78-6088-02 10/25/99

Carrierless AM/PM Rate Adaptive Digital Subscriber Line Interface Specification

1. Introduction

This document provides the specification for operating Cisco's carrierless amplitude and phase modulation (CAP) rate adaptive digital subscriber line (RADSL) based broadband access system subscriber interface.

2. Scope

The document is organized into sections that follow the Open System Interconnection (OSI) reference model, as shown in Figure 1. Only Layer 1, the physical layer, and Layer 2, the data-link layer of the OSI reference model are defined for the purposes of this document.

Figure 1 OSI Reference Model

Layer 7	Application
Layer 6	Presentation
Layer 5	Session
Layer 4	Transport
Layer 3	Network
Layer 2	Data link
Layer 1	Physical

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA

2.1 Physical Layer

Following is a list of the physical layer attributes:

- Describes the transmission technique used to support the simultaneous transport of plain old telephone service (POTS) and the upstream and downstream channels of the rate adaptive digital signals on a single twisted pair
- Defines the line codes and spectral composition of signals transmitted by the remote termination unit at the central office (RTU-C) and the one at the remote location (RTU-R)
- Specifies the receive signals at both the RTU-R and RTU-C
- Describes electrical and mechanical specifications of the network interface
- Defines the exchange of information sequences during start-up to establish the system configuration and transmission link

2.2 Data-Link Layer

Following is a list of the data-link layer attributes:

- Describes a logical interface to the Asynchronous Transfer Mode—transmission convergence (ATM-TC) and embedded operations channel—transmission convergence (EOC-TC) sublayers
- Describes the ATM layer virtual path identifier/virtual channel identifier (VPI/VCI) mapping

3. Definitions, Acronyms, Symbols, Abbreviations

This section includes the definitions, abbreviations, acronyms, and symbols that appear throughout this document.

3.1 Definitions

Following is a list of common terms found in this document:

- Bridge taps—Sections of unterminated twisted pairs that are connected in parallel across the pairs under consideration
- Downstream—RTU-C to RTU-R direction (network to customer direction)
- Upstream—RTU-R to RTU-C direction (customer to network direction)
- Loading coils—Inductors that are placed in series with the cable at regular intervals to improve the voice-band response
- POTS splitter—A low-pass/high-pass pair of filters that separate high (RADSL) and low (POTS) frequency line signals
- Voice band—0.3 to 3.4 kHz based on the Bellcore local switching system generic requirement (LSSGR) and transmission system generic requirement (TSGR) specifications

Table 1 lists the common acronyms, symbols, and abbreviations found throughout this document.

Table 1 Acronyms, Symbols, and Abbreviations

Acronyms	Description
⊕	modulo 2 addition; logical exclusive
AAL5	ATM adaptation layer 5
ADSL	asymmetric digital subscriber line
AM	amplitude modulation
ANSI	American National Standards Institute
ASC	Advanced Intelligent Network Switch Capabilities
ATM	Asynchronous Transfer Mode
ATU-C	ADSL transmission unit—central office
BER	bit error rate
BERT	bit error rate tester
CPE	customer premises equipment
dB	decibels
DOH	Digital Off-Hook
DPDU	data-link protocol data unit
DSL	digital subscriber line
DSLAM	digital subscriber line access multiplexer
EOC	embedded operations channel
FEC	forward error correction
GF	Galois field
HEC	header error control
IPC	interprocess communication
ISDN	Integrated Services Digital Network
KB	kilobyte
kbps	kilobits per second
mA	milliampere
Mbps	megabits per second
n/a	not applicable, not supported
OSI	Open System Interconnection
PM	phase modulation
PMD	physical medium dependent
POTS	plain old telephone service
ppm	parts per million
PSD	power spectral density
PSTN	Public Switched Telephone Network
PVC	permanent virtual circuit
PVP	permanent virtual path

Table 1 Acronyms, Symbols, and Abbreviations (continued)

Acronyms	Description
PTC	physical transmission convergence
QAM	quadrature amplitude modulation
RADSL	rate adaptive digital subscriber line
RS	Reed-Solomon
RTU-C	remote termination unit—central office
RTU-R	remote termination unit—remote terminal
SAR	segmentation and reassembly
SSS	self-synchronizing scrambler
SM	service module
T	interface(s) between RTU-R and RTU-C
TBD	to be determined
TC	transmission convergence
T-SM	interface(s) between RTU-R and the SM(s)
U-C	loop interface—central office
U-R	loop interface—remote terminal
UNI	User-Network Interface
V	logical interface between RTU-C and a digital network element such as one or more switching systems
VC	virtual circuit
VCC	virtual circuit connection
VPC	virtual path connection

4. System Reference Model

The system reference model shown in Figure 2 illustrates the physical blocks required to provide RADSL service.

Figure 2 RADSL Functional Reference Model

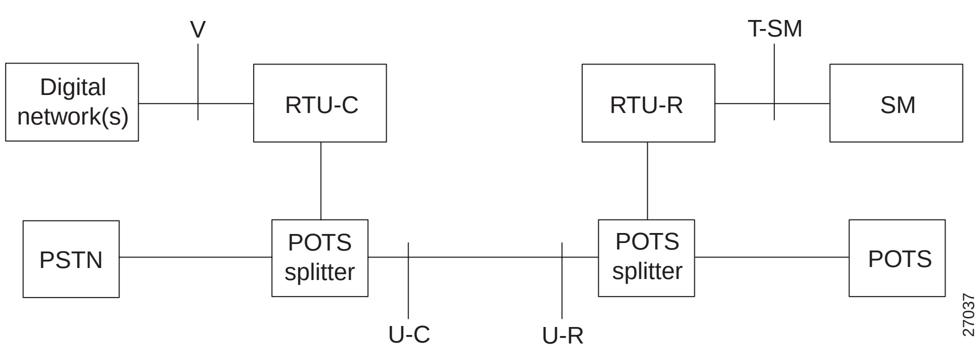

The following notes apply to Figure 2:

1 V interface is defined in terms of logical rather than physical functions.

2 V interface can consist of an interface or interfaces to one or more switches.

3 Splitter function, or a portion thereof, can be integrated within the RTU or might not be required.

5. Physical Layer—CAP

The following section discusses the CAP physical layer, which includes reference models, forward error correction (FEC), line signals, scrambling method, and signal constellations and trellis encoding.

5.1 Functional Transmitter Reference Model

For the upstream and downstream channels, the sequence of the signal processing functions defined in this specification are shown in Figure 3. The CAP physical transmission convergence (PTC) layer consists of the Reed-Solomon (RS) encoder with interleaving. The CAP physical medium dependent (PMD) layer consists of the remaining core modulation and demodulation functions including

- scrambling
- trellis encoder
- channel precoder
- modulation
- demodulation

The downstream channel contains the RS encoder with interleaving; support for the RS encoder in the upstream channel is not supported. The output bit sequence of the interleaved RS encoder is then fed to a frame-locked scrambler. The trellis encoder and Tomlinson precoder convert the serial bit stream into two dimensional symbols for modulation by the CAP transmitter.

The upstream channel has the same signal processing functions as the downstream channel, except that the RS encoder is not supported (shown in Figure 3 as a dotted box).

Figure 3 RTU-C and RTU-R Transmitter Reference Diagram

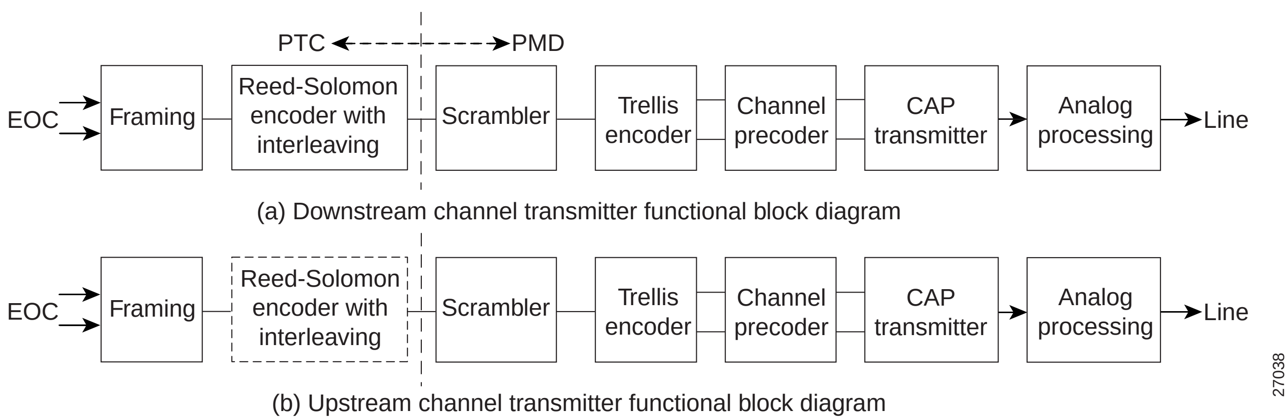

5.2 Forward Error Correction

The two types of FEC are RS encoding and interleaving. Both are discussed in the following paragraphs.

A two-symbol (byte) error-correcting RS encoder is defined for the CAP-based RADSL system. In this case, four redundant check bytes are appended to the K message bytes $m_0, m_1, \ldots, m_{k-1}$ to form an RS encoder word of size $N = K + 4$. The check bytes are computed from the message bytes using the equation:

$$C(x) = M(x) \cdot x^4 \bmod g(x)$$

where:

$M(x) = m_0 \cdot x^{k-1} \oplus m_1 \cdot x^{k-2} \oplus \ldots \oplus m_{k-1} \cdot x \oplus m_k$ is the message polynomial,

$C(x) = c_0 \cdot x^3 \oplus c_1 \cdot x^2 \oplus c_2 \cdot x^1 \oplus c_3$ is the check polynomial, and

$g(x) = (x \oplus \alpha^1)(x \oplus \alpha^2)(x \oplus \alpha^3)(x \oplus \alpha^4)$ is the generator polynomial.

The RS encoder is performed in the Galois Field GF(2^8), where α is the primitive element that satisfies the primitive binary polynomial $p(x) = x^8 \oplus x^4 \oplus x^3 \oplus x^2 \oplus 1$. The data byte is identified with the Galois Field element $d_7\alpha^7 \oplus d_6\alpha^6 \oplus \ldots \oplus d_0$.

The interleaving method for the RS encoder uses an *implied* convolutional interleaving method that inserts the parity symbols into appropriate locations within the original data sequence, and the resulting encoder words are constructed via the convolutional interleaving rules. The input data (message) symbols stay in their original sequence when transmitted on the subscriber line. Figure Figure 4 shows a functional transmitter block diagram of the FEC block with interleaving. The effect of interleaving to a depth D is provided by entering the input data symbols into the appropriate encoder in a sequence corresponding to the interleaving rules; convolutional interleaving is recommended for efficient use of memory. Each encoder has a rate of (N,K), where N is the encoder-word size and K is the number of data symbols (message size) in the encoder word. The value of N and K can be selected at start-up; the default values of N and K are 68 and 64 respectively. The multiplexer inserts the parity symbols of each encoder into the appropriate locations within the original transmit data sequence.

Figure 4 Transmitter Structure for the Reed-Solomon Encoder Interleaving

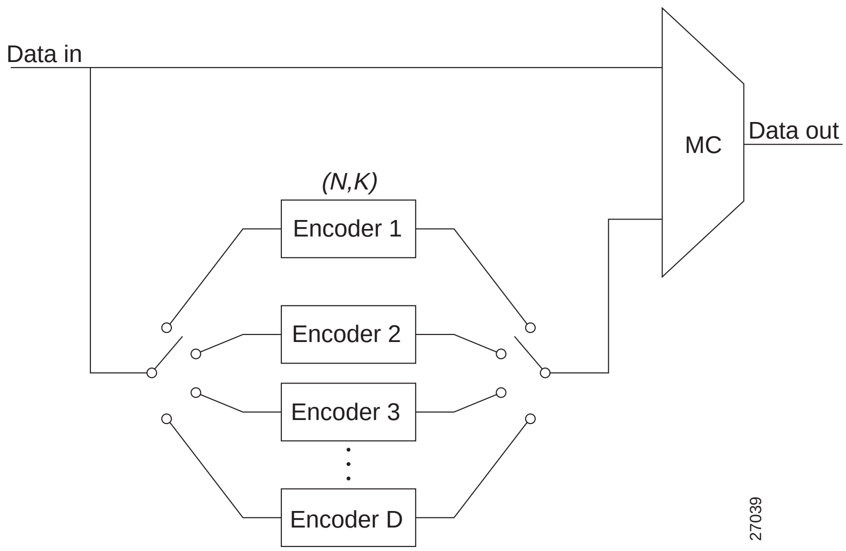

Figure 5 shows the encoder-word construction with implied interleaving. The symbol *m* represents the time index of the data block that contains both data and parity symbols. The data symbols are sequentially written into the rows of the matrix and are also read from the rows of the matrix in the same sequence. The parity symbols $P_j(k), j = 1, 2 \ldots, D$ are inserted into the matrix using the following rules for the case where $N > D$:

1. Define parameter $L_j = (j-1) \times \left[\frac{N}{D}\right]$, where $j = 2, 3, \ldots, D$ and $L_1 = N$, and where [*x*] denotes the greatest integer function of the argument *x*. The parameter L_j defines the row location of the $(N-K)$ th parity symbol for the encoder word $P_j(k)$ where the argument $k = 1, 2, \ldots, N-K$ represents the parity symbol index for the *j*th encoder word.
2. The parity symbols $P_j(k)k = 1, 2, \ldots, N-K$ of encoder word *j* are placed in matrix locations $A_m(i, j)$, where $i = L_j - (N-K) + k$ is the row location and *j* is the matrix column that also represents the encoder word.

