Wireless data: Systems, standards, services

Antonio DeSimone* and Sanjiv Nanda*

Performance Analysis Department, AT&T Bell Laboratories, Holmdel, NJ 07733-3030, USA

Abstract. Wireless data products and services being proposed today include exotic mixes of services and technologies: packet transport over cellular circuits, facsimile service over Cellular Digital Packet Data (CDPD), voice and video over wireless LANs, and everything in between. Data networking terms that seem to have a clear meaning – data-link, network and transport layers; circuit-mode and datagram; connection-less and connection-oriented – in fact have meaning only in context. Thus TCP, a reliable packet transport protocol, is being used in CDMA circuit-mode data to provide a reliable data-link layer for the error-prone wireless link. IP datagrams will be transported over cellular links using dedicated channels with call establishment, possibly per packet. Market demands for timely solutions, competition between alternative technologies and the plethora of alternative fora for standards development are driving wireless data into fragmented directions. The primary constraints come from the limited spectrum, the need for security in the presence of mobility and the size and weight of mobile terminals and devices. Often the optimization for the latter constraints is sacrificed at the altar of the former drivers. Based upon our experience and work with standards and systems we attempt to put wireless data into perspective. We compare and contrast major services and products and identify the choices that were made and why.

1. Introduction

Systems and standards for wireless data have proliferated in recent years, as the capabilities of available technologies have risen to meet the requirements of emerging applications. This paper presents a survey of technology issues regarding wireless channel access, protocols for wireless links, and networking in a mobile environment. While wireless LANs provide local area coverage, several wide area wireless networking services and standards have been developed.

A number of wireless LAN products for in-building use are on the market. They provide shared access to a high-speed channel (hundreds of Kb/s to a few Mb/s) in unlicensed frequency bands, providing an alternative to wired local-area networks: the Proxim RangeLAN and the AT&T-GIS WaveLAN are prominent examples. These products can be attractive from the in-building wire-replacement perspective, by allowing network access without new cable and by allowing some degree of reconfiguration, although limited mobility. However, many office buildings are already wired up, and given the anticipated arrival of high bandwidth multimedia applications to the desktop, office network managers are reluctant to commit to technology that gives them 1 Mbps, and perhaps even less if the unlicensed bandwidth must be shared with multiple in-building LANs or other unlicensed PBX type products. These considerations have to date resulted in a limited market for these wireless LAN products that provide only limited range and mobility. An evaluation of currently available wireless LAN products in [24] agrees with this conclusion. The IEEE 802.11 committee [1] is developing a standard based on the types of access methods used in current wireless LAN products. The 802.11 media access control mechanism is discussed in section 2.1.

For wide-area coverage, data services on cellular systems can provide either circuit-mode access, where a mobile has dedicated use of a cellular channel in the same manner as voice, or packet-mode access, where mobiles contend for access on a packet basis. In either case, the bandwidth is typically limited to around 10 kbps or less. However, the market for wide-area service appears to be enormous. Modem and fax capabilities over analog cellular channels are available today, although the quality of the circuit is lacking. For the TDMA and CDMA digital cellular standards, sophisticated circuit mode data capabilities have been designed [33,32] and will become available in the next year or two. These standards are discussed in section 3.

Cellular Digital Packet Data (CDPD) [6] is a contention-based packet access scheme for analog cellular channels. CDPD provides a packet switched backbone infrastructure that provides gateways to the Internet and X.25 packet networks. The CDPD specification provides mobility management and routing through a network of nodes comprising MDIS (Mobile Data Intermediate Stations) and MDBS (Mobile Data Base Stations). The MAC and coordination of access with cellular analog voice channels are discussed in section 2. CDPD mobility management is described in section 4.2. There is wide-spread interest in CDPD for several types of low-bandwidth vertical services, such as credit-card verification. It is also possible to use the (soon to be deployed) CDPD backbone network for packet routing

* We acknowledge the contributions of our colleagues at AT&T Bell Laboratories: Ken Budka, Mooi Choo Chuah, Bharat Doshi, Subra Dravida, Richard Ejzak and On-Ching Yue.

while using circuit mode data connections on the digital (TDMA or CDMA) air interface. Such a scheme is discussed in section 3.4.

Several more specialized services also exist, to do messaging or enhanced paging, often as part of a vertical service offer. These include Short Message Service (SMS) over the soon-to-be-deployed digital cellular systems, the Mobitex, ARDIS, and other proprietary systems, as well as Specialized Mobile Radio (SMR). These will not be discussed further in this paper. For a market survey of current wide area wireless data services, we refer the reader to [8]. Also, [18] provides a look at the Mobitex system deployed by RAM mobile data and provides a market comparison with ARDIS and CDPD.

In addition, the problem of mobility for computers, whether wireless or wired, has become more important as portable computers have become powerful enough to qualify as fully-functional systems with networking capabilities, rather than as specialized electronic organizers or portable terminals. The mobile-IP activities in the IETF [25] address mobility in the context of IP networking. Mobility for data networking is also addressed in CDPD. In cellular SMS and circuit-mode data services, mobility is handled using the standard cellular Home Location Register (HLR), Visitor Location Register (VLR) model. These mobility management issues are taken up in section 4.

2. Multiple access in wireless systems and standards

2.1. IEEE 802.11 MAC layer

The IEEE 802.11 committee is standardizing protocols for wireless local-area networks [1]. The Media-Access Control protocol, known as Distributed Foundation Wireless MAC (DFWMAC), is essentially complete. 802.11 uses a contention mechanism to allow stations to share a wireless channel, based on carrier-sense multiple access (CSMA), in the spirit of 802.3. A straightforward extension of the CSMA/CD protocol used in 802.3 LAN is not possible in the wireless environment, because a station cannot simultaneously listen on the same channel on which it is transmitting, as required for "CD" (collision detection) part of CSMA/CD. A station on a wireless LAN, therefore, will not be able to determine that a collision has occured until the end of a packet transmission, making collisions more expensive in 802.11 than in 802.3. The 802.11 MAC uses a *collision avoidance* mechanism to reduce the probability of collision. The basic protocol is known as the Distributed Coordination Function (DCF) and is based on a distributed, contention-based media access protocol known as CSMA/CA. This is described in detail below.

The Basic protocol: Distributed coordination function

CSMA/CA is a variation on the usual CSMA protocol [28], where a station listens for signal energy in the band to determine if the medium is available, and transmits only when the medium is idle. The "CA" piece (for Collision Avoidance) is a mechanism to reduce the probability of collision among stations contending for the medium at the end of a frame transmission by calculating a random idle time at each station, during which the station defers transmission, waiting to see if the medium remains idle.

Fig. 1 gives a view of the CSMA/CA access mechanism. A key notion is using various values of the inter-frame spacing (IFS) to give priorities to different types of frames. The application of the short inter-frame spacing (SIFS) and PCF inter-frame spacing (PIFS) are described below. The basic CSMA/CA mechanism, described presently, uses the distributed IFS (DIFS).

From the point of the end of a frame transmission, the access procedure behaves as follows:

1. all stations are quiet for a time DIFS;
2. stations with new data to send calculate a random starting time (in slots) in the contention window. This is known as the "Access Back-off Procedure," and the starting slot is calculated as $CW \times \eta$ where η is a random variable, uniform on [0,1).

