
Abstract

This article provides an overview of the flexible, high-performance packet data channel that has been designed for high-rate packet data services over IS-136 TDMA channels. To achieve the highest data rates in the limited 30 kHz channel bandwidth, the packet data channel is designed for adaptive modulation and, in addition to a fixed coding mode, permits operation using an incremental redundancy mode.

GPRS-136: High-Rate Packet Data Service for North American TDMA Digital Cellular Systems

**KRISHNA BALACHANDRAN, RICHARD EJZAK, SANJIV NANDA,
STANISLAV VITEBSKIY, AND SHIV SETH**
LUCENT TECHNOLOGIES

Current North American time-division multiple access (TDMA) systems support voice services and circuit data services at a rate limited to 9.6 kb/s. High-rate packet data services are desirable for short bursty transactions as well as applications such as World Wide Web/Internet access, electronic mail, and file transfer. This article provides a general overview of the proposed system, and describes flexible, high-performance medium access control (MAC) and radio resource management procedures that have been adopted by the Telecommunications Industry Association (TIA) for high-rate packet data services over IS-136 TDMA channels. The 30 kHz spectrum usage, symbol rate, and TDMA format (6 time slots every 40 ms) are maintained as in the IS-136 standard [1] for compatibility with existing mobiles and in order to minimize impact on existing infrastructure. Assuming these constraints, a new TDMA packet data channel (PDCH) is defined to carry both user data and information from other higher-layer control and management entities.

Current cellular TDMA systems are designed to achieve good coverage over most of a typical cell, and as a result, the signal to interference plus noise ratio is sufficient to support higher data rates through the use of 8- and 16-level modulations over a large portion of a typical cell. However, current systems rely on the use of only $\pi/4$ -differential quadrature phase shift keying (DQPSK) modulation to achieve modest data rates over all operating conditions. The packet data MAC and physical layers are designed to increase throughput over a significant fraction of the cell area by using coherent 8-PSK in addition to $\pi/4$ -DQPSK. The standard is designed to permit dynamic adaptation of the modulation scheme based on measured carrier-to-interference ratio (C/I). The modulation schemes are chosen to be the same as those adopted for voice services in order to simplify the development of dual-mode (voice and packet data capable) mobile stations. Hooks are provided in the standard to support a 16-level modulation such as 16-quadrature amplitude modulation (QAM), 16-PSK, or 16-DPSK in the future.

Three wireless data infrastructure options were investigated: Cellular Digital Packet Data (CDPD) infrastructure that is widely deployed in the United States; Global System for Mobile Communications (GSM) General Packet Radio Service (GPRS) infrastructure that is expected to be deployed throughout the world in 1999 and 2000; and the third option considered, an Internet service provider (ISP)-like model with Point-to-Point Protocol (PPP) tunnels to the Internet or corporate intranets. In order to achieve economies of scale and simplify evolution to third-generation systems, the upper layers (layer 3 and above) of the packet data protocol stack were chosen to be the same as that used by GSM GPRS. GPRS-136 is thus a TDMA packet data standard based on GPRS, but utilizing 30 kHz for the physical layer and allowing connection to the American National Standards Institute (ANSI)-41 network [2]. GPRS-136 utilizes most of the existing network elements from the GPRS network reference model. It adds a gateway-mobile switching center (MSC)/visitor location register (VLR) function, which allows connection of the GPRS packet data network to the ANSI-41 based mobile circuit-switched network. The relevant GPRS specification documents can be found in [3–8].

Another benefit of this choice is that the existing GPRS standard is used as a baseline, allowing for quick development. This also enables evolution to the Enhanced General Packet Radio Service (EGPRS), which has been proposed to the International Telecommunication Union (ITU) as a third-generation radio transmission technology for International Mobile Telecommunications in the year 2000 (IMT-2000). The use of EGPRS channels for TDMA packet data will be standardized during 1999 and is to be called GPRS-136HS.

Organization of the Article

The following section provides an overview of the service and network reference model. The GPRS-based protocol stack is described briefly. Provision of voice/data integration and operation with half-duplex and full-duplex terminals is discussed.

The article next describes the logical and physical structure of the PDCH. The allocation of packet paging and packet broadcast channels, and PDCH reselection procedures are described. Multiple-time-slot operation is also discussed. We then describe the MAC layer and uplink media access control procedures, including the use of the packet channel feedback (PCF) mechanism.

The article describes the procedures associated with adaptive modulation (DQPSK and coherent 8-PSK), followed by the incremental redundancy and fixed coding mode radio link protocols (RLPs). The incremental redundancy mode provides approximately 15 percent higher throughput at the cost of additional receiver memory. We provide detailed time slot formats accommodating the fields required by the RLP, MAC, and physical layers. The last section provides a brief summary.

GPRS-136 Network Reference Model and Operation

The Packet Data Network Reference Model

The network elements are as follows:

- Terminal equipment (TE), which typically interfaces with the user and contains packet data applications.
- Mobile termination (MT), which interfaces to the TE, and terminates the radio interface.
- Base station (BS), which constitutes the interface between the network and mobile station, and transfers packet data and signaling messages between serving GPRS support nodes (SGSNs) and mobile stations in its coverage area.
- SGSN, a packet data switch that routes data packets to appropriate mobile stations within its service area.
- Gateway GPRS support node (GGSN), which acts as the logical interface between the GPRS-136 network and external packet data networks. It tunnels IP packets from external networks to the SGSN using the GPRS Tunneling Protocol (GTP).
- GPRS home location register (HLR), accessible from the SGSN and GGSN, contains GPRS-136 subscription and routing information.
- ANSI-41 HLR, accessible from the serving and gateway MSC/VLR functions, contains subscription and routing information for circuit-switched service.
- ANSI-41 gateway MSC/VLR, provides functions such as circuit call routing and circuit service related paging within the GPRS-136 network.
- ANSI-41 serving MSC/VLR, provides circuit switching functions for mobile stations in its service area.
- Message center (MC), receives and accepts requests to deliver teleservice messages to the mobile subscriber.

The Protocol Stack (GPRS-136)

Figure 2 shows the GPRS-136 protocol stack. The MAC and radio resource management procedures

are significantly different from GPRS because of the fundamental differences in the IS-136 and GPRS physical layers.

The radio resource (RR) sublayer provides data transfer and control services to LLC and GPRS-136 mobility management. It consists of three cooperating entities: the MAC entity, radio resource management entity (RRME), and broadcast management entity (BME).


The MAC entity provides services directly to LLC through a SAP. It supports two multiplexed logical links of different priorities, one for normal data transfer and the other for providing expedited data delivery services to higher-layer and management entities.

The BME and RRME provide services to the GPRS mobility management layer through two other SAPs, and carry out both voice and packet-data-related signaling functions. These entities use the services of the MAC entity to communicate with their peers, and to control the MAC entity. The core MAC and RRME procedures are generic, and the interfaces can easily be modified to support other protocol stacks.

Joint Voice/Data Operation

In addition to packet data service, GPRS-136 allows subscribers to obtain both IS-136 and ANSI-41-based services (e.g., circuit-switched voice services, short messaging services, intelligent roaming), provided mobile stations support both circuit and packet modes. Two new classes of mobile stations are supported by GPRS-136: dual mode (i.e., circuit and packet mode capable) and packet mode only.