In this specification, the value of *N* is always greater than *D*. The maximum value of interleave memory is 3.8 KB and the maximum value for *D* is $D_{\max} = \left[\frac{3.8KB}{N}\right]$, where [*x*] is the greatest integer function of argument *x*, and *N* is the encoder-word size.

With the convolutional assignment of the parity symbols, the encoder words are constructed between adjacent data blocks. Because the index *m* increases with increasing time, a encoder word with parity symbols in block *m* has its associated data symbols in blocks *m* and $m-1$ for encoder words $j = 2, 3, \ldots, D$. For encoder word $j = 1$, the whole encoder word resides in the current block *m*.

Figure 5 EnCoder Word Construction

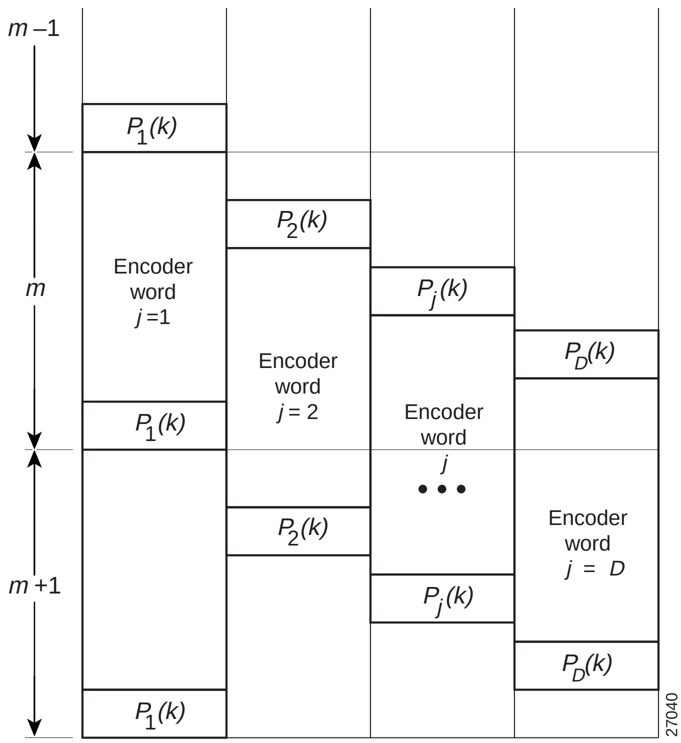

5.3 Line Signals

This section covers downstream and upstream channel symbol rates, constellation sizes, and bit rates. Transmit signal characteristics, including spectral placement and the power spectral density (PSD) mask, are also discussed.

The predefined mandatory downstream channel symbol rates follow:

- 136 kilobaud ± 25 parts per million (ppm)
- 340 kilobaud ± 25 ppm (required in the start-up procedure)
- 680 kilobaud ± 25 ppm
- 952 kilobaud ± 25 ppm

The predefined mandatory upstream channel symbol rate is 136 kilobaud ± 25 ppm (required in the start-up procedure).

The predefined optional upstream channel symbol rates follow:

- 17 kilobaud ± 25 ppm
- 34 kilobaud ± 25 ppm
- 68 kilobaud ± 25 ppm

Support for *variable* symbol rates (that is, other rates beyond the preceding predefined rates) can be defined according to the following rule:

The downstream symbol rate clock is defined as $f_{\mathrm{Baud,Down}} = \frac{34.56 \cdot N_0}{D_0 D_1}$ in MHz, where N_0 and D_1 are integers in the range from 1 to 256, D_0 is an integer in the range from 1 to 48, and $\frac{\frac{1}{2.3} < N_0}{D_0 < 3.2}$. The upstream symbol rate clock is defined as $f_{\mathrm{Baud},Up} = \frac{2 \cdot f_{\mathrm{Baud,Down}}}{D_2}$ in MHz, where D_2 is an integer in the range from 2 to 32. For the variable symbol rate case, the parameters N_0 and D_0 to D_2 allow downstream symbol rate selection in the range from approximately 64 kilobaud to 1088 kilobaud, and upstream symbol rate selection in the range from approximately 4 kilobaud to 136 kilobaud. Proper selection of N_0, D_0, D_1, and D_2 allow provisioning of arbitrarily small step sizes. The maximum downstream channel symbol rate is 1088 kilobaud ± 25 ppm, and the maximum upstream channel symbol rate is 136 kilobaud ± 25 ppm.

The symbol rates are selected in the predefined mode with all parameters passed during transceiver start-up as described in the “CAP Start-Up Procedures” section on page 21.

The system supports constellation sizes up to 256 points with an integer number of bits per symbol. Also, in the downstream direction with a 256-point constellation, the RS FEC can be bypassed to provide more payload data. This condition is referred to as a 256 unencoded constellation (256 UC). To increase reach in the upstream, a special case of the 8-point constellation, 8 extended range (8er), can be used. For 8er, the transmitter sends the same symbol three consecutive times. The receiver then resolves the intended symbol by majority rule from the three transmitted symbols.

Not all constellations are mandatory for each baud rate because of overlap of the data rates. Table 2 defines which constellations are supported for the mandatory baud rates. Table 3 shows constellations for the optional upstream baud rates.

Table 2 Supported Constellations per Mandatory Baud Rate

Downstream Baud Rate	Upstream Baud Rate	Supported Downstream Constellations	Supported Upstream Constellations
136 kilobaud	136 kilobaud	256 UC, 256, 128, 64, 32, 16, 8	256 UC, 256, 128, 64, 32, 16, 8, 8er
340 kilobaud	136 kilobaud	256 UC, 256, 128, 64, 32, 16, 8	256 UC, 256, 128, 64, 32, 16, 8, 8er
680 kilobaud	136 kilobaud	256 UC, 256, 64, 16	256 UC, 256, 64, 16
952 kilobaud	136 kilobaud	256 UC, 256, 64, 16	256 UC, 256, 64, 16

Table 3 Supported Constellations per Optional Upstream Baud Rate

Upstream Baud Rate	Supported Upstream Constellations
17 kilobaud	256 UC, 256, 128, 64, 32, 16, 8, 8er
34 kilobaud	256 UC, 256, 128, 64, 32, 16, 8, 8er
68 kilobaud	256 UC, 256, 128, 64, 32, 16, 8, 8er

The data symbol bit rates are determined by the symbol rate, and the rate of the trellis and RS encoders. The bit rate is computed by $R_b = m \cdot \left(\frac{N}{K}\right) \cdot f_{\mathrm{Symbol}}$, where *m* is the number of data bits per symbol in the constellation, *N* and *K* are the encoder word and information field sizes respectively of the RS encoder and f_{Symbol} is the symbol rate of the respective upstream and downstream channels. The downstream channel bit rates range up to approximately 8 Mbps depending on the RS encoder-word size, and the upstream channel bit rates range up to approximately 1 Mbps.

Table 4 defines the downstream user data rates given the downstream baud rate and supported constellations.

Table 4 Payload Bit Rates per Constellation (kbps) with (68, 64) RS Encoding for the Mandatory Downstream Baud Rate

	256 UC	**256**	**128**	**64**	**32**	**16**	**8**
136 kilobaud	1024	896	768	640	512	384	256
340 kilobaud	2560	2240	1920	1600	1280	960	640
680 kilobaud	5120	4480	n/a	3200	n/a	1920	n/a
952 kilobaud	7168	6272	n/a	4480	n/a	2688	n/a

Table 5 defines the upstream user payload data rates given the mandatory upstream baud rate and supported constellations.

Table 5 Payload Bit Rates per Constellation (kbps) for the Mandatory Upstream Baud Rates

	256 UC	**256**	**128**	**64**	**32**	**16**	**8**	**8er**
136 kilobaud	1088	952	816	680	544	408	272	90.7

Table 6 defines the upstream user payload data rates given the optional upstream baud rate and supported constellations.

Table 6 Payload Bit Rates per Constellation (kbps) for the Optional Upstream Baud Rates

	256 UC	**256**	**128**	**64**	**32**	**16**	**8**	**8er**
17 kilobaud	136	119	102	85	68	51	34	11.3
34 kilobaud	272	238	204	170	136	102	68	22.7
68 kilobaud	544	476	408	340	272	204	136	45.3

Figure 6 shows the representative spectral placements of the upstream and downstream channels, which are placed at frequencies above the POTS channel. The upstream channel occupies the lower frequencies and the downstream channel occupies the higher frequencies for the DSL line signals.

Figure 6 Spectral Placement of Upstream and Downstream Channels

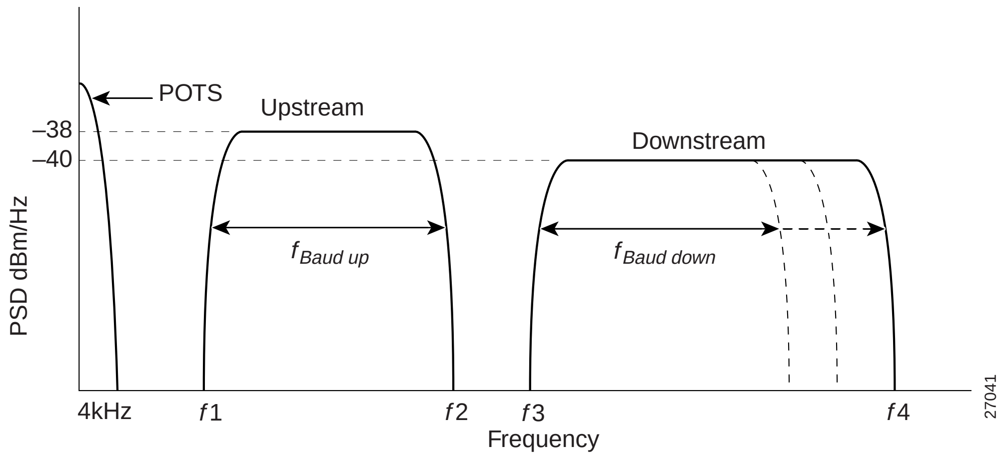

In Figure 6, frequencies *f1* and *f2* represent the start and stop frequencies for the upstream channel spectrum, and *f3* and *f4* represent the start and stop frequencies for the downstream spectrum. The values for the start and stop frequencies are defined as where the passband is down to 27 dB from the nominal PSD (–40 dBm/Hz in the downstream and –38 dBm/Hz in the upstream). The start and stop frequencies are functions of both the upstream and downstream baud rates.

The two distinct starting frequencies for the upstream are 25 kHz and 35 kHz. The 25 kHz starting frequency is the default and is predominant because of the increase in upstream performance in the presence of bridge taps. Also, because of the POTS band proximity to the 25 kHz starting frequency, a higher order high-pass filter is used compared to the 35 kHz starting frequency. Therefore, ADSL interference into the POTS band is minimized. The 25 kHz starting frequency values are displayed in Table 7.

The 35 kHz starting frequency is a supported feature, but it is no longer being used as the default upstream starting frequency. The 35 kHz starting frequency values are displayed in Table 8.

Table 7 Spectral Start and Stop Frequencies for an Upstream Starting Frequency of 25 kHz

Downstream Baud Rate	Downstream		Upstream 17 kilobaud		Upstream 34 kilobaud		Upstream 68 kilobaud		Upstream 136 kilobaud	
	Start Freq. f3 (kHz)	Stop Freq. f4 (kHz)	Start Freq. f1 (kHz)	Stop Freq. f2 (kHz)	Start Freq. f1 (kHz)	Stop Freq. f2 (kHz)	Start Freq. f1 (kHz)	Stop Freq. f2 (kHz)	Start Freq. f1 (kHz)	Stop Freq. f2 (kHz)
136 kilobaud	240	406	TBD	TBD	TBD	TBD	23	106	22	185
340 kilobaud	241	655	TBD	TBD	TBD	TBD	23	106	19	185
680 kilobaud	236	1014	n/a	n/a	n/a	n/a	n/a	n/a	14	193
952 kilobaud	245	1376	n/a	n/a	n/a	n/a	n/a	n/a	16	177

Table 8 Spectral Start and Stop Frequencies for an Upstream Starting Frequency of 35 kHz

Downstream Baud Rate	Downstream		Upstream 17 kilobaud		Upstream 34 kilobaud		Upstream 68 kilobaud		Upstream 136 kilobaud	
	Start Freq. f3 (kHz)	Stop Freq. f4 (kHz)	Start Freq. f1 (kHz)	Stop Freq. f2 (kHz)	Start Freq. f1 (kHz)	Stop Freq. f2 (kHz)	Start Freq. f1 (kHz)	Stop Freq. f2 (kHz)	Start Freq. f1 (kHz)	Stop Freq. f2 (kHz)
136 kilobaud	240	406	TBD	TBD	TBD	TBD	23	106	32	193
340 kilobaud	241	655	TBD	TBD	TBD	TBD	23	106	21	207
680 kilobaud	236	1014	n/a	n/a	n/a	n/a	n/a	n/a	17	202
952 kilobaud	245	1340	n/a	n/a	n/a	n/a	n/a	n/a	21	200

The PSD mask for the downstream channel has an upper limit of $\frac{-40\text{dBm}}{\text{Hz}}$ in the nominal passband region with no variation exceeding $\frac{-37\text{dBm}}{\text{Hz}}$. The tolerance does not infer such variability in the total signal power, but rather permits the fluctuation of individual spectral components within the passband. Total power in the passband is dependent on the symbol rate and does not exceed the limit imposed by the following equation:

$$P_{TX} = \frac{-40\text{dBm}}{\text{Hz} + 10\text{Log}(\text{SymbolRate})}$$

where P_{TX} is the total power expressed in decibels per milliwatt (dBm) and Symbol Rate is expressed in baud. The maximum downstream total power is estimated at +21.5 dBm.

