A dynamic packet reservation multiple access scheme for wireless ATM

Deborah A. Dyson [a] and Zygmunt J. Haas [b]

[a] *3Com Corporation, Boxboro, MA 01719, USA*
[b] *School of Electrical Engineering, Cornell University, Ithaca, NY 14853, USA*

The dynamic packet reservation multiple access (DPRMA) scheme, a medium access control protocol for wireless multimedia applications, is proposed and investigated. DPRMA allows the integration of multiple traffic types through a single access control mechanism that permits users to specify their immediate bandwidth requirements. The primary feature of DPRMA is the dynamic matching of the traffic source generation rates with the assigned portion of the channel capacity. This is accomplished by a control algorithm that regulates the actual amount of channel capacity assigned to users. To support multimedia communication, channel capacity assignments are prioritized by traffic type. The performance of the scheme is evaluated and the scheme is shown to perform well in a system with voice, video conferencing, and data users present. It is also shown to provide improved performance over a system with a modified version of the packet reservation multiple access (PRMA) scheme. Furthermore, several system parameters are studied and optimized.

1. Introduction

Wireless networks in today's society are receiving an ever-increasing amount of attention. In particular, wireless multimedia systems are of significant interest and are expected to be in demand soon by the general public. The current cellular systems, however, are incapable of efficiently and effectively supporting the introduction of multiple traffic types. The cellular systems were designed specifically for voice users but are already being used for other applications as well. In many cases the systems' resources have been saturated by this additional traffic load. In order to make the systems more suitable for multimedia users, several steps must be taken. In particular, the method for allocating systems' resources must be reevaluated, and a new Medium Access Control (MAC) protocol must be implemented.

The most popular multimedia computer networking protocol in use today is the Asynchronous Transfer Mode (ATM) [1]. This protocol uses statistical multiplexing as a method of dividing channel bandwidth up among all users. As ATM becomes a more and more widely accepted protocol, it is reasonable to assume that future wireless networks will be expected to easily interface with ATM networks. Therefore this work focuses on the transmission of ATM cells.

In this study we concentrate on Time Division Multiple Access (TDMA) schemes. The current wireless digital MAC standards, such as IS-54/136, are well-suited for transmission of voice traffic. A cellular frequency channel is divided into multiple slots and these slots are grouped into frames. Each user is allocated a set number of slots in every frame. This protocol provides *voice users* with the bandwidth they require and guarantees minimal delay in transmission of packets. However, the addition of different types of traffic into the system would cause significant problems, since these users may require more or less bandwidth than the fixed allocation allows for. For high-bandwidth real-time traffic sources this could cause significant delay in packet transmission. An additional consideration is that time division multiple access techniques can potentially provide high throughput for Constant Bit Rate (CBR) traffic, but not necessarily for Variable Bit Rate (VBR) traffic. Therefore, this technique is not well suited for a multimedia system.

Another protocol that has been proposed to support packetized voice communication is called Packet Reservation Multiple Access (PRMA) [2]. This protocol assumes that voice activity detectors are implemented in the system, and consequently the system can differentiate between periods of speech and periods of silence in a conversation. Users are allowed to reserve slots whenever they are in a talkspurt. When the talkspurt ends, the reservation is released. Data users can also be accommodated by this protocol, but are not permitted to make slot reservations. PRMA allows a greater number of voice users to be admitted into the system than does a strict TDMA protocol. The main drawback however is that the protocol is specifically designed for a predominantly voice system. Aside from low-bit rate data users, most other traffic types cannot be easily accommodated. Therefore, another more suitable protocol is required.

In this work we introduce the Dynamic Packet Reservation Multiple Access (DPRMA) protocol. It is inspired by the PRMA protocol, but is better suited for a multimedia traffic system. DPRMA is designed to flexibly assign bandwidth to users and update these assignments as necessary to support the changing needs of VBR users. Priority is given to the real-time traffic users, but an effort is made to accommodate non real-time traffic whenever possible.

Section 2 presents a general background on some proposals for wireless MAC protocols which have received recent attention. Section 3 provides an explanation of the PRMA protocol. This is followed by a description of Dynamic Packet Reservation Multiple Access in section 4. The similarities between the two protocols are highlighted, and the improvements that DPRMA offers over PRMA are emphasized. In section 5, the types of traffic that are included in the protocol simulations are outlined. Quality of Service (QOS) requirements are specified for each user. The simulation results are presented in section 6, where the performance of DPRMA is compared with that of a modified PRMA protocol.

2. MAC protocols for wireless networks

Beyond the basic TDMA technology, there have been several MAC protocols that have been proposed for use with wireless services [3–7]. Those that are most suited for multimedia applications are reservation-based. They provide the users with some mechanism to indicate the portion of the bandwidth that they need allocated for transmission purposes. One such protocol is the Multiservices Dynamic Reservation (MDR) TDMA format [8], an example of which appears in figure 1(a). This protocol divides the bandwidth into frames which are further divided into several regions. The first region is divided into small request slots that the mobiles use to contend for access to the system via a slotted ALOHA protocol. Based on the information obtained during this period, the remainder of the frame is divided among the CBR, VBR and data users. CBR traffic is allocated first and may be assigned a limited number of time slots. The remaining slots are dynamically allocated among the VBR and data users. Once a CBR user gains access to the system it is allocated a slot in every frame for as long as it continues to transmit. VBR and data users can request multiple slots in a frame and the protocol will attempt to comply if the bandwidth is available. If it is not, the portion of the request that is not honored will be accommodated in subsequent frames.

This protocol has several significant drawbacks. First it assumes that CBR (typically voice) traffic has the highest priority in the system and makes no attempt to provide guarantees for QOS of the other users. The DPRMA protocol will be shown to provide guarantees for VBR traffic. Another disadvantage of the MDR protocol is the delay that may be incurred by each user while it waits for the beginning of the request period. This delay may be unacceptable for real-time traffic. DPRMA eliminates this delay by allowing users to contend for reservations whenever there is an empty slot available. Another problem with this protocol is that the use of minislots makes it difficult to implement in a large system. In reality, each minislot will need its own guard band in order to account for the different propagation delays that each user will experience. Thus, the use of minislots assumes a small service area for each base station.

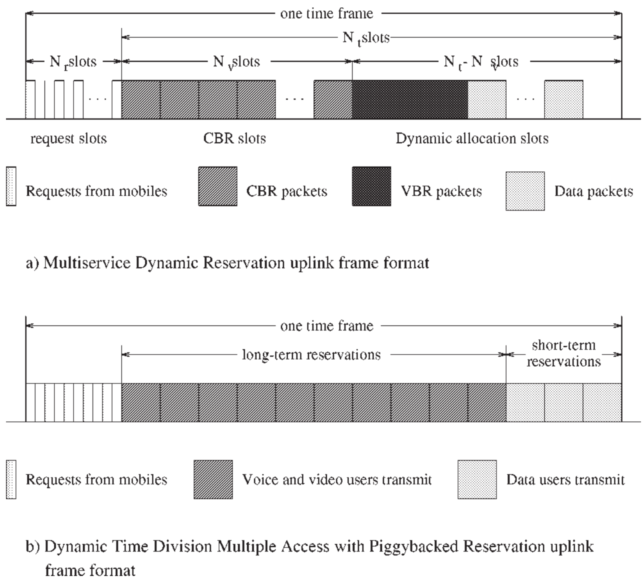

Figure 1. Format for MDR and DTDMA/PR protocols.

Another reservation-based MAC protocol is Dynamic TDMA with Piggybacked Reservation (DTDMA/PR) [9]. In this protocol, uplink frames are divided into three fields as shown in figure 1(b). The first field contains reservation minislots. The second field consists of multiple long-term reservation slots for time-dependent CBR and VBR traffic. The last field contains slots for short-term reservations, specifically for data users making reservations on a frame-by-frame basis. The boundary between the last two fields is movable. Each user initially obtains a reservation by randomly selecting to transmit in one of the reservation minislots. At the end of the reservation field the base station provides feedback about the assignment of slots in the current frame. CBR users are assigned a fixed number of slots in the long-term reservation field. VBR users are also assigned slots in this field but can update their reservation needs by using a piggybacked reservation when they transmit their packets.

