

The Need for Cross-Layer Information in Access Point Selection Algorithms

Karthikeyan Sundaresan
Georgia Tech
sk@ece.gatech.edu

Konstantina Papagiannaki
Intel Research
dina.papagiannaki@intel.com

Abstract The low price of commodity wireless LAN cards and access points (APs) has resulted in the rich proliferation of high density WLANs in enterprise, academic environments, and public spaces. In such environments wireless clients have a variety of affiliation options that ultimately determine the quality of service they receive from the network. The state of the art mechanism behind such a decision typically relies on received signal strength, associating clients to that access point (AP) in their neighborhood that features the strongest signal. More intelligent algorithms have been further proposed in the literature. In this work we take a step back and look into the fundamental metrics that determine end user throughput in 802.11 wireless networks. We identify three such metrics pertaining to wireless channel quality, AP capacity in the presence of interference, and client contention. We modify the low level software functionality (firmware and microcode) of a commercial wireless adaptor to measure the necessary quantities. We then test, in a real testbed, the ability of each metric to capture end user throughput through a range of diverse network conditions. Our experimental results indicate that user affiliation decisions should be based on metrics that do not only reflect physical layer performance, or network occupancy, but also concretely capture MAC layer behavior. Based on the acquired insight, we propose a new metric that is shown to be highly accurate across all tested network scenarios.

Categories and Subject Descriptors

C.2.1 [Computer-Communication Networks]: Network Architecture and Design - Wireless Communication

General Terms

Algorithms, Management, Measurement, Experimentation, Performance

Keywords

IEEE 802.11, Access Point Selection, Cross-Layer

1 Introduction/Motivation

IEEE 802.11 has become the de facto protocol for wireless access in urban areas capitalizing on the large deployment of 802.11 Access

Points (APs). Within such dense deployments, a wireless client has a variety of choice in its association with the wired infrastructure. The state of the art mechanism, implemented in the majority of 802.11 wireless adaptors, relies on measurements of received signal strength (RSSI); the client associates with that AP that is heard at the highest signal strength. The reason driving such a decision stems from the fact that wireless adaptors employ rate adaptation, tuning their transmission rate in response to the quality of the wireless link they experience to their AP. If the link quality is poor, then the client needs to employ more robust modulation and coding schemes, thus reducing its effective transmission rate. Affiliating with an AP featuring a high signal strength implies that the client can communicate with the AP at higher transmission rates.

Such an affiliation algorithm has received significant criticism due to its ignorance of AP load. Sole consideration of link quality in the AP affiliation process can lead to the overload of APs with high client concentration, while other APs remain unused due to their slightly longer distance from the majority of the clients. As a consequence, new algorithms were proposed that incorporate AP load in the selection process [4, 8]. Some of these algorithms rely on passive measurements collected from Beacon frames, while a recent approach advocates the use of active measurements for the identification of the “best” AP [6].

In this work we are taking a step back from previous work and look at the fundamental metrics that should drive the AP selection process in order to accurately reflect potential user throughput. Moreover, we focus on passive measurements that can guide such a process without requiring pre-existing authentication with the APs under consideration. We identify the following differentiating aspects in AP selection: (i) the *AP capacity*, that captures the capacity of an AP in the presence of interfering devices (802.11-enabled or not), (ii) the *duty cycle* of the AP, that captures the average amount of time the AP spends to serve all its users once, and (iii) the *quality* of the link between the AP and the new client, which determines the client’s instantaneous transmission rate. We propose metrics to capture the aforementioned dimensions in AP selection. Measurement of the first two metrics needs to capture the state of the MAC protocol, and is not supported by commercial cards. We believe that the need for changes in the low layer functionality of the wireless adaptors for the exposition of these two metrics has been a limiting factor in their study through measurement. Our ability to modify the microcode and firmware of the Intel 2915 ABG card puts us in a unique position to study their performance in reflecting user throughput. Using a small scale testbed we explore alternative network scenarios and the tradeoffs that different metrics may face across a diverse set of networking environments. We clearly show that the use of any one metric in isolation is not sufficient to lead to optimal decisions across all environments. For such a task we propose a new metric that is capable of capturing the cross-layer

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

IMC’06, October 25–27, 2006, Rio de Janeiro, Brazil.