The PSD mask for the upstream channel has an upper limit of $\frac{-38\text{dBm}}{\text{Hz}}$ nominal with no variation exceeding $\frac{-35\text{dBm}}{\text{Hz}}$. The tolerance does not infer such variability in the total signal power, but rather permits the fluctuation of individual spectral components within the passband. Total power in the passband is dependent on the symbol rate and does not exceed the limit imposed by the following equation:

$$P_{TX} = \frac{-38\text{dBm}}{\text{Hz} + 10\text{Log}(\text{SymbolRate})}$$

where P_{TX} is the total power expressed in dBm and Symbol Rate is expressed in baud.

The maximum upstream total power is estimated at $+13.4\text{dBm}$.

The PSD mask is shown in Figure 7.

Figure 7 **CAP RADSL PSD Mask**

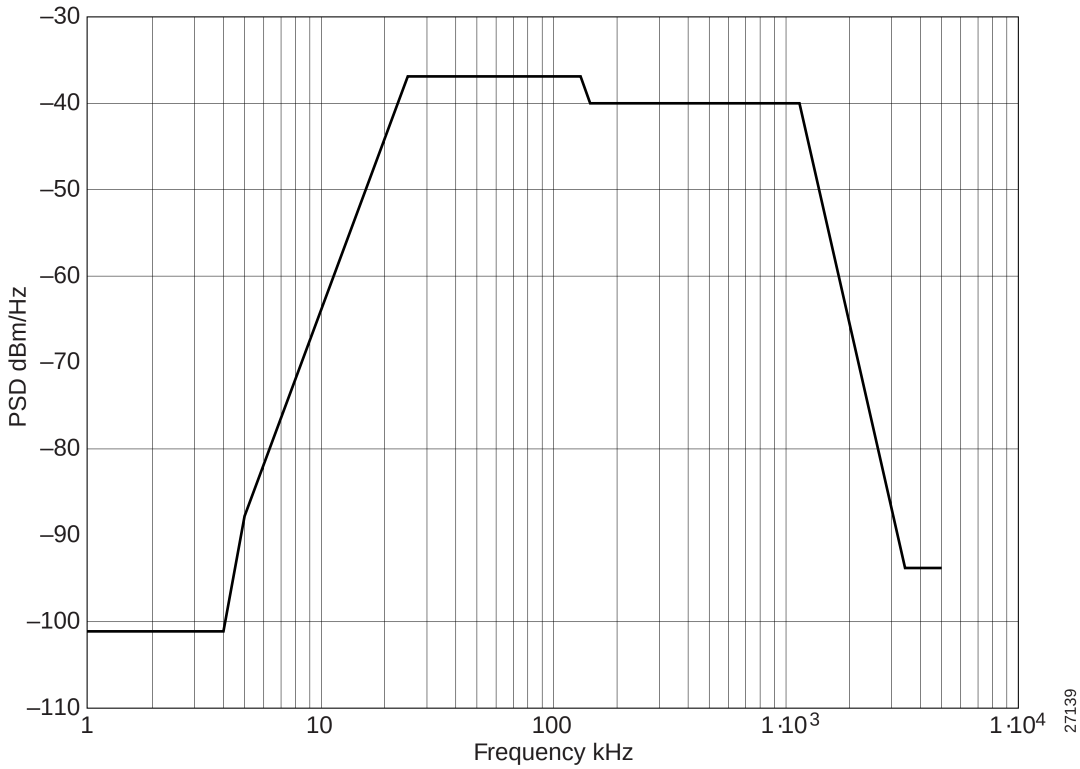

5.4 Scrambling Method

The scrambler/descrambler, included in each transceiver, is different in the two directions of transmission. The generator polynomials are as follows:

Customer premises transceiver (RTU-R) = $1 \oplus x^{-18} \oplus x^{-23}$

Exchange transceiver (RTU-C) = $1 \oplus x^{-5} \oplus x^{-23}$

Figure 8 shows the scramblers and descramblers as they operate during start-up in the self-synchronizing mode. At the transmitter, the scrambler effectively divides (modulo 2) the bits sequence by the generator polynomial. The coefficients of the quotients for this division, taken in descending order, form the data sequence that appears at the output of the data scrambler. At the receiver, the received bit sequence is multiplied (modulo 2) by the polynomial to recover the original bit stream.

During data transfer, the scramblers are locked and the scrambled sequence is added (modulo 2) at the transmitter and subtracted (modulo 2) at the receiver as shown in Figure 9. The transfer from the self-synchronizing mode to the locked mode occurs with the transmit data being all ones. The transfer to locked mode does not require synchronization of the transfer at the two ends.

Figure 8 **CAP RADSL Scrambler and Descrambler During Start-Up Mode**

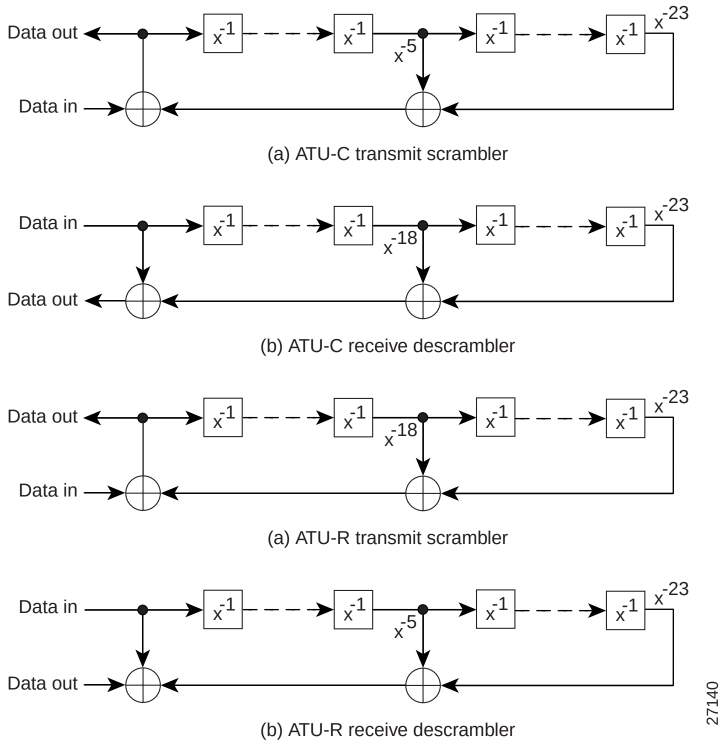

Figure 9 CAP RADSL Scrambler and Descrambler During Data Mode

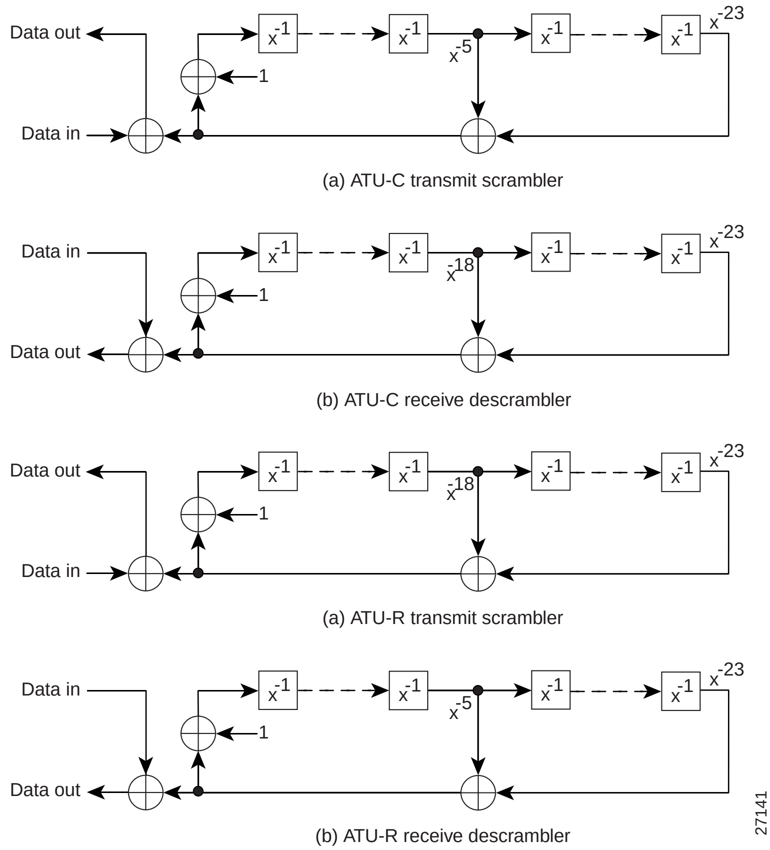

5.5 Signal Constellations and Trellis Encoder

This section covers signal constellations during start-up and the trellis encoder used during data mode.

An unencoded mode is used during the start-up procedure, as described in the “CAP Start-Up Procedures” section on page 21. An unencoded one-dimensional signal constellation is shown in Figure 10. Figure 11 shows the unencoded 16-CAP signal constellation. Figure 12 shows an unencoded 256-CAP signal constellation. In all cases, Z0 is the least significant bit.

Figure 10 Unencoded One-Dimensional 2-Point Signal Constellation

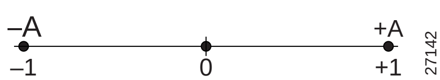

Figure 11 Unencoded 16-CAP Signal Constellation

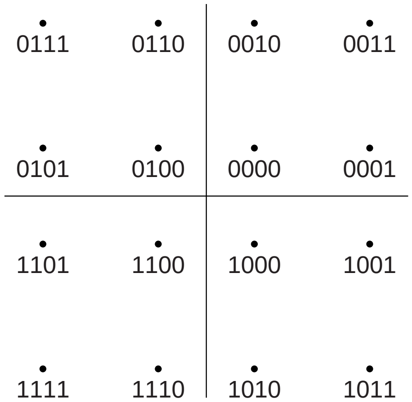

Encoder point designations are Z3 to Z0

27143

Figure 12 Unencoded 256-CAP Signal Constellation

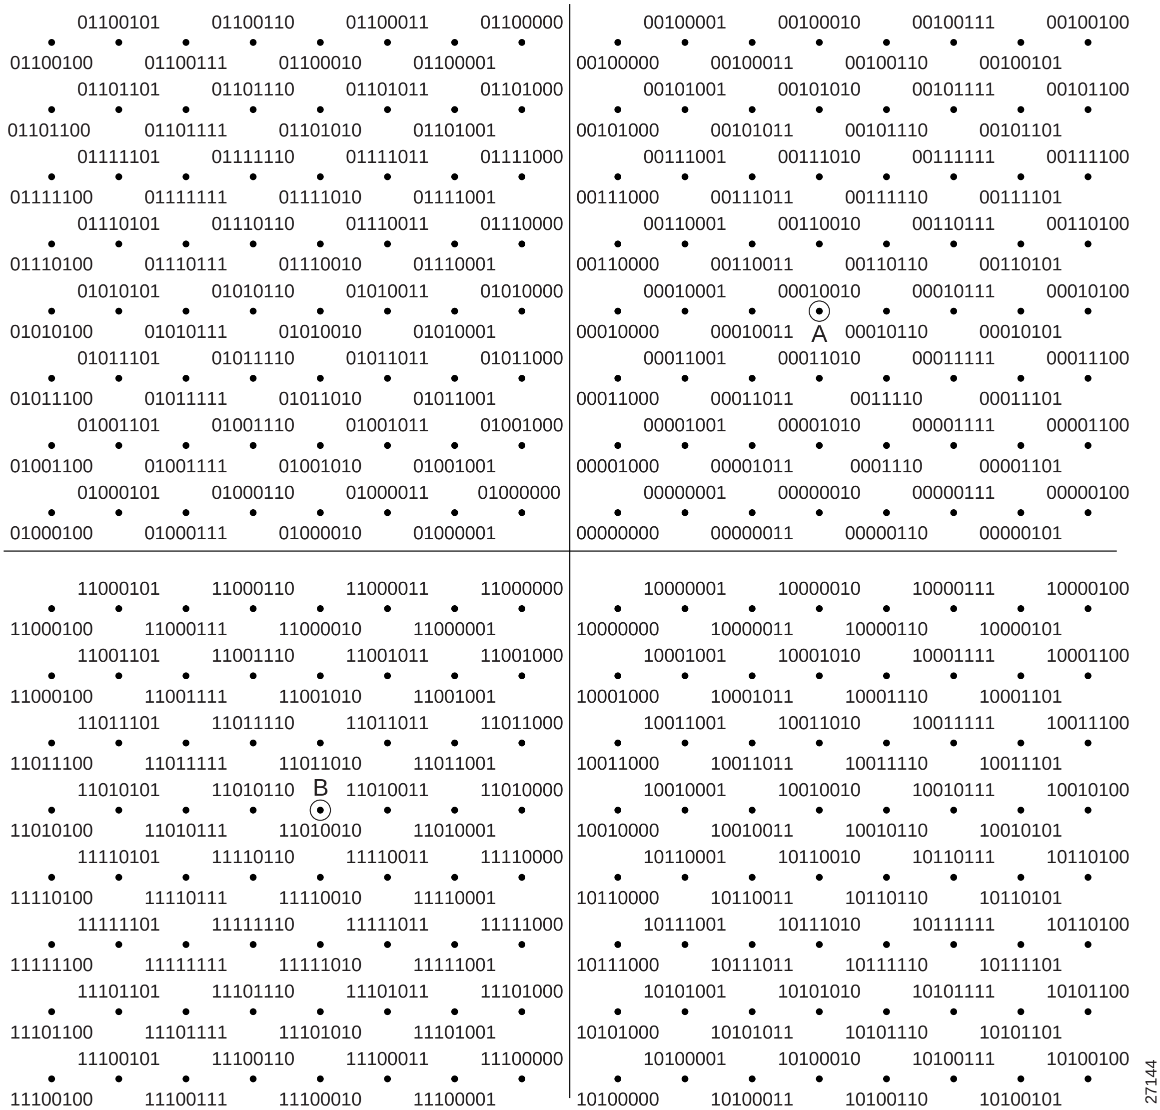

Encoder point designations are Z7 to Z0

The two-dimensional eight-state trellis encoder shown in Figure 13 is used during data mode, regardless of the number of points in the signal constellation. This trellis encoder uses the same convolutional encoder as the one defined in ITU Recommendation V.32. The trellis diagram is shown in Figure 14, and Figure 15 shows the 8-point, 32-point, and 128-point constellations during data mode. Figure 16 shows the 16-point, 64-point, and 256-point constellations during data mode.

