

Resource Estimation on Wireless Backhaul Networks

(Invited Paper)

Irfan Sheriff, Prashanth Aravinda Kumar Acharya, and Elizabeth M. Belding

Department of Computer Science, University of California, Santa Barbara
{isheriff, acharya, ebelding}@cs.ucsb.edu

ABSTRACT

The increased usage of IEEE 802.11 wireless backhaul networks and the growing popularity of real time applications, such as VoIP, presents a challenging resource management problem due to the limited capacity of wireless networks. At high traffic volumes, measurements have shown that packet collisions and interference in 802.11 networks can lead to degraded performance to the extent that users experience unacceptably low throughput, which can ultimately lead to complete network breakdown [1]. A resource management framework that limits network flows can prevent network breakdown and improve the performance of throughput and delay-sensitive multimedia applications. To address this problem, we present a measurement-driven framework that leverages wireless characteristics for intelligent admission control in a static wireless network. Experiments on a 25 node wireless testbed show that the proposed scheme can enhance network performance such that the QoS requirements of real time applications, such as VoIP, can be met.

1. INTRODUCTION

IEEE 802.11 networks have gained attention as an alternative to wired networks to provide low cost backhaul infrastructure in both office and outdoor environments. However, these networks are not robust to large traffic volumes. Complete network breakdown due to excessive traffic has been observed in real deployments [1]. The lack of a framework that can effectively limit flows hinders the ability of these networks to support the QoS requirements of real time applications.

In general, estimation of available resources in IEEE 802.11 networks is a challenging problem due to the presence of multiple factors that introduce unpredictability into the system. Node mobility can introduce drastic variations in the link behavior; the traffic usage could exhibit significant variations and the network itself is susceptible to external interference. However, there are multiple deployment scenarios, expected to be widely used, to support real time traffic with predictable characteristics that isolate the impact of client mobility. A typical office wireless network [2] for desk-bound users or a two-tier mesh network deployment [3] with the clients operating on a separate frequency spectrum to support real time applications, such as VoIP, are two instances of such networks. Our goal is to show the feasibility of resource estimation in these networks.

There has been significant research on admission control for 802.11-based wireless networks. Yet there lacks a realistic system implementation to limit traffic for real

time application support in these networks. This gap between proposed solutions and their actual deployment is, in many cases, due to the existence of unrealistic assumptions that render a system implementation infeasible.

In this paper, we present Measurement-Driven Admission Control (MDAC), a framework that leverages measurements for resource estimation in wireless networks. The framework uses measured average signal strength, packet loss and data rate values to characterize the behavior of the wireless channel. MDAC continuously monitors the resource availability in the form of the time fraction for which the wireless medium is free. This information is then used in the decision-making process for flow admission control.

Our work makes a case for measurement based admission control in wireless networks. We show that it is feasible to build a distributed framework that can draw inferences about the underlying wireless conditions from locally observable measurements. While we use the framework for resource estimation, we believe that the observations have broader implications on the design of more realistic algorithms for 802.11 wireless networks.

The remainder of the paper is organized as follows. Section 2 lists the assumptions and terminology used in the paper. We present our findings about the characteristics of wireless links in Section 3. Section 4 describes the design of the admission control scheme. We present the implementation details and results in Sections 5 and 6 respectively. In Section 7, we contrast our work with existing literature and, finally, we conclude in Section 8.

2. ASSUMPTIONS AND TERMINOLOGY

In this section, we first describe the network attributes for which we design our admission control scheme. Then, we define some of the common terms used in the paper.

2.1 Network attributes

Our admission control scheme is designed with certain network attributes in mind. For instance, we design our admission control solution for the generic setting of static multihop wireless networks such as backhaul mesh networks and static ad hoc networks. As a first step, we focus on the problem of admission control in a simplified network scenario that consists of traffic between single-hop neighbors only. However, much of the discussion in the paper is generic to networks with multihop flows.

2.2 Terminology

In this paper, we use the term ‘interference’ to refer to the impact of concurrent transmission on packet receptions. An interfering neighbor thus refers to a node that can affect the successful packet reception at another node. We use the term ‘carrier sensing’ to refer to sensing the medium to assess whether it is free. The term ‘busy-time’ refers to the time duration for which the channel is not free. A carrier sensing neighbor is thus a node that causes busy-time at a given node X and, therefore, can cause deferral of the data transmissions at X. The term Received Signal Strength Indicator (RSSI) refers to the signal strength of a packet (in dBm) relative to the noise floor at the radio. The default value of the noise floor for Atheros radios is -95dBm. Note that the above definition of RSSI is specific to Atheros radios. The term Received Signal Strength (RSS) refers to the absolute energy level of a packet and is measured in dBm. The RSS value of a packet can be computed using the RSSI and noise values reported by the Atheros radio.

3. UNDERSTANDING WIRELESS CHARACTERISTICS

Many previously proposed admission control solutions are based on a common set of assumptions that may not hold true in real-world 802.11 networks. In the following sections, we evaluate the accuracy of a few such well-known assumptions. We then incorporate the insights obtained from these studies in the design of MDAC described in Section 4.

3.1 Busy-time

A common assumption of many existing admission control techniques is that the impact of a new flow on the busy-time of the surrounding nodes is binary. In other words, in the simplest case, the packet transmission by a node has no effect on the busy-time at nodes outside the transmitter’s carrier sense range, while the increase in busy-time at nodes within the transmitter’s carrier sense range is equal to the duration of the packet transmission. Hence, the increase in busy-time at the carrier sense neighbors is equal to the transmission time of a newly admitted flow. To investigate whether this is indeed the case, we perform the following set of experiments. These experiments are conducted during the night to minimize the effect of interference from co-located WLANs.

