# ♦ Cellular Digital Packet Data Networks

*Kenneth C. Budka, Hong Jiang, and Steven E. Sommars*

*Cellular digital packet data (CDPD) is the first wide area wireless data network with open interfaces to enter the wireless services market. Operating in the 800-MHz cellular bands, CDPD offers native support of transmission control protocol/Internet protocol (TCP/IP) and connectionless network protocol (CLNP). This paper presents an overview of CDPD's network architecture, its network performance issues, and wireless data applications that take advantage of CDPD's seamless support of IP and CLNP data.*

## Introduction

Cellular digital packet data (CDPD) unites two dynamic technologies: internetworking and wireless communications. Designed as an overlay to typical 800-MHz analog cellular (Advanced Mobile Phone System [AMPS]) networks, CDPD seamlessly supports network applications based on Internet protocol (IP) or connectionless network protocol (CLNP). Native support of these popular networking protocols allows mobile data users to run familiar applications and facilitates rapid development of new applications that take advantage of CDPD's anytime, anywhere access to internets and intranets.

The CDPD system specification[1] was developed in the early 1990s by a consortium of U.S. cellular service providers, later organized as the CDPD Forum. Members of the CDPD Forum sought to create an open, nationwide wireless data service. In late 1994, CDPD entered commercial service.

After an initial connection setup procedure called *registration*, CDPD mobile subscribers can send and receive data on demand without additional connection setup delay. CDPD networks were designed to support "pay-by-the-packet" and "pay-by-the-byte" billing schemes. These schemes, shown in **Table I**, combined with attractive pricing packages, make wireless data a cost-effective option for applications that periodically send and receive relatively small amounts of data. For such applications, the cost of making a circuit-switched wireless call for each transaction or of keeping a call established could be prohibitive.

The CDPD network builds on the familiar cellular network architecture shown in **Figure 1**. CDPD includes specifications for an air link, mobility management, accounting, and internetworking.

## CDPD Network Architecture

CDPD networks consist of several major components:

- Subscriber devices,
- Infrastructure equipment provided by a cellular operator, and
- Network connections to internets and intranets, as shown in **Figure 2**.

An air link provided by CDPD efficiently supports digital data over 800-MHz cellular frequencies, with each 30-kHz cellular channel capable of serving multiple CDPD subscribers simultaneously. A subscriber registered with the CDPD network may keep a session intact for many hours, regardless of the volume of data sent or received.

CDPD uses different strategies for managing access to the *forward* (data flowing to the mobile user) and *reverse* (data flowing from the mobile user) channels. Base stations continuously transmit CDPD's forward air link, sending control and data signals to mobile units. The mobile data intermediate system (MD-IS) controls the flow of data in the forward direction and sends it serially in the form of link layer frames. A mobile data base station (MDBS) then relays the link layer frames over the forward air link. The CDPD digital sense multiple access control protocol with collision detection (DSMA/CD), discussed later in this paper, governs transmissions over the reverse air link. This Ethernet-like medium access control (MAC) protocol arbitrates reverse air link contention.

CDPD supports both unacknowledged broadcast and multicast services. Transmission of a single packet over CDPD's *broadcast* service efficiently sends the packet to all mobiles in a geographic area. CDPD's *multicast* service transmits messages to a select group of mobile end systems (M-ESs) in an area. Multicast transmissions follow subscribers in the multicast group as they roam. Only one packet is sent over the air for all members of a multicast group registered for receipt of multicast data on a particular CDPD channel, saving air link bandwidth.

### Mobile End System

The CDPD subscriber device, called a mobile end system (M-ES), takes a variety of forms: an integral part of a hand-held mobile telephone, a Personal Computer Memory Card International Association (PCMCIA) card installed in a laptop, or a hardened point-of-sale terminal.

The major components of an M-ES include a modem/radio and a processor running the CDPD protocol stack. CDPD uses Gaussian filtered minimum shift keying (GMSK) modulation[2] and Reed-Solomon[3] forward error correction (FEC) to provide wireless data service in cellular's typically harsh radio-frequency (RF) environment. M-ESs may be full duplex, with separate transmitter and receiver sections in the radio, or half duplex, with a transmitter and a receiver that share major components. Full-duplex mobile units—which, unlike half-duplex units, can transmit and receive data simultaneously—provide superior throughput and service under adverse RF conditions. As with all portable wireless devices, power management is an important function. CDPD mobile units have a range of maximum transmit power levels from 0.6 to 3 watts and dynamically change their transmit power level to conserve power and reduce interference. An optional sleep mode conserves power by allowing M-ESs to power down RF circuitry when they are not transmitting data and to periodically awaken to see if forward data is pending.

The M-ES uses information broadcast over the forward channel to determine when and how to

**Panel 1. Abbreviations, Acronyms, and Terms**

AMPS—Advanced Mobile Phone System
CDMA—code division multiple access
CDPD—cellular digital packet data
CLNP—connectionless network protocol
CM-ES—circuit-switched mobile end system
CMD-IS—circuit-switched mobile data intermediate system
CPU—central processing unit
CS-CDPD—circuit-switched cellular digital packet data
DSMA/CD—digital sense multiple access with collision detection
FEC—forward error correction
F-ES—fixed end system
GMSK—Gaussian filtered minimum shift keying
GPS—global positioning system
HDML—handheld device markup language
ICMP—Internet control message protocol
IP—Internet protocol
IS—intermediate system
LAP-D—link access procedures for D (data) channel
MAC—medium access control
M-ES—mobile end system
MDBS—mobile data base station
MD-IS—mobile data intermediate system
MDLP—mobile data link protocol
OSI—Open System Interconnection
PCMCIA—Personal Computer Memory Card International Association
PDU—protocol data unit
POTS—"plain old telephone service"
PSTN—public switched telephone network
PVC—permanent virtual circuit
RF—radio frequency
SNDCP—subnetwork dependent convergence protocol
SREJ—selective reject
TCP—transmission control protocol
TDMA—time division multiple access
TIA—Telecommunications Industry Association
UDP—user datagram protocol
WAN—wide area network

Table I. High-level comparison of packet- and circuit-switched data networks.

| Billing scheme | Network architecture | Latency for initial transmission (sec) | Typical pricing | Typical network usage |
|---|---|---|---|---|
| Packet data (for example, CDPD) | Connectionless | ≤ 1 | Based on data volume | Short, bursty transactions |
| Circuit data | Connection-oriented | 5-20 | Based on connect time | Larger amounts of data |

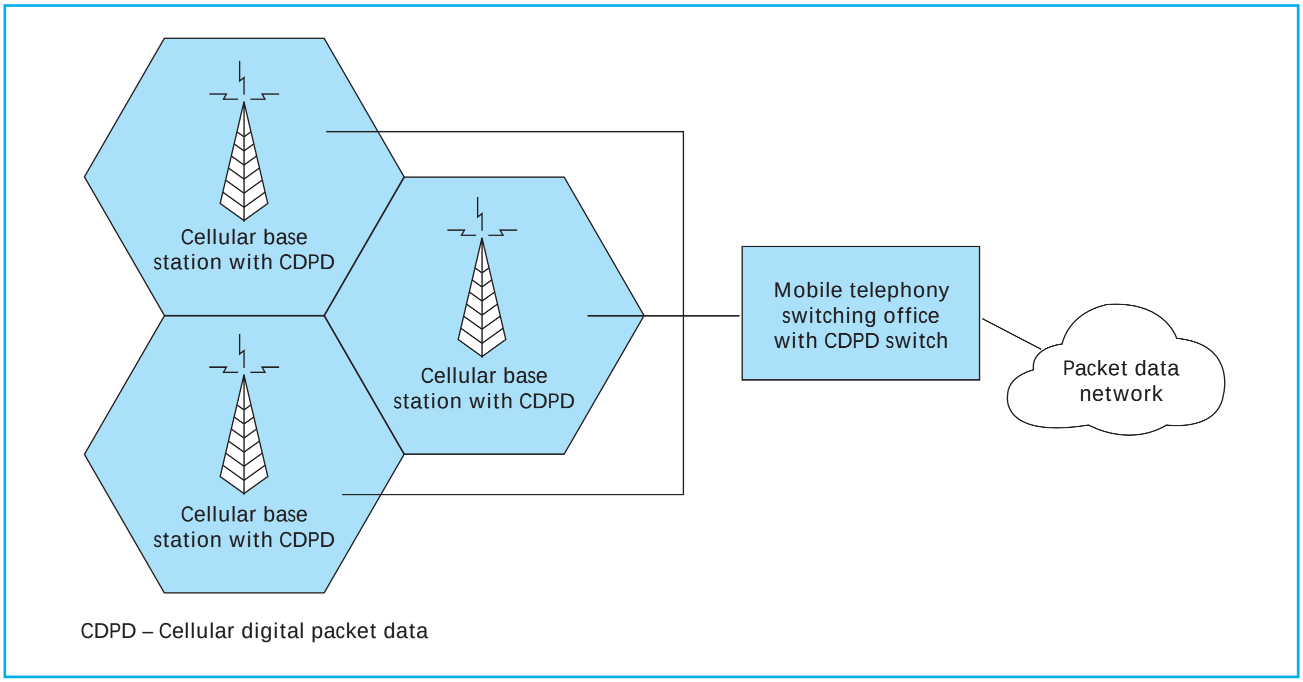


*Figure 1.*
*CDPD as an AMPS network overlay.*

search for new RF channels. RF channels may need to be changed when:

- The mobile unit travels between cell sectors,
- The current RF channel experiences fades or interference, or
- Contention with the AMPS network for RF channels triggers channel hopping, as described later in this paper.