Copyright 2006 ACM 1-59593-561-4/06/0010 ...\$5.00.

behavior that should drive AP affiliation decisions.

The rest of the paper is structured as follows. In Section 2 we present a model for the long-term 802.11 user throughput under fully saturated traffic conditions. Using our insight from Section 2 we discuss alternative metrics that could guide the AP selection process in Section 3 and propose a way to measure them using existing hardware. In Section 4 we present our experimental methodology for the assessment of the accuracy of each individual metric across a diverse set of scenarios. Our results highlight the need for the definition of a cross-layer metric which is presented in Section 5. We summarize in Section 6.

2 User throughput in 802.11 networks

For analytical tractability in the derivation of the metrics affecting user throughput in 802.11 networks we focus on the following scenario. We are looking into fully saturated wireless networks, where traffic is primarily flowing from the AP to the users and each user always waits for a packet. In this case, MAC layer modeling becomes easier since only APs are the senders in the network and the amount of interference caused by APs does not depend on the number of their users. We further assume that the wireless network is the bottleneck. Under these assumptions we know that all users will achieve the same *long-term* throughput, shown experimentally in [2]. Future work on how to relax such assumptions is discussed in Section 5.

Two users associated with the same AP do not in general receive signal with the same path loss because of different distances to the AP and of varying channel conditions. The rate adaptation mechanism of 802.11 adapts the encoding rate of the transmitter to the channel conditions; users with poorer quality links use a lower, more robust, encoding rate, thus occupying the medium for longer periods of time. A new user can estimate its instantaneous transmission rate based on his/her RSSI using measurement-based formulas such as the ones listed in [5, 7]. If we assume that information is transmitted to each user in data units of the same length S , the data unit transmission delay of user u is given by $d(u) = \frac{1}{f(\text{SINR}(u))}$, where $f(\text{SINR}(u))$ gives the *instantaneous* transmission rate on the channel from AP a to u , that is expressed in data units per second.

If AP a has other APs in its contention domain, its medium utilization $M(a)$ will not be 100% and its actual capacity will only be a fraction of the medium capacity. In such a setting the long-term throughput obtained by each user u associated with a in a reference measurement period T is given by:

$$r(a, u) = \frac{M(a) * C(a) * T}{\sum_{v \in \mathcal{U}_a} d(v)}, \quad (1)$$

where \mathcal{U}_a corresponds to the set of users associated to AP a , and $C(a)$ denotes the capacity of the AP a in data units per second (different for 802.11a/g and 802.11b networks).

Note that despite the fact that the time to transmit the same unit of information is different from one user to another in the same cell, all the users receive the same long-term throughput; in other words, each user will have received the same number of data units, in a reference period T [2]. The denominator of $r(a, u)$ (which is identical for all users associated with the same AP) will be referred as the *aggregated transmission delay* (ATD) of the AP in what follows.

From the above discussion, we find that the actual long-term throughput of a user in a 802.11 network depends on three factors: (i) the effective capacity of the AP; that is the maximum amount of traffic the AP can serve under the best conditions in the presence of other also saturated, interfering APs (ii) the quality of the wireless link from the AP to the user, which determines the amount of time

a data unit transmitted to this user occupies the medium, and (iii) the average amount of time a user will need to wait to gain access to the medium. The last metric further depends on: (1) the number of users in the cell, and (2) the quality of the wireless link from the AP to each one of the existing users. Information on these two metrics allows for the computation of the *aggregated transmission delay* which consequently impacts the frequency with which the new user will receive a packet.

