

Facilitating Access Point Selection in IEEE 802.11 Wireless Networks

S. Vasudevan[†], K. Papagiannaki[‡], C. Diot[‡], J. Kurose[†] and D. Towsley[†]

[†]*Dept. of Computer Science,
University of Massachusetts,
Amherst, MA 01003*

[‡]*Intel Research
Cambridge, UK
CB3 0FD*

{*svasu,kurose,towsley*}@cs.umass.edu {*dina.papagiannaki,christophe.diot*}@intel.com

Abstract

The performance experienced by wireless clients in IEEE 802.11 wireless networks heavily depends on the clients' ability to identify the Access Point (AP) that will offer the best service. The current AP affiliation mechanism implemented in most wireless clients is based on signal strength measurements received by the client from all the APs in its neighborhood. The client then affiliates with the AP from which it receives the strongest signal. It is well-known that such an algorithm can lead to sub-optimal performance, due to its ignorance of the load at different APs.

In this work, we consider the problem of AP selection. We identify *potential bandwidth* as the metric based on which hosts should make affiliation decisions, and define it as the (MAC-layer) bandwidth that the client is likely to receive after affiliating with a particular AP. We further limit ourselves to the use of passive measurements that do not require an end-host to affiliate with the AP, thus allowing the end-host to simultaneously evaluate the potential bandwidth to multiple APs in range. This can also facilitate more informed roaming decisions. We propose a methodology for the estimation of potential upstream and downstream bandwidth between a client and an AP based on measurements of delay incurred by 802.11 Beacon frames from the AP. Preliminary experiments conducted in a controlled environment demonstrate that the proposed methodology looks promising, yielding fairly accurate results under varying conditions.

1 Introduction

There has been an enormous growth in the adoption of IEEE 802.11 wireless networks in the last few years. The ease of installation and the low infrastructure cost of 802.11 networks makes them ideal for network access in offices, malls, airports, cafes, hotels and so on. The widespread deployment of IEEE 802.11 networks means that a wireless client is often in the vicinity of multiple APs with which to

affiliate. The selection of the AP that the client decides to affiliate with needs to be done carefully since it will dictate the client's eventual performance.

The conventional approach to access point selection is based on received signal strength measurements from the access points within range. However, it has been pointed out in several papers [1, 2, 4] that affiliation based on signal strength can lead to very bad performance for the end-host, since the signal-strength metric does not convey information regarding other attributes that affect end-host performance, such as the AP load and the amount of contention on the wireless medium.

In this paper, we describe how an end-host can take the aforementioned attributes into account while choosing an access point to affiliate with. We identify a metric that can capture *the bandwidth that an end-host is likely to receive if it were to affiliate with a given access point*, which we call *potential bandwidth*. The MAC-layer bandwidth offered by different wireless networks in the vicinity of the wireless client is a desirable metric as it takes into account the AP load, the contention on the wireless medium, as well as the signal strength.

In designing an affiliation algorithm based on potential bandwidth estimation, several constraints must be taken into consideration. The algorithm needs to be non-intrusive, i.e. it should not introduce additional overhead to the APs or their affiliated users. The algorithm should not require any changes at the AP side, if possible. More importantly, such an algorithm should be able to estimate potential bandwidth without previous affiliation with the several APs within range. Such a constraint minimizes the amount of time a client spends in the evaluation of the several choices it may have (since it does not associate and dis-associate with the different APs), while allowing for the continuous evaluation of AP performance even when an affiliation has taken place¹. The latter implies that a wire-

¹Notice that in the case of the initial affiliation the client will be able to identify the AP that offers the highest potential bandwidth on *any* frequency. In the case of roaming, the client will be able to quantify potential

less client implementing the proposed functionality will be able to make more informed and efficient roaming decisions, continuously quantifying the performance of all APs in range.

In this paper, we propose a methodology for the estimation of potential bandwidth between a given AP and an end-host that fulfills the aforementioned requirements. The proposed methodology does not require the end-host to change its current affiliation and introduces very little overhead. Unlike [1, 2], the affiliation algorithm proposed in this paper is end-host initiated and therefore, does not necessitate changes at the AP.