### Mobile Data Base Station

The mobile data base station (MDBS) resides at the cellular site and typically covers a geographical area of 0.5 to 5 km in radius. On the forward channel, the MDBS transmits status data about its transmit power level, the adjacent sectors' CDPD channels, the decode status of received reverse data, and the activity status (busy or idle) of the reverse channel. The MDBS does not maintain information about registered mobile units, nor does it play a direct role in M-ES mobility.

CDPD often uses the same base station antenna as an AMPS cell and the same cellular voice RF plans. Network planners may configure the RF channels used by CDPD in various ways at the MDBS:

- *Dedicated*. One or more 30-kHz channels are dedicated to CDPD in each sector of the cell site. This is the simplest configuration, but it uses the most RF spectrum.
- *Omnidirectional overlay*. Each cell broadcasts an omnidirectional CDPD signal, overlaid on a sectorized AMPS cell. Fewer RF channels are used for CDPD, with a corresponding reduction in CDPD capacity.
- *Channel hopping*. CDPD radios are dedicated to

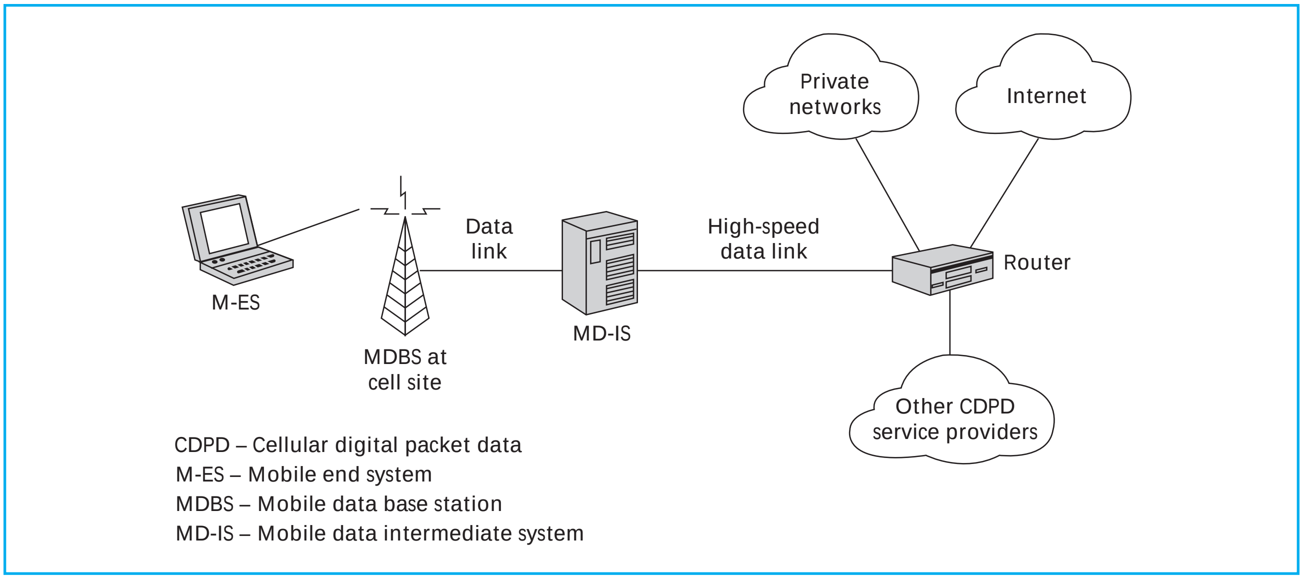


***Figure 2.***
***CDPD network architecture.***

each AMPS sector, but CDPD attempts to share 30-kHz channels with AMPS calls, where AMPS calls have priority over data. As described later in this paper, channel hopping works well when AMPS blocking rates are low to moderate.

At the cost of complicating RF planning, different cell configurations can be used in adjacent cell sites, or even within the same cell.

CDPD was designed to coexist with AMPS and to share/reuse many components such as power, enclosures, antennas, and RF amplifiers. Cellular service providers have found it easy to add CDPD equipment to existing AMPS base stations. CDPD is being used in some areas, however, as a standalone data network.

### Mobile Data Intermediate System

The mobile data intermediate system (MD-IS), typically located at the mobile telephony switching office, provides:

- Support for CDPD mobile protocols, including transmission of subscriber data. To prevent fraud, M-ESs are authenticated as part of the registration process and are denied CDPD network access if they present invalid credentials. The MD-IS and M-ES share responsibility for ensuring that user data is reliably sent. Data flowing between the MD-IS and M-ES is encrypted to protect it against eavesdropping.
- Mobility management. Mobiles must be tracked as they travel between cells or between channels within a single cell.
- Accounting. The MD-IS records detailed accounting data in a standard format.
- An interservice provider interface. As a mobile roams outside its home area, the mobile unit's home MD-IS must cooperate with the *serving* MD-IS. The home MD-IS must determine if the mobile is allowed to receive service and tunnel forward subscriber traffic from the home to serving systems.
- Connections into wide area networks (WANs). The MD-IS is typically connected to one or more conventional routers that route subscriber traffic towards its destination.

From a network and application viewpoint, CDPD is a wireless extension of the Internet. IP-based applications usually run on CDPD networks with *no* modifications. End users may find it beneficial to make some changes to their applications to improve performance and to lower network usage costs, as discussed later, in "CDPD Applications." The CDPD specification supports conventional IP (IPv4) and Open System Interconnection (OSI) mobile devices. Support of the next version of IP (IPv6)[4] is planned.

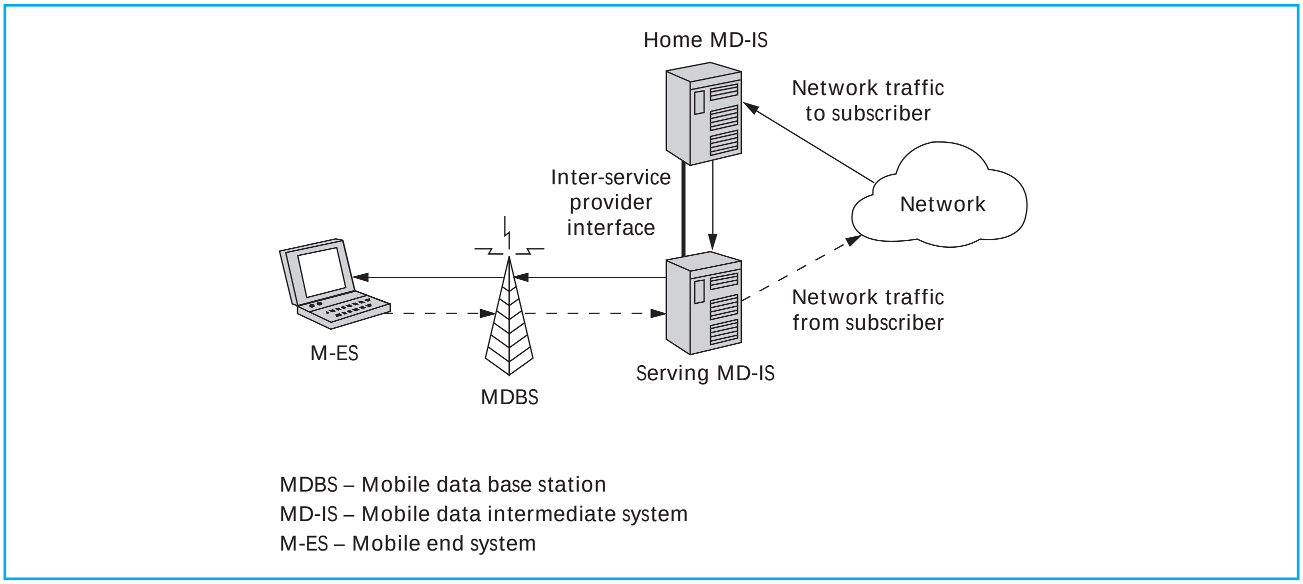


***Figure 3.***
***CDPD mobility management.***

The MD-IS is the focal point for mobility management, either within a service area or as the subscriber roams between CDPD service providers. A CDPD subscriber must be registered with the CDPD network to receive service. After the mobile unit has found an appropriate CDPD channel, it registers by sending its network credentials, based on shared secrets. These credentials validate that the mobile unit is authorized to receive CDPD service. In the simplest case, the MD-IS has direct access to a subscriber database for this authentication.

As a CDPD subscriber moves from one local AMPS sector to another, the M-ES scans for and moves to new channels to maintain service. The serving MD-IS tracks these handoffs. When the CDPD subscriber moves outside his or her home service region, MD-ISs from other service providers may be enlisted to maintain CDPD service. An M-ES has a *fixed* network address (typically an IP address), but it may receive service from any CDPD service provider that operates with the subscriber's home system, as shown in **Figure 3**.

The home MD-IS of an M-ES manages roaming by authenticating the subscriber and sending forward data to the MD-IS where the subscriber is currently receiving CDPD service. The serving MD-IS provides the air link, collects detailed accounting data, and operates with the subscriber's home MD-IS.

