To appear in *IEEE/ACM Trans. on Networking*, vol. 8, no. 1, February 2000.

A Unified Wireless LAN Architecture for Real-Time and Non-Real-Time Communication Services

Sunghyun Choi and Kang G. Shin

***Abstract*— Addressed is how to support *both* real-time and non-real-time communication services in a wireless LAN with Dynamic Time Division Duplexed (D-TDD) transmission. With D-TDD, a frequency channel is time-shared for both downlink and uplink transmissions under the dynamic access control of the base station. The base station (i) handles uplink transmissions by polling mobiles in certain order determined on a per-connection (per-message) basis for transmitting real-time (non-real-time) traffic from mobiles; and (ii) schedules the transmission of downlink packets. To handle location-dependent, time-varying, and bursty errors, we adopt the channel-state prediction, transmission deferment, and retransmission. We consider the problems of scheduling and multiplexing downlink packet transmissions, and polling mobiles for uplink transmissions depending on the channel state. We also establish conditions necessary to admit each new real-time connection by checking if the connection's delivery-delay bound can be guaranteed as long as the channel stays in good condition without compromising any of the existing guarantees. Finally, the performance of the proposed protocol is evaluated to demonstrate how the protocol works, and study the effects of various parameters of the protocol.**

***Keywords*— Wireless LAN, MAC protocol, QoS-sensitive communication, priority scheduling, location-dependent errors, polling, Dynamic Time Division Duplexing (D-TDD), admission tests.**

1 Introduction

Wireless local area networks (WLANs) are emerging as an attractive alternative, or complementary, to wired LANs [1, 2], because they enable us to set up and reconfigure LANs easily without incurring the cost of wiring. WLANs are characterized as high-speed wireless systems which cover relatively small geographical areas, as compared to other wireless systems such as cellular, PCS, and mobile data radio systems. WLANs are expected to be the solution to the problem of meeting the growing demand for mobile clients to have access to the existing high-speed wired networks. As the need for broadband multimedia communications involving digital audio and video grows, it is increasingly important for communication systems to support various traffic with Quality-of-Service (QoS) guarantees.

Depending on its distinct characteristics and QoS requirements, diverse traffic is categorized into two classes as summarized in Table 1: (1) class-I real-time traffic such as voice and video that requires bounded delays, but is usually tolerant of some packet losses; and (2) class-II traffic like the conventional data services that requires loss-free transmission, but requires no bounded delay. Class II can be divided further into two subclasses: (a) class II-A which is delay-sensitive like FTP and remote log-in; and (b) class II-B which is delay-tolerant like paging and e-mail. Class II-A is given priority over class II-B. Real-time communication service deals with the first traffic class, in which the packet delivery delay is bounded at the cost of some packet losses depending on the time-varying channel condition, while non-real-time communication service is for the second traffic class. In this paper, we consider a unified WLAN architecture, composed of a base station (BS) and a number of mobiles, which provides both real-time and non-real-time communication services.

The integrated communication services considered here can be divided into four types: (1) downlink real-time; (2) uplink real-time; (3) downlink non-real-time; and (4) uplink non-real-time services, where the downlink (uplink) is for transmitting BS-to-mobile (mobile-to-BS) traffic. To guarantee the timely delivery of packets, real-time traffic is handled by a connection-oriented service, while non-real-time traffic is handled by a connection-less service. Dynamic Time Division Duplexed (D-TDD) transmission is used, and hence, the same frequency channel is time-shared for downlink and uplink transmissions under the dynamic access control of the BS.

Due to the different nature of uplink and downlink transmissions, they are treated differently. For the downlink, the BS schedules the transmission of those packets which had already arrived at the BS. On the other hand, the exact status of uplink packets in mobiles is not known to the BS. Thus, mobiles request the BS for permission to trans-

	Class I	Class II-A	Class II-B
Name	real-time	non-real-time	
Examples	voice & video	remote login	e-mail
Delay	bounded	sensitive	tolerable
Loss	loss-tolerant	zero-loss	

Table 1: The classification of heterogeneous traffic

The work reported in this paper was supported in part by the US Department of Transportation under Grant No. DTFH61-93-X-00017. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the funding agency.

Sunghyun Choi is with Philips Research, Briarcliff Manor, New York 10510 USA (e-mail: sunghyun.choi@philips.com). This work has been done while the author was with the University of Michigan as a graduate student.

Kang G. Shin is with the Real-Time Computing Laboratory, Department of Electrical Engineering and Computer Science, The University of Michigan, Ann Arbor, Michigan 48109 USA (e-mail: kgshin@eecs.umich.edu).

mit their uplink packets, which, upon reception of the requests, schedules uplink transmission permissions. These requests are made on a *per-connection* basis for the real-time service, and on a *per-message* basis for the non-real-time service, where a message is composed of a number of packets. The BS can give a mobile a packet-transmission permission by *polling* the mobile. Real-time traffic is given priority over non-real-time traffic for its timely delivery.

The WLAN system should provide (1) timely delivery of real-time packets while minimizing packet losses; (2) virtually error-free transmission of non-real-time packets; (3) good average delay and throughput performance of non-real-time traffic utilizing the bandwidth left unused by real-time traffic; (4) fair usage of a channel for non-real-time traffic among mobiles; and (5) low latencies for non-real-time packet transmissions and real-time connection setup, and in handling hand-off requests. We design a medium access control (MAC) protocol which can support these requirements, the scheduling of uplink and downlink packet transmissions, and the admission control of newly-requested real-time connections in a dynamic environment with location-dependent and bursty errors.

The paper is organized as follows. Section 2 states the specification and assumptions of the wireless network under consideration. In Section 3, we describe the MAC protocol to support both real-time and non-real-time traffic. Section 4 presents the run-time scheduling of real-time and non-real-time traffic, and the admission-control scheme for real-time connection requests. The proposed protocol is evaluated in Section 5. Section 6 discusses related work, putting our protocol in a comparative perspective. Finally, the paper concludes with Section 7.

2 Assumptions and System Specifications

Before presenting our protocols, we state the assumptions and specification of the WLAN under consideration.

2.1 Network Specification

The WLAN of our interest is based on an infrastructure that is composed of a wired backbone network and a (possibly large) number of base stations (BSs) which cover the whole geographic area in service. Each BS handles an area, called a *cell*, where all the mobiles in the cell communicate through the BS. In this system, the downlink and uplink in a cell can respectively be thought as the end-most and front-most links of a multi-hop end-to-end connection. Since wireless links usually have much less bandwidth than the wired counterpart, the former might become the bottleneck of the end-to-end communication performance. We will, in this paper, focus on the uplink and downlink transmissions within a single cell.

Each mobile is assumed to transmit packets with some limited transmission power so as to reach the BS only, which is located at the "center" of the cell. So our system

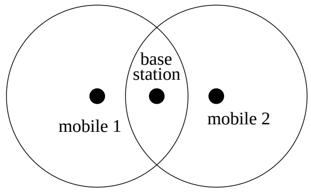

Figure 1: The hidden terminal problem.

differs from peer-to-peer communication systems which do not require any infrastructure, e.g., ad-hoc networks. The lower the transmission power, the better, since (1) mobile devices operate with batteries; and (2) the transmitted signals in a cell are nothing but inter-cell interferences to other cells. In this sense, it is desirable to transmit packets with as low a power as possible, which may, unfortunately, cause the *hidden* terminal problem [2]. In Figure 1, the transmission ranges of mobiles 1 and 2 do not allow them to hear each other, but can both be heard by the BS in between. Mobiles 1 and 2 are hidden terminals to each other. So, in a cell, the uplink (mobile-to-BS) is not a broadcast channel while the downlink (BS-to-mobile) is. Hence, mobiles are assumed not to listen directly to one another even if some can, depending on their relative locations. When a mobile wants to send a packet, regardless whether the packet is destined for another mobile in the same cell or a mobile or static node outside the cell, it first sends the packet to its BS. The BS will then forward the packet to its final destination, sometimes via a wired backbone network, i.e., when the packet is destined for a node outside the cell. When the BS transmits packets through the broadcast downlink, all but the destination mobile in the cell ignore them. Note that only the BS can determine if a collision has occurred in the uplink channel. In summary, the system under consideration can be considered as a network with the *star topology* whose center is the BS.

2.2 Dynamic Time-Division Duplexed Transmission

Dynamic Time Division Duplexed (D-TDD) transmission is used in the system, i.e., a wireless channel over a frequency slot is time-shared for both downlink and uplink transmissions under the BS's dynamic access control. We could instead use (1) Frequency Division Duplexed (FDD) transmission as in most cellular systems, in which two different frequency channels are allocated for uplink and downlink transmissions; or (2) static TDD in which a portion (usually a half) of each time frame is allocated for the uplink and the remaining portion is for the downlink. Even though these static duplexing transmissions are simpler, and adopted in most currently available cellular systems, D-TDD is shown to offer more efficient link utilization in case of unbalanced and time-varying uplink and downlink traffic [3]. Considering the growing number of applications involved with unbalanced two-way traffic (e.g., web browsing), D-TDD is a very promising and important design choice.

2.3 Location-Dependent Errors and Their Control

It is well-known that a wireless channel is inherently unreliable due mainly to noises, interferences, fadings, etc. We use a *hybrid* of Forward Error Correction (FEC) and Automatic Repeat reQuest (ARQ). That is, the receiver tries to correct errors first, and if unsuccessful, a decoding failure is declared, implying the detection of a packet error. A selective-repeat ARQ is used, so the receiver sends an ACK or NAK for each received packet according to the decoding result, and only NAKed packets will be retransmitted by the sender. However, real-time packets can be retransmitted only for a limited time due to their timing constraints or deadlines, which should be met by appropriately scheduling packet transmissions. Wireless channels, in many cases, are known to experience location-dependent, time-varying, and bursty errors, switching back and forth between good and bad states over time [4–7]. A *channel* is defined between each mobile and the BS. Channel probe and transmission deferment are adopted to handle the above-mentioned types of errors, i.e., before transmitting each real-time packet over a channel, the channel condition is probed through a hand-shaking mechanism, and if the condition is bad, the transmission is deferred, and a different packet is scheduled for transmission. This in turn reduces the need of retransmissions significantly. Since the prediction of channel condition/state cannot be perfect, a packet transmission may fail. Deferred or erroneously-transmitted real-time packets must be scheduled for (re)-transmission before their deadline in a best-effort manner so that they won't compromise the other undeferred packets' deadline guarantees. We assume that certain channel coding is used for the receiver to detect and correct errors in packets it received. However, the MAC protocol should provide the means of channel probing and feedback from the receiver to the sender for possible retransmissions.

Throughout the rest of this paper, we will ignore the packet-propagation delay, since it is usually small relative to the other delay components like queueing and transmission delays in a cell.[1] Finally, all packets, like ATM cells, are assumed to be of the same fixed size. So, a long message is fragmented into a number of packets of identical size at the transmitter. With a selective-repeat ARQ scheme, fragmentation could be effective. Otherwise, a long message received in error could result in the retransmission of the entire message, wasting a significant amount of network resources.

[1] A cell in this paper refers to a *micro-cell*, which has coverage of a few hundred meters, or a *pico-cell*, which covers small indoor areas [1].

3 Protocol Description

This section describes the proposed protocol that supports both real-time and non-real-time communication services. In this protocol, due to their different QoS requirements, real-time and non-real-time packets are treated differently. Table 2 summarizes how the two communication services are differentiated and supported. Throughout this and the next sections, we give a detailed account of each entry in the table. We will first outline the protocol, and then present in the next section the details of both the scheduling and admission-control schemes.