This protocol suffers the same problem with minislots as did the previous protocol. In addition, the base station only updates slot reservations once a frame, immediately after the reservation field. Thus, the piggybacked requests of the video users cannot be accommodated immediately. This is a major drawback of the protocol. The DPRMA protocol, however, allows immediate updating of slot assignments in order to provide CBR traffic users with the best possible QOS requirements. In addition, the DPRMA protocol allows data users to maintain reservations instead of forcing them to contend on a frame-by-frame basis. This significantly reduces the contention in the system.

3. Packet reservation multiple access

One of the widely-accepted wireless MAC protocols to date has been the Packet Reservation Multiple Access protocol. DPRMA is, in fact, inspired by this protocol. Like the other MAC protocols, PRMA allows multiple users to share the resources of one frequency channel. The channel is divided into time slots which are grouped into frames. The size of the slots and frames are designed specifically to efficiently accommodate voice users. A one-slot-per-frame reservation guarantees that a voice user's QOS requirements are met. A distinction is made between two different traffic types in the system: random traffic (typically voice) and periodic traffic (typically data). All users contend for access to the channel using a modified Reservation ALOHA protocol, R-ALOHA [10]. A new user waits for an unused slot and then transmits with probability p_p for periodic users and p_r for random users. If a collision occurs, the user waits for the next available slot and tries again with the same transmission probability. This continues until the user transmits successfully or decides to give up. If a periodic user is successful in transmitting, that slot becomes reserved for the user. A data user, however, is not allowed to obtain a reservation and instead must contend for a slot each time it has a new packet to transmit.

The values that are chosen for p_p and p_r in a system do not have to be identical. By making them different, the system can give higher priority to one traffic type over the other. Time-dependent traffic requires timely delivery to ensure that its QOS requirements are met. Data traffic can tolerate a much greater delay without any impact on its Quality of Service. Therefore, the transmission probability for periodic (voice) users is generally set higher than that for random (data) users.

Information about the availability of slots is provided to the users by the base station on the downlink channel. In addition, the base station indicates the success or failure of a new transmission attempt. Each periodic user transmits in its reserved slot until it no longer has any more packets to send. The user then leaves its next reserved slot empty, informing the base station that the user is releasing its reservation. In the future, when the periodic user has more packets to send, it will have to contend again for a reservation.

There have been many adaptations to PRMA that have been proposed in order to better accommodate VBR traffic into the system [11–14]. The DPRMA protocol that we are presenting is one such proposal.

4. Dynamic packet reservation multiple access

4.1. Medium access control

The DPRMA protocol resembles PRMA in several ways. In DPRMA, the frequency channels continue to be divided into slots and frames. The size and spacing of these remain such that a voice user needs to reserve exactly one slot in every frame. Also like PRMA, the users of the system initially contend for access using the R-ALOHA-like method. A user who is successful in acquiring a slot and who requests a single slot per frame reservation is allocated the same slot in which it successfully contended. This assignment continues for all subsequent frames until the user changes or releases its reservation. The reservation release mechanism is the same in both protocols; users simply leave the last reserved slot empty. Thus, with only voice users present, the system reverts back to PRMA.

The primary difference between the two protocols is the manner in which reservations are made and resources are allocated for VBR sources. The slots within a DPRMA frame are divided among the users based on the amount of bandwidth that each user requires. Users may reserve a number of slots within a frame or even slots in alternating frames, as long as that capacity is currently available. In addition, changes to a user's allocation request can be dynamically accommodated. All users in the system are permitted to request reservations, including data users. The manner in which the base station attempts to accommodate the data users is slightly different than the way it accommodates time-dependent users. Specific details about this issue will be provided later in this section.

The base station has the responsibility of dividing the bandwidth up among the active users. In order to accomplish this task in DPRMA, each mobile conveys its requirements to the base station via several Reservation Request (RR) bits that are included in the header of each uplink time slot. It is the user's responsibility to determine the appropriate rate required and set its rate bits accordingly.

Each user can transmit at a limited number of different transmission rates, c_i. This limits the amount of overhead introduced by the presence of the RR bits. These rates are defined as

$$c_i = 2^i \times C/n, \quad i_{\text{min}} \leqslant i \leqslant \log_2 n, \tag{1}$$

where C is the data rate of the channel in bits per second, n is the number of slots in a frame, and i is an integer. The value for i_{min} dictates the smallest possible bandwidth allocation and can be set to any value that is appropriate for the system in question. For this study it will be shown that the smallest bandwidth allocation will be 35.333 kbps, and therefore i_{min} is set to 0.

When a user has a burst of information to transmit it must first attempt to obtain a reservation. It sets the appropriate RR bits to indicate its rate request, contends for an empty slot, and monitors the downlink channel to determine its success or failure status. Success or failure is indicated by the base station via several Reservation Acknowledge (RA) bits in the header of the downlink messages. When a successful transmission has occurred, the base station immediately attempts to accommodate as much of the requested rate as is possible. If the total request cannot be accommodated, then a partial allocation is made and all empty slots are assigned to the new user. The base station keeps a record of any partial allocations, so that the remaining request can be accommodated whenever the capacity becomes available.

For time dependent traffic, as many of the available empty slots as are necessary can be used in order to accommodate a request. If there are not enough empty slots, then these users are permitted to preempt the data users from their reservations. The base station, as the resource allocator, determines how many and which data slots will be reassigned to the time dependent users. As many data users as are necessary will be preempted in order to obtain the full allocation. Data requests that have been preempted in this manner must be placed in a queue to await further service when the bandwidth becomes available.

If a full allocation of a new user's request is possible, the base station must determine which of the remaining unclaimed slots will be assigned. The base station first identifies which slots are currently unallocated and determines how many such slots exist. Next, the base station examines each of these slots in sequential order to determine if the slot will be assigned to accommodate the new request. Throughout the process, the base station maintains a record of how many slots, S_n, the user still needs in order to have its request satisfied. Every time a slot is successfully assigned, S_n is decremented. In addition, the base station keeps track of the number of available slots, S_c, that have not yet been considered for assignment. Each time a new slot is considered, S_c is decremented. As the base station considers each available slot, it assigns the slot with probability P_a, where

$$P_a = S_n/S_c. \tag{2}$$

Thus, the probability that a slot is assigned is dependent upon how many slots are still needed to satisfy a user's request. This process ensures that the user will be assigned exactly the number of slots it requires. Also, it tends to spread the allocation of slots randomly throughout the frame. An example of the operation of the slot assignment algorithm is shown is figure 2(a).

Once a user has secured a reservation, it must monitor the downlink channel to determine in which slots it is allowed to transmit. This is indicated via Slot Reservation (SR) bits that are incorporated into the downlink message header. Any changes to a user's reservation requirements are communicated by the user to the base station via the RR bits. An increase in reservation is accommodated if the resources are available. Once again, real-time traffic requests will take precedence over non real-time traffic reservations. User requests may be added and data reservations removed such that the real-time requests are accommodated whenever possible. The additional slots needed are assigned according to equation (2). A decrease in reservation is always accommodated immediately to prevent the user from running out of packets to send and, consequently, losing its reservation.

