Packet Data over Cellular Networks: The CDPD Approach

Apostolis K. Salkintzis, The University of British Columbia

ABSTRACT Cellular digital packet data is a mobile packet data technology that operates on the spectrum assigned to a telephone cellular network, such as the Advanced Mobile Phone Service. This article undertakes a thorough survey of the CDPD radio interface and explores the main functional layers of this interface. Specifically, it extensively studies the physical layer, the data link layer, and the subnetwork-dependent convergence protocol, and explains their semantics and functional characteristics. Furthermore, it emphasizes several significant aspects such as the medium access procedure, the forward and reverse channel configurations, the data multiplexing scheme, and the channel hopping procedure.

Cellular digital packet data (CDPD) is a mobile data technology that permits subordinate packet data operation on the spectrum assigned to a telephone cellular network, such as the Advanced Mobile Phone Service (AMPS). It was first introduced by IBM as a packet-switching overlay to the existing analog cellular voice network and frequencies. Later, a CDPD System Specification [1] was formed by a consortium of cellular carriers including AirTouch, McCaw Cellular, Southwestern Bell Mobile Systems, NYNEX, Ameritech, GTE, Bell Atlantic Mobile, and Contel Cellular [2]. Now, CDPD technology is being deployed by a number of cellular companies in the United States, including Bell Atlantic, Ameritech, GTE, and McCaw Cellular, and related equipment is provided by a variety of manufacturers.

An industry association that handles the shaping of CDPD technology and supports the growth of the commercial marketplace is the Wireless Data Forum [3]. According to this forum, by the end of the third quarter of 1997, CDPD was available in 195 markets in the United States — 118 metropolitan statistical areas (MSAs), 41 rural statistical areas (RSAs), and 36 international markets — and was available to 53 percent of the U.S. population.

This article focuses on the wireless interface of CDPD and explores the main functional layers of this interface. Specifically, the next section provides a general outlook on CDPD, and explains its major network elements and their interaction and functionality. The following three sections concentrate on the wireless interface and outline the physical, medium access control, and logical link control layers, respectively. In this context, several significant aspects are studied such as the medium access protocol, the forward and reverse channel configurations, and the data link establishment procedures. The article then focuses on the subnetwork-dependent convergence protocol (SNDCP), and demonstrates its major characteristics and the services it provides. The channel hopping procedure is underlined, and finally, the last section summarizes our conclusions.

A CDPD OVERVIEW

The primary elements of a CDPD network are the end systems (ESs) and intermediate systems (ISs), as shown in Fig. 1. (In Internet terminology, the ESs are known as *hosts* and ISs as *routers*.) The ESs represent the actual physical and logical end nodes that exchange information, while the ISs represent the CDPD infrastructure elements that store, forward, and route the information.

There are two kinds of ESs: The mobile ES (M-ES), which is a device used by a subscriber to access the CDPD network over the wireless interface, and the fixed ES (F-ES), which is a common host, server, or gateway connected to the CDPD backbone and providing access to specific applications and data. By definition, the location of an F-ES is fixed, whereas the location of an M-ES may change.

Typically, an M-ES consists of a *mobile terminal* (personal computer, personal digital assistant, or other standard equipment) and a CDPD *radio modem*, which attaches to the mobile terminal and manages the radio link and protocols. Usually, the communication between the radio modem and the mobile terminal is supported by standard serial protocols, such as the Serial Line Internet Protocol (SLIP) or Point-to-Point Protocol (PPP).

On the other hand, there are two kinds of ISs: a generic IS, which is simply a router (in most cases, an Internet Protocol, IP, router) with no knowledge of CDPD and mobility issues, and a mobile data IS (MD-IS), which is a specialized IS that routes messages based on its knowledge of the current location of M-ESs. The MD-IS is a set of hardware components and software functions that provide switching, accounting, registration, authentication, encryption, and mobility management functions. The mobility management software allows the switching system to track the M-ESs regardless of their location in the network and allows the M-ESs to use a single network address. The CDPD mobility management software follows the mobile IP model [4] established by the Internet Engineering Task Force (IETF).