The slot duration is media-dependent, and is conceptually similar to the minimum packet size on CSMA/CD systems – the slot is chosen so that stations starting transmissions in different slots are guaranteed not to collide. CW is the contention window, which stations initialize to CW^{min}. Retransmissions cause the contention window to grow exponentially (binary exponential back-off) up to CW^{max}.

Exactly one random starting slot is calculated by each station for each frame: the station does not draw another random starting slot if another station begins transmission while it is waiting: instead, the station continues counting down once the channel becomes idle. The starting slot can be thought of as a timer that counts down *only* when the medium is sensed idle after a DIFS. When the timer reaches zero, the station transmits its frame.

Note that, in the interest of fairness, a station cannot send back-to-back packets (separated by a DIFS), but *must* generate a random back-off for each frame. The maximum single-station throughput is then simple to

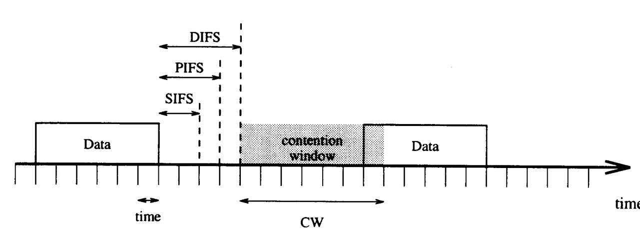

Fig. 1. CSMA/CA access mechanism.

calculate: sending a large number of maximum-size frames requires an idle time on average of $\mathrm{DIFS} + \frac{1}{2} CW^{min}$.

The augmented DFWMAC

The basic CSMA/CA mechanism for the 802.11 distributed coordination function is augmented with

- MAC-level acknowledgments and retransmissions,
- special ready-to-send/clear-to-send (RTS/CTS) frames, which implement a mechanism for reservations that addresses some of the hidden-terminal problems in wireless LANs,
- three values for the minimum inter-frame spacing, to implement three levels of priority.

Priorities are assigned to frames by allowing transmission with a smaller inter-frame spacing. An acknowledgment for a received frame uses the shortest IFS (SIFS), and has highest priority. Similarly, CTS frames in response to a just-received RTS are sent at highest priority, with a SIFS. An intermediate level (PIFS) is used for control traffic to implement the Point Coordination Function. This is a second mode of channel access in 802.11 that uses a centralized, polling mechanism. A unique station plays the role of the polling master. The centralized polling station defines a super-frame structure and using the PIFS gains priority access to the channel when contention-free access must be provided to stations requiring time-bounded service. Ordinary data frames are at lower priority, and use the DIFS.[1]

CSMA/CA performance

The 802.11 MAC is designed to operate over multiple physical layers, and does not specify various media-dependent parameters. In [34], to study protocol performance, parameter values are chosen to represent an RF spread-spectrum system like WaveLAN. In particular, results are provided for a 2 Mbps LAN speed with a packet size of 576 octets, which corresponds to a packet transmission time of approximately 2.3 msec. A key parameter is the *slot time*, which is equal to collision vulnerability period. Operationally, this means that if a station begins transmitting in a slot, every other station on the LAN should hear the transmission before the next slot. The slot time includes the propagation delay, the time to acquire the spreading code and the time to turn around the channel from receive to transmit. We reproduce results here from [34] for slot times varying between 5 and 125 μsec.

Table 1 shows the aggregate throughput performance of the 802.11 MAC, as a function of the number of contending stations and of the "vulnerability interval." In this simulation experiment, the DIFS is 2 slots and the SIFS (for Acks) is 1 slot. An important parameter is the collision window. A large collision window can improve performance by making the probability of collision small. However, even with no contention, the collision avoidance mechanism reduces the maximum throughput for a single station by enforcing an inter-frame gap. For a small number of stations, a small collision window can improve performance by making the idle time between frames small. In the simulation results, the minimum collision window (CW^{min}) is one slot, and the maximum (CW^{max}) is 10 slots. This set of parameters gives good performance for a small vulnerability interval, but poor performance for a larger value. When the vulnerability period is large, the throughput rapidly falls with increasing number of stations. The optimal value of the collision window will depend on the number of stations sharing the LAN, since the probability of collision will depend on the number of contending stations. This points out the need for sound engineering rules to configure 802.11 LANs.

Table 1
Throughput for CSMA/CA.

Vulnerability interval (msec)	Number of stations		
	1	5	10
0.005	0.78	0.71	0.71
0.025	0.78	0.57	0.42
0.125	0.78	0.11	0.02

2.2. MACs for cellular systems

In cellular systems, a central base station manages contention. Examples of cellular MACs include the access channels for cellular systems, including the analog AMPS, as well as the TDMA IS-54B digital control channel [31]. The CDPD media access control protocol also fits in this category.

As we saw earlier, in the wireless LAN environment where all transmitters share the same frequency band, collision detection is not possible. However, if the transmitter and receiver are in separate frequency bands, with adequate filtering and guard bands, the mobiles can simultaneously transmit on one band and receive base station feedback on the other. The collision detection interval is drastically reduced in this case, which improves the efficiency of the random access procedure. Thus with frequency division duplexing (FDD) in the cellular systems, busy/idle feedback from the base station provides collision detection.

Also, with busy/idle feedback from the base station, the hidden terminal problem, well known in packet radio networks, and a serious problem in wireless LANs, is avoided. Thus, even though two contending mobiles may not be able to hear each other's transmissions that result in a collision at the base station receiver, since all

[1] A fourth, still lower-priority value is suggested for "certain management frames."

mobile transmissions are to the central base station (not peer-to-peer), the base station is able to inform the mobiles of collisions through the use of busy/idle tones or flags on the feedback channel.[2]

Other common features of cellular MACs, including CDPD, include an exponential back-off mechanism and a slotted channel. To provide efficient utilization of the valuable cellular spectrum, all cellular MACs use a subset of the up-link and down-link flags discussed below.

In addition to the busy/idle flags discussed above, the feedback channel may also use an ACK/NACK flag to indicate decoding failure to request MAC level retransmissions. In CDPD (see schematic of the access channel in Fig. 2) this ACK/NACK flag is known as the decode status flag.

In the TDMA digital control channel, the reserved/available flag on the down-link is used to reserve the channel for use by a specific mobile. This further reduces contention if the base station desires to reserve the channel to get a response from a specific mobile, and wants to prevent any other mobiles from contending for the next slot.

Other up-link flags include an indication of initial/repeat (physical layer) burst and a more/final flag. The latter is used to indicate the desire by the mobile to retain the reservation for another (physical layer) burst by indicating that there is more data to send.[3] In CDPD, this is called the continuation/EOT (End Of Transmission) flag. The initial/repeat flag is used in the TDMA Digital Control Channel to indicate whether the current burst is a retransmission of a previously NACK-ed burst.

2.3. CDPD channel availability

CDPD is motivated by the low channel utilization in a typical AMPS cellular system. With a 7-cell reuse pattern, an AMPS cell may have as many as 60 voice channels, however, a more typical deployment has three sector cells with 15–20 channels per sector. When designed for good blocking performance (say 2%), with a trunk group of size 20, the channel utilization is quite low. CDPD uses "channel-sniffing" to determine when AMPS channels are unused by voice, and acquires them for CDPD. Then it defines a packet-based contention scheme to use the idle voice channels for data. The goal is to provide packet data service and connection to gateways to public data networks, using the existing cellular infrastructure. CDPD is capable of modest throughput: ignoring contention, the maximum throughput after accounting for framing and coding overhead is 11.8 kbps on the down-link and 13.3 kbps on the up-link. MAC contention further reduces the up-link throughput, and add access delays.