3.1 Downlink vs. Uplink

A BS has full control of the transmission of all downlink and uplink packets within its cell. For the downlink, the BS just schedules the transmission of the packets which had already arrived at the BS. In contrast, the exact number of pending uplink packets in each mobile is not known to the BS. Mobiles request the BS's permission to transmit uplink packets, and upon reception of these requests, the BS schedules uplink packet transmissions. These requests are made on a per-connection basis for the real-time service, and on a per-message[2] basis for the non-real-time service as explained below. The BS gives a mobile packet-transmission permission by *polling* the mobile, i.e., checking if it has a packet to transmit. That is, the BS schedules the order of both polling and downlink packet transmissions.

3.2 Real-Time vs. Non-Real-Time

The real-time communication service is provided through connection-oriented communication. Real-time traffic like audio and video usually arrives regularly or periodically and a real-time session/connection usually lasts for a long time, as compared to randomly-arriving non-real-time traffic. Moreover, it is desired to guarantee the requested delivery-delay bound for *each* real-time packet, and hence, we need to set up a real-time connection. For a connection to be admitted, a request should be made to the BS with the connection specification including the requested delay bound. Then, the BS performs an admission test to decide if it is possible to guarantee the requested delay bound without violating the existing connections' guarantees. At run-time, the BS performs priority-based scheduling on packets to guarantee the required delay bounds for all admitted real-time connections provided their channels stay in good condition. A connection is defined to be *uni-directional*, i.e., downlink or uplink. For the sake of notational simplicity, it is assumed that an *active* mobile has only one connection even though it is usually expected to have at least two connections (i.e., one for uplink and the other for downlink) at a time. Active mobile i with downlink (uplink) real-time connection i is called downlink (up-

[2] A message, in this paper, is referred to as a number of non-real-time packets, which arrive at the same time.

	Real-Time	Non-Real-Time
Connection	Connection-oriented	Connection-less
Request	Per-connection based	Per-message based
Admission Test	Yes	No
Scheduling	Deadline-driven	Round-robin-variant

Table 2: Comparison of real-time and non-real-time communication services in the proposed protocol.

link) real-time mobile i.

Non-real-time communication service, on the other hand, is connection-less (or best-effort). Whenever a non-real-time packet arrives at the BS, its transmission is scheduled by the BS. If an uplink message arrives, or is generated at a mobile, the mobile should request the BS's permission to transmit the message; the BS then schedules the transmission permission for that mobile. Some form of scheduling is also needed to give each mobile fair access to the wireless link. Use of the wireless link is prioritized; when there is a pending schedule to poll mobiles for uplink real-time packets or there are "eligible" (to be defined later) downlink real-time packets, the BS will continue to poll mobiles for uplink real-time packets or transmit downlink real-time packets. However, whenever there is no such a polling schedule nor any downlink real-time packet to transmit, the BS will transmit downlink non-real-time packets or poll mobiles for the transmission of uplink non-real-time packets. As will be described later, the prioritized channel usage becomes more complicated when location-dependent errors must be handled efficiently.

3.3 Slotted Channels

A frequency channel is slotted along the time axis. There are three types of slots: *packet-transmission slot*, *transmission-request slot*, and *control mini-slot*. For simplicity, it is assumed that both packet-transmission and transmission-request slots have the same size T_s while the control mini-slot has the size T_{ms} $(= T_s/K)$, where K is an even number much larger than 1. First, a (fixed-size) downlink or uplink packet is transmitted in a packet-transmission slot. Within a packet, the following information can be piggybacked:

- Information within a downlink packet to give a mobile permission to transmit uplink packets;
- ACK/NAK for a previously-received packet within a downlink or uplink packet;
- A message-transmission or connection-setup request by a mobile within an uplink packet.

Second, the transmission-request slot is used by mobiles to convey a request to the BS. Specifically, the usage is three-fold:

- A non-real-time message transmission can be requested with (mobile ID, Class II-A/B ID, the number of packets in the message);
- A real-time connection-setup request can be made with (mobile ID, connection specification);
- A hand-off request can be made by a mobile, which is communicating with an adjacent BS in the overlapping region of two cells, with (previous cell ID, mobile ID, connection specification).

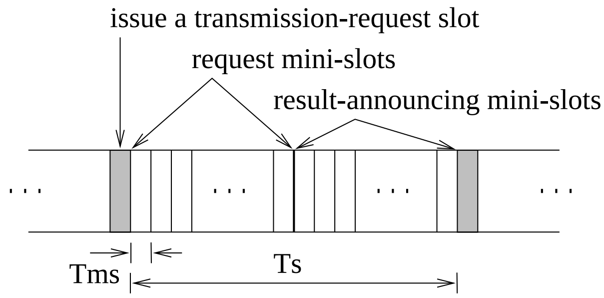

Figure 2: The structure of a transmission-request slot.

Finally, using the control mini-slot, a control packet is transmitted by the BS or by a mobile. A control packet can be used for the following cases:

- For the BS to poll a mobile for transmission of an uplink packet;
- For the BS to issue a transmission-request slot by announcing that the next slot is a transmission-request slot;
- For the BS to announce the admission-test result for a newly-requested or handed-off real-time connection;
- For a real-time mobile to return the transmission permission to the BS without transmitting any packet;
- For the BS or a mobile to send an ACK/NAK for each received packet;
- For the BS to send or receive a channel-probing control packet to predict the channel condition a mobile is experiencing.

More than one of the above can be piggybacked in a control packet, and only the last three cases use uplink control packets.

3.4 Transmission Request Access

As shown in Figure 2, a transmission-request slot is divided into two parts: (1) the first half is the set of $K/2$ *request mini-slots* used by mobiles; and (2) the second half is the set of $K/2$ *result-announcing mini-slots* corresponding to the previous request mini-slots. Then, the first K_{ho} request mini-slots are reserved for real-time hand-off requests. The request mini-slots are used by a slotted ALOHA-like random access protocol: when the BS issues a transmission-request slot, each mobile with pending requests decides whether to make a request with probability p_r or not (i.e., in a p_r-persistent manner). If it decided to send the request, the mobile randomly chooses one of the last $(K/2 - K_{ho})$ request mini-slots, then sends it in that chosen mini-slot. One difference with real-time hand-off requests is that they are made through one of the first K_{ho} request mini-slots. The result of each of $K/2$ mini-slots can be success,[3] or collision, or empty/unused. Using each of the next $K/2$ downlink mini-slots, the result of the corresponding request mini-slot is announced.

The reason why K_{ho} mini-slots are reserved for real-time hand-off requests is because without this reservation, the hand-off requests, which should be made in a timely manner to guarantee the connection's deadline, might be swamped by many requests for non-real-time traffic transmissions. If any of K_{ho} reserved mini-slots results in collision, all request mini-slots in the following transmission-request slot are dedicated to real-time hand-off requests (i.e., $K_{ho} = K/2$) so that the collided hand-off requests can be made successfully with a much higher probability. The initial value of K_{ho} can be adaptively set according to the hand-off rates. For example, if there was a collision in a reserved request mini-slot, the size can be increased by one (and this new value is used after the request slot with $K_{ho} = K/2$ as described above), then if there is no such collision for a certain time period, it can be decremented by one.

The request access probability p_r is determined by each mobile independently of others. For the first attempt of a message-transmission request, p_r is set to 1. If the request becomes unsuccessful, then p_r is decreased according to

$$p_r := \frac{p_r}{p_r + 1}, \tag{1}$$

thus following the sequence $1, 1/2, 1/3, \cdots$ (i.e., the "harmonic" back-off) [8]. Each unsuccessful request should be retransmitted p_r-persistently until it is successfully transmitted. Note that there are two ways to transmit a request: one through a transmission-request slot, and the other by piggybacking it in an uplink packet. Which of the two ways should be used depends on the request type. First, for a non-real-time message-transmission request, if a mobile has not yet transmitted all of its pending uplink packets of which transmission requests had already been made to the BS, then, upon arrival of another message, the mobile need not use a contention-based transmission-request slot. Instead, the request can be piggybacked in an uplink packet. Second, for a real-time connection-setup request, the mobile transmits the request using either of two ways, whichever becomes available first. Third, a hand-off request can be made only through a transmission-request slot since the BS is unaware of the existence of that mobile. In this case, the mobile will access transmission-request slots in a 1-persistent manner. Due to the limited packet size, one may not be able to piggyback all pending requests in a single uplink packet. In such a case, priority is given in the order of (1) real-time connection-setup request; (2) class II-A message-transmission request; and (3) class II-B message-transmission request.

[3] Due to the *capture effects* [2], a request can be transmitted successfully even in the presence of concurrent requests from other mobiles.

3.5 Error-Handling Mechanisms

To handle location-dependent, time-varying, and bursty channel errors, the channel state can be predicted via channel probing before a packet is transmitted. That is, before polling a mobile or transmitting a downlink packet to the mobile, the BS transmits a probing control packet to the mobile, which then returns the control packet to the BS. If the BS does not receive the probing control packet correctly from the mobile, the channel is considered bad, and the polling or transmission is deferred. The channel condition is estimated before each real-time packet transmission in order to reduce the need for retransmitting real-time packets, since retransmission can be really harmful in meeting deadlines while it is optional for the non-real-time case (to reduce its significant overhead).

The ACK/NAK is sent from the receiver upon packet reception.[4] The most desirable way to send an ACK/NAK is to piggyback it in another packet. ACK/NAK for an uplink packet can always be piggybacked since after each uplink packet transmission, a downlink packet (or a control packet) will always be transmitted from the BS. Note that a downlink (control) packet and the piggybacked ACK/NAK may have different destination mobiles. The packet transmitted after a downlink packet can not be a regular uplink packet from the destination mobile of the previous downlink packet, because the destination mobile transmits an ACK/NAK via an uplink control packet right after it received a downlink packet. If the destination mobile of a downlink packet does not know the packet's destination due to an error in the packet, then the mobile won't transmit NAK either. If the BS does not receive an ACK/NAK within T_{ms} after transmitting a downlink packet, the packet will be assumed lost and hence retransmitted later. In summary, (1) ACK/NAK for each packet is transmitted right after its reception; (2) ACK/NAK for an uplink packet is always piggybacked in a downlink (control) packet; and (3) after a downlink packet transmission, a control packet containing ACK/NAK is transmitted from the destination mobile unless an uplink packet transmission from that mobile is scheduled next.

When a mobile is polled for a real-time (non-real-time) packet transmission, if it does not have any real-time (non-real-time) uplink packet to transmit,[5] the mobile transmits an uplink control packet, saying that it has no packets to transmit. So, if the BS does not receive any (control) packet within the mini-slot period T_{ms}, the polling control packet had not been received by the mobile due to an error, implying that the mobile's channel is bad, and hence the polling is deferred. The issue of handling this deferment is addressed next.

[4] Actually, right after determining if the received packet can be corrected by the channel decoder.