In a nutshell, our approach to potential bandwidth estimation relies on passive measurements of the timings of beacon frames sent out by an AP. Beacon frames are broadcast by APs periodically, and are used by APs to announce their identity as well as for the synchronization of the entire network. The delay between the time when a beacon frame is scheduled for transmission and its eventual transmission captures the load of the AP and the contention inside the network, conditions that the client would face if affiliated with that AP. The corresponding delay of data frames provides an estimate for the bandwidth a client will receive from the AP downstream. Upstream potential bandwidth estimation relies on frames sent by the client to the AP in the unaffiliated state and is based on a similar methodology that quantifies the respective delays.

Our technique can be used as part of an AP selection mechanism or for the evaluation of a wireless network's health. We evaluate its accuracy using controlled experiments in a low-noise environment. Preliminary experiments indicate that our approach yields fairly accurate estimates of the actual bandwidth from the AP to end-host, indicating that our approach looks promising.

The rest of the paper is structured as follows. In the next section, we describe related work. In Section 3, we describe our potential bandwidth estimation scheme. We discuss experimental results in Section 4. Finally, we conclude and describe in detail future directions in Section 5.

2 Related Work

The conventional AP selection mechanism, based on signal strength measurements, has been shown to lead to poor user experience [1, 2, 7] and highly unbalanced load distribution among APs [4]. Due to these shortcomings there have been several alternative proposals which typically fall in one of three categories: (i) AP-assisted [3, 7, 1], (ii) centralized [2], and (iii) active [8] solutions. In this work we take a step back and look at the fundamental requirements of the AP selection problem. Based on the identified requirements, we propose a technique that *does not require*

bandwidth only for the APs residing in the *same* frequency.

the assistance of the AP, does not require previous affiliation of the client with an AP, and is initiated by the client without the need for central coordination. Such properties allow for the continuous evaluation of the “quality” of all APs within range that could also facilitate better roaming decisions.

Our work targets the estimation of the potential bandwidth and not the available bandwidth as in [5], which is defined as the maximum rate at which a host can send its data without lowering the sending rates of other already affiliated hosts. In this work, we are not interested in the bandwidth available to a client before affiliation, but the MAC-layer bandwidth the client will receive after it affiliates with the AP. In addition, we do not aim to estimate the layer-3 throughput that a client would receive once affiliated with an AP, since such an estimation would require knowledge of the client's workload and its path through the wired network. The metric of potential bandwidth can characterize the wireless part of the client's connections. In future work, we intend to look into passive measurement techniques that could allow us to extend our estimates to account for the wired part of the network, say by passively observing the performance currently experienced by other users in the same wireless network.

The closest recent work to ours is [6], where the authors propose a methodology for passive bandwidth estimation between two communicating wireless stations. However, their method does not provide an estimate of the potential bandwidth that an end-host is likely to receive on a wireless link with another host (when one of the hosts is not part of the network yet).

3 Potential Bandwidth Estimation

In this section, we describe how an end-host can estimate both the potential upstream and downstream bandwidth between the AP and itself. The final affiliation decision made by the end-host is going to be some function of the upstream and downstream bandwidth and is likely to depend on the user's requirements. For the remainder of this work, we assume that the client has credentials to associate with any AP within range and selects the AP offering the highest bandwidth in the direction the client will use for its data transfer. We begin by providing a brief background of the IEEE 802.11 MAC protocol for data transmission.

3.1 Background

The protocol for data transmission is the same regardless of whether the transmitter is an AP or an affiliated host. Each node (including the AP) that has data to transmit in an IEEE 802.11 network first senses the channel for a duration equal to *DIFS* (Distributed Inter-Frame Sequence). If the node determines the channel to be idle for this duration,

then the node enters a back-off stage, in which it delays its transmission by a random number of time slots (each slot of duration $SLOT$) chosen from the interval $[0, CW]$, where CW is called the contention window size. If the channel is still idle at the end of the back-off stage, then the node transmits a Request-to-Send (RTS) frame to the intended receiver. On receiving the RTS frame, the receiver responds back with a Clear-to-Send (CTS) frame to the sender after a delay equal to Short Inter-Frame Sequence ($SIFS$). Nodes, other than the sender or the receiver, that hear either the RTS or the CTS frame delay their transmissions until after the end of the data transmission between the sender and the receiver, as specified in the duration field of the RTS and CTS frames. Upon receiving the CTS frame, the sender waits for a duration of $SIFS$ and sends its data frame. Finally, the receiver responds back with an ACK frame to acknowledge the receipt of $DATA$ frame. The absence of either a CTS or ACK frame causes the sender to timeout and re-transmit the RTS frame or the $DATA$ frame respectively. Many implementations also allow nodes to simply turn on or disable the RTS/CTS handshake. In this case, nodes directly transmit their data frames, on determining the channel to be idle at the end of the backoff stage.