Further, because the CDPD packet data uses idle voice channels, data packets will be delayed if there is no idle voice channel available. Although, as we observed above, the average utilization of voice channels is low, the probability that all voice channels are in use is not negligible, and has impact on delay. With small trunk groups, the periods during which an AMPS channel is available to carry CDPD traffic, are long, typically, equal to several call holding times. Periods when all channels are used by voice calls correspond to CDPD channel unavailability. Such periods last a fraction of a call holding time. For example, with a (small) trunk group of 10 channels, the probability that an unavailability period lasts longer than 10% of a call holding time is around 0.5. That is, with a call holding time of 100 seconds, half the CDPD channel unavailability periods are longer than 10 seconds. This delay can have substantial impact on services. Note that with larger trunk groups, the dynamics become faster, so that the availability periods as well as the unavailability periods become shorter. Detailed studies of these issues are reported in [4] and [13].

[2] Of course, this assumes that the mobiles are able to hear feedback transmissions from the base station.

[3] The base station may deny this request by setting the busy/idle flag to idle, thus indicating loss of reservation.

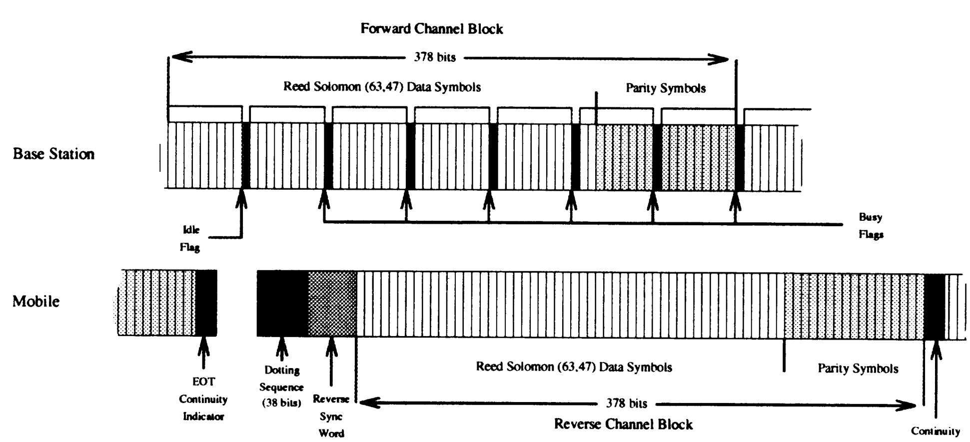

Fig. 2. The CDPD MAC.

3. Cellular data link layer standards

Both TDMA and CDMA North American digital cellular standards make provisions for circuit-mode data, motivated largely by a desire to provide good support for fax traffic on the cellular network. As shown in Fig. 3, the cellular standards define an Inter-working Function (IWF) to terminate the cellular circuit data protocol and to connect to a modem or fax over the wired network. Radio link protocols for circuit mode data for cellular systems are based on the philosophy that the error-prone cellular link must be made reliable. In particular, complete recovery is guaranteed between the IWF and the mobile.[4] As discussed next, this is not only desirable, on the wireless portion of the cellular circuit it is actually necessary.

3.1. Need for link layer retransmissions

Both the TDMA circuit-mode data standard and the CDMA circuit-mode data standard do complete recovery of errored frames between the mobile and the IWF, using different mechanisms. The TDMA standard uses periodic receiver-state feedback to recover lost data and is described below in section 3.3. CDMA uses TCP between the IWF and the mobile, with the usual timeout and retransmission mechanisms for complete recovery, augmented with a radio-link protocol that provides sequencing and will attempt limited retransmission of CDMA physical layer bursts. The CDMA scheme is described below in section 3.3. CDPD defines a link layer recovery mechanism called MDLP which is an enhanced version of LAPD recovery procedures. CDPD also does MAC level retransmissions on the up-link.

When the burst error rate [5] on the wireless channel is high, link (or MAC) level retransmissions are required for throughput. For example, if the burst error rate for a typical TDMA or CDMA burst of (approximately) 20 octets is 1% [6] then the probability that at least one physical layer burst of a 576 octet end-to-end data packet will be lost on the wireless channel, is 0.95. In the absence of link layer retransmissions, with 95% probability this long packet will have to be retransmitted by the end system. If the burst error rate on the wireless link could be improved to 10^{-6} (or better) through retransmissions [7] then the end-to-end throughput could be improved.

The actual improvement would depend on the behavior of both the link and the end-to-end retransmission mechanisms. For example, with TCP as the end-to-end mechanism, the actual throughput improvement is reduced by the fact that the link layer retransmissions increase both the mean and variance of the round trip delay. In a different context in [9] it was shown that, when the entire 576 octet end-to-end data packet is the retransmission unit on the wireless link, retransmissions to reduce error rates at the link level can in fact decrease TCP throughput. This is discussed further, next.

3.2. TCP in wireless networks

TCP is a reliable, connection-oriented transport protocol that provides correct, in-sequence communications even over unreliable networks by detecting lost and duplicate packets and requesting retransmission of lost packets. Many differences may exist in TCP implementations, with different retransmission and dynamic windowing behavior, even though they conform to the TCP specification. The behavior described here is what is seen in the BSD implementation, and similar behavior is seen in other commercial implementations with roots in BSD, the Berkeley version of Unix. The specific TCP mechanisms for loss detection, window flow control and error recovery are discussed in the appendix.

TCP performance can depend in a complex way on the characteristics of the underlying network. Fig. 4, taken from [9], shows the throughput of TCP in a simple two-link network, where one link introduces only delays and no errors, while the second is a noisy wireless link that drops packets at random. The throughput for TCP recovery only (TLR), is compared with the throughput for the case with both the wireless link (LLR) and TCP recovery (TLR). Here the assumption is that the entire

[4] The physical radio link terminates at the base station. A wired infrastructure provides connectivity from the base stations to the switch and IWF.

[5] In this paper, burst refers to the bits transmitted in one physical layer time slot. Burst error rate is the probability of erroneous reception of a physical layer burst.

[6] Note that a 1% burst error rate is the design point for the CDMA system.

[7] or through channel coding, but coding to achieve this would be rather wasteful of bandwidth.

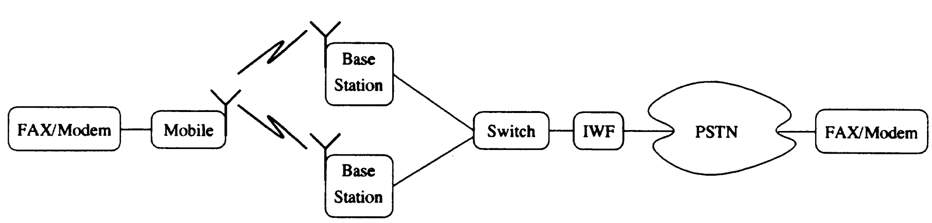

Fig. 3. Schematic for cellular circuit mode data.