Integration of the GPRS packet network with the ANSI-41 circuit-switched network provides a unique challenge. GPRS provides its own set of registration, authentication, authorization, and mobility management functions which are thoroughly integrated with other GSM services such as circuit-switched voice service. These functions are implemented very differently in the ANSI-41 network. To avoid the complexity of fully integrating these two sets of functions, GPRS-136 provides a method of "tunneling" ANSI-41 signaling messages between a mobile terminal and the gateway MSC/VLR via the SGSN. These messages are tunneled as



■ **Figure 1.** The GPRS-136 packet data network reference model. (Primed interfaces — e.g., Gb' — indicate ETSI GPRS interfaces that have been modified for GPRS-136.)

specially marked LLC frames. ANSI-41 registration, authentication, authorization, paging, and short message service (SMS) messages are delivered transparently (tunneled) through the SGSN. This allows the MSC/VLR and SGSN to implement and execute these functions independently as necessary for the proper operation of the ANSI-41 and GPRS networks, respectively. The mobile terminal provides the necessary sequencing of related functions between the ANSI-41 and GPRS networks.


The GPRS-136 mobility management layer consists of two entities: GPRS mobility management (GMM) and 136 mobility management (136MM), which support mobility management functions specific to packet data and circuit-switched services, respectively. A new tunneling function is defined in the SGSN, and the GPRS LLC and BSSGP protocols have been modified in order to route 136MM signaling messages to the mobile station over the PDCH.

Dual-mode mobile stations perform registration, authentication, authorization, and location update functions independently and in parallel with the ANSI-41 and GPRS networks while the mobile is camped on a PDCH. If an ANSI-41 SMS message needs to be delivered while a mobile

is camped on a PDCH, it is delivered on the PDCH along with other packet data as tunneled LLC data. An incoming circuit call is indicated to a mobile with an ANSI-41 page message tunneled from the gateway MSC/VLR to the mobile. The mobile moves to an IS-136 control channel to respond to the page and accept the incoming call. Similarly, to initiate a circuit call, the mobile stops any activity on the PDCH and moves to the IS-136 control channel. Circuit calls take precedence over packet data transactions. When a circuit call is completed, the dual mode mobile station returns to camping on a PDCH.

Half-Duplex Operation

Half-duplex devices have complexity, size, and battery life benefits which make them attractive for applications that do not require full duplex (i.e., simultaneous transmission and reception) capability. Full-rate operation (i.e., similar to IS-136 voice) is half-duplex by definition because of fixed time offsets between the uplink and downlink [1]. The support of downlink double- and triple-rate (multislot) operation for packet data is quite straightforward, and is enabled by occasionally scheduling uplink time slots for obtaining automatic



■ Figure 2. The GPRS-136 packet data protocol stack.

repeat request (ARQ) feedback. However, mobile stations cannot support more than full-rate operation on the uplink without simultaneous transmission and reception. GPRS-136 supports half-duplex operation on both the uplink and downlink for all mobile stations regardless of their multislot (or bandwidth) capability. Double- or triple-rate half-duplex operation on the uplink is made possible by carrying out a fixed allocation of a sequential number of time slots. No downlink transmissions are scheduled for the device during this period.

Packet Data Channel Structure

Logical Channel Structure

The PDCH consists of the following logical channels (Fig. 3):

- Packet broadcast control channel (PBCCH) for indicating generic system configuration related information
- Packet paging channel (PPCH) dedicated to delivering pages
- Downlink packet payload channel (PLCH) for delivering data generated by LLC, RRME, and GMM
- Packet channel feedback (PCF) for support of random access and reserved access on the uplink
- Uplink packet random access channel (PRACH) used by mobile stations to request packet data access to the system
- Uplink packet payload channel (PLCH) for delivering data generated by LLC and GMM

Physical Channel Structure

The PDCH uses 30 kHz RF channels and the time-slot structure specified in IS-136 [1]. Each 40 ms frame on a 30 kHz RF channel consists of six time slots (three time slot pairs), numbered 1 to 6. One or more time slot pairs may be allocated to a PDCH. The remaining time slot pairs may be allocated to a digital control channel (DCCH) and/or digital traffic channel (DTC).


Multirate Channels – A PDCH may be full-rate, double-rate, or triple-rate, depending on whether one, two, or three time slot pairs are allocated to the channel within each 40 ms frame. The channel bandwidth is indicated through system broadcast information.

Superframe Structure – Sleep mode is defined for mobile stations on the PDCH in order to improve mobile station standby time. A superframe structure similar to the IS-136 DCCH is defined on each full-rate portion of a PDCH in order to manage sleep mode. The total number of time slots per superframe is 32. The superframe phase (SFP) is a modulo 32 up-counter which increments every 20 ms and helps mobile stations find the start of the superframe. If a double- or triple-rate PDCH is allocated on a particular frequency, superframe synchronization across the individual full-rate portions is required.

Primary and Supplementary Phases – A multirate PDCH operates on a single channel frequency and consists of the following:

- Primary phase
- Supplementary phase(s)

In this context, a phase corresponds to a full-rate portion of a multirate PDCH. The primary phase always corresponds to a full-rate channel, and is the part of a multirate channel that contains logical broadcast and paging channels on the downlink. The supplementary phases on a multirate PDCH correspond to all time slots that are not part of the primary



■ Figure 3. Logical PDCHs.

phase. The possible allocation of primary and supplementary phases is provided in Table 1.

The primary phase superframe consists of time slots that are reserved for the PBCCH. Nominal paging time slots are determined from among the remaining time slots using a standard hashing algorithm that relies on mobile station (subscriber) identity. The actual number of PPCH slots per superframe are configurable through a fast broadcast message. Except for the assignment of broadcast and paging channels and SFP management, a mobile station views a double- or triple-rate PDCH as one common channel (i.e., a fat pipe) regarding data transmission or reception.


Downlink and Uplink Burst Associations – Subchannels are defined on the uplink in order to allow sufficient processing time at both the mobile station and BS in conjunction with a random access event. The full-rate PDCH is defined to consist of three subchannels; there are six subchannels in each double-rate PDCH and nine subchannels in each triple-rate PDCH (Fig. 4).

An enhanced packet channel feedback (PCF) field that relies on this association is defined on the downlink for uplink resource management. PCF flags are carried in downlink time slots to provide feedback for bursts sent previously on the uplink and to indicate subsequent assignments on the uplink. Mobile stations identify access opportunities on the uplink by reading PCF on the corresponding downlink time slots. The PCF field allows the support of contention access and reserved access on the same channel, and assumes different flag definitions depending on the context. These flags are reliably encoded to ensure good performance even under poor channel conditions.

For a full-rate PDCH, the uplink and downlink time slots are multiplexed to create three distinct access paths, as shown in Fig. 4. Assuming that path 1 (P1) in the downlink indicates that the next P1 time slot on the uplink is designated as a contention slot and is selected for a random access attempt, a mobile station sends the first burst of its access at that time

Rate	Slot numbers	Primary phase slots	Slots for supplementary phase(s)
Full-rate PDCH	1, 4	1, 4	–
Full-rate PDCH	2, 5	2, 5	–
Full-rate PDCH	3, 6	3, 6	–
Double-rate PDCH	1, 2, 4, 5	1, 4	2, 5
Double-rate PDCH	2, 3, 5, 6	2, 5	3, 6
Double-rate PDCH	1, 3, 4, 6	1, 4	3, 6
Triple-rate PDCH	1, 2, 3, 4, 5, 6	1, 4	2, 3, 5, 6

■ Table 1. Multirate PDCH phase allocations.



■ **Figure 4.** Downlink and uplink burst associations.

(24.8 ms after receiving the full P1 slot on the downlink). The mobile station then starts reading the PCF flags in the next downlink P1 slot (21.8 ms after transmitting its burst) to determine the reception status of its initial access.

Similarly, there are six access paths on each double-rate channel and nine access paths on each triple-rate channel. In general, a mobile station which has data to send transmits 24.8 ms after obtaining an assignment via PCF on the downlink. The mobile station reads PCF on the downlink 21.8 ms after transmission in order to obtain the reception status of its transmitted burst.

