Perceptive Admission Control for Wireless Network Quality of Service

Ian D. Chakeres ⋆

Boeing Phantom Works
Mathematics & Computing Technology
Seattle, WA 98124 USA
ian.chakeres@gmail.com

Elizabeth M. Belding-Royer

University of California, Santa Barbara
Dept. of Computer Science
Santa Barbara, CA 93106 USA
ebelding@cs.ucsb.edu

Joseph P. Macker

Naval Research Laboratory
Information Technology Division
Washington, DC 20375 USA
macker@itd.nrl.navy.mil

Abstract

As wireless networks become more widely used, there is a growing need to support advanced services, such as multimedia streaming and voice over IP. Traditional approaches to guarantee quality of service (QoS) work well only with predictable channel and network access. In wireless mobile networks, where conditions dynamically change as nodes move about the network, a stateless, high level approach is required. Since shared wireless resources are easily over-utilized, the load in the network must be controlled so that an acceptable QoS for real-time applications can be maintained. If minimum real-time requirements are not met, these unusable packets waste scarce bandwidth and hinder other traffic, compounding the problem. To enable high QoS for all admitted traffic, we propose the Perceptive Admission Control (PAC) protocol. PAC monitors the wireless channel and dynamically adapts admission control decisions to enable high network utilization while preventing congestion. Through discussion, simulations and testbed experiments, we demonstrate that PAC ensures low packet loss and delay for all admitted flows.

1 Introduction

Wireless devices are becoming prevalent because of their ability to provide mobile networking. Since many common applications, including voice and multimedia, require low packet loss and delay, quality of service (QoS) is an important requirement for these networks. In contrast to traditional wired networks, mobile networks operate under harsh conditions that include a shared wireless channel, limited bandwidth, and mobility.

Traditional attempts to provide guaranteed QoS [1] are unable to cope with constantly changing wireless network conditions. Similarly, hard real-time QoS constraints in wireless mobile networks are unrealistic because of shared medium access and mobility. Solutions that provide a stateless service and offer better than best-effort packet delivery for high priority packets, such as DiffServ [2] and IEEE 802.11e [3], are more successful. Unfortunately, these solutions may fail to provide the low loss and delay that real-time applications require if the network becomes congested.

QoS for high priority flows is achievable without fully coordinated channel and network access. The wireless channel must be kept from reaching the congestion point since loss and delay increase rapidly once this point is reached. Keeping the utilization below the congestion point is difficult because the channel is shared between nodes that may not be able to communicate directly.

To control the amount of traffic in the network and provide high quality service to all admitted traffic, we introduce the Perceptive Admission Control (PAC) protocol. PAC ensures that the network does not admit a flow that will cause congestion. To make an admission decision, sources consider not only the limited area within their transmission range, but the entire region that a new flow's transmissions will impact. We show that with a proper carrier signal detection range, the time that the wireless channel is sensed as busy is a good estimate of the utilization and available bandwidth. Using this measure, PAC performs admission control for new flows to avoid admitting flows that would cause congestion. Our discussion focuses on single hop admission control.

The rest of this paper is organized as follows. Section 2 provides background on wireless transmissions, including methods for determining the available bandwidth and previous approaches for providing high packet delivery and low delay in wireless networks. In Section 3 we describe PAC, our approach for admission control. We present the performance of PAC in simulation (Section 4), describe how it avoids the shortcomings of previous approaches (Section 5), and discuss experimental results from tests performed using our Mica2 mote implementation (Sec-

⋆ The research discussed in this articale was performed while Ian Chakeres was a graduate student at University of California, Santa Barbara.

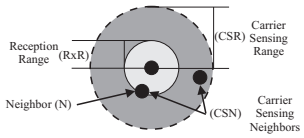

Fig. 1. Approximation of wireless communication ranges, specifically, the reception range (RxR) and carrier sensing range (CSR). Nodes within RxR are called neighbors (N), while carrier sensing neighbors (CSN) are all nodes within CSR.

tion 6). Section 7 presents our conclusions.

2 Background

To perform admission control in wireless networks, it is important to understand how a wireless transmission impacts other nodes in distributed wireless networks. In this section, we describe the important distances for packet transmission and reception. We then examine several methods for calculating the available bandwidth, since admission control decisions depend on accurate estimation of the available bandwidth. We also elaborate on the effect of a less-deterministic propagation model including fading and multipath. We then discuss related work and why most proposed solutions are insufficient. Subsequently, we describe the solution most closely related to our approach.

2.1 Impacted Area

There are a number of important ranges for wireless communication. Each of these is important for MAC layer protocol operation, measurement of channel utilization and prediction of available bandwidth, as well as other wireless mechanisms including perceptive behaviors [4]. At short range, we assume that nodes are capable of direct communication. We refer to the maximum separation between a sender and receiver for successful packet reception as the reception range (RxR), shown in Figure 1. Nodes within RxR of a particular sender are called as its neighbors (N).

Nodes that are within carrier sensing range (CSR) of a sender can detect packet transmissions. These nodes are called its carrier sensing neighbors (CSN). All CSN are able to detect a transmission but they will not be able to decode the contents of the packet if they are outside RxR. The relationship between the RxR and CSR

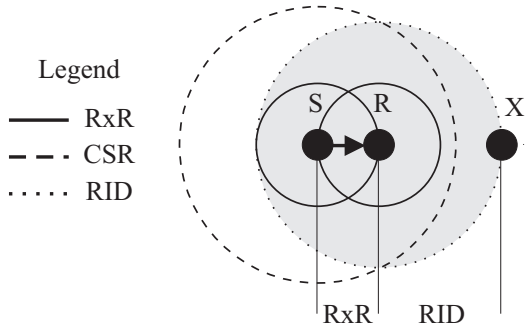

Fig. 2. The receiver interference distance (RID) is the distance between a receiver (R) and another sender (X), such that the receiver can successfully receive node S's packets and node X can simultaneously send a packet to another receiver.

is defined by the physical layer encoding (PHY), and is often configurable to encourage or discourage spatial reuse. The CSR is typically much larger (two to four times larger) than the RxR to avoid collisions.

In wireless MAC protocols based on CSMA, such as IEEE 802.11 [5] and IEEE 802.15.4 [6], the CSN of the sender are not allowed to initiate a packet transmission while another node is transmitting because they sense the channel is busy. In CSMA networks, a large CSR prevents multiple transmissions from simultaneously occurring close together and helps avoid interference at receivers. In contrast, a smaller CSR allows for more spatial reuse, though more collisions and interference may occur.

For correct packet reception, the channel surrounding a receiver must be free of multiple interfering transmissions. If another node close to the receiver transmits a packet, it may interfere with an ongoing packet reception, even if the two senders are outside each others' CSR. To quantify this effect, we define the receiver interference distance (RID) as the distance between a receiver and another sender, such that this receiver's ability to decode a packet from its sender is not affected. For example, in Figure 2, if node X is outside node R's RID, node X can transmit at the same time as node S without affecting packets received by node R from node S. If node X is inside node R's RID and transmits a packet at the same time as node S, node R is unable to successfully receive the packet from node S because the two packet transmissions collide. In both cases, node X is not prohibited from transmitting because node S is outside node X's CSR, and it cannot sense an ongoing transmission between nodes S and R.

The exact size of the RID depends on many factors, including transmission power, minimum reception power, propagation model, and hardware capture capabilities. In Figure 2, note that the CSR (dashed line) is larger than the RID (dotted line) and the RID is larger than the RxR (solid line). These line styles will be used throughout

the PAC discussion to denote the different ranges.

In the transmitter, the RxR and CSR depend on the characteristics of the wireless physical layer transmission protocol. A physical layer transmission protocol can be engineered to create a specific RxR and CSR. For example, in IEEE 802.11 [5], multiple physical layer encoding schemes exist, such as IEEE 802.11b [7], and IEEE 802.11g [8]. The CSR for all data rates is the same, while the RxR shrinks as the data rate rises. However, CSR will always be much greater than RxR. At the receiver, the ability to receive or detect a transmission depends on the sender's physical layer encoding, the propagation model and the receiver sensitivity. The propagation model defines the loss in power a transmission incurs as it travels from the transmitter to the receiver. The receiver sensitivity is based on the receivers' ability to decipher a transmission or detect a carrier signal. The receiver will not be able to decode or detect a packet if the received signal is close to the noise level. In most CSMA protocols, including IEEE 802.11, the signal power at RxR and CSR are well above the noise level, so noise does not significantly impact these ranges.

For two simultaneous transmissions to be successfully received by different receivers, the transmitting nodes must be separated in space. The distance between two senders that ensures proper packet reception at a receiver is $RxR + RID$. This distance holds for all possible network scenarios and could be considered the worse-case scenario. At any distance smaller than $RxR + RID$, it is possible that the transmissions of two senders will interfere with a receivers' ability to properly decode a packet.