3 Implementation

While some of the aforementioned metrics can be obtained from driver-level statistics, some others are low-level metrics and need to be obtained right from the firmware. From the previous section we have identified the following metrics as critical factors in the determination of 802.11 user throughput.

Received Signal Strength Indicator (RSSI): Received signal strength is measured by the firmware upon each successful packet reception. This information is further propagated to the driver, where it is typically used for user affiliation decisions.

Aggregated Transmission Delay (ATD): The ATD metric was introduced in the previous section to capture the average amount of time an AP needs to serve one unit of information to each one of its users (assumes fully saturated, downlink traffic). This kind of information is not readily available at the driver level and needs to be obtained from the firmware. We modify the firmware to collect information on the amount of time needed to serve each individual client of the AP. More precisely, we measure the time elapsed between the queuing of a packet at the MAC layer until we receive a MAC-level ACK for each packet transmission to each individual client. Such a measure incorporates the effects of rate scaling as well as any possible retransmissions. The MAC address of the client along with the client's transmission delay are sent to the driver. The driver then computes the average transmission delay to serve one "round" of users, and updates ATD using a weighted average filter¹.

AP Capacity (APC): The AP capacity comprises two different metrics: (i) its nominal capacity, and (ii) the fraction of time the AP gains access to the medium given the existence of other APs (or even non-802.11 devices) operating on its frequency (or overlapping ones). The former metric can be computed upon inspection of the supported physical layer, conveyed in the Beacon frames (802.11a/g corresponding to approximately 30 Mbps Layer 3 capacity and 802.11b to 5 Mbps). The fraction of time that the AP gains access to the medium however requires access to the firmware. Every AP measures the number of slots it spends in the (i) transmission/reception, (ii) backoff and (iii) idle states. Our measurement period is defined such that it encounters five transmission/reception events. Longer durations were found unsuitable due to the wrapping around of the counters. At the start of the measurement period, the three counters are initialized and at the end of the measurement period they are read and reset. The read values are passed up to the driver where the channel utilization fraction is estimated (busy slots/total slots) and maintained as a weighted moving average.

The metrics requiring firmware-support are AP centric, i.e. ATD, and AP capacity. In order for clients to make use of them in their affiliation decisions, APs need to propagate them to the clients. For their propagation to the clients we use the Beacon frames. We introduce additional elements in the Beacon template that carry values for ATD, and APC. Upon reception of the Beacon frames, our modified clients can decode the additional Beacon fields and base their affiliation decisions on the additional information.

¹In our implementation the most recent measurement is weighed by 0.8.

4 Experimental Methodology

We deploy three APs and four clients in an office environment. Our experimental methodology is rooted in the identification of different experimental topologies that can expose the weaknesses of the three metrics when used in isolation. These same topologies will be tested later using our proposed metric in order to demonstrate its potential to deal with cases that may be problematic when using simple metrics. To achieve this result we select the locations and traffic loads for the APs and the clients in a way that can stress test particular dimensions in the problem of AP selection. We then introduce a new client in the environment and collect measurements for all three metrics listed above with respect to two candidate APs (the two choices at the disposal of the client). The client is then instructed to affiliate with each one of the two candidate APs. Upon each association a 4 MBps CBR traffic stream is sent from the chosen AP to the client lasting for 60 seconds and with a packet size equal to 1500 Bytes. We record the resulting throughput, identifying the AP that offered the best throughput performance. We then compare with the AP selections that would have been made using each individual metric, if the client selected (i) the AP with the highest RSSI, (ii) the lowest ATD, or (iii) the highest capacity.

Most of the experiments involve downlink traffic since it appears to be the dominant mode of wireless usage (most servers still reside in the wired infrastructure). We have, however, performed uplink experiments and the trends and observations are no different from the downlink ones. All experiments are performed in the evening when no other wireless activity takes place in the building.