Figure 13 **Two-Dimensional Eight-State Trellis Encoder**

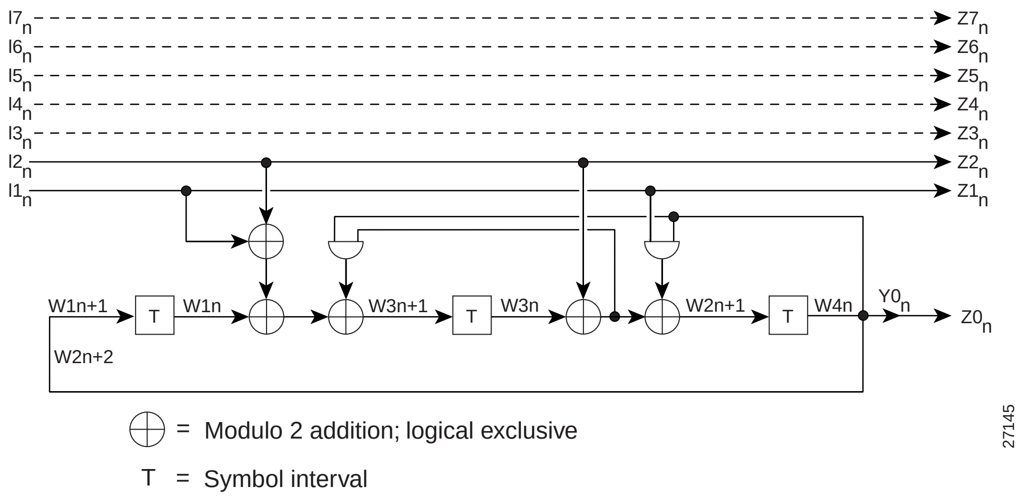

Figure 14 **Trellis Diagram of a Two-Dimensional Eight-State Encoder**

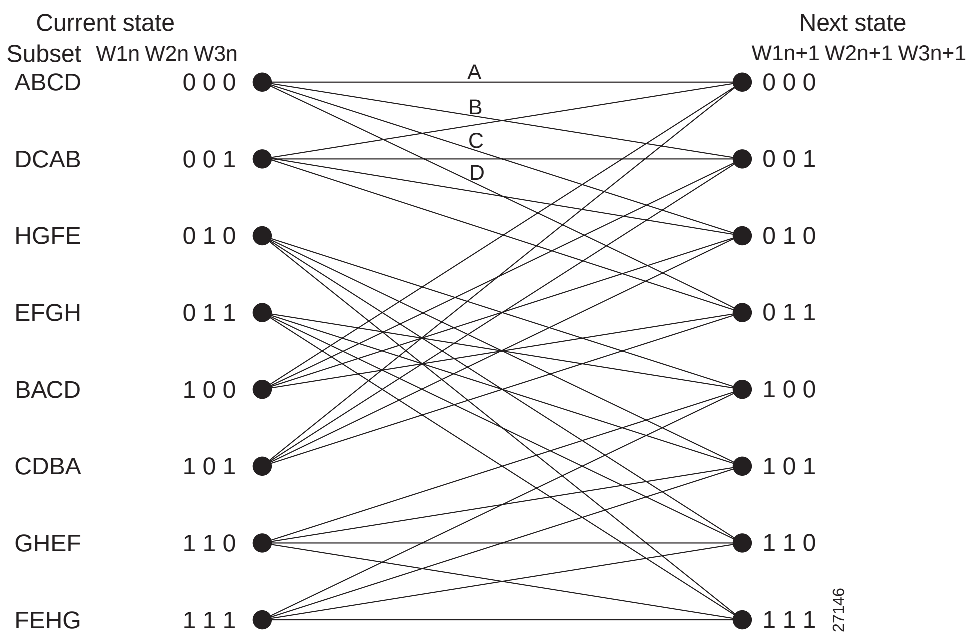

Figure 15 8-Point, 32-Point, and 128-Point Constellations During Data Mode

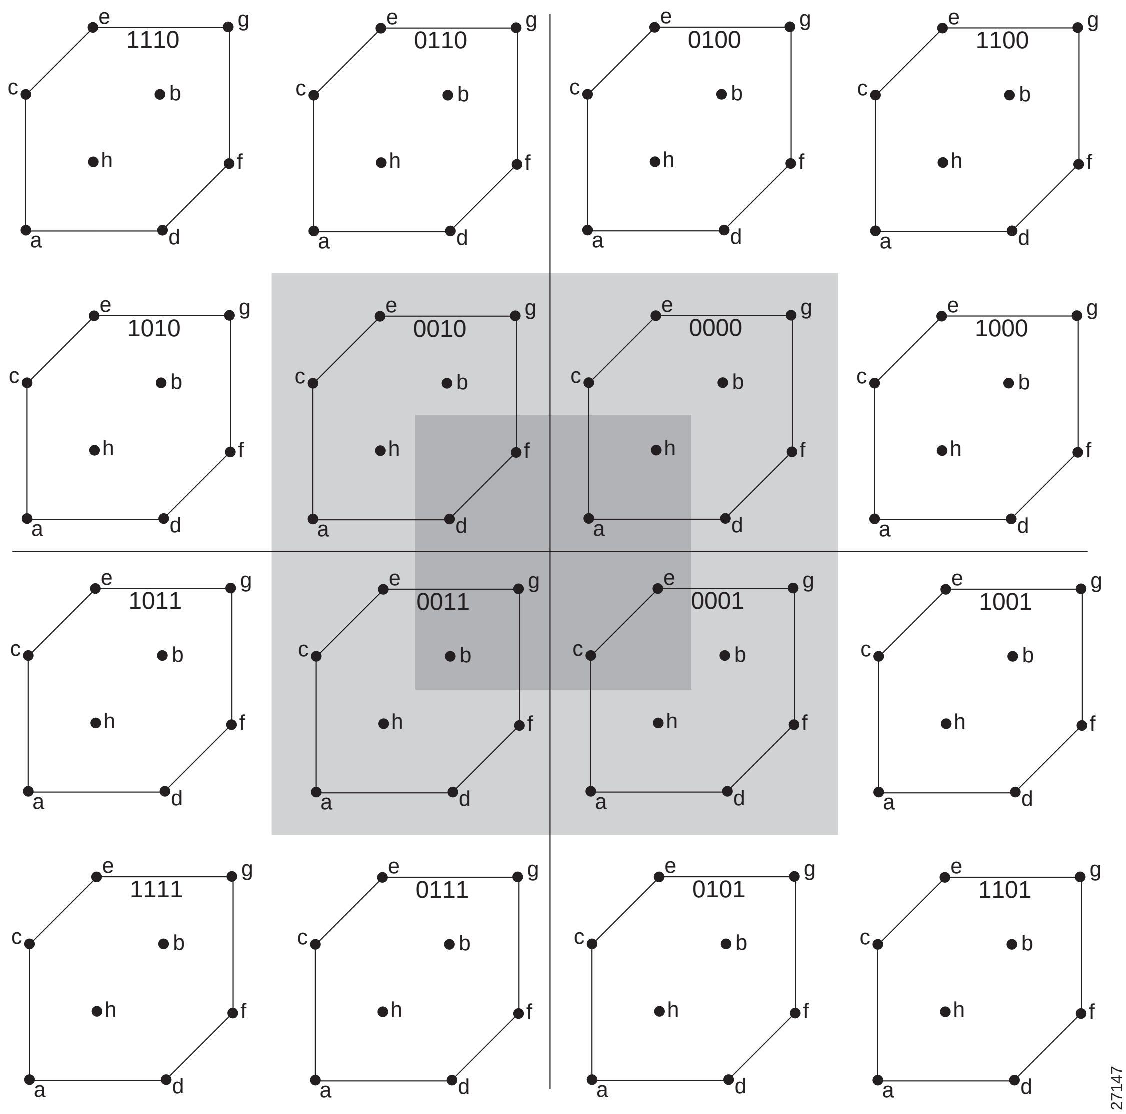

Subset bit mapping
SBS[2:0]
a—000
b—010
c—100
d—110
e—011
f—101
g—111
h—001

8 points SBS[2:0]
32 points—A[4:3], SBS[2:0]
128 points—A[6:3], SBS[2:0]
A[ ] defines the MSB values

Figure 16 **16-Point, 64-Point, and 256-Point Constellations During Data Mode**

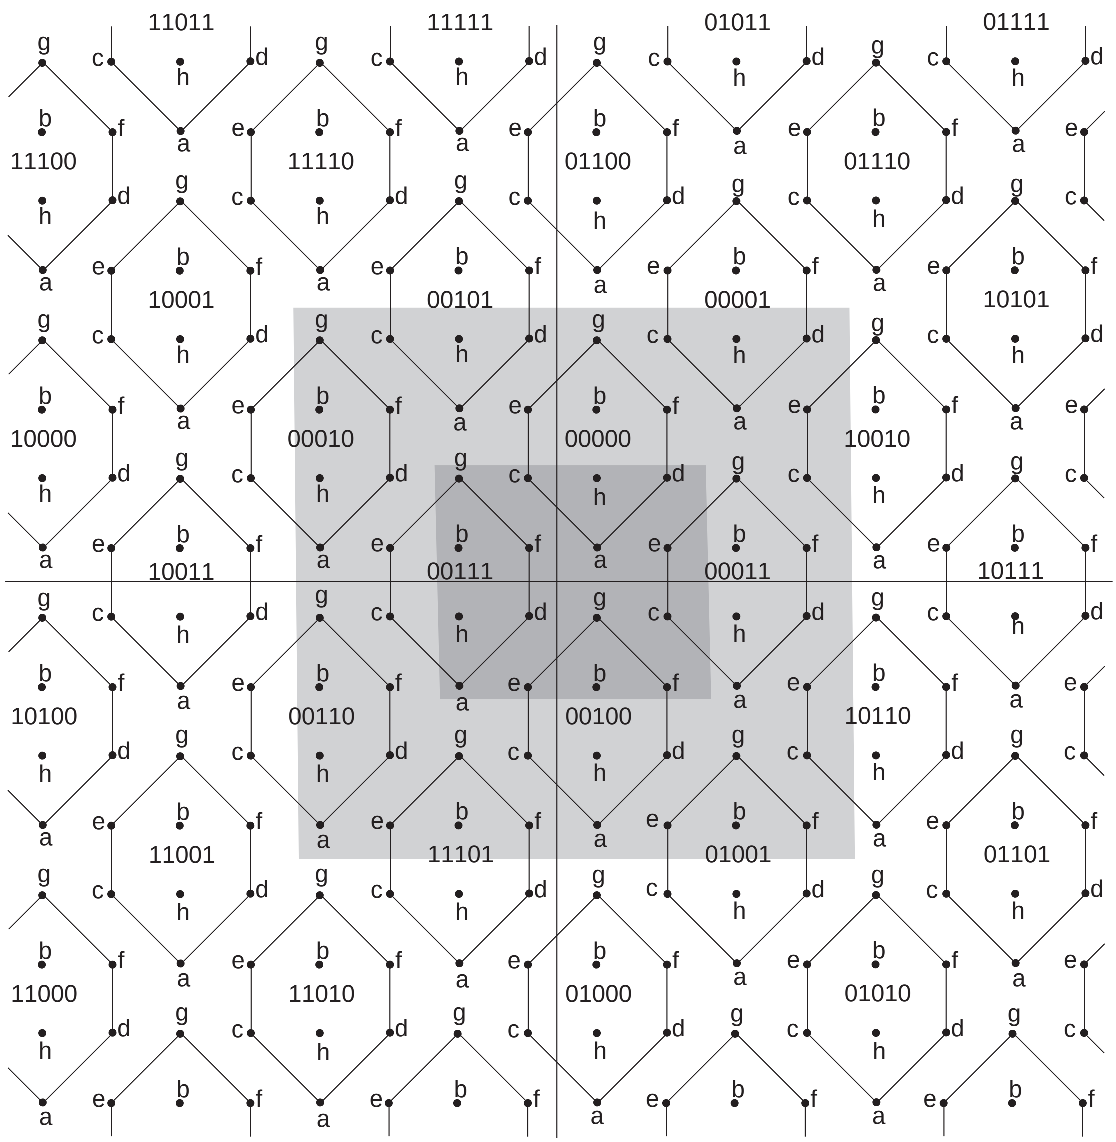

Subset bit mapping
SBS[2:0]
a—000
b—010
c—100
d—110
e—011
f—101
g—111
h—3001

16 points—A3, SBS[2:0]
64 points—A[5:3], SBS[2:0]
256 points—A[7:3], SBS[2:0]
A[ ] defines the MSB values

27148

5.6 CAP Start-Up Procedures

This section describes the various signals and bit sequences used during start-up. Various start-up modes are defined as follows:

- Initialization at installation (cold start)—accomplished at service installation where the modem assumes no prior channel knowledge, and performs an exhaustive search of the best configuration (highest bit rate for a specified margin) for the loop and noise environment.
- Session—transceiver starts up with the configuration of the previous session. The transceiver determines if the bit rate should be increased or decreased based on the signal quality received.
- Specific configuration—forces a specific configuration and does not search for an alternative.
- Warm start-up procedure—defines a faster method to reestablish a communication session between local and remote transceivers by eliminating the configuration negotiation phase and starting up with the same specific configuration as the previous session.

