

Seeker: A bandwidth-based association control framework for wireless mesh networks

Dhruv Gupta · Prasant Mohapatra ·
Chen-Nee Chuah

Published online: 17 May 2011
© Springer Science+Business Media, LLC 2011

Abstract The rapid deployment of wireless mesh networks across universities and enterprises, and the pervasiveness of mobile devices equipped with Wi-Fi connectivity, has resulted in a scenario wherein end users have the option to choose from a multitude of access points at any given location. Moreover, with the increasing availability of rich online content, there has been a steady increase in mobile Internet traffic. Since the choice of access point that a user associates with will directly impact his performance, it is imperative that there exist an efficient association control mechanism, in order to enhance the end user's experience. As part of this work, we propose Seeker, a novel framework for association control in wireless networks that utilizes "available bandwidth" as the design metric. The goal of Seeker is to assist the mesh network in making an intelligent decision regarding which access point a client should associate with. As part of our scheme, we implement and evaluate a passive tool to estimate available bandwidth in wireless networks. We then describe how we use this tool to implement our association control scheme, and evaluate it via extensive experiments on an outdoor testbed. Seeker takes into consideration the performance of the mesh backhaul, in addition to the client-to-AP link quality, thereby achieving significant

advantages over traditional association control schemes for wireless-LANs.

Keywords Wireless mesh · Available bandwidth · Association control

1 Introduction

Wireless mesh networks (WMNs) are increasingly becoming the network of choice for cities [6], universities [3] and enterprises [23]. Thanks to their significant advantages in terms of increased coverage, ease of deployment, and reduced costs, it has become increasingly common for a wireless user to see multiple wireless access points (APs) in any given area. According to authors in [29], high-density population areas have over 300 APs per square kilometer. In contrast to the case of wireless LANs, where different APs may belong to different individuals or organizations, in a mesh network, a user might be able to hear from multiple APs belonging to the same organization. This increased availability of wireless connectivity has also been complemented by the increasing popularity of mobile devices such as smartphones, music players, and e-book readers, which come equipped with Wi-Fi connectivity. The result has been a tremendous surge in mobile Internet traffic, a large portion of which traverses the high speed Wi-Fi networks, rather than the bandwidth-constrained cellular networks [5, 7].

Given this increasing usage of Wi-Fi based networks, it is imperative that there exist an efficient association control scheme to assist the client in selecting the right access point. This is all the more important for wireless mesh networks where the client's traffic will traverse through multiple wireless hops, and hence, the choice of wireless

D. Gupta (✉) · P. Mohapatra
Department of Computer Science, University of California
Davis, Davis, CA, USA
e-mail: dhgupta@ucdavis.edu

P. Mohapatra
e-mail: pmohapatra@ucdavis.edu

C.-N. Chuah
Department of Electrical Engineering, University of California
Davis, Davis, CA, USA
e-mail: chuah@ucdavis.edu

AP will have a definite impact on the client's performance. Present day devices such as laptops, and Wi-Fi enabled smartphones, primarily utilize a signal strength based mechanism for ranking APs, wherein users associate with the AP that has the maximum received signal strength (RSSI). However, several previous studies have already identified the short-comings of the RSSI-based approach. Selecting APs merely on the basis of signal strength can cause increased congestion at few selected APs, thereby resulting in poor performance for the end users [30]. Moreover, it has been widely studied that signal strength is not an accurate indicator of channel conditions in a wireless network [9]. Therefore, even if the received signal strength from an AP is high, it may not correctly indicate the current performance level of that AP. Even though several solutions have been proposed for association control and AP selection in 802.11-based wireless LANs (WLANs), they suffer from some fundamental limitations. Existing solutions that utilize different network metrics such as AP load, estimated throughput, and others, cannot be applied to mesh networks, which comprise of a multi-hop wireless backhaul. For example, merely knowing the round-trip time to the AP with which the client wishes to associate may not be enough, as the client's performance will also depend on the conditions inside the mesh backhaul.

As part of this work, we identify three main design considerations for an association control scheme.

1. First, apart from the channel quality of the client-to-AP link, we also need to be aware of the network performance within the mesh backhaul. Specifically, the client needs to be aware of the status of the network path between the APs of interest and the mesh gateway node.
2. Secondly, we contend that mesh network itself should aid in the association process by providing the client with the necessary network statistics. The network can keep track of the current link conditions and provide the client with an accurate image of the network performance.
3. Thirdly, we propose that the various metrics of interest should be gathered as efficiently as possible, while introducing the least possible amount of overheads. The measurement process should not significantly impact the network resources, nor result in performance degradation for the end user.

Based on the above design criteria, we propose *Seeker*—a novel network-assisted association control scheme for multi-hop wireless networks. *Seeker* utilizes bottleneck available bandwidth as the metric for association control. To this end, we first provide a comprehensive evaluation of existing available-bandwidth estimation schemes for

wireless mesh networks. We evaluate the performance of these schemes in the presence of various factors such as interference, packet loss, and rate-adaptation, on an outdoor testbed. We show how previous works in this area are severely lacking in their evaluation of bandwidth estimation tools. We further implement our own passive tool for bandwidth estimation in wireless networks and demonstrate its advantages over existing tools and techniques.

We then adopt the passive tool to implement our proposed association control scheme. As part of *Seeker*, each node in the mesh network estimates the available bandwidth for its respective links and exchanges this information with all the mesh nodes along its backhaul route. In order to avoid the communication overheads, our scheme utilizes the control packets available as part of the OLSR routing protocol [2], which is used on our testbed. In this manner, each mesh AP keeps track of the bottleneck node along its route to the mesh gateway, and makes this information available to the client device during the association phase. As the conditions within the mesh network change, the mesh APs will update their respective information, and the client devices will receive this new data during the next scanning phase. The clients can choose to re-associate with new APs based on the updated information. We perform extensive experiments on our mesh testbed to evaluate the efficacy of our proposed scheme. Based on our results, we observed that *Seeker*, which uses available bandwidth as the design metric, provides significant performance improvement for the end user, over signal strength based association schemes. In brief, we make the following contributions in this paper:

- We evaluate various existing available bandwidth estimation tools for multi-hop wireless networks. We also implement our own passive tool and demonstrate its advantages over existing schemes.
- We propose *Seeker*, a new association control scheme for multi-hop wireless networks. Our scheme utilizes the passive tool to estimate available bandwidth in mesh networks, which is shown to provide the maximum accuracy, and lowest overheads, among the evaluated tools.
- We evaluate the performance of *Seeker* via extensive experiments, for varying scenarios. We provide concrete results that demonstrate the advantages of our scheme over existing schemes based on signal strength.

The rest of the paper is structured as follows. Section 2 outlines the limitations of some previously proposed available bandwidth estimation schemes and association control schemes, followed by an overview of the problem addressed in this paper. In Sect. 3, we describe the implementation of our passive tool, along with the comparison results of existing bandwidth estimation

techniques. In Sect. 4, we outline the architecture of our association control scheme. The implementation details are described in Sect. 5, along with the experimental methodology and results. Section 6 outlines the related work, and finally Sect. 7 concludes this paper.

2 Motivation and problem overview

2.1 Limitations of bandwidth estimation tools

While several tools and techniques have been proposed for measuring the available bandwidth in a network, their focus has primarily been on wired networks. These tools lack a thorough evaluation in the realm of wireless networks, especially in the multi-hop setting. The passive tools have mostly been evaluated via simulations only. We identify the following two primary drawbacks with the existing solutions:

Measurement accuracy: The probe-based tools [28, 41, 44] are based on certain assumptions such as fixed capacity, and fluid model of cross-traffic, about the underlying network. However, the contention-based random access nature of IEEE 802.11 protocol may not conform to these assumptions. Moreover, it is unclear as to how these tools perform in the presence of factors such as rate-adaptation, which are specific to IEEE 802.11-based wireless networks.

Network overheads: The probe-based tools will involve some overheads in the measurement process that may adversely impact the network performance. This problem is further compounded in wireless networks due to a shared transmission medium. Previous studies have not focussed on this aspect of bandwidth estimation.

2.2 Limitations of association control schemes

Several previous schemes have been proposed for association control in wireless networks. However, their primary focus has been on single-hop wireless LANs. Moreover, many of the proposed schemes have been evaluated only via analytical methods or simulations, and are based on approximations that do not conform to a real-world setting. We identify the following main drawbacks of such schemes:

Association control metric: Previous works on association control and AP selection have proposed using metrics such as AP load [30], potential bandwidth [14], and TCP throughput [39]. However, for the case of multi-hop wireless networks, the client's performance depends not only on the channel conditions between the client device and the AP, but also on the backhaul path between the APs and the mesh gateway. While it may be easy to obtain the

above mentioned measurements for a single wireless hop, integrating them over multiple wireless hops may not be straight forward. Hence, it is necessary to use a metric that provides an accurate view of the complete path from the client to the mesh gateway.

Measurement overheads: Several previous approaches [36, 38, 39] have proposed using probe-based active schemes to estimate the network conditions for association control. These approaches can be applied to multi-hop wireless networks in order to make the clients aware of the mesh backhaul performance. However, these approaches suffer from a major drawback. Performing end-to-end tests that involve data transfer would add overheads in the wireless network, which will contend with the transmissions of other clients, leading to potential performance degradation. We aim to minimize the overheads involved in both the measurement, and reporting, of network data. Metrics such as ETX [21] and ETT [23] also involve measurement overheads and are hence not suitable.