[5] Even though this should not happen with the non-real-time case if everything went correctly because the polling is based on transmission requests from mobiles

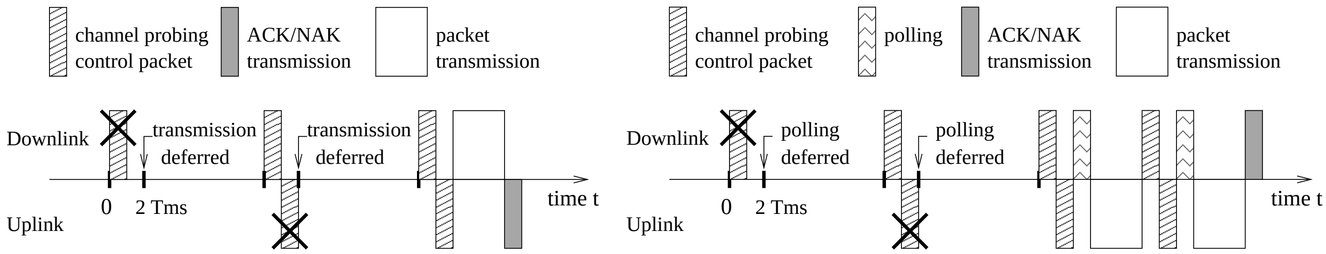

(a) Downlink transmission (b) Uplink transmission via polling

Figure 3: Real-time communication slot structures with channel probing.

3.6 Real-Time Communication Service

Real-time connection i is specified by a triplet (M_i, T_i, D_i), where

- (M_i, T_i): M_i is the maximum number of packets that can arrive in an interval of length T_i;
- D_i: the packet delivery delay (i.e., from the mobile to the BS, or from the BS to the mobile, hence a fraction of the target end-to-end delay) bound for real-time packets of connection i, which has the following relationship with T_i:

$$D_i \geq D_i^{min} = \begin{cases} T_i, & \text{for downlink,} \\ 2T_i, & \text{for uplink.} \end{cases} \quad (2)$$

As will be clear later, the BS will admit new connections and schedule packet-transmissions and channel-polling based on the minimum delay bound D_i^{min} so that each packet's delay is bounded by D_i^{min} as long as the channel condition continues to be good. Once a packet transmission is deferred or fails, the BS will attempt to deliver the packet within D_i in a best-effort manner. Note that the larger D_i, the more likely packets will be delivered in time under a time-varying channel condition.

To set up a real-time connection, the source node (which is a mobile if it is an uplink connection) will request the BS to set up the connection with its specification triplet. The BS will then perform the admission test given in Section 4.2. Depending on the test result, the new connection will be admitted or rejected. This test result is conveyed to the mobile through a downlink control packet. If it is admitted, a connection is set up: if it is an uplink connection, the BS starts to schedule the polling for that mobile, or if it is a downlink connection, it starts to schedule downlink real-time packets arriving at the BS. Even though it was described in the context of a new connection request, the same procedure can be used for a hand-off connection request. For an uplink connection, the BS just polls mobiles assuming the contracted traffic characteristics because the BS does not know the exact status of a mobile's pending real-time packets. The control packet for polling a mobile contains the mobile ID. When a mobile is polled, it can transmit one pending real-time packet through the uplink.

Figure 3 shows how the channel probing and deferment work for both downlink and uplink transmissions, where the label 'X' represents loss of a control packet. The BS will attempt to transmit each deferred polling or downlink packet later so that its delivery bound can be met if the channel condition improves before the deadline expires without compromising other undeferred packet transmissions as described in Section 4.1. The channel-state probing can be used adaptively; that is, a real-time connection can be serviced without this process initially, but later if this connection is determined to suffer from time-varying errors, this process can be invoked. Especially, a mobile which stays close to the BS will be experiencing a very good channel, so it will not need this process at all. Note that this channel probing imposes bandwidth overhead; $2T_{ms}$ is consumed for each channel probing. As shown in Figure 3, the transmission of a real-time packet (both downlink and uplink) consumes $3T_{ms} + T_s$. Note that the ACK/NAK of an uplink packet can always be piggybacked in a downlink (control) packet, so it was not counted.

3.7 Non-Real-Time Communication Service

The BS schedules the transmission of downlink non-real-time packets or the polling for uplink non-real-time packet-transmission permissions with lower priority than the scheduled real-time traffic. Important concerns about non-real-time communication service include:

- Virtually error-free transmission via our retransmission-based error control mechanism;
- (Long-term) fair access to the link of mobiles: a round-robin-based packet/polling scheduling is used for fairness;
- Maximizing aggregate throughput: our non-real-time scheduling is geared toward reduction of control packet overheads associated with ACKs/NAKs and polling. Moreover, channel probing is not used as long as a mobile's channel continues to be good.

4 Run-Time Scheduling and Admission Control

As mentioned earlier, D-TDD requires uplink and downlink real-time packets to be multiplexed and scheduled over the same frequency channel. We need two different packet/polling scheduling policies for real-time and non-real-time services due to their distinct requirements. To set up a real-time connection, the BS performs an admission test to check if it can guarantee the timely delivery of every real-time packet from the new channel without compromising other existing guarantees.

4.1 Real-Time Traffic Scheduling

The conventional first-in-first-out (FIFO) scheduling is unable to bound the packet-delivery delay. So, we adopt a non-preemptive earliest-deadline-first (EDF) scheduling algorithm — the earlier the deadline the higher the priority. Liu and Layland [9] proved that the deadline-driven scheduling is optimal among all dynamic-priority scheduling policies when the deadline of each task is equal to the end of its period. Our real-time scheduling and admission control are based on the real-time channel protocol in [10].

4.1.1 Downlink Packet Scheduling

The *logical arrival time* for the n-th packet, arrived at the BS at time $t_i(n)$, of real-time connection i is defined as:

$$l_i(n) = \begin{cases} t_i(n), & \text{for } 1 \le n \le M_i, \\ \max\{l_i((\lfloor n/M_i \rfloor - 1)M_i + 1) + T_i, \ \ t_i(n)\}, \\ \qquad \text{for } n > M_i. \end{cases} \tag{3}$$

The *logical deadline* this packet is then $l_i(n) + D_i^{min}$. A packet with $t_i(n) < l_i(n)$ (i.e., arrived too early) is not transmitted until it becomes 'current' since its immediate transmission might result in buffer overflow at the destination mobile. By using this logical deadline, one can protect well-behaving connections from misbehaving ones. For example, much more than M_i packets could arrive during T_i from a misbehaving connection with the contract (M_i, T_i, D_i), which can cause violation of the bounded-delivery guarantees of packets belonging to other well-behaving connections if their logical deadlines are not enforced. Note that D_i^{min} is added to the logical arrival time instead of D_i in the calculation of the logical deadline. In this way, the BS attempts to deliver each real-time packet within its minimum delay bound as long as the destination mobile's channel remains in good condition. A deferred or erroneously-transmitted packet will be (re)-transmitted when the channel is predicted and its actual deadline D_i can be met.

When the n-th packet, $p_i(n)$, of real-time connection i, C_i, arrives at the BS from the wired network at time $t_i(n)$, it is placed in a connection-specific FIFO queue $Q[i]$, and its logical deadline is calculated. Then, the packet's connection ID is fed into queue **R** (meaning 'ready') if $t_i(n) \ge l_i(n)$, and the entries of this queue are sorted according to the corresponding packets' deadlines. On the other hand, if $t_i(n) < l_i(n)$, this connection ID will be held in another queue **H** (meaning 'hold'), and will later be fed into **R** at time $t = l_i(n)$ (i.e., when it becomes current). When **R** is selected for service, the connection, say C_j, whose ID is at the head of **R**, is served, and then the connection ID is removed from **R**. To serve C_j, the BS first (i) drops all those packets in queue $Q[j]$ whose deadlines cannot be met, (ii) probes the destination mobile's channel, and finally, (iii) transmits the packet at the head of queue $Q[j]$ if the channel condition is predicted to be good; else defers its transmission and places the connection ID in a FIFO queue **D** (meaning 'deferred').

If the packet transmission upon prediction of a good channel condition results in a NAK, the connection ID is placed into a FIFO queue **B** (meaning 'back-logged'). Both queues **D** and **B** are served similarly to **R**. For convenience, the process of serving queue **X** is referred to as a "queue **X** service," where **X** is **R** or **D** or **B**. Note that the (re)-transmission of an once-deferred or unsuccessfully-tried packet should not cause other undeferred packets to miss their deadlines. To achieve this, we use a book-keeping counter, called the *credit counter* (CC), which records the time both queues **D** and **B** can be served continually without delaying queue **R** service. The CC is updated as follows:

- $CC := CC + 5T_{ms} + T_s$, if a queue **R** service results in a deferment;
- $CC := CC + 2T_{ms}$, if a queue **R** service results in a packet transmission;
- $CC := CC - 2T_{ms}$, if a channel is probed for a queue **D** or **B** service;
- $CC := CC - (T_{ms} + T_s)$, if a packet is transmitted via a queue **D** or **B** service;
- $CC := CC - T_{tr}$, if the link is used for non-real-time traffic or transmission-requests for the duration of T_{tr};
- $CC := 0$, if $CC < 0$ after any of the above updates.

As will be described later in Section 4.2, $5T_{ms} + T_s$ is assumed necessary for each packet transmission during the admission-control phase, because packet transmission without deferment takes $3T_{ms} + T_s$ and additional $2T_{ms}$ is reserved for one extra channel probing for each packet. This is why $5T_{ms} + T_s$ ($2T_{ms}$) is added when a queue **R** service results in a deferment (packet transmission).

When $CC \ge 3T_{ms} + T_s$, queue **D** is given priority over queue **R**. Then, **B** is given next priority after **D**. Basically, the channel usage is prioritized in the order of **D**, **B**, and **R** if $CC \ge 3T_{ms} + T_s$, and **R**, **D**, and **B** if $CC < 3T_{ms} + T_s$. Whenever the wireless link becomes free (after completing an on-going transmission), the BS will choose the non-empty highest-priority queue for the next channel usage. When queue **D** is chosen, the connection specified by the D_{index}-th entry of **D** is served. Unless the queue **D** service results in a deferment, the D_{index}-th entry is removed, and all the following entries are shifted toward the head within **D**, i.e., the $(i+1)$-th entry becomes the i-th entry for all $i \ge D_{index}$. The index D_{index} is updated as follows, with the initial value $D_{index} = 0$:

- $D_{index} := 0$, if **D** becomes empty;
- $D_{index} := D_{index} + 1$, when a queue **D** service results in a deferment;
- $D_{index} := 1$, when (1) an entry in queue **R** is fed into the empty queue **D**, or (2) the last entry in queue **D** is served, or (3) CC becomes less than $3T_{ms} + T_s$.

In fact, if $D_{index} = 1$, there is another condition — that is, the queue **D** service flag $D_{flag} = 1$ — to serve queue **D**.

Otherwise, the link is used to serve the queue with the next priority. The value of D_{flag} is updated as follows, with the initial value $D_{flag} = 0$:

- $D_{flag} := 1$, if any packet (other than control packets) is transmitted over the channel;
- $D_{flag} := 0$, when an entry in queue **R** is fed into the empty queue **D**, or a queue **D** service results in a deferment when $D_{index} = 1$.

Note that by using D_{flag}, once a packet transmission is deferred, at least one packet is transmitted before serving this packet again. This is reasonable since trying to continuously serve a deferred packet can waste bandwidth significantly due to continuous channel probes without its successful transmission. The same rule is used for queue **B** by replacing D_{index} and D_{flag} with B_{index} and B_{flag}, respectively.

Note that the entries in real-time queues **R**, **D**, and **B** specify connection IDs, not packet IDs, and the packet with the smallest deadline in each connection — irrespective of which queue specifies this connection for a service — is transmitted. In this way, we can maximize the chance to deliver packets before their deadlines. Figure 6 summarizes how to determine which queue to serve, including the non-real-time queues that will be covered later.