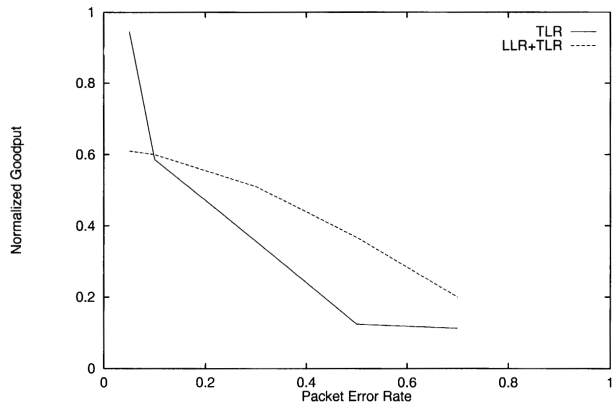

Fig. 4. TCP throughput.

TCP segment is retransmitted on the wireless link. This is of interest in wireless LAN environments.[8] The plots show the sensitivity of the throughput for TCP as a function of the packet loss rate on the slow wireless link, with deterministic link delays on all links. When the probability of packet loss is higher than 0.1, link layer recovery is beneficial in this example. Depending on the values of the parameters this tradeoff will be different in different scenarios.

Note also that this type of performance is peculiar to TCP, which was designed for shared packet networks, and its approach of jointly performing flow control, error recovery and congestion control (see appendix). A LAPD-like end-to-end retransmission schemes that is designed for circuit mode operation does not attempt to do recovery and congestion control simultaneously through the use of round trip acknowledgment delay estimates. Instead, fixed size windows and explicit NACKS are used for selective retransmission. More robust performance over the range of operating conditions is observed for circuit mode data applications using LAPD style recovery. The CDMA link layer retransmission scheme, discussed in section 3.3, uses TCP for link layer recovery on the wireless link between the IWF and the mobile.

TCP congestion control is based on the principles outlined in [14]. The approach has become wildly successful and pervasive because it can be implemented with no changes to existing networks. This approach to congestion control is also strongly encouraged by RFC 1122 [3]. As described in [14] and as implemented in many commercial products, TCP uses implicit indications of congestion to control traffic sources and relies on the cooperative behavior of the TCP sessions sharing a link to provide some degree of fairness and efficiency in sharing bottleneck resources. Very simply, the implicit indication of congestion assumes that any detected packet loss is due to congestion, and the requirement to share a bottleneck requires a rather aggressive response to congestion. These two premises fail to hold in many situations relevant to wireless networks. Losses due to errors can be significant on wireless links. Also, some applications, such as circuit-mode data connections, are dedicated to a single user so that the resource sharing issues are different than in a land-line backbone network.

3.3. Recovery in cellular circuit mode data

Two approaches are possible to do complete recovery between the IWF and the mobile in circuit-mode data services. The first approach is to use a reliable link layer ARQ scheme with moderate size link layer frames that are accommodated into fixed or variable size physical layer (radio link) bursts. When a radio link burst is lost, recovery is done through the retransmission of the link layer frame. Let us refer to this scheme as the single recovery level scheme to distinguish it from the two levels of recovery in the next scheme. This approach is taken in the TDMA standard [33], and is discussed in section 3.3. The GSM approach is similar and is discussed in [15].

In the alternate approach the basic unit of retransmission is the physical layer burst. The physical layer burst is of small (and variable) size (around 20 octets or less). To keep the overhead small, the CRC and other header fields to specify data types, sequence numbers, feedback etc. must be minimized. The physical layer retransmission scheme (to recover physical layer bursts) is then neither flexible (small header) nor reliable (short CRC). To obtain a reliable link layer up to the IWF an additional recovery mechanism is required in this case. Since partial recovery provided by the physical layer recovery mechanisms can recover in most cases, the second recovery level only occasionally retransmits. This means that the packet size for the second recovery level can be made large to reduce the overhead. This second recovery level can provide additional CRC and header fields to provide flexible services and reliability that was missing from the first recovery level. This approach is standardized in CDMA [32] (see below).

For moderate (physical layer) burst error rates (1% to 5%) with independent burst errors, the two level scheme with lower overhead per burst can provide better throughput since exactly one physical layer burst is retransmitted for each lost burst, as long as the second level of recovery is not invoked. At lower burst error rates, the difference between the two schemes becomes negligible, but the single level scheme has lower complexity.

The cellular link is subject to long durations of fades when the burst error rates are large, or disruptions during hand-offs when several consecutive physical layer bursts are lost. Also, occasionally, the limited CRC in the short physical layer burst can fail. In all these cases, the long packet of the second recovery level will have to

[8] In cellular networks with narrow bandwidth channels, wireless link layer recovery is done on smaller radio link frames, as discussed below.

be retransmitted. At those times the end to end throughput will experience a sharp drop and large delays will be seen. Although, throughput and delay hits will be observed using the single recovery scheme as well, both the variability of the throughput, as well as the variability of the delay are smaller for the single recovery level scheme.

When the second level of recovery is chosen to be one of the standard versions of TCP, these variations are observed to be greatly magnified. This is because the slow start and congestion control mechanisms built into TCP further magnify the effects of the observed values of the round trip acknowledgment delays through back-off procedures for the retransmission timeout (RTO) and the congestion window. Finally, the problem may be further magnified if the two level recovery scheme is used for circuit mode data at the IWF, and the end system application adds one more level of end to end recovery (often TCP).

TDMA receiver state feedback ARQ scheme

The TDMA receiver state feedback scheme [33] is a selective retransmission scheme that relies on explicit feedback from the receiver. Per packet timers are not required either at the transmitter or at the receiver.

The Receiver State consists of NR, NL and a Bitmap. NR is the sequence number of the next link layer frame required at the receiver for in-sequence delivery to the next higher layer at the receiver. NL is the largest sequence number frame received so far. Note that largest in this case is interpreted as within the window and modulo the highest sequence number. In particular, for the TDMA standard the sequence number modulus is picked to be 127. The Receiver State Bitmap is the the receive status of the frames with sequence numbers $NR - 1$ to NL, with a 1 in position k denoting that frame number $NR + k$ is received and stored in the receiver buffer, and a 0 denoting a missing frame. Note that $NL \geq NR - 1$.

The Receiver State is updated on the receipt of each frame. The receiver periodically transmits the complete receiver state i.e., NR, NL, Bitmap, or partial receiver state by placing the Bitmap into octets called Bitmap groups. The details of the partial receiver state feedback are given in the standards as well as in [23]. The crucial point to note here is that the implementations with complete and partial receiver state feedback are completely consistent. In fact, complete and partial feedback may be used at the same time, as the feedback channel bandwidth availability permits.

The scheme relies on the in-sequence delivery of frames over the underlying physical layer. That is, a frame transmitted later cannot be received prior to one that was transmitted earlier. The transmitter maintains its most current knowledge of Receiver State. In addition the transmitter remembers the order in which frames were transmitted or retransmitted. When a frame that is transmitted later is acknowledged by the receiver through the feedback, the transmitter determines that all unacknowledged frames that were last transmitted prior to the acknowledged frame, have been lost. The transmitter retransmits only these "lost" frames, with priority over new frames. Also, it updates the transmit order of the retransmitted frame to reflect when the particular frame was last transmitted.