The communication distances and related thresholds described above are for networks where all nodes use omnidirectional antennas and transmit packets with the same transmission power on the same channel. We also assume no obstacles and that only simple fading occurs. In Section 2.3 we explore a more complex propagation model. We do not look at relaxation of the other assumptions here.

2.1.1 MAC Layer Acknowledgments

Acknowledgments (ACKs) are used in many MAC layer protocols, such as IEEE 802.11 [5], to immediately inform the sender that successful unicast packet reception has occurred. If an ACK is not received, the sender will retransmit the packet multiple times. The Data-ACK mechanism is used to combat packet loss at the MAC layer caused by collisions and errors introduced by the wireless channel. Generally, carrier sensing is not performed by the receiver prior to sending an ACK. The channel state is not checked at this time because carrier sensing might silence a receiver. If the receiver is silenced, it would prevent transmission of the ACK and require the sender to retransmit the packet. Retransmission of the packet, in turn, would waste wireless resources and increase delay.

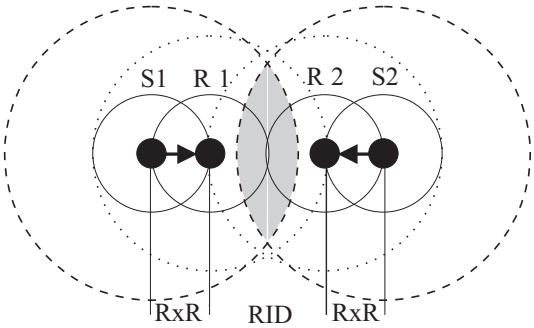

Fig. 3. This figure shows the spacing requirement for simultaneous transmissions in wireless networks that utilize MAC layer acknowledgments.

When receivers do not perform carrier sensing prior to sending an ACK, they must be separated by RID to ensure that no collision occurs. In this type of network, the separate sets of data and ACK transmissions should not overlap. If they do overlap, the data transmissions and ACKs will cause a collision, which will result in unsuccessful packet reception.

Given that the two receivers are separated by RID and each sender-receiver pair is separated by RxR, the distance between two senders for successful simultaneous transmissions is

$$2 * RxR + RID \tag{1}$$

A network topology illustrating this distance is shown in Figure 3. In this worst-case scenario, if the two senders are closer than $2 * RxR + RID$ and the transmissions overlap in time, the data and ACK pairs will collide and communication will suffer.

2.2 Determining the Available Bandwidth

The goal of our work is to allow nodes to depend on their estimation of the available bandwidth to make correct admission control decisions. In this section, we examine several methods to determine the available bandwidth.

The most common way to calculate available bandwidth (B_{avail}) is to measure network utilization (U). Given the network utilization and the maximum bandwidth (B_{max}), the available bandwidth is estimated using the following equation [9]:

$$B_{avail} = (1 - U) * B_{max} \tag{2}$$

There are many techniques to measure the network utilization. Some metrics of

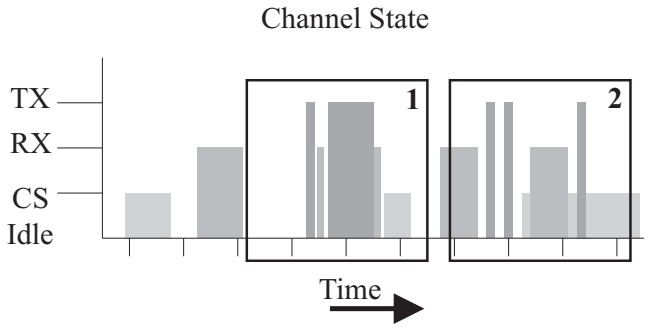

Fig. 4. An example of the channel state versus time. At different times a node may either be idle, sensing a packet transmission (CS), receiving a packet (RX) or sending a packet (TX). In window 1, the channel is busy half the time. In window 2, the channel is busy nearly 100%.

network utilization are queue length, MAC layer congestion window, number of collisions, delay, and channel busy time. The queue length, MAC layer congestion window, and number of collisions provide little or no information regarding network utilization if the network is not congested. For example, packets will not be queued unless the network is in a congested state, so nodes cannot accurately measure current utilization using these methods. Since these three techniques are not adequate for determining the available bandwidth, we explore only the two remaining techniques, delay and channel busy time, in more detail.

Delay is one of the most widely used metrics for determining available bandwidth. In general, approaches to measure this metric inject probe packets into the network that solicit responses from another node. The other node then returns either the packets or a measurement from the packets received. While many advanced probing techniques exist, they present a number of problems (a comparison of many of these techniques may be found in [9]). The primary disadvantage of probing delay as a measure of available bandwidth is overhead since bandwidth is scarce in wireless networks. Additionally, probing provides only an instantaneous value; the probe must be repeated several times to create an average value, which in turn further increases overhead. Also, since probes are an active measurement technique, the probes may not be able to determine an accurate value if packet loss occurs. Losses thus reduce the quality of the measurement. Finally, since probing is performed between each pair of nodes, measurements from probe messages may be significantly different between different node pairs. These measurements are highly dependent on many factors including the two nodes' spatial location and the network load.

The second technique for determining network utilization, called busy time, is a direct measure of channel utilization. In wireless networks, a node can detect three states: transmitting, receiving, and busy. If the node detects a carrier signal, it senses that the channel is busy, and it is only able to decode and receive the packet contents if the packet is transmitted by a node within RxR. By measuring the amount

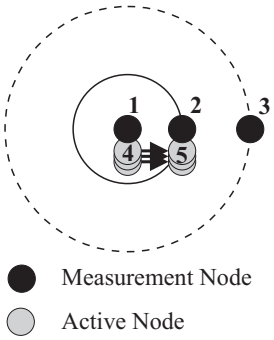

Fig. 5. Network topology for utilization metric testing and available bandwidth calculation.

of time the channel is sensed as busy (CS), receiving (RX) or sending (TX), a node can measure not only transmissions that occur within its RxR, but also those within its CSR. By measuring the portion of time the medium is busy, a node can estimate the current utilization. We define the busy time to be the total time within an interval that a node is transmitting packets, receiving packets, or sensing packet transmissions. For example, in Figure 4, the channel is busy half the time in window 1. In window 2, the channel is nearly always busy.

Network simulations were performed in order to demonstrate the ability to determine the available bandwidth using the busy time measurement. IEEE 802.11 [5] was the PHY and MAC layer protocol for the experiments. The data rate was set to 2 Mbps. A network consisting of three measurement nodes and ten sender-receiver pairs (Active Nodes) was created, as shown in Figure 5. Ten senders were chosen so that the wireless channel usage could be stressed. Node 1 and all senders were co-located. Likewise, node 2 and all receivers were co-located just inside the RxR of the senders. Node 3 was located just inside the CSR of all the senders. Each active sender-receiver pair transmitted constant bit rate (CBR) traffic. The measurement nodes were not the source or destination of any CBR traffic. Simulations with an aggregate traffic load from zero to 2 Mbps were performed. Each node monitored every packet it transmitted, received or sensed to calculate the busy time.

The busy time metric correctly measured the utilization; it varied from zero (fully idle) to almost one (fully busy). In this scenario, the maximum achievable throughput (B_{max}) was 1200 kbps, which is close to the theoretical value in IEEE 802.11 networks [10]. This maximum bandwidth was used, along with the measured utilization and Equation 2, to calculate the available bandwidth. Figure 6 shows the network utilization and available bandwidth using the busy time measure for nodes one through five. Only one line is visible in the graph because all five nodes detected the same utilization and calculated the same available bandwidth.

With any measurement technique, it is common for instantaneous values to vary, sometimes widely. For our approach, we utilize an equally weighted sliding win-

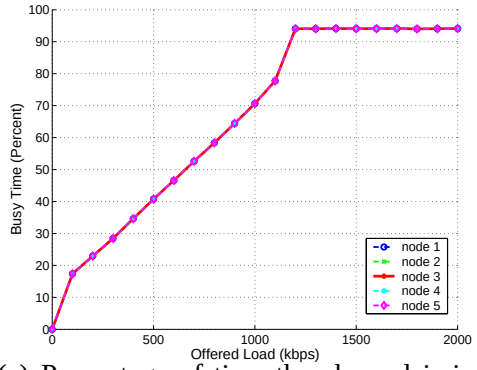

(a) Percentage of time the channel is in use versus increasing load.

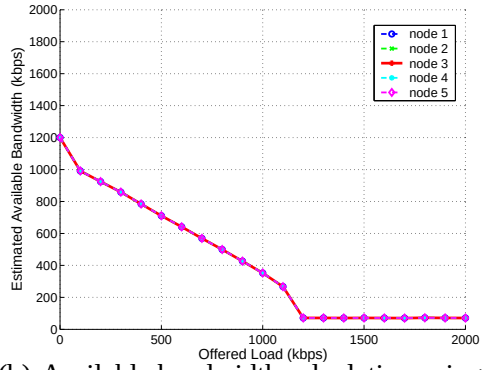

(b) Available bandwidth calculation using busy time measurement.