4.1.2 Polling Order

As mentioned earlier, the polling order for uplink real-time mobiles is determined based on the contracted traffic characteristics as follows:

1. The BS generates a polling request for uplink real-time connection i with (M_i, T_i, D_i), periodically once every T_i.
2. The request generated at time t_i is sorted in queue **R** according to the deadline $t_i + D'_i$ along with downlink connection IDs, where $D'_i = D_i^{min} - T_i = T_i$. Conceptually, a polling request for connection i is considered as a packet of size $M_i(3T_{ms} + T_s)$ including the channel probing overhead.
3. When the request is at the head of queue **R**, the BS will serve[6] uplink real-time mobile i up to M_i times consecutively until the channel for this mobile is predicted to be bad, or this mobile does not have any more pending packets to transmit.

The deadline-driven scheduling plus the admission control described in Section 4.2 will ensure that the BS can start probing the channel for uplink connection i as late as $D'_i - M_i(3T_{ms} + T_s)$ after generation of the last polling request. Then, for the k-th polling request generation time G_k (where $G_{l+1} - G_l = T_i$ for all l), all of M_i packets generated during the interval of $[G_{k-1},\ G_k]$ will be transmitted by time G_{k+1} as long as the channel condition continues to be good. So, the worst-case delay will be $G_{k+1} - G_{k-1}$, which is $2T_i$ or D_i^{min}.

[6] By 'serving' an uplink real-time mobile, we imply that this mobile's channel is probed, and the mobile is polled if the channel is predicted to be good; otherwise, the polling is deferred.

If the BS sends a channel probing control packet to a mobile when this mobile does not have any pending real-time packet, the mobile will specify that it does not have any real-time packet via the returning control packet so that the BS need not poll the mobile. On the other hand, if a mobile completes the transmission of all of its pending real-time packets after being polled N_i times (where $N_i < M_i$), the BS will be informed of this via the N_i-th packet so that the BS need not poll the mobile again. When polling for connection i in queue **R** is deferred after polling it N_i times, where $0 \le N_i < M_i$, the polling request is queued up in **D** with $P_i = M_i - N_i$, which is the number of polls to be taken, as well as the connection ID. When this request is to be served later, mobile i will be polled up to P_i times consecutively. The credit counter is updated with polling as follows:

- $CC := CC + 2T_{ms}N_i + (3T_{ms} + T_s) + (M_i - N_i - 1)(5T_{ms} + T_s)$, if polling for connection i in queue **R** is deferred after it is polled N_i times;
- $CC := CC - 2T_{ms} + M_i(5T_{ms} + T_s)$, if mobile i had no pending packet to transmit;
- $CC := CC + 2T_{ms}N_i + (M_i - N_i)(5T_{ms} + T_s)$, if mobile i completed the transmission of its pending packets after being polled N_i times (where $0 < N_i \le M_i$), i.e., it had N_i pending packets;

The other rules associated with serving a polling request are the same as those for real-time downlink packet transmissions. For each erroneously-transmitted uplink packet, a polling request is queued up in **B**.

4.1.3 Transmission-Request Slot Scheduling

To achieve low latencies for message transmissions or connection setup/hand-off requests, it is necessary to issue transmission-request slots regularly and frequently. When it is desirable to have a transmission-request slot once every T_{req} time units, this slot is issued with the same scheduling used for polling with $T_i = T_{req}$, $M_i = 1$, and $D_i = D_i^{min} = 2T_{req}$. Note that no channel probe is made before issuing a transmission-request slot since there is no specific destination mobile in this case. For the rest of this paper, we "imagine" there exists a virtual uplink connection with $(M_i, T_i, D_i) = (1, T_{req}, 2T_{req})$, which is equivalent to transmission-request slots. However, the scheduling discussed so far is valid only when the system is busy, i.e., there always exist some downlink packets or uplink pollings scheduled continuously. If not, the BS basically has nothing to do. Then, the BS continues to issue transmission-request slots, thus reducing the transmission-request access latency as much as possible. With the above scheduling, during the system busy period, a transmission-request slot will appear, on average, once every T_{req} time units, and the time span between two consecutive transmission-request slots is bounded by $2T_{req}$.

Note that the packet scheduling itself cannot guarantee a bounded delivery delay. A proper admission control should be performed to determine if a set of connections can be guaranteed to have their required delivery-delay bounds using the above scheduling mechanisms.

4.2 Admission Test

We now consider the admission test for a new connection request. If the other communicating party of a mobile is outside of the cell, the entire communication path will be divided into one or two wireless links and a route within the wired network. The admission test should consider end-to-end communication. However, here we consider the admission test for a wireless link only since the transmission over the wired links can be handled by other known schemes like the one in [10]. (Admission tests for the wired-network part have been studied by many other researchers, and depend on the scheduling policy used in the intermediate nodes.) For a new connection to be established, the results from admission tests on all components (on the communication path) should be positive.

As mentioned already, the admission control is applied so that the delay of each packet in connection i for all i is bounded by D_i^{min} as long as mobile i's channel stays in good condition. Let's consider a set of connections $\{C_i = (M_i, T_i, D_i) \mid i = 1, 2, \cdots, N_r\}$. The admission test is composed of two phases, i.e., a bandwidth test phase, and a delay-bound test phase. For the newly-requested connection new, the bandwidth test is to check if the sum of reserved bandwidth for each connection including overheads satisfies the following condition:

$$(T_s + 5T_{ms}) \left(\sum_{i=1}^{N_r} M_i/T_i + M_{new}/T_{new} \right) \leq 1 - \Delta_r, \quad (4)$$

where Δ_r is the portion of the link capacity reserved for real-time packet retransmissions and non-real-time traffic. Note that $(5T_{ms} + T_s)$ instead of $(3T_{ms} + T_s)$ is needed for a packet transmission to reserve additional $2T_{ms}$ for an extra channel probing. The larger Δ_r, the more packets will make their deadlines even under bad channel conditions. The value of Δ_r can be adapted depending on the channel conditions of the existing connections and the observed packet-dropping probability.

If a new connection does not satisfy the first condition, it is rejected while it is considered for the second phase delay-bound test only if it satisfies the first condition. The delay-bound test is set up as follows. First, suppose the elements of set $\{C_i\}$ are arranged in ascending order of D_i', where

$$D_i' = \begin{cases} D_i^{min}, & \text{if downlink,} \\ D_i^{min} - T_i, & \text{if uplink.} \end{cases} \quad (5)$$

In fact, $D_i' = T_i$ for both uplink and downlink packets. Note that when a packet (or polling request) from each of N_r connections arrives at the same time, the packet (or polling request) from a smaller-numbered connection will have higher priority. To establish a real-time connection, the BS needs to check its schedulability with the connection's worst-case delay. We can derive an admission test similar to that in [10]. For downlink connection i, M_i packets are assumed to arrive periodically at the beginning of each virtual packet-arrival period T_i, which represents the worst case in terms of the delivery delay of these packets. Then, the last of these M_i packets should finish the transmission within D_i' as long as the channel continues to be good. If this works correctly, uplink packets are also transmitted by the minimum delay bound as explained in Section 4.1.2.

Now, for the schedulability test, we consider the transmission time of the last of M_i packets of downlink connection i. The worst-case packet-delivery delay occurs when its arrival coincides with all other higher-priority packet arrivals. This "critical" instant is denoted as time $t = 0$. The transmission of the packet can be delayed by all instances of higher-priority packets (and pollings) at $t = 0$ and by the subsequent instances of these packets (and polling requests). Arrival times of these instances within the delay bound D_i' are given as:

$$A_i = \{D_i'\} \bigcup \quad (6)$$
$$\{kT_j \mid j = 1, 2, \cdots, i-1; k = 1, 2, \cdots, \lfloor (D_i'/T_j) \rfloor\}.$$

The total time required for transmitting this packet and others with priority equal to, or higher than, this packet (or polling) is given by:

$$W_i(t) = T_{max\_poll} + \quad (7)$$
$$M_i(5T_{ms} + T_s) + \sum_{j=1}^{i-1} \{M_j(5T_{ms} + T_s)\} \lceil t/T_j \rceil.$$

Note that the first and second terms of Eq. (7) denote, respectively, an upper bound of the time to complete in-progress packet transmission or polling and the time to transmit M_i packets of real-time connection i. The third term denotes the time to transmit all packets with the same or higher priority. Note that $(5T_{ms} + T_s)$ instead of $(3T_{ms} + T_s)$ was again assumed necessary for a packet transmission. In fact, T_{max_poll} is affected by uplink connections since the time to complete an in-progress downlink packet transmission will be just $3T_{ms} + T_s$. An uplink real-time mobile i will be polled up to M_i times consecutively once it is polled:

$$T_{max\_poll} = \max\{2T_s, \max_{i \in \mathbf{S}_u}\{M_i(3T_{ms} + T_s)\}\}, \quad (8)$$

where $\mathbf{S}_u$ is the set of uplink connections' indices. By definition, T_{max_poll} will be at least $2T_s$ even when there is no uplink real-time connection. This follows from the maximum possible non-real-time traffic transmission without interrupting real-time packet transmissions as explained in Section 4.3. A set of connections $\{C_i\}$ is said to be *schedulable* (i.e., all packet minimum delivery-delay constraints will be met as long as the channel conditions continue to be good) if the following condition holds:

$$W_i(t) \leq t, \qquad \text{for some } t \in A_i, \quad (9)$$
$$\text{for all } i = 1, 2, \cdots, N_r,$$

where N_r is the number of connections, and A_i is given as in Eq. (6). From this schedulability condition, we now derive the second phase delay-bound test of a new connection when there already exists a set of connections $\{C_i\}$ in the system. Note that the schedulability condition in Eq. (9) already holds for the set of existing connections $\{C_i\}$.

1. Upon arrival of a new connection request, the BS inserts the newly-requested connection new in the set of connections $\{C_i\}$ according to the ascending order of D'_{new}.
2. If new is a downlink connection, the admission test is to check if the following condition holds:

$$W_i(t) \leq t, \qquad \text{for some } t \in A_i, \qquad \text{for } i = new, new+1, \cdots, N_r. \tag{10}$$

3. If new is an uplink connection, T_{max_poll} is re-calculated using Eq. (8).
   - If T_{max_poll} hasn't changed, the admission test is given by Eq. (10);
   - If T_{max_poll} has changed, the admission test is given by Eq. (9).

Note that if T_{max_poll} hasn't changed, the performance guarantees of connections with smaller D'_i's are not affected by the admission of connection new.

4.3 Non-Real-Time Traffic Scheduling

The transmission of non-real-time traffic is based on a modified round-robin scheduling policy with a set of two prioritized queues. Using the round-robin scheduling, it is possible to give each mobile fair channel access as long as each mobile's channel condition continues to be good. When there are M_o mobiles in the cell, the BS is equipped with the following two service queues (in that order of priority): (1) a class II-A round-robin queue (with M_o entries), called queue **RR.A**, and (2) a class II-B round-robin queue (with M_o entries), called queue **RR.B**. The i-th entry of a round-robin queue, corresponding to mobile i, is composed of two parallel queues for the downlink and uplink, respectively. The uplink queue has a number of uplink packets requested for transmission permissions, which is the sum of the number of packets in transmission permission-requested messages. On the other hand, the downlink queue buffers the downlink packets in a FIFO manner which had arrived at the BS. Figure 4 shows a round-robin queue for M_o mobiles in the cell.

