

PAC: Perceptive Admission Control for Mobile Wireless Networks

Ian D. Chakeres
Dept. of Electrical & Computer Engineering
University of California, Santa Barbara
idc@engineering.ucsb.edu

Elizabeth M. Belding-Royer
Dept. of Computer Science
University of California, Santa Barbara
ebelding@cs.ucsb.edu

Abstract

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 approach is required. As wireless networks become more widely used, there is a growing need to support advanced services, such as multimedia streaming and voice over IP. 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 data packets waste bandwidth and hinder other traffic, compounding the problem. To address this issue, 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 and simulations, we show that PAC achieves this goal and ensures low loss and delay for all admitted flows.

1. Introduction

Wireless devices are becoming prevalent because of their ability to provide mobile communication. Since many common applications, including voice and multimedia, require low packet loss and delay, quality of service (QoS) is becoming an important requirement for these networks. In contrast to traditional wired networks, mobile networks operate under harsh conditions that include mobility, a shared wireless channel and limited bandwidth. Traditional attempts to provide guaranteed QoS are unable to cope with the constantly changing network conditions. Meeting hard real-time QoS constraints in wireless mobile networks is unrealistic because of node mobility and shared medium access. Instead, solutions that provide a stateless service and offer better than best-effort packet delivery for high priority packets are more successful, such as DiffServ and IEEE 802.11e. Unfortunately, these solutions still fail to provide

the low loss and delay that real-time applications require if the network becomes congested.

High quality of service without fully coordinated channel and network access is achievable. The wireless channel must be kept from reaching the congestion point, since loss and delay increase rapidly once this point is reached. Maintaining the utilization below the congestion point is difficult because the channel is shared between nodes that may not be able to communicate directly; therefore, nodes need to passively determine the network utilization. Once the amount of available bandwidth is determined, nodes can then adapt their data traffic to keep the channel from becoming congested.

We propose the Perceptive Admission Control (PAC) protocol to control the amount of traffic in the network and provide high quality service to all admitted traffic. PAC ensures the network congestion point is not reached through the requirement of call admission for all new flows. To make an admission decision, PAC considers not only the limited area within a sender's transmission range, but the entire area that a new flow may impact. We show that the time that the wireless channel is sensed as busy is a good estimator of available bandwidth. Using this measure, PAC performs admission control for new flows to avoid congestion. We begin our discussion by focusing on single hop admission control. We then describe how to easily extend PAC for multihop paths.

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. In Section 4 we demonstrate the performance of PAC in simulation and describe how it avoids the shortcomings of previous approaches. Finally, Section 5 concludes the paper.

2. Background

To perform admission control in wireless networks it is important to understand how a wireless transmission impacts other nodes. In Section 2.1 we describe the important distances for packet transmission and reception. Since ad-




Figure 1. Approximation of reception range (RxR) and carrier sensing range (CSR). Nodes within reception range are called neighbors (N), while carrier sensing neighbors (CSN) are all nodes within carrier sensing range.

mission control decisions depend on accurate estimation of the available bandwidth, we examine several methods for calculating the available bandwidth in Section 2.2. In Section 2.3 we categorize related work and discuss why most proposed solutions are insufficient. In Section 2.4 we describe the solution most closely related to our proposed approach.

2.1. Impacted Area

For admission control purposes, there are multiple notable ranges for wireless communication. Each distance is important for the measuring channel utilization and predicting the available bandwidth. At a 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 RxR , as shown in Figure 1. Nodes within RxR of a particular sender are considered its neighbors (N).

Nodes that are within carrier sensing range of a sender can sense packet transmissions. The nodes inside a sender's carrier sensing range are called carrier sensing neighbors (CSN). These nodes detect a transmission but may not be able to decode the packet. The maximum distance that a node can detect an ongoing packet transmission (carrier signal) is called the carrier sensing range (CSR). This range is typically much larger than the reception range. In wireless MAC protocols based on CSMA, such as IEEE 802.11, all CSN of the sender are unable to initiate a packet transmission while the sender 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.