Fig. 6. Simulation results of the busy time measurement technique.

dow to obtain the wireless utilization. Through testing, we determined a window size that was large enough to obtain an accurate estimate and small enough to quickly adjust to changing traffic conditions. An alternate weighting technique, such as a weighted average that favors recent measurements, may provide a good estimation of utilization and available bandwidth and quick adaptation to flows entering and leaving the system.

In this section, we examined busy time as a way to determine the utilization. The validation simulations were run using the two ray ground propagation model which provides deterministic RxR and CSR. The two ray ground model, however, may not realistically represent actual wireless transmission propagation. For this reason, we analyzed the effect of a more realistic propagation model, called the shadowing model, presented in the next section.

2.3 Realistic Wireless Propagation

The *shadowing* model [11] in the NS-2 simulator represents signal propagation more realistically than the two ray ground model. The shadowing model includes path loss and multipath fading. Path loss reflects the drop in signal power with distance and is defined by the path loss exponent. Multipath (i.e. Raleigh and Ricean) fading is simulated using a log-normal random variable. The combination of these two properties results in a probabilistic distribution of packet reception. For example, two nodes are able to communicate with some probability at a fixed distance. This varied communication distance is in contrast to the two ray ground model, where communication between two nodes is deterministic, based on their separation distance.

Using the shadowing model, the RxR varies probabilistically with distance, as

Propagation Model	Path Loss Exponent	% Packets Received	RxR
Two Ray Ground	Free Space	100%	250m
Shadowing	Free Space	99%	95m
Shadowing	Free Space	95%	130m
Shadowing	Free Space	90%	155m
Shadowing	Free Space	80%	190m
Shadowing	Free Space	70%	225m

Table 1. Reception range with various propagation model parameters.

shown in Table 1. For fixed receiver sensitivity and high reception probability, the shadowing model results in a short effective range. For high probability of reception (greater than 90%) while utilizing the shadowing model, the RxR is much shorter than the two ray ground model. If only nodes with a high probability of reception communicate, the RxR will be much shorter while utilizing the shadowing model.

To examine the behavior of the shadowing model on the available bandwidth calculation, simulations similar to those in Section 2.2 were performed. For these simulations, free space propagation is assumed. Higher path loss exponents, such as those for scenarios in urban areas, do not alter the fact that a more realistic propagation model does not significantly impact the ability to determine the available bandwidth.

In these simulations, there were a few modifications to the original simulation scenario exploring the busy time measure. In the new simulations, the receivers nodes were placed 155m from the sources. This allowed the receivers and Node 2 to receive 90% of the packets sent. Node 3 was placed 550m from Node 1. This location is the same location as the previous experiment using the two ray ground model. At 550m node 3 senses the carrier signal of more than 99% of the transmissions. The simulation results matched those shown in Figure 6 because all the nodes sensed more than 99% of the transmissions. The closer a node is to a sender, the more packets it detects when using the realistic propagation model. Similarly, as the distance between nodes increases, the ability to detect transmissions decreases. This ability to detect transmissions is exactly the behavior expected and it helps nodes correctly calculate the available bandwidth. Overall, a more realistic propagation model does not negatively impact the ability to use carrier signal sensing to calculate available bandwidth. For this reason and to achieve deterministic simulation and analysis, we utilized the two-ray ground in our simulations of PAC in Section 4.1.

2.4 Related Work

The shared nature of the wireless channel presents a challenge to QoS protocols that does not exist in wired networks. For this reason, QoS approaches that require

MAC layer synchronization (i.e. TDMA) [12–15], network-wide information dissemination [16–18] or reservations [19,1] do not work well in mobile networks where the network topology changes frequently. Similarly, measurement-based admission control protocols [20–22] cannot be applied directly since they ignore the impact of the wireless medium. PAC is a measurement-based admission control protocol, and we define the requirements to properly measure the current utilization in wireless networks.

INSIGNIA [23] and SWAN [24] are both protocols that enable a high QoS by limiting the traffic in the network. INSIGNIA uses in-band signaling by piggybacking control information on data. This in-band signaling allows INSIGNIA to quickly restore flow state when topology changes occur. In SWAN multiple traffic classes and explicit congestion notification (ECN) help give priority packets better than best effort QoS. The main drawback to both of these protocols is that they rely on active probe messages to determine the current available bandwidth and ignore many characteristics of wireless networks.

The Contention-aware Admission Control Protocol (CACP) [25] is one strategy that addresses admission control for wireless networks and considers the shared nature of the wireless channel. However, CACP has significant flaws and lacks support for node mobility. CACP is described in detail in the next section and qualitatively compared with our solution in Section 5.

Our admission control protocol, PAC, was designed specifically to be used in wireless mobile networks. PAC considers the shared nature of the wireless channel and the receiver's reception requirements. In addition, it is a stateless approach that does not need network-wide synchronization or control message dissemination. Finally, node mobility and its effect on the shared channel is also taken into account.

2.4.1 Contention-aware Admission Control Protocol

The Contention-aware Admission Control Protocol (CACP) [25] is one strategy that addresses admission control for wireless networks and considers the shared nature of the wireless channel. CACP shares many characteristics with our admission control strategy. To make an admission control decision in CACP, each node considers not only the resources of its immediate neighborhood, but also the resources of all nodes within its CSR. CACP is contention-aware in that each node passively monitors the amount of time the channel is sensed as busy. This busy time includes the time a carrier signal is detected, as well as when a packet is transmitted or received. The available bandwidth is calculated by taking the inverse of the current channel utilization, as described in Section 2.2.

CACP consists of two main operations: an admission control decision that is performed on a hop-by-hop basis, and a multihop routing protocol. Before a new data

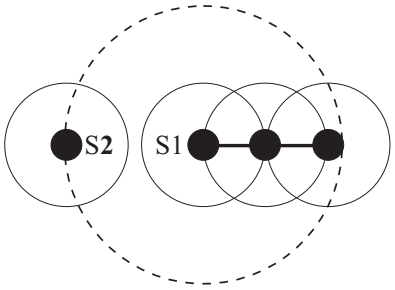

Fig. 7. This network presents a scenario where node S2, a CSN of node S1, cannot be reached via any multihop path.

flow is admitted, the available bandwidth must be checked. Since the available bandwidth calculation does not include all nodes that may be impacted by a new flow, a query message must be sent to all nodes within CSR. If all CSN detect enough available bandwidth, then the flow is admitted.

In making a single-hop admission control decision, CACP describes two methods to query the available bandwidth at the CSN of a node prior to flow admission. The first method is a multihop approach that floods query messages using a limited hop count. The CACP authors acknowledge that this approach operates inaccurately in networks where a node within CSR is not reachable via any path. For example, in Figure 7, node S2 must be queried to see whether the new flow can be admitted; however, it cannot be reached because it is outside of transmission range any node. Using this query method, node S1 cannot ensure enough network bandwidth is available at node S2.

In the second approach, a sender issues an available bandwidth query using a high power packet transmission. Through the high power transmission, all nodes within CSR of the new sender are contacted. If any node that receives the query does not have enough available bandwidth to support the new flow, it sends a rejection message which acts negative acknowledgment (nack), again using a high power packet transmission. This query-nack procedure results in poor behavior as the network utilization increases.

To explain CACP's single hop admission control decision operation, an example is provided. Consider the network shown in Figure 8 and an admitted traffic flow between nodes Z and Y that consumes half the network bandwidth. The network state is shown in Table 2 at time T1. Only nodes X, Y and Z detect the flow; node W does not detect the communication between Z and Y since it is outside of the measurement range. Later, node W wants to introduce a new traffic flow requiring 25% of the bandwidth. Node W checks its available bandwidth and discovers enough bandwidth is available. Node W then sends a query message to all nodes inside its CSR, i.e. nodes X and Y, using a high powered message. Both X and Y check their available bandwidth measurement. Since enough bandwidth is available, they do

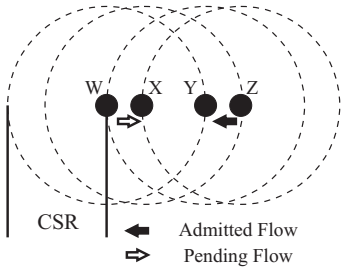

Fig. 8. CACP single-hop admission control example.

Time/Node	W	X	Y	Z
T1	100%	50%	50%	50%
T2	75%	25%	25%	50%

Table 2. CACP available bandwidth estimation.

not send a rejection message to node W. After a timeout, node W admits the new traffic flow. After a short time the available bandwidth measurement of each node adjusts to the newly admitted traffic, shown as time T2 in Table 2. Later, node W has another flow to admit. This flow requires 50% of the bandwidth. Node W checks its available bandwidth measurement and enough bandwidth is available, so node W sends a query message. Nodes X and Y receive the query and check their available bandwidth. Enough bandwidth is not available, so they both send a rejection message to node W. When node W receives a rejection message from either node X or Y, the pending admission request is denied.