We first describe how the round-robin queues work assuming that there is no channel error, then address how to handle channel errors. For the i-th entry's turn, up to two packets for mobile i are served. To reduce the ACK/NAK-transmission and polling overheads, it is desirable to sequentially serve one downlink and one uplink packets. In this case, a downlink packet destined for mobile i is transmitted first with up to two different pieces of piggybacked information: (1) ACK/NAK for a previously-transmitted uplink packet if the previous packet was uplink; and (2) the transmission permission for mobile i to transmit an uplink packet (i.e., a polling to mobile i). Then, mobile i will transmit an uplink packet with ACK/NAK for the

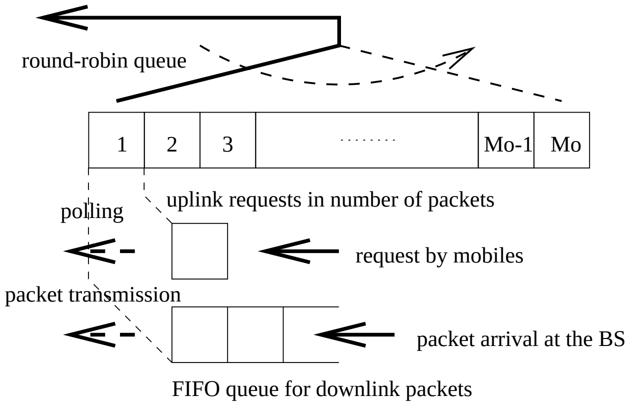

Figure 4: A round-robin queue for non-real-time communications.

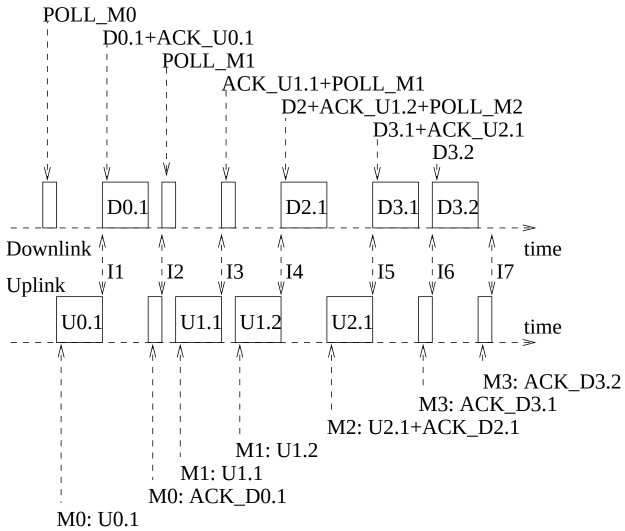

Figure 5: An example of non-real-time transmissions.

previously-received downlink packet. Note that through this transmission order, two mini-slots were saved, i.e., for one downlink control packet to poll mobile i, and for one uplink control packet for the ACK from mobile i. If the i-th entry does not have both downlink and uplink requests, up to two downlink or uplink packets for mobile i can be transmitted consecutively using two mini-slots for polling or ACK transmissions. An uplink packet and a downlink packet can also be transmitted sequentially if a downlink packet arrives during the transmission of an uplink packet.

An example of this non-real-time scheduling is shown in Figure 5, where a square and a small rectangle represent a packet transmission slot and a mini-slot, respectively. The notation used is: $U_{i.j}$ is the j-th uplink packet transmission of mobile i; $D_{i.j}$ is the j-th downlink packet transmission to mobile i; $POLL\_M_i$ is a polling to mobile i; $ACK\_\{x\}$ is the ACK/NAK for the previous packet transmission x; M_i represents the transmission from mobile i, and '+' represents piggybacking of the following information in the same (control) packet. Four different cases are shown in chronological order: (1) an uplink packet from, and a downlink packet sequentially to, mobile 0; (2) two uplink packets from mobile 1; (3) a downlink packet to, and an uplink packet sequentially from, mobile 2; and (4) two downlink packets to mobile 1. Note that two control packets are used to transmit two packets from/to mobiles 0, 1, and 3. However, no control packet is used for mobile 2 to receive one downlink and transmit one uplink packet.

Note that real-time traffic or higher-priority non-real-time traffic (such as class II-A traffic if the traffic considered is class II-B) can take over the channel according to the priority rule when the wireless link becomes free after (1) the completion of the ACK transmission for a non-real-time downlink packet, or (2) the completion of a non-real-time uplink packet transmission. For the latter case, the BS needs to piggyback the ACK/NAK within the following (control) downlink packet. So, the BS can switch to the real-time traffic service or a higher-priority non-real-time traffic service at the marked (with 'I_n') moments in Figure 5. The worst-case delay before switching to the real-time traffic service is shown in the figure, i.e., from I_4 to I_5. To interrupt at I_5, there is a delay of up to $2T_s$ time units. This is why T_{max_poll} should be at least $2T_s$ in Eq. (8). Even though it was not shown in the figure, the request for an uplink message transmission can also be piggybacked in an uplink packet.

Now, let's assume that there was a NAK associated with a packet transmission from/to mobile i. Then, the entry for this mobile in the particular round-robin queue is marked "back-logged" indicating that the mobile's channel condition is bad. For this entry's next service turn, the mobile's channel is probed first. If the channel is predicted to be good, the mobile's packets are served, and the entry is marked "active." Otherwise, the entry stays back-logged without the mobile being served at all. For the turn of the i-th entry, which is back-logged, the i-th *non-real-time compensation counter* $NCC.(*)[i]$ is increased by the number of packets to be served if the entry was active, where $(*)$ is A (B) if the queue is for class II-A (class II-B). This counter is increased when an entry becomes back-logged as well. For example, if uplink class II-A packet $U_{0.1}$ in Figure 5 is erroneously transmitted, the 0-th entry of queue **RR.A** is marked back-logged, and $NCC.A[0]$ is increased by two since both $U_{0.1}$ and $D_{0.1}$ could be transmitted if the channel was good. For the turn of the i-th entry which becomes active, as many as $NCC.(*)[i]+2$ packets of mobile i are served continually. Then, $NCC.(*)[i]$ is reset to zero. In this way, we can achieve long-term fairness among mobiles. Note also that before serving the first entry in a round-robin queue, if all the entries are marked back-logged, the round-robin queue service flag $RR^{(*)}_{flag}$ is set to zero, where $(*)$ is again A (B) if the queue is for class II-A (class II-B). $RR^{(*)}_{flag}$'s are set to 1 whenever a packet is transmitted over the channel. A round-robin queue can be served only when its flag is 1. This way we can reduce the overhead associated with the channel probing for non-real-time traffic while being able to serve class II-B traffic when all mobiles with class II-A traffic are experiencing bad channel conditions. Note that the channel is probed for back-logged connections only in order to reduce unnecessary bandwidth waste for channel probing.

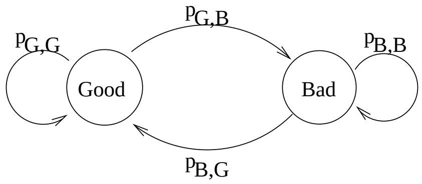

Figure 7: Two-state Markov chain modeled channel.

An ACK/NAK can be lost due to errors. Because the feedback should be received right after the packet transmission, the sender assumes, in case of ACK/NAK loss, that the packet transmission was not successful and retransmits the previous packet later. The receiver, in turn, knows that the previous ACK/NAK was lost when it receives a duplicate packet. If the receiver is the BS, the BS needs to update a round-robin queue's entry by increasing the number of uplink packet requests by one since it received the same packet again. Finally, recall that the above-mentioned scheduling is applied whenever (1) all the real-time queues **R**, **D**, and **B** are empty, or (2) queue **R** is empty, $D_{flag}=0$, $D_{index}=1$, $B_{flag}=0$, and $B_{index}=1$. Figure 6 provides a pseudo-code to determine which queue to be served whenever the wireless link becomes free.

5 Experimental Results and Discussion

This section evaluates and discusses the performance of the proposed scheme. We first state the assumptions and the system specification used in the evaluation study.

5.1 Assumptions and System Specification

We assume that there are N_u $(=10)$ mobiles in the cell, and the duration T_{ms} of a mini-slot is used as a basic *time unit*. Only one request can be transmitted in a transmission-request slot or piggybacked in an uplink packet by a mobile. The wireless channel is modeled as follows:

C1. The channel between the BS and a mobile is modelled by a Markov chain as shown in Figure 7. The channels of two different mobiles are assumed to be independent.

C2. A transition between the good and bad states can happen at mini-slot boundaries with the transition probability $p_{G,B}$ $(p_{B,G})$ from the good to bad states (from bad to good states).

C3. For a given average time duration T_G (T_B) at the good (bad) state, $p_{G,B}=1/T_G$ $(p_{B,G}=1/T_B)$. We assume that $T_G=2000$, and $T_B=100$.

C4. Each request could be received in error (1) when it is corrupted due to channel errors, or (2) when it collides with another request within the same request mini-slot.

C5. If a packet transmission duration overlaps with a time in the bad state, the packet is received in error. The

whenever (the link becomes free) {
  if ($C \geq 3T_{ms} + T_s$ and ($D_{index} > 1$ or ($D_{index} = 1$ and $D_{flag} = 1$)))
    serve queue **D**;
  else if ($C \geq 3T_{ms} + T_s$ and ($B_{index} > 1$ or ($B_{index} = 1$ and $B_{flag} = 1$)))
    serve queue **B**;
  else if (queue **R** is non-empty)
    serve queue **R**;
  else if ($D_{index} > 1$ or ($D_{index} = 1$ and $D_{flag} = 1$))
    serve queue **D**;
  else if ($B_{index} > 1$ or ($B_{index} = 1$ and $B_{flag} = 1$))
    serve queue **B**;
  else if (queue **RR.A** is non-empty and $RR^{A}_{flag} = 1$)
    serve queue **RR.A**;
  else if (queue **RR.B** is non-empty and $RR^{B}_{flag} = 1$)
    serve queue **RR.B**;
  else issue a transmission request slot;
}

Figure 6: A pseudo-code to determine which queue to serve when the wireless link becomes free.

packet is received correctly otherwise.

C6. Every packet received in error is assumed to be detected by the channel decoder.

The real-time traffic of interest is specified as follows:

R1. Four types of real-time connections are considered with the connection specification (M_i, T_i, D_i): (1) type-1 downlink $(1, 200, 300)$; (2) type-1 uplink $(1, 200, 500)$; (3) type-2 downlink $(1, 500, 600)$; and (4) type-2 uplink $(1, 500, 1100)$. Note that type-1 (or 2) downlink and uplink connections have the same M_i and T_i.

R2. Real-time connections arrive according to a Poisson process with rate λ_{rt}. A half of them are hand-off connections which arrive from adjacent cells.

R3. The lifetime of each connection i is distributed geometrically with mean T_{life} $(= 50T_i)$.

Finally, the assumptions on non-real-time traffic are as follows:

N1. Downlink (uplink) messages arrive according to a Poisson process with rate λ_d (λ_u) (messages/mini-slot). The overall message-arrival rate λ_{nrt} is defined to be $\lambda_{nrt} = \lambda_d + \lambda_u$.

N2. The downlink class II-A (II-B) messages, destined for each of N_u mobiles, arrive independently at the BS with rate $0.9\lambda_d/N_u$ $(0.1\lambda_d/N_u)$.

N3. The uplink class II-A (II-B) messages arrive independently with rate $0.9\lambda_u/N_u$ $(0.1\lambda_u/N_u)$ at each of N_u mobiles.

N4. The number of packets in a class II-A (II-B) message is geometrically-distributed with average 2 (18).