When a rate decrease is requested by a user, the base station first determines which slots are currently assigned to that user. The base station then considers each of these slots one at a time in sequential order, for deallocation purposes. The number of slots yet to be released, S_d, and the number of slots yet to be considered for release, S_r, are constantly

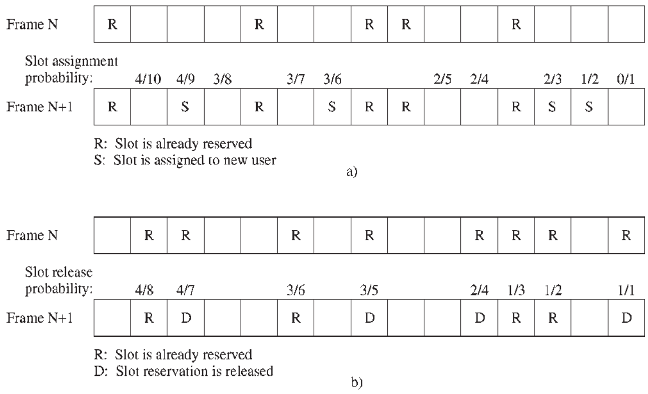

Figure 2. a) New user arrives and requests 4 slots. New slots are assigned according to $P_a = S_n/S_c$. b) A user with 8 slots wants to decrease its reservation to 4. Assigned slots are released according to $P_d = S_d/S_r$.

updated throughout this process. Each slot is released with probability

$$P_d = S_d/S_r. \tag{3}$$

As with the slot assignment algorithm, this process ensures that the necessary number of slots are always released and continues to maintain a random spread of slot assignments throughout the frame. An example of how the slot deallocation protocol operates is shown in figure 2(b). The base station uses this same deallocation process whenever data users must be preempted to accommodate real-time traffic users.

After the slots have been deallocated, the base station determines if there are any users waiting for additional slots to be assigned. Real-time traffic users are considered first. The backlog of such users is handled in a first come, first served manner. If there are any resources remaining after the real-time traffic users are satisfied, then data users may be assigned. Data user assignments are made using the first come, first served method. Any additional backlog remains queued and is handled whenever additional slot reservations are released. It should be clear thus far that DPRMA places a higher priority on the timely delivery of real-time traffic than on the delivery of data traffic. For real-time users, if a packet is not transmitted within the guaranteed specification, it is dropped. In this work, excessive transmission delay is the only cause of packet dropping for real-time traffic. Data traffic users suffer packet loss only when their buffers overflow. All packets that are transmitted are assumed to be received by the base station without transmission error.

4.2. Rate requests

In the DPRMA protocol, it is up to each user to attempt to reserve the appropriate rate that ensures timely delivery of traffic. For CBR traffic this is simply a matter of reserving the rate that is as close to the generation rate as possible. At times it may be impossible to exactly match the generation rate to the reservation rate due to bandwidth quantization. Therefore, every so often, the user may have

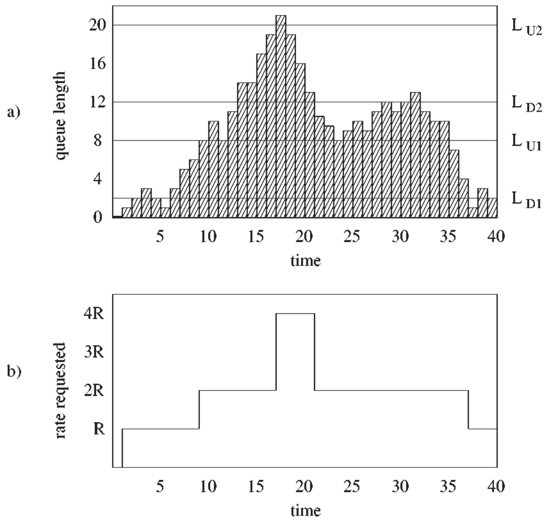

Figure 3. a) User queue length in relation to threshold levels. b) Resulting rate request.

to increase or to decrease the reservation rate to avoid running out of packets or delaying transmission.

The rate selection method proposed here allows newly generated packets to be queued in a buffer as they await transmission. As the size of the queue grows, the user increases its reservation request to avoid excessive transmission delay or buffer overflow. If the queue length subsequently decreases, the user then requests a lower reservation rate to avoid running out of packets. The buffer size that corresponds to an increase or decrease is defined as a threshold. The thresholds are set to produce a system with hysteresis.

This threshold level protocol can also be adopted by users with VBR traffic. These traffic types, however, can have a large range of possible packet generation rates. Therefore, a system with multiple threshold levels may be required. An example of this rate request method with multiple thresholds can be seen in figure 3.

For this example the user in question is allowed to specify one of three possible rates, R, 2R, and 4R. Initially the user sets its request at R. When the user's queue length first crosses the L_{U1} threshold, the user doubles its rate request to 2R. It can maintain this rate request until the queue length crosses the L_{U2} threshold or the L_{D1} threshold. Crossing L_{U2} causes the user to double its rate request to 4R, and crossing L_{D1} causes the user to halve its rate request to R.

A more detailed discussion of the selection of appropriate threshold levels appears in section 6.5.

5. Traffic types

5.1. Voice traffic

The voice traffic model used here is based on the work done by Brady [15] and assesses that speech sources generate periods of talkspurts and gaps. By assuming that a voice activity detector can be used to differentiate between

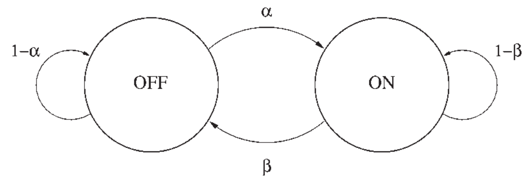

Figure 4. Two-state Markov model for voice.

principle talkspurts and principle gaps, voice traffic can be characterized by the two-state Markov chain model displayed in figure 4. The system alternates between the ON and OFF states, which correspond to the talkspurts and idle periods of speech. In the ON state, voice packets are generated at a constant rate. No packets are generated in the OFF state due to silence detection. Time spent in each state is exponentially distributed with means α^{-1} for the OFF state and β^{-1} for the ON state. A voice source would therefore require a reservation while in the ON state, but then could release the reservation during the OFF state when no packets are available for transmission. The effects of this traffic model are addressed by Goodman in [2] in his study of PRMA.

The ON-OFF voice traffic sources used in this study are modeled with the parameter values $\alpha^{-1} = 1.35$ s and $\beta^{-1} = 1.0$ s [16]. Since voice cells must be delivered in real time, there is a maximum transmission delay allowed; any voice cell that has not been transmitted within 40 ms of its time of generation will be discarded at the source. For this study, the number of cells that are lost in this manner must not exceed 1%.

5.2. Video conferencing traffic

The model used to describe video conferencing traffic is based upon work done by Heyman et al. [17]. In this study of actual video conferencing traffic, video frames (VF) were found to be generated periodically and to contain a varying number of cells in each video frame. The number of cells per video frame was determined to be approximately characterized by the negative binomial distribution. A Markov chain model can be constructed that demonstrates the transition from one state to the next. The transition matrix is calculated using the following equation:

$$P = \rho_a I + (1 - \rho_a)Q, \tag{4}$$

where I is the identity matrix, ρ_a is the autocorrelation coefficient and each row of the Q matrix is composed of the probabilities $(f_0, \ldots, f_K, F_K)$. The quantity f_k has the negative binomial distribution and represents the probability that k cells are present in one video frame. The value of K in the Q matrix represents the peak cell rate and $F_K = \sum_{k>K} f_k$. The Markov chain for this model is displayed in figure 5.

The statistics for video conferencing traffic that were obtained in [17] resulted from coding a video sequence with a modified version of the H.261 coding standard. The results showed a peak cell generation rate of 220 cells/VF

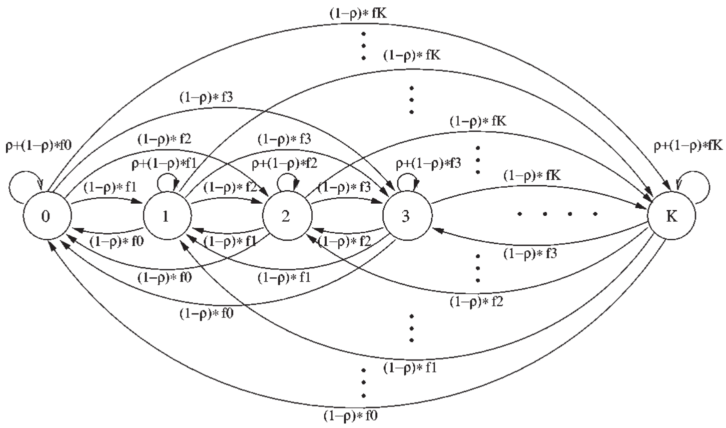

Figure 5. Markov model for video.

and an average cell generation rate of 104.8 cells/VF. A cell size was 48 bytes, which is equivalent to the payload in an ATM cell. New VF's arrived every 40 ms with the cells being generated at the peak rate of $220/0.04 = 5500$ cells/s.