The procedures discussed in the following sections define a set of primitive start-up sequences, which can be selected by the system operator. Definitions of the preceding start-up mode(s) are constructed with the passing of appropriate data within the start-up procedures.

5.6.1 Start-up Procedure Definitions

This section covers the definitions used with start-up procedures.

Transparency

Prior to the completion of activation, transmission on the subscriber line is not transparent; the signals that are present at the line interface are specific to the start-up patterns generated by the RADSL transceiver. The transceiver provides transparent transmission of the payload data after termination of the activation procedure.

Signal Quality

The signal quality parameter is estimated at the receivers in the RTU-C and RTU-R. This value is used to estimate the bit error rate (BER) or signal-to-noise ratio margin of the received data. It takes into account the total signal-to-interference ratio, where the interference includes background noise, cross talk, residual intersymbol interference, residual echo from the neighboring upstream or downstream channel, and distortion.

T

Symbol interval of the transmit signal.

5.6.2 Transmitted Signal Definitions

The following definitions describe the transmitted signals during activation.

Silent

No signal is transmitted to the line during this state.

S0 Signal

The S0 signal is used as an alerting or wake-up sequence to initiate transceiver activation. The sequence uses the generator polynomial of $1 \oplus x^{-5} \oplus x^{-6}$ seeded with bits [000001]. The bit to symbol mapping for the S0 sequence is two points, namely ± A as defined in Table 9.

CS0 Signal

The S0 signal transmitted from the RTU-C is labeled CS0.

RS0 Signal

The S0 signal transmitted from the RTU-R is labeled RS0.

S0’ (CS0’ and RS0’) Signal

CS0’ and RS0’ are the same physical signals as their corresponding CS0 and RS0, however, their duration is at least $252 \cdot T(4 \cdot 63 \text{ symbol intervals})$ where T is the symbol interval of the corresponding upstream or downstream (T_C) signal. For the upstream channel, the symbol rate is 85 kilobaud, and for the downstream channel, the symbol rate is 136 kilobaud.

S1 (CS1 and RS1) Signal

The S1 signal is an unencoded 16-CAP signal sequence. In the downstream channel, this signal is called CS1 and in the upstream channel, this signal is called RS1.

CS1 Signal

CS1 uses the generator polynomial, $1 \oplus x^{-5} \oplus x^{-23}$, and the line-signal symbol rate is the one selected in the trailer-data exchange field. This sequence is only transmitted in the downstream direction.

RS1 Signal

RS1 uses the generator polynomial, $1 \oplus x^{-18} \oplus x^{-23}$, and the line-signal symbol rate is the one selected in the trailer-data exchange field. This sequence is only transmitted in the upstream direction.

16-CAP Signal

This signal can be any arbitrary unencoded 16-CAP signal sequence. This signal is also referred to as S5.

S2 (CS2 and RS2) Signal

The S2 signal is an unencoded 16-CAP signal sequence. CS2 and RS2 define the respective sequences in the downstream and upstream channels.

CS2 Signal

In the downstream direction, the RTU-R transmitter initiates the data mode scrambler, as shown in Figure 9. The scrambler-generator polynomial for CS2 is $1 \oplus x\mathrm{r} - 5 \oplus x^{-23}$, seeded with all zeros. The signal constellation is 16-CAP.

RS2 Signal

In the upstream direction, the RTU-R transmitter initiates the data mode scrambler, as shown in Figure 9. The scrambler-generator polynomial for RS2 is $1 \oplus x^{-18} \oplus x^{-23}$, seeded with all zeros. The signal constellation is 16-CAP.

16-CAP (S5) Signal + Precode

The 16-CAP (S5) signal as previously defined with the Tomlinson precoder enabled.

CS3 Signal

The CS3 signal is an encoded 16-CAP signal that includes the data mode scrambler, as shown in Figure 9 in the frame-locked condition, that is, with the all-zeros data sequence transmitted. The trellis encoder and Tomlinson precoder are enabled for the transmission of CS3. The corresponding symbol rate is the one chosen in the negotiation phase of the alerting sequence.

RS3 Signal

In the upstream direction, the RS3 signal is an encoded *N*-CAP signal that includes the data mode scrambler, as shown in Figure 9 in the frame-locked condition, that is, with the all-zeros data sequence transmitted. The trellis encoder and Tomlinson precoder are enabled for the transmission of CS3. The corresponding symbol rate and constellation size are those chosen in the negotiation phase of the alerting sequence.

CS4 Signal

CS4 is an encoded *N*-CAP signal transmitted in the downstream direction that includes the data mode scrambler, as shown in Figure 9 in the frame-locked condition, that is, with the all-zeros data sequence transmitted. The trellis encoder and Tomlinson precoder are enabled for the transmission of CS4. The corresponding symbol rate and constellation size are those chosen in the negotiation phase of the alerting sequence.

S6 (CS6 and RS6) Signal

The S6 signal is the unencoded 16-CAP signal with Tomlinson precoding enabled. In the downstream direction, the RTU-C uses the generator polynomial $1 \oplus x^{-5} \oplus x^{-23}$; the corresponding transmit signal is referred to as CS6. In the upstream direction, the RTU-R uses the generator polynomial $1 \oplus x^{-18} \oplus x^{-23}$; the corresponding transmit signal is referred to as RS6.

S7 (CS7 and RS7) Signal

The S7 signal is the unencoded 16-CAP signal with Tomlinson precoding enabled and the scrambler reset. In the downstream direction, the RTU-C transceiver uses the generator polynomial $1 \oplus x^{-5} \oplus x^{-23}$; the corresponding transmit signal is referred to as CS7. In the upstream direction, the RTU-R uses the generator polynomial $1 \oplus x^{-18} \oplus x^{-23}$; the corresponding transmit signal is referred to as RS7.

5.6.3 Timers

This section defines the timers that are used in the activation sequence. It also describes the values for those timers, which correspond to the activation sequences.

T1 ($T1_W$)

Time-out or maximum response time for an RTU-C replying to an RS0 signal from the RTU-R.

T2 ($T2_W$)

The time it takes for an RTU-C transmitter to send a CS1 signal upon completion of the CS0 signal plus the trailer sequence. The RTU-R can use this timer to determine the receive time of the CS1 sequence.

T3 ($T3_W$)

This defines the time it takes for the RTU-R receiver to detect a CS1 signal and respond with the RS1 sequence.

T4 ($T4_W$)

The RTU-R timer that is used to determine the reset time for the data mode scrambler in the RTU-R.

T5

Start time for transmission of the Tomlinson coefficients in the RTU-R.

T6 ($T6_W$)

This timer starts the RS encoder frame. This is an exact number used in the activation sequence in that it has no tolerance specified.

T7 ($T7_W$)

Transmission time of 16-CAP + precode followed by the CS4 signal in the RTU-C, and 16-CAP + precode followed by the RS3 signal in the RTU-R during the RS encoder locking phase.

Timer Values

Table 9 defines the timer values for the activation sequences. Note that some timer values are expressed in terms of symbol intervals, where T_{340} is the symbol interval for the 340 kilobaud symbol clock. The symbol f_B is the baud rate selected during the negotiation phase and is inversely related to T_B, which is the symbol interval selected during the negotiation phase. Table 10 defines the equivalent timer values for the warm start-up activation sequence.

Table 9 Timer Values for the Activation Sequence

Timer	Lower Bound	Upper Bound
T1	$64 \cdot T_{340}$	$4000 \cdot T_{340}$
T2	100 ms	180 ms (<340 kilobaud) 600 ms (>340 kilobaud)
T3	100 ms	180 ms (<340 kilobaud) 600 ms (>340 kilobaud)
T4	$(210 \text{ kilobaud} \pm 10,000) \cdot T_{340}$	
T5	$(150 \text{ kilobaud} \pm 1,000) \cdot T_{340}$	
T6	$200,000 \cdot T_B$ for $f_B \le 340$ kilobaud $400,000 \cdot T_B$ for 340 kilobaud $< f_B \le 680$ kilobaud $600,000 \cdot T_B$ to $f_B > 680$ kilobaud	
T7	$(180,000 \pm 10,000) \cdot T$ for $f_B \le 340$ kilobaud $(360,000 \pm 10,000) \cdot T$ for 340 kilobaud $< f_B \le 680$ kilobaud $(540,000 \pm 10,000) \cdot T$ to $f_B > 680$ kilobaud	

Table 10 Timer Values for the Warm Start-up Activation Sequence

Timer	Lower Bound	Upper Bound
T1	$64 \cdot T_{340}$	$4,000 \cdot T_{340}$
T2	100 ms	180 ms($\le$ 340 kilobaud) 600 ms (>340 kilobaud)
T3	100 ms	180 ms($\le$ 340 kilobaud) 600 ms (>340 kilobaud)
T4	$(210,000 \pm 10,000) \cdot T_{340}$	
T5	350 ± 5 ms	

Table 10 Timer Values for the Warm Start-up Activation Sequence

Timer	Lower Bound	Upper Bound
T6	$200,000 \cdot T$ for $f_B < 340$ kilobaud $400,000 \cdot T$ for $340 < f_B \leq 680$ kilobaud $600,000 \cdot T$ to $f_B > 680$ kilobaud	
T7	$(180,000 \pm 10,000) \cdot T$ for $f_B \leq 340$ kilobaud $(360,000 \pm 10,000) \cdot T$ for $340 < f_B \leq 680$ kilobaud $(540,000 \pm 10,000) \cdot T$ to $f_B > 680$ kilobaud	

5.6.4 Customer-Initiated Activation Sequence

This section covers the customer-initiated activation sequence as shown in Figure 17 and the alerting or wake-up stage. It also covers trailer data bit information.

Alerting is initiated by the RTU-R that is transmitting the RS0 sequence followed by a trailer block that contains information for system configuration and rate negotiation. The structure of the alerting phase is shown in Figure 18. In the first event of RS0, the trailer data is configured as a *poll* packet. Once it detects the alerting RS0 sequence, the system responds with the CS0 sequence. The trailer data to the CS0 sequence is configured as a *reply* packet, responding to the *poll* from the RTU-R.