Though we do not focus on multihop networks in this work, we should mention that CACP includes a multihop routing protocol that determines the bandwidth required for a new data flow at each hop along a path. The amount of bandwidth required at each node is a function of the number of neighbors on the path within CSR of the node. By requiring the available bandwidth to be large enough to support the local transmission of the flow and all other retransmissions of the same flow in its neighborhood, enough bandwidth for the complete path is ensured. A similar protocol may be used to extend PAC to multiple hops [26]. For a detailed description of CACP's multihop routing protocol, please refer to [25].

Though CACP works well in some networks, there are multiple problems with the protocol. Most importantly, CACP control packet losses lead to erroneous admission decisions, and the frequency of this event is directly correlated with the network load. Second, CACP does not have any mobility support. To achieve acceptable performance in a mobile network, CACP reserves extra capacity and leverages the routing protocol. Also, since each node relies on exchanging messages with its CSN to determine whether enough bandwidth is available, mobility support is prohibitively expensive. Finally, in CACP, conservative admission decisions lead to

lower aggregate network throughput by prohibiting some acceptable spatial reuse.

3 Perceptive Admission Control Operation

To maintain a high QoS for traffic in wireless mobile networks, we introduce the Perceptive Admission Control (PAC) protocol. The core idea for our admission control algorithm is to allow nodes to depend on their own estimation of the available bandwidth to make correct admission decisions. We change the range of the available bandwidth calculation to include all possible interfering sources. By including all nearby transmissions, admission control decisions are accomplished without the need to communicate with any other nodes.

3.1 Available Bandwidth Calculation and Admission Control

We alter the sensing range so that transmissions are sensed at a distance large enough to enable correct local admission decisions. As shown in Section 2.1, the distance between two senders (using CSMA with ACKs) to avoid any possible receiver interference is $2 * RxR + RID$. By changing the sensing measurement range to be at least the distance $2 * RxR + RID$, each node can itself make admission control decisions. At any distance greater than $2 * RxR + RID$, two ongoing transmissions will not interfere with packet receptions [1] . Therefore, when a node has to make an admission control decision, its PAC-based available bandwidth measurement is sufficient to make the correct decision. If the available bandwidth is more than the bandwidth required by the new flow, then the new flow can be admitted.

After a new flow is admitted, the flow immediately begins consuming network bandwidth. Since the available bandwidth calculation is continuously updated, it takes the newly admitted traffic into consideration for future admission control decisions. Similarly, when a flow stops, the increase in available bandwidth is quickly incorporated into the network utilization measurement so that other flows can be admitted.

For example, in Figure 9, assume there is an admitted traffic flow between nodes Z and Y that consumes half the network bandwidth. The current network state is shown in Table 3 at time T1. Since node Z is within $2 * RxR + RID$ of nodes W, X and Y, all nodes estimate the available bandwidth to be 50%. Node W wants to

[1] Using a PAC sensing range of $2 * RxR + RID$ will discourage spatial reuse, though since the MAC layer CSR is unmodified some spatial reuse may occur. Using a shorter PAC sensing range could increase spatial reuse, but it might also result in persistent collisions and congestion. The tradeoff between protocol correctness and higher performance by allowing more spatial reuse is not examined in this paper.

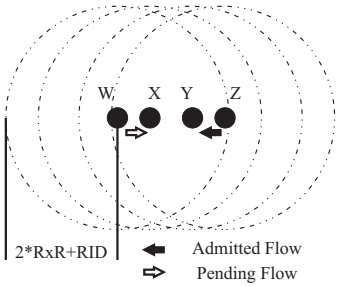

Fig. 9. PAC single-hop admission control example.

Time/Node	W	X	Y	Z
T1	50%	50%	50%	50%
T2	25%	25%	25%	25%

Table 3. PAC available bandwidth estimation.

introduce a new traffic flow requiring 25% of the maximum bandwidth. Node W checks its available bandwidth and determines that enough bandwidth is available. Hence it admits the new traffic flow. After a short time, shown as time T2 in Table 3, the available bandwidth measurement of each node adjusts to incorporate the newly admitted traffic. Later, node W has another flow to admit. This flow requires 50% of the bandwidth. Node W checks its available bandwidth measurement and determines that there is not enough bandwidth available. Hence node W does not admit the traffic flow. In contrast to CACP, PAC is able to determine the correct available bandwidth without requiring any inter-node communication.

In wireless CSMA networks, throughput drops once the network becomes congested [10]. To prevent the channel congestion, PAC ensures that the quantity of admitted traffic is below the network saturation point by reserving a small portion of the bandwidth. We call this amount the reserved bandwidth. The reserved bandwidth is also useful to detect changes in the available bandwidth due to mobility.

To admit a new flow, the required bandwidth (B_{req}) for the new flow must meet the following condition: $B_{avail} - B_{rsv} > B_{req}$. The inclusion of the reserved bandwidth prevents the channel from becoming congested and allows all admitted traffic to receive high delivery rates and low delay. The amount of reserved bandwidth can be varied based on the conditions of the channel, but for the purpose of our experiments the reserved bandwidth is fixed. There is a tradeoff between the amount of bandwidth reserved and the aggregate throughput attainable by admitted flows.

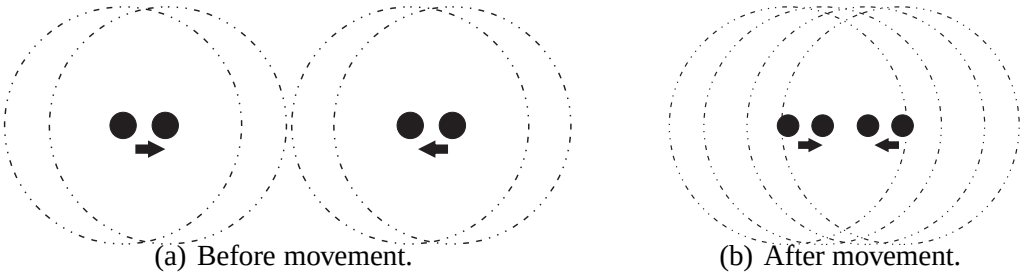

(a) Before movement. (b) After movement.

Fig. 10. In this example network, two flows are admitted outside impact range of each other. As the two sources get closer to each other they will interfere. To maintain a high QoS, the sources must throttle or reject some traffic.

3.2 Mobility

When a node and its traffic flows move within a wireless network, the area impacted by its traffic changes with the node's location. Therefore, it is important to not only admit flows, but also to throttle or reject them as network conditions change.

The following example illustrates the importance of determining whether the network is close to over-utilization. In Figure 10(a), suppose two flows, each consuming more than half of the maximum bandwidth, are admitted at nodes far enough apart that each participating node pair is outside the CSR of the other. Later, as shown in Figure 10(b), if the nodes participating in the network flows move into interference range of each other, the network will become saturated since it is not possible to support the two flows simultaneously. Using PAC, each source can detect the ensuing network congestion when another sender enters its PAC measurement range, and then each source can notify its applications to throttle or stop their traffic flows. In the example, if both flows are allowed to continue at their present transmission rate, neither flow will receive the high QoS necessary.

Therefore, to handle mobility, each source monitors the available bandwidth. If a source has an ongoing packet flow and the available bandwidth drops below a threshold value (B_{min}) when a packet is to be sent, then the flow source should throttle or stop the flow. After a random backoff time, a source with a throttled or rejected flow can attempt to increase or re-admit the traffic flow. By using this method, admitted flows backoff and the network remains in an un-congested state. For this study, we assume all flows require a minimum level of service such that the flow cannot be throttled. Therefore, PAC rejects flows to avoid congestion when the available bandwidth drops below a threshold value.

To avoid throttling multiple flows in response to mobility-induced congestion, some randomness should be introduced. Throttling multiple flows is discouraged because often only one flow must be throttled to avoid congestion. For our implementation,

each source checks only the state of the available bandwidth after a random time and when it has a packet to send. If the channel is congested at this time, this source throttles or stops the flow. Since the random timeout is large compared to the window size, it is unlikely that two sources will sense the channel and detect congestion before the available bandwidth calculation automatically adjusts.

3.3 Multihop Routing

The PAC admission decision can be utilized to create multihop routes during reactive route discovery in wireless multihop networks by checking the required bandwidth at each hop [26]. However, calculating the bandwidth required at intermediate nodes along multihop paths is difficult.

This difficulty arises from the fact that a wireless transmission impacts all nodes within CSR, but nodes can only effectively communicate with nodes inside RxR. To coordinate with nodes inside CSR, but perhaps outside RxR, other measures must be used, such as [27] and [25].