Video traffic, like voice, must be delivered within a certain time period. The maximum transmission delay for this system is set to be 40 ms with packets being dropped when this deadline is reached. The allowed packet dropping probability is set to 0.01%.

5.3. Data traffic

The data users in the system are modeled to generate packets according to a Poisson process. The average generation rate, λ, is 94.34 packets per second (which corresponds to the rate of 40 kbps). Since data traffic is generally non real-time traffic, no packets are discarded due to excessive delay. However, a buffer of 10 kBytes is allocated to each data user. Data packets are lost when this buffer overflows. Data traffic cannot tolerate lost packets, and therefore, essentially no data packets should be lost in this manner.

6. Simulations and results

6.1. The transmission channel

In order to accommodate the traffic types described above, we selected a channel with a capacity, C, of 9.045 Mbps. Propagation delays on this channel are neglected in this study. The channel is divided into ATM-sized slots of 53 bytes each. The 5 bytes of header are modified to include all of the overhead associated with the MAC protocols. This is a reasonable assumption since the Virtual Circuit Identifier and Virtual Path Identifier fields in an ATM header are considerably longer than necessary to accommodate the number of active users that would be present in a single wireless cell. Thus we assume no additional overhead is necessary to implement the DPRMA scheme. The frame size was selected so that a reservation of one slot in a frame provides 35.333 kbps which accommodates 32 kbps of payload. This is exactly the bit rate for voice traffic. Thus, a voice user requires a reservation of one slot per frame. In order to accomplish this, the channel is divided into frames consisting of 256 slots. The frame length, T, is 12 ms. The system parameters are summarized in table 1.

Table 1
Values of the basic system parameters.

Channel capacity (Mbps)	C	9.045
Frame length (ms)	T	12
Slots per frame	N	256
Slot size (bytes)	D	53
Average time spent in ON state for voice (s)	β^{-1}	1.0
Average time spent in OFF state for voice (s)	α^{-1}	1.35
Video frame length (ms)	F	40
Peak cell generation rate for video (cells/VF)		220
Average arrival rate for data (cells/s)	λ	94.34
Max allowed transmission delay for voice (ms)	D_v	40
Max allowed transmission delay for video (ms)	D_{vc}	40
Max transmission buffer size for data (kBytes)	B	10
Max allowed cell dropping probability for voice	P_{dropv}	10^{-2}
Max allowed cell dropping probability for video	P_{dropvc}	10^{-4}
Max allowed cell dropping probability for data	P_{dropd}	0.0

6.2. PRMA*

Since the protocol we have based our design on is Packet Reservation Multiple Access, it is desirable to compare the performance of our protocol to that of the pure PRMA. Since PRMA cannot accommodate the multimedia traffic that we will be simulating, we have made some simple modifications to the protocol to make it more suitable for use in a multimedia network. We will refer to this modified PRMA protocol as PRMA*.

PRMA* is adapted to accept VBR users by allowing a user to reserve multiple slots in each frame. The user must monitor both the number of slots it needs and the number of slots it currently has reserved. The users must contend for each slot individually, and there is no communication between the mobile and the base station about the user's reservation requirements. Each time the user successfully contends for and gains a slot reservation, it is allocated the same slot in subsequent frames. In addition, the user continues to maintain a reservation in all the slots it had reserved in previous frames. A user decreases the number of slots it has reserved, by leaving the appropriate number of slots empty. An empty slot does not indicate that the user's entire reservation is being released, but only that a single slot is being given up. When the user needs to release all of its reservation, all slots must be left empty by that user.

This modified PRMA protocol requires that the mobiles determine their own reservation requirements. Therefore, it is assumed that the PRMA* users assess their rate requirements using the same threshold level method that was presented in section 4.2. The PRMA* users will be allowed to attempt to obtain reservations for the bit rate intervals specified in equation (1). During rate increases, the user will contend until it has successfully reserved the appropriate number of slots. When a decrease is required, the user will cease transmission until the appropriate number of its reserved slots have been released.

The PRMA* protocol differs from DPRMA in several important ways. In PRMA*, the base station has no control over the reservation of slots. Therefore, no priority can be given to different traffic types. In addition, the PRMA* system does not allow data users to make reservations and all other users must contend for multiple slot reservations. All DPRMA users, however, are able to adjust reservations by submitting rate requests to the base station in a contention-free manner.

In this paper, we evaluate the performance of the DPRMA protocol and present a comparison with the PRMA* protocol. We emphasize that the comparison with the PRMA*, by itself, is not of primary importance. Rather, such a comparison allows us to understand the effects of some DPRMA-specific features on the performance of MAC schemes. Thus, the PRMA* simulation serves as a "control" scheme to gain insight into the effects of the DPRMA features.

6.3. Performance comparison

6.3.1. Voice/video system

In both the PRMA* and the DPRMA systems the video conferencing users control their rate requests using multiple threshold levels. The thresholds used for this simulation and their associated rates are indicated in table 2. These thresholds are set up such that in the DPRMA system there is a very low probability that a video conferencing user's queue will empty and the user will later need to reacquire a reservation. More details about the selection of threshold levels appear in section 6.5.

Since DPRMA video users rarely lose their reservations, the principle time that video users will contend for access to the system is when a new connection is requested. For this study, video users obtained reservations prior to the beginning of the simulation and were active for its entire duration. Therefore, the transmission probability for video users, P_{tvc}, did not have a significant impact on the performance of the overall system. For this reason, in the DPRMA simulation P_{tvc} was arbitrarily set to 0.3.

Table 2
Threshold levels for video conferencing sources.

Queue length	Transition from	Transition to
6	70.667 kbps	141.333 kbps
11	141.333 kbps	282.667 kbps
16	282.667 kbps	565.333 kbps
21	565.333 kbps	1.131 Mbps
26	1.131 Mbps	2.261 Mbps
31	2.261 Mbps	4.523 Mbps
26	4.523 Mbps	2.261 Mbps
21	2.261 Mbps	1.131 Mbps
16	1.131 Mbps	565.333 kbps
11	565.333 kbps	282.667 kbps
6	282.667 kbps	141.333 kbps
1	141.333 kbps	70.667 kbps

Unlike the video conferencing users, the performance of the voice users is affected by the value of the transmission probability, P_{tv}. If this value is too low or too high, the users may suffer excessive delay and consequently may be forced to discard an unacceptable number of packets. The objective here is to determine the value for P_{tv} that will allow the most voice users access to the system while maintaining QOS guarantees.

The number of voice users that can be permitted in the system, N_v, is dependent on the number of video users that have been admitted, N_{vc}. Therefore, the system performance was analyzed with 0 to 5 video conferencing users present. For each case, the maximum number of voice users permitted in the system was determined for a range of P_{tv} values.

The PRMA* protocol was simulated using the same DPRMA system parameters as specified in tables 1 and 2. Unlike DPRMA, however, it was found that the value selected for P_{tvc} has a significant impact on the operation of the PRMA* system. This behavior is expected since the video conferencing users in PRMA* must contend for new slots every time they need to increase their transmission rates. DPRMA eliminates this contention altogether. For PRMA* we found that for the set of parameters used, the best performance results when $P_{\text{tvc}} = 0.3$.

The comparisons of the DPRMA results and the PRMA* results are depicted in figures 6 and 7. These results were obtained by running each protocol for 400 simulated seconds. In figure 6, the maximum number of voice users allowed in the system is plotted as a function of P_{tv}. From this figure it is clear that the advantage of implementing DRPMA comes when more and more video conferencing users are introduced. The DPRMA protocol outperforms PRMA* by allowing more voice users access to the system resources. PRMA* reaches a performance limit with 4 video users and about 95 voice users. DPRMA, on the other hand, can accommodate up to 5 video users and still about 70 voice users.