It has been shown that this basic receiver state feedback scheme does no spurious retransmissions, and is therefore very efficient. Only frames that are lost are retransmitted. This means that the ratio of successful frames to transmitted frames is very high.

The actual data rate across the link can however, be increased by increasing the total utilization of the channel, with some spurious retransmission, and less efficient use of the channel. Since the TDMA circuits are dedicated per user, higher throughput at the expense of less efficient channel use is acceptable (even desirable).

In the TDMA standard, pre-emptive retransmissions are introduced during periods that the window is closed or when there is no new data to transmit. Using pre-emptive retransmission, the transmitter "pre-emptively" retransmits any unacknowledged data frames in each slot that there are no new or retransmitted data frames to be transmitted by the basic receiver state feedback scheme. These pre-emptively retransmitted frames increase the throughput during times that the window is closed to new transmissions, and complete and partial feedback frames are lost. Multiple copies of these frames may be received, thus resulting in decreased efficiency compared to the basic protocol, where no more than one copy of a frame is correctly received.

Without pre-emptive retransmissions, there are occasions when deadlock might occur. Since lost frames are detected only after frames that are transmitted later are received successfully, if the transmitter has no new frames to transmit after a frame that is lost, deadlock will occur. With pre-emptive retransmission, these frames will be pre-emptively retransmitted without waiting for feedback and will be eventually received and acknowledged. Without pre-emptive retransmission, to avoid deadlock, the transmitter must send a timer based poll to inform the receiver about the transmission of these trailing frames. The receiver can then update its state to show those frames as lost. Once the feedback is received at the transmitter those trailing lost frames will be retransmitted.

Recovery in CDMA circuit mode

As discussed earlier, using TCP, with a minimum segment size of 576 octets requires the use of link level recovery on the wireless link. However, link layer recovery increases both the mean and the variance of the round trip delay, which could lead to reduced end-to-end throughput, especially when TCP retransmissions are required. The variability of the round trip time due to link level recovery is especially detrimental, as TCP timers at the hosts will expire and retransmit TCP segments that the link layer retransmission scheme is still attempting to recover.

To reduce the variability of the round trip delay, a Radio Link Protocol (RLP) with partial recovery is used [17]. RLP attempts to do partial link level recovery through the use of receiver timers. The receiver transmits NACKs for lost frames and sets NACK timers. The NACKs are retransmitted if the frames are not received prior to the expiry of the NACK timers. An abort timer is set on the retransmission of the NACK. If the abort timer expires recovery of the lost frame is aborted and successfully received data is sent to TCP. The lost data will require recovery by the TCP. However, because of the *fast recovery* mechanism (see Appendix), this is preferable to a retransmission timeout. Thus the RLP abort is better suited for overall throughput, rather than full recovery by RLP.

The use of RLP for partial link layer recovery is appropriate if TCP is being used to provide end-to-end recovery. For the circuit-mode cellular data application, TCP is being placed in the IWF for recovery at an intermediate network node. The reasoning is that for subsequent packet data services with TCP possibly running at the end points, RLP can be reused for link layer recovery with the host TCP, and no TCP in the IWF (as discussed in section 3.4). Nevertheless, for circuit mode connections to hosts running TCP, as well as for fax machines that run their own recovery schemes, the CDMA circuit-mode service may give rise to unexpected and anomalous performance. With TCP at the hosts running over TCP between the IWF and the mobile, running over RLP, strange behavior will be observed when multiple TCPs retransmit the same data simultaneously at the host and the IWF, or at the TE and the MT.

Performance results

Extensive simulation studies of the TDMA and CDMA circuit mode data recovery schemes were performed as part of standardization activities. The physical layers and interference environment in the two systems are very different, hence comparisons of the two are not justified. A discussion of the constraints and the rationale behind certain design decisions that were made in the two standards is in [10]. Interestingly, both standards are defined to provide around 8 kbps user throughput. Higher data rates can be provided using multiple channels, with techniques already specified for TDMA and in consideration for CDMA standards.

In the TDMA scheme, the $R = 1/2$ convolutional coder used for the most sensitive voice bits, is punctured to $R = 5/6$ to maximize throughput for data. It is shown in [23] that the ARQ scheme is extremely efficient. Given a frame error rate (FER), the scheme achieves throughput very close to the $(1 - \text{FER})$ upper bound with reasonably large window size. In particular, the window is chosen to be between 32 and 64 link layer frames, which corresponds to around a kilobyte. The TDMA scheme provides a throughput between 7800 and 8200 bps over 90% of the coverage area, with hundreds of msec delay on the link.

In the CDMA scheme the physical layer is left unchanged from the voice channel, with sophisticated power control used to obtain a burst error rate around 1%. With RLP recovery only, the CDMA scheme gets a throughput in the 7500–8000 bps range up to 5% burst error rate. The maximum TCP window size is constrained to be 4 segments of maximum size 536 octets. Consequently the mean delays are again in the range of hundreds of msec.

On occasions that RLP aborts, or the physical layer CRC fails to detect an error, TCP recovery is required. This involves shrinking the TCP congestion window and expanding retransmission timeouts. Hence, recovery of the TCP following deep fades, or following hard handoffs that require RLP to be reset, result in periods of low throughput and large delays. It was observed through simulation that the recovery delay following a fade that lasts one second, can be several times longer than a second [11]. Fortunately, these TCP recovery events are rare, since RLP recovers with high probability. These results, and comparisons with an alternative approach that was proposed for the standard, will be reported in a paper that is in preparation [7].

3.4. Connecting via cellular to public packet data networks

As noted earlier, CDPD defines a multiple access procedure to be used on idle analog cellular (AMPS) channels. Additionally, CDPD provides connectivity through gateways to the Internet or X.25 public data networks. Once deployed, the CDPD backbone infrastructure will be used to connect to public data networks even after the cellular system has transitioned to digital (TDMA or CDMA). Packet-based multiple-access may be done using the digital control channel procedures of TDMA IS-54B, or on yet to be defined procedures for CDMA.

Providing a shared access channel in cellular has some shortcomings: inefficient use of the bandwidth, complexity at the physical layer, non-trivial impact on voice users and new security concerns. A simpler intermediate approach to packet network access, re-using the CDPD infrastructure, is to provide circuit connections over the wireless link for packet data users. Of course, the approach of assigning circuits to non-idle packet data users, and holding up the circuits until time-outs can be inefficient as well.

Whether the wireless access is packet or circuit, the general approach is that the packet data mobile would first register with the CDPD MDIS (IWF) connected to the local cellular system. This registration establishes packet forwarding from the Home MDIS, as discussed in section 4.2. It also establishes a virtual circuit for the registered mobile, between the IWF and the cellular network. Cellular calls (with packet data option) between the mobile and the IWF are originated by the network side or the mobile whenever there is back-logged data at either end.[9] If no data transfer occurs for a specific timeout period the cellular call is terminated, while the registration at the IWF remains valid and the virtual circuit remain open but becomes inactive.

When packet data from the land side arrives at the IWF for a registered mobile whose virtual circuit is inactive, the IWF can forward the data to the cellular network, which in turn will set up a packet data call through the usual means of paging and call establishment. Alternately, if the IWF knows when the packet data call to the mobile is down (i.e. the virtual circuit is inactive), then it may buffer the data and send an origination request for a packet data call to the cellular network. In particular, if the transport protocol is TCP, the TCP congestion window size remains at the value reached prior to call termination. Given, the 8 kbps rate on the channel, the TCP congestion window is likely to be no larger than 2000 octets or so (roughly a screen-full). Thus the host may dump up to 2000 octets of data destined for the mobile, on the IWF.