Besides the ESs and ISs, there is also another element, the mobile data base station (MDBS), which is analogous to a common cellular base station. An MDBS performs no networking functions, but rather relays data link information between a number of M-ESs and their serving MD-IS (it is a data link functional element). Furthermore, it performs radio resource management procedures, the most important being the hopping of the CDPD radio frequency (RF) channel in response to voice network activity (see [5] for details). In summary, the MDBS creates and manages the air interface between the M-ESs and the CDPD backbone under the constraints arising out of the underlying voice network.

The CDPD backbone provides connectionless transport services, also called *datagram* services. This means that the network individually routes packets based on the destination address each packet carries and the knowledge of the current network topology. For the routing of packets, CDPD supports both IP and the Connectionless Network Protocol (CLNP), which is an open systems interconnection (OSI) standard protocol.

0163-6804/99/$10.00 © 1999 IEEE

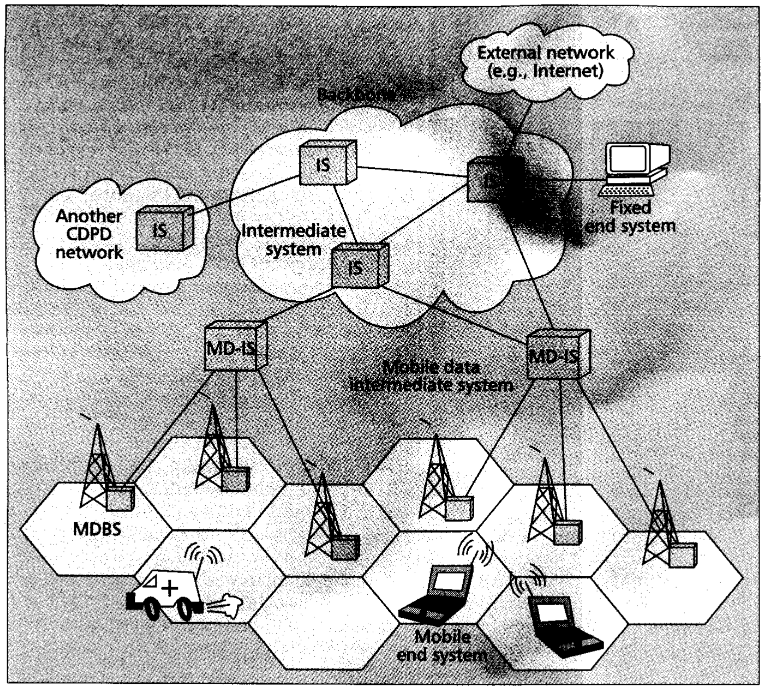

■ **Figure 1.** *General CDPD network architecture.*

THE PHYSICAL LAYER

As indicated in Fig. 2, the physical (PHY) layer in CDPD corresponds to a functional entity that accepts a sequence of bits from the medium access control (MAC) layer and transforms them into a modulated waveform for transmission onto a physical 30 kHz RF channel.

Communications between an MDBS and an M-ES take place over a pair of such RF channels (having a fixed frequency separation). The first channel, called the *forward* channel, accommodates transmissions in the direction from the MDBS to the M-ESs and is either dedicated to CDPD use or shared with the voice cellular network. In any case, transmission on the forward channel is continuous as long as it is in use for CDPD. The second channel, called the *reverse* channel, accommodates transmissions in the direction from the M-ESs to the MDBS and is shared among all M-ESs communicating with the same MDBS. A pair of associated reverse and forward channels forms a CDPD *channel stream*.