The Medium Access Control Function

PDCH Selection, Reassignment, and Reselection

Mobile stations are directed to a PDCH through IS-136 DCCH broadcast information. In cases where there are multiple PDCHs per sector, mobile stations are directed to a beacon PDCH. The beacon PDCH broadcast information indicates the number of PDCHs, as well as the bandwidth (full-, double-, or triple-rate) of each PDCH supported. Mobile stations then hash onto a particular PDCH depending on their identity and the number of PDCHs. Load balancing may be carried out by reassigning mobile stations across radio resources, and the performance may be further improved by maintaining MAC/RLP state across reassignments. Cell reselection procedures ensure continuity of service across cell boundaries since mobile stations autonomously perform PDCH reselection when they detect a stronger signal from a neighbor cell.


Active Mobile Identity Management

Each mobile station is assigned a 7 bit temporary local identifier called an active mobile identity (AMI) which remains valid for one or several closely spaced transactions. The AMI

is used to identify uplink time slot assignments and identify the recipient of data on the downlink. Of the 128 possible AMI values, only 89 are allowed for mobile stations engaged in point-to-point transactions. The all-zero AMI is excluded for PCF-related functions and also for identifying point-to-multipoint information on the downlink. The remaining 38 AMIs closest to the all-zero codeword are excluded in order to increase PCF reliability.

AMI assignment procedures are executed for both uplink and downlink transactions spanning more than one time slot. If a valid AMI has not already been assigned, the AMI assignment is carried out as a part of the transaction initiation procedure. Once an AMI has been assigned to a mobile station, it is used for transactions in both directions (i.e., the AMI is assigned on the initiation of an uplink or downlink transaction, whichever begins first, and remains assigned until released). AMI release is based on the expiration of timers at both the mobile station and BS. The AMI release procedures are designed to ensure the following:

- The availability of the same AMI to the mobile station for



■ **Figure 5.** Logical format of the PCF. The SLF is context-dependent and contains either an AMI for providing feedback to contention slots, or three fields for providing feedback to reserved slots.

another downlink or uplink transaction which is initiated before the expiration of the timer

- The release of the AMI first by the mobile station and subsequently by the BS in order to avoid any potential hijacking scenarios.

Transaction Initiation

A new transaction is initiated by the transmit controller when a transmission opportunity is identified and if the transmit buffer contains new data. Downlink transactions may be acknowledged or unacknowledged, but uplink transactions are always acknowledged. The initiation of a new transaction is carried out through the transmission of a BEGIN PDU. The BEGIN PDU handshake (i.e., acknowledged transfer of the BEGIN PDU) is used to initialize the AMI and an RLP in either incremental redundancy (IR) or fixed coding (FC) mode for the transaction.

For downlink transactions, the AMI and mode are assigned using the BEGIN PDU. For uplink transactions, the mobile station suggests an AMI and mode using the BEGIN PDU. If the suggested AMI and mode are acceptable to the base station, it provides an acknowledgment using the packet channel feedback (PCF) field. The PCF acknowledgment is treated as an implicit AMI and mode assignment by the mobile station. If the suggested AMI and/or mode are unacceptable to the base station, it provides a negative acknowledgment using PCF and subsequently assigns an AMI and/or mode for the transaction using a supervisory ARQ Status PDU.

Packet Channel Feedback

Functions – The IS-136 digital control channel [1] uses a shared channel feedback (SCF) field to provide acknowledgment and assignment functions. The SCF uses three flags — Received/Not Received (R/N), Busy/Reserved/Idle (BRI), and a coded partial echo identifier (CPE) — in order to manage feedback and assignment functions. The DCCH procedures and fields as specified in IS-136 are not well suited for long packet data transactions.

As described above, the uplink and downlink burst associations on the PDCH result in three subchannels per full-rate channel, six subchannels per double-rate channel, and nine subchannels per triple-rate channel. A PCF field is associated with each subchannel, and 24 bits are used in each downlink time slot for the PCF. The PCF is designed to be more reliable than the SCF and allows efficient management of contention access and reserved access on the same channel. For a particular subchannel, the PCF provides acknowledgment for a transmission on the previous time slot and also indicates assignment of the next time slot. Consistent canonical feedback and assignment fields are defined as follows: acknowledgment/negative acknowledgment (ACK/NAK) feedback identifying the transmitter when acknowledging a contention slot, and assignment of the next slot as idle (for contention-based access) or reserved (for access by a specific user). The approach presented simplifies the state machine since the acknowledgment for the previous slot and the assignment for the next slot are unambiguously identified with specific mobiles.

Figure 5 shows the logical format of the PCF field associated with each subchannel. The main purpose of the PCF field is to acknowledge an access in the previous slot and to assign the next slot to a particular mobile. The acknowledgment and assignment functions are handled independently for each subchannel through two logical fields, subchannel feedback (SLF) and subchannel

assignment (SLA). The use of the SLF mechanism is context-dependent (i.e., the SLF field takes different values depending on whether feedback is being provided for a contention slot or a reserved slot). This provides greater reliability for acknowledgments to reservation-based transmissions. In addition, the same field may also be used for reliably reassigning the subchannel to the same user and providing continuous uplink channel quality feedback to the mobile station.

Assignments to successful contending users or users sniffing for reservation-based access may be carried out by explicitly identifying the user being assigned the next slot (i.e., by setting the SLA field to a valid coded active mobile identity). This approach is useful when there are several users on the channel, and subchannels (time slots) are constantly being assigned to new users. However, with fewer active users on the channel, subchannels may continuously be reserved for (reassigned to) the same users. Figure 6 shows that if there are three or fewer users on a full-rate channel, a round-robin assignment scheme is equivalent to reassigning subchannels to the same users. Similarly, Fig. 7 shows that if there are nine or fewer users on a triple-rate channel, a round-robin assignment scheme is equivalent to reserving one or more subchannel(s) for each user. In such cases, it is not necessary to use a coded AMI to reassign a subchannel to the same user. If there are fewer users on the channel, it is possible to provide a series of reservation-based transmission opportunities on each subchannel through the use of a 1-bit CONTINUE/STOP (C/S) indicator. The CONTINUE indication is similar to BRI = BUSY [1], but is encoded more robustly. Since the CONTINUE (or BUSY) indication is associated with feedback for a reserved slot and is relevant only to the user who transmitted on the previous time slot on a particular subchannel, the C/S flag is included in the SLF field.

PCF Encoding


Subchannel Feedback – The 12-bit SLF field is used to acknowledge bursts transmitted in the previous slot. For feedback corresponding to contention slots SLF is defined as follows:

- SLF = valid coded AMI (provides an implicit ACK to the mobile with that coded AMI suggested or assigned)
- SLF = E-NAK (explicit NAK to all mobiles which attempted access)


The (12,7) code described below is used to encode this field. The all-zero codeword is reserved for indicating an explicit negative acknowledgment (E-NAK) to all mobiles which attempted contention access.

For feedback corresponding to reserved slots, SLF is further divided into the following:

- R/N (5 bits): A (5,1) repetition code is used. A 1 indicates



■ **Figure 6.** Round-robin assignments for users *a*, *b*, and *c* on a full-rate channel. This is equivalent to reservation of subchannels 1, 2, and 3 for users *a*, *b*, and *c*, respectively.



■ **Figure 7.** Round-robin assignments for users *a*, *b*, ..., *i* on a triple-rate channel. This is equivalent to reservation of subchannels 1, 2, ..., 9 for users *a*, *b*, ..., *i*, respectively.