We use the reverse-engineered Open HAL¹ implementation of the MadWifi driver to understand the relationship between busy-time and communication range. Atheros maintains register counters to track the ‘medium busy time’ and the ‘cycle time’. The cycle time counter is incremented at every clock tick of the radio and the medium busy counter represents the clock tick count for which the medium was sensed busy. The ratio of these two counters thus represents the busy-time fraction. Note that the

¹<http://madwifi.org/wiki/OpenHAL>




Figure 1: Packet reception versus busy fraction.

busy-time fraction includes the time spent for transmission and reception of packets.

In our experiment, a receiver node is placed at a fixed location. A sender node is placed at different locations to vary the packet reception rate at the receiver. Keeping the receiver stationary ensures that the variation in environmental noise at the receiver is minimal. The sender transmits fifty 100 byte broadcast packets per second at 1Mbps. The receiver tracks the packet reception rate and also estimates the busy-time fraction caused by the sender. The busy-time is estimated as the difference in busy-time values reported by the radio before and after the experiment. To obtain the busy-time fraction, the measured value of busy-time is normalized with respect to the difference in the reported busy-time when the sender is close to the receiver with 100% packet reception. The experiment lasts 30 seconds at each sender location.

Figure 1 shows a plot of the calculated busy-time fraction at the receiver versus the fraction of packets received by the receiver for different locations of the sender node. The graph shows that medium busy-time fraction of a node on its neighbor ranges anywhere from 0 to 1 and is not a binary relationship. For example, when 40% of the transmitted packets are received, the busy-time fraction varies from 0.4 to 0.5.

The experimental results imply that, in a larger network, the bandwidth consumption of a flow on each carrier sense neighbor is not 100% of the bandwidth requirement of the flow. On the contrary, the impact is likely to vary based on the location of the neighbor, as well as other environmental factors. Thus, existing admission control schemes that assume a new flow will consume 100% of the bandwidth requirement can be highly pessimistic and are likely to overestimate the resource requirements. We incorporate the above result into MDAC by computing the fractional increase in busy-time relative to the transmission rate of the sender at the carrier sensing neighbors of the transmitting node.

3.2 Communication with carrier sensing neighbors

Many existing admission control techniques require sender nodes to communicate with nodes outside of their reception range, but within their carrier sensing range.

These solutions often assume the carrier sense range is about twice the transmission range [4, 5] and, therefore, direct communication with carrier sense neighbors is not possible. Some techniques use high power transmissions or multi-hop forwarding to communicate with carrier sense neighbors [4]; others propose techniques to adjust a node's carrier sense range so that it can hear the transmissions of potentially interfering neighbors [6]. However, recent results show that the simple assumption of carrier sense range being twice the transmission range does not hold in real-world networks [7]. Based on the observations in Section 3.1, we are interested in understanding whether nodes that cause transmission deferral at each other can communicate, and if so, the extent of communication achievable between these carrier sensing neighbors.

We again refer to Figure 1 to understand the possibility of communication between carrier sensing neighbors. The graph shows that the receiver node is able to receive a non-zero number of packets from a transmitting neighbor that induces more than 5% of medium busy fraction. The probability of receiving a packet from a neighbor increases with the increase in busy-time the neighbor induces at the receiver node. For example, a node can receive about one of five packets from a neighbor that induces 40% busy-time and about three of five packets from a neighbor that induces 80% busy-time.

The same experiment repeated at a higher data rate would result in much lower packet reception rates for similar values of busy-time fraction. Additionally, the reception capability of a radio depends on its sensitivity, thermal noise level and other environmental factors. However, multiple experiments with different off-the-shelf radios showed that the packet reception at 1Mbps is achievable at low busy-time values for a majority of the radios.

To understand the reason for successful packet reception inside carrier sense range, we refer to the datasheet for the Atheros radios. The data sheet states that a packet sent at 1Mbps can be successfully received when the RSSI is as low as zero. In contrast, the Clear Channel Assessment (CCA) threshold for Atheros radios is reported to be about -81dBm, which corresponds to an RSSI of 14 (assuming the noise floor to be the default -95dBm) [8]. The CCA threshold is used to determine whether the channel is busy. The significant difference between the CCA threshold and the minimum RSSI required for packet reception at 1Mbps indicates that nodes can communicate with carrier sensing neighbors using low data rate packets. CCA values for Prism radios have been reported to be as high as 40dB above the noise floor².

Based on the above results, we have shown that a node can communicate with all the carrier sensing neighbors that have non-negligible busy-time impact on it. The different values of packet reception indicate that this communication may be unreliable. In MDAC, we propose to use this unreliable reception of low data rate packets to periodically communicate the resource availability in-




Figure 2: RSS-packet reception piecewise linear curves.

formation of a node to its carrier sensing neighbors. We add sufficient redundancy in the communication to ensure that, with a high probability, all the carrier sense neighbors that have non-negligible busy-time impact receive resource availability information.

3.3 The RSS-packet reception relationship

We now study the relationship between the packet reception rate and the average RSS of the received packets at different data rates. Reis et al. [8] study the behavior of wireless links in static networks and show that the packet reception probability is a function of the RSS of the received packets, and this relationship is specific to each node in the network. They use RSS values of received probe packets to derive a mathematical function in the form of a piecewise linear curve that models the reception probability at different RSS values³. We use a similar measurement-based model for characterizing packet reception rates in our testbed. The model in [8] measures and predicts link quality using the RSS values of packets sent at the broadcast rate. However, the SINR (Signal to Interference-Noise Ratio) requirements for packet reception are different for each of the supported 802.11 data rates. For example, the SINR required for packet reception at the 54Mbps data rate is much higher (≈ 24 dB) than that at 6Mbps (≈ 6 dB) [9]. Therefore, link quality predictions based solely on measurements that use broadcast packets may suffer from significant inaccuracies.