We first describe our methodology to estimate the downstream bandwidth from an AP to an end-host in the absence of RTS/CTS handshake and then describe how the RTS/CTS handshake mechanism can be accommodated into the estimation scheme. We also discuss how an end-host can determine its upstream bandwidth to an AP. We initially ignore losses and subsequently, describe how losses can be accounted for in Section 3.5.

3.2 Beacon Delays

In order to estimate the downstream bandwidth from the access point to an end-host, we propose a methodology that allows the end-host to estimate the delays of the periodic *Beacon* frames sent from an access point. Figure 1 illustrates how beacon frame transmissions are handled at an access point. As seen from the figure, an access point sched-

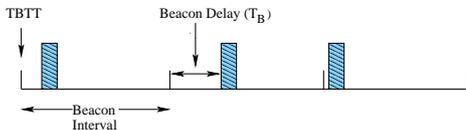


Figure 1: Beacon Transmissions at an Access Point

ules a *Beacon* frame every *beacon interval* (typically, 102.4 ms). The time instant at which the access point schedules the next beacon message is referred to as the *Target Beacon Transmission Time (TBTT)*. As per the 802.11 standard, time zero is defined to be a $TBTT$. Given the value of the beacon interval, the end-host knows the exact time

instants when beacon messages are scheduled for transmission. Once a beacon message is scheduled, it is transmitted according to the normal frame transmission rules. In this paper, we assume that beacon frames are not prioritized over other frames, as implemented in the APs used in our experiments. Handling beacon frame prioritization is an interesting extension and will be considered in our future work. The time difference between the instant when a beacon message transmission begins (as obtained from the timestamp field of the *Beacon* frame) and the $TBTT$ yields an estimate of the beacon delay, T_B , which is the total time spent by a beacon frame at the access point waiting for transmission. Since we assume that beacon frames are not prioritized over other frames, T_B provides an estimate of the total queuing delay plus the contention delay that will be experienced by a data frame transmitted by the AP. Note that beacon delays are computed solely based on timestamps provided by the access point and thus, synchronization issues do not arise.

We now proceed to describe how we can use observed beacon delays to estimate the downstream bandwidth from an AP to a mobile host.

| | |
|------------------|--|
| T | total delay incurred by a data frame from an AP |
| T_D | delay incurred between the instant when a data frame is scheduled for transmission to the instant that the frame is received at the receiver |
| T_A | delay of the ACK frame from the receiver to the sender |
| T_B | total contention delay experienced by a data frame from the AP |
| $DATA, RTS, CTS$ | size of the data, RTS, CTS frame respectively |
| R | data rate at which the sender transmits the data frame |
| R_b | basic rate at which control frames are transmitted |
| B | potential bandwidth from the AP to the end-host |

Table 1: Notations for the computation of downstream bandwidth

3.3 Downstream Bandwidth estimation in the absence of RTS/CTS

The total delay incurred by a data frame from an AP in the absence of RTS/CTS handshake is given by: the contention and transmission delay of the data frame plus the respective ACK delay.

$$T = T_D + T_A \quad (1)$$

T_D in turn can be estimated from the beacon delay T_B , estimated as in Section 3.2, and the transmission delay of

the frame², and is given by:

$$T_D = T_B + \frac{DATA}{R} \quad (2)$$

Upon receiving the data frame, the receiver sends an *ACK* frame after a delay of *SIFS*. *ACK* frames are fixed in length and are typically sent at the same rate as the data frame. Hence, knowing the sender rate, T_A can be easily determined as:

$$T_A = SIFS + \frac{ACK}{R} \quad (3)$$

The potential bandwidth B from the *AP* to the end-host is then given by:

$$B = \frac{DATA}{T} \quad (4)$$

3.4 Downstream Bandwidth estimation in presence of RTS/CTS

With the *RTS/CTS* handshake, each data frame transmission incurs a total delay (T) given by Eq. (5), the sum of delays incurred by the *RTS*, *CTS*, data and *ACK* frames respectively.