This suggests that up to 2 kilobytes of buffer per registered mobile is required, either at the IWF or on the cellular side. Even though the data is buffered, TCP round trip timers are likely to expire while awaiting the several second call establishment delay. Thus whether the congestion window's worth of data is buffered at the IWF or in the cellular network, the buffered data will be retransmitted by TCP.[10] To minimize the amount of data that TCP has to retransmit, the TCP window must be reduced prior to making the virtual connection inactive. There are at least two possible solutions.

If the inactivity timer is at the mobile, then on the expiry of the timer the mobile TCP, sends an ACK to the host TCP shrinking the maximum window size, and then terminates the packet data call. Thus, the TCP maximum window size is kept at the small value (perhaps zero), while the virtual connection is inactive. Although implementing this feature requires enhancement of the TCP implemented at the mobile, no changes are required in the host. The host implements only known TCP mechanisms in response to window updates.

Alternately, the TCP congestion window at the host may be shrunk by generating Source Quench packets at the IWF or other intermediate node. Source Quench packets are introduced by intermediate routers in an IP network to provide explicit congestion information to the end-system TCP. There are two problems. First, several details such as how the control mechanism must react to multiple Source Quench packets, as well as how frequently an IP router was permitted to transmit such messages, were never quite worked out. So many – perhaps most – TCP implementations simply ignore Source Quench packets. Second, Source Quenching may become obsolete in the future implementations of TCP. New mechanisms for Forward and Backward congestion indication in TCP connections are being introduced. These techniques are more sophisticated and the control mechanisms are being designed to work with high bandwidth connections. Given that Source Quench mechanisms never became popular and newer congestion control mechanisms are being standardized, the Source Quench mechanism might be abandoned.

Neither of the fixes suggested here are completely satisfactory, and the problem is still open. The primary issue (other than wasted network bandwidth), is the likely retransmission of user data up to the congestion window size when the packet data call is not up. Every time data is sent by a host to an inactive virtual connection, up to a congestion window's worth of data will be retransmitted, accompanied by an expansion of the round trip timers and reduction in throughput. A solution to this problem would provide improved performance and would permit reducing the inactivity timer and consequent wasted resources. We point out that shrinking the receive window would be even more valuable for future larger bandwidth cellular circuits where the TCP congestion window would be several times larger than 2 kilobytes.

3.5. Summary of link layer recovery

As mentioned above, the TDMA and CDMA circuit-mode data standards use link-layer retransmission (ARQ) to provide a reliable wireless link. The 802.11 MAC also uses an acknowledgment and retransmission mechanism to recover from errors on the link, whether caused by collisions or by channel impairments. The general issue of link-layer retransmission is complicated. Focusing only on the link, there is a tradeoff between ARQ and coding (FEC) that depends in a non-trivial way on the error characteristics of the link as well as the delay requirements for the service. Further, the overall, end-to-end performance is determined not just by the link performance, but by the performance over multiple links in the connection and by the behavior of the transport protocol. In particular, for protocols like TCP that use round-trip time measurements to control data flow, link-layer recovery may increase link throughput but the variable delays introduced by link-layer recovery will interact in a complex way with the

[9] For CDMA, it is proposed that only RLP is used for recovery under the packet data option, with the assumption that there is end-to-end recovery at the transport layer.

[10] Note that if the data from the host is in response to a transaction initiated by the TE, this would not apply since the virtual connection is activated by the request. Of course, if the host's response arrives after the expiry of the inactivity timer and the subsequent tear-down of the call, the same problem occurs.

transport protocol. On timeout, the transport protocol may retransmit data which the wireless link layer is still attempting to retransmit. These retransmissions will be discarded as duplicates after transmission across the wireless link, thus reducing throughput. The CDMA circuit data protocol hints at some of this complexity by limiting the level of retransmission in the RLP [17].[11] TCP performance over wireless links is currently of great interest [5,20,2].

4. Mobility management

Mobility management for wireless networks is a topic of much interest. Several papers reporting on research into database and resource management for mobility are available in the literature for example [22]. The cellular approach applicable to voice calls and circuit-mode data is to make routing decisions prior to call establishment through the use of separate signaling [19,30]. Thus a Home Location Register (HLR) contains the most current knowledge of the user location, updated through the use of location update procedures.

In this section, we briefly discuss mobility management for packet networks. Route optimization per call through signaling is appropriate for the circuit mode, but cannot be done per packet in packet networks. Several approaches to mobility in packet networks have been proposed and implemented. In the local area, one can expect a surfeit of capacity, and tracking mobiles can rely on link-level broadcasts and flooding techniques, similar to the 802.1D bridging approach used for wired networks. This approach has problems in scaling to large networks – as evidenced by the traffic problems in large bridged networks – hence other strategies are used in wide-area networks.

4.1. Mobile-IP

The mobile-IP protocol [25] is particularly notable, since IP is available on all the usual computing platforms. The mobile-IP protocol is still in draft stage, but the mechanisms are fairly well-understood. The operation of the protocol is sketched in Fig. 5. Briefly, the operation is as follows:

1. A Mobile Node (MN) discovers a Foreign Agent and obtains a Care-Of-Address.
2. The MN registers with its Home Agent. Registration messages carry cryptographic authentication.
3. Traffic for the MN is encapsulated by the Home Agent, and forwarded to the Care-of-Address of the MN.

[11] The CDMA circuit mode protocol is an odd example because TCP is *not* used end-to-end, rather "end-to-middle."

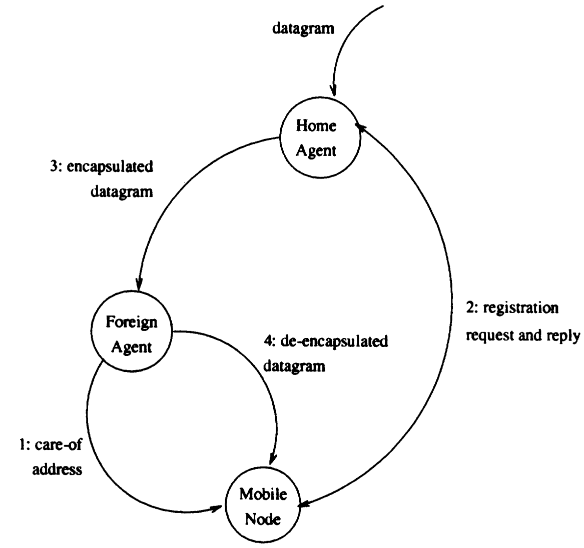

Fig. 5. Operation of the mobile-IP protocol.

4. The Foreign Agent de-encapsulates the datagram for forwarding to the MN.

The procedures do not preclude the MN being its own Foreign Agent, if it can obtain a legitimate Address on the visited network. Key exchange between the MN and the Home Agent can be done securely when the MN is in the home network, and this key can be used later for mutual authentication between the MN and the Home Agent. No authentication is required between the MN and the Foreign Agent.