## Circuit-Switched CDPD

During the early stages of CDPD deployment, some AMPS cell sites may not be equipped with MDBSs. In addition, CDPD's usage-based accounting may not be cost-effective for applications exchanging large amounts of data. Both situations are addressed by introducing circuit-switched CDPD (CS-CDPD).[5] Although CDPD's air link is not used, its mobility model and subscriber management are. CS-CDPD works using a dedicated connection between the subscriber and service provider. A CS-CDPD session, shown in **Figure 4**, requires a circuit-switched connection that includes, for example, a cellular data call using an AMPS-specific modem, a land-line public switched telephone network (PSTN) call using a conventional 28.8-kb/s modem, or an integrated services digital network (ISDN) data link.

Today, CDPD service is available in most metropolitan areas. When subscribers travel to an area without CDPD coverage, they can use the same applications, the same IP network address, and typically the same modem to establish a CS-CDPD cellular connection.

Although CS-CDPD shares much with CDPD, such as accounting, encrypted subscriber traffic, and

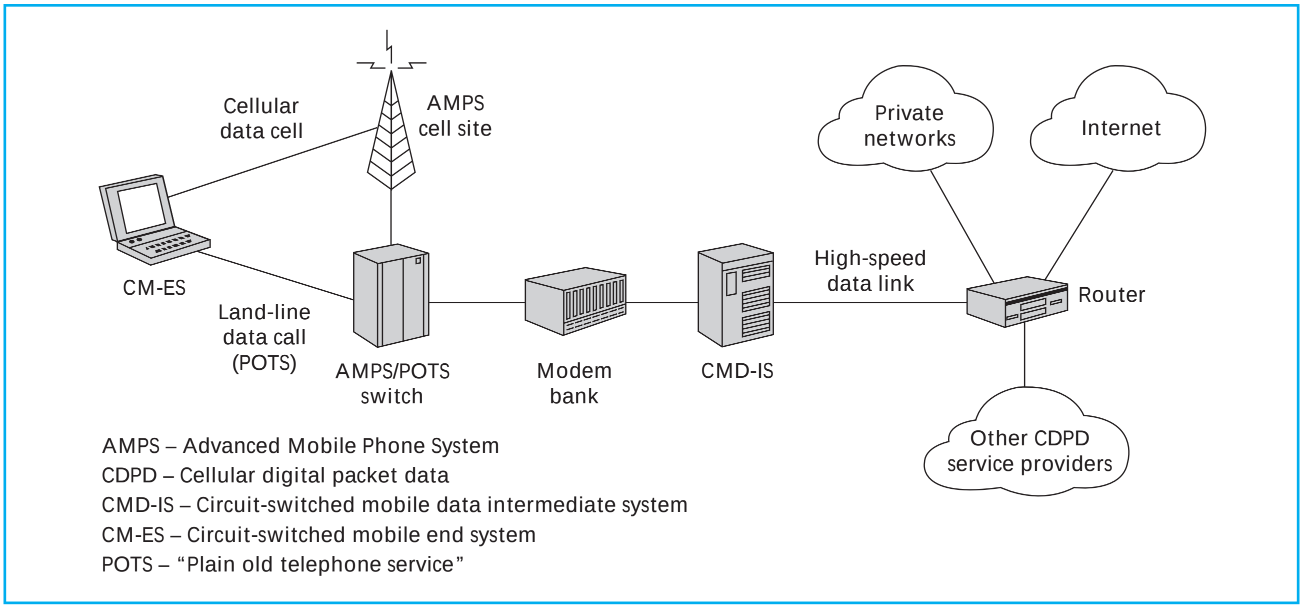


***Figure 4.***
***Circuit-switched CDPD network architecture.***

reliability, there are some differences. The circuit-switched MD-IS (CMD-IS) shares the same base protocols with CDPD, but instead of CDPD's MAC protocol, it uses a standard serial data transfer protocol over the circuit. A circuit-switched mobile end system (CM-ES) may either have an active connection to the CMD-IS or may be in a suspended state. When the CMD-IS serves a suspended CM-ES, it may call the mobile unit to reestablish a connection if it receives traffic destined for that CM-ES. Alternatively, the CM-ES may initiate the reconnection. CS-CDPD relies on the underlying cellular (or wired) network to handle CM-ES mobility within a service region and CDPD's support of roaming when the CM-ES moves between service areas.

## The CDPD Protocol Stack

The air link is the most valuable resource in narrow bandwidth wireless data networks. As such, the CDPD protocol stack was designed for efficient use of air link bandwidth. **Figure 5** shows a high-level profile of the CDPD protocol stack, including the network layer, the subnetwork-dependent convergence protocol (SNDCP), the mobile data link protocol (MDLP), the digital sense multiple access with collision detection protocol (DSMA-CD), and the physical layer.[1]

### Network layer

CDPD networks support both Internet (IP) and OSI (CLNP) network layer protocols. M-ESs are assigned a unique network address by their CDPD network service provider. M-ESs are anchored at a home MD-IS to support roaming.

### Subnetwork-Dependent Convergence Protocol

Network layer packets can carry fairly long headers. CDPD's SNDCP layer compresses TCP/IP packets using Van Jacobsen[6] header compression. A similar technique compresses CLNP's verbose network layer packet header. As a result, standard 40-octet TCP/IP protocol headers are compressed to an average of 3 octets by CDPD's SNDCP layer. CLNP headers with 57 octets are replaced with a 1-octet compressed header. Packet payloads can be further compressed using CDPD's optional V.42bis data compression feature.

After the header and payload are compressed, packets are segmented into 128-byte-long protocol data units (PDUs), which are encrypted for transfer between the M-ES and MD-IS. Encryption keys are exchanged between the M-ES and MD-IS using the Diffie-Hellman public key encryption algorithm.[7] New keys are generated periodically.

### Mobile Data Link Protocol

CDPD's MDLP is similar to link access procedures for D (data) channel (LAP-D),[8,9] a popular link protocol. A number of modifications have been made to

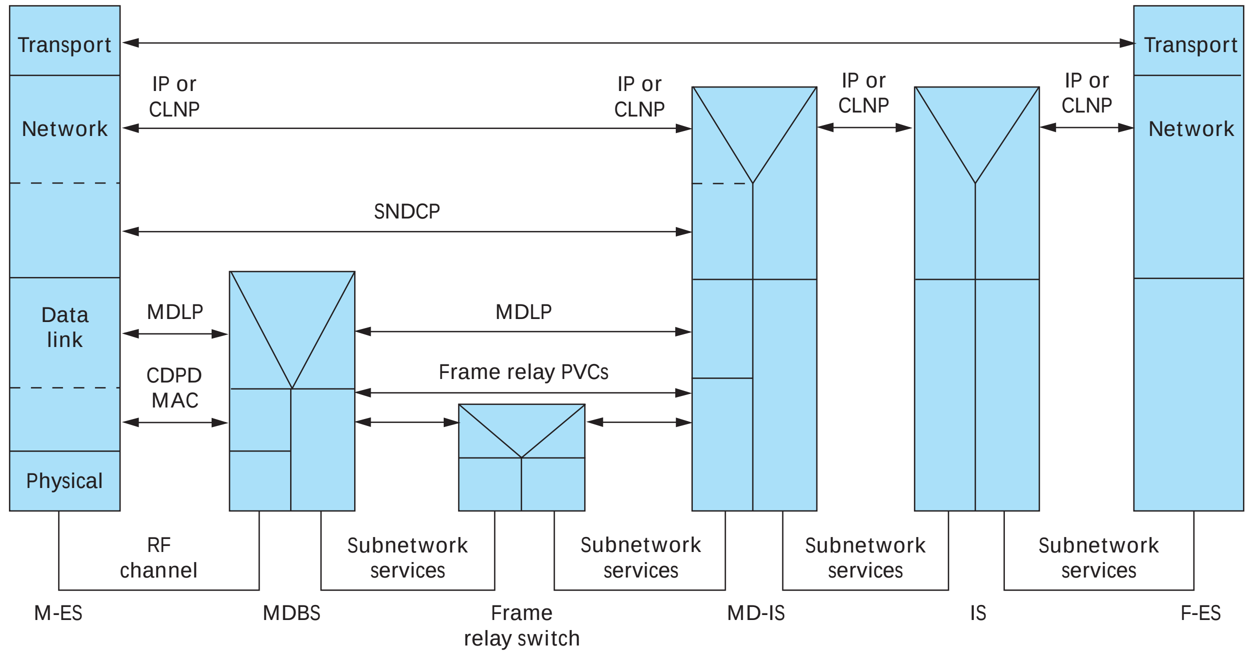


CDPD – Cellular digital packet data
CLNP – Connectionless network protocol
F-ES – Fixed end system
IP – Internet protocol
IS – Intermediate system
M-ES – Mobile end system
MAC – Medium access control
MDBS – Mobile data base station
MD-IS – Mobile data intermediate system
MDLP – Mobile data link protocol
PVC – Permanent virtual circuit
RF – Radio frequency
SNDCP – Subnetwork dependent convergence protocol



***Figure 5.***
***Profile of the CDPD protocol stack.***

tailor the link layer for the wireless environment, however. Unnecessary link layer retransmissions waste air link bandwidth. For bandwidth efficiency, MDLP selectively requests retransmission of lost packets using selective reject (SREJ) packets. MDLP also supports multicast and broadcast addressing over its unacknowledged data service. Only one copy of a link layer frame is sent over the forward link to all M-ESs in a multicast or broadcast group, saving air link bandwidth.

