\frac{1}{R_j} + \sum_{k=1}^{|A_j|} \frac{1}{R_k} \right) \cdot (|N_j| + 1) + \sum_{l=1}^{|B_j|} \frac{1}{R_l}} \quad (4)$$

C. Traffic Intensity Estimation

In most cases, the consideration of saturated traffic is not realistic. Practically, nodes run different applications that generate traffic with varying rates. Assuming that all nodes in a network generate traffic with the same rate can only estimate performance under the worst case scenario. In order to model realistic scenarios we have to characterize each transmitter according to the traffic pattern that it follows. For this purpose, we define an activity indicator, denoted as f_n , for each node $n \in M \cup N$. Each node n measures the rate of packets arriving to its transmission queue during a constant time interval, capturing its arrival rate (λ_n). Moreover it can estimate its affordable service rate (μ_n) using expression (1),

which is approximately the rate at which packets leave its queue for transmission. In the case that the rate of packets arriving to the transmission queue is higher than the rate at which packets leave the queue, only the number of backlogged packets increases, while the traffic injected in the network remains constant.

Based on the above, n estimates its maximum affordable traffic rate, by setting its activity indicator f_n as follows:

$$f_n = \min\{\lambda_n, \mu_n\} \quad (5)$$

Each node of the network announces its f_n to all its neighbors. This way, every node that receives reports about ongoing transmissions in its neighborhood, manages to create a list of all its "1-hop" neighbors and their corresponding activity indicators. Moreover, it has to detect its "2-hop" neighbors. For this purpose, all nodes have to exchange their lists of neighbors. The above activity estimation procedure is performed by each node of the network, either operating as an AP or as a STA. We now use the activity indicators to transform equation (4) as follows:

$$T_{ij}^{down} = \frac{1}{\left(\frac{f_j}{\bar{R}_j} + \sum_{k=1}^{|A_j|} \frac{f_k}{R_k}\right) \cdot (|N_j| + 1) + \sum_{l=1}^{|B_j|} \frac{f_l}{R_l}} \quad (6)$$

where \bar{R}_j is now calculated regarding the percent of traffic destined to each individual $STA_m \in N_j$ and STA_i accordingly.

A similar approach can be used for uplink communication as well. When a STA transmits on uplink, all of its frames are destined to the AP it is associated with. The transmitter, STA_i in this case, shares the medium with its "1-hop" neighbors (A_i), while its "2-hop" neighbors (B_i) are the nodes that are located in the AP's neighborhood but not in the STA's ($A_j \setminus A_i$). For the uplink case, we arrive at the following expression:

$$T_{ij}^{up} = \frac{1}{\frac{f_i}{R_{ij}} + \sum_{k=1}^{|A_i|} \frac{f_k}{R_k} + \sum_{l=1}^{|B_i|} \frac{f_l}{R_l}} \quad (7)$$

III. PROPOSED ALGORITHMS

A. Association Mechanism

The above analysis concludes in two expressions (6) and (7), that estimate throughput performance for uplink and downlink communications accordingly. In our model, we have assumed that APs are statically assigned predefined channels, that do not change during operation time. We do not consider channel allocation in this work, since our focus is on devising throughput-efficient access point association mechanisms. As long as the operating channel remains constant for the APs, they can constantly monitor their "1-hop" and "2-hop" neighborhoods. On the other hand, the STAs are able to change their operating channels by performing handoffs between APs that operate on different channels. This way, the set of detected neighbors depends on the channel the STA operates on. During the scanning period, each STA has to remain on each channel

for duration equal to the Neighbor Reports' interval, in order to collect all the reports transmitted by its neighbors. We denote this time period as t_r . At the end of each t_r , each STA has to store the list of neighbors that it detects on each channel. We use A_i^c and B_i^c , to denote the sets of "1-hop" and "2-hop" neighbors, detected by STA_i on each channel c . This scanning procedure is repeated $|C|$ times, so that each STA can estimate its neighbors on all the available channels.

Generally in wireless communications, downlink connections dominate the overall communication load. However most real-time applications such as VoIP or video conferencing require suitable Quality of Service (Qos), in both the uplink and the downlink. We indicate the uplink-to-total-link ratio as u_r and similarly for downlink as d_r . Each STA can determine its own ratios, concerning the type of application it is running.