A frequently-used parameter, the *offered load*, is defined as the overall expected number of the packets to be transmitted within a slot. For the overall non-real-time message-arrival rate λ_{nrt} (messages/mini-slot), the offered load from non-real-time traffic, L_{nrt}, is given by

$$\begin{aligned} L_{nrt} &= (0.9 \cdot 2 + 0.1 \cdot 18) K \lambda_{nrt}, \\ &= 3.6 K \lambda_{nrt}, \quad \text{(packets/slot)}, \qquad (11) \end{aligned}$$

of which a half is for class II-A and the other half is for class II-B. On the other hand, the offered load from real-time traffic, L_{rt}, is given by

$$\begin{aligned} L_{rt} &= \lambda_{rt} T_{life} E\left\{ \frac{M_i T_s}{T_i} \right\}, \\ &= \lambda_{rt} T_{life} \sum_{j=1}^{2} \alpha_j \frac{\hat{M}_j T_s}{\hat{T}_j}, \quad \text{(packets/slot)}, \qquad (12) \end{aligned}$$

where $E\{\cdot\}$ represents the averaging operation, α_j is the fraction of type-j (both downlink and uplink) connections out of the total real-time connection arrivals, $\hat{M}_j$ is the M_i of type-j connection i, and $\hat{T}_j$ is the T_i of type-j connection i.

5.2 Interworking of Real-Time Connections with Non-Real-Time Traffic

Throughout this subsection, the following parameters are used: $K = 20$ and $T_{req} = 200$, so there is 10 $(= K/2)$ request mini-slots within a transmission-request slot. Three request mini-slots are reserved for hand-off requests, i.e., $K_{ho} = 3$. We first consider the admission region of type-1 and 2 connections depending on the reserved link capacity portion Δ_r for real-time packet retransmissions and non-real-time traffic. Using Eqs. (4) and (9), each pair of the numbers of type-1 and 2 connections is tested; a pair (n_1, n_2) represents a set of n_1 type-1 connections and n_2 type-2 connections. Figure 8 shows the curves representing the admission regions for three different Δ_r values. For each line, a set of connections for each pair under or on a curve can be admitted into the system. For example, 6 type-1 uplink connections and 2 type-2 uplink connections can co-exist in the system at a time when $\Delta_r = 0 - 0.1$. We observe that the larger Δ_r, the smaller the admission

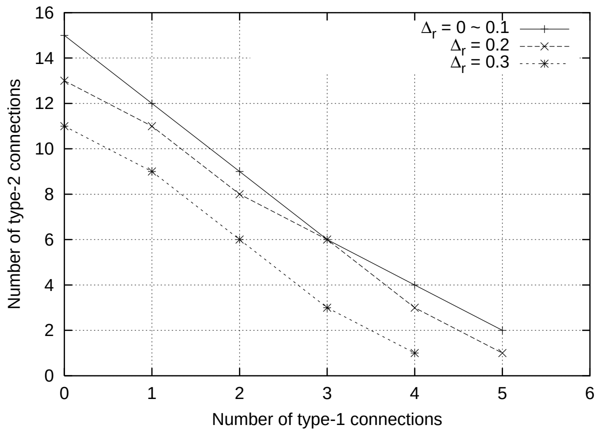

Figure 8: The admission region for real-time connections: $K = 20$ and $T_{req} = 200$.

Real-time arrival rate (λ_{rt})	0.0001	0.0005	0.001
Real-time offered load (L_{rt})	0.10	0.50	1.00
Real-time throughput	0.099	0.359	0.439
Conn.-blocking probability	0.0013	0.273	0.563
Average packet delay	164.55	181.76	185.09
Maximum packet delay	500	500	500
Packet-dropping probability	0.0031	0.0040	0.0042

Table 3: Performance of real-time traffic when non-real-time uplink traffic arrives with $\lambda_u = 0.0005$: $K = 20$ and $T_{req} = 200$.

region as it should be. The admission regions for $\Delta_r = 0$ and $\Delta_r = 0.1$ are the same is due to the effect of T_{max_poll} in Eq. (7) for the delay-bound test. Note that the admission regions of both downlink and uplink connections are the same since both downlink and uplink connections of the same type have the same M_i and T_i.

Now, let's consider the performance of real-time traffic. For further study, we assume that $\Delta_r = 0$, and only type-1 uplink connections were generated, where each of them is a constant bit rate (CBR) connection, i.e., M_i packets are generated every T_i time units. We learned from the previous admission-region figure that up to 5 type-1 uplink connections can be admitted at a time. The reason why we consider only uplink connections is to examine the access latency of uplink connection requests, which is a very important performance measure. In the simulation, non-real-time uplink traffic arrives with rate $\lambda_u = 0.001$. Note that the performance of real-time traffic is almost unaffected by non-real-time traffic because real-time traffic is given priority over non-real-time traffic. (Some small effect can appear due to non-preemptive transmission of non-real-time traffic, and packet deferments of real-time traffic.) Table 3 shows the real-time offered load L_{rt}, real-time traffic throughput, connection-blocking probability, average and maximum packet delays, and packet-dropping probability (i.e., packet drops due to the deadline expirations) for three different arrival rates λ_{rt} of real-time connections. The blocking probability represents both new-

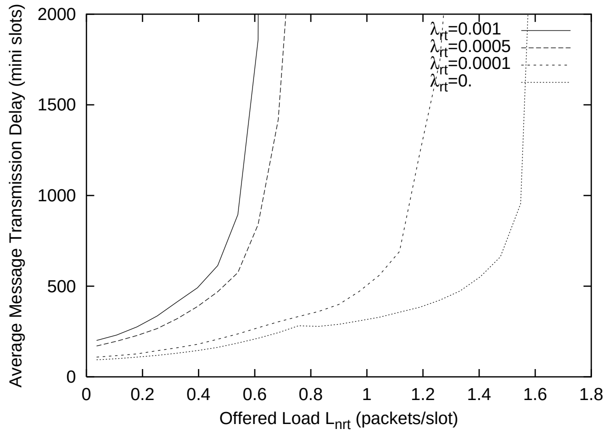

Figure 9: Average uplink class II-A message-transmission delay vs. non-real-time offered load: for different real-time connection-arrival rates λ_{rt}, $K = 20$ and $T_{req} = 200$.

connection blocks and hand-off connection drops since a half of the entire real-time connection arrivals are assumed to be hand-offs in our simulation. As the real-time connection arrival rate increases, more connections are blocked. We also observed that the larger λ_{rt}, the larger the average packet delay and the packet-dropping probability, since the more real-time connections exist, the longer a scheduled polling is likely to be delayed. However, the maximum delay is observed to be always 500, which is the delay bound D_i. Note that this delay is bounded at the cost of some packet drops. We can easily expect that the larger Δ_r, the larger the connection-blocking probability, and the smaller packet-dropping probability.

Figure 9 shows the average uplink class II-A message-transmission delay for four different real-time traffic-arrival rates. We observed that the larger λ_{rt}, the larger the delay is for a given non-real-time offered load L_{nrt}, because real-time traffic is given priority over non-real-time traffic. Note that a half of non-real-time offered load is from class II-B. This is the reason why the delay of class II-A could not be infinite even for the offered load is more than one. Even though it is not shown, the delay of class II-B is surely infinite for the offered load greater than one. The next subsection will elaborate on the performance of non-real-time communication service.

Next, we consider the request-access latency, which is defined as the time span from the arrival of a request to its successful reception by the BS. Recall that a request access is made through the transmission request slot or piggybacking for (1) non-real-time (class II-A and II-B) message-transmission requests, (2) a new real-time connection request, and (3) a hand-off request. The request-access latencies for four different cases are plotted in Figure 10 when $\lambda_{rt} = 0.001$, and only uplink non-real-time traffic exists. We observe that the request-access latency for real-time and hand-off connections are reasonably low (less than 150) throughout the whole non-real-time offered load examined. Especially, the request-access latency for hand-off connections are saturated by 120 beginning $L_{nrt} \approx 0.4$.

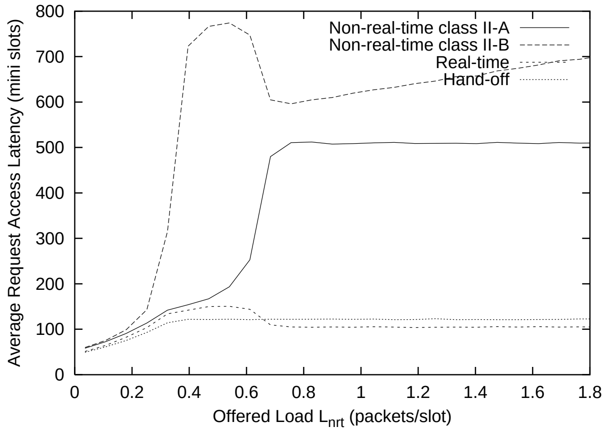

Figure 10: Average request-access latency vs. non-real-time offered load: $\lambda_{rt} = 0.001$, $K = 20$, and $T_{req} = 200$.

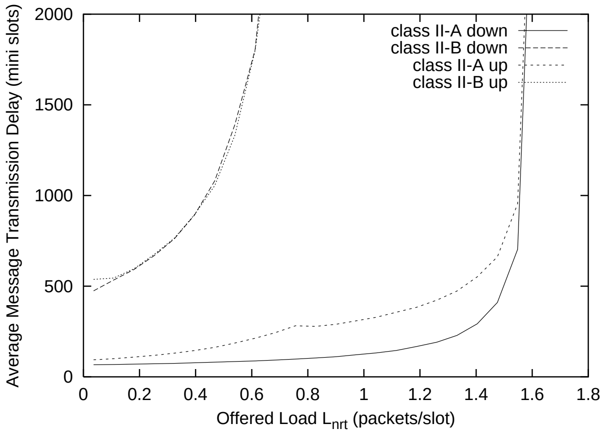

Figure 11: Average message-transmission delay vs. non-real-time offered load when downlink or uplink traffic exists: $K = 20$ and $T_{req} = 200$.

Note that the request-access latency of hand-off connections should be kept low because otherwise, the hand-off will not be smooth. To give a detailed account, we divide the entire non-real-time offered loads examined into three regions: (1) lightly-loaded region: $L_{nrt} = 0.0$ to 0.5; (2) moderately-loaded region: $L_{nrt} = 0.5$ to 0.8; and (3) heavily-loaded region: $L_{nrt} = 0.8$ to 1.8.

First, we consider the request-access latency of non-real-time traffic. In the lightly-loaded region, both the latencies of class II-A and II-B traffic increase monotonically. Class II-B traffic's access latency is larger since class II-A's requests are given priority. In the moderately-loaded region, class II-B traffic's access latency starts to decrease due to the effect of piggybacked requests. Class II-A traffic's access latency, on the other hand, continues to increase further due to the effect of piggybacked requests. That is, in this region, as the offered load increases, more requests will be made via piggybacked requests. The piggybacked request-access delay will increase since the time span between two consecutive services of a mobile will increase due to the increasing number of mobiles with pending requests in round-robin queues. In the heavily-loaded region, class II-A request-access latency is saturated at around 500 while class II-B request-access latency increases slowly beginning at around 600. In this region, non-real-time requests are likely to be made only though the piggybacked requests since the non-real-time round-robin service queues in the BS are unlikely to be empty. In this case, a mobile is supposed to have a chance to transmit two uplink packets once every 440 time units when there is no real-time traffic since there are 10 mobiles and 2 packets (plus 2 control packets) can be transmitted at each round-robin turn. So, the request-access latency would be about 220 when there is no real-time traffic. Due to real-time connections, these chances are delayed more, depending on the real-time arrival rate. It is found to be about 500 for class II-A traffic when $\lambda_{rt} = 0.001$. The request access of class II-B traffic is likely to be delayed more since only one request for a traffic class can be piggybacked in a packet.