Figure 17 Activation Sequence with RTU-R Initiated Calling

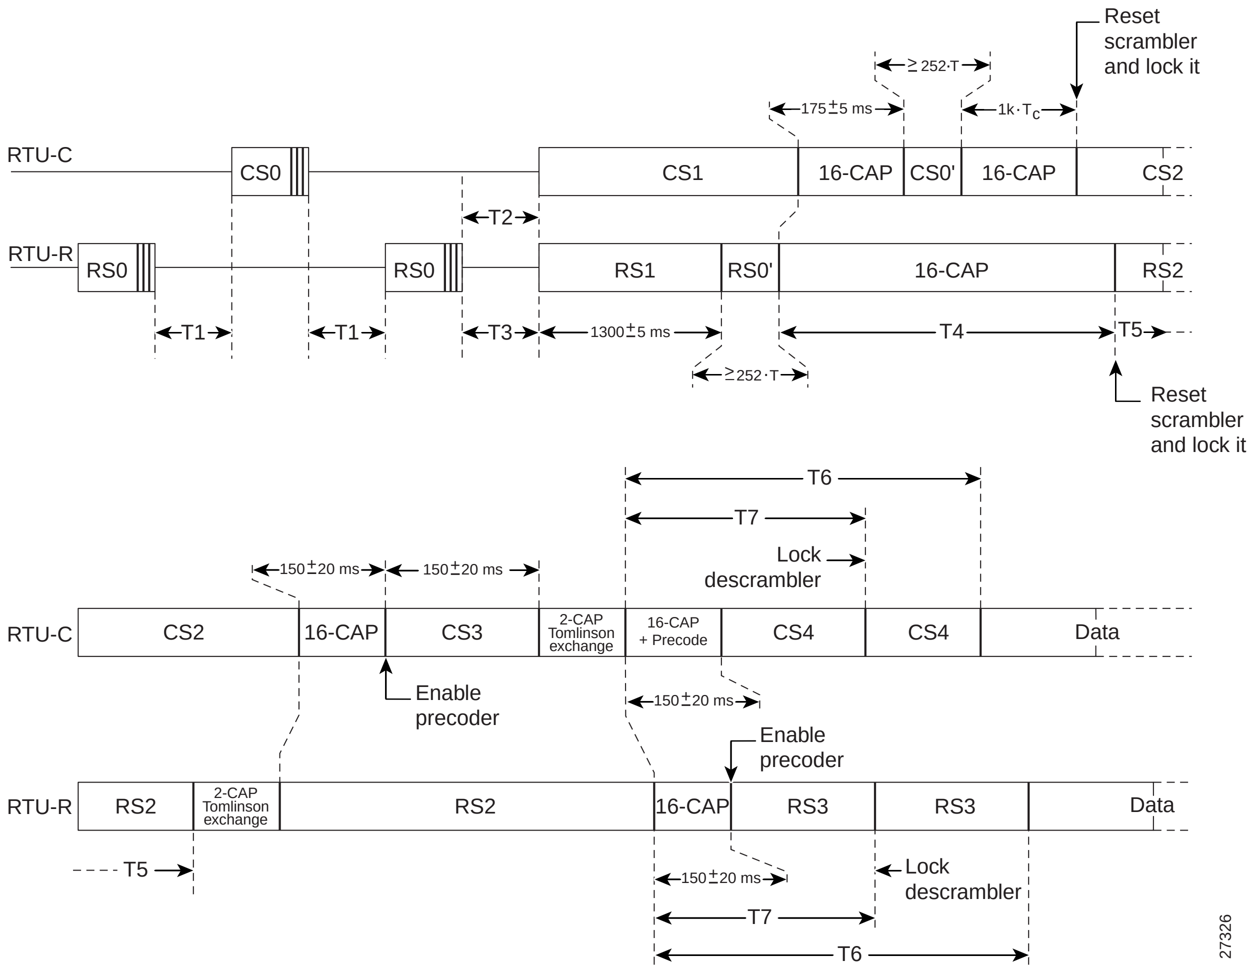

The symbol rates for the upstream and downstream channels in this phase are 136 kilobaud for the upstream channel and 340 kilobaud for the downstream channel. A logic *one* in a trailer data bit is encoded by sending one complete cycle of the $2^6 - 1$ pseudorandom sequence at the corresponding channel baud rate using the constellation in Figure 10.

A logic *zero* bit is encoded by sending the 0 symbol defined in Figure 10.

The duration of S0 (CS0 or RS0) is $63 \cdot 60 \cdot T$, where T is the respective downstream or upstream symbol interval.

The duration of the trailer data is $(n + 5) \cdot 63 \cdot T$, where $n + 5$ is the number of trailer data bits including the start bit (ST), parity bit (P), two stop bits (SP0 and SP1), and T is the respective downstream or upstream symbol interval. The value of n for this version of the specification is 63.

Figure 18 Alerting Phase of Activation Sequence

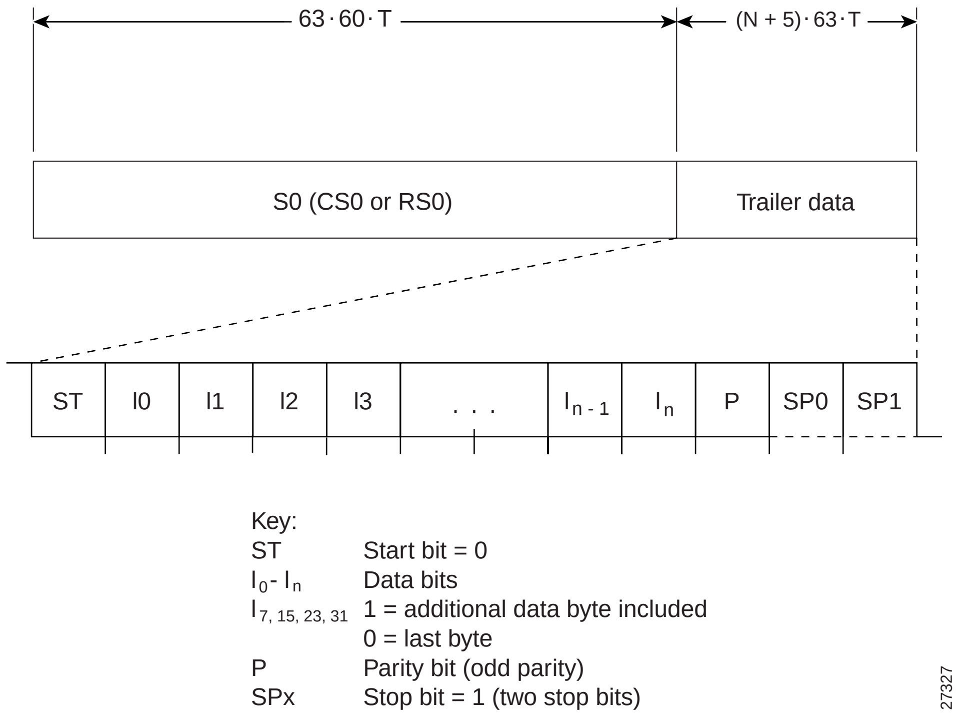

The first byte in the alerting phase trailer consists of bits I0 to I7. Bits I1,I0 are defined as follows:

- 00 = data mode
- 01 = run 511 bit error rate tester (BERT) for 10 seconds
- 10 = run BERT for 30 seconds, 11 = run BERT for 2 minutes.

Bits I2 and I3 define the transmit signal power reduction. Bits I2,I3 range from 00 to 11, where 00 identifies no power reduction (0 dB), 01 = 6 dB reduction, 10 = 12 dB reduction, and 11 = 18 dB reduction.

Bit I4 defines the trailer packet as either a poll (0) or a reply (1).

Bit I5 defines the start-up request as a normal start-up (0) or a warm start-up (1).

Bit I6 defines the start-up capability as a normal start-up (0) or a warm start-up (1).

Bit I7 is a continuation bit that, if set to 1, is followed by a byte.

The second byte consists of bits I8 to I15. The lower four bits (I8 to I11) define the specification version supported. Bit I12 is defined as the variable baud clock mode request, where 1 = variable baud mode and 0 = predefined baud mode. Bit I13 is defined as the variable baud clock mode capability where 1 = variable baud rate and 0 = predefined baud mode. Bit I14 defines auto selection (0) or predefined (1) constellation mode. Bit I15 is the continuation bit.

The third byte consists of bits I16 to I23. Bits I16 to I19 define the symbol clock selection, which is summarized in Table 11.

Table 11 Symbol Clock Selection

I19, I18, I17, I16	Downstream/Upstream Baud Clocks
0000	340 kilobaud/136 kilobaud
0001	680 kilobaud/136 kilobaud
0010	952 kilobaud/136 kilobaud
0011	1088 kilobaud/136 kilobaud

Bits I20 to I22 are reserved for future growth and set to zero. Bit I23 is a continuation bit.

The fourth byte contains the symbol rate capability, as shown in Table 12, and consists of bits I24 to I31, where bit I31 is reserved as the continuation bit. In the fourth byte, a 1 identifies that the symbol rate capability exists.

Table 12 Symbol Rate Capability Selection

Bit	Downstream/Upstream Symbol Rate Capability
I24	340 kilobaud/136 kilobaud
I25	680 kilobaud/136 kilobaud
I26	952 kilobaud/136 kilobaud
I27	1088 kilobaud/136 kilobaud
I28	136 kilobaud/136 kilobaud
I31	continuation bit

The fifth byte defines the first seven bits of the baud clock variable N0. Bits I32 to I38 = N0[6:0], and I39 = continuation bit.

Table 13 CO Upstream Request—Fifth Byte

Kilobaud	I38	I37	I36
136	0	0	0
68	0	0	1
34	0	1	0
17	1	0	0

Table 14 CPE Upstream Request—Fifth Byte

Kilobaud	I34	I33	I32
136	0	0	0
68	0	0	1
34	0	1	0
17	1	0	0

The sixth byte defines the eighth bit of N0 and the baud clock variable D0. Bits I40 to I45 = D0[5:0], bit I46 = N0[7], and bit I47 = continuation bit.

Table 15 CO and CPE Upstream Capabilities—Sixth Byte

Downstream/ Upstream	I46	I45	I44	I43	I42	I41	I40
136/136	0	0	0	0	0	0	0
340/136	0	0	0	0	0	0	0
340/68	0	0	0	0	0	0	1
340/34	0	0	0	0	0	1	0
340/17	0	0	0	0	1	0	0
136/68	0	0	0	1	0	0	0
136/34	0	0	1	0	0	0	0
136/17	0	1	0	0	0	0	0

The seventh byte defines the first seven bits of the baud clock variable D1. Bits I48 to I54 = D1[6:0], and I55 = continuation bit.

The eighth byte defines the eighth bit of D1 and the baud clock variable D2. Bits I56 to I60 = D2[4:0], bit I61 = reserved, bit I62 = D1[7], and bit I63 = continuation bit.

Transceiver Training

Transceiver training is defined in the activation sequence of Figure 17.

Tomlinson Precoder Coefficient Exchange

This section covers the Tomlinson coefficients and the data found in the data block field.

Transfer of the Tomlinson coefficients is done by switching to the 2-CAP signal constellation, which is defined using points A and B in the unencoded 256-point constellation in Figure 12. The data frame structure that is used to transmit the coefficients is shown in Figure 19.

The size of the data block is one word or 16 bits. As shown in Figure 19, the first 8 bits (1 byte) in the data block field define the size of the data block in terms of the number of bytes. The number in this field is a binary number that identifies the number of bytes following the size byte. The most significant bit is transmitted first.

Figure 19 Tomlinson Coefficient Exchange Frame Structure

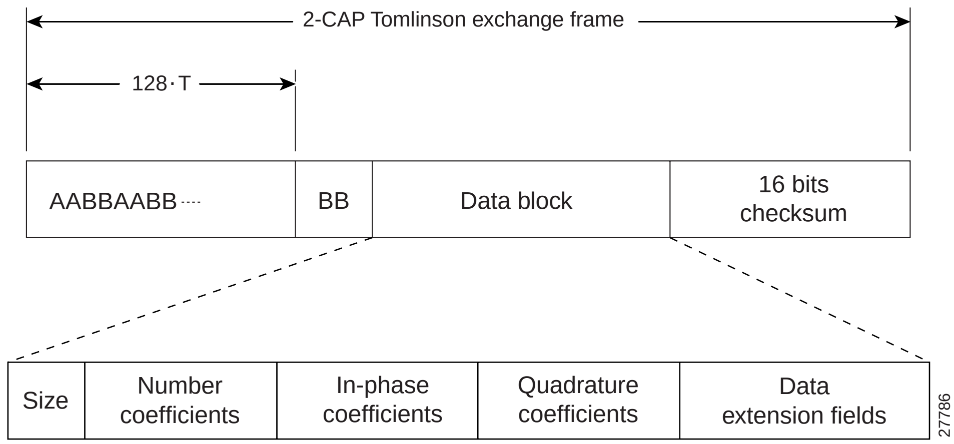

The number of coefficients in the Tomlinson precoder is one word. The lower 8 bits of this 16-bit word define the number of complex coefficients in the Tomlinson precoder. The minimum number of taps is 16 for the downstream channel and 3 for the upstream channel.

The in-phase coefficients are equal to N words. The coefficients are transmitted in the sequence $c_0, c_1, \ldots, c_{N-1}$, where N is the number of coefficients. This field contains the in-phase coefficients to the precoder feedback filter.

The quadrature coefficients are also equal to N words. This field contains the quadrature phase coefficients to the precoder feedback filter.

The requested CAP constellation is two words, which are reserved for the constellation size request; one for the downstream channel, as shown in Table 16, and the other for the upstream channel, as shown in Table 17.

Table 16 Constellation Request for Downstream Channel

Bit Location	Description
B0=0	reserved
B1	8-CAP, 2D8S trellis encoder enabled
B2	16-CAP, 2D8S trellis encoder enabled
B3	32-CAP, 2D8S trellis encoder enabled
B4	64-CAP, 2D8S trellis encoder enabled
B5	128-CAP, 2D8S trellis encoder enabled
B6	256-CAP, 2D8S trellis encoder enabled
B7	256-CAP, 2D8S trellis encoder disabled
B8 to B15	reserved; bits are set to 0

The two-word channel constellation capability identifies which constellations the transceiver can support in the upstream and downstream channels. Each bit position corresponds to a constellation configuration. A *one* in the bit field indicates that the system supports the associated constellation; a *zero* indicates that the system does not support the associated constellation. Note that bits 0 and 8 to 15 are reserved and are all set to the default zero value.

Table 17 Upstream Channel Constellation Capability

Bit Location	Description
B0=0	reserved
B1	8-CAP, 2D8S trellis encoder enabled
B2	16-CAP, 2D8S trellis encoder enabled
B3	32-CAP, 2D8S trellis encoder enabled
B4	64-CAP, 2D8S trellis encoder enabled
B5	128-CAP, 2D8S trellis encoder enabled
B6	256-CAP, 2D8S trellis encoder enabled
B7	256-CAP, 2D8S trellis encoder disabled
B8 to B15	reserved; bits are set to 0

Transmit Power is a 16-bit word that describes the transceiver transmit power setting in decibels per milliwatt.

Receiver Gain1 is a 16-bit word that describes the associated receiver input gain setting in decibels.

Received Signal Quality is a 16-bit word that describes the associated receiver output signal-to-noise ratio in decibels during start-up. The signal quality is equal to $49.34 - 10\log(x)$ in decibels. If x = 0 then the signal quality is >49.34 dB.