The reason for the inferior performance of the PRMA* system can be seen clearly in figure 7. This figure shows

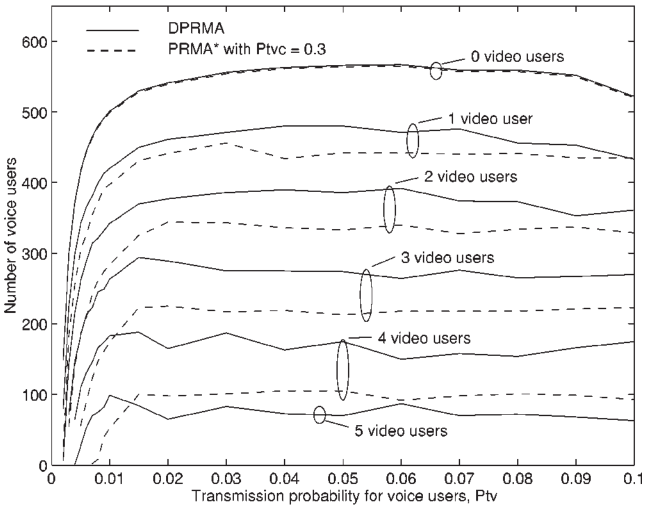

Figure 6. Maximum number of voice users vs. P_{tv}, $P_{\mathrm{vdrop}} = 0.01$, $P_{\mathrm{vcdrop}} = 0.0001$, and $P_{\mathrm{tvc}} = 0.3$.

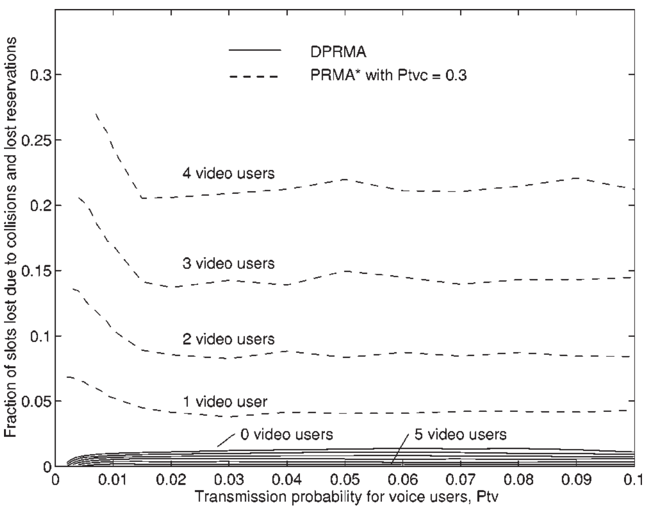

Figure 7. Fraction of slots wasted vs. P_{tv}, $P_{\mathrm{vdrop}} = 0.01$, $P_{\mathrm{vcdrop}} = 0.0001$, and $P_{\mathrm{tvc}} = 0.3$.

the fraction of slots that are wasted in the system. This waste is caused by two mechanisms: slots that are left empty because a user has no more packets to transmit and slots in which the transmissions of contending users collide. These results were also obtained in conjunction with those in figure 6. Consequently, they indicate the fraction of slots wasted when the system is operating at full capacity. In the PRMA* protocol, a slot is left empty each time the user wishes to reserve one fewer slot. Therefore, each time rate requirements decrease, the user could leave multiple slots empty within the frame. Users in the DPRMA system, on the other hand, provide the base station with this rate information in advance. This allows the base station to reallocate the unneeded slots to other users or to declare that they are available for access by new users. The only time slots are left empty is when the user actually runs out of packets to transmit. Figure 7 indicates that many slots are wasted in the PRMA* system due to the two mechanisms. In fact, up to 22% are lost when 4 video users are present. DPRMA proves to be a much more efficient protocol with only up to 1.5% of the slots wasted.

Now that the improved performance that DPRMA offers over PRMA* has been demonstrated for the voice and video example we use the results shown in figure 6 to determine an appropriate value for P_{tv}. We choose a value of 0.05 for future simulations, since this value on average allows the greatest number of voice users to be admitted into the system. It also avoids operating a system with multiple equilibrium points, a common characteristic of ALOHA-based systems [18]. This undesirable situation appears to result in our system when $P_{\mathrm{tv}} \geqslant 0.06$.

More details about the performance of the PRMA* and DPRMA protocols in a voice and video system appear in [19].

6.3.2. Introduction of data users

When data users are introduced, several new parameters must first be evaluated in order to optimize system performance. For the DPRMA system, the parameters in question are transmission probability of data users, P_{td}, and the threshold level pairs. A simple queueing analysis presented in section 6.5 demonstrates that a single threshold level pair provides very good performance for this traffic type. The threshold levels used for this study are $L_{\mathrm{U}} = 30$ and $L_{\mathrm{D}} = 10$, and the allowable transmission rates that are associated with these levels are $c_2 = 70.667$ kbps and $c_1 = 35.333$ kbps. The levels are selected such that the data users lose their reservations very infrequently. The selection of P_{td} is then simplified. It was found that in a data-only system nearly identical performance resulted when $0.006 < P_{\mathrm{td}} < 0.1$. Since data users are given the lowest priority in this study, we choose the lowest value for P_{tv} in the range specified above, $P_{\mathrm{tv}} = 0.007$. Thus, an individual data user is the least likely to transmit a packet in any empty slot.

In the PRMA* system, data users are unable to obtain reservations. They simply attempt transmission whenever a packet arrives, and therefore there is no need to implement threshold levels. The only parameter then that needs optimization is P_{td}. This parameter was varied over a range of values in a voice/data and video/data system. Optimal results varied depending on the number of voice or video users present. However in all cases, reasonably good results were obtained with $P_{\mathrm{td}} = 0.02$. This value is lower than P_{tv} and P_{tvc} in the PRMA* system, and therefore the data users are given lowest priority.

A combined voice, video and data system was simulated for both the PRMA* and DPRMA protocols to determine the maximum capacity that each system is capable of supporting. Several series of simulations were run in which the number of voice users was fixed at 0, 50, and 100 users. In each case, the number of video conferencing users was varied from 0 to 5 and the maximum number of data users that could be supported by the system was determined. Results of this study can be seen in figure 8. In each case, DPRMA

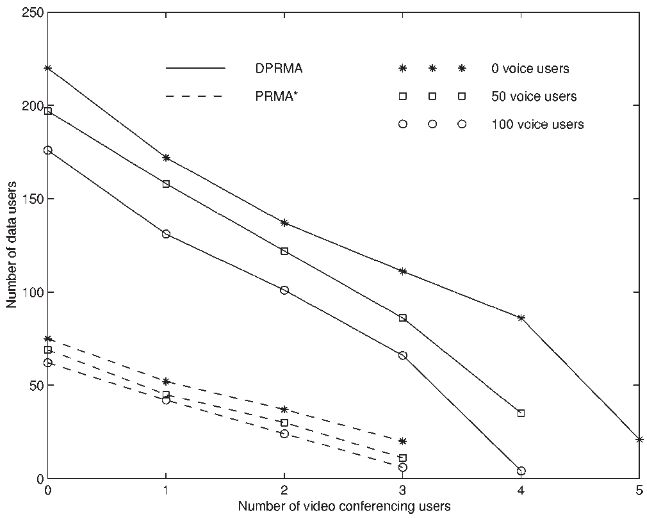

Figure 8. Maximum number of data users permitted vs. number of video conferencing users present. $P_{\rm td} = 0.02$ for PRMA* and $P_{\rm td} = 0.007$ for DPRMA.