The naïve approach to rectify this problem is to send probe packets at all possible rates. This approach, however, is highly inefficient and causes excessive load on the network. An alternative approach is to use the existing unicast traffic in the network to measure the link quality at different rates. To this effect, we note that the rate selection algorithm SampleRate [10] maintains statistics about the average number of packet retransmissions at different rates between each neighbor pair. Thus, the traffic utilized by SampleRate is an ideal candidate on which to piggyback the link quality measurements at different

²<http://pdos.csail.mit.edu/roofnet/doku.php?id=interesting>.

³Please refer to [8] for details of the model.

rates. An advantage of integrating the link quality measurement with SampleRate is that SampleRate maintains up-to-date statistics on all the rates that a node uses or is likely to use. Also, SampleRate does not send frequent probe packets for data rates that it is not likely to use. Therefore, this integration provides a scalable method of link quality estimation at different data rates.

We modify the MadWifi driver to maintain statistics about the RSS of each received packet on a per-rate basis for each neighbor. Each node periodically exchanges its SampleRate statistics with its neighbors. Each node then computes the RSS-packet reception relationship for each rate. We deploy the modified MadWifi driver throughout our testbed and each node obtains the RSS-packet reception piecewise linear curves for different data rates.

Figure 2 shows the RSS-packet reception at different rates for an example node in the testbed. Due to lack of space, we present data for a single representative node while the inferences below are based on data for the entire testbed. We observe from the graph that each data rate exhibits significantly different characteristics for the same RSS value. The graph also shows the different RSS requirements for the 802.11b and 802.11g rates that result from the difference in the underlying physical layer technology. Another observation from the graph is that there does not exist an obvious correlation between the curves for different rates. Therefore, extrapolating link quality based only on broadcast packets leads to errors in prediction. Note that some nodes in the network do not have packet reception data for all possible rates. This gap is because SampleRate measures link quality only for those data rates that it is likely to use, which may be a subset of all the available data rates. Because these data rates are not likely to be used for transmission, the absence of this data does not affect the link quality predictions.

From the above result we learn the importance of measuring the link quality at each data rate individually. The implementation of MDAC uses the modified MadWifi driver. The RSS-packet reception curves obtained from the driver are then used to estimate the impact of a new flow on packet reception.

3.4 Temporal behavior

Previous work on measurement of packet receptions in static wireless networks has shown that, although packet losses occur in bursts, the quality of a majority of links in the network is stable when measured over long time intervals (on the order of tens of seconds) [8, 11]. This observation implies that the packet reception rate of a link measured during one time-window of suitable granularity can be used to accurately predict the packet reception probability during the subsequent time-window.

We explore the degree of stability of the busy-time metric and the appropriate time scale to measure and predict the busy-time metric. We conduct experiments similar to those in Section 3.1 for all the links in the testbed to measure the packet reception and corresponding busy-time. In this case, however, the experiments are con-

ducted for one hour during various times of the day. We then compute Pearson’s linear correlation coefficient to study the relationship over time between packet reception rate and busy-time. We find that the two entities are highly correlated for all the links with 0.895 as the median value of the correlation coefficient. This implies that, similar to packet reception, the busy-time behavior of links can also be predicted for appropriate time scales (on the order of tens of seconds). However, the predictability of busy-time is affected by atypical events such as operation of a microwave.

4. ADMISSION CONTROL

We build on the results from Section 3 and present the design of our Measurement-Driven Admission Control (MDAC) system. The use of MDAC results in the reduction of network congestion that is likely to be created by the unrestricted admission of flows.

To perform admission control, we consider the bandwidth allocation, delay and jitter metrics. By controlling bandwidth allocation, delay and jitter can also be controlled [4]. Therefore the metric of primary interest for MDAC is the available bandwidth. In a shared wireless medium, the two factors that determine the bandwidth available on a link are 1) the fraction of time for which the medium is free and hence available for transmissions, and 2) the data rate used for transmissions. Each node has the transmission data rate information available locally. However, the fraction of time for which the medium is free must be computed.

We next present a discussion of 802.11 node behavior and interaction with neighbors. An understanding of the node interaction with neighbors is essential to identify the factors that affect the packet transmissions and receptions in the network upon the admission of a new flow.

4.1 Node behavior

An 802.11 wireless node can be in any one of the following states: TX, RX, ChannelBusy and Free. We represent the fraction of time spent by a node in each of these states as t_{TX} , t_{RX} , t_{CB} and t_F , respectively. TX and RX represent the transmit and receive states. ChannelBusy represents the duration for which a radio cannot attempt a packet transmission. Thus t_{CB} includes the time during which the channel is busy (\bar{t}_{CCA}), and the additional time the 802.11 MAC spends for the DIFS, SIFS and backoff periods (t_{MAC}). Note that $\bar{t}_{CCA} = t_{CCA} - t_{RX} - t_{TX}$, since t_{CCA} , the fraction of time the Clear Channel Assessment (CCA) indicates the channel to be busy, includes the transmission and receive time duration. In short, t_{CB} constitutes the unusable fraction of medium time that could be due to any of: channel noise, thermal noise, neighbor-to-neighbor packet transmissions or the MAC overhead. The remainder of the time constitutes t_F . The Free state time fraction, t_F , can be calculated as $t_F = 1 - (t_{CCA} + t_{MAC})$.