In this scenario, all traffic to a mobile node must be routed via the Home Agent. A proposal exists for triangle-leg elimination in certain cases, but is not part of the base standard. Also, no provision is made for moving to a new Foreign Agent (a hand-off) without re-registering with the Home Agent. Again, a proposal for such a capability exists. As with the triangle-leg elimination proposal, the hand-off proposal introduces new security concerns since both rely on mechanism for manipulating packet routing that might be exploited by an attacker. Mobility compromises the existing IP security models in the following ways:

- Existing Internet firewalls rely on hierarchical IP addressing.
- Mobility introduces another mechanism for manipulating packet routing.
- Efficient routing under mobility (triangle leg elimination) requires intermediate nodes to change routing, which introduce still more points that can be compromised.

Much work is in progress in the IETF to improve the security mechanisms in IP. Key distribution is currently being standardized. With the availability of key management any two entities will be able to mutually authenticate each other and several more sophisticated routing procedures may be added to enhance the service. Some discussion of proposed enhancements and alternative approaches are found in [16,26].

4.2. CDPD mobility management

Mobility management in CDPD [6] follows the mobile-IP model, except that network entities are assumed to be trusted. Several Mobile Data Base Stations (MDBS) are connected to an Mobile Data Intermediate Station (MDIS). The MDIS as well as other Intermediate Stations (IS) comprise the backbone routing network. Some IS act as gateways to and are routers in the Internet or the X.25 public data network.

The Mobile End Station (MES) authenticates and registers with a visited MDIS which establishes a forwarding entry in the routing table at the Home MDIS for the MES. Roaming between the MDBS supported by the same MDIS is handled by the MDIS. Roaming outside the area covered by an MDIS requires re-registration with the Home MDIS and an update of the forwarding entry in the Home MDIS. Forwarding between one visited MDIS and another is not permitted. Packets arriving at an MDIS for an MES not registered with it, are dropped.

5. Discussion on applications

Complexity of wireless data applications increases with the mobility-bandwidth product. Today, applications of wireless data range from vertical services in the area of sales and delivery, which often use low-speed wide area networking from a licensed service provider; to wire replacement applications for mobile multi-media computing in an in-building environment, which use high-speed networking, often in unlicensed spectrum.

Various cellular-based services discussed in this paper may be tailored to the requirements of some specific set of applications. While the cellular approach enables wide-area mobility, limits to the available spectrum make high-speed services unlikely.

Wireless LANs typically offer limited mobility and higher bandwidths in the range 100 kbps to 1 Mbps. Higher bandwidths are desirable, and although several experiments, prototypes and standards are in progress, products are still several years away. In particular, wireless LAN speeds are also limited due to the limited bandwidth available in the unlicensed spectra. Although more bandwidth is available at 5 GHz and higher frequency bands, the hardware is more expensive.

Wireless data users today must therefore make do with products and services that are subject to low bandwidths and limited networking. Research to provide higher rate access continues at the physical and MAC layers. Simultaneously, new networking protocols and standards are being designed to manage mobility, addressing, and quality of service for wireless data. As motivation, we conclude this paper with a discussion of a particularly interesting application of wireless data networking that requires sophisticated mobility management (although in a local area) and large bandwidth. This application is to provide messaging and database access on a stock exchange floor. In particular, a study carried out by AT&T found the following requirements for the New York Stock Exchange (NYSE) floor:

- Around 2000 wireless hand-held device users distributed in an area approximately one city block.
- Total shared bandwidth requirement between 5 and 20 Mbps.
- Stringent delay requirements: less than 100 msec.
- High mobility.
- Hostile indoor propagation environment that does not permit easy spectrum reuse due to multiple reflections.
- Security concerns.

Products currently available on the market do not satisfy all of the above requirements, although it is possible to come up with solutions that can cover up to 50 or so users. In fact a variety of security exchanges, including the American Stock Exchange and the Commodities Exchange (COMEX) in New York have conducted limited trials. A limited production system was deployed and evaluated at the COMEX [21]. Meeting the requirements of a comprehensive system on the scale of the NYSE is at the limits of current technologies.

Appendix: TCP description

In this appendix, we discuss the mechanisms used to implement the congestion control in the BSD 4.3 "Reno" TCP implementation. TCP has gone through several refinements in various releases of the BSD networking code. The core mechanisms, which provide for loss detection, window flow control, error recovery and congestion avoidance, are taken up in turn in the sections below.

Loss detection

The basic mechanism to detect packet loss in TCP is per-TCP-segment timers maintained at the transmitter. A TCP transmitter will maintain an exponentially weighed-moving average as an estimate of the mean round-trip acknowledgment delay. According to [14], the retransmission timeout (RTO) is nominally set to a value of the mean plus a multiple of the deviation of the round-trip acknowledgment delay. This is an improvement on the proposal in the TCP specification [27] to use a small multiple of the mean round-trip. In TCP-Reno, the RTO is kept using a coarse-grained timer with a resolution of 500 milliseconds.

TCP also uses an auxiliary mechanism to detect loss: TCP's acknowledgments are cumulative, and repeated acknowledgments are interpreted as an indication that a segment has been lost, and later segments are causing the receiver to send acknowledgments for the last correctly-received segment. TCP-Reno retransmits when it receives three duplicate acknowledgments. This is sometimes called the *fast retransmit* procedure.

Since TCP does not distinguish packet losses due to congestion from those due to transmission errors, congestion control procedures are set into motion whenever a loss is detected. This involves shrinking the congestion window and expanding the RTO, as discussed below.

Window flow control

TCP uses window flow control. A very basic (and perfectly legal) TCP would use a fixed-size window, controlled by the receiver to guarantee that receiver's buffers are not over-run. On establishment, TCP connection end-points or hosts negotiate a receiver window size, that is usually chosen to be equal to the receive buffer size at the host. There is no requirement that the maximum window size in the two directions be equal. End-to-end flow control is managed through this maximum window. Through TCP acknowledgments, the receiver informs the transmitter about the available receive buffer space. This dynamically controls the amount of data that the transmitter can send to the receiver, so that the receiver buffers never overflow.

The insight of [14] is to augment the simple *receiver window* in TCP with a *congestion window* that is controlled by the transmitter and that attempts to prevent loss due to over-run of buffers in the network. This is a much harder problem than setting the receiver window. The congestion window size is determined by the congestion avoidance algorithms.

Error recovery

TCP-Reno error recovery can recover from one loss in a round-trip time. Once the loss has been acknowledged, the transmitter can advance the "left window edge" past all acknowledged data and resume sending under the constraints of the current congestion window.

Congestion avoidance

The TCP congestion avoidance algorithm proposed in [14] is colloquially known as "slow-start." The basic algorithm is very simple: For *each* lost segment detected, TCP reduces its congestion window, and then increases the window as acknowledgments are received. A retransmitted segment will "back off" the retransmission timer, doubling the timer for each retransmission. These two aspects – rapid shrinking and slow increase of the window, and binary exponential back-off for retransmitted segments – are considered essential for the equitable sharing of bottleneck resources by multiple TCP sessions.

The detailed operation of TCP is described well in [29] and [12]. A connection begins in *Slow Start*, with a congestion window of one segment. Each acknowledgment increases the congestion window by one segment, so the congestion window is doubled each round-trip time, until the congestion window reaches the slow-start threshold. The sender then enters the congestion avoidance phase, increasing the congestion window by roughly one segment per round-trip time, up to the receiver window size. A loss due to timeout puts the transmitter into slow-start, and a fast retransmit – triggered by three duplicate acknowledgments – puts the transmitter into the congestion avoidance state. This response to duplicate acknowledgments is sometimes called *Fast Recovery*.