A major benefit of terminating MDLP at the MD-IS is efficient management of handoffs. By monitoring the MDLP addresses of link layer frames received over each CDPD channel stream it serves, an MD-IS passively detects M-ES handoffs and updates internal routing tables. MDLP frames that may have been lost during a handoff are retransmitted using MDLP's reliable data service.

### Digital Sense Multiple Access with Collision Detection Protocol

Access to the CDPD reverse link is governed by CDPD's DSMA/CD protocol, discussed in "Tuning CDPD's Reverse Link MAC Protocol." The protocol parameters share a number of similarities with the MAC protocol used by Ethernet. The DSMA/CD parameters are configurable, allowing CDPD service providers to tune the reverse link for desired performance.

### Physical layer

As an AMPS overlay network, CDPD uses the same 30-kHz channels as AMPS. The physical air link bit stream is transmitted using GMSK[2] at a raw data rate of 19.2 kb/s. Data sent over the air link is protected using a (63, 47) Reed Solomon code[3] and is transferred in a series of physical layer blocks. The

**Table II. Accumulation of overhead in the CDPD protocol stack.**

| Packet length | Bytes | Cumulative bytes for 1,000 bytes of network layer data |
|---|---|---|
| After header and data compression | $N + h$ | 1,000 |
| After segmentation | $\left\lceil \frac{N+h}{128} \right\rceil + (N+h)$ | 1,008 |
| After framing | $8\left\lceil \frac{N+h}{128} \right\rceil + (N+h)$ | 1,064 |
| After blocking | Forward channel: $\left(8\left\lceil \frac{N+h}{128} \right\rceil + (N+h)\right)\frac{420}{282}$ Reverse channel: $\left(8\left\lceil \frac{N+h}{128} \right\rceil + (N+h)\right)\frac{385}{282}$ | Forward channel: 1,585 Reverse channel: 1,453 |

coding scheme is well-suited to combating the burst errors common on the air link that result from Rayleigh fading and other channel impairments.

**Table II** shows the amount of overhead added by the CDPD protocol stack for a network layer packet containing a compressed header and payload a total of ($N$+$h$) bytes long. As Table II shows, nearly one-third of the 19.2 kb/s throughput of the air link physical layer is spent on error correction coding redundancy and, in the case of the forward channel, the addition of channel bits and collision feedback bits. Packets sent over the forward channel carry more overhead than reverse packets because of the in-band control bits sent over the forward channel to mediate access to the reverse channel.

Typical maximum network layer throughputs without V.42bis enabled are roughly 13.2 kb/s (reverse) and 12.1 kb/s (forward). Headers of user datagram protocol (UDP) and Internet control message protocol (ICMP) packets are not compressed by CDPD's SNDCP layer. Packet payloads that have been either encrypted or already compressed by an application may not be further compressed by SNDCP's V.42bis algorithm. Higher maximum throughputs are attainable when packet headers and payloads are readily compressible.

## Channel Hopping

AMPS networks with three-sectored cells and a reuse factor of seven—a popular configuration for AMPS networks in North America and other parts of the world—are typically equipped with 10 to 25 channels per sector. To offer tolerable call-blocking probabilities, the channels in these sectors must be used with moderation.

When operating in the MDBS channel-hopping mode, described earlier, an MDBS "borrows" an idle AMPS channel to transfer data to and from M-ESs. If an AMPS call starts using the channel being borrowed by the MDBS, the MDBS quickly moves the affected channel stream to another idle AMPS channel. If no channel is available, the CDPD channel "blacks out" until an AMPS channel becomes idle. To help M-ESs relocate the CDPD channel after a channel hop, the MDBS periodically broadcasts messages to inform M-ESs of AMPS channels that are likely to be used after preemption by an AMPS call. If data applications and users can tolerate occasional blackouts, the CDPD channel-hopping feature can squeeze additional revenue from idle AMPS channels for AMPS service providers.

Operating in this parasitic mode does not guarantee that an MDBS will be able to find an idle AMPS channel. In sectors with heavy AMPS call loads, channel hops and blackouts can be frequent, adversely affecting CDPD network performance. As a result, in some cells channel hopping may not offer adequate performance for delay-sensitive applications.

An M/G/c/c queuing system[10] can analytically quantify idle AMPS capacity and determine how AMPS call loads, AMPS holding times, and the number of AMPS channels per sector influence the CDPD

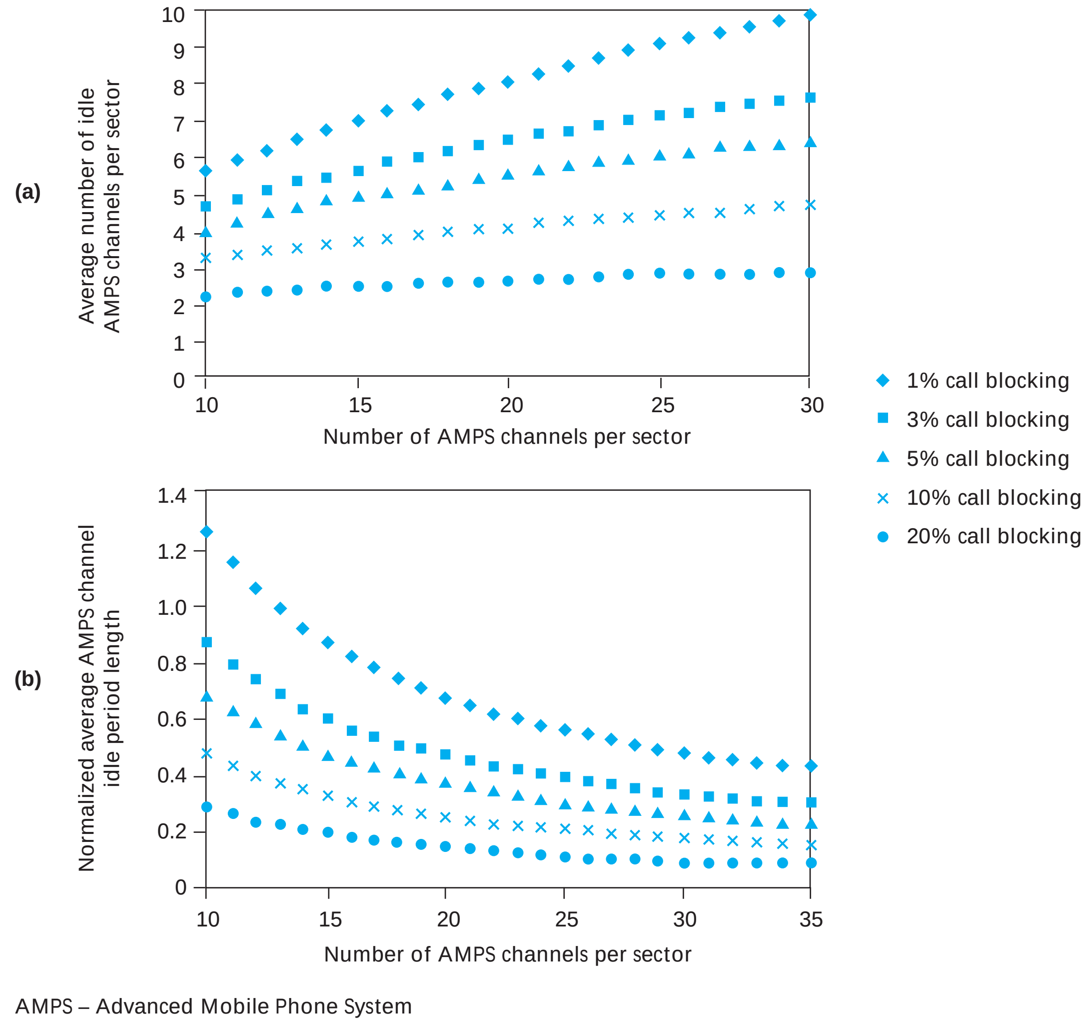


***Figure 6.***
***(a) Average number of idle AMPS channels per sector as a function of AMPS channels per sector. (b) Average duration of AMPS channel idle periods.***

channel-hopping feature. The model can also help determine when it is appropriate to use the CDPD channel-hopping feature.[11]