To understand the impact of a new flow in the network, we need to examine its effect on both the TX and RX

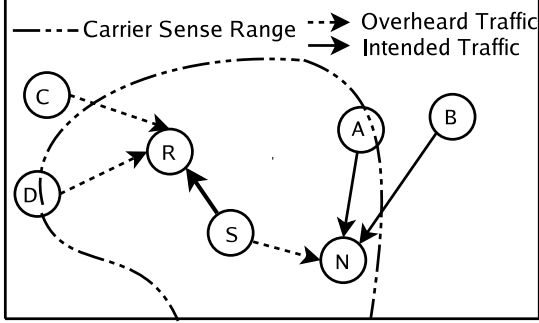


Figure 3: A node pair S - R admitting a flow.

states at a node. If the medium time consumption of the new flow is higher than the available time t_F at a node, the t_{TX} time at the node could be restricted. This behavior is due to the physical carrier sense mechanism of the IEEE 802.11 MAC protocol. Packets are transmitted only when medium is sensed free. The implication of this behavior is that packet transmissions can be throttled due to an increase in the ChannelBusy time, t_{CB} , caused by a new flow.

The impact of a new flow on the RX state depends on the signal strength of the current receptions. An increase in channel noise or neighbor packet transmissions decreases the reception capability of a node. Packet losses during RX are likely to result in retransmissions from the sender, leading to an overall increase in t_{RX} at the node.

Thus there are two factors we consider important at every node during admission control: sufficient t_{TX} to allow packet transmissions as a sender, and sufficient tolerance to interference to limit the packet loss as a receiver.

In the remainder of the paper, we define Busytime as the fraction of time spent in the TX, RX or ChannelBusy states. Therefore Busytime includes the additional MAC overhead caused by the silence periods due to DIFS, SIFS and backoff. We can represent Busytime as:

$$\text{Busytime} = t_{CCA} + t_{MAC} \quad (1)$$

Note that throughout the paper, we have used the term ‘busy-time’ to refer to t_{CCA} .

4.1.1 Example scenario

We now explain neighbor interaction with the help of an example. Consider the example shown in Figure 3, where a node pair S - R needs to admit a new flow in the network. For now, consider that the flow only consists of data packets in one direction from S to R ; we postpone the discussion on the impact of ACKs until later in this section. Assume that the flow requires t_{flow} fraction of the medium time. We examine the impact of the new flow on the different nodes in the network.

Once the new flow is admitted, t_{TX} at node S increases by t_{flow} . If there is sufficient free medium time at node S ($t_F > t_{flow}$), the impact on existing transmissions at S is negligible. However, without sufficient medium time ($t_{flow} > t_F$), packet transmissions at S could be throttled with the admission of the flow.

A similar argument holds for node N , a carrier sensing neighbor of node S . As a carrier sensing neighbor, the Busytime at node N increases by up to t_{flow} with the admission of the new flow. The impact of this increase depends on the available medium time t_F at N . Without sufficient t_F , the existing packet transmissions could be throttled at N .

Now, suppose that nodes A and B both have an existing flow to node N . The reception of packets at node N from senders A and B can be affected due to the collisions caused by the overheard traffic from node S . Note that node A is partially outside the carrier sense range of S and node B is completely outside the carrier sense range. Thus, nodes S and A share the medium and node B is a hidden terminal to node S . The rate of collisions is thus expected to be higher on transmissions from node B , assuming other factors such as data rate, packet size and packet rate to be the same.

Packet collisions result in retransmissions from A and B . These retransmissions increase t_{RX} at node N and t_{CB} at other neighbors of A and B . This can cause a potential congestion *ripple effect*. The reasoning is as follows. Typical response of a rate selection scheme to packet loss is to reduce the data rate for more robust transmissions [10]. When packets from nodes A and B experience collisions due to transmissions from node S , lower data rates are used. This leads to a further increase in t_{RX} at N and t_{CB} at other neighbor nodes of A and B . It also results in the packets from nodes A and B being more susceptible to overlap with other packets in the medium because the lower rate traffic consumes additional medium time. The impact of collisions can thus quickly spread throughout the network. Thus, it is important to control the collisions at the neighbors of N to prevent such a ripple effect.

Ensuring that the existing flows are not affected by the new flow does not suffice. Suppose nodes C and D are senders with existing flows to other nodes in the network (not shown in the figure). Node R is a neighbor of both nodes C and D and can overhear their transmissions. At node R , the new flow from S causes an increase in t_{RX} . This increase in t_{RX} makes the packets received at node R , including packets of the new flow, more susceptible to collisions caused by transmission from nodes C and D . In particular, the packets are more susceptible to collisions from hidden terminals such as node C . The above discussion about the impact of collisions at node N is also applicable in this case, with node R as the receiver. In this case, the transmissions are from neighbor nodes C and D , and the intended traffic is from node S .

In summary, ensuring sufficient t_{TX} to allow neighbor packet transmissions and ensuring the loss imposed on neighbor packet receptions is minimal involves predicting the increase in Busytime and the increase in packet collisions that can result with admission of a new flow.

IEEE 802.11 requires transmission of an ACK for every data packet received. Due to the smaller packet size, the Busytime impact of ACK is expected to be much less

than that of data traffic. However, in case of VoIP traffic, we cannot ignore the impact of ACKs due to the comparable packet transmission durations [12].

The packet collision impact of ACK packets does not depend on the carrier sense range of the node R . This is because ACKs are always sent after a SIFS wait upon a data reception. This is in contrast to a data packet transmission that occurs after carrier sensing that the channel is free. This property of transmission of ACKs implies that R acts as a hidden terminal to nodes outside the carrier sense range of the node S . This is because neighbors of the node S that can overhear the data packet set the NAV value so as to not collide with the ACK packet.

We next present the design and operation of MDAC.

4.2 Operation

At a high-level, MDAC decides whether to admit a new flow into the network based on two considerations: first, the node and its neighbors have sufficient resources (medium time) to accommodate the new flow; and second, the admission of the new flow does not negatively affect the performance of existing flows. The two considerations of the scheme are evaluated with the help of five operational steps outlined below. Each of these steps is discussed in detail in the following sections.