References

[1] IEEE 802.11, *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, Institute of Electrical and Electronics Engineers, Inc. (November 1993), draft standard.

[2] B.R. Badrinath, A. Bakre, T. Imielinski and R. Marantz, Handling mobile clients: A case for indirect interaction, *Fourth Workshop on Workstation Operating Systems* (1993). See also ftp://paul.rutgers.edu/pub/badri/itcp-tr314.ps.Z.

[3] R.T. Braden, Requirements for Internet hosts - communication layers, RFC 1122 (October 1989).

[4] K. Budka, Cellular digital packet data: Channel availability, *Proc. PIMRC '95*.

[5] R. Caceres and L. Iftode, Improving the performance of reliable transport protocols in mobile computing environments, IEEE J. Select. Areas Commun. (1995), Special Issue on Mobile Computing Networks.

[6] CDPD, P.O. Box 97060, Kirkland, WA 98083, USA, Cellular Digital Packet Data System Specification: Release 1.0 (July 19, 1993) Vols. 1–5.

[7] M.C. Chuah, S. Nanda and B.T. Doshi, Performance comparisons of two CDMA radio link protocols, in preparation (1995).

[8] *Special Report on Wireless Data*, Data Communications, 21 March 1995 (McGraw-Hill, 1995).

[9] A. DeSimone, M.C. Chuah and O.-C. Yue, Throughput performance of transport-layer protocols over wireless LANs, *Proc. IEEE Global Telecommunications Conference – Globecom '93* (1993) pp. 542–549.

[10] B.T. Doshi, R.P. Ejzak and S. Nanda, Radio link protocols for cellular data, *Fifth WINLAB Workshop* (1995), papers to be published by Kluwer Academic Publishers.

[11] B.T. Doshi, R.P. Ejzak and S. Nanda, Working papers, CDMA Development Group (1994).

[12] S. Floyd, TCP and explicit congestion notification, Comp. Commun. Rev. 24 (1995) 8–23.

[13] J.M. Jacobsmeyer, Capacity of channel hopping channel stream on cellular digital packet data (CDPD), *Fifth Virginia Tech Symposium on Wireless Personal Communications* (1995).

[14] Van Jacobson, Congestion avoidance and control, in: *Communications, Architectures & Protocols*, ACM SIGCOMM (August 1988) pp. 314–329.

[15] J. Jayapalan and M. Burke, Cellular data services architecture and signaling, IEEE Personal Commun. Mag. 1(2) (1994) 44–55.

[16] D.B. Johnson, Scalable support for transparent mobile host internetworking, *Ninth Annual IEEE Workshop on Computer Communications (1994)*, Wireless Networks, this issue.

[17] P. Karn, *The Qualcomm CDMA Digital Cellular System*, `ftp://ftp.qualcomm.com/pub/cdma`.

[18] M. Khan and J. Kilpatrick, Mobitex and mobile data standards, IEEE Commun. Mag. (March 1995) pp. 96–101.

[19] Y.-B. Lin and S.K. DeVries, PCS network signaling using SS7, IEEE Personal Commun. Mag. (July 1995).

[20] P. Manzoni, D. Ghosal and G. Serazzi, Impact of mobility on TCP/IP: An integrated performance study, IEEE J. Select. Areas Commun. (1995), Special Issue on Mobile Computing Networks.

[21] J. Mason, Pork bellies and spread spectrum, Commun.: The Magazine of Mobile Radio Professionals (1993).

[22] S. Mohan and R. Jain, Two user location strategies for personal communications services, IEEE Personal Commun. 1(1) (1994) 42–50.

[23] S. Nanda, B.T. Doshi and R.P. Ejzak, A retransmission scheme for circuit mode data on wireless links, IEEE J. Select. Areas Commun. JSAC-12 (1994) 1338–1352.

[24] D. Newman and K. Tolly, Wireless LANs: How far? how fast? Data Commun. (21 March 1995) 77–87.

[25] C. Perkins, *IP Mobility Support*, Internet Engineering Task Force (21 March 1995), Internet Draft (work in progress). See also the material at `ftp://software.watson.ibm.com/pub/mobile-ip/mobile-ip.html`.

[26] C.E. Perkins and K. Luo, Using DHCP with computers that move, *Ninth Annual IEEE Workshop on Computer Communications* (1994), Wireless Networks 1 (1995) 341–353.

[27] J.B. Postel, Transmission control protocol, RFC 793 (September 1981).

[28] M. Schwartz, *Telecommunications Networks: Protocols, Modeling and Analysis*, Addison-Wesley Series in Electrical and Computer Engineering (Addison-Wesley, 1987).

[29] R. Stevens, *TCP/IP Illustrated* (Addison-Wesley, 1994).

[30] Telecommunications Industry Association, *Cellular Radio-Telecommunications Intersystem Operations, IS-41*.

[31] Telecommunications Industry Association, *Cellular System Dual-Mode Mobile Station – Base Station: Digital Control Channel (to be published as IS-136.1)*, Draft (November 1994).

[32] Telecommunications Industry Association, *Data Services Option Standard for Wideband Spread Spectrum Digital Cellular System, IS-99*, V&V Draft (September 1994).

[33] Telecommunications Industry Association, *TDMA Radio Interface: Radio Link Protocol 1*, V&V Draft (August 1994).

[34] O.-C. Yue, M.C. Chuah and A. DeSimone, A performance characterization of the IEEE 802.11 MAC, work in progress.

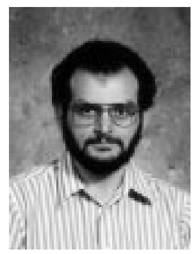

Antonio DeSimone is with Performance Analysis Department at AT & T Bell Laboratories, Holmdel, NJ. He is engaged in the design and analysis of local and wide-area communications networks, computer systems and data communications protocols. He holds a Ph.D. (1986) and an M.Sc. (1983) in physics from Brown University, and a B.S. (1981) in physics from Rensselaer Polytechnic Institute.
E-mail: tds@hoserve.att.com

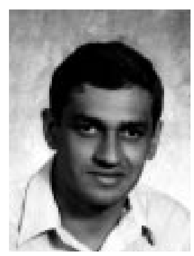

Sanjiv Nanda received the B. Tech. degree in electrical engineering from the Indian Institute of Technology, Kanpur, India, in 1983; the M.S. degree in mathematics in 1986, and the M.S. and Ph.D. degrees in electrical engineering in 1985 and 1988, respectively, all from the Rensselaer Polytechnic Institute, Troy, NY. During 1989 and 1990 he was with the Wireless Information Network Laboratory (WINLAB) at Rutgers University, Piscataway, NJ. He joined the Performance Analysis Department at AT & T Bell Laboratories at Holmdel in 1990, where he is currently a Member of Technical Staff. He is currently involved in design, performance study and modeling of wireless communications systems of the future. Dr. Nanda won the Jack Neubauer Award of the IEEE Vehicular Technology Society for the best systems paper published in the VT Transactions in 1991.
E-mail: nanda@hoserve.att.com