When a carrier signal is sensed by a receiver, packet reception from another sender may not be impacted. For correct packet reception, the area surrounding a receiver must be free of multiple interfering transmissions. If another node transmits a packet close to the receiver it may interfere with




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

an ongoing packet reception, even if the two senders are outside each others carrier sensing range. 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 at the same time as node S , node R is unable to successfully receive packets from node S . In both cases, node X is not prohibited from transmitting because node S is outside its carrier sensing range; 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 abilities. Note that CSR (dashed line) is larger than RID (dotted line) and RID is larger than RxR (solid line), as shown in Figure 2. These line styles will be used throughout the paper to denote the different ranges.

For two simultaneous transmissions transmitted to be successfully received by different receivers, the transmissions (and nodes) must be separated in space. The distance between two senders to ensure proper packet reception at a receiver is $RxR + RID$. This distance holds for all possible network scenarios. 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. If the distance is larger than $RxR + RID$, by definition, the receiver and another sender cannot be closer than RID .

These communication distances are for networks where all nodes use omnidirectional antennas and transmit packets with the same transmission power on the same channel. Further we assume that there are no obstacles and only simple fading occurs. We plan on exploring relaxation of these conditions as future work.




Figure 3. Spacing requirement for simultaneous transmissions in wireless networks that utilize acknowledgments. Since the sender's RID and receiver's CSR are not important to the interference calculation, they are not shown.

2.1.1. MAC Layer Acknowledgments

Acknowledgments (ACKs) are used in many MAC protocols, such as IEEE 802.11, to immediately inform the sender that successful reception has occurred. If an ACK is not received the sender will retransmit the packet a maximum number of times. The Data-ACK mechanism is used to combat packet loss at the MAC layer due to collisions and errors introduced by the wireless channel. Generally carrier sensing is not performed by the receiver prior to sending an ACK. This is because carrier sensing might silence a receiver, upon successful data reception, and therefore require the sender to retransmit the packet. This in turn would waste wireless resources and power and increase delay.

When receivers do not perform carrier sensing prior to sending an ACK after successful data reception, the receivers must also be separated by RID. 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 collisions. These collisions 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

$$2 * RxR + RID \quad (1)$$

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

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 the available bandwidth (B_{avail}) is to measure the network utilization (U). Given the network utilization and the maximum bandwidth (B_{max}),

the available bandwidth is estimated using the following equation [11]:

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

where $0 \leq U \leq 1$. There are many techniques to measure the network utilization. Some metrics of network utilization are:

- MAC Layer Congestion Window
- Queue Length
- Number of Collisions
- Delay
- Channel Busy Time

The first three methods provide little or no information regarding network utilization if a node is not actively transmitting packets. For example, a collision only occurs if a packet fails to send. If a node does not send any packets, it cannot determine the current state of the channel. The same holds true for the MAC layer congestion window and the queue length. Since these techniques are not adequate for determining the available bandwidth, we explore the two remaining techniques, delay and channel busy time, in more detail.

Both delay and channel busy time can be used to determine the current bandwidth usage; however, channel busy time has several advantages. These include no additional overhead, no measurement gaps and adaptable measurement range. The simulation results comparing these two techniques are omitted due to space limitations; see [2] for additional detail.

Channel busy time is a direct measure of the channel utilization. In wireless networks, carrier sensing enables nodes to detect three states; transmitting, receiving and busy. If the node detects a carrier signal it senses that the channel is busy, but it is only able to decode the packet contents if it is within RxR. By measuring the amount of time the channel is sensed busy (CS), sending (TX) or receiving (RX), a node can measure not only transmissions that occur within its reception range, but also those within its carrier sensing range. Using this metric, more transmissions result in a busier channel. 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.

With any measurement technique it is common that instantaneous values vary, sometimes widely. For our approach we utilize an equally weighted sliding window to obtain the wireless utilization. Through testing, we determined a window size that was large enough to make 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 better estimation of the utilization and available bandwidth.




Figure 4. Network scenario with unreachable CSN. In this figure node S1 cannot contact node S2 via any multihop path.