By using these factors, STA_i under association, can calculate the combined metric, considering the achievable performance when it is associated with AP_j , as follows:

$$T_{ij}^{total} = u_r \cdot T_{ij}^{up} + d_r \cdot T_{ij}^{down} \quad (8)$$

Having calculated the above metric, considering every $AP_j \in A_i$, STA_i estimates the potential performance both on uplink and downlink for each available association and then decides to associate with the AP, that provides the maximum calculated metric. A brief pseudocode description of the implemented association algorithm is given in Table Algorithm 1.

B. Handoff Mechanism

The above analysis should be extended in the handoff mechanism. A handoff in 802.11 is the process that allows a STA to change the AP that is associated with, because it detects degradation of the communication quality. According to the IEEE 802.11 standard, when a STA moves away from the AP it is associated with, the SNR of the link drops, and if the Cell Search Threshold is reached, the MAC Layer Scan function starts to search for potential APs. The Cell Search Threshold is not explicitly defined in the standard. Implementation of appropriate triggering mechanisms is typically left to the wireless card manufacturer, and is therefore proprietary.

As previously explained, several more factors than the signal strength affect communication quality. The key feature that our mechanism supports is the consistent monitoring of all these factors jointly, by calculating the proposed metrics. This way, each STA can monitor the throughput performance that its current association offers and consequently decide whether a handoff to another available AP is required. In the proposed scheme, the triggering of the scanning procedure is based on a throughput percent threshold denoted by H_1 , instead of RSSI-based thresholds.

Moreover, in our model we introduce a periodical scanning window, during which each STA triggers the scanning procedure, so as to be able to estimate potential performance considering APs that operate on the other available channels. We define as t_m the period of this periodical procedure. In addition, we set a time threshold denoted by H_2 , to determine the validity of the results in our scan cache. This threshold

Algorithm 1 ASSOCIATION MECHANISM

Require: TIME := $|C| * t_r$
Require: INCOMING NEIGHBOR_REPORT OF EACH $k \in A_i^c$
Require: INCOMING BEACON OF EACH $j \in M_i$
Require: TIME OF CHANNEL
Ensure: ASSOCIATION DECISION FOR STA_i

- 1: **while** TIME < $|C| * t_r$ **do**
- 2: **for** $c \in C$ **do**
- 3: WAIT IN RECEIVE MODE FOR t_r
- 4: **for** $k \in A_i^c$ **do**
- 5: COLLECT NEIGHBOR_REPORT OF k
- 6: SAVE NEIGHBOR_LIST OF k, R_k, f_k
- 7: **end for**
- 8: CALCULATE LIST OF A_i^c, B_i^c
- 9: **for** $j \in M_i$ **do**
- 10: **if** $c_j = c$ **then**
- 11: COLLECT BEACON OF AP_j
- 12: ESTIMATE R_{ij} Using $RSSI_j$
- 13: SAVE N_j
- 14: CALCULATE LIST OF A_j, B_j
- 15: CALCULATE $T_{ij}^{down}, T_{ij}^{up}$ Using (6), (7)
- 16: CALCULATE T_{ij}^{total} Using (8)
- 17: **end if**
- 18: **end for**
- 19: **end for**
- 20: **end while**
- 21: **for** $j \in M_i$ **do**
- 22: STA_i SELECTS AP_j That Maximizes (8)
- 23: **end for**
- 24: STA_i ASSOCIATES WITH AP_j

is used to avoid the overhead induced by inefficient scanning procedures that lead to similar results. A brief pseudocode description of the implemented handoff algorithm is given in Table Algorithm 2.

C. Implementation Details

In this section we describe the key challenges encountered in the driver implementation and the corresponding solutions. For the implementation of our mechanism, we used the MAD-WiFi open source driver [14]. Our proposed mechanism assumes that each node is able to receive information about ongoing transmissions from all the BSSs taking place on the channel it is operating on. However, all packets received by the network adapter are filtered out, so that the ones with a destination address different than the local MAC address of the adapter are discarded. Only unicast packets that are destined to the adapter's MAC address, multicast and broadcast packets can be captured. A solution for our needs would be the passive approach of capturing unicast data packets that would contain the required information in specially generated fields in the header of the data packets had to be avoided, due to the large amount of information that has to be exchanged between STAs.