Parameter	Description
<i>Access_Count_Max</i>	Maximum number of times a mobile may backoff and reattempt contention access before declaring an access failure.
<i>Access_Count</i>	Counter indicating the number of times contention access has been attempted. This parameter can take values 0, 1, ..., <i>Access_Count_Max</i> .
<i>T_Retry</i>	Number of idle (contention) slots that the mobile must wait before another transmission attempt. <i>T_Retry</i> is a random integer that lies in the closed interval [0, <i>T_Retry_Max</i>].
<i>T_Retry_Max</i>	Value computed from other access parameters, determines the distribution used for generating random backoff.
<i>T_Retry_Init</i>	Initial value of <i>T_Retry_Max</i> to be used for <i>Access_Count</i> = 1.
α	Configurable parameter which is used along with <i>T_Retry_Init</i> to compute <i>T_Retry_Max</i> .

■ **Table 2.** Parameters for the random access procedure

<i>Access_Count</i>	<i>T_Retry_Max</i>
0	0
1, 2, ..., <i>Access_Count_Max</i>	$T_Retry_Init (2)^A$, where $A = \frac{Access_Count - 1}{\alpha}$

■ **Table 3.** Expressions for computing *T_Retry_Max*

that the transmission was received (R), while a 0 indicates that the transmission was not received (N). The mobile declares R/N = R if the Hamming weight of the received 5-bit word is strictly greater than 3; otherwise, it declares R/N = N.

- C/S (5 bits): A (5,1) repetition code is used to encode this flag. A 1 indicates CONTINUE or C; that is, the mobile is assigned the subsequent time slot on the same sub-channel. A 0 indicates STOP or S; that is, the mobile must read the SLA field to determine subsequent assignments on that sub-channel. The mobile station declares C/S = C if the Hamming weight of the received 5-bit word is strictly greater than 2; otherwise, it declares C/S = S.
- Channel quality feedback (2 bits): This field provides feedback on uplink channel quality to the mobile. A mobile capable of operation on multiple modulations can use this feedback to propose a different modulation for subsequent reservation based transmission opportunities.

Subchannel Assignment – The 12-bit SLA field can take the following values:

- SLA = valid coded AMI (assigns the subchannel to a mobile with that coded AMI)
- SLA = IDLE (identifies a contention opportunity)

The (12,7) code described below is used to encode this field. The all-zero codeword is used as an IDLE indicator (i.e., to indicate a contention slot).

AMI Encoding – Each mobile station is assigned a 7-bit AMI value for the duration of each transaction. The AMI assignment is initiated by an uplink or downlink transaction, whichever begins first.

The 7 bit AMI is encoded using a (12,7) code. This code is derived from a (15,11) Hamming code as follows. As an intermediate step, the (15,11) Hamming code is shortened to an (11,7) code as specified in [1]. This code is used for encoding the CPE on the IS-136 DCCH; it has a minimum Hamming distance of 3 and guarantees single error correction. However,

false error correction (i.e., interpretation of one codeword as another) may occur if two or more errors occur in the channel. Two error detecting capability is provided by increasing the minimum Hamming distance of the code to 4 through the addition of a single parity bit. The resulting (12,7) code may be viewed as an extended (11,7) code.

Mobile Station Procedures –

Mobile stations assume an appropriate SLF structure depending on whether they are examining feedback corresponding to a contention slot or a reserved slot. The decoding rules at the mobile station provide unequal error protection to

the different flag values.

A mobile which transmits in a particular *contention slot* reads the corresponding SLF field to determine if its transmission was successful:

- On receiving an ACK (i.e., SLF/coded AMI match), it reads the SLA field on all subchannels it is capable of operating on to determine if it is granted a reservation.
- On declaring a NAK (i.e., SLF/coded AMI mismatch), it follows random access procedures for attempting another contention access.

A mobile which transmits in a particular *reserved slot* reads the corresponding SLF field to determine if its transmission was successful:

- On receiving an ACK (i.e., R/N = R), the mobile reads the C/S indicator to determine if it can transmit in the subsequent slot.
 - On decoding C/S as CONTINUE, the mobile ignores the SLA field and assumes that it has permission to transmit in the subsequent time slot associated with the same subchannel.¹
 - C/S = STOP implies that the mobile being acknowledged must continue sniffing (reading SLA) on that subchannel for reservation-based transmission opportunities.
- On receiving a NAK (i.e., R/N = N), the mobile reads the SLA field to determine if it has been assigned the subsequent slot (i.e., it tries to match SLA with its coded AMI).²

Other active mobiles which are sniffing for reservations ignore the SLF field and examine the SLA field for time slot assignments. With this scheme, it is possible to reserve subchannels for some users and carry out round-robin assignments for other subchannels.

Contention Access

Contention slots are provided on the uplink PDCH in order to allow mobile stations to initiate packet data transactions. The mobile station identifies a contention opportunity by reading the PCF. If the mobile station attempts transmission in a contention slot and is negatively acknowledged by the BS, it starts a timer and waits for an AMI and/or mode assignment from the BS. If no AMI/mode assignment is received and the timer expires, the mobile station must wait for *T_Retry* idle (contention) slots before making another access attempt. The parameters used by the random access procedure are listed in Table 2.

T_Retry is uniformly distributed in the closed interval [0, *T_Retry_Max*] where *T_Retry_Max* is a function of *Access_Count*, *T_Retry_Init*, and α . The parameter

T_Retry_Max is computed using the expressions in Table 3. An access failure is declared if the mobile station is unsuccessful after $Access_Count_Max$ access attempts.

Adaptive Modulation

The FC and IR RLPs rely on adaptive modulation to achieve the best throughput under delay constraints. The RLP blocks are chosen to be of fixed length and a variable integer number of blocks are packed into a MAC PDU depending on the modulation. This ensures that blocks transmitted using a particular modulation can be retransmitted even if the modulation has changed. The receiver provides periodic channel quality feedback (CQF) which indicates the maximum constellation size allowable under the prevailing channel conditions. The transmitter uses the CQF along with knowledge of the offered load in order to carry out modulation adaptation. The modulation and coding mode for the data in each time slot are indicated to the receiver through a separately coded field called the coded data field type (CDFT). This allows efficient multiplexing of IR and FC mode transactions with dynamic (slot-by-slot) modulation switching without any resulting ambiguity at the receiver.

Channel Quality Feedback

The receiver MAC layer provides CQF, which indicates the modulations that are allowable under the prevailing channel conditions. Based on this information, the transmitter chooses the modulation scheme for subsequent data transfer.

Table 4 shows the scheme for determining the maximum allowable constellation size based on the signal-to-interference-plus-noise ratio, $S/(I + N)$, estimated at the input to the decoder. The downlink adaptation threshold, θ_L , is configurable as a broadcast parameter, and $S/(I + N)$ estimation may be carried out using the technique described in [9, 10].

On the downlink, CQF is transmitted as a part of PCF and supervisory ARQ status PDUs. On the uplink, CQF is transmitted as a part of a supervisory ARQ status PDU. A valid CQF may be used along with the offered load to change the downlink modulation.

Data Field Type

The DFT is a 3-bit field indicating the mode (FC or IR) and modulation ($\pi/4$ -DQPSK or 8-PSK) of the data fields in each time slot, and also whether the data fields are associated with the broadcast or nonzero (i.e., mobile-specific) AMI. This field is always transmitted using $\pi/4$ -DQPSK.

The mode indication allows fixed coded slots to be inter-

DFT value	Modulation	Mode	AMI type
000	$\pi/4$ -DQPSK	Fixed coding	Idle (all zero)
001	$\pi/4$ -DQPSK	Fixed coding	Mobile-specific
010	$\pi/4$ -DQPSK	Incremental redundancy	Mobile-specific
011	8-PSK	Fixed coding	Mobile-specific
100	8-PSK	Incremental redundancy	Mobile-specific
101	Reserved		
110	Reserved		
111	Reserved		

■ **Table 5.** Data field type values (uplink and downlink). No 16-level modulation has been specified at this time.