One limitation in our testbed stems from the use of prototype APs that host the modifiable wireless cards. In particular, our wireless adaptors do not implement per-client rate adaptation, a function supported by all commercial APs. Throughout the experiments we use a fixed rate of 11 Mbps, which results in an effective maximum throughput of 5 Mbps. Due to our inability to use rate adaptation on our AP the effective rate supported by a link can no longer adapt to the quality of the wireless channel. The absence of rate adaptation in the downlink experiments does not impact the AP capacity metric but will influence to some extent the representativeness of RSSI and ATD metrics to reflect user throughput, as will be extensively discussed in Section 5. For this reason, our first three experimental topologies attempt to restrict the quality of the communication channel to regions where rate adaptation is not required. We further saturate the channel such that good quality links can be used up to their maximum rate. The effect of rate adaptation is tested on a separate experiment which focuses on uplink traffic, capitalizing on the ability of the clients to perform such a function.

4.1 Experiment 1 - RSSI

Summary of setup: One AP is closer to the new client but has an associated client along with co-channel interference. The other AP has no associated clients nor interference but is farther.

Our first experiment is designed to test the RSSI metric. We deploy 3 APs in the office environment. Two of the APs (AP1, and AP2 in Figure 1(a)) operate on the same frequency and have one client each to which they send traffic at a constant rate equal to 4 MBps. The third AP, AP3, is operating on an orthogonal frequency and has no associated client. The path loss from the new client (C2) to AP1 is smaller than the path loss to AP3. If the affiliation algorithm operated on signal strength alone, the client would select AP1 that features both a client, as well as a busy co-channel AP. We instruct client C2 to affiliate with AP1 and AP3 in turn and measure the throughput achieved by a 4 MBps CBR stream sent from the AP to the client. Table 1 lists the value for the different candidate met-

rics observed by the client with respect to the two candidate APs, as well as the resulting throughput. The throughput obtained from the affiliation with AP3 is 25% higher than the one achieved through AP1, even though the signal strength received from AP1 is higher than the one from AP3. Table 2 lists the decision that the other 2 metrics would have led to if used in the affiliation decision. It can be seen that the impact of the co-channel AP along with the existing traffic on the co-channel APs is much greater than the impact of lower RSSI with respect to AP3.

4.2 Experiment 2 - Aggregated Transmission Delay

Summary of setup: One AP has no clients but faces co-channel interference and is farther away from the new client. The other AP is closer but has an associated client receiving moderate traffic.

The second topology is presented in Figure 1(b). The focus of this topology is to stress the performance of the aggregated transmission delay metric. ATD is affected by the number of clients, as well as their transmission delays. As before, the client has two affiliation choices, AP1 or AP3. AP1 operates on channel 10, featuring co-channel AP2 with client C2 constantly receiving traffic at 4 Mbps rate. On the other hand, AP3 operates on an orthogonal frequency and sends traffic to its associated client at 1 Mbps rate. The quality of the link from C3 to AP1 is worse than towards AP3. The affiliation of client C3 to each AP in turn leads to the results reported in Table 1. The throughput achieved by affiliation with AP3 is six times higher than with AP1. The aggregation transmission delay for AP1 is zero, since AP1 has no clients. However, the co-channel AP with the high workload, along with the worse link quality leads to very poor throughput.

4.3 Experiment 3 - AP Capacity

Summary of setup: One AP is closer to the new client and has no associated clients but faces co-channel interference from a busy AP. The other AP faces no interference but is farther away and has an associated client.

The third topology tests the performance of the APC metric and is shown in Figure 1(c). AP1 and AP2 are operating on the same channel, with AP2 sending traffic to C3 at a rate of 2 Mbps. AP3 operates on an orthogonal channel and constantly sends traffic to C4 at a 4 Mbps rate. The signal quality to AP1 is better than to AP4. The AP capacity reported by AP1 is smaller than the one of AP3 due to the co-channel AP. If the client decides to affiliate with the AP that occupies the medium for the greatest amount of time it obtains half the throughput than if it were to affiliate with AP1 (Table 1). It can be seen that the impact of lower RSSI and contending clients (associated with the same AP) is greater than the impact of lesser capacity by AP1. Note, that the capacity is not significantly low since AP2 sends a light traffic stream to C3 and is also far from AP1.