Energy considerations: In the last few years, Wi-Fi based hand-held devices have become extremely popular. However, these devices are usually resource constrained, especially in terms of available power. It has been shown in previous works [24, 25] that transmission and reception of data via the wireless interface causes significant energy consumption at the end device. Moreover, sending small sized probes results in a larger wastage of energy [15] due to the high cost of keeping the interface up. Hence, we contend that using probe-based active schemes for estimating network conditions across multiple APs is not a suitable approach, as it will cause extra energy consumption.

2.3 Problem characterization

In order to overcome the drawbacks of existing approaches, we propose *Seeker*, a novel association control scheme for multi-hop wireless networks. An overview of our problem is depicted in Fig. 1. The figure shows a multi-hop wireless network in an indoor setting. The network comprises of pre-existing clients that are generating some background traffic. A new client arrives and can associate with one of the several possible choices of APs. From the end user's perspective, the AP that can provide the best end-to-end throughput or delay will be the best choice. However, from the network's perspective, we may also be interested in evenly distributing the traffic across the network, to achieve better network utilization. In order to help the client in making the right decision, we propose that the mesh network should provide the client device with the necessary information regarding the network status. For our scheme, we use available bandwidth as the decision metric, as it incorporates both the performance, and the




Fig. 1 Problem Overview. A mobile user has the option to choose from multiple APs

load-balancing, factors. Instead of ranking the APs in order of received signal strength, we propose to rank the APs based on their bottleneck available bandwidth values. The AP that provides a backhaul path with the minimum bottleneck available bandwidth will be given the highest rank.

Our proposed scheme addresses the previously mentioned design challenges and drawbacks in the following manner:

Design metric: We use bottleneck available bandwidth as the design metric for our association control scheme. As is shown in later sections, available bandwidth incorporates both the link conditions and the network load, thereby providing us with an accurate metric for association control.

Measurement overheads: We use a passive technique to estimate available bandwidth in the mesh network, thereby avoiding any measurement overheads. Further, we utilize existing control packets in the network to disseminate the necessary information, thereby reducing the communication overheads as well.

Energy efficiency: By avoiding any active scanning or probing on part of the end user, we help conserve power on mobile devices. This is especially important with the prevalence of Wi-Fi enabled devices such as PDAs, e-book readers, and music players.

We first describe the implementation and evaluation of our passive bandwidth estimation tool, followed by the description of our association control scheme.

3 Estimation of available bandwidth

3.1 Passive bandwidth estimation tool

Our passive bandwidth estimation technique is based on the concept of monitoring the channel for packet

transmissions and using this information to calculate the channel utilization. Consider a wireless network with RTS/CTS disabled. For each packet transmission, the sender has to wait for DIFS amount of time (T_{DIFS}), followed by which the sender will transmit the preamble and the PLCP header. After this, the actual data packet is transmitted (T_{DATA}). The receiver waits for a SIFS (T_{SIFS}) amount of time and then transmits the link-layer ACK (T_{ACK}) for the packet. The transmission may also involve some random back-off time ($T_{backoff}$), giving the following equation:

$$t = T_{DIFS} + T_{backoff} + T_{DATA} + T_{SIFS} + T_{ACK} \tag{1}$$

For a given choice of protocol (802.11 a/b/g), all the values are fixed, except for the back-off time and the data transmission time. The data transmission time depends on the frame size and the modulation rate (physical data rate) used. [1] lists the various parameter values associated with 802.11a scheme (which is used in our testbed). The idea behind the passive scheme is for each node to monitor the channel and listen for ongoing packet transmissions. For each packet that a node hears, it calculates the total transmission time based on the frame size and the data rate used to transmit the packet. Thus, each node can estimate the channel utilization (μ), and multiplying the complementary of this quantity ($1 - \mu$) by the capacity of the channel gives us an estimate of the available bandwidth:

$$BW = (1 - \mu) \times C = (1 - \frac{t}{T}) \times C \tag{2}$$

where t is the total time for which the channel was busy (Eq. 1), in a sampling period of T . The capacity (C) of the channel can be estimated by using tools such as Pathrate [31]. Table 1 lists the capacity estimates for our testbed for various transmission rates. It should be noted that in wireless networks, the capacity of a link may also vary over time. We can capture this variation by periodically measuring the capacity of the network. The periodicity of capacity measurement will be dependent on the network environment. Further implementation details of our tool are provided in [26].

While the above technique can be used to estimate the available bandwidth for each individual node, further modifications are required in order to calculate the available bandwidth for a wireless link between two nodes. In [48], the authors have discussed the importance of estimating the overlap probability of the “free” channel time

Table 1 Capacity estimates (Mbps) for outdoor testbed for 802.11a transmission rates

Trans. Rate	6	9	12	18	24	36	48	54
Cap.	4.05	6.6	7.1	7.1	7.2	7.3	7.3	7.3

between two nodes, in order to obtain accurate estimates. In order to achieve this, the channel busy time at each node is divided into two components. We have the BUSY time (T_B) during which a node is transmitting or receiving data, and the SENSE BUSY time (T_S), during the which the node can sense transmissions from other nodes, but is not involved in any transmission or reception itself. Hence, Eq. 2 can be modified as:

$$BW = \left(1 - \frac{T_B + T_S}{T}\right) \times C \quad (3)$$

If the two nodes have similarly overlapping periods of IDLE, BUSY, and SENSE BUSY times, then we can use Eq. 3 to estimate the available bandwidth on the link between the two nodes. However, if these time periods are non-overlapping, then the actual bandwidth will be less than the estimated bandwidth. In order to integrate this in our bandwidth estimation scheme, we use the approach outlined in [48]. For two nodes a and b , let p_1 (p_2) be the probability that node a (node b) is in the SENSE BUSY (IDLE) state and node b (node a) is in the IDLE (SENSE BUSY) state. Also, let T_S^i be the time for which node i is in SENSE BUSY state in the interval T . Then we have the following estimates of each node's available bandwidth:

$$BW_a = \left(1 - \frac{T_B^a + T_i^a \times (p_2 \times T_S^b/T)}{T}\right) \times C \quad (4)$$

$$BW_b = \left(1 - \frac{T_B^b + T_i^b \times (p_1 \times T_S^a/T)}{T}\right) \times C \quad (5)$$

Hence, the available bandwidth for the link between nodes a and b is given by:

$$BW_{a,b} = \min(BW_a, BW_b) \quad (6)$$

3.2 Comparison of existing tools

Apart from the passive tool that we implemented, we selected four existing probe-based tools: Pathload [28], Spruce [44], PathChirp [41], and PTR [27]. These tools are representative of the various bandwidth estimation techniques proposed so far. Pathload and PathChirp use the probe-rate model to estimate available bandwidth. They are based on the concept of self-induced congestion, wherein they estimate the available bandwidth by searching for the turning point at which the probe sending and receiving rates start matching. Spruce and PTR are based on the probe-gap model and they estimate available bandwidth based on the time gap between the arrivals of two successive probes at the receiver.

A preliminary comparison study of these tools appears in [26], where we evaluated these tools in a controlled setting on an indoor testbed. As part of this work, we extend our study by performing a more detailed evaluation

of these tools on an outdoor testbed, in an uncontrolled setting. The evaluation methodologies and metrics are the same as ones used in our previous work [26].

Evaluation testbed: Our testbed is deployed in a wildlife reserve [4]. The outdoor testbed consists of about thirty nodes spread over two thousand acres of forest land. The wireless nodes in our testbed support both 802.11b/g and 802.11a modulation schemes. This testbed has the advantage of being free from any external interference in the region. We use the OLSR routing protocol for normal operation, but also have the option of statically configuring the routes. The hardware consists of Soekris net4826 embedded devices running a customized version of Linux. We use the open-source Madwifi wireless device driver on these nodes. Our testbed also supports ecological research in the area and the traffic comprises of audio and video data generated by sensors deployed in the network. This traffic is generated periodically at a fixed rate and uses fixed sized packets.