1. Each node in the network maintains a profile for each of its neighbors. The neighbor profile includes statistics of average RSS, packet reception rates, number of packets received per second, busy-time, and link data rates. The neighbor profile is generated and maintained using either broadcast probes or the data traffic transmitted by the nodes. In addition to a neighbor profile, nodes also maintain statistics about the consumption of local resources, such as the Busytime, the average number of packets sent and received per second, and the average data rate.
2. A node that wishes to admit a new flow estimates the availability of sufficient local bandwidth for packet transmissions. The available local bandwidth is predicted using the local Busytime information.
3. The node then estimates whether the new flow will throttle on-going transmissions at any of its neighbors. The Busytime information from the neighbor profile and the Busytime impact of the new flow is used for the decision.
4. Next, the node estimates whether the packet collisions caused by the new flow at any of the neighboring nodes exceed a threshold. The Busytime and packet statistics from the neighbor profile of the two-hop neighborhood are used for this decision.
5. Finally, the node estimates whether the packet collisions caused by the existing traffic on the new flow exceed a threshold. This decision is based on the Busytime from the neighbor profile and packet statistics collected at the receiver node.

A new flow is admitted into the network only if all of the above decision steps indicate that the new flow will not lead to congestion in the network. We next discuss each step of the admission control operation in turns.

4.2.1 Profiling the neighbors

Neighbor profiling involves collection of statistics required for the admission control decision. The following statistics about the neighborhood are essential: RSS and packet reception at each data rate; packets received per second at each data rate; data rate usage on links; and the Busytime impact of a node on its neighbors.

In order to minimize the network overhead, we collect all our measurement data from the observed data traffic in the network. The MadWifi driver was modified to track the RSS of the received packets as well as the count of packets received and sent different data rates from all the neighbors. This data is aggregated and periodically communicated to the neighbors. The communication of the aggregated data occurs with a periodic broadcast packet. MDAC needs neighbor statistics from a two-hop neighborhood. The one-hop neighbor statistics obtained during the current cycle of broadcast are, therefore, piggy-backed in the next cycle of packet broadcast.

With the received packet count statistics from neighbors, nodes can determine the packet reception rate at different data rates to each neighbor. Nodes can then construct piecewise linear curves at each data rate using the reception rate and RSS statistics. Nodes also compute the average medium time consumption from the average RSS statistics using the PHY Deferral model proposed by Reis et al. [8]. Reference [8] provides a detailed description of the PHY Deferral model.

To summarize, we maintain the following information at each node in the network:

- RSS and reception statistics per neighbor averaged over a time window. We represent the average RSS from a node S at R as \overline{R}_{SR} .
- Piecewise linear curves representing the relationship between RSS and data reception at different data rates. Piecewise curves predict the reception rate at a certain RSS for a given data rate at a node. At a node R , we represent this function of \overline{R}_{SR} and data rate d as $\hat{p}_R(\overline{R}_{SR}, d)$.
- Packets received per second and data rate statistics of the two-hop neighborhood. We represent total packets per second at a node R as PPS_R ; packets per second from S at R as PPS_{SR} ; and packets per second from S at R at a data rate d as $PPS_{SR}(d)$.
- Fractional Busytime impact of a node on its neighbors. We represent the fractional Busytime impact of a node S on node R as BI_{SR} , and it can be understood as: for every y packets transmitted by node S , $BI_{SR} * y$ packets are carrier sensed at R .
- The most likely data rate to be used between a node S and its neighbor R , and the expected number of transmissions to send a packet successfully at the data rate. DR_{SR} is the data rate used most during the previous interval of time. The expected number of transmissions at a data rate d on link S - R is represented as ETX_{SR}^d .

4.2.2 Busytime estimation

We now describe steps 2 and 3 of MDAC operation.

Local Busytime: As discussed in Section 4.1, Busytime is the fraction of time spent in the TX, RX or ChannelBusy states. The MAC overhead, t_{MAC} , can be approximated by measuring the number of overheard data and ACK packets. For a network operating on the 2.4 GHz band, we can account for t_{MAC} by promiscuously measuring the packets heard on 802.11b and 802.11g. t_{MAC} calculation also accounts for the DIFS, SIFS and the average backoff for a packet. The average backoff can be computed as a function of the number of neighbors n at a node [9]. Note that the Busytime estimation can only be a close approximation as it is difficult to model backoff accurately and account for all the packets that cause medium busy-time at a node.

Let the Busytime requirement of the data and ACK packets of a new flow f to be admitted between nodes S and R be BR_f^b and BR_f^a . The Busytime requirement values can be estimated with the knowledge of the bit rate of the flow, the likely data rate on link S - R (D_{SR}) and the expected number of transmissions at data rate D_{SR} . The bit rate information for the flow can be obtained from the application layer. For instance, a typical value of 64Kbps with a packet size of 160 bytes can be assumed for VoIP flows using G.711 codec. If δ represents the margin for variations, the first check for admission of a flow f from S to R is given by:

$$BR_f^d + BR_f^a + \delta < 1 - \text{Busytime}_S \quad (2)$$

$$BR_f^d + BR_f^a + \delta < 1 - \text{Busytime}_R \quad (3)$$

Busytime at neighbors: The Busytime impact on the neighbor nodes of S and R needs to be assessed in order to prevent the throttling of packet transmissions at these nodes. Consider a neighbor node N for the sender S . Let $BI_{SN} \geq 0$ and $BI_{NS} \geq 0$ represent the fractional Busytime impacts of one node on the other. Consider that the neighbor node N has $t_F = 0$; its medium time is completely saturated by its transmission and receptions. Let the local admission check at S and R in Equations 2 and 3 pass for a new flow f . This indicates that there is sufficient free time to accommodate all the transmissions at S and R . However, this check does not suffice to prevent the transmission at N from being throttled by the new flow f , even for the case of S being able to carrier sense all the packet transmissions from N . This is because additional transmissions from the new flow f contend for the medium with node N 's transmissions. The increase in contention results in increasing the overall transmit duration for the node N . However, since node N is already completely saturated, node N is likely to experience throttling of its transmissions if the new flow is admitted.