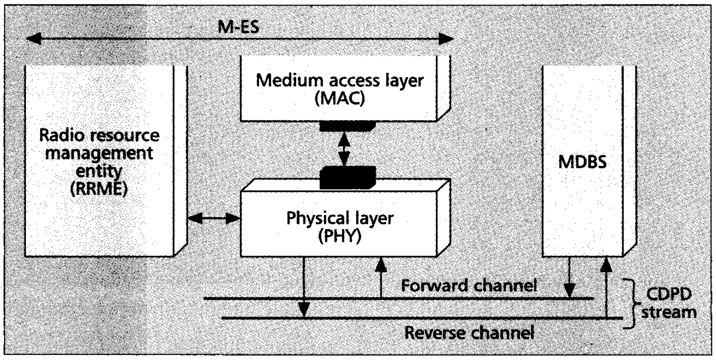

■ **Figure 2.** *The PHY layer model: operation and interaction with other functional entities.*

As illustrated in Fig. 2, the PHY layer interfaces with another entity, the radio resource management entity (RRME). Through this interface the RRME can:

- Tune the PHY layer to a specific RF channel pair
- Set the transmission power level to the desired value
- Measure the received signal level of an RF channel and estimate its potential to offer acceptable communication
- Suspend and resume operation of the PHY layer in cases where power saving facilities are required

The modulation employed on a RF channel stream is Gaussian minimum shift keying (GMSK) [6] with BT = 0.5. A frequency greater than the central carrier frequency represents a logical 1, whilst a logical 0 is represented by a frequency less than the central carrier frequency. The modulation rate on both the forward and reverse RF channels is 19.2 kb/s.

MEDIUM ACCESS CONTROL

As shown in Fig. 3, the MAC layer models a functional entity logically operating between the PHY and link layer control (LLC) layers. The MAC layer within an M-ES cooperates with the corresponding MAC layer within the MDBS. The purpose of this layer is to convey information, namely link protocol data units (LPDUs), between peer LLC entities across the CPDP air interface. For this purpose, the MAC layer provides the following services:

- Encapsulates LPDUs into frame structures to ensure LPDU delimiting, frame synchronization, and data transparency
- Encodes LPDUs to provide error protection against mobile channel impairments
- Detects and corrects bit errors within received frames
- Arbitrates access to the shared reverse channel
- Synchronizes with the forward channel transmissions to make feasible the reception of data as well as control information transmitted in every CDPD cell

The MAC layer communicates through an implementation-dependent interface with the RRME. Through that interface, the MAC layer notifies the RRME whether it has acquired synchronization with the currently selected forward channel (next section), and also passes to the RRME status information regarding the number of received bit and block errors. In this way, the RRME may estimate the acceptability of a given CDPD channel and provide the radio resource management functionality [5].

FORWARD CHANNEL CONFIGURATION

The forward channel characteristics are among the most essential issues required to understand CDPD operation. In order to

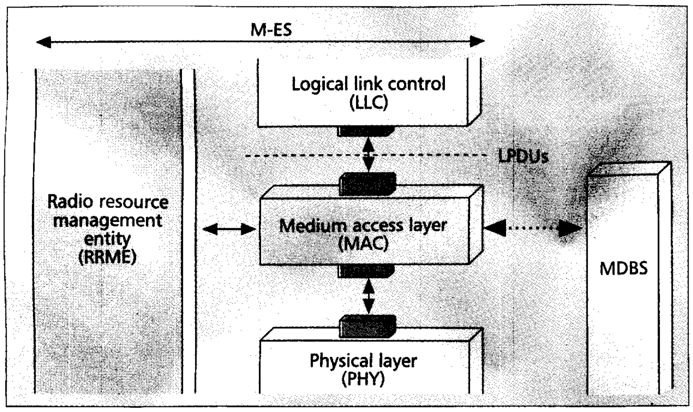

■ Figure 3. *An operational model of the medium access control layer.*

outline the configuration and the semantics of the forward channel, we will discuss how the LPDUs in an MDBS are transformed within a sequence of stages before they construct a continuous bitstream for transmission.

As shown in Fig. 4, the sequence of LPDUs pending for transmission are first flag-delimited (using the well-known bit pattern 01111110) and zero-stuffed, and then linked together to form a continuous frame data bitstream. The continuity of this bitstream is ensured because, even when there are no data LPDUs for transmission, either control LPDUs or sequences of contiguous flags are transmitted.