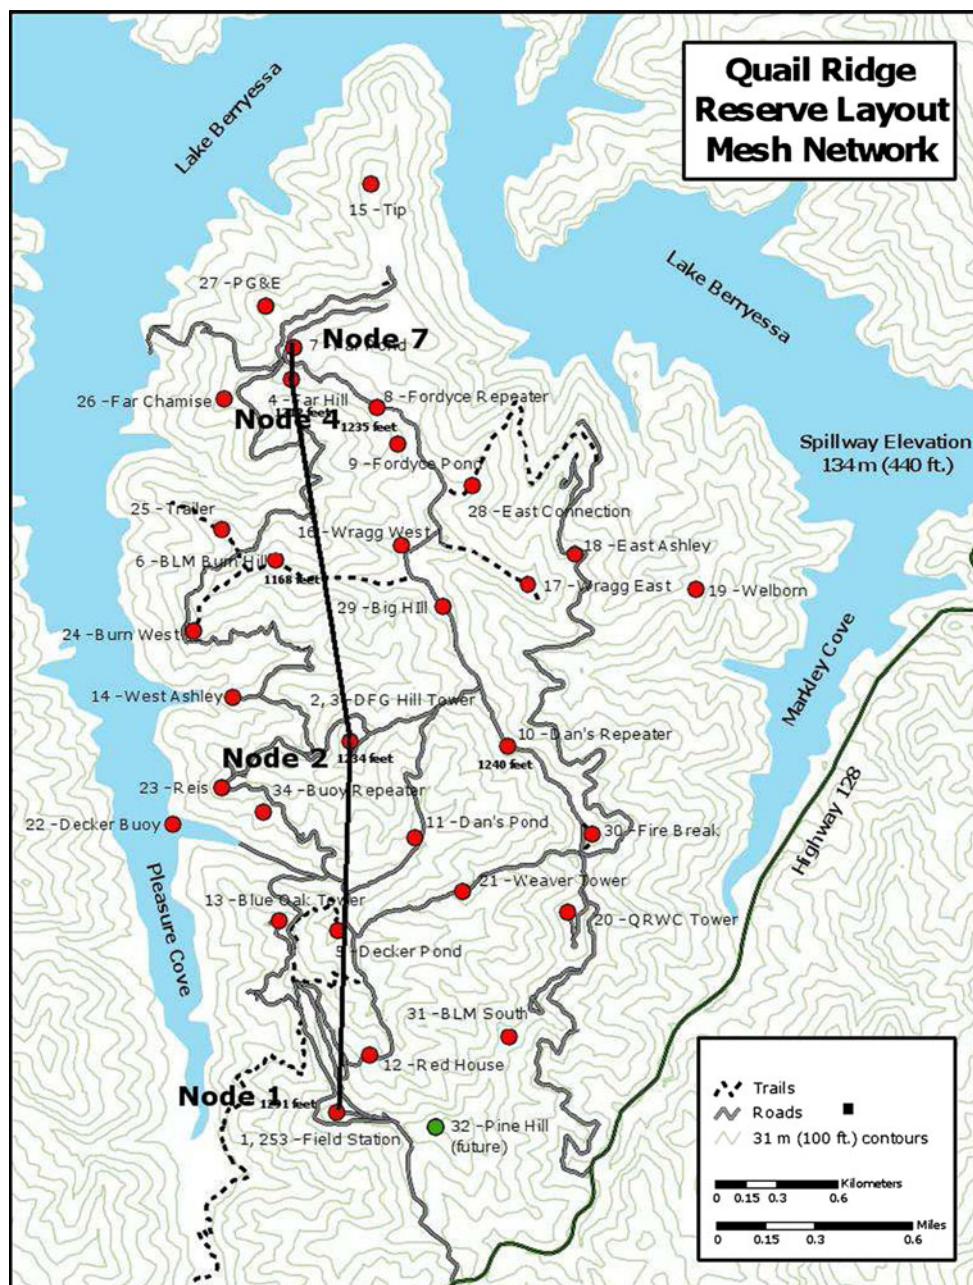
3.2.1 Impact of interference

Figure 3 shows the results for the single-hop case (link 1–2 in Fig. 2) for various IEEE 802.11a transmission rates. We observed that the measured estimate (referred to as the “actual” estimate) saturated with increasing transmission rate. Our passive tool returned fairly accurate estimates. The other tools under-estimated the available bandwidth, especially for higher transmission rates. This happens because due to the random access nature of IEEE 802.11 protocol, the initial probe-gap for these tools is high, which translates to a low initial sending rate. As a result, the tools do not ramp up their sending rate enough, and hence underestimate the available bandwidth. PathChirp provided more accurate estimates (average of 12% error) due to its ability to change the sending rate in the middle of a packet train. The passive tool returned the most accurate estimate (within 6% of the actual estimate).

3.2.2 Impact of cross-traffic

We also investigated the behavior of these tools in the presence of cross-traffic. While interference is caused by transmissions on neighboring links, cross-traffic refers to packets flowing between the same two nodes on which the bandwidth estimation is being performed. In order to study the responsiveness of the tools to cross-traffic, we use the D-test [44]. We first measure the available bandwidth using each tool, denoted by M_1 . We then introduce additional cross-traffic on that link at the rate of $\tau = \frac{M_1}{2}$. We then again use the tools to estimate the available bandwidth,

Fig. 2 Layout of QuRinet. For single hop, we use link 1–2. For multi-hop, we use link 1–2–4–7 (shown as highlighted)



denoted by M_2 . We then calculate the relative error of the measurements as:

$$\epsilon = \frac{M_2 - \tau}{\tau} \quad (7)$$

We first fixed the transmission rate of the link with cross-traffic at 6 Mbps, and then at 54 Mbps. Figure 4 shows the CDF of the relative error for the two scenarios. Ideally, the CDF plot should be a step function at the zero mark. The passive tool seems to be very accurate, while only slightly over-estimating the bandwidth. It gave a relative error of less than 0.1 for 99% of the cases. Spruce seems to respond well to cross-traffic (as shown by the step

function) but tends to either over or under-estimate the bandwidth. This is because while Spruce is still able to send out the probe-pairs, the delay estimates are being impacted by the cross-traffic, resulting in inaccurate estimates. Pathload and PathChirp did not respond well to cross-traffic and over-estimated the bandwidth. It seems that these tools are estimating the fair-share of the bandwidth, which is more than the actual available bandwidth. Similar observations were made in [18] and [22] for the performance of these tools in wired networks. Our results confirm these findings for wireless networks. The performance of the tools seems to improve for the second scenario because there is more available bandwidth

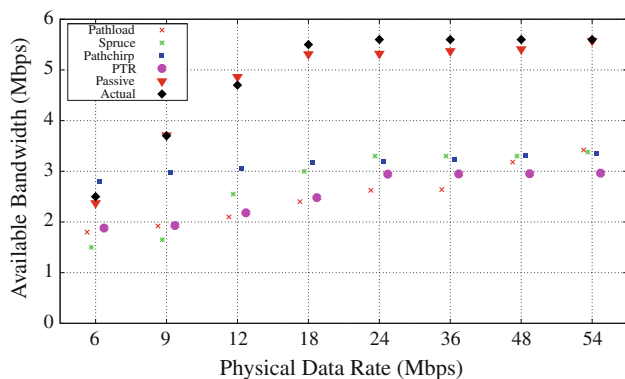


Fig. 3 Performance of passive scheme on outdoor testbed for single hop with interference

for the tools and hence the impact of cross-traffic is not as much.

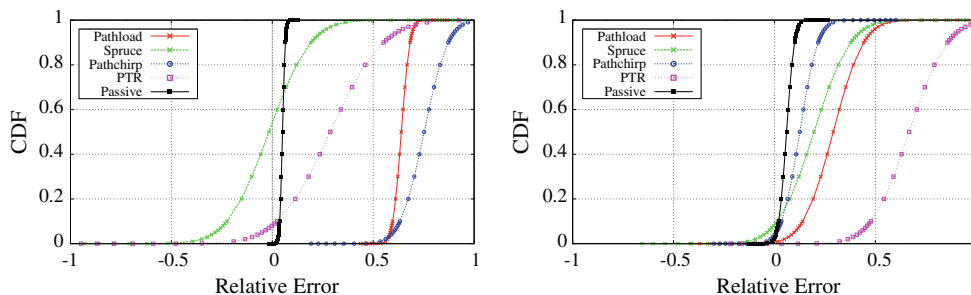
3.2.3 Impact of rate adaptation

Current IEEE 802.11 implementations incorporate a rate-adaptation feature that enables wireless radios to vary their transmission rates based on current channel conditions. If the data rate of a node changes, then the channel occupancy time of the node will also change, which will impact the available bandwidth estimation. Most previous studies assume fixed transmission rates when evaluating bandwidth estimation tools, thereby not characterizing the performance of these tools accurately.

In order to model the rate-adaptation behavior, we disabled the default rate-adaptation algorithm available in the Madwifi wireless driver. We modified the wireless driver to change the transmission rate at different instants of time, irrespective of the channel conditions. We started the experiment with the transmission rate of link 1–2 fixed at 6 Mbps, and changed it to 54 Mbps after a few seconds. Such an approach enabled us to repeat the experiments with similar variations in transmission rate.

Figure 5 shows the corresponding results. The passive tool reacted very quickly to the variations in the transmission rate. On the other hand, since PathChirp uses packet trains with varying rates, there is some latency

Fig. 4 Responsiveness of tools to cross-traffic for single hop with 6 and 54 Mbps transmission rate



involved before it can adjust its sending rate to the new transmission rate. Pathload and PTR calculate an average dispersion rate (ADR) before the actual estimation, and since this value does not get updated when the transmission rate changes, these tools are unable to reflect the variations in the transmission rate. Spruce did not give any results. The change in transmission rate seemed to impact the Spruce algorithm and it would stop sending packets when the data rate was changed.

3.2.4 Impact of packet loss

Probe-based estimation tools rely on successful delivery of probe packets from the sender to the receiver. The delay measurements obtained from these probes are used for estimating available bandwidth. Hence, if a significant number of probe packets are lost, then the tool may give an inaccurate estimate. We thus decided to investigate the impact of packet loss on the performance of these tools.