The additional medium contention introduced depends on the Busytime impact values BI_{SN} and BI_{NS} . For the best case of $BI_{SN} = 0$, there is no transmission throttling at node N . However, for any $BI_{SN} > 0$, the impact

on the transmissions at N depends on the extent of the transmission period that node S 'steals' from node N . This depends on multiple factors: the data rates selected by the nodes, the packet scheduling and the Busytime impact of one node on the other. We consider the free time left at node N if flow f steals all the medium time as predicted by the Busytime impact of node S and R .

If δ represents the margin for variations, the next check for admission of a flow f from S to R is given by:

$$BI_{SN} * BR_f^d + BI_{RN} * BR_f^a + \delta < 1 - \text{Busytime}_N \quad (4)$$

$BI_{SN} * BR_f^d$ represents the Busytime impact of the data packets from S at N , and $BI_{RN} * BR_f^a$ represents the Busytime impact of the ACK packets from R at N .

The final piece of the framework involves estimation of the impact of packet collisions in the neighborhood. We consider the impact of packet collisions from the new flow on the existing traffic in the medium as well as the impact of packet collisions on the new flow in our implementation. For space constraints, we do not present the analysis in this paper.

5. IMPLEMENTATION

We implemented the MDAC solution on UCSB Mesh-Net. Our implementation consists of two main modules at each node in the network: a measurement-based profiling module and the decision control framework. The profiling module collects and maintains the statistics at each node about its neighbors. The decision control framework makes flow admission decisions based on these statistics.

Measurement-based profiling: Our measurement framework consists of modifications to the MadWifi driver along with a user-space module to collect and exchange statistics between nodes. As discussed in Section 3.3, we leverage statistics from the SampleRate module to track the ETX values at different data rates for each neighbor. This information is periodically extracted through the /proc filesystem interface and broadcast to neighbors every 10 seconds. Additionally, as discussed in Section 4.2.1, we modified the driver to track the average RSS of all received packets and the number of packets from each neighbor. The measured packet reception and average RSS values are used to incrementally build the piecewise linear curves. The last piece of maintaining the neighbor profile involves measuring the Busytime impact of a node on its neighbor. The Busytime impact is computed from the average RSS statistics of the received packets as described in Section 4.2.1.

Decision control framework: The decision control framework is implemented as a user space library that can be imported into the application code. As discussed in Section 4, we implement the four operations of estimating local medium time, neighbor Busytime, packet collisions caused by the new flow and packet collisions on the flow.

For our evaluations, we use a Busy-time margin of 10% for the δ parameter in Equations 2, 3 and 4. In other

words, Busy-time is considered fully utilized at 90% usage. The above value of δ permits aggressive admission of flows and improves the medium utilization. In our implementation, we set parameters for use with the SampleRate algorithm as follows. SampleRate switches to a data rate expected to consume the shortest time for a successful packet transmission. We can thus compute the packet reception rate required at each data rate before a switch to a lower data rate occurs. We restrict flows when the predicted collision losses at the current data rate reach within a margin of $\delta=10\%$ of the packet losses at which SampleRate switches to a lower data rate.

6. EVALUATION

All the experiments described in this paper are conducted over the UCSB MeshNet which is a 25-node indoor wireless testbed deployed in three floors of the engineering building [13]. All nodes in the testbed use 802.11b/g cards based on the Atheros chipset. Each node is also equipped with an Ethernet interface that is used to control the node during experiments, thus ensuring that the experiment control traffic does not affect the wireless network experiments. We use the MadWifi [14] driver v0.9.2 to control the cards. RTS/CTS is disabled for all the radios. The testbed co-exists with an 802.11b WLAN that provides Internet connectivity throughout the building. We evaluate MDAC on UCSB MeshNet. The focus of our evaluation is to understand the Busytime impact and collision prediction, the ability of MDAC to limit flows in the network, and to ensure QoS to VoIP flows.

6.1 Busytime prediction

We first understand the accuracy of Busytime prediction in the presence of traffic. This helps us understand the effectiveness of computing local and neighborhood Busytime as part of the admission control metric.

In these experiments our objective is to understand the accuracy of Busytime prediction. We do not yet perform admission control, and hence we expect packet collisions, particularly from hidden terminals, to affect the Busytime prediction. Packet losses due to collisions lead to variation in the data rate used by the rate selection algorithm. This can in turn affect the Busytime usage at the nodes that overhear the traffic. The prediction is thus expected to be less accurate in the scenarios when the Busytime is affected by packet collisions.

We adopt the following procedure to filter the results in situations where we expect the variation in data rate due to collisions to have caused a significant impact on the Busytime at nodes. With the collected statistics, we can isolate instances where nodes outside the carrier sense range of a node that has initiated a flow experience significant variation in the Busytime for the duration of the flow. In other words, we separate instances where nodes hidden from the initiator node experience significant Busytime variation during the flow. One of the factors that may have caused this variation is the packet collisions that result from the new flow admission. Other factors






Figure 4: CDF of predicted Busytime error.



(a) Without collision estimation



(b) With collision estimation

Figure 5: Per-flow throughput with MDAC.

include link quality variations or traffic variations. However, the traffic on our testbed does not vary significantly due to the constant bit rate used in the experiments. We therefore believe that packet collisions account for a majority of these instances of variation in Busytime.