The access latency of real-time and hand-off connection requests is found to be lower than that of non-real-time traffic over the entire offered load region. This is because real-time connection requests are given priority over non-real-time transmission requests, and K_{ho} (which is 3 in our simulations) request mini-slots are reserved for hand-off requests. For the very lightly-loaded region (say $L_{nrt} \leq 0.2$), the latencies for all types of requests are about the same since almost every request is made through a transmission-request slot and is mostly successful because collisions rarely occur. As the offered load increases, all the latencies increase while the hand-off request latency is saturated. In the moderately-loaded region, the access latency for real-time connections starts to decrease. This is again because some of transmission requests are piggybacked. Then, in the heavily-loaded region, real-time and hand-off request-access latencies are saturated at around 100 and 120, respectively. The real-time request-access latency is smaller than the hand-off request-access latency since a hand-off request can only be made through the transmission-request slot while a new connection request can be made through either a transmission-request slot or piggybacking. They are about 100 since the period of issuing the transmission-request slot $T_{req} = 200$ while they are about the same since the effect of the piggybacked request is negligible. The exact request-access latencies will depend on the real-time connection-arrival rate λ_{rt}, but the overall tendency observed remains the same irrespective of the rate.

5.3 Non-Real-Time Communication Service

Here we examine in detail the performance of non-real-time communication service. To understand the non-real-time communication service better, we generated only non-real-time traffic for the simulations, i.e., $\lambda_{rt} = 0.$. Figure 11 shows the average message-transmission delays for both class II-A and II-B for two extreme cases: one with downlink traffic arrivals only (marked with 'down') and the other with uplink traffic arrivals only (marked with 'up'). We ob-

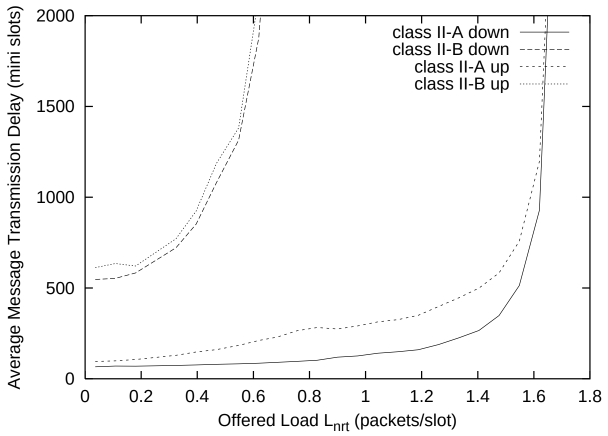

Figure 12: Average message-transmission delay vs. non-real-time offered load when downlink and uplink traffic arrive with the same rate, i.e., $\lambda_d = \lambda_u$: $K = 20$ and $T_{req} = 200$.

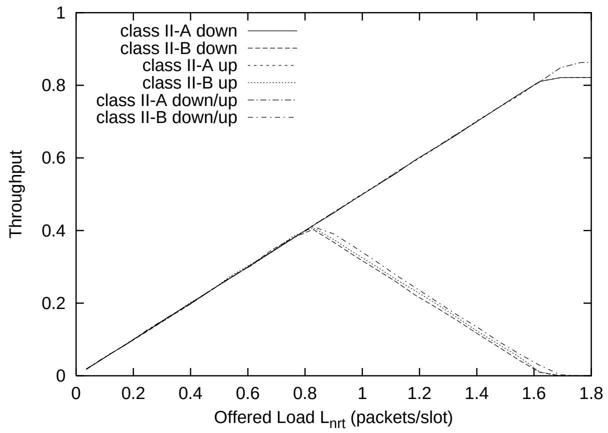

Figure 13: Throughput vs. non-real-time offered load: $K = 20$ and $T_{req} = 200$.

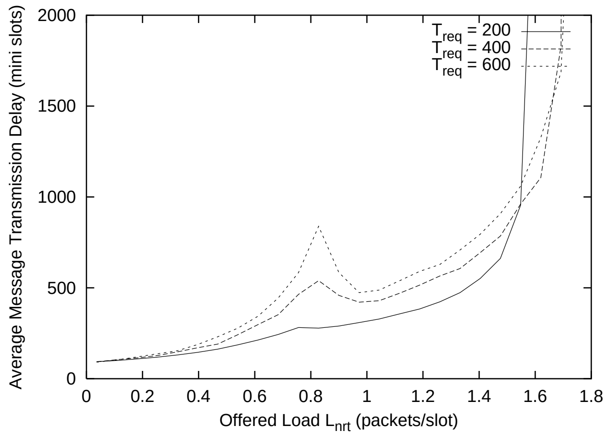

Figure 14: Average uplink class II-A message transmission delay vs. non-real-time offered load: with different values of T_{req}: $K = 20$.

serve that the delays with downlink traffic arrivals are a bit smaller than those with uplink traffic arrivals. This indicates the effect of the uplink transmission-request access latency. For both cases, class II-B's delay goes to infinity at around $L_{nrt} = 0.7$ while class II-A's delay goes to infinity at around $L_{nrt} = 1.6$. This means that class II-B traffic is rarely transmitted for $L > 0.7$. This is reasonable since a half of the offered load is from class II-B, and the maximum achievable throughput is less than one. The fact that class II-A's delay goes to infinity at around $L_{nrt} = 1.6$ means that the maximum achievable throughput is about 0.8. Figure 12 shows the average message-transmission delays for both class II-A and II-B when a half of arrived messages are downlink and the other half are uplink, i.e., $\lambda_d = \lambda_u$. The general tendency is the same as in the previous figure, but the maximum achievable throughput is found to increase since the delay of class II-A is still about 2000 (i.e., within the figure) even for $L_{nrt} = 1.65$. This shows the effect of sequential transmission of a downlink packet and an uplink packet without any control packet at a mobile's turn in the BS round-robin queue.

These facts can be seen clearer in Figure 13. First, beginning at $L_{nrt} = 0.8$, the throughput of class II-B starts to decrease while that of class II-A continues to increase. When the maximum achievable throughput is reached, the throughput is saturated there, i.e., about 0.8 for both only downlink (marked with 'down') and only uplink (marked with 'up') arrival cases. Note that the throughputs for both downlink and uplink arrival case (marked with 'down/up') for a given offered load are larger than, or equal to, those for the other two extreme cases. The maximum achievable throughput for the latter case is found to be about 0.86. The reason why the maximum achievable throughput is less than one is due to the overhead of transmission-request slots, polling mini-slots, ACK/NAK-transmitting mini-slots, and channel probing processes.

So far, we used the value of $T_{req} = 200$, i.e., the average interval between two transmission-request slots is about 200. Figure 14 shows the average transmission delays of uplink class II-A messages for three different T_{req} values. Note that the larger T_{req}, the larger the maximum achievable throughput could be. The delay is found to increase abruptly at around $L_{nrt} = 0.8$ for a large T_{req}. This is because the larger T_{req}, the more chance of a mobile having a message to request upon appearance of a transmission-request slot. So, the larger T_{req}, the more likely a request is to collide with others, and the larger the delay eventually. This effect is lessened dramatically after $L_{nrt} = 0.8$ due to the increases of piggybacked requests. Even though it is not shown here, the access latency of real-time connection requests also has a similar abrupt increase at some regions of offered load with large T_{req}. This observation implies that too large a T_{req} is not desirable.

Figure 15 shows the effect of K values when $T_{req} = 600$. As K increases, the peak delay at $L_{nrt} = 0.8$ is found, on average, to decrease while the delay increases. A large K means a large packet size and an increased number of request mini-slots in a transmission-request slot. The decreased peak delay is due to the second fact, and the in-

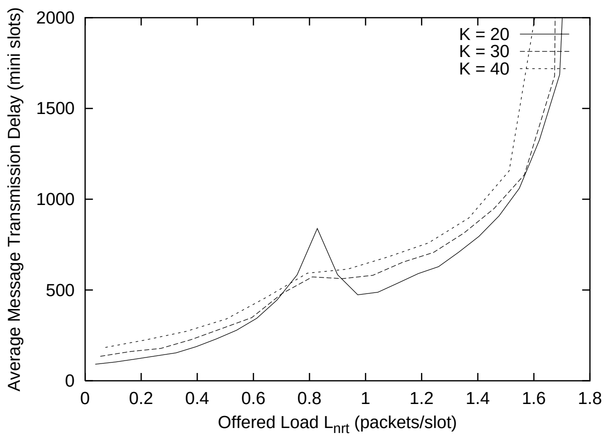

Figure 15: Average uplink class II-A message-transmission delay vs. non-real-time offered load: for $T_{req} = 600$ and with different values of K.

creased average delay is due to the first fact. Note that too large a K is not good even though the peak delay is reduced because a large fixed-size packet is likely to be transmitted with some empty portion, thus lowering the system utilization. From the above two figures, one can conclude that the choice of K and T_{req} is closely related to each other, and hence, they should be considered together so that there is no undesirable peak in the average delay.

6 Related Work

Recently, there have been significant research efforts to support QoS guarantees in wireless networks. A wireless MAC protocol in [11] can be considered as a polling scheme using stop-and-go queueing for real-time traffic, and FIFO queueing for non-real-time traffic. However, this scheme can support only a finite number of delay bounds due to the inherent limitation of stop-and-go queueing. The authors of [12] dealt with the scheduling and admission control for a TDMA system to support variable bit rate (VBR) connections with different packet-dropping probabilities. However, it deals with the uplink only, and connections are allowed to have only two different delay bounds. Both of the schemes do not provide proper means to handle wireless channel errors. Remote-Queueing Multiple Access (RQMA) in [13] also addresses how to provide QoS for heterogeneous traffic using existing packet-scheduling algorithms designed for wired networks. This scheme also handles channel errors via a FEC/ARQ hybrid like ours. However, it did not consider how to handle location-dependent errors.

An emerging WLAN MAC standard, IEEE 802.11, provides both real-time and non-real-time communication services using polling and CSMA/CA (Carrier Sense Multiple Access/Collision Avoidance) [14]. The CSMA-based protocol is not suitable for a network environment with the star topology considered in this paper, since carrier sensing is not possible in such an environment. The IEEE 802.11 also does not specify the polling order, admission control, and how to handle channel errors which are all essential for real-time communication.

QoS-supporting wireless systems have recently been studied in the context of wireless ATM. The distributed queueing request update multiple access (DQRUMA) is used for BAHAMA as a MAC protocol [15, 16]. DQRUMA is a demand-assignment access scheme, which has some limited support for real-time traffic due to its contention-based access of transmission requests. In fact, the polling for uplink non-real-time transmission in our system can be classified as this demand-assignment access scheme because transmission requests are made on a per-message basis, and the permission is given via polling. The MAC protocol for a wireless ATM in [17] uses a polling scheme. However, it deals with uplink accesses only. The Seamless Wireless ATM Network (SWAN) [18] also uses a polling scheme as a MAC protocol, but such details as polling scheduling and admission control were left unaddressed. WATMnet [19] uses a dynamic TDD/TDMA as the MAC to support heterogeneous traffic. The MAC protocol in WAND (MASCARA) [20] is another MAC protocol for a wireless ATM system (named WAND) using reservation and contention-based TDD/TDMA techniques. The authors of [21] proposed a wireless MAC protocol for wireless ATM, in which the access and scheduling procedures are distributed. Basically, none of the above wireless ATM schemes deals with location-dependent and bursty errors explicitly.