In order to have a similar rate of packet loss across experiments with different tools, we did not rely on interference or weak signal quality. Instead, we modified our wireless driver to drop one packet for every fifty packets correctly received. These packets were dropped at the receiver end. The physical layer data rate for link 1–2 was fixed at 6 Mbps (rate-adaptation was disabled). The results are summarized in Table 2. Pathload and Spruce did not report any estimates. This is because both the tools require certain number of probes to get through. For the packet loss rate we choose, these tools do not give an estimate. For PTR, the loss of probe packets causes it to lower the sending rate, and as a result, it underestimates the available bandwidth. However, the passive tool provides an accurate estimate since there are no probes involved.

3.2.5 Impact of multiple hops

Our final experiment involved validating the accuracy of the various tools over multiple wireless hops. Multiple wireless hops lead to several issues such as self-interference, which will impact the performance of the selected tools. We used a three hop scenario (Fig. 2), with each link




Fig. 5 Performance of passive scheme on outdoor testbed for single hop with rate-adaptation

Table 2 Impact of packet loss on estimation accuracy

Tool	Estimate (Mbps)
Actual	2.34
Pathload	0
Spruce	0
Pathchirp	1.75
PTR	1.625
Passive	2.2

configured to be on the same channel. The center link (2–4) was fixed at the lowest 6 Mbps transmission rate, thus acting as the bottleneck link. Figure 6 shows that the passive technique provides very close estimates to the “actual” value. While PathChirp consistently over-estimated the bandwidth, Pathload and PTR provided lower estimates.

3.2.6 Measurement overheads

Apart from evaluating these tools for accuracy, we also evaluated the overheads involved in the measurement process. In order to quantify the overheads of these tools, we look at *measurement latency and intrusiveness*. Measurement latency is the amount of time after which the tool returns the estimate. For some tools, even though a first estimate is obtained quickly, it takes longer for the tools to stabilize and provide a more accurate estimate. For such tools, we use the latter time as the measure of latency. A lower latency is desirable, in order to quickly predict the available bandwidth. Intrusiveness [43] is defined as the ratio of the average probing rate of the tool to the available bandwidth. The higher the intrusiveness, the more bandwidth it consumes during the estimation process, and more will be the impact on existing user traffic.




Fig. 6 Performance of various tools on outdoor testbed for multiple hops with self-interference

Figure 7 shows the latency and intrusiveness for the single-hop scenario. We observed that Pathload took a longer time (nearly double for all transmission rates) to return the first estimate, which may not be desirable in a wireless network where the channel quality varies at much smaller granularity. PTR was seen to have higher intrusiveness and also caused about 4% packet loss for the cross-traffic. It should be noted that the passive tool does not involve any overheads in the measurement process itself. We observed a similar result for the multi-hop scenario (Fig. 8).

4 Available bandwidth based association control

In this section, we describe Seeker, our proposed scheme that utilizes available bandwidth as the metric for association control. We describe how to estimate the channel conditions between the AP and the client, and how to use the passive available bandwidth measurement tool to estimate the conditions in the mesh backhaul.

4.1 Client-to-AP link quality

This parameter indicates the performance of the first wireless hop between the client device and the mesh AP. Traditional AP selection schemes for wireless LANs have focused on evaluating this particular parameter. Signal strength, though important, has been shown to be an insufficient metric, and previous works have highlighted the importance of incorporating more metrics in the AP selection process.

We propose using the *estimated throughput* metric to estimate the channel link quality between the client and the mesh AP. We derive this metric using the concepts outlined in Sect. 3.1 (Eq. 2). Consider an AP a that has a set of users U associated with it. Let L_u^a be the load generated by each user $u \in U$ on AP a per unit time. If R_u^a is the

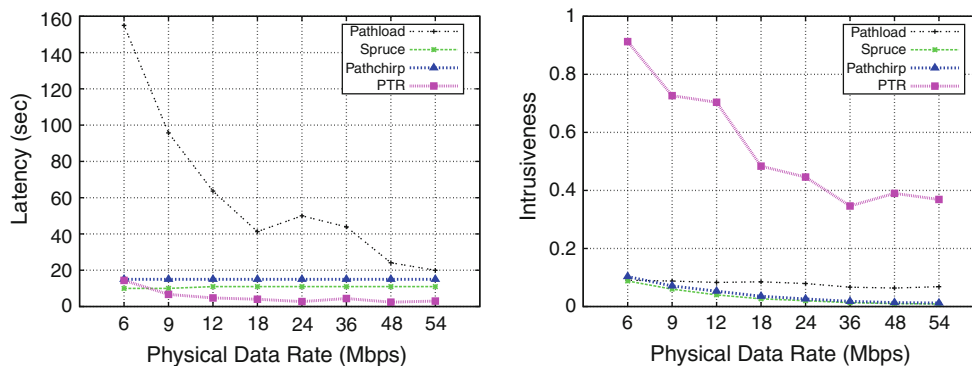


Fig. 7 Latency and Intrusiveness for single hop with interference

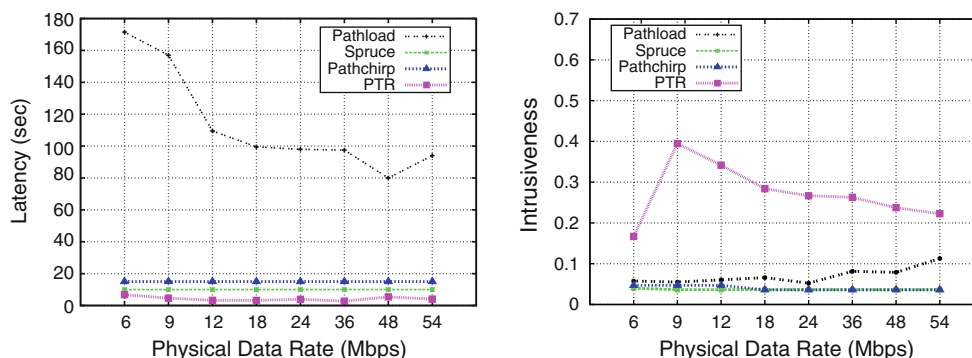


Fig. 8 Latency and Intrusiveness for three hops with self-interference

maximum transmission rate achieved by each user, then we define the utilization of AP a as:

$$\mu_a = \sum_{u \in U} \frac{L_u^a}{R_u^a} \tag{8}$$

This utilization gives us an estimate of the fraction of the AP’s capacity that is being used to serve the existing clients. Hence, the remaining capacity is available for the new client, and can be used to estimate the throughput between the AP and the new client. If R_c^a denotes the transmission rate between AP a and the new client c , then the estimated throughput is given by:

$$Thr_c^a = (1 - \mu_a) \times R_c^a \tag{9}$$

This value is made available by the respective APs to the clients during the scanning process.

4.2 Performance of mesh backhaul

We attempt to derive a metric that captures the performance of the client’s flows within the mesh backhaul. This approach is significantly different from existing association control schemes for wireless LANs, which do not need to measure the performance of the wired backhaul network. On the other hand, in a mesh network, even if the

client-to-AP link is good, the route between that AP and the gateway node may suffer from poor link quality.

There are various techniques to capture this performance metric. The simplest approach would be to extend the current IEEE 802.11 mechanism that uses signal strength for AP selection. We can obtain the signal quality measurements for each link along a backhaul path and identify the links with the minimum signal quality along each route. The route that has the best bottleneck link in terms of signal strength will be chosen to carry the client’s traffic. We evaluate this approach on our testbed and compare it with our proposed scheme.

Another approach is to utilize a probe-based scheme to measure certain network metrics such as packet loss, RTT, and so on. The APs can then exchange this information amongst themselves in order to estimate the end-to-end path metric. However, such a scheme involves extra measurement and communication overheads that may adversely impact the end users’ performance.

Our approach involves utilizing the passive bandwidth estimation tool that we described earlier. Our tool measures the available bandwidth on each link in a passive manner, without adding any extra overheads. As shown in Sects. 3.2.1 and 3.2.4, our tool provides accurate estimates even in face of interference and packet loss. As part of our

scheme, each mesh node estimates the available bandwidth on its respective links. It then exchanges this information with other mesh nodes that lie along its route to the gateway node. Once this information is available, each mesh node uses Eq. 6 to compare the available bandwidth data for all the mesh nodes along its path to the mesh gateway, thereby identifying the bottleneck point in terms of available bandwidth. This information will be made available to the client device during the scanning phase. The client device will receive similar information from multiple mesh nodes (APs), and can then make an intelligent decision regarding which AP to associate with.

Using available bandwidth as the metric for association control provides two significant advantages:

- It automatically incorporates channel diversity in the association control mechanism. If a particular path has links on different channels, then there will be reduced interference among the links, resulting in increased available bandwidth.
- It also incorporates the current link conditions and load conditions in the association control mechanism. Since the estimation of available bandwidth involves both the current node utilization, and the current transmission capacity (which is determined by the prevailing link conditions), we are able to provide the client with the most accurate and up-to-date status of the network.