4.4 Experiment 4 - RSSI (Uplink)

Summary of setup: One AP is closer to the new client but has an associated client sending traffic in the uplink direction. The other AP has no associated clients but is farther away.

All experiments so far have focused on the downlink traffic scenario. The topology in Figure 1(d) focuses on the performance of the RSSI metric on the uplink direction in the presence of rate adaptation, since the clients have such a capability. For a crisper exposition of the impact of rate scaling, this is the only experiment where we use 802.11g due to the greater range of available transmission rates. Client C2 has two affiliation options: (i) AP1 that receives

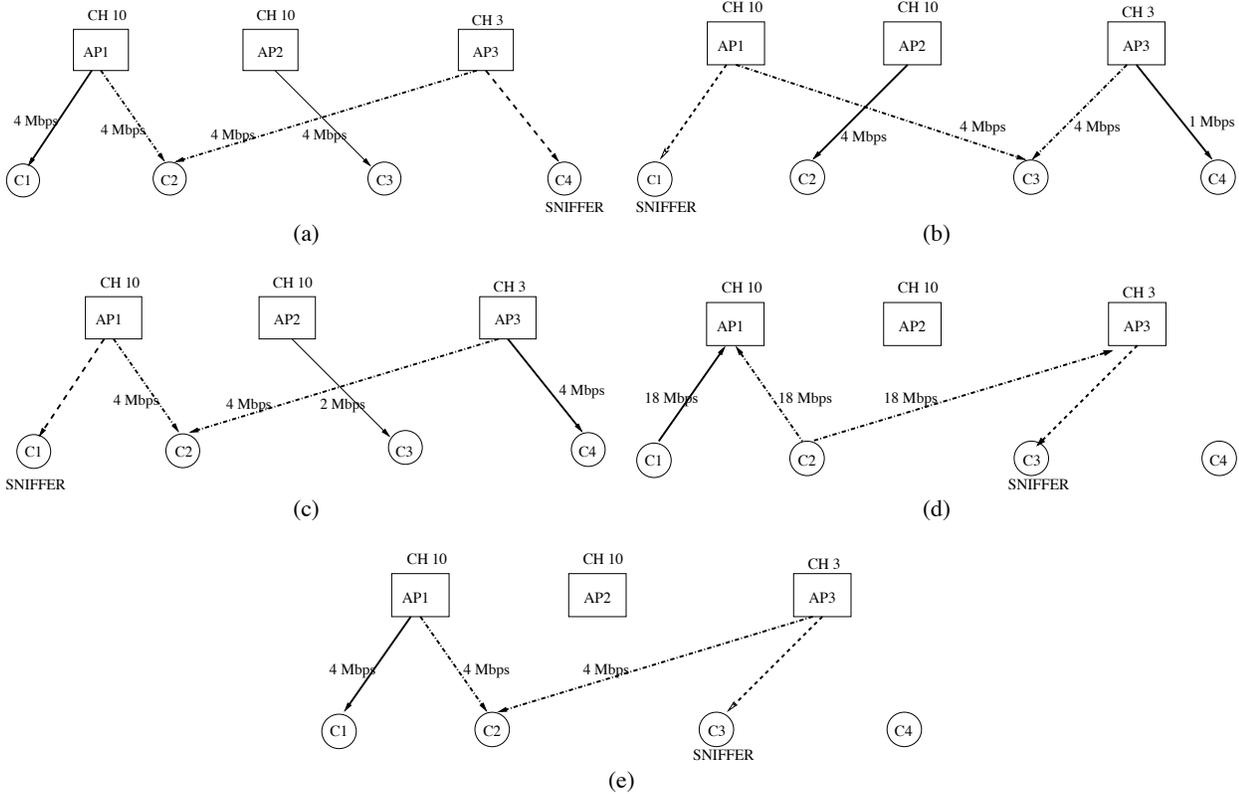


Figure 1: Experimental topologies: (a) RSSI (802.11b), (b) Aggregated Transmission Delay (802.11b), (c) AP Capacity (802.11b), (d) RSSI - Uplink (802.11g), and (e) Impact of Absence of Rate Adaptation on Downlink (802.11b).

traffic from client C1 at 18 Mbps rate, and (ii) AP3 that is farther away but has no clients. Affiliation of C2 with the closest AP1 results in 33% loss in throughput due to the fact that AP1 already supports one client (Table 1). Affiliating with the slightly worse quality AP3 can actually lead to better performance in the presence of rate adaptation. While we had observed this on the downlink, we see that the impact of the wrong decision is greater on the uplink since the clients are capable of rate adaptation which reduces the impact of poor channel quality to a great extent.

5 Toward a more inclusive metric: Expected Throughput

It can be clearly seen from the above experiments, that none of the metrics considered so far is able to help make the right association decision under all circumstances. This is inherently due to the fact that a single metric cannot capture all factors that influence throughput performance. According to Section 2 we can estimate the long-term throughput of a client based on all three metrics. We call the new metric “expected throughput” (ET). The ET metric for a client u with respect to an access point a according to Equation 1 is given by:

$$ET(a, u) = \frac{APC(a)}{ATD(a) + d(u)} \quad (2)$$