The frame data bitstream is divided into segments of 274 consecutive bits, each of which is prefixed by an 8-bit color code. Hence, a series of consecutive *data blocks* is formed, each with a fixed length of 282 bits. The color code is a special pattern assigned to every individual CDPD channel stream and is used for cochannel interference detection. Three bits within this pattern are MD-IS-specific; they have the same value in all channel streams transmitted in the set of cells controlled by a given MD-IS. On the other hand, the other five bits of the color code are MDBS-specific; they specify an individual channel stream within the set of cells controlled by a given MD-IS. Inside a cell, all RF channels available for CDPD use are assigned the same value of color code.

Data blocks are encoded using a systematic (63, 47) Reed-Solomon error correcting code. From an encoding point of view, each data block represents an information field of 47 6-bit symbols (or codewords). The encoding of this information field generates a 16-symbol parity field (96 bits), which is appended at the end of the information field. In this manner, a consecutive sequence of Reed-Solomon (RS) encoded blocks is generated, as shown in Fig. 4. These encoded blocks, each with a fixed length of 378 bits, form the basic transmission units of the forward channel. The (63, 47) Reed-Solomon encoding is common to both the forward and reverse channels and typically is capable of correcting as many as 8 bits within each encoded block.

Prior to actual transmission on the forward channel, each RS block is passed through a ninth-order scrambler with a generator polynomial, $g(x) = x^9 + x^8 + x^5 + x^4 + 1$. This process reduces the likelihood of having long strings of binary 1s and 0s within the transmission bitstream. Such long strings are generally avoided because they are difficult to track by certain