Our algorithm for association control using available bandwidth involves the following steps:

- **Step 1:** The client device first scans for the available APs, and obtains their estimated downlink throughput values.
- **Step 2:** The client sorts the APs in order of their downlink throughput values and selects the top three APs.
- **Step 3:** The client queries the top three APs for their respective bottleneck backhaul bandwidth information and selects the best AP to associate with.
- **STEP 4:** The mesh APs continuously monitor the channel conditions and update their throughput and available bandwidth values.

The client device can periodically scan the channels and query the available APs for the updated information. Based on these new measurements, the client device can decide to associate with a new AP. A similar process can be used if the client moves to a new location.

When a new client enters the system and associates with an AP, other associated clients may experience varying degrees of performance degradation. This might result in some other clients re-associating with a new AP, causing a potential cascade of re-associations. In order to avoid such frequent client re-associations in the network, we adopt the

following solution. When a client gets the updated information from the mesh network, it will re-evaluate the rankings of the APs it can hear from. Only if the new AP provides a significant performance improvement over the existing AP (an improvement of 5% in our scheme), will the client initiate the re-association process.

5 System implementation and evaluation

5.1 Implementation details

As mentioned in Sect. 4, we need to gather certain network statistics from the mesh nodes for our association control scheme. We have already described the implementation of our passive bandwidth estimation scheme. We now outline how we measure the estimated throughput value for the client-to-AP link, and how we estimate the bottleneck link in terms of available bandwidth in the mesh backhaul.

The design architecture of Seeker is shown in Fig. 9. The first metric that we need to estimate is the channel quality between the client and the APs. This function is performed by the client-AP link quality estimator. We currently assume that most of the traffic flows in the downlink direction from the mesh gateway to the clients, since most servers currently reside on the external wired network. Hence, we only estimate the downlink channel quality between the APs and the clients. We use the estimated throughput metric for this purpose. Each mesh node keeps track of the total number of bytes that it transmits to clients associated with itself in every sampling period. It also keeps track of the maximum possible number of bytes that it can transmit to each client (transmission rate) in a given time period. Using these two parameters, each mesh AP can calculate its downlink utilization. Based on this information, we can calculate the estimated throughput that the client will achieve in the downlink direction using Eq. 9. This data is periodically estimated by the client-AP link quality estimator and passed on to the Seeker algorithm.

The second component of our scheme involves the backhaul performance estimator, which estimates the quality of the backhaul routes in terms of bottleneck bandwidth. Using the technique described in Sect. 3.1, each mesh node estimates the available bandwidth on its respective links. The data disseminator component is responsible for the exchange of available bandwidth data among the mesh nodes. This is done using the information from the underlying OLSR routing protocol. We modify the OLSR HELLO message and add a field to carry the available bandwidth value for each node, thus avoiding the overhead of transmitting extra packets. When a mesh node receives the bandwidth data from another node, it checks to

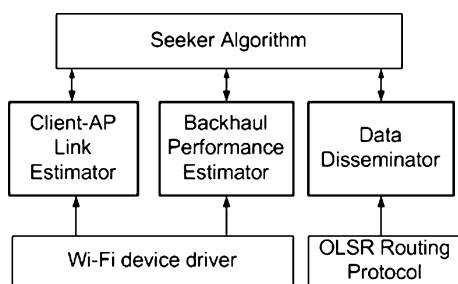


Fig. 9 Design architecture for Seeker

see whether that node lies on its path to the mesh gateway. Using this information, each mesh node can estimate the bottleneck link on its path to the mesh gateway in terms of available bandwidth.¹

As part of our scheme, we currently require slight modifications at the client end. In order to obtain the necessary information from the mesh APs during the scanning phase, the clients need to query an OLSR plugin. We implemented a customized plugin for OLSR that runs on each AP, and keeps track of all the data measured by that AP. The clients themselves do not need to run the OLSR daemon. When a client performs the AP scanning process, it simply sends a GET query message to the desired mesh APs, and the OLSR plugin on the respective APs responds with the measured statistics. The client first requests for the downlink throughput data from all the APs, and then the available bandwidth data from only the top three APs.²

5.2 Evaluation results

In this section, we describe the results of the experimental evaluation of our proposed association control scheme. The testbed has already been described in Sect. 3.2. Figure 10 shows the overlapping coverage of the various APs at Quail Ridge. We performed war-driving at the reserve using a Linux-based laptop, attached with a USB device. We used the *kismet* software to detect various available access points, and recorded their location using the GPS data. At most locations, we were able to hear from multiple mesh APs, giving us the option of choosing from multiple available access points. Similar studies conducted at various other universities and urban areas further motivate the need for an efficient association control scheme. By choosing the right AP to associate with, the client will be able to significantly improve the achieved performance.

¹ We use OLSR as it is the routing protocol deployed on our testbed. Our scheme will work with any other routing protocol.

² In the absence of OLSR, this information can be obtained via other means, such as by generating small control packets, or by using the IEEE 802.11k protocol [35, 47].

The background traffic in our testbed was generated by the various audio, video, and weather data sensors supported by our testbed. These sensors generate periodic data such as weather-related information and images of various flora and fauna, and transmit them to a central server. We do not add any other background traffic for our experiments.

We compare our approach against a modified RSSI-based approach. In the IEEE 802.11 protocol, the clients associate with the AP that has the maximum received signal strength. We extend this approach over multiple wireless hops. We keep track of the signal strength measurements for each link in the mesh network. Similar to the bandwidth based approach, mesh nodes along the same backhaul route exchange their respective data with each other. In this manner, each mesh AP is aware of the bottleneck link along its backhaul route, in terms of signal strength. The mesh AP with the best link in terms of bottleneck RSSI along its backhaul route will be chosen by the client.

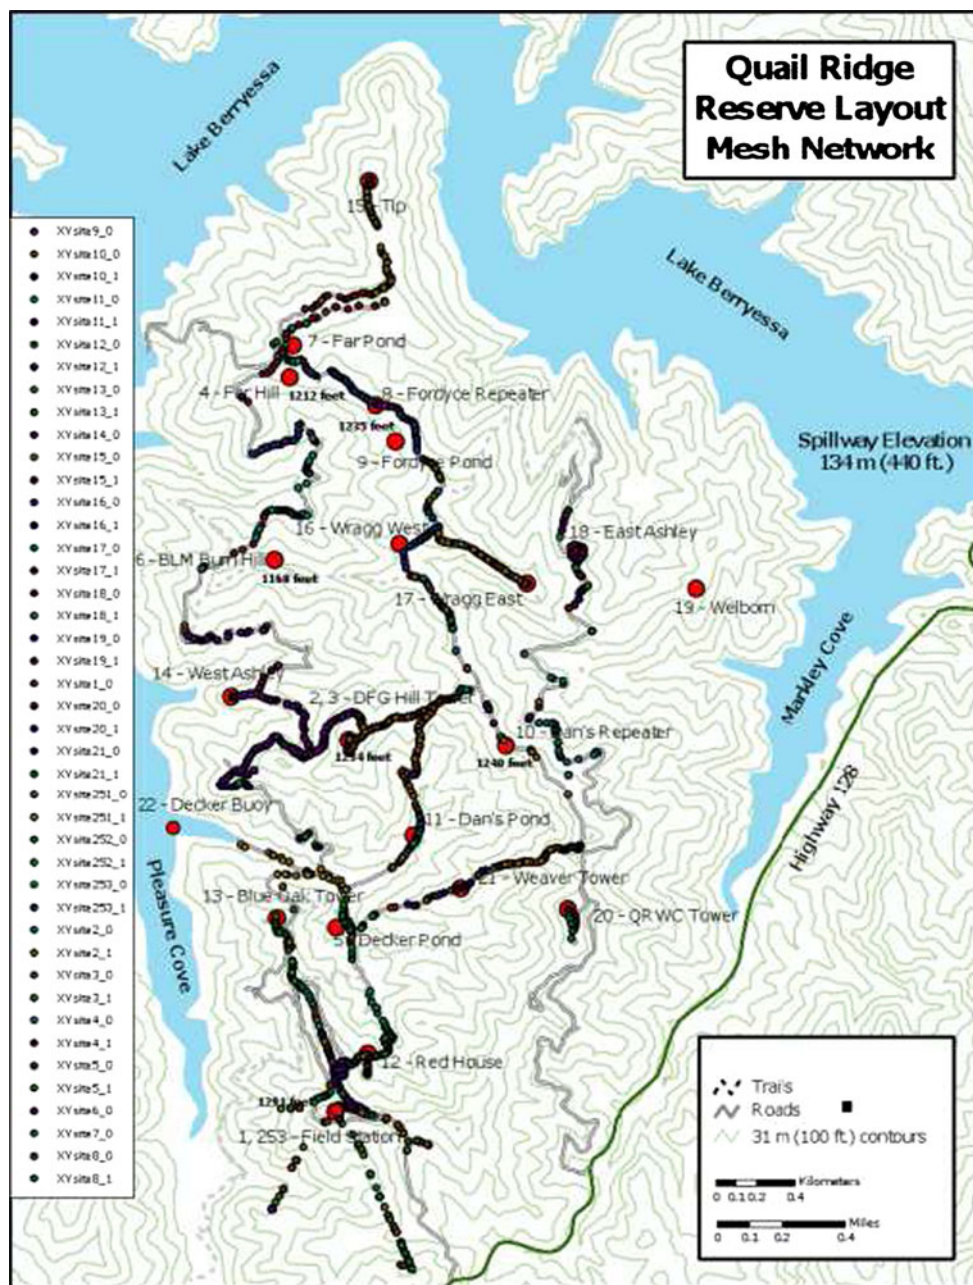
5.2.1 Connection quality

Our first step was to evaluate the connection quality of the mesh APs that were selected by our bandwidth based association control scheme. We chose three different locations in our testbed and used a laptop to associate with the mesh APs available in the neighborhood of each location. The client device first uses the RSSI-based scheme, and then our proposed scheme, to select the best AP in the neighborhood. Once the AP was selected, we used a throughput measurement tool (ThruRay) to measure the available bandwidth from the selected mesh AP to the mesh gateway node. Figure 11 shows the corresponding results. The results are averaged over five runs of the above process. From the results, it was confirmed that our bandwidth based association control scheme selects the better AP as compared to the RSSI-based AP selection scheme.