where $APC(a)$ is the capacity of a , ATD is the aggregated transmission delay as reported by the AP, and $d(u)$ is the estimated transmission delay for a data unit by client u as estimated using his/her RSSI. Notice that Equation 2 does not include T due to the dependence of our measurement reference period on the trans-

mission/reception events for each AP (Section 3). However, the weighted moving average nature of APC and ATD should still offer a common basis for comparison across APs. Notice that this limitation stems from the sensitivity of the firmware code to sustain a continuous stream of measurements. Upon the calculation of the ET metric for different APs, the client will decide to affiliate with the AP that features the highest ET value.

The corresponding results for the ET metric are shown in Tables 1 and 2. It can be clearly seen that the ET metric is able to make the right decision in all 4 experiments, thereby confirming its effectiveness. However, it must be noted that the ET metric presented above relies on two key assumptions, the use of rate adaptation and saturated downlink traffic, in the absence of which modifications will be required.

5.1 The impact of rate adaptation

The ET metric assumes that the transmission rate and hence the aggregated transmission delay is a function of the RSSI, which in turn assumes the use of some form of rate adaptation. While the ET metric will certainly be a good throughput indicator in uplink experiments where rate scaling is implemented, we found that the ET metric was also a good indicator in all the downlink experiments conducted so far. This is because as long as the difference in RSSI between the two APs being considered is small, the other metrics play a dominant role in the choice of the best AP and hence there is no significant impact on the decision reached by the ET metric. However, when the RSSI difference is large between the two APs being considered, the absence of rate scaling could potentially force the ET metric to reach the wrong decision. We highlight this scenario through our final experiment.

Topologies	APC (Mbps)	ATD (ms)	RSSI	ET (Mbps/T)	Tput (Mbps)
Topo1-AP1	4.9	1.92	-46	1.62	2
Topo1-AP3	5	0	-55	3.3	2.5
Topo2-AP1	1	0	-57	0.12	0.48
Topo2-AP3	5	2.25	-50	1.4	2.86
Topo3-AP1	4.6	0	-49	3.61	2.5
Topo3-AP3	5	1.92	-57	1.37	1.23
Topo4-AP1	26	3.95	-50	6.9	10.5
Topo4-AP3	26	0	-56	16.4	15.7
Topo5-AP1	5	2.78	-40	1.35	1.83
Topo5-AP3	5	0	-68	1.89	1.33