$$T = T_R + T_C + T_D + T_A \quad (5)$$

Since the frame transmission rules for an *RTS* and beacon frames are the same, the delay incurred by an *RTS* frame can be estimated using Eq. (6), as the sum of T_B and transmission delay (all MAC control frames are transmitted at the base rate).

$$T_R = T_B + \frac{RTS}{R_b} \quad (6)$$

Upon receiving a *RTS* frame, a receiver waits a duration of time equal to *SIFS* and transmits a *CTS* frame, again at the base rate R_b . The *CTS* frame is transmitted at the base rate R_b and its delay is given by:

$$T_C = SIFS + \frac{CTS}{R_b} \quad (7)$$

The delay incurred by the data frame is given by:

$$T_D = SIFS + \frac{DATA}{R} \quad (8)$$

Lastly, the computation of T_A remains the same across both schemes and is given by Eq. (3). The potential bandwidth B is then obtained using Eq. (4).

²If the AP has multi-rate support, then the current sending rate R of the AP can easily be inferred from the duration fields in the data frames transmitted by the AP.

3.5 Loss Probability Estimation

So far, the potential bandwidth estimation methodology assumed no packet losses. Losses occur due to collisions when multiple wireless stations transmit simultaneously and also due to environmental effects such as multipath, fading etc. Packet losses reduce the bandwidth between communicating stations, since they cause nodes to double their contention window and thereby, backoff for longer durations before retransmitting their data.

Thus, in order to estimate the potential downstream bandwidth from a given AP, an end-host needs to estimate the loss rate on the wireless link from the AP to itself. We propose that nodes infer frame losses, by exploiting the 12-bit sequence number field present in the 802.11 data and management frames. An end-host passively monitors all frames transmitted by the AP for a certain duration. The end-host can then infer data frame losses based on gaps in sequence numbers during the monitoring period. It is possible that the monitoring node may hear a data frame from an AP that is a retransmission of an earlier frame, which it did not hear. In this case, the monitoring node can detect retransmissions by looking at the Retry bit in the Frame Control field of the received frame. If this bit is set, it indicates that the frame is a retransmission of an earlier frame. Since the Retry bit does not indicate the number of retransmissions of a frame, we make a simplifying assumption that the probability of more than two successive retransmissions of a frame between an AP and a host affiliated to that AP is negligible.

The above described method of inferring loss rate, is useful both in the presence of *RTS/CTS* and in its absence. In the presence of *RTS/CTS*, the probability of an *RTS* frame loss differs from the probability of a data frame loss, since an *RTS* frame is transmitted at the base rate. An *RTS* frame loss can be inferred by a monitoring end-host, if the monitoring host overhears a data frame transmission from an AP to an end-host, but does not hear the *RTS* frame transmission from the AP to the end-host preceding the data transmission. Just as in the case of a data frame, an *RTS* frame retransmission can be detected from the Retry bit in the Frame Control field of the frame. Data frame losses can be detected from the missing sequence numbers over the monitoring period.

The estimated loss probability can be used to calculate the expected delays incurred by the *RTS* frames and data frames transmitted by an AP. For simplicity, we assume that *CTS* and *ACK* frames from the end-host to the AP are transmitted loss-free. This may be a reasonable assumption since *CTS* and *ACK* frames are very short. Furthermore, *CTS* frames are transmitted at the base rate and the *ACK* frames are transmitted collision-free. This assumption means that *CTS* and *ACK* frames always incur fixed delays. Losses then only impact the *RTS* and data frames

in our model. The estimated loss probability can easily be incorporated to obtain the expected back-off delay and the corresponding frame delay, using the analysis shown in [6].

When there are no affiliated hosts, a monitoring node does not overhear any transmissions except the beacon frames transmitted by an AP. Absence of a beacon frame in a *beacon interval* indicates that the beacon frame was lost. A monitoring host can estimate the loss probability of data frames to be the loss probability of the beacon frames. The *RTS* frames are transmitted at the base rate and can be assumed to be transmitted loss-free, especially given that there is no contention for the medium and that the probability of a collision is zero.