5.2.2 Client performance

We evaluated the end users' performance by introducing new clients into the network at varying locations. The number of new clients was varied from 1 to 5. Traffic from different clients would traverse different number of wireless hops, depending upon the client's location. Each client performs five UDP based CBR transfers between itself and a reference server attached to the mesh gateway. Figure 12 shows the average client throughput and round trip delay. Our proposed association control scheme based on available bandwidth provides a significant improvement over the RSSI-based approach.

Fig. 10 Overlapping AP coverage at QuRiNet. Different colors correspond to different mesh APs



We repeated the above experiments for TCP-based traffic. Each client performs a file transfer operation between itself and the reference server attached to the mesh gateway node. Figure 13 shows the average throughput and round trip delay achieved by various clients at different locations in our testbed. Once again, our proposed association control scheme based on available bandwidth provides significant improvement over the RSSI-based approach.

The primary reason for our approach performing better than the existing RSSI-based mechanism is that signal strength does not accurately reflect the conditions of the network links. Moreover, it is also not a good indicator of

the amount of load on different mesh APs. On the other hand, the available bandwidth metric combines both the link conditions, and the load on the nodes, into a single metric, thereby achieving a much better performance. Another observation that provides credence to our association policy being a better choice is the number of clients that associate with different mesh APs. When using the RSSI-based approach, all the clients were distributed only amongst two mesh APs, as opposed to utilizing four different mesh APs with the bandwidth based scheme. Moreover, it was observed the mesh APs selected by the clients using our proposed scheme had much more channel diversity along their backhaul routes, as compared to the




Fig. 11 Connection quality of mesh APs selected by different schemes

RSSI-based approach. This is a direct implication of clients being spread over a larger number of mesh APs using our scheme, which results in the client traffic being more evenly distributed in the mesh backhaul.

5.2.3 Mobility experiments

The previous results have been presented for stationary clients. In this section, we describe how our scheme caters

to nomadic clients. We consider clients that change their location from time to time, while spending certain amounts of time at each location. As mentioned previously, our scheme periodically estimates the network conditions and provides the clients with the updated information during the client scanning process. The client device compares the performance of the new available APs with the existing AP. If the performance improvement upon re-association will be greater than 5%, only then the client initiates the re-association process. Such an approach also enables us to prevent excessive re-associations in the network.

In order to test our scheme for mobile clients, we perform the following experiment. We setup a client at a particular location and perform five file transfers between the client and a server attached to the mesh gateway. We then move the client to a new location, and perform the same file transfers after a short time interval. The same experiment is repeated at a third location. Figure 14 shows the performance achieved by client at each of the three locations, both using our scheme and the RSSI-based approach. At the first location, both schemes choose the same mesh AP for the client to associate with. It was observed that with the RSSI-based approach, the client continues to remain attached with the first AP, even after it was moved to the new location. On the other hand, with our bandwidth based scheme, the client device re-associated

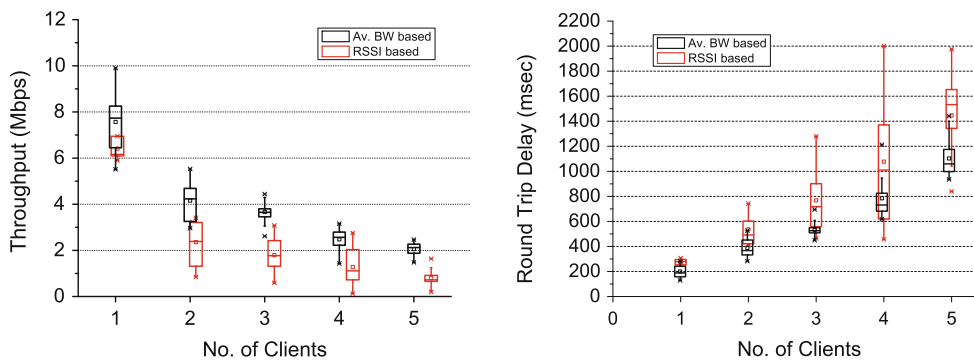
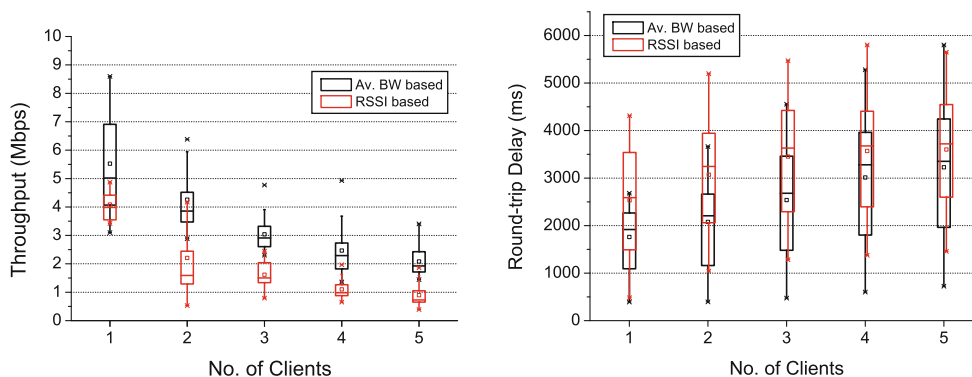


Fig. 12 Client throughput and round-trip delay for UDP CBR traffic. The *box plot* shows the mean, the 1st and 3rd quartile values, and the max and min values

Fig. 13 Client throughput and round-trip delay for FTP traffic



with a new mesh AP after receiving a new set of measurements at the new location.

5.2.4 Load-balancing

As mentioned previously, one of the main disadvantages of using signal strength for client association is that it does not reflect the current load conditions on the mesh APs. This problem is even more severe for wireless mesh networks, as the load on multiple mesh nodes needs to be accounted for. Using the RSSI-based association policy can result in several clients being connected to the same mesh AP, resulting in uneven traffic distribution in the mesh network. In this section, we describe how our scheme helps achieve a degree of load-balancing in the mesh network.

Since the estimation of available bandwidth involves estimating the utilization of various mesh nodes, it helps us in incorporating network load in the association control process. In order to illustrate this, we perform the following experiment. We use two clients that are initially attached to the same mesh AP. One of the clients starts a CBR UDP flow at the fixed rate of 8 Mbps (Fig. 15). This traffic flow roughly simulates the playback of a video stream. A few seconds later, we start a similar traffic flow from the second client. After a certain delay, the first client receives updated information from the neighboring mesh APs. Based on the updated information, client 1 estimates that it will receive better performance if it re-associates with a new mesh AP. The time interval at which the client receives the new network information is a tunable parameter and can be set based on the network settings.

We repeated the above experiments with TCP-based traffic (Fig. 15). We used three clients for this experiment. They were placed at different locations, but were associated with the same mesh AP. Each of the three clients starts a FTP file transfer between itself and the reference server. As the file transfer progresses, the clients receive updated network information from the neighboring mesh APs. Based on this information, Client 1 decides to re-associate with one of the neighboring APs, resulting in improved

network performance. Similarly, after certain time interval, client 2 also re-associates with another AP based on the updated network information it receives. It is clearly visible from these experiments how using available bandwidth can help achieve a degree of load-balancing in the mesh network, resulting in improved performance for the end users.

5.2.5 Client overhead

Apart from evaluating the performance improvement for the client device, we also evaluated the overheads introduced by our scheme, in terms of association time. At three different locations in our testbed, we evaluate the time it takes for the AP probing and selection process. Figure 16 shows the comparison of the time taken by our scheme and the RSSI based scheme. For our scheme, the client device scans for the mesh APs, queries the top three APs for their respective measurement data, and associates with the best AP. For the RSSI-based scheme, the client device simply scans for the available APs and associates with the best AP.

Even though our proposed scheme incorporates some delay in selecting the AP, the increase in delay is not much. Moreover, this slight increase in association delay is compensated for by performance gains in client throughput, as well as by better network load balancing. Moreover, this delay can easily be reduced by utilizing multiple radios at the client device [12], or by creating multiple virtual interfaces on a single wireless card [20]. With such an approach, the client device can either scan for multiple APs at the same time, or use one interface for scanning for new APs in the background, while performing data transfer with the other.