Table 1: Measurements collected by the client in the different topologies with respect to the two candidate APs. Measurements are reported for AP Capacity (APC), Aggregated Transmission Delay (ATD), Received Signal Strength Indicator (RSSI), and Expected Throughput (ET). The last column lists the throughput achieved through the affiliation with each candidate AP. Discrepancies between the last two columns are due to non-saturating loads. Under fully saturated conditions the estimates are accurate.

Topologies	Gain (%)	Best AP	APC	ATD	RSSI	ET
Topo1	25	AP3	AP3 ✓	AP3 ✓	AP1 x	AP3 ✓
Topo2	495	AP3	AP3 ✓	AP1 x	AP3 ✓	AP3 ✓
Topo3	103	AP1	AP3 x	AP1 ✓	AP1 ✓	AP1 ✓
Topo4	50	AP3	AP1 x	AP3 ✓	AP1 x	AP3 ✓
Topo5	38	AP1	AP1 ✓	AP3 x	AP1 ✓	AP3 x

Table 2: AP selected according to the four criteria. Gain reports the relative improvement offered by the “best” AP.

The topology in Figure 1(e) is the same as in the fourth experiment with a difference in that all traffic is now being sent in the downlink direction and hence limited to 4 Mbps. Further, since the APs do not perform rate scaling, the importance of RSSI is increased in this case. Once again C2 connects to AP1 and receives a CBR traffic stream of 4 Mbps for one minute, after which it connects to AP3 and receives a traffic stream of 4 Mbps for another minute. The throughputs are noted for each case. The values for the association metrics for the two APs are further listed in Table 1. The corresponding AP choice made using the different metrics is listed in Table 2. It can be seen that the RSSI metric leads to the right decision, while the ET metric fails. The reason is as follows. First, the AP capacity of AP1 equals that of AP3. Hence, the client choice is determined by the values of ATD and RSSI. Further, there is a large difference in RSSI between the two APs (28 dBm), consequently favoring AP1 with a higher throughput. Since we are considering downlink, APs do not perform rate adaptation. Hence, the transmission delays do not take into account the impact of lower RSSI. For example, at lower RSSI, a rate scaling mechanism would reduce the transmission rate from the AP to the client, thereby increasing the transmission delay. Similarly, the aggregated transmission delay would decrease in the presence of very good RSSI due to the high transmission rates used. Since such a feature is missing, the ET metric is unable to make the right decision. However, if we were to assume that rate scaling is supported by the APs, AP1 would use a high rate of 11 Mbps for a signal strength of 40 dBm or less. This means an increase of rate by more than a factor of 10 (from 11Mbps down to 1 Mbps) or accordingly a decrease of transmission delay by a factor of 10. When this factor is incorporated, the ET metric for AP1 turns out to be smaller than that of AP3, thereby leading to the right decision.

All the experiments conducted thus far corroborate that the ET metric is in fact a very good indicator of client throughput as long as rate adaptation based on RSSI is employed and traffic is saturated.

5.2 Limitations: saturated traffic

The ET metric in Equation 2 could be considered to capture user throughput under the worst case scenario, i.e. the AP is fully saturated, all users require service at all times, and they use a common data unit size. In other words, by design the ET metric is meant to capture the minimum throughput the user should expect from an AP. If other users use smaller packet sizes, their transmission delay is going to be smaller. Moreover, if not all users require service at all times then the user can gain access to the medium more often than what we expect through our model.

In future work we would like to look at the potential of extending the definition of the ET metric to capture scenarios that may not fully abide to the aforementioned assumptions. Modeling user throughput in non-saturated 802.11 networks, that feature both uplink and downlink traffic, has only received attention recently [3] due to the inherent difficulty in dealing with a purely randomized access scheme such as CSMA/CA. In future work we intend to explore the potential of an analytical model like the one presented in [3] to drive the design of a measurement scheme that can operate under non-saturated conditions.