There is one reserved word. This word is used for future growth and contains zeros.

The RS encoder configuration request and capability are two words. The RS encoder is a two-symbol error-correcting encoder. The default size for both the downstream and upstream is $(N,K) = (68,64)$.

A *short* interleave depth is defined as size $D = 4$. The first word contains the RS encoder configuration request information and the second word contains the RS encoder configuration capability. Table 18 defines the bit assignments for the RS encoder request word and capability word.

Table 18 RS Encoder Configuration Request (Capability)

Bit No. (0 to 15)	Description
0	downstream short interleave requested (supported)
1	downstream long interleave depth requested (supported)
2	downstream variable interleave depth requested (supported)
3	downstream default encoder word size requested (supported)
4	downstream variable encoder word size requested (supported)
5	upstream short interleave requested (supported)
6	upstream long interleave depth requested (supported)
7	upstream variable interleave depth requested (supported)
8	upstream default encoder word size requested (supported)
9	upstream variable encoder word size requested (supported)
10 to 15	reserved

The downstream RS encoder interleaving depth size is one 16-bit word. The upstream RS encoder variable encoder-word size also is one 16-bit word.

The Vender Identification is a 16-bit word that is set to 0034 hex. The four vender-specific 16-bit words are used for determining compatibility and configuring feature sets. Table 19 defines the bits in the vender-specific message that initiate from the RTU-R, and Table 22 defines the bits in the vender-specific message that initiate from the RTU-C.

Table 19 Vender-Specific Encoders for the Upstream

Word	Bits	Description
Word 0	15	odd parity
Word 0	14:8	RTU-R feature set
Word 0	7	reserved for future use; set to 0
Word 0	6:0	RTU-R hardware version
Word 1	15:0	reserved for future use; set to 0000 hex
Word 2	15:0	reserved for future use; set to 0000 hex
Word 3	15:0	reserved for future use; set to 0000 hex

The odd parity bit in word 0 is used for error detection in control word 0. It is valid only for word 0.

The RTU-R feature set is an incremental count for each significant Cisco software feature release. For all versions less than and including release 2.0, the value is set to zero. These values are defined in Table 20.

Table 20 RTU-R Feature Set

RTU-R Feature Set Value	Description
0	baseline unit; functionality controlled at RTU-R
1	support for session and idle timers set at Cisco 6100

RTU-R hardware version is an incremental count for each new Cisco hardware product. These values are defined in Table 21.

Table 21 RTU-R Hardware Version

RTU-R Hardware Version Value	Description
0	Cisco 675
1	Cisco 675
2	Cisco 605

Table 22 Vender-Specific Encoders for the Downstream

Word	Bits	Description
Word 0	15:0	upstream rate in kbps for ScalaRate
Word 1	15	odd parity
Word 1	14:13	reserved for future use; set to 0
Word 1	12:10	timer type

Table 22 Vender-Specific Encoders for the Downstream

Word	Bits	Description
Word 1	9:0	timer value
Word 2	15:0	reserved for future use; set to 0000 hex
Word 3	15:0	reserved for future use; set to 0000 hex

Word 0 for the downstream vender-specific encoders defines the upstream ScalaRate in 1 kbps increments. The value is in hexadecimal notation.

The odd parity bit in word 1 is used for error detection in control word 1. It is valid only for word 1.

The timer type coupled with the timer value in Word 1 is used by the Cisco 6100 Series system to establish timeouts on the RTU-R. There are two distinct timers that can be set by the Cisco 6100 Series system into the RTU-R. These are session and idle timers. The values for timer type are defined in Table 23.

Table 23 Timer Values for Downstream Vender-Specific Encoders

RTU-R Timer Type	Description
0	no timer used
1	idle timer active
2	session timer active

A session timer is the total amount of time a link stays active once it is trained regardless of traffic. This timer starts counting after modem initialization and continues until the link is dropped by you or until it counts down to zero, at which time the RTU-R drops the ADSL physical layer link. This timer continuously counts and is only reset when the ADSL link goes active.

The idle timer is used to drop the ADSL physical layer link with the lack of ethernet packet activity on your ethernet port. This timer starts counting after modem initialization and continues until it counts down to zero, at which time the RTU-R drops the ADSL link. The timer continuously counts and is reset upon detection of incoming traffic on your data port.

In word 1, the timer value is the number of minutes that the active timer type uses for the timeout value.

Receiver Gain2 is a 16-bit word that describes the associated receiver input-gain setting in linear scale, where 0x800 = 1.

The downstream and upstream RS Size is 1 16-bit word. The upper byte contains the upstream encoder-word size and the lower byte contains the downstream encoder-word size.

A 16-bit word defines the service class, framer presence, and configuration. Bit 0 defines the service class configuration of 0 = class 1 and 1 = class 2. Bit 1 defines the framer presence of 0 = framer not present and 1 = framer present. Bit 2 identifies the class 1 payload type of 0 = ATM-based and 1 = packet-based payload. Bits 3 to 15 are reserved for future growth and are set to 0.

The checksum field is 1 word and contains the "twos" complement of a zero-run checksum.

RS Time Stamp

Transfer of the RS time stamp is done by switching to the 2-CAP signal constellation structure used for the Tomlinson exchange; that is, the transmission of points A and B as defined in the unencoded 256-point constellation of Figure 12.

The structure of the RS time stamp sequence is shown in Figure 20. The data frame starts with the transmission of $256 = 128 \cdot 2\ \mathrm{symbols}$ symbols (2 points) with the repetitive sequence of points AABBAABB---AABB for synchronization and identification of the start of the coefficient data frame. This is followed with the RS marker symbol sequence, BB, which identifies the time stamp.

Figure 20 RS Time Stamp Structure

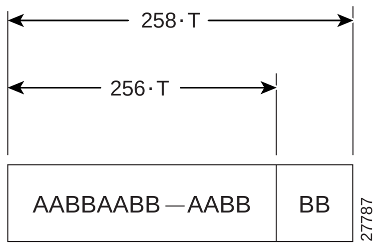

After completion of the time stamp sequence, the RTU-C starts the T6 timer. The RTU-R receiver starts its T6 timer when it finishes receiving the time stamp.

Lock Descrambler

Immediately after sending the RS time stamp, the RTU-C transmits 150 ± 20 milliseconds of 16-CAP + Precode. The RTU-C then sends the CS4 sequence for the duration of the T7 timer and locks the descrambler.

The RTU-C sends the 16-CAP + precode signal for 8 150 ± 20 milliseconds immediately after receiving the RS time stamp. The RTU-R then sends the RS3 signal for the duration of the T7 timer and locks the descrambler.

5.6.5 Central-Office Initiated Activation Sequence

Initiation from the central office location is provided by the RTU-C initiating a poll packet, that is, the RS0 sequence is followed by a trailer with a poll identifier bit that is set accordingly. This poll initiates the activation procedure that is defined in the "Customer-Initiated Activation Sequence" section on page 25.

Warm Start-Up

The activation sequence for a customer-premises initiated warm start-up is shown in Figure 21.

Figure 21 Warm Start-Up Activation Sequence with RTU-R Initiated Calling

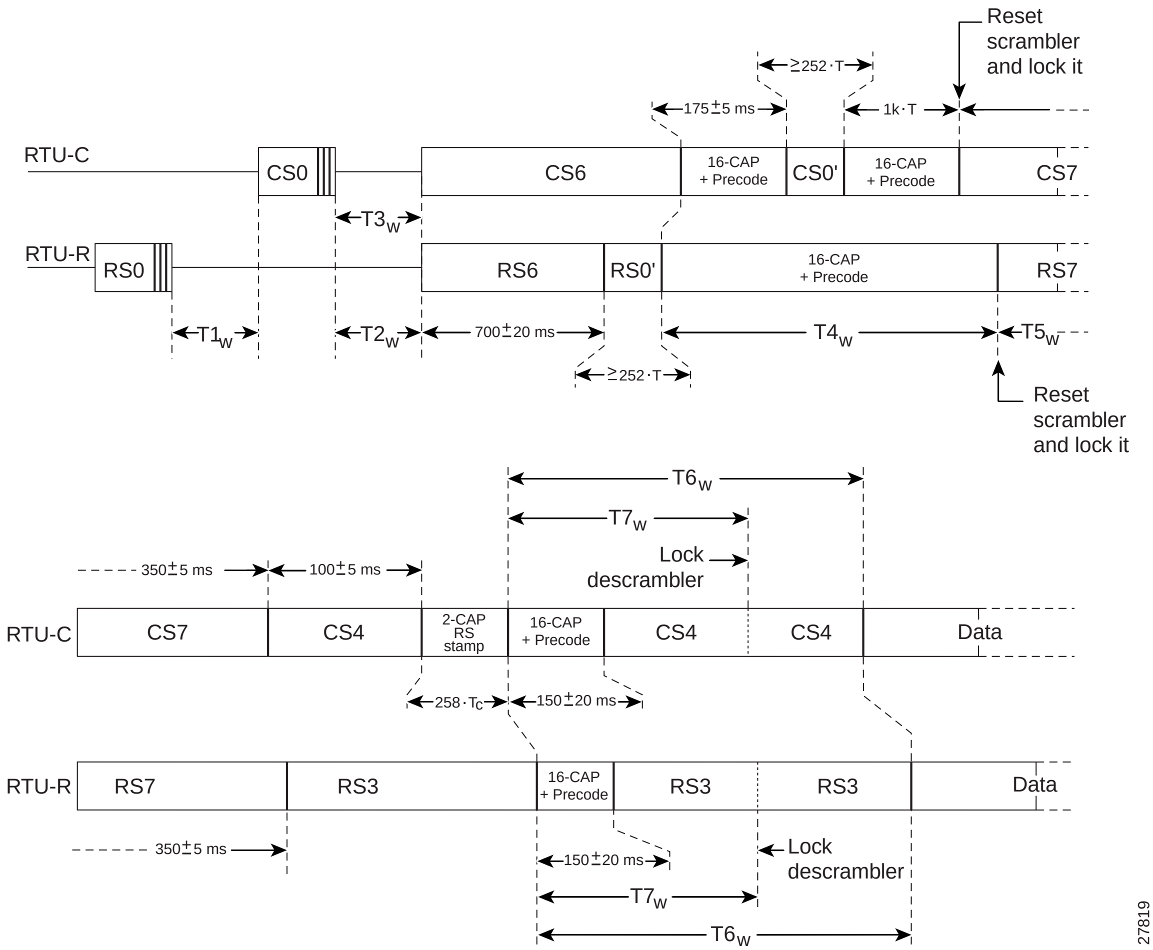

Warm start-up initiation from the central office is provided by the RTU-C initiating a poll packet. Once initiated, the warm start-up sequence is the same as previously defined.

Digital Off-Hook Signaling

In a system configured for Digital Off-Hook (DOH) signaling, in which subscriber lines are oversubscribed to RTU-C modems, in-band signaling in the form of sine tones is used to convey both busy (all modems allocated) or alert (the network-initiated activation of a subscriber modem that is on hook, or off line).

DOH Busy Signal

The DOH busy signal is a 300 kHz ±15 kHz tone transmitted from the RTU-C to the RTU-R at a level of 16.5 dBm into a nominal impedance of 100 ohms for a duration of 500 milliseconds (–200 milliseconds, +500 milliseconds). The tone is initiated at the RTU-C interface during and instead of the normal training initialization when the RTU-R begins training and no modem is available in the network.

DOH Alert Signal

The DOH alert signal is a 300 kHz ±15 kHz tone transmitted from the RTU-C to the RTU-R at a level of 16.5 dBm into a nominal impedance of 100 ohms for a duration of 500 milliseconds (–200 milliseconds, +500 milliseconds). The tone is initiated by the RTU-C if network traffic is available and the RTU-R is on hook or idled.

6. Data-Link Layer

This section discusses the second layer of the OSI model.

6.1 ATM Mapping

In a transmission frame adaptation, ATM cells are carried over the U interface as a continuous stream of cells in a cell-based format. Cell structure is the same as the one recommended in ITU-I.361. However, with a contiguous transmission, cells are transmitted contiguously with no special or intervening symbols.

Cell header delineation provides identification of the cell boundaries at the receiver. Full cell header delineation is required in the TC layers at both RTU-C and RTU-R receivers. It should be implemented as specified in ITU-I.432.4.5, with alpha = 7 and delta = 8, as recommended by ITU-I.432 for cell based streams. (There is also an option to leave alpha and delta parameters variable, and up to the implementor because their modification does not affect interoperability.)

The downstream path (U-C to U-R) should use the self-synchronizing scrambler (SSS) as specified in ITU-I.432; the upstream path (U-R to U-C) should also use the SSS. Because the SSS is used in both paths, no modifications are needed to the transmit header error-check sequence. Header verification at the receiver is not required to perform single-bit error corrections. Header error control (HEC) generation, where necessary, is as described in ITU-I.432 section 4.3.2 including the recommended modulo 2 addition of the pattern 0101010101b to the HEC bits. The generator polynomial coefficient set that is used and the HEC sequence generation procedure are in accordance with ITU-I.432.

Idle cells are inserted by each DSL transmitter for the purposes of cell rate decoupling, and are simply discarded at the receiver. The cell header fields for idle cells are encoded as specified in ITU-I.432.

6.2 Control Plane

Permanent virtual circuits (PVCs) and permanent virtual paths (PVPs) are used to provide up to 4 virtual circuit connections (VCCs) and up to 1 virtual path connection (VPC) as specified in ITU-I.361 and the ATM Forum User-Network Interface (UNI) Specification, version 3.1, using fixed VPI/VCI addresses. Allowable addresses and allocated use are shown in Table 24.