3.6 Upstream Bandwidth Estimation

Our proposed approach to estimating the upstream bandwidth requires that the end-host sends data frames to an access point in the unaffiliated state and records the time elapsed between the instant when a frame is scheduled for transmission and the time when the end-host receives an *ACK* message. It is interesting to note that the IEEE 802.11 standard allows a station in an unassociated state to send data frames to an access point. By sending several such frames and measuring the delays incurred by the frames, an end-host gets an estimate of the expected delay of a data frame. The potential upstream bandwidth can then be estimated using Eq. (4).

The implementation of the upstream bandwidth estimation scheme requires modifications to the wireless driver to allow a station to send frames in the unaffiliated state and is currently being investigated.

4 Experimental Results

In this section, we describe results from controlled experiments of our downstream bandwidth estimation scheme. All our experiments were conducted in an *anechoic chamber* that is designed to provide a very low noise environment, suitable for controlled experimentation. We configured a linux box with a Netgear MA 311 wireless PCI card to function as an access point running the *hostap* driver. The *RTS/CTS* handshake was disabled and the card was operated at a fixed rate of 11 Mbps.

4.1 Beacon delays in contention-free environments

In a contention-free environment and when the AP has no load, the mean beacon delay can be expressed as: Mean Beacon Delay = $DIFS + E[CW_{min}] \times SLOT + PLCP$, where $DIFS$ is the duration for which an AP senses the channel before transmitting a beacon frame; $E[CW_{min}] \times SLOT$ is the back-off delay once the AP has sensed the

channel to be idle for a duration $DIFS$; and $PLCP$ is the Physical Layer Convergence Protocol overhead associated with every transmitted frame. The IEEE 802.11b standard specifies the various parameter values as follows: $DIFS = 50\mu s$, $SLOT = 20\mu s$, $CW_{min} = 31$, $PLCP = 192\mu s$. From these values, we obtain the mean beacon delay to be $552\mu s$.

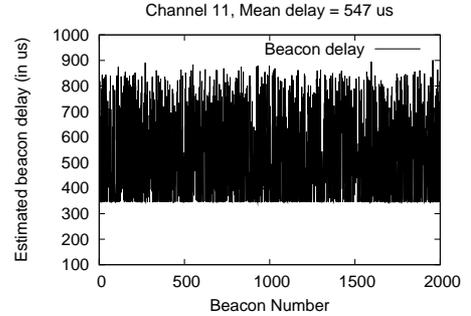


Figure 2: Beacon Delays when the AP has no load

We conduct a number of experiments to estimate the beacon delays using the methodology described in Section 3. Figure 2 shows that the mean estimated beacon delay value is $547\mu s$, which is close to the expected value of $552\mu s$. We next perform experiments to determine whether the bandwidth estimated through the beacon delay measurements closely approximates the actual bandwidth obtained by the end-host upon affiliation with the AP.

4.2 Bandwidth Estimation

In a collision-free environment, we know from Section 4.1 that the mean beacon delay is $552\mu s$. For a packet of size L bytes and data rate R , the potential downstream bandwidth is then given by (Eq. 4):

$$B = \frac{8L}{552 + \frac{8L}{R} + T_A}$$

where $T_A = 213\mu s$. For instance, when $L = 640$ and $R = 11$ Mbps, the potential downstream bandwidth yields an estimate $B = 4.16$ Mbps.

We performed a simple experiment to verify whether the actual bandwidth observed on the downlink from AP to an end-host compares with the estimated value obtained above. A UDP session is initiated from the AP to an affiliated end-host. The duration of the transfer was 200 seconds and the AP was constantly backlogged. The actual bandwidth B_m from the AP to the end-host for the duration of the transfer was measured to be 4.3 Mbps, which closely agrees with the estimate B obtained above.