Assume that in a sector with $c$ AMPS channels, AMPS calls—that is, calls that originate in the sector and calls that are handed off to a sector—are generated in accordance with a Poisson process with normalized rate $a$ calls/unit. The length of time an AMPS call holds onto an AMPS channel is assumed to be generally distributed with unit mean. If we use arguments from renewal theory, the number of idle channels per sector, $n_{idle}$, can be expressed as

$$n_{idle} = c - a(1 - B(c,a)), \tag{1}$$

where $B(c,a)$ denotes the Erlang blocking formula for an M/G/c/c system with offered load $a$

$$B(c,a) = \frac{a^c}{\sum_{j=0}^{c}\left(\frac{a^j}{j!}\right)}. \tag{2}$$

Under the assumption that each channel receives an equal fraction of the AMPS call load, we may also calculate $T$, the average length of time an AMPS channel is idle, where

$$T = \frac{c - a(1 - B(c,a))}{a(1 - B(c,a))}. \tag{3}$$

**Figure 6** shows the average number of idle AMPS channels per sector (Figure 6a) and average idle period

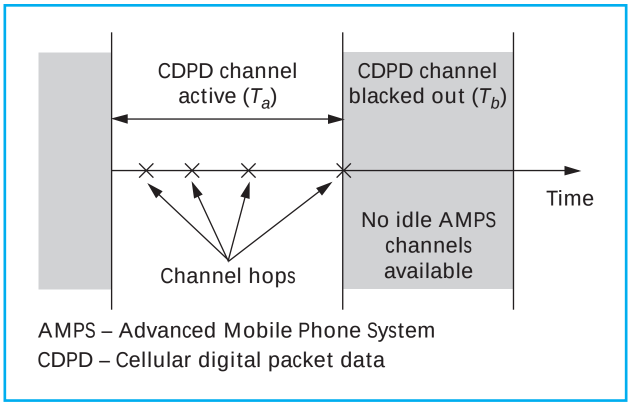


*Figure 7.*
***CDPD channel active and blackout periods.***

per channel (Figure 6b) for a variety of systems of interest. Normalization assumes AMPS channel holding time of one unit. At AMPS call blocking rates of 3 to 5%, which are typical call blocking rates for mature networks, a fairly large number of AMPS channels are idle, on average. Furthermore, at a typical AMPS channel holding time of 90 seconds, AMPS idle periods are fairly long, implying that a parasitic data network should have ample time to detect and use idle AMPS channels.

Using the CDPD channel-hopping feature does not come without cost. As depicted in **Figure 7**, periods of time may occur when all AMPS channels are occupied by AMPS calls and the CDPD channel will not be available. If these periods are short and infrequent, some data applications may be able to tolerate the disruptions.

With the further assumption that AMPS channel-holding times are exponentially distributed, the Laplace transforms of the distributions of active and blackout periods can be calculated.[12,13] For the case of an AMPS sector equipped with one CDPD channel stream, the mean length of the active periods $T_a$ can be expressed as

$$T_a = \frac{1 - B(c, a)}{c\, B(c,a)}. \qquad (4)$$

The mean length of blackout periods $T_b$ for the case of one CDPD channel stream per sector is independent of the AMPS call load and can be expressed as

$$T_b = \frac{1}{c}. \qquad (5)$$

**Figure 8** shows the mean length of active (Figure 8a) and blackout (Figure 8b) periods for several systems of interest for sectors equipped with one CDPD channel. (The length of a blackout period in this configuration is independent of AMPS call loads.)

Note from Figure 8a that as AMPS call blocking periods increase, active period lengths shorten. As a result, channel streams spend a greater fraction of time "blacked out." Determining whether the CDPD channel-hopping feature provides adequate service depends on the delay sensitivity of the applications running on a CDPD network and the rates charged for CDPD service. For AMPS blocking rates of 5% or less in AMPS sectors with fewer than 25 channels, the CDPD channel-hopping feature may be a viable alternative to an AMPS channel dedicated to CDPD. At higher blocking rates or when applications cannot tolerate periodic blackouts, CDPD must be deployed on dedicated channels.

## Tuning the CDPD Reverse Link MAC Protocol

The delay and throughput performance of the CDPD air link will dominate the performance observed by most applications running on a CDPD network. The air link will likely be the lowest throughput leg on a packet's journey between an M-ES and the wireline network. In addition, link layer retransmissions between an M-ES and an MD-IS may be needed to recover from air link transmission errors, adding further delay.

M-ESs also experience air link delay as they wait to transmit packets over the reverse air link. To control access to the reverse air link, CDPD uses a MAC protocol similar to the one employed by Ethernet. CDPD's reverse link MAC protocol has a number of tunable parameters that strongly influence the throughput and delay performance of the reverse link. Proper tuning of the reverse link is important because, during periods of reverse air link congestion, small, delay-sensitive control messages, such as those sent by an M-ES during registration, must still reach the MD-IS with relatively small delay.

As **Figure 9** shows, the reverse air link is a slotted channel consisting of a series of microslots, each 60 bits long. The forward channel carries a stream of con-

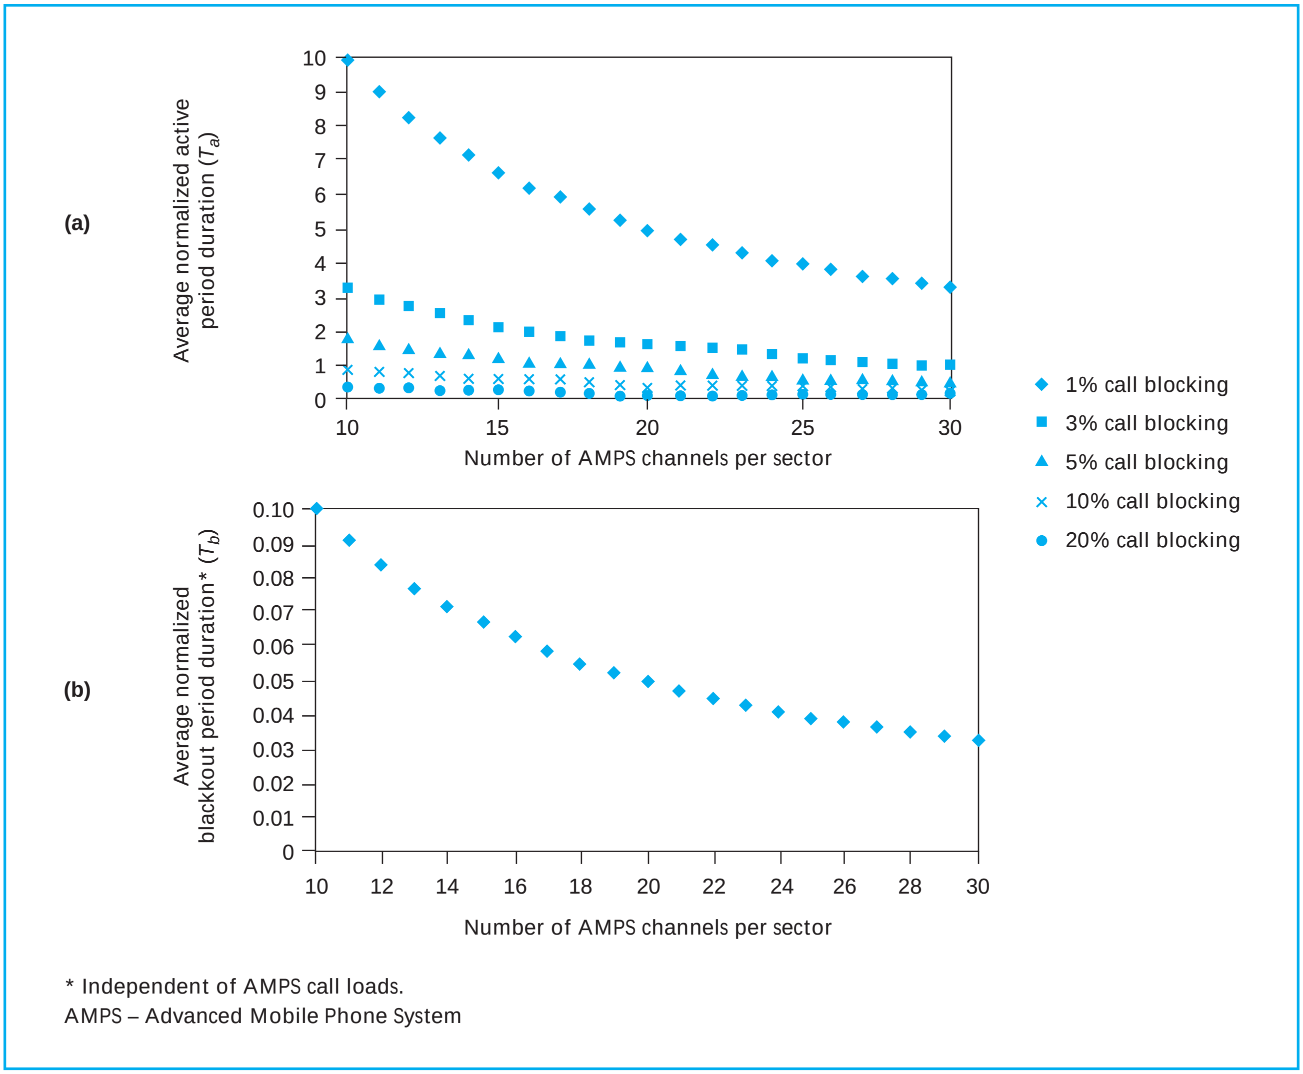


***Figure 8a.***
***Average normalized active period duration ($T_a$) for sectors equipped with one CDPD channel.***
***Figure 8b.***
***Average CDPD channel blackout period duration ($T_b$) for sectors equipped with one CDPD channel.***

trol bits that inform M-ESs of the busy/idle status of each reverse air link microslot, as well as a regular bit pattern that allows an M-ES to easily determine microslot boundaries.