A primary reason for testing the different bandwidth estimation tools on our testbed was that unlike a testbed in a campus or urban scenario, our testbed allows us to perform a comprehensive sensitivity analysis of various parameters involved, due to the lack of interference from external sources. At the same time, there is enough traffic generated in our network to model other networks.

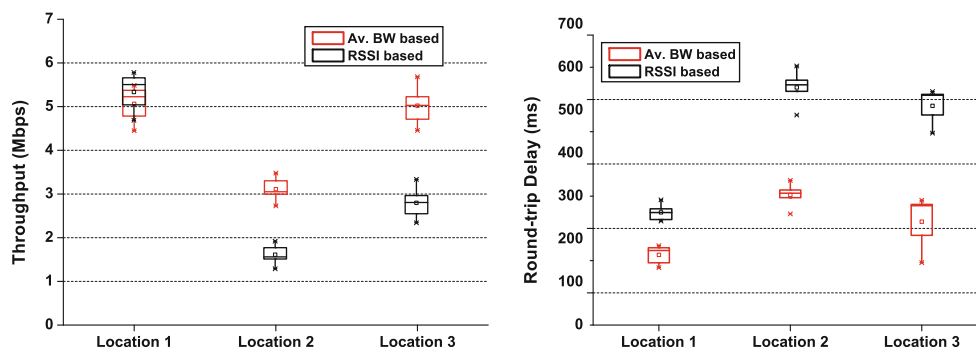


Fig. 14 Performance of nomadic client at different locations

Fig. 15 Load balancing of **a** two CBR UDP flows and **b** three TCP flows

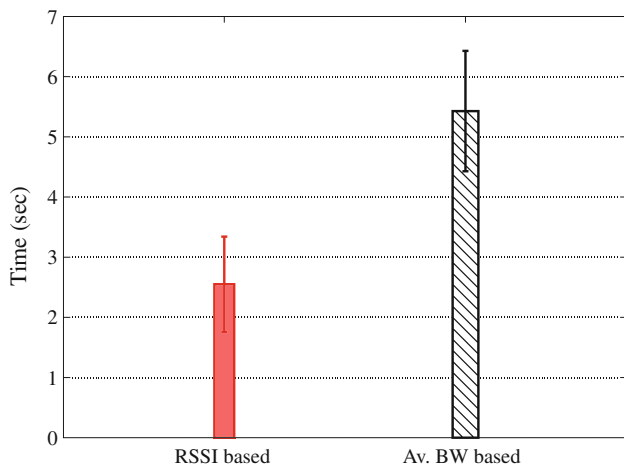



Fig. 16 Average time for each scheme to scan and test the available APs and associate with the best AP

For the perspective of association control, there are two primary considerations for urban scenarios that can affect the performance of our scheme. The first is the impact on the measurement and reporting of necessary network metrics for the mesh network. For our proposed scheme, since the measurement of bandwidth is done in a passive manner, and existing routing control packets are used for transmitting this information, we believe that no extra overheads will be added in this scenario.

The second consideration is that of increased number of clients and its impact on AP selection. Increased number of clients will translate to heavier AP load and increased interference. However, our scheme will automatically account for this during the measurement phase, and will select the APs with the best performance, even though the absolute network performance (in terms of achieved throughput and delay) may be slightly less. Further, frequent changes in the number of clients associated with an AP may result in faster variations in AP load, and hence in the condition of mesh backhaul links. As a result, clients may have to perform AP selection at an increased rate. However, since our scheme is de-centralized, there should not be an increased load on any node in the mesh network.

Moreover, as mentioned in the previous section, utilizing techniques such as multiple radios on client device or multiple virtual interfaces on a single radio, can help eliminate delays and perform faster association control.

6 Related work

Bandwidth estimation: A detailed description of various tools and techniques for estimating available bandwidth in wired networks can be found in [26]. In the realm of wireless networks, some previous works have proposed various passive techniques for bandwidth estimation in wireless networks [33, 42]. The most recent work in this area appears in [48]. However, the proposed schemes have been evaluated only via analytical methods or simulations and lack a comprehensive experimental study. Some other studies [37] have presented limited experiments using the passive bandwidth estimation technique. However, they lack a thorough evaluation of how various factors such as interference, rate-adaptation, and cross-traffic impact the performance of these tools.

Association control: Several previous works have proposed improving the AP selection process for single-hop wireless LANs, by exposing some channel quality metrics to the clients. Such works primarily fall into two categories. The first class of works are client-driven solutions, wherein the clients either passively monitor the wireless channel to estimate metrics such as bandwidth [46], perform a series of probe-based tests to estimate the network performance [39], or submit performance reports of various APs to a database, which can be utilized by future clients [40]. The second class of works involve network-assisted schemes, wherein the APs themselves monitor metrics such as expected throughput or bandwidth, and provide this information to the clients [14, 32, 36, 45]. Other works such as [16, 34] model the AP selection problem as a max-min fairness problem and present heuristics for the same. In [17], authors perform an analytical comparison of selfish (client-based) and centralized (network-assisted)

AP selection algorithms. Works such as [8, 47] propose using the measurement-reporting mechanisms outlined in IEEE 802.11k draft for AP selection. In [13, 19], authors propose dynamically changing the transmission power of the APs (cell-breathing) to control client association.

The work that is closest to ours appears in [11]. Here, the authors propose an AP selection scheme for multi-hop wireless networks, based on the channel airtime metric proposed in IEEE 802.11s protocol. However, their approach has been evaluated via simulations only. Moreover, they consider load-balancing only for the client-AP wireless hop, and not in the mesh backhaul. Another simulation-based approach appears in [10], wherein a game-theoretic approach for association control in mesh networks is proposed that aims to minimize the association cost (maximize association quality) for each client.

7 Conclusion

The wide-spread deployment of wireless mesh networks, and the increasing popularity of Wi-Fi enabled hand-held devices, has resulted in a scenario where users have to choose from a multitude of APs at any given location. This is even more true for users in urban locations and university campuses. The user's performance will be determined by his choice of AP selection. In this paper, we present a new association control scheme, called Seeker, for Wi-Fi enabled devices operating under a large-scale wireless mesh network. We identify the different measurement components that need to be incorporated in an association control scheme. Our scheme makes use of the information available in the mesh backhaul in a passive manner, thereby incurring minimum overheads as compared to active probe-based measurement schemes. We have implemented a passive tool for estimating available bandwidth in wireless mesh networks and utilize this tool to implement our scheme. We have implemented, and evaluated, the proposed scheme on a thirty node outdoor test-bed. Our scheme provides a significant improvement in terms of end-to-end user performance. We further provide results demonstrating the advantages of our scheme for mobile clients, and also for achieving load-balancing in the mesh backhaul.

References

1. IEEE 802.11a standard. (1999). <http://standards.ieee.org/getieee802/download/802.11a-1999.pdf>, <http://standards.ieee.org/getieee802/download/802.11a-1999.pdf>.
2. Optimized link state routing (olsr) protocol. (2007). <http://www.olsr.or>.
3. Orange mesh. (2008). <http://orangemesh.ucdavis.edu>.
4. Quail ridge reserve wireless mesh network. (2008). <http://spirit.cs.ucdavis.edu/quailridg>.
5. AT&T Faces 5,000 Percent Surge in Traffic. (2009). <http://www.internetnews.com/mobility/article.php/384300>.
6. Wifi Google. (2009). <http://wifi.google.co>.
7. Smart phones are making Wi-Fi hotspots hot again. <http://www.physorg.com/news180286896.htm>.
8. Abusubaih, M., & Wolisz, A. (2008). Interference-aware decentralized access point selection policy for multi-rate IEEE 802.11 wireless lans. In: *Proceedings of the IEEE PIMRC*.
9. Aguayo, D., Bicket, J., Biswas, S., Judd, G., & Morris, R. (2004). Link-level measurements from an 802.11b mesh network. In: *Proceedings of ACM SIGCOMM*.
10. Argento, A., Cesana, M., Gatti, N., & Malanchini, I. (2010). A game theoretical study of access point association in wireless mesh networks. *Computer Communications* (In Press, Corrected Proof).
11. Athanasiou, G., Korakis, T., Ercetin, O., & Tassioulas, L., (2009). A cross-layer framework for association control in wireless mesh networks. *IEEE Transactions on Mobile Computing*, 8, 65–80.
12. Bahl, P., Adya, A., Padhye, J., & Walman, A. (2004). Considering wireless systems with multiple radios. *SIGCOMM Comput Commun Rev*, 34, 39–46.
13. Bahl, P., Hajiaghayi, M., Jain, K., Mirrokni, S., Qiu, L., & Saberi, A. (2007). Cell breathing in wireless lans: Algorithms and evaluation. *Mobile Computing*, IEEE Transactions on 6.
14. Balachandran, A., Bahl, P., & Voelker, G. M. (2002). Hot-spot congestion relief and service guarantees in public-area wireless networks. *SIGCOMM Comput Commun Rev*, 32, 59.
15. Balasubramanian, N., Balasubramanian, A., Venkataramani, A. (2009). Energy consumption in mobile phones: A measurement study and implications for network applications. In: *IMC '09: Proceedings of the 9th ACM SIGCOMM conference on internet measurement conference*.
16. Bejerano, Y., Han, S. J., & Li, L. (2007). Fairness and load balancing in wireless lans using association control. *IEEE/ACM Trans Netw* 15.
17. Berg, M., & Hultell, J. (2006). On selfish distributed access selection algorithms in IEEE 802.11 networks. In: *Proceedings of IEEE VTC*.
18. Bredel, M., & Fidler, M. (2008). *Lecture notes in computer science*, Springer, chap A Measurement study of bandwidth estimation in IEEE 802.11g wireless LANs using the DCF, pp. 314–325.
19. Brickley, O., Rea, S., & Pesch, D. (2005). Load balancing for qos optimisation in wireless lans utilising advanced cell breathing techniques. In: *Proceedings of the IEEE VTC*.
20. Chandra, R., & Bahl, P. (2004). Multinet: Connecting to multiple IEEE 802.11 networks using a single wireless card. In: *Proceedings of IEEE Infocom*.
21. Couto, S. J. D., Aguayo, D., Bicket, J., & Morris, R. (2005). A high-throughput path metric for multi-hop wireless routing. *Wireless Networks*, 11, 419–434.
22. Croce, D., Mellia, M., & Leonardi, E. (2009). The quest for bandwidth estimation techniques for large-scale distributed systems. In: *Second Workshop on Hot Topics in Measurement & Modeling of Computer Systems (HotMetrics)*.
23. Draves, R., Padhye, J., & Zill, B. (2004). Routing in multi-radio, multi-hop wireless mesh networks. In: *Proceedings of ACM MobiCom*, Philadelphia, PA.
24. Feeney, L., & Nilsson, M. (2001). Investigating the energy consumption of a wireless network interface in an ad hoc networking environment. In: *Proceedings of IEEE Infocom*.
25. Gupta, A., & Mohapatra, P. (2007). Energy consumption and conservation in wifi based phones: A measurement-based study. In: *Sensor and Ad Hoc Communications and Networks (SECON)*.