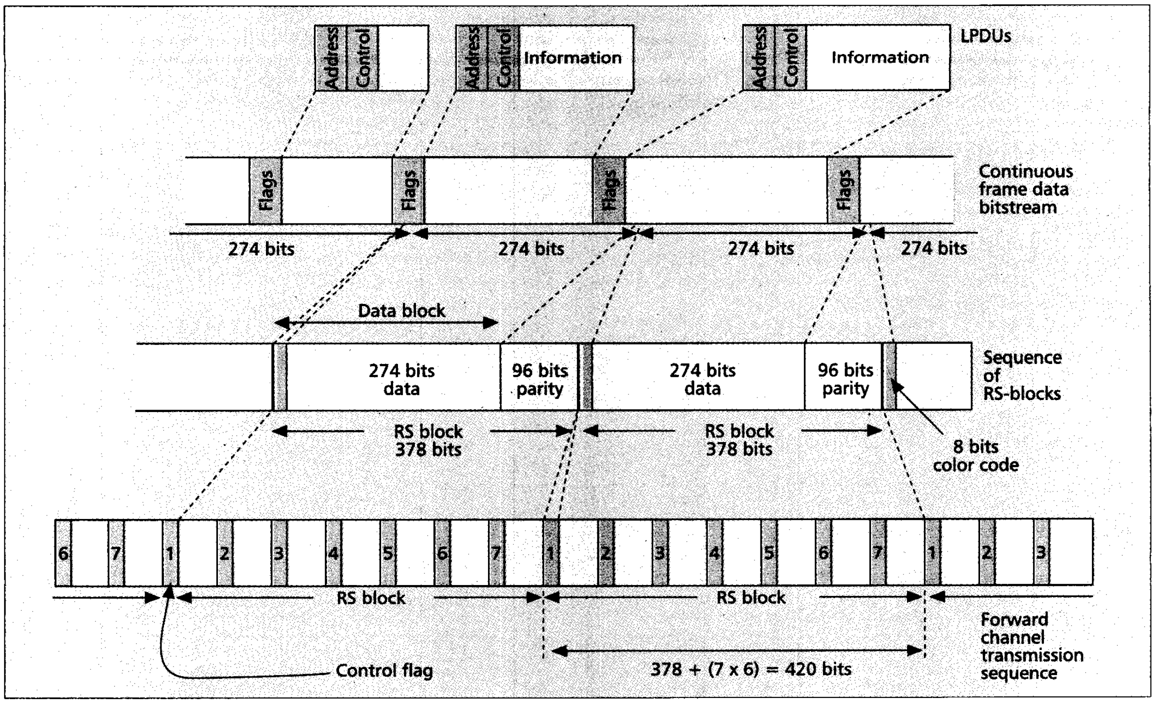

■ Figure 4. *Formulation of the forward channel data stream.*

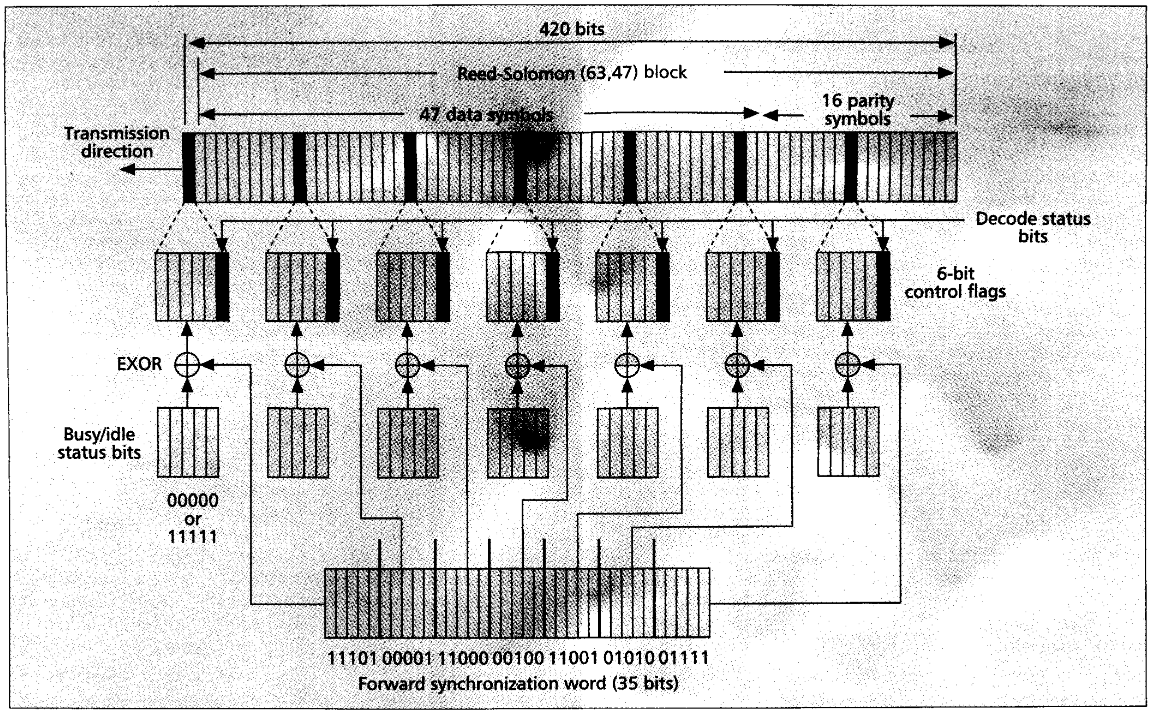

Figure 5. *Detailed construction of a forward channel data stream.*

types of demodulators (e.g., PLLs) and may result in reduced performance or increased implementation complexity.

As illustrated in Fig. 4, what is actually transmitted on the forward channel is the contiguous sequence of RS blocks (after scrambling), interleaved with special control flags. These flags carry synchronization information that helps M-ESs acquire block synchronization and decode the forward channel, as well as MAC-level control information that helps M-ESs effectively share the common reverse channel.

Figure 5 illustrates in detail the forward channel transmission structure. It shows how the control flags are constructed and how they are interleaved with the forward channel RS blocks. Each control flag is composed of one *decode status* bit plus five more bits, which derive from an exclusive-OR operation between a 5-bit section of the *forward synchronization word* (FSW) and *busy/idle status* bits. The FSW is a 35-bit sequence that provides a reference marker within the forward channel bitstream to discriminate between control flags and RS block boundaries. Additionally, it provides a timing reference for the reverse channel *microslot clock*. This microslot clock as well as the decode status and busy/idle status bits form the primary elements of the MAC procedure and are discussed later.

Since the FSW is transmitted after being exclusively-ORed with the busy/idle status bits, its discrimination within the transmitted bitstream would be impossible if the value of busy/idle status bits was not known. Fortunately, as discussed later, the busy/idle status can be either 11111 or 00000; therefore, the 5-bit portion of the FSW carried within every control flag is either inverted (when the busy/idle status is 11111) or not inverted (when the busy/idle status is 00000).