2.3. 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) [3, 5, 7, 17], network wide information dissemination [8, 12, 13] or reservations [6, 9, 16] do not work well in mobile networks where the network topology changes frequently.

The contention-aware admission control protocol (CACP) [15] 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 Section 2.4, and qualitatively compared with our solution in Section 4.3.

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 receivers 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. Contention-Aware Admission Control Protocol

When admission control decisions are made in CACP, each node considers not only the resources of its immediate neighborhood, but the resources of all nodes within its carrier sensing range. CACP is contention-aware in that each node passively monitors the amount of time the channel is sensed as busy. This includes the time a carrier signal is detected, as well as when a packet is transmitted or received. The available bandwidth is calculated 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 flow over one hop 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 carrier sensing range. If all CSN detect enough available bandwidth then the flow is admitted.

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 network scenarios where a node within carrier sensing range is not reachable via any path. For example, in Figure 4 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 carrier sensing range 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 using a high power packet transmission.

To better explore CACP operation, an example is provided. In the network in Figure 5(a), 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 1(a) at time T1. Only nodes X, Y and Z detect the current flow; node W does not detect the communication between Z and Y since it is outside of measurement range, CSR. 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 carrier sensing range, i.e. nodes X and Y. Both X and Y check their available bandwidth measurement. Since enough bandwidth is available, they do not send a rejection message to node W. After a timeout, node W admits the new traffic flow. After a short time, shown as time T2 in Table 1(a), the available bandwidth measurement of each node adjusts to 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 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, the pending admission request is denied.

Though we do not focus on multihop networks in this paper, 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 carrier sensing range 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. For a detailed descrip-




Figure 5. Single-hop admission control decision example.

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

(a) CACP.

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

(b) PAC.

Table 1. Single-hop example available bandwidth.

tion of CACP’s multihop routing protocol please refer to [15].

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 correlated with the network load. Second, CACP does not have any mobility support. To achieve acceptable performance it 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. These problems are further discussed in Section 4.3.

To address the shortcomings of previous solutions, we propose a simple perceptive admission control protocol, described in the following section.

3. Perceptive Admission Control

To perform admission control in wireless mobile networks, we propose a 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 propose changing the range of the available bandwidth measurement so that each node can make admission control decisions without communicating with any other nodes. In the following sections, we describe our admission control protocol, as well as mechanisms to handle mobility.

3.1. Available Bandwidth Measurement Range and Admission Control Decisions

In Section 2.2 we showed that the channel busy time calculation is a good measure of the network utilization. For PAC, we change the sensing range so that transmissions are sensed at a distance large enough to allow 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 carrier 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 the packet receptions. 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 take the newly admitted traffic into consideration for future admission control decisions. Likewise, when a flow stops, the increase in available bandwidth is quickly incorporated into the network utilization measurement so that other flows can be admitted.

To further describe the operation of PAC an example is provided. In Figure 5(b), 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 1(b) 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 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 1(b), 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 there is not enough bandwidth available. Hence node W does not admit the traffic flow. In contrast to previous work, 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 [1]. 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.




Figure 6. Example of admission control in a mobile network that requires sources to throttle or reject traffic.

To admit a new flow, the required bandwidth (B_{req}) for the new flow must meet the following condition:

$$B_{avail} - B_{rsv} > B_{req} \quad (3)$$

This 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 this paper it is fixed.

3.2. Mobility

When a node, and consequently 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 throttle or reject them as network conditions change.

The following example illustrates the importance of this property. In Figure 6(a), suppose two flows, each consuming 75% of the maximum bandwidth, are admitted at nodes far enough apart that each participating node pair is outside CSR of the other. Later, as shown in Figure 6(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. Using PAC, the sources detect the ensuing network congestion and throttle or reject the offending traffic flows when another sender enters the PAC measurement range. If both flows are allowed to continue at their present transmission rate, neither flow will receive its needed quality of service.

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 uncongested state. For this study, we assume all flows require a minimum level of service such that the flow cannot be throttled. Therefore, we reject flows to avoid congestion.

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 only checks 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 adjusts.