There have also been remarkable research efforts for QoS-provisioning CDMA systems, e.g., [8, 22]. While the CDMA technology is expected to dominate for the third-generation (3G) wireless systems, it is not suitable for typical WLAN environments since the data rate of an individual user can hardly be commensurate with the system bandwidth due to the spreading factor. The authors of [23] provide a good survey of MAC protocols covering both TDMA and CDMA-based schemes for wireless multimedia networks while a good survey of wideband local access schemes comparing both wireless LAN and wireless ATM technologies can be found in [24].

Recently, there have been significant research efforts on packet scheduling to handle location-dependent and bursty errors based on channel-prediction mechanisms and packet-transmission deferments [25–30]. Most of the work has focused on how to provide (long-term) fair access to a wireless link in the presence of location-dependent errors [26–30]. Moreover, some of these have limitations, such as the issue of expediting the retransmission of packets was not addressed [26–28]. We adopted two different types of packet-scheduling policies to handle both real-time and non-real-time traffic, in which EDF scheduling is used for real-time traffic to bound packet-delivery delays at the cost of some possible packet drops, and round-robin scheduling is used for non-real-time traffic to achieve both error-free transmissions and long-term fairness among different mobiles.

Our scheduling for real-time communication guarantees that each connection receives the service it was promised as long as its channel continues to be good while a connection with deferred transmissions is served in a best-effort manner. In that sense, our scheduling is different from those in [26, 28, 30], which compensate the bandwidth of connections which experienced deferments by reducing the bandwidth of other connections which received more bandwidth than promised before. Our non-real-time scheduling, however, is based on these scheduling algorithms. The scheduling policy in [29] does not have an explicit concept of compensation like our real-time scheduling, and it does not address how to bound packet delay bounds either. Finally, our real-time scheduling is deadline-driven, so it can bound delays without addressing the fair access issue explicitly, but fairness is also a main concern in our non-real-time scheduling.

We adopted the polling scheme as in [11, 17, 18] for both real-time and non-real-time traffic with different polling and scheduling strategies. The D-TDD technique in [11, 13, 18–20] was also used for more efficient and flexible utilization of a frequency channel. We showed how to support the distinct QoS requirements of heterogeneous traffic (i.e., bounded delivery-delay for real-time and virtually loss-free transmission for non-real-time traffic), including the transmission/polling scheduling and admission control for real-time connections, in the dynamic environment with location-dependent, time-varying, and bursty errors.

7 Conclusion

This paper considered how to support both real-time and non-real-time communication services in a WLAN. We proposed a MAC protocol based on polling mobiles for their uplink accesses with slotted ALOHA-based and piggybacked transmission requests. A variation of EDF scheduling is used for real-time traffic to bound delivery-delays at the cost of some packet losses depending on the channel conditions while a variation of round-robin scheduling is used for non-real-time traffic to achieve error-free transmissions and long-term fairness among mobiles. Also addressed are the problems of (1) multiplexing downlink and uplink traffic for D-TDD transmission, and (2) handling location-dependent, time-varying, and bursty channel errors. An admission test for each new real-time connection request has been established, determining if the requested packet delivery-delay bound can be guaranteed without violating the existing guarantees.

An extensive evaluation study has shown the proposed scheme to work well and meet the design goals. We first considered the admission regions of real-time connections for four types of connections, and the average performance of real-time traffic. We also examined the inter-working of real-time and non-real-time traffic while emphasizing the access delays of non-real-time messages, real-time connection requests, and hand-off requests. Finally, the performance of non-real-time communication service was evaluated in detail.

References

[1] K. Pahlavan and A. H. Levesque, *Wireless Information Networks*. New York, NY: Wiley-Interscience, 1995.

[2] H. Ahmadi, A. Krishna, and R. O. LaMaire, "Design issues in wireless LANs," *Journal of High-Speed Networks*, vol. 5, no. 1, pp. 87–104, 1996.

[3] S. Choi and K. G. Shin, "Centralized wireless MAC protocols using slotted ALOHA and dynamic TDD transmission," *Performance Evaluation*, vol. 27 & 28, pp. 331–346, Oct. 1996.

[4] D. Duchamp and N. Reynolds, "Measured performance of a wirelss LAN," in *Proc. 17th Conference on Local Computer Networks*, pp. 494–499, Sept. 1992.

[5] F. Swarts and H. Ferreira, "Markov characterization of digital fading mobile VHF channels," *IEEE Trans. on Vehicular Technology*, pp. 977–985, Nov. 1994.

[6] H. Wang and N. Moayeri, "Finite state morkov channel – a useful model for radio communication channels," *IEEE Trans. on Vehicular Technology*, pp. 163–171, Feb. 1995.

[7] D. Eckhardt and P. Steenkiste, "Measurement and analysis of the error characteristics of an in-building wireless network," in *Proc. ACM SIGCOMM'96*, pp. 243–254, 1996.

[8] S. Choi and K. G. Shin, "An uplink CDMA system architecture with diverse QoS guarantees for heterogeneous traffic," *IEEE/ACM Trans. on Networking*, vol. 7, Oct. 1999.

[9] C. L. Liu and J. W. Layland, "Scheduling algorithms for multiprogramming in a hard-real-time environment," *Journal of ACM*, vol. 20, pp. 46–61, Jan. 1973.

[10] D. D. Kandlur, K. G. Shin, and D. Ferrari, "Real-time communication in multi-hop networks," *IEEE Trans. on Parallel and Distributed Systems*, vol. 5, pp. 1044–1056, Oct. 1994.

[11] S. Choi and K. G. Shin, "A cellular wireless local area network with QoS guarantees for heterogeneous traffic," *ACM Mobile Networks and Applications (MONET)*, vol. 3, no. 1, pp. 89–100, 1998.

[12] J. M. Capone and I. Stavrakakis, "Delivering diverse delay/dropping QoS requirements in a TDMA environment," in *Proc. ACM/IEEE MobiCom'97*, pp. 110–119, 1997.

[13] N. R. Figueira and J. Pasquale, "Remote-queueing multiple access (RQMA): providing quality of service for wireless communications," in *Proc. IEEE INFOCOM'98*, 1998.

[14] B. P. Crow, I. Widjaja, J. G. Kim, and P. T. Sakai, "IEEE 802.11 wireless local area networks," *IEEE Communications Magazine*, Sept. 1997.

[15] K. Y. Eng *et al.*, "BAHAMA: A broadband ad-hoc wireless ATM local-area network," in *Proc. IEEE ICC'95*, pp. 1216–1223, 1995.

[16] M. J. Karol, Z. Liu, and K. Y. Eng, "Distributed-queueing request update multiple access (DQRUMA) for wireless packet (ATM) networks," in *Proc. IEEE*

ICC'95, pp. 1224–1231, 1995.

[17] C.-S. Chang *et al.*, "Guaranteed quality-of-service wireless access to ATM networks," *IEEE Journal on Selected Areas in Communications*, pp. 106–118, Jan. 1997.

[18] P. Agrawal *et al.*, "SWAN: A mobile multimedia wireless network," *IEEE Personal Communications*, pp. 18–33, Apr. 1996.

[19] D. Raychaudhuri *et al.*, "WATMnet: A prototype wireless ATM system for multimedia personal communication," *IEEE Journal on Selected Areas in Communications*, pp. 83–95, Jan. 1997.

[20] N. Passas, S. Paskalis, D. Vali, and L. Merakos, "Quality-of-service-oriented medium acess control for wireless ATM networks," *IEEE Communications Magazine*, pp. 42–50, Nov. 1997.

[21] M. Listanti, F. Mascitelli, and A. Mobilia, "D^{2}MA: A distributed access protocol for wireless ATM networks," in *Proc. IEEE INFOCOM'98*, 1998.

[22] I. F. Akyildiz, D. Levine, and I. Joe, "A slotted CDMA protocol with BER scheduling for wireless multimedia networks," *IEEE/ACM Trans. on Networking*, vol. 7, pp. 146–159, Apr. 1999.

[23] I. F. Akyildiz *et al.*, "Medium access control protocols for multimedia traffic in wireless networks," *IEEE Network Magazine*, vol. 13, pp. 39–48, July–August 1999.

[24] K. Pahlavan, A. Zahedi, and P. Krishnamurthy, "Wideband local access: wireless LAN and wireless ATM," *IEEE Communications Magazine*, pp. 34–40, Nov. 1997.

[25] P. Bhagwat, P. Bhattacharya, A. Krishna, and S. K. Tripathi, "Using channel state dependent packet scheduling to improve TCP throughput over wireless LANs," *ACM Wireless Networks*, vol. 3, pp. 91–102, 1997.

[26] S. Lu, V. Bharghavan, and R. Srikant, "Fair scheduling in wireless packet networks," in *Proc. ACM SIGCOMM'97*, pp. 63–74, Sept. 1997.

[27] C. Fragouli, V. Sivaraman, and M. B. Srivastava, "Controlled multimedia wireless link sharing via enhanced class-based queuing with channel-state-dependent packet scheduling," in *Proc. IEEE INFOCOM'98*, Mar. 1998.

[28] T. S. E. Ng, I. Stoica, and H. Zhang, "Packet fair queueing algorithms for wireless networks with location dependent errors," in *Proc. IEEE INFOCOM'98*, pp. 1103–1111, Mar. 1998.

[29] P. Ramanathan and P. Agrawal, "Adapting packet fair algorithms to wireless networks," in *Proc. ACM/IEEE MobiCom'98*, pp. 1–9, Oct. 1998.

[30] S. Lu, T. Nandagopal, and V. Bharghavan, "A wireless fair service algorithm for packet cellular networks," in *Proc. ACM/IEEE MobiCom'98*, pp. 10–20, Oct. 1998.

Sunghyun Choi (S'96–M'00) is a Senior Member of Research Staff at Philips Research, Briarcliff Manor, New York. He received his B.S. (summa cum laude) and M.S. degrees in electrical engineering from Korea Advanced Institute of Science and Technology (KAIST) in 1992 and 1994, respectively, and received Ph.D. at the Department of Electrical Engineering and Computer Science in the University of Michigan, Ann Arbor in September, 1999. His research interests are in the area of wireless/mobile networks with emphasis on the QoS guarantee and adaptation, in-home multimedia networks, connection and mobility management, multi-media CDMA, and link layer protocols.

Dr. Choi was a recipient of the Korea Foundation for Advanced Studies Scholarship and the Korean Government Overseas Scholarship during 1997-1999 and 1994-1997, respectively. He is also a winner of the Humantech Thesis Prize from Samsung Electronics in 1997. His email address is sunghyun.choi@philips.com.

Kang G. Shin (S'75–M'78–SM'83–F'92) received the B.S. degree in Electronics Engineering from Seoul National University, Seoul, Korea in 1970, and both the M.S. and Ph.D degrees in Electrical Engineering from Cornell University, Ithaca, New York in 1976 and 1978, respectively.

He is currently a Professor and the Director of the Real-Time Computing Laboratory, Department of Electrical Engineering and Computer Science, The University of Michigan, Ann Arbor. He has authored and coauthored approximately 600 technical papers and numerous book chapters in the areas of distributed real-time computing and control, computer networking, fault-tolerant computing, and intelligent manufacturing. He has also co-authored (jointly with C. M. Krishna) a textbook "Real-Time Systems," McGraw Hill, 1997.

His current research focuses on Quality of Service (QoS) sensitive computing and networking with emphases on timeliness and dependability. He has also been applying the basic research results to telecommunication and multimedia systems, intelligent transportation systems, embedded systems, and manufacturing applications. Dr. Shin was a Distinguished Visitor of the Computer Society of the IEEE, an Editor of *IEEE Trans. on Parallel and Distributed Computing*, and an Area Editor of *International Journal of Time-Critical Computing Systems*.