In addition to finding a multihop path that can support a flow's required bandwidth, congestion due to mobility should be monitored and detected. When congestion is detected, the source must be notified so that it can throttle or reject its traffic. Congestion detection may be performed continuously, periodically, or on-demand.

4 Simulation-based Evaluation

In this section, we demonstrate that PAC controls flow admission to avoid congestion and maintain a high QoS for all admitted flows. We present simulation results that show PAC performs admission control efficiently and effectively.

4.1 Simulation Environment

To evaluate PAC, we used the NS-2 simulator [11]. Our simulation parameters are listed in Table 4. We used IEEE 802.11 as the MAC layer protocol. A packet was considered receivable if its reception power was above a threshold value, called the reception power threshold (RX_{thresh}). Likewise, if a packet was received and its power was above the carrier sensing power threshold (CS_{thresh}), the channel was sensed busy during the packet transmission. Given a threshold value, transmission power and propagation model, a specific maximum distance for packet reception or detection was determined [28]. For our simulations, the propagation model was two ray ground and no obstacles were considered. We utilized this simple propagation

Parameter	Value	Parameter	Value
Simulator	NS-2	Queue Size	50 packets
Propagation Model	Two Ray Ground	Data Packet Size	512 bytes
Antenna	Omni Directional	CBR Data Rate	128 kbps
MAC Protocol	IEEE 802.11	Packets per second	31.25
Transmission Power	30mW	Network Area	1000m x 1000m
Frequency	2.4GHz	Mobility Model	Random Waypoint
MAC Layer Data Rate	2 Mbps	Speed	0-5 m/s
Reception Range	250m	Pause Time	20 s
Carrier Sensing Range	550m	Number of nodes	50
Capture Factor	10.0	Simulation Time	200 s
Receiver Interference Distance	440m	Number of Runs	10

Table 4. Simulation parameters.

model instead of a more realistic propagation model to enable deterministic ranges and simplify analysis. This propagation model results in a RxR of 250m and a MAC layer carrier sensing range (CSR) of 550m. This MAC layer CSR does allow some spatial reuse to take place in the simulations.

The reception power threshold, propagation model, and capture factor must be known to determine the RID. The capture factor defines the minimum power ratio between the received power of two packets such that the packet with the higher power can be received successfully. The capture factor is a hardware specific value; for our simulations, we used 10.0. To further explain the calculation of RID, we provide the following example: Given a packet received with the minimum reception power (RX_{thresh}) and a second packet that is transmitted simultaneously, the received signal strength of the second packet must be less than $RX_{thresh}/10.0$ for the first packet to be successfully received. Otherwise, neither packet can be decoded by the receiver. Given our simulation parameters, if the sender and receiver were separated by RxR, another sender must be at least 440m away for its transmission to be able to take place simultaneously. Therefore, the RID was 440m for our simulations; at this distance, the received power of another sender was guaranteed to be less than $RX_{thresh}/10.0$.

With a RxR of 250m, a RID of 440m, and IEEE 802.11, the sensing range for PAC was 940m, as calculated by Equation 1. Given the propagation model and other simulation parameters, we calculated the minimum reception power threshold at this distance [28]. In our simulations, if a packet was received with a power above this threshold value, the packet was considered in the available bandwidth calculation. The carrier sensing mechanism for the MAC layer behaved as if the minimum reception threshold was not modified. If the carrier sensing mechanism were changed, the collision avoidance attributes, spatial reuse [29–31] and medium access [32] would be affected.

PAC Sensing Range	940 m
Busy Time Window Size	250 ms
B_{max}	1200 kbps
B_{rsv}	240 kbps
B_{min}	120 kbps
T_{retry}	1 to 2 seconds
$T_{backoff}$	1 to 2 seconds

Table 5. PAC parameters.

Table 5 lists the parameter values used by PAC in our simulations. To perform the available bandwidth calculation, a maximum effective bandwidth (B_{max}) of 1200 kbps was assumed. We determined this value experimentally and it is close to the analytical value derived in [10]. We reserved 20% (240 kbps) of the maximum bandwidth to avoid congestion, allow for temporary fluctuations, and detect mobility before congestion. The same reserved bandwidth was used during simulation of CACP. In PAC, if the detected available bandwidth dropped below 120 kbps (10% of the maximum bandwidth), we assumed over-utilization was imminent. We utilized a sliding window to calculate the PAC-based available bandwidth. The size of the window was 250 ms. We found this window size sufficient to quickly adjust the available bandwidth according to the usage of admitted flows, but still a large enough time scale to avoid overreacting to a short burst of packets. In our simulations, the backoff time between flow admission attempts after flow rejection was between 1 and 2 seconds. The time interval between congestion detection checks was also between 1 and 2 seconds. The simulation results show these values were adequate; no two flows were rejected in response to any imminent congestion event. Optimization of the backoff and detection intervals are left for future work.

4.2 Local Admission Control Performance

In our experiments, we studied networks where the sender and receiver are always within range of each other to emphasize the effect of the admission control decision. Under these conditions no routing protocol was needed. There were 25 sender-receiver pairs and every five seconds another sender started sending traffic. Therefore, after 125 seconds of simulation time, all senders were active. The traffic was modeled as a CBR flow with 512-byte packets at a rate of 128 kbps. This traffic model represents a multimedia flow.

A summary of the results is presented in Table 6. Without admission control, all 25 sender-receiver pairs became active transmitters. With admission control, only 12 were admitted. By limiting the traffic in the network, PAC and CACP kept the network from becoming congested. By avoiding congestion, all admitted traffic

Admission Control Protocol	Number of Flows Admitted	Packet Losses	Packets Delivered	Average Delay (s)	Utilization (% busy)	Standard Deviation of Packets Received Per Second Per Flow
None	25	26778	81825	0.973	97%	6.10
PAC	12	0	58173	0.005	80%	0.79
CACP	12	0	51182	0.004	78%	0.75

Table 6. Overall admission control performance results.

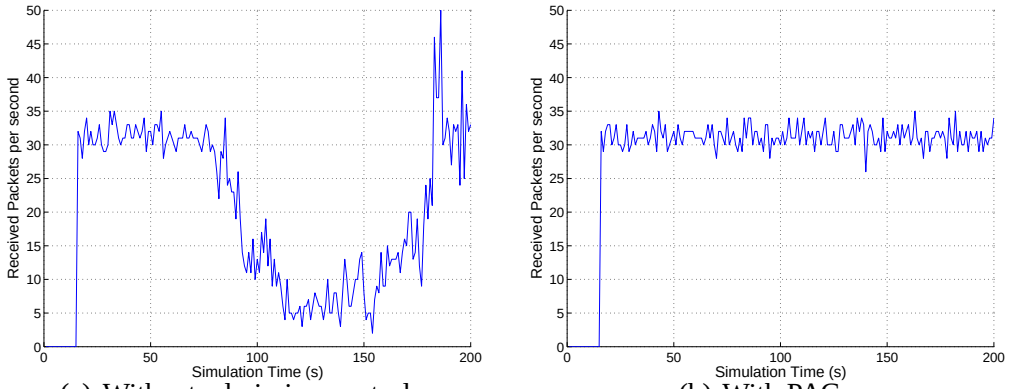

(a) Without admission control. (b) With PAC.

Fig. 11. Throughput of a single representative receiver in one particular simulation run.

received a high packet delivery ratio (nearly 100%) and low delay. Without admission control, significant packet loss and delay occurred due to congestion and queue overflow. Notice that the number of packets delivered with PAC enabled was approximately 20% lower than the number of packets delivered without admission control. This directly relates to the bandwidth reserved to detected congestion, as presented in Table 5. By decreasing B_{rsv}, more aggregate throughput would be available for admitted traffic, but more flows would be throttled or rejected as nodes move around the network.

In addition to packet loss and delay, Table 6 shows the standard deviation of packets received per second per source while all admitted sources are active. Note that a single flow transmitted approximately 30 packets per second. The standard deviation without admission control was very large and each flow received markedly different QoS. With admission control the standard deviation was low, indicating all admitted flows received a high QoS.