3.3. Multihop Routing

The PAC admission decision can be utilized to create multihop routes during route discovery using a method such as CACP's multihop routing protocol. However, instead of CACP's admission control decision, PAC's available bandwidth measurement and admission control decision process should be used. For more details on CACP's multihop routing protocol please see [15].

In addition to a multihop routing protocol that performs admission control, 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. This should be performed continuously, periodically or on-demand. Since multihop routing is simply an application of PAC's admission control decision to a multihop routing protocol, it is not discussed further in this paper.

4. Performance Evaluation

In this section, we demonstrate that PAC effectively controls traffic admission to avoid congestion and maintain quality of service. Furthermore, PAC allows high network utilization and spatial reuse without degrading QoS. First we describe simulation results that show PAC performs admission control efficiently and effectively. We then describe how PAC has addressed a few problems with CACP.

4.1. Simulation Environment

To evaluate PAC we use the NS-2 simulator. Our simulation parameters are listed in Table 2. In our simulations, a packet is considered receivable if its reception power is above a threshold value, called the reception power threshold. Likewise, if a packet is received and the power is above the carrier sensing power threshold, the channel is sensed busy during this packet transmission. Given a threshold value, transmission power and propagation model, a specific maximum distance for packet reception or detection can be determined [14]. For our simulations, the propagation model is two ray ground and no obstacles are considered. This results in a reception range of 250m and a carrier sensing range of 550m.

The reception power threshold, propagation model and capture factor must be known to determine the receiver interference distance (RID). The capture factor defines the minimum power ratio between the received power of two

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	20s
Carrier Sensing Range	550m	Number of nodes	50
Capture Factor	10.0	Simulation Time	200 seconds
Receiver Interference Distance	440m	Number of Runs	10

Table 2. Simulation Parameters

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

Table 3. PAC Parameters

Admission Control Protocol	Packet Losses	Packets Delivered	Average Delay (s)
None	26778	81825	0.973
PAC	0	58173	0.005
CACP	0	51182	0.004

Table 4. Performance

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 use 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 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 are separated by RxR , another sender must be at least 440m away for its transmission to be able to take place simultaneously. Therefore, for our simulations RID is 440m; at this distance the received power of another sender is guaranteed to be less than $RX_{Thresh}/10.0$.

With a reception range of 250m and a RID of 440m, the range for PAC is 940m, as calculated by Equation 1. Given the propagation model and other simulation parameters we calculated the minimum reception power threshold at this distance [14]. In our simulations, if a packet is received with a power above this threshold value, the packet is considered in the available bandwidth calculation. An adjustable sensing range will likely become a common feature in new hardware to improve performance [18]. The carrier sensing mechanism for the MAC layer is filtered so that it behaves as if the minimum reception threshold was not changed. If the carrier sensing mechanism was changed, the collision avoidance attributes, spatial reuse [4, 18] and medium access are affected. For other more challenging propagation models (i.e. shadowing) a larger measurement range may be used to ensure proper operation.


Table 3 lists the values used by PAC in our simulations. To perform the available bandwidth calculation, a maximum effective bandwidth (B_{max}) of 1200 kbps is as-

sumed¹. We determined this value experimentally in Section 2.2 and it is close to analytical value derived in [1]. We reserve 20% (240 kbps) of the maximum bandwidth to avoid congestion, allow for temporary fluctuations and detect mobility before congestion. The same reserved bandwidth is used for CACP in the simulations. If the detected available bandwidth drops below 120 kbps (10% of the maximum bandwidth), we assume over-utilization is imminent. We utilize a sliding window to calculate the PAC-based available bandwidth. The size of the window we utilize is 250ms. 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. The backoff time between flow admission attempts after flow rejection is between 1 and 2 seconds. The time interval between congestion detection checks is also between 1 and 2 seconds. The simulation results in Section 4.2 show these values are adequate, since in our simulations no two flows were rejected in response to the same congestion event. Tuning or dynamically adjusting these parameters will further increase PAC's performance and is a subject of further work.