M-ESs listen for idle microslots during a series of transmission attempts. If an M-ES determines that the microslot is idle, it transfers data to the MDBS in short bursts, as shown in **Figure 10**. Each reverse channel burst begins with a dotting sequence, which helps the MDBS easily detect reverse link transmissions, followed by a reverse channel synchronization word. After transmission of this 60-bit preamble, the M-ES transmits an integral number of physical layer blocks. Each reverse link block is 385 bits long and carries roughly 32 bytes of uncompressed network layer data. As soon as the MDBS detects a dotting sequence on the reverse air link, it sets the reverse channel busy/idle status flags to busy and signals other M-ESs to refrain from sending data over the reverse link while the burst is being transmitted. Block decode status bits carried over the forward channel acknowledge each physical layer block that the MDBS receives and correctly decodes. If the MDBS determines that the first block in a burst is in error, it assumes that the

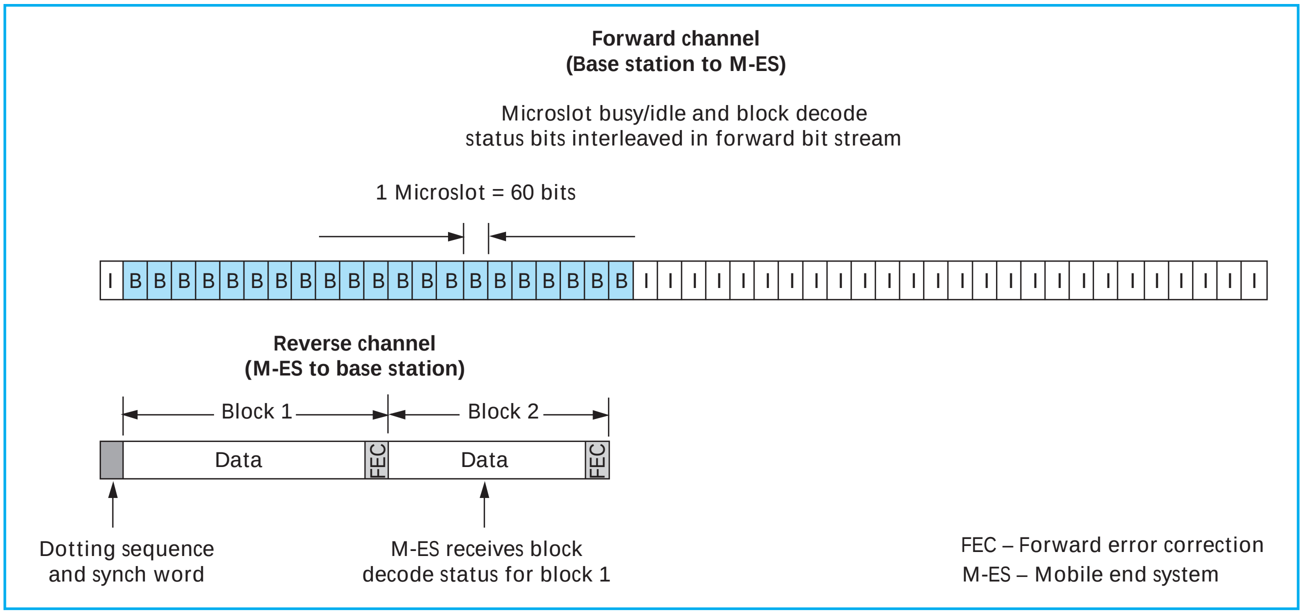


***Figure 9.***
***Slotted structure of the reverse air link.***

error is due to a collision and informs the M-ESs to stop transmitting their bursts and to reschedule their transmission attempts.

Full-duplex M-ESs are capable of sending more than one block in a reverse link burst. The reverse link MAC parameter `MAX_BLOCKS` controls the maximum number of blocks that a full-duplex M-ES may send in a single burst. If the M-ES receives any block in the burst in error, it halts transmission and retransmits the last link layer frame it believes the MDBS has not received. Half-duplex M-ESs, however, can only send one block at a time; as soon as they send a block over the reverse air link, they quickly tune to the forward air link to determine whether the block was correctly received at the MDBS.

Because half-duplex M-ESs send only one block per transmission attempt, they experience longer packet transmission delays (and lower achievable throughputs) than full-duplex M-ESs. **Figure 11** shows average round-trip packet transmission delays for a series of ICMP echo (“ping”) messages[14] sent from full-duplex and half-duplex M-ESs to a fixed end system (F-ES) on an otherwise unloaded, dedicated CDPD channel stream with no V.42bis compression. MAC parameters used for the measurements were `MAX_BLOCKS=10`, `MIN_IDLE_TIME=46`, `MAX_ENTRANCE_DELAY=66`, `MAX_TX_ATTEMPTS=30`, `MIN_COUNT=7`, and `MAX_COUNT=9`. The field measurements were collected using the network configuration shown in **Figure 12**. CDPD’s V.42bis data compression feature was disabled for the field measurements.

After an M-ES has sent a burst, the CDPD reverse link MAC protocol forces the M-ES to relinquish the channel for a period of time. The delay between bursts is controlled by the MAC parameter `MIN_IDLE_TIME`, which specifies the minimum number of microslots an M-ES must wait from the end of one burst to the start of the next. The purpose of both MAC parameters `MIN_IDLE_TIME` and `MAX_BLOCKS` is to allow CDPD service providers to prevent M-ESs that are sending large amounts of data from “hogging” the air link. Forcing an M-ES to relinquish the reverse channel between reverse air link bursts gives other M-ESs an opportunity to vie for use of the reverse channel. This preventive measure decreases the maximum network layer throughput that an individual full-duplex M-ES can achieve. Limiting an individual user’s maximum throughput is beneficial for the “common good” of all M-ESs sharing a CDPD channel. In this way, even M-ESs with small packets to send have an equal chance of seizing the reverse air link.

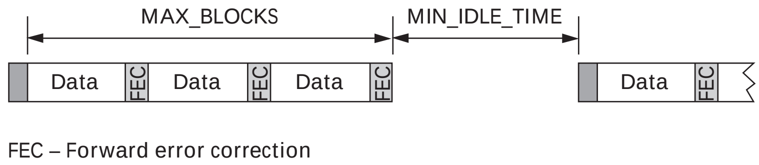


***Figure 10.***
***Influence of the reverse link MAC parameters* MAX_BLOCKS *and* MIN_IDLE_TIME *on reverse link transmissions for full-duplex M-ESs.***

The maximum average reverse link throughput that a single full-duplex M-ES can achieve at its IP or CNLP network layer without V.42bis data compression is simply the maximum network layer throughput of the reverse air link (13.2 kb/s) multiplied by the fraction of time the M-ES is using the reverse air link to send network layer data. Hence, the maximum reverse channel throughput that a full-duplex M-ES can achieve (assuming link or transport layer flow control does not constrain throughput) is related to the parameters MAX_BLOCKS and MIN_IDLE_TIME, with $A$ as the maximum reverse air link network layer throughput for a full-duplex M-ES (without V.42bis and SNDCP header compression), as follows:

$$A \approx 13.2\left(\frac{MAX\_BLOCKS \cdot 385}{MAX\_BLOCKS \cdot 385 + 60 \cdot (MIN\_IDLE\_TIME + 1)}\right) \text{kb/s.} \quad (6)$$

If the M-ES senses the channel busy during a transmission attempt, it waits a random number of microslots before it initiates another transmission attempt. If the M-ES has not been involved in any collisions since it began attempting to send a burst, it selects a delay between transmission attempts that is uniformly distributed between 0 and the MAC tunable parameter MAX_ENTRANCE_DELAY microslots.

The CDPD reverse link MAC protocol incorporates an exponential backoff algorithm to protect the link in the event of overload. If at any point an M-ES is involved in a collision, the M-ES sets an internal counter named count to the value of the MAC parameter MIN_COUNT. The delay selected by an M-ES for subsequent retransmission attempts is uniformly distributed between 0 and $2^{count}$ microslots. After the initial collision, each time the M-ES either finds the reverse channel busy or experiences another collision, the M-ES increments count by one to at most MAX_COUNT, another MAC parameter.

To help tune the reverse link MAC protocol, a simulation model of M-ESs accessing the CDPD reverse air link on a dedicated channel was constructed using Lucent's discrete event simulation package *Q*+.[15,16] Each simulated M-ES used the MAC protocol described above to access the reverse link. The tool simulated a variety of load scenarios, from M-ESs performing bulk data transfers to periodically sending short messages. It also measured collision rates, message delay statistics, throughputs, and other key performance parameters under a variety of loading scenarios. The simulation model was used to determine a good candidate set of values for MAC parameters.

Field experiments were conducted to validate the simulation model and to observe reverse link MAC performance in realistic loading scenarios. In one set of experiments, several full-duplex M-ESs were registered on a common, dedicated CDPD channel stream. All M-ESs continuously transferred large volumes of data to an F-ES using the network configuration shown in Figure 12. During the transfers, a CDPD air link analyzer captured the flow of link layer frames from M-ESs to the MDBS. The CDPD air link analyzer decoded, time-stamped, and recorded each link layer frame transmitted during the experiments for later processing.