Channel conditions	CQF	Allowable modulations
$S/(I + N) < \theta_L$	00	$\pi/4$ -DQPSK
$S/(I + N) \geq \theta_L$	01	$\pi/4$ -DQPSK, 8-PSK
–	10	Reserved
–	11	Reserved

■ **Table 4.** CQF values on uplink and downlink. The threshold θ_L is a broadcast parameter. CQF = 10 is intended for 16-level modulation. No 16-level modulation has been specified at this time.

persed between IR slots. This allows, for example, expedited control/feedback frames to be transmitted in FC mode while an IR mode transaction is in progress. The modulation type indication enables slot-by-slot adaptation. The indication of mode and modulation in each time slot is also useful in order to quickly detect handshake failures and take recovery action.

The DFT values are indicated in Table 5. Only five DFT values are currently used; additional DFT values may be used for the support of FC and IR with 16-level modulation. No 16-level modulation has been specified at this time.

The DFT is encoded reliably, and soft decision decoding is recommended at both ends in order to ensure that the error rate is significantly lower than that of the data under a wide range of operating conditions. On the uplink, the 3-bit DFT field is encoded using a (6,3) code to form the CDFT. On the downlink, the 3-bit DFT is combined along with the 5-bit superframe phase (SFP), and the resulting 8-bit word is encoded using a (12,8) code to obtain the CSFP/CDFT field. While decoding, the mobile station can use its prior knowledge of SFP for each time slot in order to achieve better performance.

Radio Link Protocol


The radio resource (RR) sublayer provides data transfer and control services to LLC and GMM. It consists of three cooperating entities: the MAC entity, RRME, and BME.

The MAC entity accepts data from LLC and RRME. It supports two multiplexed logical links of different priority, one for providing expedited data delivery services (MAC logical link 0, or MLL0), and the other for normal data transfer (MAC logical link 1, or MLL1). For acknowledged data transfer, each logical link uses an RLP. The RLP for MLL1 operates in two modes: IR and FC. The RLP for MLL0 uses only FC mode. A header including a higher-layer protocol identifier (HLPI) is attached to each MAC service data unit (SDU), and the result is converted into an octet stream using standard byte-oriented HDLC framing. The octet streams for each MAC logical link are buffered independently, and are converted into RLP segments of different size depending on the mode.

For both IR and FC, the retransmission scheme is based on the TDMA circuit data RLP IS-130 [11, 12]. Thus, retransmissions are given priority over new data, and retransmissions are determined by the transmitter based on receiver-state feedback (bitmap feedback). Retransmissions are not based on timers. Extensions to the

¹ The cell always sets SA to the coded AMI of the user being assigned a particular slot. Therefore, if the mobile erroneously declares a NAK when an ACK was sent, it can still obtain the assignment of the subsequent slot.

² When the BS transmits $R/N = N$, it sets $C/S = S$.



■ Figure 8. An IR segment encoder.

IS-130 RLP have been defined for operation:

- At the MAC layer for packet data transactions
- With modulation adaptation
- Using the incremental redundancy mode

The RLP is initialized in IR or FC mode using the transaction initiation procedure described earlier. A BEGIN PDU is used for initializing a transaction, and subsequent data transfer is carried out through the transmission of a series of CONTINUE PDUs. Each transaction may be bounded or unbounded, and unbounded transactions may be either gracefully ended or aborted on detecting protocol violations. Supervisory ARQ Status PDUs are used to periodically provide the peer transmission controller with knowledge about the receiver state.

A transmission opportunity identified during normal protocol operation results in the transmission of one of the following:

- ARQ Status PDU
- IR or FC CONTINUE PDU

CONTINUE PDU formats for IR and FC will be described in the following sections. Seamless operation with adaptive modulation is achieved by packing fixed-size blocks within each CONTINUE PDU. Different integer numbers of these fixed-size blocks can be accommodated in a CONTINUE PDU for different modulations. In IR mode, multiple fixed-size data and parity blocks are derived from each RLP segment. In FC mode, each data block is an RLP segment.

MAC PDUs transmitted on the uplink are acknowledged via the PCF mechanism. Optionally, full or partial bitmap feedback may be provided by the BS via ARQ status PDUs. The transmitter may poll the receiver for bitmap feedback at any time. For downlink data transfer, the BS provides the mobile station with reservation opportunities to obtain bitmap feedback.

The local variables and tables maintained at the transmitter and receiver, the update procedures, and a detailed description of the protocol using the Specification and Description Language (SDL) are provided as a part of the

MAC-layer specification [13].

Incremental Redundancy Mode RLP

The IR RLP achieves higher throughput since redundant bits are transmitted only when necessary. It exploits the fact that the extra redundancy introduced by heavy coding is not required in most cases. IR mode places more memory requirements on the receiver, but achieves more throughput than FC mode. IR mode is to be used only for regular data transfer (i.e., MLL1), not the transfer of expedited control information (MLL0).


Data Formats and Encoding – The byte stream derived from the upper-layer data is segmented into fixed-length RLP segments of length L . To each RLP segment, a cyclic redundancy check (CRC) sequence computed over the data is added. The data and CRC are encoded using a rate 1/2, 32-state, tailbiting, nonsystematic, maximal free distance convolutional encoder with octal generators (65, 57). The rate 1/2 convolutional encoder has two outputs, one of which is suitably punctured, as discussed below. Without loss of generality, the unpunctured encoder output is referred to as *data* and the punctured output as *parity*. In reality, the data output contains no redundancy and represents a one-to-one mapping to the actual data. The data bits at the output of the encoder are interleaved and segmented into D blocks of length L/D . These blocks are called *data blocks* and denoted by D_{ij} ($j = 1, 2, \dots, D$). Of the L parity bits at the encoder output, Dh parity bits are punctured, and the remaining

parity bits are segmented into D parity blocks of equal size denoted by P_{ij} ($j = 1, 2, \dots, D$). Here, h is the size of the header required for each parity block. At the receiver, soft information from data and parity blocks corresponding to the same segment is combined for decoding. Figure 8 shows the mapping of data segments to data and parity blocks. Based on delay/throughput trade-offs, the parameter D was chosen to be 3. The output bits of the encoder are mapped in an interleaved manner to maximize code/time diversity.

To each parity block, a coded parity/control block header (CPCBH) is appended. For parity blocks, PCBH contains the following:

- 1-bit PCBH header type (= 1 for parity block)
- 10-bit block sequence number (BSN)

The PCBH is encoded using a punctured rate 1/2 convolutional code to obtain the CPCBH. Since parity blocks are identified through a header, the data blocks D_{ij} are assigned the same BSNs as the parity blocks, P_{ij} , for $j = 1, 2, \dots, D$. The blocks to be transmitted in a CONTINUE PDU are determined (Table 6) as follows. Adaptive modulation is employed to achieve the best throughput under delay constraints. The protocol is designed so that a variable integer number of RLP data and/or parity



■ Figure 9. The composition of an IR CONTINUE PDU (layer 1 SDU) for $\pi/4$ -DQPSK.


blocks may be packed into a single CONTINUE PDU, depending on the modulation. Retransmittable blocks are given priority over new RLP data blocks, and data blocks in each CONTINUE PDU are required to be in sequence. NULL blocks or END control blocks (indicating the last valid sequence number) may be inserted in place of parity blocks.

For the PDCH, two, three, or four (combinations of data and parity/control) blocks are transmitted in each CONTINUE PDU corresponding to the use of 4-, 8-, or 16-level modulation, respectively. Figures 9 and 10 show CONTINUE PDU formats as a function of the modulation type.