Algorithm 2 HANDOFF MECHANISM

Require: TIME := $k * t_m, k \in \mathbb{N}$
Require: HANDOFF THRESHOLDS H_1, H_2
Require: ASSOCIATED WITH AP_0 WITH T_{i0}^{old}
Require: OPERATION ON CHANNEL $c = c_{old}$
Require: INCOMING NEIGHBOR_REPORT OF EACH $k \in A_i^c$
Require: BEACON OF AP_0
Ensure: HANDOFF DECISION

- 1: **while** TIME < $k * t_m$ **do**
- 2: **for** $k \in A_i^c$ **do**
- 3: SAVE NEIGHBOR_LIST OF k, R_k, f_k
- 4: **end for**
- 5: CALCULATE LIST OF A_i^c, B_i^c
- 6: CALCULATE LIST OF A_0^c, B_0^c
- 7: CALCULATE $T_{i0}^{down}, T_{i0}^{up}$ Using (6), (7)
- 8: CALCULATE T_{i0}^{new} Using (8)
- 9: **if** $(T_{i0}^{new} - T_{i0}^{old})\% > H_1$ **then**
- 10: **if** $Scan_Invalid(H_2)$ **then**
- 11: $Scanning_Procedure()$
- 12: **end if**
- 13: $Association_Decision()$
- 14: **end if**
- 15: **end while**
- 16: **if** $Scan_Invalid(H_2)$ **then**
- 17: $Scanning_Procedure()$
- 18: **end if**
- 19: $Association_Decision()$

Instead of using this scheme, an active information approach has been followed. More specifically, the first modification we made in the driver is the generation of an 802.11 broadcast frame, that is transmitted periodically. This special control packet, called *Neighbor Report*, includes the PHY rate used in the last transmission of each node and its activity indicator, computed using expression (5). We further modified the driver in a way that each node n that receives these broadcast frames, estimates its "1-hop" neighborhood and subsequently creates a list with the MAC addresses of each node $k \in A_n$, their PHY rates and their f_k accordingly. Upon the reception of Report packets by all its "1-hop" neighbors, node n estimates its "2-hop" neighborhood, as described in the previous section.

A third modification of the driver was the extension of the Beacon frames transmitted by the APs, by adding an extra field that contains the number of associated users and its mean PHY transmission rate, as calculated upon observation according to the percent of traffic that is destined to each associated STA. In addition, the PHY rates (R_{ji} and R_{ij}) that will be used by STA_i for its communication with each AP_j are estimated considering the strength of Beacon and Probe-Response Frames, transmitted by the neighboring APs. Our final modification was made in the scanning procedure, where we set the interval that STAs have to remain on each channel equal to the Neighbor Reports' interval (t_r).

IV. EXPERIMENTAL CONFIGURATION

In order to evaluate the performance and study the behavior of the association scheme that we implemented, we used a large scale programmable testbed of wireless nodes, called NITOS. By testing the proposed scheme in a real large scale testbed, we were able to measure the performance under real conditions.

A. NITOS Testbed

NITOS (Network Implementation Testbed for using Open Source platforms) is a wireless testbed, that is designed to achieve reproducibility of experimentation. Users can perform their experiments by reserving slices (nodes, frequency spectrum) of the testbed through NITOS scheduler, that together with OMF [15] management framework, support ease of use for experimentation and code development. It is *remotely* accessible and currently consists of 40 wireless nodes, outdoor located in a non-RF-isolated environment. The nodes are equipped with 2 wireless interfaces using Wistron CM9 - mPCI Atheros 802.11a/b/g 2.4 and 5 GHz cards, that run MAD-WiFi open source driver. NITOS is deployed at the exterior of the University of Thessaly campus building.

B. Measurement Methodology

The throughput performance of the experiments, is measured by using Iperf [16], which is a powerful tool for traffic generation and measurement. A typical experimental setup for experiments considering only downlink transmissions, would be to run an Iperf client at the nodes, that act as *APs*, in order to generate traffic streams, having an Iperf server residing on each STA, receiving the traffic and collecting the measurements. To remove any random effect and short-term fluctuation, we run each experiment 5 times and each run lasts for 10 minutes. In order to get final results we average the results of the five experiments.

V. EXPERIMENTAL EVALUATION

Based upon the testbed described in Section IV, numerous experiments were conducted, and the results obtained are reported and analyzed in this section. The first three experiments, were performed in two discrete phases. In the first phase, we use the unmodified MAD-WiFi driver, which follows the RSSI approach to determine association decisions, while in the second phase we used the modified driver that implements our mechanism. The two first experiments have been designed to evaluate performance for downlink, while the third one for uplink. We did not run experiments measuring combined throughput on purpose, as we expect the performance to be relevant. A representation of the testbed's topology that illustrates the associations following the RSSI approach is depicted in Figure 1. All the figures that illustrate association decisions, use a solid line to represent the association of the first interface, while a distinctive line is used to represent their association of the second interface. In all the conducted experiments, the default Rate adaptation algorithm of the driver has been used.