4.2. Local Admission Control Performance

In this section we show that PAC results in a high quality of service for all admitted flows, whereas lack of admission control leads to high packet loss and delay. We also compare the performance of PAC to that of CACP. We study 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 is needed; the sender-receiver pairs move together. There


¹ An accurate prediction of the maximum achievable throughput in ad hoc networks is very difficult. Since nodes may not all be within reception or carrier sensing range of each other this further complicates analysis.



(a) Without admission control.

(b) With PAC.

Figure 7. Throughput of a single representative receiver in one simulation.




(a) Without admission control.

(b) With PAC.

Figure 8. Delay of a single representative receiver in one simulation.

are 25 sender-receiver pairs and every five seconds another sender starts sending CBR traffic. Therefore, after 125 seconds of simulation time, all senders are active.

A summary of the results is presented in Table 4. It is evident from the results that lack of an admission control protocol results in significant packet loss and delay. Figure 7(a) shows the 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 progresses and more sources become active, the channel becomes congested. After 80 seconds have elapsed, the throughput for this receiver decreases significantly. At 180 seconds the node gains unfair advantage in channel access and again experiences acceptable throughput. This temporary unfairness is a well known behavior in IEEE 802.11 [10]. This results in a spike in throughput as queued packets are delivered. In addition to experiencing degraded throughput for most of the simulation, the delay experienced by received packets is often unacceptable for real-time applications. Figure 8(a) presents the delay for the



(a) Without admission control.

(b) With PAC.

Figure 9. Throughput for all flows.

received packets without admission control. Once the channel becomes congested, the delay value increases sharply. This is particularly high since all packets traverse only a single hop from the source to destination.

In contrast to the poor performance without admission control, PAC enables admitted sessions to experience much better service. Figures 7(b) and 8(b) show the number of packets received per second and the delay for the same receiver as in Figures 7(a) and 8(a). The figures show that traffic throughput for this session is nearly constant. In addition, the delay is extremely small. Note that the difference in the scale of the y-axis between Figures 8(a) and 8(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 demonstrated by 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 9 shows the packet receptions per second for all 25 flows with and without PAC; each vertical line represents the start of a new flow. In Figure 9(a), we see that without admission control each flow experiences notably different throughput. In contrast, with PAC each flow experiences nearly the same throughput, as shown in Figure 9(b). This is possible because PAC limits the number of admitted flows.

In terms of delay and throughput for admitted flows, CACP performs similarly to PAC, as shown in Table 4. One difference is the number of packets delivered. Since CACP has messaging overhead for every admission decision attempt, this consumes a part of the bandwidth that would otherwise be available for data packet delivery. In the random network topologies simulated, the conditions, discussed in Section 4.3, where CACP performs improperly or overly conservatively were not present. Hence CACP performed well in these scenarios.

To summarize the results of these simulations, through admission control PAC is able to minimize packet loss and

delay. Further, the bandwidth is fairly shared between all admitted flows. Without PAC, the channel is susceptible to congestion, resulting in large packet loss and delay.

4.3. Qualitative Comparison

Although CACP performs well in some cases, the protocol has many weaknesses. First, CACP is likely to make erroneous admission decisions when certain common network topologies exist. Second, CACP does not include support for mobility. Adding mobility support would be difficult, since it would require periodic messaging along the whole path. Third, CACP is conservative in area in its admission decisions, hindering legitimate spatial reuse. In PAC each of these problems is mitigated. In [2], we present the general scenarios where the performance of CACP degrades and describe how these scenarios are addressed in PAC.

5. Conclusions

In this paper we present PAC, a perceptive admission control protocol for use in wireless mobile networks. PAC addresses two issues: shared wireless bandwidth and node mobility. PAC is able to compute its available bandwidth and determine whether a flow can be admitted by sensing all transmissions that may interfere. Also, since calculating the available bandwidth is a simple, passive technique, each source can quickly adapt its admitted traffic flows to changing wireless channel use. Simulation results illustrate that PAC effectively limits the amount of data traffic to avoid congestion. This results in consistent throughput, low packet loss and delay for all admitted flows. PAC is useful in wireless networks with applications that require high quality of service, such as multimedia applications.