Table 24 ATM Address Assignments

Connection Type	VPI/VCI address	Usage
VCC	1 / 0	data
VCC	1 / 1	data
VCC	1 / 2	data
VCC	1 / 3	data
VPC	1 / *	data

Table 24 ATM Address Assignments

Connection Type	VPI/VCI address	Usage
VCC	1 / 4	in-band management

6.3 Management Plane

A mechanism is needed to transport the management information between the DSLAM (6100) and the CPE. This can be done by creating an in-band management channel between CPE and 6100, which helps in configuring the CPE, software download for CPE, subscriber notifications, and CPE performance monitoring. This section describes the complete design for the CPE in-band management channel, but the implementation in release 2.3 is done only to support the Display and Terminate commands.

Following is a list of design goals for the CPE in-band management channel:

- CPE and DSLAM communicate on a well-known virtual circuit (VC; VPI/VCI = 1/4).
- VCs are created and deleted dynamically based on subscriber connectivity—a maximum of 64 VCs (1 per modem) can exist at any time.
- Either end can initiate the messages—CPE or DSLAM.
- Messaging protocol caters to forward and backward compatibility.
- Any application residing in DSLAM, which can communicate with a system controller, can send or receive messages to or from the CPE.
- Messages sent as ATM adaptation layer 5 (AAL5) payload over ADSL.
- Application able to send a broadcast message to all active CPEs.
- Application able to send a message with a request for acknowledgment.
- Application able to send a message with multiple AAL5 PDUs.

Following is a list of design considerations for the CPE in-band management channel:

- CPE in-band VCs are not network management system (NMS) controlled VCs—not created or deleted by the management.
- VCs are created or deleted dynamically (at the DSLAM), based on state of subscriber line—no need for storing data in persistence database.
- CPE has a well-known VC opened permanently for the in-band channel—in addition to user data VCs.
- Messaging between CPE and DSLAM—predefined protocol.
- Preexisting software segmentation and reassembly (SAR) used for AAL5 SAR at the DSLAM end.
- Hardware SAR already being used in CPE.
- Terminal controller at DSLAM and CPE—responsible for creating and deleting VCs and transporting incoming and outgoing messages to the respective applications.
- Applications at both ends handle incomplete messaging gracefully—ADSL line might go up or down in the middle of messaging.
- Application is expected to have knowledge of the line state before trying to send a message to a CPE—if message is sent to controller and line is not up, controller discards message and sends a message back to the originating application so that originator can try again later.

- Terminal controller has an aging timer to clean up message registration when message response is lost or delayed.
- Protocol version is encoded in each message for forward and backward compatibility.
- Multiple packet response can be sent back as separate command.
- CPE SAR driver currently supports AAL5 payloads of up to 2048 bytes in length. The applications are expected to
  - Perform fragmentation if required.
  - Set the End of PDU indication in each Common Part Convergence Sublayer Protocol Data Unit or fragment generated.
- In-band VCs are not visible to the NMS, although it can be seen using the network interface debug menu.
- Total number of available VCs in the DSLAM are reduced because of an extra VC per active subscriber for the CPE in-band channel. The total number of transit VCs are reduced by 24 and the total number of network VCs are reduced by 16.

Following is the functional structure of the management plane:

- Connection Make—The application subscriber control is the CPE in-band channel controller. Subscriber control creates the in-band cross connects in the switch as soon as the ADSL link is created between a CPE and the ADSL Transmission Unit-central office (ATU-C). This cross connect is in addition to the provisioned VCs for your data traffic. This cross connect terminates at the DSLAM. The cross connect VPI/VCI on the CPE end is predefined, and the fiber channel selects the VPI/VCI on the SAR end. The connections terminating at the SAR are created or deleted by calling the SAR through interprocess communication (IPC).
- Connection Break—The subscriber control deletes the in-band cross connect as soon as the ADSL link is broken between the CPE and the ATU-C. This protects against wasting the bandwidth allocation for VCs that are not in use.
- Messaging—Messages are sent across as an AAL5 payload. The messages can be initiated from either the CPE or by any application in the DSLAM. There is a CPE in-band channel controller at each end, CPE and DSLAM. The DSLAM controller creates or deletes the cross connects when required. The controller also forwards or receives in-band messages to or from the local applications. The controlling application at DSLAM is the subscriber control.

Following is the message flow of the management plane:

- Message originated from DSLAM—The message destined for a CPE can be generated by any application in DSLAM. The application sends the message to the subscriber control through an IPC message queue. The subscriber control fills all the fields of the data link protocol data unit (DPDU).

  Subscriber control registers the source address and Advanced Intelligent Network Switch Capabilities (ASC) of this message, places the Invoke ID (INV ID) handle into the DPDU, and sends it to the CPE. This associates the response with the source application (this does not apply to release 2.3). If the subscriber control receives an outgoing message for a CPE, but the in-band channel is not up, it discards the message and notifies the application. Registration is needed only if the response is requested by the originating application.

  CPE reads the command-specific data based on the command field and processes it. The CPE application recognizes either ACK, RESPONSE, NAK, or no acknowledgement at all based on the message type. CPE might forward the message to various applications based on the command and data type. The ACK or NAK is sent back with just the command type field altered and the command-specific field eliminated. (This paragraph does not apply to release 2.3.)

Upon receipt of a message, subscriber control finds the actual source of the message based on the INV ID handle, and then forwards it to the source application through IPC. If the response does not come or is delayed, the subscriber control unregisters the message. The response from CPE is expected within 5 seconds.

- Message originated from CPE—CPE creates the message and sends it to the subscriber control of the DSLAM. The subscriber control detects that this message is a command and not a response to a command, based on the command sequence (C/SEQ) field.

  Subscriber control forwards the message to the respective application based on the command and object type. There is no source registration required in this direction of messaging.

  The application responds to the command and shows either RESPONSE, ACK, or NAK based on the message type field.

- Broadcast messages—The subscriber control receives all the broadcast messages on a different ASC and sends the message to each active CPE.

6.4 Data Structure and Code modules

The data structure and its processor is divided into two sections based on the functionality of the two layers. The first section defines the protocol existing between the AAL5 and the application layer, which is the data link layer in the DPDU. The second section is the application layer protocol, which sits at the data part in the DPDU. All the DPDU processing is done at the terminal controller. In the case of the DSLAM, the terminal controller is the subscriber control application.

One process done by the terminal controller consists of a DPDU being sent as an AAL5 payload to the other end. Following are the functions of this layer:

- Route the incoming command messages to the applications based on the type of command.
- Register the outgoing messages with their source application and ASC (not in release 2.3).
- Route the incoming response messages to the applications based on the message registration.
- Send the message or response to the CPE.
- Resend the message if a response is not received (not in release 2.3).
- Acknowledge message by sending it back with the command changed to a response and no data field (not in release 2.3).
- Send a negative acknowledgment to the originating application if the ADSL link is not active (not in release 2.3).

Figure 22 Data Structure

PID	VER	C/SEQ	SIZE	INV ID	DATA

27820

The 1-byte Protocol Identifier (PID) tells CPE which protocol is running (Point-to-Point Protocol, 1483 bridged, or in-band management). The information given by the PID is redundant because the in-band data is always received on a well-known VC. The value for this release is 01.

The 1-byte Protocol Version (VER) identifies which version of the protocol format is running. This accommodates forward and backward compatibility. The first version is 01.

The 2-byte Command/Sequence (C/SEQ) field has two sections. The most significant bit signifies whether it is a command or response, and the remaining 15 bits signify the sequence number. (This field is not used in release 2.3.)

The 1-bit C field is used to identify the data-link command or response. If the response is not received, the message is sent again. The responding terminal sends a response back with this field altered and with no data enclosed (make sure that the SEQ remains the same). The SEQ field is the receiving end; if the message is received twice, it discards the second message. This field is initialized by the originating terminal. An example follows:

- 0 COMMAND
- 1 RESPONSE

Size (SIZE—U16) is an unsigned 16-bit field that is filled by the originating terminal to give the size of the DPDU in bytes. The minimum size of the message is 14 bytes (when it is an acknowledgment message).

Invoke ID (INV ID—U16) is an unsigned 16-bit field that is filled by the terminal controller. This works as a handle and helps the terminal controller route the message back to the originating application on responses or acknowledgments. The originating terminal registers the originating application address and the receiving ASC. This field is very flexible and the usage entirely depends on the originating terminal controller. (This field is not used in release 2.3.)

DATA is the application layer data field, which is passed to the receiving application. The detailed data structure of DATA is defined in the next section.

6.4.1 Application Data

Application Data is processed by the application. The structures in this section are application layer data structures that reside in the DATA portion of the DPDU. Each structure has two sections, the common part and the command-specific section. The size of DATA is variable but the minimum size is the size of the common part. Following is the function of this application layer:

- Break the message into more than 1 AAL5 packet if the size is greater than 2048 bytes or if the Server Message Block limitation applies (720 bytes).
- Check whether response or acknowledgment is needed.
- Command-specific processing.

A common part of the application data is the common part structure (6 bytes). Following is the required data:

- Common_part{
- u8 command,
- u8 message_type,
- u16 reserved_1, /*not used in release 2.3*/
- u16 size,
- u16 reserved_2,} /*not used in release 2.3*/

The command field identifies the kind of message in the command-specific data. For each command, the command-specific structure should be defined. The terminal controller forwards the message based on the command and the object value defined in the command-specific part. This way, each terminal does not have to carry the information about the supporting applications at the far end. Following is a list of messages and their functions:

- 01 DISPLAY—displays the string in the far end

- 02 TERMINATE—command for displaying the line termination
- 03 GET DATA—gets a value for an object attribute defined in the information model
- 04 SET DATA—sets a value for an object attribute defined in the information model
- 05 NOTIFY—Notifies the far end about an event

Only DISPLAY is supported in release 2.3.

The Message_Type-U8 field signifies the type of message. An application responds based on the value set at this field by the originator. If the command is unknown, the responder responds with a negative acknowledgment to avoid the originator resending the message, which would cause an infinite loop or deadlock. This field gives the flexibility of not hard coding the acknowledgment requirement for each command. That is, a command can be changed from NO ACK to ACK without affecting the other end terminal. Following is a list of messages as they appear in this field:

- 00 COMMAND-NO ACK REQD /*not used in release 2.3*/
- 01 COMMAND ACK REQD
- 02 RESPONSE
- 03 ACK
- 04 NAK

The reserved_1 and reserved_2 data are for future use.

The Size field shows the size of the DATA.

6.4.2 Command-Specific Data

The following structures are command specific and are added when required. This is optional data.

The DISPLAY command displays a message on a window at the CPE end. This command uses 7-bit ASCII character coding. The data structure for this command follows:

- {
- u8 filter_level,
- string of bytes. /* NULL terminated */
- }

The filter level defines the priority level of the message that displays. At the CPE end, you are able to filter the messages based on the priorities. Following are the filter levels:

- 00 Filter Level 0 (used in release 2.3)
- 01 Filter Level 1 (not used in release 2.3)
- 02 Filter Level 2 (not used in release 2.3)
- 03 Filter Level 3 (not used in release 2.3)

The Terminate command is a special form of the Display command in that it displays the accompanying text message and then causes the CPE to disconnect the physical link.

The strings present in both Display and Terminate commands are 7-bit ASCII and are null terminated.

6.4.3 Software Restrictions and Configurations

Following is a list of software restrictions and configurations:

- The CPE in-band channel is not configurable or readable through the Management Information Base.
- Messages from the network interface or any module other than the system controller should have all the Server Message Block-related limitations applied.
- Software applications should make sure that the traffic from each application is self-policed. That is, there is no mechanism to detect a message that is spinning in an infinite loop and taking up ADSL bandwidth.
- The message protocol bytes are BIG ENDIAN oriented.
- The members of a structure are tightly packed—no padding.

7. DC Characteristics

All requirements of this document are met in the presence of all POTS loop currents from 0 mA to 100 mA and differential loop voltages as follows:

- DC voltages of 0V to 105V
- Ringing signals of 40V to 150V root mean square (rms) at any frequency from 15.3 Hz to 68 Hz with a DC component in the range from 0V to 105V

The DC resistance from tip to ring at the Public Switched Telephone Network (PSTN) interface with the U-C interface shorted, or at the POTS interface with the U-R interface shorted, is less than or equal to 25 ohms. The DC resistance from tip to ground and from ring to ground at the PSTN interface with the U-C interface open, or at the POTS interface with the U-R interface open, is greater than or equal to 5 Mohms.

8. POTS Splitter Characteristics

The POTS signal occupies nominal passband frequencies up to 3.4 kHz. The POTS splitter, if installed, provides the following functions:

- Combines RTU-R signals and POTS signals for output to the U-R interface
- Separates RADSL and POTS signals input at the U-R interface
- Protects the POTS network from RADSL signal interference
- Protects the RADSL network from POTS ringing, dial pulses, and ring trip

The POTS splitter is passive and is capable of performing these functions in the absence of power at the RTU-R or RTU-C.

9. References

ANSI T1E1/97-104R2a, *Draft Physical Layer Specification for CAP/QAM-Based Rate Adaptive Digital Subscriber Line (RADSL).*

ANSI T1.413-1995, *Telecommunications—Network and Customer Installation Interfaces: Asymmetric Digital Subscriber Line (ADSL) Metallic Interface.*