Fig. 1: RSSI-based Associations

A. Downlink Experiment 1

In this experiment, we set up a network that consists of 3 APs and 14 STAs. Two of the APs operate both on channel 48 of the 5GHz band, while the other one operates on channel 36 of the same band. The APs generate UDP traffic of varying rate, using 14 Iperf clients that run simultaneously, while the corresponding Iperf servers run at each STA. We activate the STAs one by one, introducing a fixed time interval of 5 seconds. When each STA associates with a certain AP, its Iperf server starts receiving data and measuring the actual throughput. For both the RSSI and the proposed scheme, Figure 2(b) illustrates how the average throughput achieved per STA changes with respect to the load applied.

When the first phase of the initial associations ends, each STA starts checking for potential handoffs, enabling the handoff mechanism. The combined result of the association and the handoff mechanism is that the STAs manage to detect the contention for channel usage between the two APs that operate on the same channel, which leads them to associate as it is shown in Figure 2(a). The proposed scheme enables APs to deliver substantial higher throughput, for all the cases studied. More specifically, the RSSI approach reaches the maximum throughput value when the load reaches the 15 Mbps/flow, while our approach reaches the maximum throughput for load of 20 Mbps/flow. Once the traffic rate per flow increases above the value of 30 Mbps/flow, all stations in both the approaches invariably start to witness significant packet drop and throughput deterioration. However, our scheme continues to achieve higher performance than the standard approach.

Another significant issue, that has to be considered, is the fairness feature that our scheme provides. As portrayed in Figure 2(c), the proposed approach manages to provide nearly equal sharing of available throughput to all the STAs of the network, for all the cases studied. On the other hand, the standard approach, leads to associations that favor only a subset of the STAs with high performance, while letting the rest of the STAs with low throughput values.

Concluding, we see that our mechanism manages to balance the traffic load of the network not only among the available APs, but moreover among the available channels that the APs operate on. This characteristic of our scheme is the one that leads to equal sharing of the available throughput between the corresponding receivers.

B. Downlink Experiment 2

In our second experiment, we set up the same network as the one in the first experiment, but we change the number of the STAs to 10. Moreover, this scenario differentiates relative to the first one, as an extra flow of varying traffic rate is added to






Fig. 2: Downlink Experiment 1






Fig. 3: Downlink Experiment 2






Fig. 4: Uplink Experiment

the network, by a pair of nodes that belongs to another adjacent cell. This extra flow is activated between two nodes operating on the same channel with the two of the APs, remains active during the whole experiment and it is used as a contention flow for these two APs. We follow the same activation procedure as in the previous experiment. The traffic rate for each Iperf client at the APs, is set constant to 10 Mbps/STA, so that the influence of the contending pair of nodes can be clearly depicted.

Figure 3(a) depicts the resulting associations. In this scenario most of the STAs decide to associate with AP4 that operates on channel 36 and it does not contend with the other two APs that operate on the same channel. Figure 3(b) illustrates how the average throughput achieved per STA changes with respect to the traffic load injected by the contending flow. As clearly shown in the figure, our approach significantly outperforms the RSSI approach for all the different values of traffic rate that the transmitter of the contending pair uses.

The fairness that our mechanism achieves, is depicted in Figure 3(c). In this Figure, we notice that nearly equal bandwidth sharing is achieved for our scheme, till the value of 5 Mbps traffic rate of the contending flow. Although the fairness index values of our scheme decrease after the rate of 10 Mbps, its superiority in terms of performance and fairness is still maintained.

C. Uplink Experiment

The network setup for this experiment consists of 2 APs and 8 STAs. One of the APs operates on channel 48 of the 5GHz band, while the other one operates on channel 36 of the same band. For the uplink case, the STAs generate UDP traffic of varying rate, each one using an Iperf client, while the corresponding Iperf servers run at each AP, receiving traffic generated by all the associated STAs. We follow the same activation procedure here as in the two previous scenarios. When a STA associates with a certain AP, its Iperf client starts transmitting data and the actual throughput is measured by the AP that is associated with. The reason for designing a smaller topology in this scenario is that our intention has been to investigate how the proposed scheme scales in the case of multiple high load traffic flows, which would not be sustained in a network with a lot of flows. The average throughput per STA measured by the APs with respect to the traffic rate per flow, is illustrated in Figure 4(b). Both our association and handoff mechanisms are activated in this scenario. The resulting associations are represented in Figure 4(a).

Figure 4(b), clearly shows that the proposed scheme enables STAs to deliver higher throughput, in the cases studied. In the cases of 2 and 3 Mbps/flow there is no significant difference in the average throughput/STA. However, great increase in throughput performance is seen at the rate of 5 Mbps/flow.