The time-stamped link layer frame traces yield a wealth of data on the performance of the CDPD reverse link MAC protocol. Statistics of interest include the network layer throughputs achieved by each M-ES during an experimental run and net throughput of the reverse air link. Observing the link layer traces enables users to record the time epochs at which M-ESs seize and relinquish the reverse air link. The delay from the time an M-ES seizes a channel until the next time it seizes the same channel ($T_s$) is a measure of how well

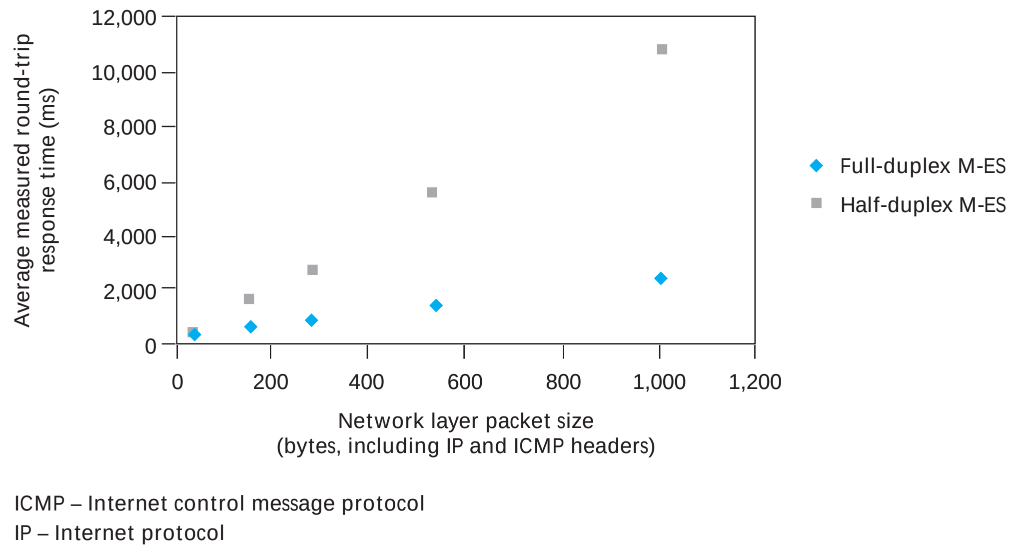


***Figure 11.***
***Average round-trip response times for ICMP echo messages sent to an F-ES over an otherwise unloaded, dedicated CDPD channel stream.***

the reverse air link is shared among M-ESs competing for reverse air link bandwidth.

Table III shows the results of one series of experiments in which seven full-duplex M-ESs were used to load the reverse air link at different values of `MAX_BLOCKS`. The results show the impact of `MAX_BLOCKS` on the mean, median, and 90$^{th}$ percentiles of $T_s$, as well as the air link throughput observed by each M-ES. Clearly, as `MAX_BLOCKS` increases, M-ESs must wait longer to seize the reverse air link. Keeping mean and tails of the reverse channel relinquish and seize delay $T_s$ small ensures that delay-sensitive control messages can be carried over the reverse link in times of congestion. Large values for the 90$^{th}$ percentiles of $T_s$ relative to the mean also imply that data transfer with large values of `MAX_BLOCKS` will appear "jerky" to users running interactive applications.

The experimental data shown in Table III illustrates some of the tradeoffs faced when fine-tuning the CDPD reverse link MAC protocol. Small values of `MAX_BLOCKS` decrease the length of time an M-ES must wait before transmitting data over the reverse air link. Setting `MAX_BLOCKS` too small, however, creates excessive segmentation of network layer packets and less efficient use of the reverse air link under heavy loads.[17] In the interest of keeping reverse link transmission delays small for users who periodically send small packets, it is important to set `MIN_IDLE_TIME` long enough so that M-ESs transmitting bulk data do not dominate the reverse air link. As shown earlier in (6), forcing an M-ES transferring bulk data to relinquish the reverse link decreases maximum throughputs, even when the M-ES is the only M-ES transferring data over the reverse link.

## CDPD Applications

As competition in the wireless data market continues, prices for CDPD service have dropped significantly. In contrast to connection-time-based pricing schemes common in cellular voice services, current CDPD pricing schemes are based on a fixed monthly fee, plus a charge related to the volume of network layer data sent and received by a user over a CDPD network. This type of charge encourages efficient use of the network. Choosing the appropriate transport layer protocol (for example, TCP vs. UDP), tuning protocol parameters, and compressing data before sending it over a CDPD network could help users reduce their service charges.

CDPD supports IP, allowing most applications using TCP or UDP to run on CDPD networks without

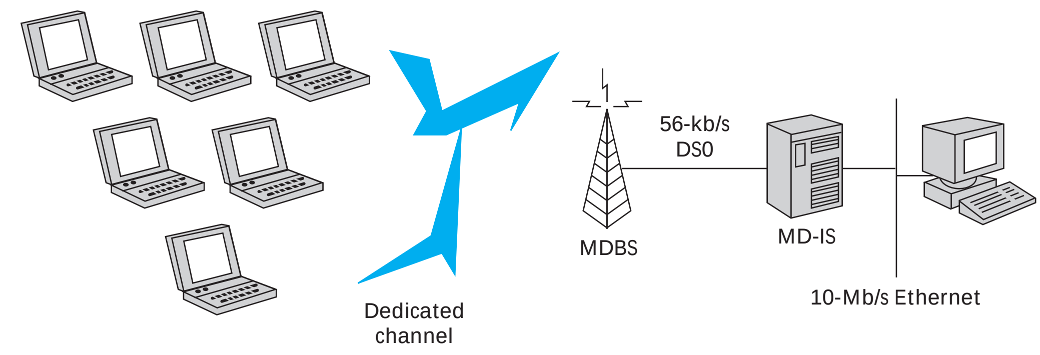


***Figure 12.***
***Network configuration used for field experiments.***

modifications. TCP is a connection-oriented protocol that detects lost, misordered, or duplicate packets and assures reliable end-to-end communications by detecting and retransmitting lost packets. UDP, a simple unacknowledged protocol, relies on the application layer to achieve reliable communications. Because of low overhead and cost, UDP is suitable for applications that have short queries and responses. On the other hand, TCP is suitable for applications that need to send a large amount of data reliably.

Parameter tuning and slight implementation modifications can dramatically enhance the delay and throughput performance of standard TCP over CDPD and, at the same time, reduce users' service charges. Many off-the-shelf TCP stacks assume wired transmission lines as the underlying communication media. In addition, many TCP parameters, such as retransmission timers and sliding window size, are set for robust performance for a variety of applications over high-speed, relatively error-free wired lines. If a packet is lost, TCP assumes that network congestion is the culprit. Although most TCP congestion control mechanisms are designed to work well under this assumption, it is invalid for most wireless applications.

Air link interference and signal fading can cause errored packets to be dropped. Owing to the discontinuity in end-to-end communications, factors such as channel hopping, channel blackouts, and mobility can all cause packet loss, and transmission time-outs can result in a high variation in response time. Several papers have discussed how to adjust TCP parameters and apply congestion control methods over the slow and unreliable CDPD links with little or no modification in the TCP/IP protocol stack.[18,19] Commercial TCP/IP stacks optimized for CDPD are also available.

Current usage-based pricing schemes and the CDPD air link make CDPD well-suited for applications with short, bursty messages. For bandwidth-intensive applications, circuit-switched CDPD might be a better alternative. CDPD provides significant economic value in vertical market segments such as public safety, public utility, retail, and transportation. CDPD enables police to access motor vehicle files and other criminal data sites within a few seconds from a terminal anywhere on the road. Because of its low cost and quick response times, CDPD has been very useful for point-of-sale transactions, especially in areas where it is uneconomical to wire telephone lines, such as food stands in stadiums or exhibition booths in convention centers. Combined with the Global Positioning System (GPS), CDPD provides a powerful capability to fleet management companies by locating and dispatching vehicles in transit. In telemetry, remotely located stationary machines, such as gas meters or vending

Table III. Influence of the MAC parameter MAX_BLOCKS on the sharing of reverse link bandwidth when several M-ESs are simultaneously transmitting bulk data over the reverse airlink for a period of 15 minutes.

| MAX_BLOCKS | Performance metric | M-ES$_1$ | M-ES$_2$ | M-ES$_3$ | M-ES$_4$ | M-ES$_5$ | M-ES$_6$ | M-ES$_7$ |
|---|---|---|---|---|---|---|---|---|
| 5 | throughput (bytes/sec) | 178.56 | 157.61 | 73.26 | 170.86 | 139.11 | 146.14 | 134.40 |
| | mean $T_S$ (sec) | 0.88 | 0.98 | 1.69 | 0.87 | 0.99 | 1.08 | 1.08 |
| | median $T_S$ (sec) | 0.55 | 0.57 | 0.64 | 0.55 | 0.59 | 0.56 | 0.56 |
| | 90$^{th}$ percentile $T_S$ (sec) | 1.64 | 1.77 | 4.49 | 1.52 | 1.79 | 1.83 | 1.83 |
| 10 | throughput (bytes/sec) | 194.77 | 227.45 | 129.63 | 167.64 | 207.01 | 123.62 | 161.66 |
| | mean $T_S$ (sec) | 1.64 | 1.37 | 2.29 | 1.78 | 1.49 | 2.27 | 1.66 |
| | median $T_S$ (sec) | 0.95 | 0.86 | 1.04 | 0.96 | 0.87 | 1.01 | 0.93 |
| | 90$^{th}$ percentile $T_S$ (sec) | 3.11 | 2.44 | 4.06 | 3.38 | 3.10 | 5.95 | 2.46 |
| 20 | throughput (bytes/sec) | 127.26 | 172.17 | 167.38 | 192.70 | 227.71 | 245.46 | 132.39 |
| | mean $T_S$ (sec) | 4.01 | 3.32 | 3.33 | 3.02 | 2.89 | 2.57 | 3.86 |
| | median $T_S$ (sec) | 1.85 | 1.76 | 1.68 | 1.50 | 1.66 | 1.65 | 1.94 |
| | 90$^{th}$ percentile $T_S$ (sec) | 8.49 | 7.70 | 8.27 | 6.24 | 6.43 | 4.65 | 9.60 |
| 30 | throughput (bytes/sec) | 203.50 | 233.57 | 147.90 | 183.27 | 238.70 | 172.34 | 164.20 |
| | mean $T_S$ (sec) | 4.53 | 3.65 | 5.12 | 3.92 | 3.65 | 4.33 | 4.75 |
| | median $T_S$ (sec) | 2.61 | 2.18 | 2.53 | 2.13 | 2.13 | 2.18 | 2.32 |
| | 90$^{th}$ percentile $T_S$ (sec) | 9.23 | 7.53 | 12.48 | 8.02 | 9.33 | 11.04 | 10.42 |