Methodology: At a central server, we generate a flow sequence file that lists a random sequence of 50 one-hop flows between node pairs in the network and indicates the initiation time and duration of the flow. The average inter-flow arrival time is 20s and the average duration of each flow is 300s. The random sequence of flows enables us to study the busy-time behavior with different levels of network utilization. The flow sequence file is distributed to all the nodes. Prior to the experiment, the nodes are time-synchronized with a common NTP server to the accuracy of about 5ms. The nodes initiate 64Kbps UDP

flows using a fixed packet size of 160 bytes in accordance with the initiation and duration time specified in the sequence file. Additionally, all nodes collect the available neighborhood Busytime statistics two seconds prior to and 10 seconds after a flow initiation between any pair of nodes. The nodes also record the predicted value of increase in Busytime for each flow prior to its initiation.

We repeat the experiment five times and collect the Busytime statistics from all nodes. With the help of the Busytime statistics collected prior to every flow initiation, we compute the deviation in the observed Busytime from the predicted increase in Busytime of a flow.

Results: Figure 4 plots a CDF of the deviation of the predicted medium Busytime from the observed Busytime for all the results and the filtered results. The median Busytime variation for the filtered results is about 6%, while it is about 18% for the complete results. We make two inferences from the results. First, collisions have a significant effect on the accuracy of Busytime prediction. This implies that the performance of admission control schemes that consider only busy-time to limit flows may suffer due to packet collisions in the network. Second, the Busytime prediction, in the absence of hidden terminal collisions, is fairly accurate (<10% deviation).

6.2 Admission control

We now evaluate the effectiveness of MDAC for admission control. In particular, we study the throughput performance of the flows in a network that uses MDAC.

Methodology: The basic methodology of network setup and the flow initiation is similar to the methodology used in Section 6.1. We initiate 30 flows with a flow interval of 10s between random sets of 1-hop node pairs on the testbed. Each flow is a 64Kbps constant bit rate UDP flow with fixed packet size of 160 bytes and lasts for 300s. We track the average throughput (in packets per second) and the average delay experienced by the individual flows. We study the performance of MDAC in two system configurations: in the first, we consider only the Busytime estimation for admission control decisions; in the second, we consider both the Busytime estimation and the collision estimation for admission control decisions. The second scenario represents the complete MDAC solution.

Results: Figure 5(a) shows a timeline of the throughput performance of 15 admitted flows using only the Busytime estimate, i.e. without collision estimation. We find that the throughput performance of four flows is significantly affected (>25% drop in throughput). Repeated experiments with higher Busytime margins (δ) of up to 50% did not prevent the throughput reduction of these flows. The average packet delay for the four affected flows is 1.2, 2.3, 1.5 and 2.1 seconds. These results illustrate that Busytime predictions alone do not suffice for admission control.

Figure 5(b) shows the timeline of performance using the complete MDAC system, i.e., with both the Busytime and collision estimates. The graph shows that 12 flows were admitted and the network could sustain all




Figure 6: Number of VoIP flows admitted.

the admitted flows. Measurements showed that the average packet delay of all the 12 admitted flows was under 60ms. This result, along with the results discussed next, shows that MDAC is effective in limiting the number of flows in the network and ensuring bandwidth availability to admitted flows.

6.3 Ensuring QoS

We now evaluate the performance of MDAC in terms of providing the desired quality of service for the example application of VoIP. A practical guideline to evaluate the quality of VoIP flows in a network includes the following two constraints: a network delay budget of 80ms, and packet loss rate below 5% assuming the presence of error concealment algorithms [15].

Methodology: We follow a similar methodology as in Section 6.2 with a flow interval rate of 10s and duration of 300s. Additionally, we implement a handshaking protocol that mimics the call-setup phase before the start of each VoIP flow. The duration of the experiment is 30 minutes and we run seven trials of the experiment with different flow sequence files. For each flow topology, the experiment is performed with the MDAC system as well as without admission control.

Results: Figure 6 compares the number of flows admitted with admission control (AC) and no admission control (No AC) for different flow topologies. Each topology corresponds to a different experiment trial. The figure also shows the number of flows for which the QoS requirements were satisfied with admission control and no admission control. Of the possible 180 flows, the maximum number of flows that were admitted during any experiment is only 93. This is because heavy congestion causes failure of the call-setup handshake and prevents establishment of the flow. The results show that without admission control, only a small fraction of the flows have their QoS requirements satisfied. On the other hand, MDAC allows initiation of only about 50-60% of the number of flows initiated without admission control. More importantly, the QoS requirements are met for about 95% of the admitted flows, thus demonstrating the effectiveness of MDAC.

7. RELATED WORK

The problem of ensuring QoS in 802.11-based wireless networks presents several challenges and has been a topic of active research for several years. The channel variations and the contention-based access of the 802.11 protocol makes the problem of ensuring the QoS requirements of delay-sensitive and bandwidth-sensitive multimedia applications a challenging problem.

In prior research, TDMA-style approaches [16, 17] have been proposed to enable fine-grained control of the medium and provide QoS guarantees to real-time applications. The controlled access of the medium simplifies the problem of reserving bandwidth on a per-flow basis. These approaches, however, require effective synchronization among the nodes in the network and are thus difficult to realize in a real-world network.

Some QoS solutions, such as [18, 19], propose a joint admission control and routing scheme for ad hoc wireless networks. They provide detailed computations that can be used to estimate available bandwidth in a multi-hop wireless network. The bandwidth computations for a node assume a 100% Busytime impact for all the neighboring nodes. However, in our work we show that this assumption leads to overestimation of the bandwidth requirements in the network.