26. Gupta, D., Wu, D., Mohapatra, P., & Chuah, C. N. (2009). Experimental comparison of bandwidth estimation tools for wireless mesh networks. In: *The 28th IEEE International Conference on Computer Communications (INFOCOM) mini-conference*.
27. Hu, N., & Steenkiste, P. (2003). Evaluation and characterization of available bandwidth probing techniques. *Selected Areas in Communications, IEEE Journal on*.
28. Jain, M., & Dovrolis, C. (2002). Pathload: A measurement tool for end-to-end available bandwidth. In: *Passive and Active Measurements workshop (PAM)*.
29. Jones, K., & Liu, L. (2007). What where wi: An analysis of millions of wi-fi access points. In: *Proceedings of IEEE PORTABLE*.
30. Judd, G., & Steenkiste, P. (2002). Fixing 802.11 access point selection. *Computer Communication Review* 32.
31. Lakshminarayanan, K., Padmanabhan, V., & Padhye, J. (2004). Bandwidth estimation in broadband access networks. In: *The 4th ACM SIGCOMM Internet Measurement Conference (IMC)*.
32. Lee, H., Kim, S., Lee, O., Choi, S., & Lee, S. J. (2008). Available bandwidth-based association in ieee 802.11 wireless lans. In: *Proceedings of ACM MSWiM, Vancouver, Canada*.
33. Lee, H. K., Hall, V., Yum, K. H., Kim, K., & Kim, E. (2006). Bandwidth estimation in wireless lans for multimedia streaming services. In: *IEEE ICME*.
34. Lu, M., & Wu, J. (2007). Localized access point selection in infrastructure wireless lan. In: *Proceedings of IEEE MILCOM*.
35. Mangold, S., & Berlemann, L. (2005). Ieee 802.11k: Improving confidence in radio resource measurements. In: *Proceedings of IEEE PIMRC*.
36. Murty, R., Padhye, J., Chandra, R., Wolman, A., & Zill, B. (2008). Designing high performance enterprise wi-fi networks. In: *Proceedings of USENIX NSDI*.
37. Neilsen, M., Ovsthus, K., & Landmark, L. (2006). Field trials of two 802.11 residual bandwidth estimation methods. In: *IEEE MASS*.
38. Nicholson, A. J., & Noble, B. D. (2008). Breadcrumbs: Forecasting mobile connectivity. In: *Proceedings of the ACM MobiCom, San Francisco, CA*.
39. Nicholson, A. J., Chawathe, Y., Chen, M. Y., Noble, B. D., & Wetherall, D. (2006). Improved access point selection. In: *Proceedings of the ACM MobiSys, New York, NY, USA*.
40. Pang, J., Greenstein, B., Kaminsky, M., McCoy, D., & Seshan, S. (2009). Wifi-reports: Improving wireless network selection with collaboration. In: *Proceedings of the ACM MobiSys*.
41. Ribeiro, V. J., Riedi, R., Baraniuk, R., Navratil, J., & Cottrell, L. (2003). Pathchirp: Efficient available bandwidth estimation for network paths. In: *Passive and Active Measurements Conference (PAM)*.
42. Sarr, C., Chaudet, C., Chelius, G., & Lassous, I. G. (2008). Bandwidth estimation for ieee 802.11-based ad hoc networks. *IEEE Transactions on Mobile Computing*, 7, 1228–1241.
43. Shriram, A., Murray, M., Hyun, Y., Brownlee, N., Broido, A., Fomenkov, M., & Claffy, K. (2005). Comparison of public end-to-end bandwidth estimation tools on high-speed links. In: *Passive and Active Measurements Conference (PAM)*.
44. Strauss, J., Katabi, D., & Kaashoek, F. (2003). A measurement study of available bandwidth estimation tools. In: *Proceedings of 3rd IMC*.
45. Sundaresan, K., & Papagiannaki, K. (2006). The need for cross-layer information in access point selection algorithms. In: *Proceedings of the IMC*.
46. Vasudevan, S., Papagiannaki, K., Diot, C., Kurose, J., & Towsley, D. (2005). Facilitating access point selection in ieee 802.11 wireless networks. In: *Proceedings of the IMC*.
47. Villegas, E. G., Ferre, R. V., & Aspas, J. P. (2006). Load balancing in wlans through ieee 802.11k mechanisms. In: *Proceedings of the IEEE ISCC*.
48. Zhao, H., Garcia-Palacios, E., Wei, J., & Xi, Y. (2009). Accurate available bandwidth estimation in ieee 802.11-based ad hoc networks. *Computer Communications*, 32, 1050–1057.

Author Biographies



Dhruv Gupta is currently working at AT & T Labs Inc. He received his Doctoral degree in Computer Science from University of California Davis in 2010. His research interests include monitoring and management of wireless mesh networks.



Prasant Mohapatra is currently a Professor at and Chair of the Department of Computer Science at the University of California, Davis. In the past, he has been on the faculty at Iowa State University and Michigan State University. He has also held Visiting Scientist positions at Intel Corporation, Panasonic Technologies, Institute of Information Research (I2R), Singapore, and National ICT Australia (NICTA). Dr. Mohapatra received his Ph.D. in Computer Engineering from the Pennsylvania State University in 1993. He was/is on the editorial board of the IEEE Transactions on Mobile Computing, IEEE Transactions on Computers, IEEE Transactions on Parallel and Distributed Systems, ACM WINET, and Ad Hoc Networks. He was the Program Vice-Chair of INFOCOM 2004, Program Co-Chair of QSHINE 2006, and the Program Co-Chair of the First IEEE International Conference on Sensor and Ad Hoc Communications and Networks (SECON 2004). He has been a Guest Editor for IEEE Network, IEEE Transactions on Mobile Computing, IEEE Communications, and the IEEE Computer. Dr. Mohapatra's research interests are in the areas of wireless mesh and ad hoc networks, sensor networks, Internet protocols and QoS. Dr. Mohapatra's research has been funded through grants from the National Science Foundation, Department of Defense, Intel Corporation, Siemens, Panasonic Technologies, Hewlett Packard, Raytheon, and EMC Corporation.



Chen-Nee Chuah (S'92, M'01, SM'06) is currently an Associate Professor in the Electrical and Computer Engineering Department at the University of California, Davis. She received her B.S. in Electrical Engineering from Rutgers University, and her M.S. and Ph.D. in Electrical Engineering and Computer Sciences from the University of California, Berkeley in 1997 and 2001, respectively. Before joining UC Davis, she spent 9 months as a

visiting researcher at Sprint Advanced Technology Laboratories. Her

research interests lie in the area of computer networking and distributed systems, ranging from Internet measurements, network management, overlay/peer-to-peer systems, network security, wireless/mobile networking, to opportunistic communications. Chuah received the NSF CAREER Award in 2003 and the UC Davis College of Engineering Outstanding Junior Faculty Award in 2004. She has served on the technical program committee of several ACM and IEEE conferences and workshops. She received the ACM Recognition Award for co-organizing the First Workshop on Vehicular Ad Hoc Network (VANET), which was held in conjunction with ACM MobiCom 2004. She also served as the TPC vice-chair for IEEE GLOBECOM 2006.