Figure 11(a) shows the number of packets successfully received per second for a single receiver during one simulation. In this graph, admission control was not used. The graph illustrates that, as the simulation progressed and more sources became active, the channel became congested. After 80 seconds elapsed, the throughput for this receiver decreased significantly due to congestion. At 180 seconds, the sending

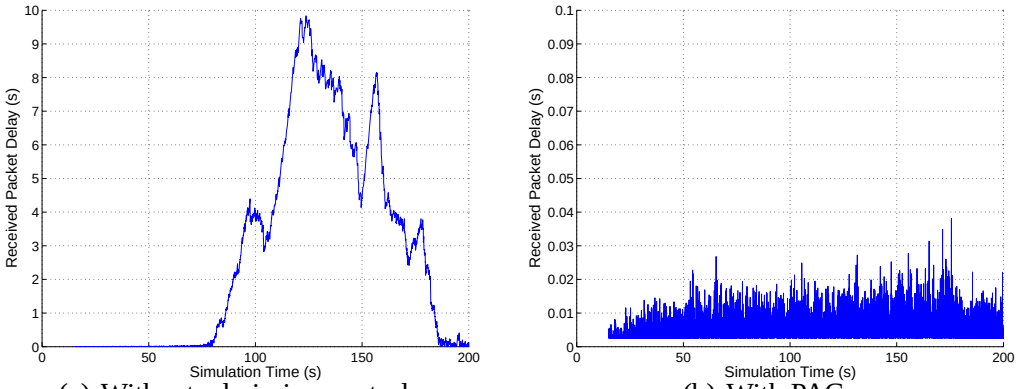

(a) Without admission control. (b) With PAC.

Fig. 12. Delay of a single representative receiver in one particular simulation run.

node gained an unfair advantage in channel access and the receiver again experienced acceptable throughput. The short-term unfairness is a well known behavior in IEEE 802.11 [33]. Unfair channel access resulted in a spike in throughput as queued packets were delivered. In addition to experiencing degraded throughput for most of the simulation, the delay experienced by received packets was often unacceptable for real-time multimedia applications. Figure 12(a) presents the delay for the received packets without admission control. Once the channel became congested, the delay value increased sharply. The delay experienced was particularly high since all packets traversed only a single hop from the source to destination.

In contrast to the poor performance without admission control, PAC enabled admitted sessions to experience a consistently high QoS. Figure 11(b) shows the number of packets received per second for the same receiver as in Figure 11(a). Similarly, Figure 12(a) shows the delay for the same receiver with PAC. These figures show that traffic throughput for this session was nearly constant with PAC. In addition, the delay was extremely small and constant. Note that the difference in the scale of the y-axis between Figures 12(a) and 12(b) is two orders of magnitude. The short packet delay, consistent packet delivery rate, and low packet loss statistics demonstrate that PAC can be used for networks to sustain real-time traffic applications, such as voice or multimedia. The results displayed for this particular flow are characteristic of other flows in the simulation.

In addition to the throughput and delay experienced by a single flow, the performance experienced by all flows is important. Figure 13 shows the packet receptions per second for all traffic flows with and without PAC; each vertical line represents the start of a new flow. In Figure 13(a), we see that without admission control each flow experienced notably different throughput. In contrast, with PAC, each admitted flow experienced nearly the same throughput, as shown in Figure 13(b). This was possible because PAC limits the number of admitted flows.

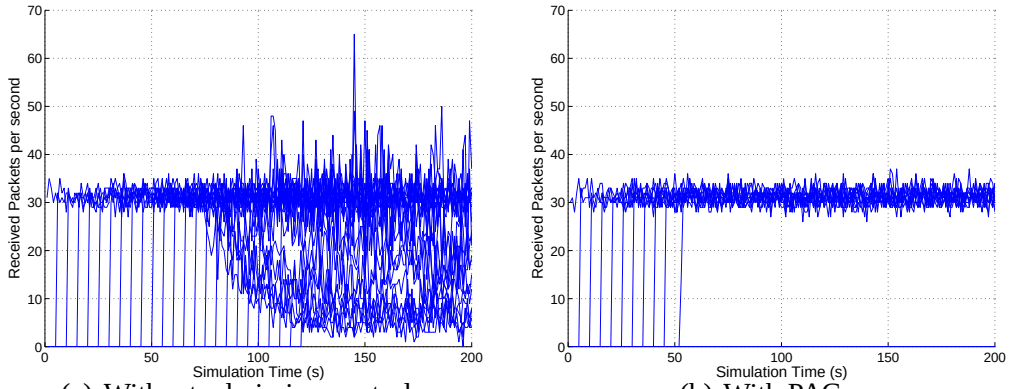

(a) Without admission control. (b) With PAC.

Fig. 13. Throughput for all flows

In terms of delay and throughput for admitted flows, CACP performs similarly to PAC, as shown in Table 6. One difference was the number of packets delivered. Since CACP has more messaging overhead for every admission decision attempt, a part of the bandwidth was consumed by overhead that would otherwise be available for data packet delivery.

In the random network topologies simulated, CACP performed well. There are several network topologies and conditions where CACP performs improperly or is overly conservative. We discuss these issues in the next section.

To summarize the results of these simulations, PAC was able to minimize packet loss and delay using admission control to ensure the channel did not become congested. Further, the bandwidth was fairly shared between all admitted flows. Without PAC, the channel was susceptible to congestion, resulting in large packet loss and delay.

5 Qualitative Comparison

Although CACP performs well in some cases, the protocol has many weaknesses. In this section, we present general scenarios where the performance of CACP degrades and describe how these scenarios are addressed in PAC.

5.1 Control Packet Losses and Erroneous Admission Decisions

Although CACP performs well in some cases, the protocol has many weaknesses. The most important weakness of CACP is that it may erroneously admit new flows

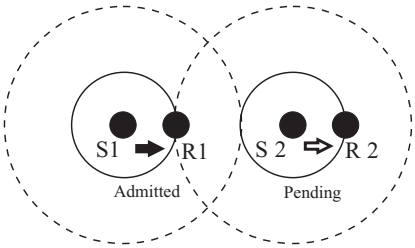

Fig. 14. In this scenario, the likelihood of an erroneous admission decision in CACP is proportional to the load induced by node S1.

when the network load is high. Prior to the admission of a new flow, CACP performs a local available bandwidth check. If enough local available bandwidth exists, the sender broadcasts a query to all of its CSN. If no rejection packet is received within a short period of time, the new flow is admitted. The reliance on a rejection message, which is essentially a negative acknowledgment, results in a poor default failure condition. For example, if a query or rejection message is lost (i.e. due to collision), the sender may make an incorrect admission decision by admitting more traffic than the channel can accommodate. Additionally, since query and rejection messages are sent using high power, the probability of collision is correlated with the utilization in the area around its CSN.

Consider the network shown in Figure 14. Node S2 is a sender that is attempting to admit a new flow. Node S1 is currently transmitting to node R1. Since node S1 is outside the CSR of node S2, it sends packets without regard for the state of node S2. Similarly, because node S2 is outside the CSR of node S1, it also sends packets without regard for node S1. Given this network, the probability that a query packet from node S2 collides with a data transmission at node R1 is directly proportional to the amount of traffic node S1 is sending to node R1.

To further investigate this behavior, we performed a set of simulations. In the simulations, the network was configured as shown in Figure 14 and the CBR traffic rate from node S1 to node R1 was varied from zero to 2 Mbps, the maximum data rate. Node S2 attempted to admit a new traffic flow that requires 5 Mbps, which was more bandwidth than was available. Therefore, in this scenario, the new flow should never be admitted. Figure 15 illustrates that as the flow rate from node S1 to node R1 increased, the frequency at which node S2 erroneously admitted new flows also increased. Each data point represents an average of 40 admission decision attempts. When the channel was highly loaded, CACP almost always wrongly admitted the new traffic flow. This faulty admission control decision occurred because of the reliance on a rejection message from a CSN during the query-reject mechanism.

In this example network, PAC's available bandwidth measurement would include node S1's traffic and no message exchange would be required. Since PAC's avail-

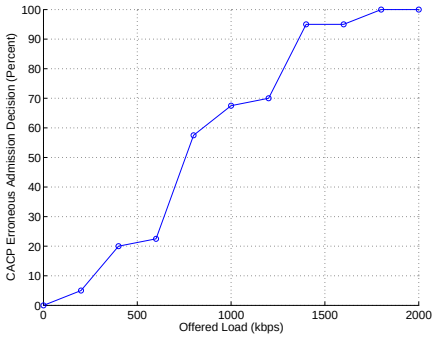

Fig. 15. CACP erroneous admission control decision as load increases.

able bandwidth measure considers all senders that might be impacted by the addition of a new traffic flow, PAC ensures that the correct admission decision is made.

5.2 Lack of Mobility Support

CACP does not address node mobility. To achieve acceptable performance in mobile networks, CACP depends on two mechanisms: conservative reservation of network capacity and route errors. CACP reserves extra capacity (B_{rsv}) to allow some flows to move within range of each other without causing network congestion.

If B_{rsv} is insufficient, or if the network becomes congested for any reason, CACP relies on the routing protocol to generate a route error since the routing protocol believes a neighbor has moved and the link to that neighbor is now broken. In CACP, link breaks are detected by the inability to transmit a unicast packet to its next hop. If congestion occurs, a packet fails to be sent to its next hop and the routing protocol issues a route error. The route error removes the route and eventually causes the source to re-initiate the admission control procedure. The use of a route break to indicate congestion is undesirable because it requires the channel to become highly congested and packet loss to occur before the error message is generated. Note that by the time the routing protocol generates the error message, the QoS has already degraded significantly.