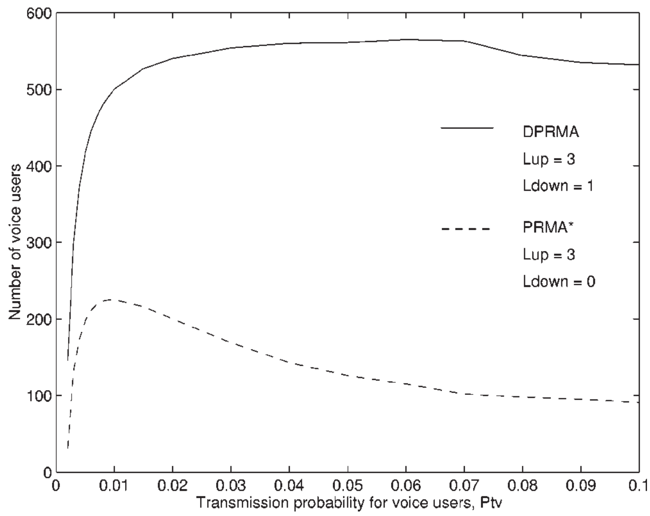

Figure 9. Maximum number of users vs. of $P_{\rm tv}$.

can admit significantly more users into the system than can PRMA*. The number of data users allowed differs by over 70 users in all cases. The limitation of the PRMA* system lies in the increased amount of contention that each user introduces. For video and data users, this contribution is quite significant. Data users add to the contention because they are not permitted to obtain reservations. Video users also increase contention because they frequently are changing their slot reservations. In DPRMA, both of these user types maintain reservations for a significant portion of the simulation time introducing almost no contention into the system.

6.4. Voice only system

So far we have only considered a system where the packet generation rate for voice users is exactly equal to the transmission rate corresponding to one slot in every frame. This situation is not always a realistic one. We consider a situation where a user needs more than one but less than two slots per frame in order to match generation rate to transmission rate. A simulation was run for 400 seconds of simulation time, in a system that had only voice users present. The channel capacity (9.045 Mbps) and the bit rate of the voice users (32 kbps) from the previous system were maintained. For this simulation, however, the channel was divided into frames of 320 time slots. This made the frame length 15 ms.

With this information, we use equation (1) to select transmission rates of $c_0 = 28.267$ kbps and $c_1 = 56.533$ kbps. These two rates straddle the packet generation rate, and they correspond to reserving one or two slots in every frame. Since the transmission rate required by the users would only take on two possible values, only one threshold level pair is required. Several combinations of threshold levels were tested to determine which would permit the greatest number of users to access the system. The best results for PRMA* were $L_{\rm U} = 3$ and $L_{\rm D} = 0$ and for DPRMA were $L_{\rm U} = 3$ and $L_{\rm D} = 1$. These were the values used for this simulation study.

The results of the simulations can be observed in figure 9. The maximum number of users permitted with PRMA* is only 225 compared to 563 permitted with DPRMA. This same system was simulated using the standard PRMA protocol but this resulted in an unacceptable packet dropping probability of over 18%. *Clearly both PRMA and PRMA* are not well suited for use with CBR traffic sources whose packet generation rate is not an integer multiple of the transmission rate of one slot per frame. DPRMA, on the other hand, provides very good results for this case.*

6.5. Threshold levels

One set of system parameters that directly affects the performance of the DPRMA protocol is the transmission queue threshold levels that the mobiles use to update their rate requests. It is important to select these levels so that good system performance can be achieved. The appropriate choices for thresholds may actually be very difficult to determine if the statistics of the generated traffic cannot be well specified. In these cases, a reasonable *estimation* of the required threshold levels will have to be implemented. If, however, the traffic characteristics for a given traffic type are well understood, it may be possible to select the levels that will produce near-optimal system performance.

6.5.1. Voice traffic

A simple example of a traffic type, whose behavior is known is the voice traffic model used in this study. For the system in section 6.4 the packet generation rate of the voice users cannot be exactly matched by the transmission rate. Therefore, each user will have to update its reservation request occasionally to ensure that packets are transmitted in a timely fashion. In order to do this, we choose two

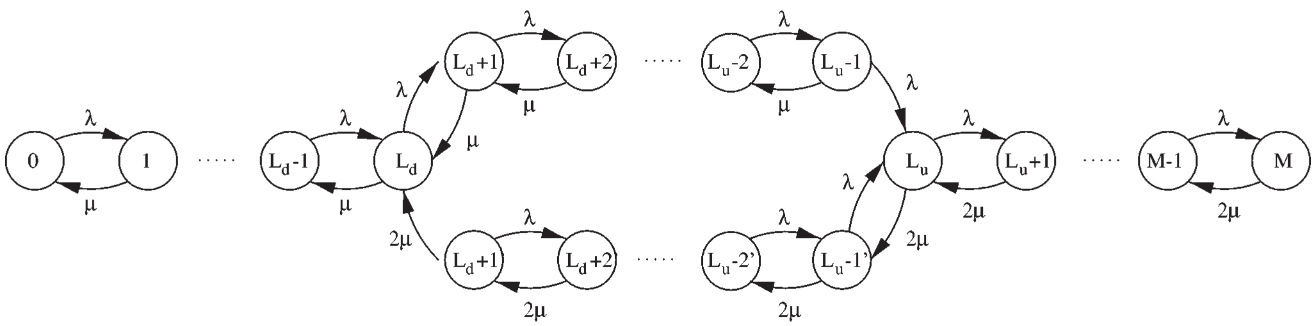

Figure 10. Queueing analysis of data traffic.

rates at which these users can transmit: one just below the packet generation rate and one just above it.

Users with new bursts of information to transmit will initially submit a request of c_0. The system must be set up, however, such that the increase to c_1 occurs before packets become too old. (Recall that packets are considered too old after 40 ms.) When the user is transmitting at the rate c_0, successive packets are sent 15 ms apart. Therefore, when the queue reaches three packets the worst case delay for the third packet will be just under $3 \times 15 = 45$ ms. This is greater than the maximum allowable delay. Thus, to ensure timely delivery of all packets, an increased rate request must be submitted when the queue length grows to 3.

The selection of the lower threshold is more involved. This value should be selected to minimize the probability that the queue level falls to zero and a reservation is lost. In addition, it should be chosen so that the lower threshold is spaced as far from the upper threshold as possible. This reduces the number of times a user will have to submit new rate requests, and the base station will have to reallocate the slots. The smallest threshold level possible is 1. Choosing this value for the lower threshold gives us the maximum distance between our threshold levels and it ensures that reservations are never lost. This second point results from the fact that voice packets are generated at regular intervals, and for this particular system, if a voice user's queue length has fallen to 1, it will always be able to generate a new packet before the next reserved slot arrives.

The choice of the threshold level pair, 1 and 3, proves to be the optimal one for this system and this traffic type. This was tested via simulation by selecting different threshold levels and comparing the performance results. In this situation, this threshold level pair outperformed all other combinations. The explanation provided in this section further justifies the selection of these values.

6.6. Data traffic

The data users modeled in this study generate VBR traffic. Since packets arrive according to a Poisson process, over a small time interval the arrival rate may span an unlimited range of values. This could indicate that multiple threshold level pairs are required. However, since data traffic is not time dependent, its packets can suffer a larger amount of delay than that of real-time users. Therefore, the subsequent queueing analysis will demonstrate that one threshold level pair is adequate to maintain QOS requirements.

The Poisson arrival process allows us to approximate the queueing behaviour of the data user's transmission queue using a Markov chain. For simplicity, we assume that data users obtain reservations whenever they are needed, that all rate increases can be handled immediately, and no data users are ever preempted in favor of real-time traffic users. An example of the Markov chain is depicted in figure 10. In this example, only one threshold level pair is implemented. The L_d state in this example represents the lower threshold, the L_u state the upper threshold, and the M state the maximum number of packets permitted in the queue. The following global balance equations correspond to this figure:

$$\begin{aligned}
\Pi_i &= \rho^i \Pi_0, \quad 0 \leqslant i \leqslant L_u, \\
\mu\Pi_i + 2\mu\Pi_{i'} &= \lambda\Pi_{i-1}, \quad i = L_d + 1, \\
\mu\Pi_i + 2\mu\Pi_{i'} &= \lambda\Pi_{i-1} + \lambda\Pi_{i-1'}, \\
& L_d + 1 < i < L_u, \\
2\mu\Pi_i &= \lambda\Pi_{i-1} + \lambda\Pi_{i-1'}, \quad i = L_u, \\
\Pi_i &= \tfrac{\rho}{2}\Pi_{i-1}, \quad L_u < i \leqslant M,
\end{aligned} \tag{5}$$