Note that if the total amount of traffic served by the AP still exceeds the AP capacity, then the ET estimation error is likely to be small in magnitude. However, if the AP itself is under-utilized then there is little understanding about how the individual client throughput can be modeled, how long the client workload should be measured and what the ultimate optimization criterion should be. If the individual client workload changes through time, should the user select the AP that maximizes his/her throughput at the time of the selection or should he/she rely on some kind of historical measurements?

Similarly, non-saturated traffic can also result in the aggregate utilization of all co-channel APs to be less than the channel capacity. In such a case, the utilization term should not only include the fraction of time that is currently being used by the AP, but must also include the additional fraction of channel capacity that is not being used by any of the co-channel APs. This could potentially be com-

puted by an AP using information on its own utilization, as well as utilization information conveyed in the beacon frames of other co-channel APs. These are some directions we intend to explore in future work.

6 Summary

Despite the fundamental need for accurate mechanisms for the association of 802.11 clients with APs, there has been little work that addresses the fundamental parameters that should drive such a process. In [4] the authors discuss the benefits of intelligent AP selection algorithms, but do not make a solid proposal as to how different metrics can be used in such a task. In this work we take a step back and look at the metrics that need to be incorporated in such a decision. Using an actual implementation we study the tradeoffs offered by different metrics. We then propose the metric of “expected throughput” that combines information from the physical and MAC layer to assist clients in their association decisions. Our metric relies on the accurate measurement of (i) AP capacity in the presence of interference, (ii) the aggregated transmission delay of all existing clients, and (iii) the instantaneous transmission rate of the new client. We modified the microcode and firmware of the Intel 2915ABG card to study their impact in an experimental testbed. Our results show great promise.

Our approach relies on the advertisement of new metrics in Beacon frames. Such a change has been already proposed within the IEEE 802.15 and 802.11e working groups in the standardization of the Quality Basic Service Set (QBSS) load metric [1]. The QBSS load metric comprises 6 elements, capturing among others (i) the channel utilization (traffic served/capacity), and (ii) the portion of total time available to a QBSS for non-silent periods under the BSS overlap mitigation procedure. Nonetheless, no recommendation is made as to how such metrics can be combined in the selection criterion. We feel that consideration of such metrics by IEEE carries promise in that future generation 802.11 devices will have the proposed measurement support.

7 REFERENCES

- [1] J. Allen. Joing proposal for 802.11e qos enhancements, February 2003. IEEE P802.15 Working Group for Wireless Personal Area Networks (WPANs).
- [2] S. Choi, K. Park, and C. Kim. On the performance characteristics of WLANs: revisited. In *ACM Sigmetrics*, 2005.
- [3] P. Clifford, K. Duffy, J. Foy, D. Leith, and D. Malone. Modeling 802.11e for data traffic parameter design. In *WiOpt*, April 2006.
- [4] D. Larson, R. Murty, and E. Qi. An Adaptive Approach to Wireless Network Performance Optimization. *Technology@Intel Magazine*, feb/mar 2004.
- [5] V. Mhatre and K. Papagiannaki. Using smart triggers for improved user performance in 802.11 wireless networks. In *ACM Mobisys*, June 2006.
- [6] A. Nicholson, Y. Chawathe, M. Chen, B. Noble, and D. Wetherall. Improved access point selection. In *ACM Mobisys*, June 2006.
- [7] Cisco Systems. Deployment guide: Cisco aironet 1000 series lightweight access points.
- [8] S. Vasudevan, K. Papagiannaki, C. Diot, J. Kurose, and D. Towsley. Facilitating Access Point Selection in IEEE 802.11 Wireless Networks. In *ACM Sigcomm IMC*, Berkeley, October 2005.