To handle mobility, PAC detects the onset of congestion by monitoring the available bandwidth. In PAC, when a source detects that congestion is about to ensue, it throttles or stops enough of its admitted data flows to avoid network congestion. This approach is not feasible in CACP since it would be too expensive to proactively monitor the available bandwidth. In CACP, handling mobility would require many periodic message exchanges at high power, and when the network load is high, CACP would make an incorrect measurement with a high probability.

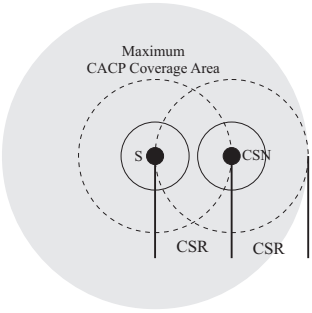

Fig. 16. CACP queries all CSN to check whether there is enough available bandwidth to support a new flow. This may result in an area up to $2 * CSR$ being queried.

5.3 Spatial Reuse

Another issue addressed by PAC is spatial reuse. In CACP, the measurement range considered by the admission control query-reject messages may be as large as $2 * CSR$ in dense networks, as shown in Figure 16. Initially, the source checks the available bandwidth within its CSR. Then the source queries all its CSN, which are at most the distance of CSR away. The CSN check the available bandwidth within their CSR. This range, indicated in Figure 16, is larger than needed to make the correct admission decision. The minimum distance between two simultaneously transmitting sources to prevent receiver interference is $RxR + RID$ (or $2 * RxR + RID$, in CSMA networks that utilize ACKs). Therefore, in a network as shown in Figure 3, if a node exists inside the CSR of both sources (the shaded region), CACP will not allow two simultaneous flows to be admitted if each flow consumes more than half of the network bandwidth. In contrast to CACP, PAC allows both flows to be admitted, resulting in twice the aggregate network throughput.

6 Experimental Evaluation

In this section, we present our PAC implementation and the experimental results from a simple scenario. In order to experimentally evaluate PAC, the MAC layer sensing information must be exposed for monitoring. To date, only one mass produced hardware platform provides this information, the Berkeley/Crossbow mote [2] . Therefore, we implemented PAC on the mote platform. In the rest of this section, we discuss the mote hardware, mote sensing capabilities and the Tiny OS PAC implementation.

The Berkeley Mica2 mote consists of a 7.38MHz Atmel 128 microprocessor with a

[2] Crossbow Inc. Mica2 Sensor Platform. http://www.xbow.com/.

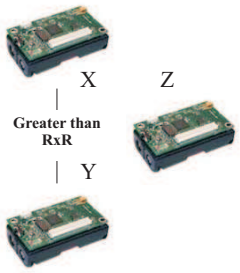

Fig. 17. Mote network scenario for experimental tests.

4KB EEPROM, 128KB program flash memory and 512KB flash data memory. For wireless communication, a CC1000 [3] radio operates at 914MHz. In addition to the processing/radio unit, a sensor board capable of monitoring light, sound, motion, and temperature is attached.

Tiny OS [4] is the operating system for the mote platform. Tiny OS provides core code to access most of the capabilities of the mote and sensor board. Of importance to our application, a CSMA MAC layer and a simple interface to the received signal strength indicator (RSSI) are available. A high RSSI indicates that a nearby node is transmitting. The RSSI is used by the MAC layer to perform carrier sensing and avoid collisions.

PAC requires that each node sense transmissions at a range of at least $RxR + RID$. To ensure this property held for our experiments, the MAC layer was modified to reduce the RxR. By only accepting packets with a RSSI value above a certain threshold a short RxR was ensured. Limiting the RxR allowed the sensing requirements for PAC to be met. Since motes excel at capturing nearby packets [34] the RID is also short. The proper RSSI threshold value to ensure PAC's sensing requirements was experimentally determined and well above the noise floor.

To calculate the available bandwidth using the motes, each node monitored the status of the wireless channel using the MAC state (i.e. transmitting, receiving, busy, or idle) and the measured RSSI during each byte communication interval. To accommodate the limited processing capabilities of the motes, the number of idle intervals was counted at a rate of 2400Hz over each quarter second window (slot). Several of these measurements were then averaged using a moving weighted average. This utilization information was then used to calculate the available bandwidth. With the available bandwidth calculation, each node determined whether a new flow could be supported by the network without interfering with any on-going flows.

[3] Chipcon CC1000. http://www.chipcon.com/.
[4] University of California, Berkeley. Tiny OS. http://www.tinyos.net/.

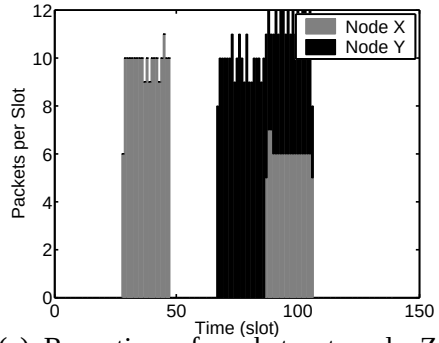

(a) Reception of packets at node Z without admission control.

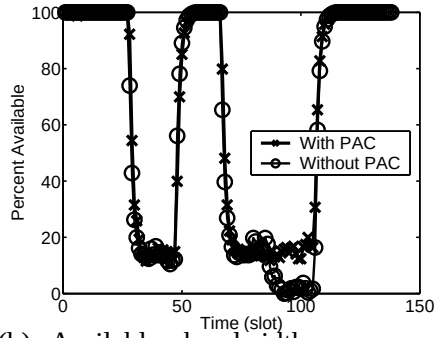

(b) Available bandwidth measurements during the experiment.

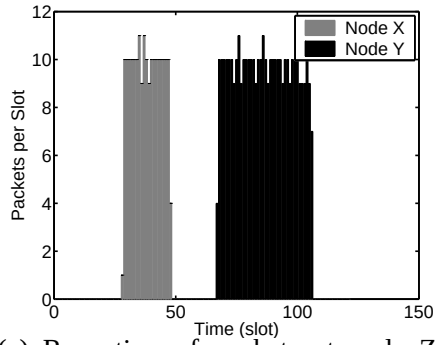

(c) Reception of packets at node Z with PAC.

Fig. 18. Mote experimental results.

6.1 Mote Experiments

To explore the effectiveness of PAC, a simple application scenario was emulated. Three nodes were all located within $RxR + RID$ of each other. Nodes X and Y were sensor nodes and were separated by more than RxR but less than the $RxR + RID$, as shown in Figure 17. Both nodes X and Y monitored their microphone sensor and, upon detecting an interesting sound, attempted to transmit the sound data stream in real-time to node Z. Since the bandwidth requirement of a sound data stream flow was approximately 80% of the maximum achievable bandwidth, two simultaneous flows to node Z could not be supported in this scenario. When both nodes X and Y transmitted their streams at the same time, significant packet loss occurred, preventing either flow from being useful to node Z. To enable a high QoS for admitted flows, PAC was used by nodes X and Y to perform admission control prior to transmitting a packet flow to node Z. If the available bandwidth was too low, a node did not admit its data flow.

During the experiment, node X detected a sound from time slot 25 until slot 45 (5 seconds). From time slot 65 to 105, node Y detected a sound. Also, during slots 85

to 105 node X detected a sound. Without admission control, both nodes X and Y transmitted during slots 85 to 105. This simultaneous transmission of both flows resulted in significant packet loss since neither sender could achieve the required throughput of approximately ten packets per slot, as shown in Figure 18(a). Since node Z received approximately half of the packets generated by each stream, neither stream was able to be reconstructed and the sound data was discarded.

When PAC was utilized in the same experiments, the sources checked their available bandwidth measurement prior to admitting a data stream. Figure 18(b) shows the percentage of the bandwidth available at node Z during the test scenario with and without PAC. Since all nodes were within $RxR + RID$ of each other, their available bandwidth measurements were nearly the same. At time 25, the available bandwidth dropped below 20%. This drop was due to node X's sound data flow transmissions that consumed more than 80% of the bandwidth. After node X stopped transmitting, the available bandwidth again returned to 100%. When node Y started transmitting at time 65, the available bandwidth again dropped below 20%.

At time 85 node X attempted to admit its new data flow. Node X decided that the wireless channel could not support the new data stream since the available bandwidth was less than 20% and the new flow required more than 80%. Therefore, node X rejected the new flow's admission request. This rejection of admission for the flow from node X allowed node Y to continue sending its data to node Z without congestion and packet loss.

By rejecting node X's new flow at time 85, congestion and the resulting packet loss were avoided. By avoiding congestion and packet loss, node Z properly received all node Y's packets in a timely manner, as shown in Figure 18(c). Since all of node Y's packets were received, node Z could reconstruct the sound stream and correctly process the data.