where Π_i is the probability of being in state i and $\rho = \lambda/\mu$. With knowledge of L_d, L_u, and M, the above equations can be solved to determine the system's steady state probabilities. In particular, we are interested in Π_0 and Π_M since these will tell us much about the performance of the system. If Π_0 is high, that indicates that the queue is emptying very often. Whenever the queue length falls to zero, there is the possibility that a reservation will be lost, slots will be wasted, and more users on average will be contending for empty slots. Therefore Π_0 should be small. Likewise, if Π_M is too high, then many of the new packets arriving into the system find the transmission queue full. When this happens, data packets are lost. An additional consideration is the maximum queueing delay, $d_{\max}$, that packets will suffer in the system. This maximum delay is dependent on the threshold levels and is determined by

$$d_{\max} = \max\left[\frac{(L_u - 1)Db}{c_k}, \left(\frac{(M - L_d)Db}{c_{k+1}} + \frac{L_d Db}{c_k}\right)\right], \tag{6}$$

where D is the number of bytes is an ATM packet (53), b is the number of bits per byte (8), and c_k is the lowest rate at which the user may request to transmit. Although data traffic is not time dependent, the delay suffered by data packets should be kept within certain limits whenever possible. The end users of the system are generally unwilling to wait for an extensive amount of time for data traffic to be transmitted.

In order to select the appropriate threshold levels for the data users, Π_0, Π_M, and $d_{\max}$ were calculated for all possible combinations of L_d and L_u, where $2 < L_u < L_u^{\max}$, $1 \leqslant L_d < L_u$, and $L_u^{\max}$ is the largest possible value for the upper threshold level. Since the maximum buffer size, B, used in the simulations is 10 kBytes, this value is

$$L_u^{\max} = \lfloor (B/D) \rfloor. \tag{7}$$

Solving this equation results in $L_u^{\max} = 188$. The selected threshold levels had to meet the following requirements: $\Pi_0 < 10^{-2}$ and $\Pi_M < 10^{-8}$ and $d_{\max} < 1.25$ seconds. The following threshold level pair met all three criteria: $L_d = 10$ and $L_u = 30$. This pair was implemented for the data users and was found to produce very good results in the simulations that were conducted in this study.

6.6.1. Multiple threshold level pairs

The selection of appropriate threshold levels for the video conferencing users is more complicated than was the case for the voice or data users. The wide variation in the number of cells that can be generated in one video frame indicates that multiple threshold levels are required for optimal performance. For this study, the exact behavior of the video conferencing traffic is quantifiable, and this knowledge is used to aid in the selection of threshold levels.

The information that is required for the threshold level selection process is the peak bit rate and the minimum bit rate of the traffic type. For the video conferencing traffic used in this study, the peak bit rate is 2.332 Mbps and the minimum bit rate is 0 kbps. These two figures correspond to packet generation rates of 220 cells and 0 cells per video frame, respectively. The model of the video conferencing traffic, however, allows us to determine that the probability that fewer than 7 packets are generated in one cell is

$$\sum_{k=1}^{6} f_k = 1.696 \times 10^{-9}. \tag{8}$$

Since these events occur with such a small probability, we neglect them when choosing the threshold levels for this system. The new minimum bit rate that we use to determine threshold levels for this traffic type is $7Db/F = 7 \times 53 \times 8/0.04 = 74.2$ kbps. The minimum rate request then becomes $c_1 = 70.667$ kbps and the maximum request is $c_7 = 4.523$ Mbps.

Since we are dealing with 7 pairs of threshold levels, c_1 through c_7, the determination of appropriate values becomes very involved. In order to simplify the level selection process we specify that $L_{\mathrm{U}j} = L_{\mathrm{D}(j+1)}$ for all thresholds and $L_{\mathrm{D}1} = 1$. The latter choice was made because selecting 1 for $L_{\mathrm{D}1}$ allows the threshold levels to be spread out over a wider range of values. This should decrease the

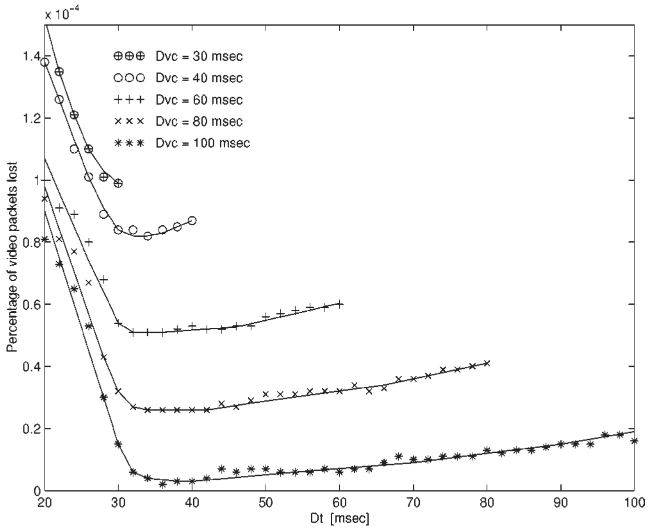

Figure 11. Percentage of packets lost vs. D_t.

number of changes in rate requests that must be submitted. The selection of correct levels is then simply a function of the maximum allowable packet delay, D_{vc}, and the other threshold levels. Each threshold level is selected to ensure that D_{vc} is not exceeded. This is done through the following equation:

$$L_{\mathrm{U}j} = \begin{cases} \frac{D_t - \sum_{i=1}^{j-1} T(L_{\mathrm{D}i} - L_{\mathrm{D}(i-1)})/2^i}{T/2^j} + L_{\mathrm{D}(j-1)}, \\ \quad 1 \leqslant j \leqslant 7, \\ 0, \quad \text{otherwise}, \end{cases} \tag{9}$$

where $L_{\mathrm{U}j}$ is rounded to the nearest integer and initially $D_t = D_{\mathrm{vc}}$. One scenario that this equation does not account for is the amount of time that it takes the DPRMA protocol to accommodate an increased reservation request. If this time interval is substantial, setting $D_t = D_{\mathrm{vc}}$ in equation (9) does not guarantee that D_{vc} is not exceeded. In addition we would like to determine if there are other factors which affect the performance of the system and impact the selection of threshold levels. This would imply that D_t in equation (9) is a function of several variables, including D_{vc} and the reallocation delay, D_r.

We have investigated the effects of different threshold levels by simulating the DPRMA system with only video conferencing users present. In these simulations, the values of D_{vc} and D_t were varied. The corresponding threshold levels for each value of D_t were calculated using equation (9). For each repetition, 6 video conferencing users were active in the system for 400 simulated seconds. The traffic generated by these users was identical for all simulations thereby producing the same traffic load in each case. The results of this study are presented in figure 11.

These results show an interesting system behavior for this traffic type. In all cases, the optimal system performance results when D_t is set to 36 ms. For values of D_t less than 36 ms, the threshold levels are set at values that are too low to prevent the frequent loss of reservations. These reservation losses result in increased packet loss and decreased system performance. When D_t rises above 36 ms, the threshold levels continue to cause the number of lost reservations to go steadily down to zero. However, as is explained in the following, other effects come into play that prevent the percentage of packet loss from decreasing any further.

One of the effects is the actual reallocation delay. Since this delay can cause packets to be queued longer than the maximum value that the threshold levels are designed for, additional packet loss will result. However, since the traffic load for all simulations is identical, the reallocation delay should stay fairly constant throughout the study. This would indicate that we expect the optimal value of D_t to shift for different values of D_{vc}. Since this is not the result we observe, we conclude that D_r is not the only factor that we should be concerned with when selecting appropriate threshold levels.

Closer consideration of the traffic type we are simulating produces some insight as to why the optimal D_t remains constant for different values of D_{vc}. The video conferencing traffic simulated in this study begins generating video frames every 40 ms, with each frame containing between 0 to 220 packets. When D_t is set close to this 40 ms value, the protocol attempts to transmit the majority of the new packets before the next video frame arrives. It does this by emptying the transmission queue fairly quickly and then decreasing its rate request to a lower value until more packets arrive. In doing so, the user's packet transmission behavior is very closely tracking the packet generation behavior. The main difference is that the packet transmission rate ideally never falls to zero. Thus, the rate reservation is maintained and reasonably good results are expected.