Implicit Addressing and Header Encoding – For IR, a separately coded data segment header (DSH) is transmitted along with each PDU (two, three, or four blocks). The DSH contains the following (Fig. 11):

- BSN — 10 bits (associated with the first data block in a time slot; if there are no data blocks, this field is given a default value which is ignored by the receiver)
- Poll indicator — 1 bit (to request an ARQ Status PDU)
- Parity/control block pointer (PCBP) — 2 bits (to indicate the composition of the PDU in terms of data and parity/control blocks)

Packet data channels are shared by multiple users, and each time slot must contain either an implicit or explicit address to indicate the transmitter identity and recipient of the data in both directions. The inclusion of an address on the uplink is useful for verification that the user assigned an uplink time slot was indeed the one whose transmission was received. Downlink addressing is necessary since multiple active mobiles are constantly listening for downlink transmissions, and the intended recipient must be identified. Furthermore, in order to provide protection from co-channel interference, it is desirable that the base station identity be indicated in both directions. The use of explicit addresses in each time slot requires large overhead and leads to degradation in performance. As a result, the addresses are included implicitly in the computation of CRC bits on both the uplink and downlink. Addressing is carried out using the 7-bit AMI along with a 5-bit abbreviated digital verification color code (ADVCC) as the identifier for the BS. The DSH, ADVCC, and AMI are used to compute the CRC. However, only the



■ **Figure 10.** The composition of an IR CONTINUE PDU (layer 1 SDU) for 8-PSK.



■ **Figure 11.** The structure of the data segment header; the number of bits required for each field is indicated.

13-bit DSH and 12-bit CRC are transmitted. The resulting 25-bit frame (DSH+CRC) is encoded using a punctured, tailbiting, non-systematic, 32 state, rate 1/2 convolutional code to obtain the coded DSH (CDSH).

Receiver Operation – A received RLP block is declared invalid and discarded if it corresponds to a segment which does not lie within the receiver window or to one that was already successfully decoded. The receiver stores soft information in the form of quantized Euclidean distance metrics or quantized *a posteriori* log likelihood ratios for each valid data or parity block received. The memory requirements at the receiver are larger than in the fixed coding case since the receiver must store soft information for all data or parity blocks corresponding to a segment until decoding is successful. Soft decision decoding is attempted if all data blocks corresponding to a particular segment have been received. If all data blocks have been received and the segment has not yet been decoded, the receiver attempts decoding after receiving each additional valid parity block corresponding to the segment. The receiver stores the receipt status of data or parity

Direction	RLP segment size (octets)	CRC size (bits)	RLP block size (bits)	Modulation RLPs	Number of blocks per PDU	CDSH size (bits)	PDU size (bits)	Peak triple-rate throughput (kb/s)
Downlink	37	16	104	$\pi/4$ -DQPSK	2	48	256	29.6
				8-PSK	3	36	348	44.4
				16-level	4	96	512	59.2
Uplink	38	14	106	$\pi/4$ -DQPSK	2	50	262	29.7
				8-PSK	3	39	357	45.6
				16-level	4	100	524	59.4
Uplink (abbreviated)	30	15	85	$\pi/4$ -DQPSK	2	48	218	24
				8-PSK	3	36	327	36
				16-level	4	96	436	48

■ **Table 6.** IR CONTINUE PDU format sizes and corresponding peak triple-rate PDCH throughput in kilobits per second. No 16-level modulation has been specified at this time.

Direction	Modulation	Number of RLP blocks	PDU size (bits)	Rate 1/2 punctured convolutional encoder output (bits)	Peak triple-rate PDCH throughput (kb/s)
Downlink	$\pi/4$ -DQPSK	2	205	256	26.4
	8-PSK	3	293	348	39.6
	16-level	4	381	512	52.8
Uplink	$\pi/4$ -DQPSK	2	205	262	26.4
	8-PSK	3	293	357	39.6
	16-level	4	381	524	52.8
Uplink (abbreviated)	$\pi/4$ -DQPSK	2	173	218	21.6
	8-PSK	3	245	291	32.4
	16-level	4	317	436	43.2

■ **Table 7.** Fixed coding mode CONTINUE PDU format sizes and corresponding peak triple-rate PDCH throughput in kilobits per second. Each RLP block consists of 11 octets. No 16-level modulation has been specified at this time.

blocks corresponding to an undecoded segment.

When a segment is successfully decoded, the receiver updates its state information. The receiver generates full or partial bitmaps using this information when it provides ARQ feedback to the transmitter. For uplink transactions, bitmap feedback is still required in order to acknowledge successfully decoded IR segments since PCF only acknowledges the receipt of blocks transmitted in the previous time slot. The bitmap consists of a primary bitmap and two optional incremental redundancy bitmaps which indicate the following:

- Primary bitmap: This bitmap indicates the acknowledgment status (ACK or NAK) for each segment. Run length coding (as specified in ITU-T T.4 [14]) is applied in order to reduce the primary bitmap size whenever compression is possible.
- IR bitmap 1: For each segment negatively acknowledged by the primary bitmap, this bitmap indicates if all data blocks corresponding to that segment have been received.
- IR bitmap 2: If IR bitmap 1 indicates that all data blocks corresponding to a particular segment have not been received, this bitmap provides additional information indicating the receipt status of each data block corresponding to that segment.

Transmitter Operation – The transmitter procedures are somewhat more complex than the receiver and more state information must be stored for each RLP block at the transmitter. The RLP delivers data in-sequence by initially sending

just the RLP data blocks, D_{i1} , D_{i2} and D_{i3} , followed by sending additional RLP parity blocks, P_{ij} , $j = 1, 2, 3$ whenever the receiver fails to decode the RLP segment. The transmitter cannot discard the data and parity bits corresponding to a segment until it has received a positive acknowledgment from the receiver for the segment, hence the protocol operates by maintaining in a table, the receipt status of blocks that have been transmitted but not yet acknowledged. A block is not considered retransmittable unless it is explicitly negatively acknowledged, and sufficient time has elapsed to ensure that the status of the block is reflected in the feedback. Upon obtaining feedback from the receiver the transmitting protocol can update the table and decide on which data or parity block to schedule for transmission at the next opportunity.

Fixed Coding Mode RLP

FC mode transactions are established through a BEGIN PDU handshake as described in an earlier section. In the FC mode, RLP segments (or blocks) are of fixed length (shorter than the IR mode RLP segment). The CONTINUE PDU carries 2, 3, or 4 RLP blocks depending on the use of 4-level, 8-level (Table 7), or 16-level modulation, respectively, a CONTINUE PDU Header, and a 16-bit CRC. This ensures that blocks transmitted using a particular modulation can be retransmitted even if the modulation has changed.

The PDU header consists of the following:

- 1 bit indicating that it is a CONTINUE PDU
- 1 bit identifying the logical link (i.e., MLL0 or MLL1)

Field	Size (symbols)	Function
SYNC	14	Synchronization field (see IS-136 [1]). SYNC is always modulated using $\pi/4$ -DQPSK.
CSFP/CDFT	6	Superframe Phase (SFP) helps mobile stations find start of superframe. Data Field Type specifies the mode and modulation for the DATA field in the time slot. CSFP/CDFT is always modulated using $\pi/4$ -DQPSK.
PCF	12	Packet Channel Feedback provides subchannel feedback and subchannel assignment for the associated uplink time slots. PCF is always modulated using $\pi/4$ -DQPSK.
DATA	128 ($\pi/4$ -DQPSK) 116 (8-PSK)	Carries layer 1 SDU
PILOT	12 (8-PSK)	Provides phase reference for coherent demodulation of 8-PSK.
PRAMP	2	Facilitates downlink power control for adjacent time slots on same 30 kHz channel.

■ **Table 8.** Downlink field functions and sizes.