The mote implementation proves that PAC can be implemented on off-the-shelf hardware. We expect that given the ability to determine the RSSI on a slot-by-slot basis, PAC can easily be developed on IEEE 802.11 [5] and IEEE 802.15.4 [6] hardware with little or no change to the sensing behavior. The experimental results show that PAC can ensure high quality service in real networks with realistic propagation effects, such as multipath and fading.

7 Conclusion

Our Perceptive Admission Control protocol for use in wireless mobile networks addresses two issues often overlooked by previous research: shared wireless bandwidth and node mobility. PAC is able to correctly compute the available bandwidth by sensing all important nearby transmissions.

In addition to the direct application of PAC, the utilization calculation is applicable to an even wider set of problems. The available bandwidth calculation can be used in conjunction with an advanced packet dropping strategy, such as RED [35]. By working with known TCP behavior, this packet dropping strategy should increase overall performance and fairness in wireless networks even with node mobility.

Another opportunity for increased control in the available bandwidth calculation is the inclusion of received signal strength. Nodes can calculate the available bandwidth using various sensing ranges simultaneously. By examining the available bandwidth with various thresholds, nodes can determine the amount of traffic and approximate distance to those traffic sources. This usage information can be used for admission control, congestion control, or to dynamically adjust MAC layer parameters, such as CSR or transmit power.

Additional mechanisms that utilize carrier signal information (such as perceptive behaviors [4]) can be coupled with PAC to support additional network features. For example, supporting bursty traffic sources, handling multiple priority traffic, and calculating the contention count [27]. These and several other communication and coordination problems can be addressed using perceptive behavior.

Acknowledgments

This work was supported in part by NSF Career Award CNS-0347886, by NSF NeTS Award CNS-0435527, and by a gift from the Intel Corporation.

References

[1] L. Zhang, S. Deering, D. Estrin, RSVP: A New Resource ReSerVation Protocol, IEEE Networks 7 (5) (1993) 8–18.

[2] K. Nichols, S. Blake, F. Baker, D. Black, Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers, RFC 2474.

[3] P. Garg, R. Doshi, R. Greene, M. Baker, M. Malek, M. Cheng, Achieving Higher Throughput and QoS in 802.11 Wireless LANs, in: Proceedings of the International Performance Computing and Communications Conference (IPCCC), Phoenix, AZ, 2003.

[4] I. D. Chakeres, Supporting Multimedia in Wireless Multhop Networks, Ph.D. thesis, University of California Santa Barbara, Santa Barbara, CA (July 2005).

[5] IEEE Computer Society, IEEE 802.11 Standard, IEEE Standard For Information Technology (1999).

[6] IEEE Computer Society, IEEE 802.15.4 Standard, IEEE Standard For Information Technology (2003).

[7] IEEE Computer Society, IEEE 802.11b Standard, IEEE Standard For Information Technology (1999).

[8] IEEE Computer Society, IEEE 802.11g Standard, IEEE Standard For Information Technology (2003).

[9] R. S. Prasad, M. Murray, C. Dovrolis, K. Claffy, Bandwidth Estimation: Metrics, Measurement Techniques and Tools, IEEE Network 17 (6) (2003) 27–35.

[10] G. Bianchi, Performance Analysis of the IEEE 802.11 Distributed Coordination Function, IEEE Journal Selected Areas in Communications 18 (3) (2000) 535–547.

[11] K. Fall, K. Varadhan, ns Manual, http://www.isi.edu/nsnam/ns/doc/The VINT Project.

[12] T.-W. Chen, J. T. Tsai, M. Gerla, QoS Routing Performance in Multihop Multimedia Wireless Networks, in: Proceedings of IEEE International Conference on Universal Personal Communications (ICUPC), San Diego, CA, 1997.

[13] Y.-C. Hsu, T.-C. Tsai, Bandwidth Routing in Multihop Packet Radio Environment, in: Proceedings of the 3^{rd} International Mobile Computing Workshop, 1997.

[14] C. R. Lin, J.-S. Liu, QoS Routing in Ad hoc Wireless Networks, IEEE Journal on Selected Areas in Communications 17 (8).

[15] C. Zhu, M. S. Corson, QoS Routing for Mobile Ad hoc Networks, in: Proceedings of the IEEE Conference on Computer Communications (INFOCOM), New York, NY, 2002.

[16] H. Luo, S. Lu, V. Bharghavan, J. Cheng, G. Zhong, A Packet Scheduling Approach to QoS support in Multihop Wireless Networks, ACM Journal of Mobile Networks and Applications, Special Issue on QoS in Heterogeneous Wireless Networks 9 (3).

[17] R. Ramanathan, M. Steenstrup, Hierarchically-organized, Multihop Mobile Wireless Networks for Quality-of-service Support, ACM Journal of Mobile Networks and Applications 3 (1).

[18] R. Sivakumar, P. Sinha, V. Bharghavan, CEDAR: a Core-Extraction Distributed Ad hoc Routing Algorithm, IEEE Journal on Selected Areas in Communications 17 (8).

[19] S. Murthy, J. J. Garcia-Luna-Aceves, A Routing Architecture for Mobile Integrated Services Networks, ACM Journal of Mobile Networks and Applications 3 (4).

[20] M. Grossglauser, D. N. C. Tse, A Time-Scale Decomposition Approach to Measurement-Based Admission Control, IEEE/ACM Transactions on Networking 11 (4) (2003) 550–563.

[21] S. Jamin, P. B. Danzig, S. J. Shenker, L. Zhang, A Measurement-based Admission Control Algorithm for Integrated Services Packet Networks, IEEE/ACM Transactions on Networking 5 (1) (1997) 56–70.

[22] K. Shiomoto, N. Yamanaka, T. Takahashi, Overview of Measurement-Based Connection Admission Control Methods in ATM Networks, IEEE Communications Surveys (1999) 2–13.

[23] S.-B. Lee, G.-S. Ahn, X. Zhang, A. Campbell, INSIGNIA: An IP-Based Quality of Service Framework for Mobile Ad Hoc Networks, Journal of Parallel and Distributed Computing, Special Issue on Wireless and Mobile Computing and Communications 60 (4) (2000) 374–406.

[24] G.-S. Ahn, A. Campbell, A. Veres, L.-H. Sun, SWAN: Service Differentiation in Stateless Wireless Ad hoc Networks, in: Proceedings of the IEEE Conference on Computer Communications (INFOCOM), New York, NY, 2002.

[25] Y. Yang, R. Kravets, Contention-aware Admission Control Protocol for Ad Hoc Networks, Tech. Rep. 2003-2337, University of Illinois at Urbana-Champaign (April 2003).

[26] A. Lindgren, E. M. Belding-Royer, Multi-Path Admission Control for Mobile Ad hoc Networks, in: Proceedings of Mobiquitous, 2005.

[27] K. Sanzgiri, I. D. Chakeres, E. M. Belding-Royer, Pre-Reply Probe and Route Request Tail: Approaches to Calculate Intra-Flow Contention in Multihop Wireless Networks, To appear in a special issue of the ACM/Kluwer Mobile Networks and Applications Journal (MONET).

[28] W. Ye, Calculating the receiving threshold, RXThresh_ for Phy/Wireless, threshold.cc, NS-2.27 (2000).

[29] X. Guo, S. Roy, S. Conner, Spatial Reuse in Wireless Ad hoc Networks, in: Proceedings of the IEEE Vehicular Technology Conference (VTC), Orlando, FL, 2003, pp. 1437–1442.

[30] F. Ye, B. Sikdar, Improving Spatial Reuse in IEEE 802.11 Based Ad hoc Networks, in: IEEE Global Telecommunications Conference (GLOBECOM), San Francisco, CA, 2003, pp. 1013–1017.

[31] J. Zhu, X. Guo, L. Yang, S. Conner, Leveraging Spatial Reuse in 802.11 Mesh Networks with Enhanced Physical Carrier Sensing, in: Proceedings of the IEEE International Conference on Communications (ICC), Paris, France, 2004.

[32] Z. Li, S. Nandi, A. K. Gupta, Improving Fairness in IEEE 802.11 based MANETs using Enhanced Carrier Sensing, in: Proceedings of NETWORKING, Athens, Greece, 2004.

[33] T. Nandagopal, T. Kim, X. Gao, V. Bharghavan, Achieving MAC Layer Fairness in Wireless Packet Networks, in: Proceedings of the 6^{th} Annual International Conference on Mobile Computing and Networking (MobiCom), Boston, MA, 2000, pp. 87–98.

[34] A. Woo, K. Whitehouse, F. Jiang, J. Polastre, D. Culler, The Shadowing Phenomenon: Implications of Receiving During a Collision, Tech. Rep. UCB//CSD-04-1313, University of California, Berkeley (March 2004).

[35] S. Floyd, V. Jacobson, Random Early Detection Gateways for Congestion Avoidance, IEEE/ACM Transactions on Networking 1 (4) (1993) 397–413.