In addition to admission control, we feel that PAC is applicable to a number of other load-aware network applications. We expect that insight into the spatial location of nodes can be gained through consideration of not only the amount of time the channel is sensed as busy, but also the length and received power level of each transmission. Also, we plan to explore multiple priority MAC layers, i.e. IEEE 802.11e, and extend PAC to determine the relative utilization of each priority. By using multiple priorities un-admitted flows may share any unreserved capacity and avoid starvation. Furthermore, we plan to implement PAC in a real system to prove its feasibility.

Acknowledgments

This work is supported in part by a gift from Intel Corporation.

References

- [1] G. Bianchi. Performance Analysis of the IEEE 802.11 Distributed Coordination Function. *IEEE Journal Selected Areas in Communications*, 18(3), March 2000.
- [2] I. D. Chakeres and E. M. Belding-Royer. PAC: Perceptive Admission Control for Mobile Wireless Networks. Technical Report UCSB-CS-2004-17, University of California, Santa Barbara, May 2004.
- [3] T.-W. Chen, J. T. Tsai, and M. Gerla. QoS Routing Performance in Multihop Multimedia Wireless Networks. In *Proceedings of IEEE International Conference on Universal Personal Communications (ICUPC)*, San Diego, October 1997.
- [4] X. Guo, S. Roy, and S. Conner. Spatial Reuse in Wireless Ad hoc Networks. In *Proceedings of the IEEE Vehicular Technology Conference (VTC)*, Orlando, FL, October 2003.
- [5] Y.-C. Hsu and T.-C. Tsai. Bandwidth Routing in Multihop PacketRadio Environment. In *Proceedings of the 3rd International Mobile Computing Workshop*, March 1997.
- [6] S.-B. Lee, G.-S. Ahn, X. Zhang, and A. Campbell. IN-SIGNIA: 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, 2000.
- [7] C. R. Lin and J.-S. Liu. QoS Routing in Ad hoc Wireless Networks. *IEEE Journal on Selected Areas in Communications*, 17(8), November 1999.
- [8] H. Luo, S. Lu, V. Bharghavan, J. Cheng, and 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), June 2002.
- [9] S. Murthy and J. J. Garcia-Luna-Aceves. A Routing Architecture for Mobile Integrated Services Networks. *ACM Journal of Mobile Networks and Applications*, 3(4), 1999.
- [10] T. Nandagopal, T. Kim, X. Gao, and V. Bharghavan. Achieving MAC Layer Fairness in Wireless Packet Networks. In *Proceedings of the Sixth Annual International Conference on Mobile Computing and Networking (MobiCOM'00)*, Boston, MA, August 2000.
- [11] R. S. Prasad, M. Murray, C. Dovrolis, and K. Claffy. Bandwidth Estimation: Metrics, Measurement Techniques and Tools. *IEEE Network*, 17(6), November 2003.
- [12] R. Ramanathan and M. Steenstrup. Hierarchically-organized, Multihop Mobile Wireless Networks for Quality-of-service Support. *ACM Journal of Mobile Networks and Applications*, 3(1), 1998.
- [13] R. Sivakumar, P. Sinha, and V. Bharghavan. CEDAR: a Core-Extraction Distributed Ad hoc Routing Algorithm. *IEEE Journal on Selected Areas in Communications*, 17(8), August 1999.
- [14] W. Ye. Calculating the receiving threshold, RXThresh_ for Phy/Wireless, 2000. threshold.cc, NS-2.27.
- [15] Y. Yang and R. Kravets. Contention-Aware Admission Control for Ad Hoc Networks. Technical Report 2003-2337, University of Illinois at Urbana-Champaign, April 2003.
- [16] L. Zhang, S. Deering, and D. Estrin. RSVP: A New Resource ReSerVation Protocol. *IEEE Networks*, 7(5), September 1993.
- [17] C. Zhu and M. S. Corson. QoS Routing for Mobile Ad hoc Networks. In *Proceedings of the IEEE Conference on Computer Communications (INFOCOM)*, New York, New York, June 2002.
- [18] J. Zhu, X. Guo, L. Yang, and 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, June 2004.