- A 1-bit poll indicator
- The 10-bit BSN of the first data block in the PDU

The PDU is encoded using a punctured, 32-state, rate 1/2, tailbiting, nonsystematic convolutional code. The puncturing (coding rate) at each modulation has been determined from the number of RLP blocks and the number of transmission symbols available for the modulation format.

Note that since the CRC bits are computed from two, three, or four RLP blocks in addition to other fields, all blocks within a PDU are lost if the PDU is not successfully decoded, and need to be recovered through retransmissions. Assigning a CRC sequence to each "small" RLP block is inefficient due to excessive overhead.

On the uplink, a CONTINUE PDU may contain an END control block which provides the last valid BSN for the transaction in lieu of a data block. This information helps the BS to schedule the required number of uplink time slot assignments in order to gracefully complete the transaction.

Implicit Addressing – Implicit addressing is carried out using the 7-bit AMI along with an 8-bit digital verification color code (DVCC) as the BS identifier. The CRC is computed using the PDU header, RLP blocks, AMI, and DVCC. Only the PDU header, RLP blocks, and CRC are provided to the convolutional encoder; the AMI and DVCC are not explicitly included.

Receiver Operation – A received RLP block is declared invalid and discarded if it does not lie within the receive window or was received previously. When a valid RLP segment is received, the receiver updates its state information. Full or partial bitmaps are generated using this information when the receiver provides ARQ feedback.

Transmitter Operation – The transmitter procedures are more complex, and more state information must be stored at the transmitter. The transmitter uses full or partial feedback information from the receiver in order to update its state information.


The data blocks to be transmitted are determined as follows. A block is not considered retransmittable unless it is explicitly NAKed, and sufficient time has elapsed to ensure that the status of the block is reflected in the feedback. Retransmittable blocks are given higher priority than new data blocks. Furthermore, the data blocks in each CONTINUE PDU are required to be in sequence. The earliest retransmittable data block is transmitted along with other retransmittable data blocks that follow in sequence. If there are no retransmittable data blocks, new data blocks within the window may be transmitted.

Radio Link Protocol End Procedures

For uplink transactions, the mobile station is required to commit to the end of the transaction if the buffer backlog is smaller than a threshold. This allows the BS to dynamically schedule uplink reservation opportunities without wasting any slots.

While initiating a transaction, the mobile station indicates the length of the transaction in the BEGIN PDU through the transaction size (TS) field, which can take the following values:

- TS = 0 indicates that the transaction ends with the BEGIN PDU.
- TS = 1, 2, ..., 62 indicates that there are TS data blocks in



■ Figure 12. The downlink time slot format.

the transaction.

- TS = 63 indicates that the transaction is unbounded.

In cases where the BEGIN PDU indicates an unbounded transaction, an explicit END control block must be transmitted when nearing the end of the transaction (i.e., when the amount of data in the transmit buffer falls below a threshold). The transmitter commits to the end of the transaction by providing the last valid sequence number for the transaction within the END block. For FC transactions, the END block is identified by an escape sequence and transmitted with other data blocks as part of a CONTINUE PDU. For IR transactions, the END block is identified by the PCBH and may be transmitted along with other data/parity blocks as part of a CONTINUE PDU. The END block is acknowledged through the end block received (EBR) indicator in the ARQ Status PDU. The transmitter and receiver continue retransmission procedures for recovering data only up to and including the last sequence number. Any new data that arrives in the transmit buffer requires the initiation of another transaction.

Radio Link Protocol

Partial Recovery and Error Handling

In both modes of operation, the transmitter and receiver can detect anomalies in protocol operation and trigger recovery action. Recovery occurs by aborting the current transaction and starting another one as necessary.

The transmitter RLP may abort the current transaction at any time by transmitting a BEGIN PDU to initiate a new transaction. Events such as the acknowledgment of previously acknowledged data, acknowledgment of data not yet transmitted, and excessive number of retransmission attempts are flagged as protocol anomalies by the transmitter RLP. If the number of anomalies that occur within a sliding time window exceeds a predetermined threshold, the transaction is aborted.

The receiver RLP may trigger a transaction abort by sending an ARQ Status PDU with a null bitmap or an abort indication. Events such as the receipt of data outside the receiver window or the receipt of multiple copies of the same data are flagged as protocol anomalies by the receiver RLP. If the number of anomalies that occur within a sliding time window exceeds a predetermined threshold, the transaction is aborted.

The transaction is also aborted if a long period of inactivity is detected at either the transmitter or receiver.


Time Slot Formats

The 30 kHz spectrum usage, symbol rate, and TDMA format (six time slots every 40 ms) are maintained as in the IS-136 standard [1] in order to minimize impact on existing infrastructure. However, the time slot contains new fields that are required for PDCH functions.


Downlink Time Slot Format

Field	Size (symbols)	Function
G	3	Guard field (see IS-136 [1]). Guard is always modulated using $\pi/4$ -DQPSK.
R	3	Ramp field (see IS-136 [1]) allows mobile stations to increase transmit power. Ramp is always modulated using $\pi/4$ -DQPSK.
PREAM	8	Preamble (see IS-136 [1]).
SYNC	14	Synchronization field (see IS-136 [1]). SYNC is always modulated using $\pi/4$ -DQPSK.
CDFT	3	Data Field Type specifies the mode and modulation for the DATA field in the time slot. CSFP/CDFT is always modulated using $\pi/4$ -DQPSK.
DATA	131 ($\pi/4$ -DQPSK) 119 (8-PSK)	Carries Layer 1 Service Data Unit.
PILOT	12 (8-PSK)	Provides phase reference for coherent demodulation of 8-PSK.

■ Table 9. Uplink field sizes and functions.



■ Figure 13. The uplink time slot format.



■ Figure 14. The uplink abbreviated time slot.

The downlink time slot formats for $\pi/4$ -DQPSK and coherent 8-PSK are shown in Fig. 12. The slot format is identical for both the IR and FC modes. The shaded fields are always modulated using $\pi/4$ -DQPSK. The different field sizes and functions are summarized in Table 8.

Uplink Time Slot Format

The uplink time slot formats for $\pi/4$ -DQPSK and coherent 8-PSK are shown in Fig. 13. The shaded fields are always modulated using $\pi/4$ -DQPSK, and the slot format is identical for both the IR and FC modes. The different field sizes and functions are summarized in Table 9.

Uplink Abbreviated Time Slot Format

Abbreviated uplink time slot formats are defined for packet data service deployments in large cells. The only difference from the regular uplink time slot format is that an abbreviated guard (AG) field of size 22 symbols is defined at the end of the time slot (Fig. 14). The DATA field contains only 109 symbols in the $\pi/4$ -DQPSK case and 97 symbols in the case of coherent 8-PSK. The G, R, PREAM, SYNC, and CDFT field

sizes are maintained as in the uplink time slot case, and these fields are modulated using $\pi/4$ -DQPSK.

Summary

This article has described the specification of a high-data-rate packet data service for the North American digital TDMA cellular standard, TIA/EIA IS-136 [1]. Details of the RLP, MAC, and physical layers are provided. The specification has the following features:

- Upper layers (L3 and above) consistent with GPRS
- Aggregation of multiple time slots
- Flexible allocation of uplink and downlink bandwidth
- Adaptive modulation
- Incremental redundancy and fixed coding modes

The standard is likely to be approved by the time this article is published. Implementations are expected in early 2000.