An M-ES is actually synchronized with a forward CDPD channel as long as it can receive and decode the forward blocks of this channel with some acceptable error rate. If the level of error rate rises above a threshold, which is implementation-dependent (this is where commercial CDPD modems may defer), the M-ES starts searching through a series of RF channels to find another more suitable CDPD data stream (see [5] for details).

REVERSE CHANNEL CONFIGURATION

The structure of the data transmitted by M-ESs in a CDPD network (i.e., the structure of transmissions on the reverse channel) is now discussed.

Consider an M-ES having, for example, three LPDUs pending transmission, as illustrated in Fig. 6. These LPDUs are flag-delimited and zero-stuffed, and thereafter are joined together to form a frame data bitstream. The first 274 bits of this bitstream are prefixed by an 8-bit color code (which is the same as the color code transmitted on the forward channel) and form a 282-bit data block. The rest of the frame data bitstream is padded with interframe time fill (usually a consecutive sequence of 1s) and divided into an integer number of sequential data blocks, each 282 bits long.

All the data blocks formed this way are thereafter subject to (63, 47) Reed-Solomon coding, and thus a sequence of contiguous RS blocks, each with a fixed length of 378 bits, is constructed. After these RS blocks are scrambled, they are ready to be transmitted on the reverse channel. However, prior to their transmission are transmitted a) a 38-bit sequence of alternating 1s and 0s (the preamble), which helps the MDBS detect the transmission start and acquire timing synchronization; and b) a reverse synchronization word (RSW), which is a 22-bit pattern that helps the MDBS acquire block synchronization. The transmission of RS blocks follows right after the RSW.

As shown in Fig. 6, a 7-bit *continuity indicator* is interleaved with each RS block; 1 bit every nine 6-bit symbols. This continuity indicator is a sequence that signals whether the reverse transmission burst is completed or not. A sequence of all 1s indicates that more RS blocks follow, whereas a sequence of all 0s marks the final transmission block. Note that since the continuity indicator is not error-protected, it features high redundancy and time-diverse transmission (which effectively uncorrelates the errors that may occur within its 7 bits). Note also that the continuity indicator carries in-band information regarding the transmission progress, which results in more robust and less complex reception. Should the continuity indicator not be used, the moment a transmission ends would be derived from received signal analysis — a more complex and time-costly procedure.

THE MEDIUM ACCESS PROCEDURE

An M-ES can access the reverse channel using a slotted non-persistent digital sense multiple access with collision detection (DSMA/CD) algorithm. This is similar to carrier sense multiple access with collision detection (CSMA/CD) used in Ethernet. However, in CDPD because the M-ESs cannot sense the status of the reverse channel directly (because they employ different reception and transmission frequency bands), a different collision detection scheme is applied.

DSMA/CD makes use of the busy/idle and decode status flags. As previously stated, the busy/idle flag is a 5-bit sequence transmitted on the forward channel once every 60 bits (i.e., once every *microslot* period). This flag provides periodic binary information with one microslot resolution indicating whether the reverse channel is busy or idle.

On the other hand, the decode status flag is a 5-bit[1] sequence that carries binary information indicating whether the MDBS has decoded the preceding block on the reverse channel successfully or not. On successful decoding the decode status flag is 00000, on unsuccessful decoding 11111.

An M-ES wishing to transmit senses first the busy/idle flag — actually, a locally stored version of it which is updated once every microslot period. If the reverse channel is found busy, the M-ES defers for a *random* number of microslots and then repeats the sensing of the busy/idle flag. Because the M-ES does not persist in continuously sensing the busy/idle flag, the access scheme is referred to as *nonpersistent*. Once the reverse channel is found idle, the M-ES may initiate transmission. Note that a transmission may initiate only at a microslot boundary, which is why the access scheme is termed *slotted*. As soon as the MDBS detects a transmission start on the reverse channel it sets the busy/idle flag in order to prevent further transmissions.