In the fairness measurements, demonstrated in Figure 4(c), we can see that our approach maintains high fairness index values even in high load per flow, while the RSSI approach faces a decrease that starts when the traffic rate per flow reaches the value of 5 Mbps.

VI. CONCLUSIONS AND FUTURE WORK

In this work, we have proposed a novel association scheme that is based on innovative metrics, capturing the effects of contention and interference in the neighborhood both on uplink and downlink. A key feature of the proposed scheme is that it manages to adapt to realistic traffic conditions. Compared with the standard RSSI-based approach, our mechanism exhibits a far better performance in terms of throughput. Another important effect of our algorithms is that they manage to provide nearly equal sharing of throughput among the intended receivers, even in high load conditions.

Possible extensions of this work include the study of an altruistic extension, where each node will be considering the overall performance of the network as well. As a second extension, we intend to study the problem of user association and frequency selection jointly, as proposed in [17]. Furthermore, a future objective is to test new mechanisms that jointly perform power and association control, in order to see how the difference in transmission range can affect the "1-hop" and "2-hop" neighborhoods, which in turn affect the corresponding associations.

VII. ACKNOWLEDGMENTS

The authors acknowledge the support of the European Commission through STREP project OPNEX (FP7-ICT-224218).

REFERENCES

- [1] IEEE 802.11-2007 Wireless LAN Medium Access Control and Physical Layers Specifications.
- [2] Y. Bejerano, S. Han, and E. L. Li. Fairness and load balancing in wireless lans using association control. In *Proceedings of MOBICOM*, pages 315–329. ACM, 2004.
- [3] Y. Bejerano and Randeep S. B. Mifi: a framework for fairness and qos assurance for current ieee 802.11 networks with multiple access points. In *IEEE/ACM Trans. Netw.*, 14(4):849–862, 2006.
- [4] T. Korakis, O. Ercetin, S. Krishnamurthy, L. Tassiulas, and S. Tripathi. Link quality based association mechanism in ieee 802.11h compliant wireless lans. In *Proceedings of RAWNET*, 2006.
- [5] B. Kauffmann, F. Baccelli, A. Chaintreau, V. Mhatre, K. Papagiannaki, and C. Diot. Measurement-based self organization of interfering 802.11 wireless access networks. In *Proceedings of INFOCOM*, pages 1451–1459. IEEE, 2007.
- [6] O. Ekici and A. Yongacoglu. A novel association algorithm for congestion relief in ieee 802.11 wlans. In *Proceedings of IWCMC*, pages 725–730, New York, NY, USA, 2006. ACM.
- [7] G. Athanasiou, T. Korakis, O. Ercetin, and L. Tassiulas. Dynamic cross-layer association in 802.11-based mesh networks. In *Proceedings of INFOCOM*, pages 2090–2098. IEEE, 2007.
- [8] K. Sundaresan and K. Papagiannaki. The need for cross-layer information in access point selection algorithms. In *Proceedings of SIGCOMM*, pages 257–262, New York, NY, USA, 2006. ACM.
- [9] M. Abusubaih and A. Wolisz. An optimal station association policy for multi-rate ieee 802.11 wireless lans. In *Proceedings of MSWiM*, pages 117–123, New York, NY, USA, 2007. ACM.
- [10] M. Abusubaih and A. Wolisz. Interference-aware decentralized access point selection policy for multi-rate ieee 802.11 wireless lans. In *Proceedings of PIMRC*, pages 1–6. IEEE, 2008.
- [11] M. Heusse, F. Rousseau, G. Berger-Sabbatel, and A. Duda. Performance anomaly of 802.11b. In *Proceedings of INFOCOM*, volume 2, pages 836–843, San Francisco, USA, March 30–April 3, 2003.
- [12] 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/ACM Trans. Netw.*, 15(3):588–601, 2007.
- [13] G. Bianchi. Performance analysis of the ieee 802.11 distributed coordination function. In *IEEE JSAC*, 18(3):535–547, March 2000.
- [14] "Madwifi Open Source Driver", <http://madwifi-project.org/>.
- [15] "OMF-control & Management Framework", <http://omf.mytestbed.net/>.
- [16] "Iperf", <http://dast.nlanr.net/Projects/Iperf/>.
- [17] I. Koutsopoulos and L. Tassiulas. Joint optimal access point selection and channel assignment in wireless networks. In *IEEE/ACM Trans. Netw.*, 15(3):521–532, 2007.