References

- [1] TIA/EIA IS-136, "TDMA Air Interface-Mobile-Station-Base Station Compatibility," Arlington: TIA, 1995.
- [2] ANSI-41, "Cellular Radio Telecommunications Intersystem Operation," Mar. 1997.
- [3] ETSI GSM 03.60: "Digital cellular telecommunications system (Phase 2+); General Packet Radio Service (GPRS); Service Description; Stage 2."
- [4] ETSI GSM 04.64: "Digital cellular telecommunications system (Phase 2+); General Packet Radio Service (GPRS); Logical Link Control (LLC)."
- [5] ETSI GSM 04.65: "Digital cellular telecommunications system (Phase 2+); General Packet Radio Service (GPRS); Subnetwork Dependent Convergence Protocol (SNDCP)."
- [6] ETSI GSM 08.16: "Digital cellular telecommunications system (Phase 2+); General Packet Radio Service (GPRS); Base Station System (BSS) — Serving GPRS Support Node (SGSN) interface; Network Service."
- [7] ETSI GSM 08.18: "Digital cellular telecommunications system (Phase 2+); General Packet Radio Service (GPRS); Base Station System (BSS) — Serving GPRS Support Node (SGSN); BSS GPRS Protocol (BSSGP)."
- [8] ETSI GSM 09.60: "Digital cellular telecommunications system (Phase 2+); General Packet Radio Service (GPRS); GPRS Tunneling Protocol (GTP) across the Gn and Gp Interface."
- [9] K. Balachandran, S. Kadaba, and S. Nanda, "Channel Quality Estimation Over Mobile Radio Channels With Applications," *Proc. Conf. Info. Sci. and Sys.*, Princeton, NJ, Mar. 1998.
- [10] K. Balachandran, S. Kadaba, and S. Nanda, "Channel Quality Estimation and Rate Adaptation for Cellular Mobile Radio," to appear, *IEEE JSAC*.
- [11] B. T. Doshi, R. P. Ejzak, and S. Nanda, "A Retransmission Scheme for Circuit-Mode Data on Wireless Links," *IEEE JSAC*, 1994.

- [12] TIA/EIA IS-130, "TDMA Air Interface — Mobile-Station-Base Station Compatibility — Data Services Option," Arlington: TIA, 1995.
- [13] TIA/EIA 136-332, "Packet Data Service — Medium Access Control Layer," Mar. 1999.
- [14] ITU-T Rec. T.4, "Standardization of Group 3 Facsimile Terminals for Document Transmission," 1996.

Additional Reading

- [1] TIA/EIA 136-123, "Impacts on DCCH Layer 3 From GPRS-136," Mar. 1999.
- [2] TIA/EIA 136-330, "Packet Data Service - Overview," Mar. 1999.
- [3] TIA/EIA 136-331, "Packet Data Service - Physical Layer," Mar. 1999.
- [4] TIA/EIA 136-335, "Packet Data Service - Radio Resource Management," Mar. 1999.
- [5] TIA/EIA 136-336, "Packet Data Service - Mobility Management," Mar. 1999.
- [6] TIA/EIA 136-932, "Packet Data Service - Stage 2 Description," Mar. 1999.
- [7] AT&T and Lucent Technologies, "IS-136+ Data Services," TIA TR45.3.2.03.09.15.
- [8] Lucent Technologies, "Fixed Coding and Incremental Redundancy Modes with 4 and 8 Level Modulation for 136+ Packet Data," TIA TR45.3.2.03.09.16.
- [9] K. Balachandran et al., "Medium Access Control Feedback Mechanism for Uplink Resource Management in Wireless TDMA Packet Data Systems," *Proc. IEEE VTC*, May 1999.
- [10] K. Balachandran et al., "Medium Access Control and Radio Resource Management for Packet Data Services Over IS-136 Channels," *Proc. IEEE VTC*, May 1999.
- [11] IETF RFC 1661, "The Point-to-Point Protocol," July 1994.

Biographies

KRISHNA BALACHANDRAN (krishnab@lucent.com) received a B.E. (Hons.) degree in electrical and electronics engineering from Birla Institute of Technology and Science, Pilani, India, in 1992, and an M.S. in computer and systems engineering, an M.S. in mathematics, and a Ph.D. in computer and systems engineering from Rensselaer Polytechnic Institute, Troy, New York, in 1994, 1995, and 1996, respectively. He held graduate teaching and research assistantship positions at RPI from 1993 to 1996, and was recognized as a Master Teaching Fellow by the dean of the graduate school in 1995. He spent the summers of 1993 and 1994 at Pacific Communication Sciences, Inc., San Diego, California, where he worked on equalization problems related to digital cellular radio. Since 1996, he has been a member of technical staff in the Performance Analysis Department at Bell Laboratories, Lucent Technologies, Holmdel, New Jersey, where he has been working on algorithm development, protocol design, and performance analysis for cellular systems. He has published several papers and contributed significantly to standardization activities in the TIA and ETSI.

RICHARD EJZAK (ejzak@lucent.com) received a B.S. degree in electrical engineering from the University of Michigan in 1976, and M.S. and Ph.D.

degrees in computer science from Northwestern University in 1982 and 1984, respectively. He has been with Lucent Technologies since 1983 and is now a distinguished member of technical staff in the Systems Engineering organization of the Network Wireless Systems Business Unit. Since joining the Network Wireless Systems Business Unit of Lucent Technologies in 1991, he has focused on various aspects of wireless data systems design, including system architecture, protocol design, and system performance. He has contributed to the development of several TDMA, CDMA, and GSM standards for data services. He was the lead architect in the initial development of infrastructure equipment for CDPD, CDMA, and TDMA data services. He holds four U.S. patents.

SANJIV NANDA (nanda@lucent.com) received a B. Tech. degree in electrical engineering from the Indian Institute of Technology, Kanpur, in 1983, an M.S. degree in mathematics in 1986, and M.S. and Ph.D. degrees in electrical engineering in 1985 and 1988, respectively, all from the Rensselaer Polytechnic Institute, Troy, New York. During 1988–1990 he was with the Wireless Information Network Laboratory (WINLAB), Rutgers University, Piscataway, New Jersey. Since 1990 he has been with the Performance Analysis Department at Bell Laboratories, Lucent Technologies, Holmdel, New Jersey, where he is currently technical manager. At Lucent he has worked on many aspects of cellular system design, including system architecture, protocol design, system performance, and link performance. He has made contributions to Lucent's wireless products through algorithms, performance analysis, and participation in standards and architecture teams. He has published several papers and holds eight U.S. patents. He co-authored a paper on AAL-2 that was chosen as the Best Paper in Networking in the *Bell Labs Technical Journal* in 1997. His paper on PRMA with Goodman and Timor won the IEEE Vehicular Technology Society's Jack Neubauer Award for the best systems paper published in *IEEE Transactions on Vehicular Technology* in 1991.

STANISLAV VITEBSKIY (svite@lucent.com) received a Dipl.Eng.(Hons) degree from Moscow Institute of Electronics Engineering (MIET), Russia, in 1992, and M.S. and Ph.D. degrees in electrical engineering from Polytechnic University, Brooklyn, New York, in 1994 and 1996, respectively. From 1995 to 1996 he was a research assistant at Duke University, Durham, North Carolina, working on ultra-wideband short-pulse radar modeling. From 1996 to 1997 he was a senior engineer at Harris Corporation, Rochester, New York, where he designed and evaluated MAC protocols for wireless networks. In 1997 he joined the Wireless Networks Group, Lucent Technologies, Whippany, New Jersey, as a member of technical staff. His current interests include algorithm development, protocol design, and performance analysis for wireless voice and high-speed data systems.

SHIV SETH (sseth@lucent.com) received a B.S. degree from Punjab University, Chandigarh, India, an M.S. degree in mathematics from Kanpur University, India, and an M.S. degree in computer science from Washington State University, Pullman. He was a systems programmer at Sperry Corporation for three years. Since 1986 he has been a member of technical staff at Lucent Technologies, Naperville, Illinois. He is interested in packet data switching for wireless communication systems, video conferencing, and automatic speech recognition.