After an M-ES has started a transmission, it checks the decode status flag in every forward channel block it receives (this assumes full duplex operation[2]) and resumes or suspends transmission depending on the value of this flag. This flag provides "real-time" information regarding the progress of its ongoing transmission. The M-ES continues transmission if the decode status flag indicates that the MDBS encountered no decoding errors so far, whereas it ceases transmission in the opposite case (note that the MDBS cannot distinguish between errors due to collision and those due to channel impairments). In the latter situation, the M-ES attempts to regain access to the reverse channel after an appropriate

---

[1] *Although M-ESs decode a 5-bit decode status flag, the MDBS transmits 6 or 7 bits of decode status per reverse channel block; the exact number depends on the relative timing difference between the forward and reverse channels [1].*

[2] *M-ESs may be either full-duplex or half-duplex. Half-duplex M-ESs are subject to certain restrictions regarding access to the reverse channel due to their inability to receive the decode status flag while transmitting.*

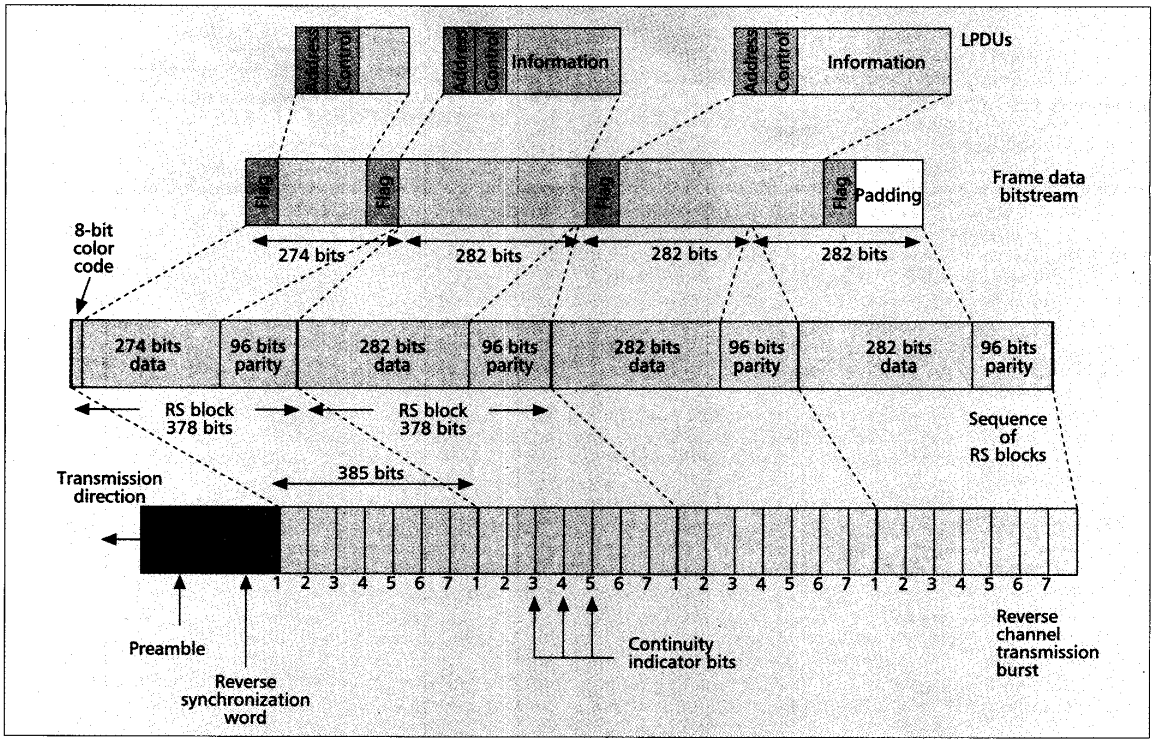

■ **Figure 6.** *Detailed construction of the reverse channel bitstream.*

DOCKET ALARM

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts

Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research

With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips

Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.

WHAT WILL YOU BUILD? | **sales@docketalarm.com** | **1-866-77-FASTCASE**

fastcase®
Smarter legal research.