In a second experiment, we place one AP and two wireless hosts $H1$ and $H2$ in the anechoic chamber. Host $H1$ is affiliated to the AP. A UDP session is initiated from the AP

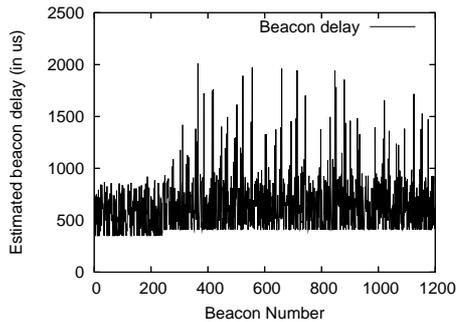


Figure 3: Beacon Delays when the AP is loaded

to the host $H1$. The UDP session consists of CBR traffic generated at the rate of 100 packets/second, each packet of size 576 bytes (640 bytes including the headers). Another host $H2$ is configured in “monitor” mode and records the delays observed for the beacon frames (beacons numbered 300 and higher in Figure 3). The monitoring host $H2$ estimates the mean beacon delay from the AP over the duration of the UDP transfer to be $687\mu s$. Using our bandwidth estimation methodology, $H2$ estimates the potential bandwidth from the AP to itself to be $B = 3.74$ Mbps. We then affiliate $H2$ with the AP and initiate another UDP session between the AP and host $H2$, simultaneously with the UDP session between the AP and $H1$. The AP is always backlogged with packets for $H2$. The actual bandwidth B_m from the AP to $H2$ is measured to be 4.06 Mbps, which agrees with the potential bandwidth estimate of 3.74 Mbps, obtained by $H2$ prior to affiliation with the AP. Thus, the experimental results suggest that our approach is promising.

5 Conclusions and Future Work

In this paper, we argued for potential bandwidth between AP and end-host as an important metric in the process of AP selection. We described a methodology for estimating the potential bandwidth based on delays experienced by beacon frames from an AP. We also presented results from experiments conducted in a low-noise environment, which showed that the proposed approach yields fairly accurate estimates of the bandwidth.

The described work is in progress and is currently looking at the following issues: - In this paper, we showed results from experiments in a noise-free environment. We plan to evaluate our bandwidth estimation scheme in the presence of noise, using controlled experiments. - The issue of how frequently should nodes estimate bandwidth to various APs in range and the duration over which estimations need to be carried out is currently under investigation. - We evaluate our bandwidth estimation with Netgear MA 311 Wireless cards that use the DCF (Distributed Co-ordination Function) mode of operation. Extending our scheme to be able to estimate the bandwidth between PCF-based (Point Co-ordination Function) APs and end-hosts is

another interesting future question. - In the Netgear 311 wireless cards, the beacon frames were transmitted with the same priority as the data frames. We wish to consider the case, when beacons are prioritized over other frames. - Finally, our estimation depends on the assumption that time zero at the AP is the time instant when the first beacon frame is scheduled for transmission, as specified in the IEEE 802.11 standard. While we observed this is very likely the case with the Netgear cards we experimented with, different vendors can be expected to implement beaconing differently. Inferring TBTTs by observing inter-beacon times remains a topic for further investigation.

References

- [1] A. Balachandran, P. Bahl, and G. Voelker. Hot-spot congestion relief and service guarantees in public-area wireless networks. *SIGCOMM Computer Communication Review*, 32(1), 2002.
- [2] Y. Bejerano, S. Han, and L. Li. Fairness and load balancing in wireless LANs using association control. In *Proceedings of ACM Mobicom*, Philadelphia, Oct 2004.
- [3] Cisco Systems Inc. Data sheet for cisco aironet 1200 series, 2004.
- [4] G. Judd and P. Steenkiste. Fixing 801.11 access point selection. In *Poster in Proceedings of ACM Mobicom*, Pittsburgh, Aug 2002.
- [5] K. Lakshminarayanan, V. Padmanabhan, and J. Padhye. Bandwidth estimation in broadband access networks. In *Proceedings of ACM Internet Measurements Conference*, Taormina, Oct 2004.
- [6] S. Lee, S. Banerjee, and B. Bhattacharjee. The case for a multi-hop wireless local area network. In *Proceedings of IEEE Infocom*, Hong Kong, Mar 2004.
- [7] R. Murty and E. Qi. An adaptive approach to wireless network performance optimization. Technical report, Corporate Technology Group (CTG), Intel Corporation, Technical Report, 2004.
- [8] S. Shah, K. Chen, and K. Nahrstedt. Available bandwidth estimation in IEEE 802.11-based wireless networks. In *Proceedings of 1st ISMA/CAIDA Workshop on Bandwidth Estimation (BEst)*, San Diego, Dec 2003.