Contention-aware Admission Control Protocol [4] and Perceptive Admission Control [6] attempt to address the challenges of admission control in a multihop wireless network. These approaches focus their attention on the problem of communication with neighbors in the carrier sense range and propose solutions that cannot be implemented with current hardware. In contrast, we have demonstrated the possibility of communication with carrier sense neighbors using commodity hardware.

SoftMAC [15] proposes a software framework that employs coarse-grained control to regulate network load and ensure QoS. The authors suggest the use of a Fraction of Air Time metric to estimate available bandwidth at a link. SoftMAC, however, does not account for either the impact of collisions or the partial Busytime impact.

To the best of our knowledge, ours is the first admission control scheme that relies on real network measurements for its decision-making. The use of measured Busytime and link quality metrics enable MDAC to account for the impact of collisions and the partial busy-time impact and facilitate effective admission control.

8. CONCLUSION

With the growing popularity and increasing usage of 802.11 wireless networks, there is an urgent need for mechanisms to prevent network breakdown in the presence of heavy traffic. We present MDAC, a measurement-driven admission control framework that leverages wireless characteristics for admission control in a static wireless network. While our work is a critical first step towards realistic flow control in wireless networks, much

work remains to address the challenges of a mixed traffic network and make MDAC robust enough to handle all network and traffic scenarios.

Our work shows the feasibility of a distributed framework to predict wireless behavior, based on inferences about wireless conditions from local measurements. This makes a strong case for measurement-based approach to protocol design in wireless networks. As part of our future work, we intend to design MAC layer solutions that leverage local busy-time information about carrier sense neighbors for efficient packet scheduling.

Acknowledgements

This work was funded in part by NSF Career Award CNS-0347886 and by NSF NeTS Award CNS-0435527.

9. REFERENCES

- [1] A. P. Jardosh, K. Mittal, K. N. Ramachandran, E. M. Belding, and K. C. Almeroth, "IQU: Practical Queue-Based User Association Management for WLANs," in *Proc. of MobiCom*, Los Angeles, CA, Sep 2006.
- [2] J. Eriksson, S. Agarwal, P. Bahl, and J. Padhye, "Feasibility Study of Mesh Networks for All-Wireless Offices," in *Proc. of MobiSys*, Uppsala, Sweden, June 2006.
- [3] J. Camp, J. Robinson, C. Steger, and E. Knightly, "Measurement Driven Deployment of a Two-Tier Urban Mesh Access Network," in *Proc. of MobiSys*, Uppsala, Sweden, June 2006.
- [4] Y. Yang and R. Kravets, "Contention-Aware Admission Control for Ad Hoc Networks," *IEEE Transactions on Mobile Computing*, vol. 4, no. 4, pp. 363–377, Jul 2005.
- [5] K. Xu, M. Gerla, and S. Bae, "How effective is the IEEE 802.11 RTS/CTS handshake in ad hoc networks," in *Proc. IEEE GLOBECOM '02*, Taipei, Taiwan, Nov 2002.
- [6] I.D.Chakeres and E.M.Belding-Royer, "PAC: Perceptive Admission Control for Mobile Wireless Networks," in *Proc. of QShine*, Dallas, TX, Oct 2004.
- [7] J. Padhye, S. Agarwal, V. N. Padmanabhan, L. Qiu, A. Rao, and B. Zill, "Estimation of Link Interference in Static Multi-hop Wireless Networks," in *Proc. of Internet Measurement Conference*, Berkeley, CA, Oct 2005.
- [8] C. Reis, R. Mahajan, M. Rodrig, D. Wetherall, and J. Zahorjan, "Measurement-Based Models of Delivery and Interference in Static Wireless Networks," in *Proc. of SIGCOMM*, Pisa, Italy, Sep 2006.
- [9] X. Yang and N. Vaidya, "On the Physical Carrier Sense in Wireless Ad Hoc Networks," in *Proc. of INFOCOM*, Miami, FL, Mar 2005.
- [10] J. Bicket, "Bit-rate Selection in Wireless Networks," Master's thesis, Massachusetts Institute of Technology, 2005.
- [11] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris, "Link-Level Measurements from an 802.11b Mesh Network," in *Proc. of SIGCOMM*, Portland, OR, Aug 2004.
- [12] C.-H. Lin, H. Dong, U. Madhow, and A. Gersho, "Supporting Real-Time Speech on Wireless Ad hoc Networks: Inter-Packet Redundancy, Path Diversity, and Multiple Description Coding," in *Proc. of WMASH*, Philadelphia, PA, Oct 2004.
- [13] (2006, December) UCSB MeshNet. [Online]. Available: <http://moment.cs.ucsb.edu/meshnet/>
- [14] (2006, December) MadWifi. [Online]. Available: <http://madwifi.org/>
- [15] H. Wu, X. Wang, Y. Liu, Q. Zhang, and Z. Zhang, "SoftMAC: Layer 2.5 MAC for VoIP Support in Multi-hop Wireless Networks," in *Proc. of SECON*, Santa Clara, CA, Sep 2005.
- [16] C. Lin, "Admission Control in Time-Slotted Multihop Mobile Networks," *IEEE JSAC*, vol. 19, no. 10, pp. 1974–1983, Oct 2001.
- [17] C. Zhu and M. Corson, "QoS routing for mobile ad hoc networks," in *Proc. of ICC*, Miami, FL, Jun 2002.
- [18] Q. Xue and A. Ganz, "Ad hoc QoS on-demand routing (AQOR) in mobile ad hoc networks," *Journal of Parallel and Distributed Computing*, vol. 63, no. 2, pp. 154–165, 2003.
- [19] L. Chen and W. B. Heinzelman, "QoS-aware routing based on bandwidth estimation for mobile ad hoc networks," *IEEE JSAC*, vol. 23, no. 3, pp. 561–572, Mar 2005.