MAX_ENTRANCE_DELAY = 66 MIN_COUNT = 7
MIN_IDLE_TIME = 66 MAX_COUNT = 9
MAX_TX_ATTEMPTS = 30

machines, periodically report their status to a centralized location. This practice avoids sending maintenance personnel to the field to collect data and prevents the depletion of stock in vending machines.

In horizontal markets, wireless electronic mail has been one of the applications driving the demand for wireless data. Recently, browsers using the handheld device markup language (HDML) have been incorporated in some cellular phones to allow cellular users to browse the Internet. The HDML-based browsers are optimized for limited memory and wireless capacities by filtering out bandwidth-intensive graphics on the Web pages and efficiently using the cellular phone's compact display size. This technology could stimulate demand for wireless Web applications and wireless data services.

The types of CDPD applications used by subscribers have a significant impact on the CDPD system performance and capacity. Different types of applications consume more of one CDPD resource and put more stress on one component of a CDPD system than others. For instance, some applications generate short, bursty transactions but stay registered for a long time, while others have long messages to transmit and only stay registered during transmission. Some applications have little or no mobility, such as telemetry, while others have high mobility, such as transportation.

MD-IS memory is required to maintain detailed state information on every registered M-ES. M-ESs that remain registered for long periods of time tend to consume more MD-IS memory on average than do users who register for short periods of time. In addition, M-ESs that perform frequent registrations and deregistrations tend to require more CDPD network central processing unit (CPU) resources, because registrations involve CPU-intensive processes such as authentication and encryption key generation. Those M-ESs that tend to transfer large volumes of data while remaining registered for short periods will probably generate more revenue with a smaller amount of system resources. Knowing the mix of application types helps a system architect balance resource allocation in the CDPD system and create a flexible system architecture that can be configured to adapt to changes in the application mix.

## Future Wireless Data Networks

As this paper is being published, the CDPD specification is being transferred from the auspices of the CDPD Forum to the Telecommunications Industry Association (TIA). This move does not signify the end of CDPD standards work, however. Code division multiple access (CDMA) and time division multiple access (TDMA) groups are actively involved in defining wireless packet data standards. While each air link is fundamentally different from the CDPD air link, other aspects of the CDPD architecture being considered as a basis for these digital networks are:

- Support of standard applications, in which no application-specific "gateways" would be required;
- Mobility model (that is, CDPD's home and serving MD-IS and the interservice provider interface);
- Subscriber management; and
- Accounting.

By being the first wide area wireless data with open interfaces to enter the wireless services market, CDPD will play a key role in stimulating demand for future wireless data networks.

### Acknowledgments

The authors acknowledge the contributions of their colleagues Stan Stachowiak and Thad Kobylarz of Lucent Technologies for conducting the field experiments described in "Tuning the CDPD Reverse Link MAC Protocol." For the use of its CDPD network to run these experiments, the authors thank Bell Atlantic NYNEX Mobile.

### References

1. *Cellular Digital Packet Data System Specification: Release 1.1,* CDPD Forum, Inc., Chicago, 1995.
2. K. Murota and K. Hirade, "GMSK Modulation for Digital Mobile Radio Telephony," *IEEE Trans. Communications,* Vol. 29, No. 7, July 1981, pp. 1044-1050.
3. I. S. Reed and G. Solomon, "Polynomial Codes over Certain Finite Fields," *J. Soc. Ind. Appl. Math.*, Vol. 8, June 1960, pp. 300-304.
4. S. Deering, and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification," *Internet RFC 1883,* DDN Network Information Center, SRI, Ravenswood, Calif., Aug. 1996.
5. *Circuit Switched Cellular Digital Packet Data Implementor Guidelines, Release 1.5,* CDPD Forum Inc., Chicago, June 13, 1995.
6. V. Jacobsen, "Compression TCP/IP headers for low-speed serial links," *Internet RFC 1144,* DDN Network Information Center, SRI, Ravenswood, Calif., Feb. 1990.
7. W. Diffie and M. E. Hellman, "New Directions in Cryptography," *IEEE Trans. Information Theory,* Vol. 22, No. 6, Nov. 1976, pp. 644-654.
8. *ISDN User-Network Interface Data Link Layer—General Aspects,* International Telegraph and Telephone Consultative Committee Recommendation Q.920, 1988.
9. *ISDN User-Network Interface Data Link Layer Specification,* International Telegraph and Telephone Consultative Committee Recommendation Q.921, 1988.
10. L. Kleinrock, *Queueing Systems,* Wiley, New York, 1979.
11. K. Budka, "Cellular Digital Packet Data: Channel Availability," *IEEE Trans. Vehicular Technology,* Vol. 46, No. 1, Feb. 1997, pp. 31-40.
12. S. Halfin and M. Segal, "A priority queuing model for a mixture of two types of customers," *SIAM J. Appl. Math.*, Vol. 23, No. 3, Nov. 1972, pp. 369-379.
13. M. Segal, "A preemptive priority model with two classes of customers," *Proc. ACM/IEEE 2nd Symp. Problems Optimization Data Comm. Syst.,* Palo Alto, Calif., 1971, pp. 168-174.
14. W. R. Stevens, *TCP/IP Illustrated,* Vol. 1, Addison-Wesley, Reading, Mass., 1994.
15. C. Funka-Lea, B. Gotz, and O-C. Yue, "Interactive Visual Modeling for Performance Analysis and Simulation of Queueing Systems," *Intl. J. Computer Simulation,* Vol. 4, No. 2, Apr. 1994, pp. 135-154.
16. C. Funka-Lea, T. Kotogiorgos, R. Morris, and L. Rubin, "Interactive Visual Modeling for Performance," *IEEE Software,* Vol. 8, No. 5, Sept. 1991, pp. 58-68.
17. R. Rom and M. Sidi, *Multiple Access Protocols: Performance Analysis,* Springer-Verlag, New York, 1990.
18. M-C. Chuah, O-C. Yue, and A. DeSimone, "Performance of Two TCP Implementations in Mobile Computing Environments," *Globecom '95,* Nov. 13-17, 1995, Singapore, pp. 339-344.

19. H. Balakrishnan, S. Seshan, E. Amir, and R. Katz, “Improving TCP/IP Performance over Wireless Networks,” *Proc. 1st ACM Intl. Conf. Mobile Computing and Networking (Mobicom)*, Nov. 1995, pp. 2-11.

## Further Reading

– M. Taylor, W. Waung, and M. Banan, *Internetwork Mobility: The CDPD Approach*, Prentice-Hall, Englewood Cliffs, N.J., 1997.
– M. Sreetharan and R. Kumar, *Cellular Digital Packet Data*, Artech House, Boston, 1996.
– J. Agosta and T. Russel, *CDPD: Cellular Digital Packet Data Standards and Technology*, McGraw-Hill, New York, 1997.

***(Manuscript approved August 1997)***

*KENNETH C. BUDKA is a member of technical staff in the Performance Analysis Department at Bell Labs in Holmdel, New Jersey. He works on the performance analysis and design of communication systems and is part of Lucent’s Cellular Digital Packet Data (CDPD) Development Team. Mr. Budka received a B.S. in electrical engineering from Union College in Schenectady, New York, and M.S. and Ph.D. degrees in engineering science from Harvard University in Cambridge, Massachusetts.*

*HONG JIANG, who received M.S. and Ph.D. degrees in electrical engineering from Northwestern University in Evanston, Illinois, is a member of technical staff in the Performance Analysis Department at Bell Labs in Holmdel, New Jersey. She is working on the performance aspects of cellular digital packet data (CDPD) and conducting economic studies of architectural alternatives for next-generation cellular and personal communication services (PCS) products. Ms. Jiang is also designing cost-minimization algorithms for a wireless infrastructure design tool.*

*STEVEN E. SOMMARS, a distinguished member of technical staff in the CDPD Systems Engineering Group of Lucent’s Network Systems in Naperville, Illinois, works on the development of Lucent’s CDPD product and various wireless data technologies. His other interests include computer and communications security, system performance, and floating point approximations. Mr. Sommars received a Ph.D. in physics from the State University of New York at Stony Brook.* ◆