When the threshold levels are set up for larger values of D_t a different behavior emerges. In this case, the user does not attempt to transmit each video frame before the next one arrives. The transmission of the frames are allowed to be spread out over a period of two or three video frames. Thus the user's packet transmission behavior is no longer closely matched to its generation behavior. This allows the maximum transmission delay to be exceeded, and an increase in packet loss results. This is easily seen in figure 11 for values of $D_t > 36$ ms.

The results obtained here indicate that the optimal D_t value is actually a function of the video frame length and the reallocation delay. Since the reallocation delay is dependent on system load, an estimated delay should be selected that closely matches D_r in a heavily loaded system. This will anticipate the worst case performance of the system. The results in figure 11 indicate

$$D_r \approx 0.1F \tag{10}$$

for a system under heavy load, where F is the time between video frames. Therefore the optimal value for D_t is

$$D_t^* \approx F - D_r \approx 0.9F. \tag{11}$$

The results shown in these simulations indicate that the selection of appropriate threshold levels is very much dependent on the traffic type. For other video traffic types, equation (11) can be applied if the traffic generates periodic video frames. Other traffic types with much different behaviors will have to be investigated further before appropriate threshold level selection can be implemented.

7. Conclusion

The results shown in this study demonstrate that the DPRMA protocol is well-suited for use in a multimedia wireless network. The protocol is capable of simultaneously providing QOS guarantees to multiple users and to multiple traffic types. In addition, it has been shown to offer improved performance over the PRMA* protocol. In a voice-only system with users requiring a one slot per frame reservation, the two protocols exhibit a similar performance. When the system is changed such that users require more than one slot per frame, the DPRMA protocol provides significantly improved performance over PRMA*. Our result demonstrates the flexibility of the DPRMA protocol for different and differing traffic type characteristics.

The performance improvement that DPRMA offers can be seen particularly in one of the systems studied here that has a combination of different user types present. In the voice and video system, the DPRMA protocol was shown to be capable of admitting as many as 70% more voice users than could the PRMA* protocol. In the combined voice, video, and data system, the disparity was even more pronounced, with DPRMA admitting up to 200% more data users into the system than did PRMA*. Much of the improvement of the DPRMA protocol, as compared with the PRMA*, comes from the following two DPRMA features: (1) the ability to adjust the reservation with minimum overhead, while minimizing the probability of losing the reservation altogether and (2) the priority mechanism in assigning capacity to different user types. The first feature is mostly the result of the multi-threshold control algorithm, a design procedure which was presented here.

References

[1] U. Black, *ATM: Foundation for Broadband Networks* (Prentice Hall, Englewood Cliffs, NJ, 1995).

[2] D.J. Goodman, R.A. Valenzuela, K.T. Gayliard and B. Ramamurthi, Packet reservation multiple access for local wireless communications, IEEE Trans. Comm. 37(8) (August 1989).

[3] N. Amitay and L.J. Greenstein, Resource auction multiple access in the cellular environment, IEEE Trans. on V. Tech. 43(4) (November 1994) 1101–1111.

[4] A. Acampora, Wireless ATM: A perspective on issues and prospects, IEEE Personal Comm. 3(4) (August 1996) 8–17.

[5] M. Conti, C. Demaria and L. Donatiello, Design and performance evaluation of a MAC protocol for wireless local area networks, MONET 2(1) (1997) 69–87.
[6] M.J. Karol, Z. Liu and K.Y. Eng, An efficient demand-assignment multiple access protocol for wireless packet (ATM) networks, Wireless Net. 1(3) (1995) 267–279.
[7] Z. Zhang and A.S. Acampora, Performance of a modified polling strategy for broadband wireless LANs in a harsh fading environment, in: *Proc. GLOBECOM '91*, Vol. 3 (1991) pp. 1141–1146.
[8] D. Raychaudhuri and N.D. Wilson, ATM-based transport architecture for multiservices wireless personal communication networks, IEEE JSAC 12(8) (October 1994) 1401–1414.
[9] X. Qiu, V.O.K. Li and J.-H. Ju, A multiple access scheme for multimedia traffic in wireless ATM, MONET 1(3) (December 1996) 259–272.
[10] W. Crowther, R. Rettberg, D. Walden, S. Ornstein and F. Heart, A system for broadcast communication: Reservation-ALOHA, in: *Proc. 6th Hawaii Int. Conf. Syst. Sci.* (January 1973) pp. 596–603.
[11] M. Eastwood, L. Hanzo and J.C.S. Cheung, Packet reservation multiple access for wireless multimedia communications, Elec. Let. 29(13) (June 24, 1993) 1178–1180.
[12] J.G. Kim and I. Widjaja, PRMA/DA: A new media access control protocol for wireless ATM, IEEE ICC 1 (1996) 240–244.
[13] G. Bianchi, F. Borgonovo, L. Fratta, L. Musumeci and M. Zorzi, C-PRMA: The centralized packet reservation multiple access for local wireless communications, in: *Proc. GLOBECOM '94*, Vol. 3 (1994) pp. 1340–1345.
[14] X. Qiu and V.O.K. Li, Dynamic reservation multiple access (DRMA): A new multiple access scheme for personal communication systems (PCS), Wireless Net. 2(2) (1996) 117–128.
[15] P.T. Brady, A model for on-off speech patterns in two-way conversation, Bell Syst. Tech. J. 48(7) (September 1969) 2445–2472.
[16] W.C. Wong and D.J. Goodman, A packet reservation multiple access protocol for integrated speech and data transmission, IEE Proc.-I 139(6) (December 1992) 607–612.
[17] D. Heyman, T.V. Lakshman, A. Tabatabai and H. Heeke, Modeling teleconference traffic from VBR video coders, in: *Proc. Int. Conf. on Comm. '94*, Vol. 3 (1994) pp. 1744–1748.
[18] A.B. Carleial and M.E. Hellman, Bistable behavior of ALOHA-type systems, IEEE Trans. on Comm. COM-23(4) (April 1975) 401–409.
[19] D.A. Dyson and Z.J. Haas, A dynamic packet reservation multiple access scheme for wireless ATM, in: *Proc. MILCOM '97* (1997).

Deborah A. Dyson received the B.S. degree in electrical engineering from the University of Vermont, Burlington, Vermont, in 1992 and the M.S. degree in electrical engineering from Cornell University in 1998. Her area of research was in medium access control protocols for wireless communication. From 1992 to 1995, she was with the Naval Undersea Warfare Center in Newport, Rhode Island, where she worked with guidance and control system upgrades for the Navy's torpedo systems. She is presently with the Switching Division of 3Com Corporation in Boxborough, Massachusetts.
E-mail: Deb_Dyson@ne.3com.com

Zygmunt J. Haas received his B.Sc. in EE in 1979 and M.Sc. in EE in 1985. In 1988, he earned his Ph.D. from Stanford University and subsequently joined AT&T Bell Laboratories in the Network Research Department. There he pursued research on wireless communications, mobility management, fast protocols, optical networks, and optical switching. From September 1994 till July 1995, Dr. Haas worked for the AT&T Wireless Center of Excellence, where he investigated various aspects of wireless and mobile networking, concentrating on TCP/IP networks. As of August 1995, he joined the faculty of the School of Electrical Engineering at Cornell University as Associate Professor. Dr. Haas is an author of numerous technical papers and holds twelve patents in the fields of high-speed networking, wireless networks, and optical switching. He has organized several workshops, delivered tutorials at major IEEE and ACM conferences, and serves as editor of several journals. He has been a guest editor of three IEEE JSAC issues ("Gigabit Networks," "Mobile Computing Networks," and "Ad-Hoc Networks"). Dr. Haas is a Senior Member of IEEE, a voting member of ACM, and a Vice Chair of the IEEE Technical Committee on Personal Communications. His interests include: mobile and wireless communication and networks, personal communication service, and high-speed communication and protocols.
E-mail: haas@ee.cornell.edu
WWW: http://www.ee.cornell.edu/∼haas/wnl.html