INTERNATIONAL TELECOMMUNICATION UNION

ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

G.989.2

(11/2001)

SERIES G: TRANSMISSION SYSTEMS AND MEDIA, DIGITAL SYSTEMS AND NETWORKS

Digital sections and digital line system – Optical line systems for local and access networks

---

Phoneline networking transceivers – Payload format and link layer requirements

ITU-T Recommendation G.989.2

ITU-T G-SERIES RECOMMENDATIONS

TRANSMISSION SYSTEMS AND MEDIA, DIGITAL SYSTEMS AND NETWORKS

INTERNATIONAL TELEPHONE CONNECTIONS AND CIRCUITS	G.100–G.199
GENERAL CHARACTERISTICS COMMON TO ALL ANALOGUE CARRIER-TRANSMISSION SYSTEMS	G.200–G.299
INDIVIDUAL CHARACTERISTICS OF INTERNATIONAL CARRIER TELEPHONE SYSTEMS ON METALLIC LINES	G.300–G.399
GENERAL CHARACTERISTICS OF INTERNATIONAL CARRIER TELEPHONE SYSTEMS ON RADIO-RELAY OR SATELLITE LINKS AND INTERCONNECTION WITH METALLIC LINES	G.400–G.449
COORDINATION OF RADIOTELEPHONY AND LINE TELEPHONY	G.450–G.499
TESTING EQUIPMENTS	G.500–G.599
TRANSMISSION MEDIA CHARACTERISTICS	G.600–G.699
DIGITAL TERMINAL EQUIPMENTS	G.700–G.799
DIGITAL NETWORKS	G.800–G.899
DIGITAL SECTIONS AND DIGITAL LINE SYSTEM	G.900–G.999
General	G.900–G.909
Parameters for optical fibre cable systems	G.910–G.919
Digital sections at hierarchical bit rates based on a bit rate of 2048 kbit/s	G.920–G.929
Digital line transmission systems on cable at non-hierarchical bit rates	G.930–G.939
Digital line systems provided by FDM transmission bearers	G.940–G.949
Digital line systems	G.950–G.959
Digital section and digital transmission systems for customer access to ISDN	G.960–G.969
Optical fibre submarine cable systems	G.970–G.979
Optical line systems for local and access networks	**G.980–G.989**
Access networks	G.990–G.999
QUALITY OF SERVICE AND PERFORMANCE	G.1000–G.1999
TRANSMISSION MEDIA CHARACTERISTICS	G.6000–G.6999
DIGITAL TERMINAL EQUIPMENTS	G.7000–G.7999

For further details, please refer to the list of ITU-T Recommendations.

ITU-T Recommendation G.989.2

Phoneline networking transceivers – Payload format and link layer requirements

Summary

This Recommendation specifies a payload format and link layer procedures for Phoneline Networking Transceiver (PNT) devices. It assumes the use of Medium Access Control (MAC) and encapsulation procedures defined in ITU-T Rec. G.989.1.

Requirements in this Recommendation are intended to ensure interoperability between devices. They are written from the perspective of a conforming transmitter, although some minimum performance requirements are established for receivers.

Source

ITU-T Recommendation G.989.2 was prepared by ITU-T Study Group 15 (2001-2004) and approved under the WTSA Resolution 1 procedure on 29 November 2001.

FOREWORD

The International Telecommunication Union (ITU) is the United Nations specialized agency in the field of telecommunications. The ITU Telecommunication Standardization Sector (ITU-T) is a permanent organ of ITU. ITU-T is responsible for studying technical, operating and tariff questions and issuing Recommendations on them with a view to standardizing telecommunications on a worldwide basis.

The World Telecommunication Standardization Assembly (WTSA), which meets every four years, establishes the topics for study by the ITU-T study groups which, in turn, produce Recommendations on these topics.

The approval of ITU-T Recommendations is covered by the procedure laid down in WTSA Resolution 1.

In some areas of information technology which fall within ITU-T's purview, the necessary standards are prepared on a collaborative basis with ISO and IEC.

NOTE

In this Recommendation, the expression "Administration" is used for conciseness to indicate both a telecommunication administration and a recognized operating agency.

INTELLECTUAL PROPERTY RIGHTS

ITU draws attention to the possibility that the practice or implementation of this Recommendation may involve the use of a claimed Intellectual Property Right. ITU takes no position concerning the evidence, validity or applicability of claimed Intellectual Property Rights, whether asserted by ITU members or others outside of the Recommendation development process.

As of the date of approval of this Recommendation, ITU had received notice of intellectual property, protected by patents, which may be required to implement this Recommendation. However, implementors are cautioned that this may not represent the latest information and are therefore strongly urged to consult the TSB patent database.

CONTENTS

Page

Page

ITU-T Recommendation G.989.2

Phoneline networking transceivers – Payload format and link layer requirements

1 Scope

This Recommendation defines the payload modulation, payload framing, and link layer protocols for Phoneline Networking Transceivers (PNT) with the following features:

– PHY-layer payload transmission rates of 4 to 32 Mbit/s, with a nominal effective throughput rate equivalent to 10BASE-T Ethernet, with provision for higher rates in future Recommendations.

– Rate adaptive transceivers that optimize data rates and packet error rates for dynamically varying channel conditions on a per-packet basis.

– A modulation technique that provides for robust communication over highly frequency-selective channels.

– Optional Limited Automatic Repeat reQuest (LARQ) which conceals impulse noise errors.

– Backward compatible with existing PNT devices in the field.

– Compatibility with other phoneline services such as POTS, V.90/V.92, ISDN and G.992.2.

– Uses the MAC and encapsulation procedures defined in ITU-T Rec. G.989.1.

2 References

The following ITU-T Recommendations and other references contain provisions which, through reference in this text, constitute provisions of this Recommendation. At the time of publication, the editions indicated were valid. All Recommendations and other references are subject to revision; users of this Recommendation are therefore encouraged to investigate the possibility of applying the most recent edition of the Recommendations and other references listed below. A list of the currently valid ITU-T Recommendations is regularly published.

– ITU-T G.989.1 (2001), *Phoneline networking transceivers – Foundation*.

– ISO/IEC 8802-3:2000, *Information technology – Telecommunications and information exchange between systems – Local and metropolitan area networks – Specific requirements – Part 3: Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications*.

– ISO/IEC 15802-3:1998, *Information technology – Telecommunications and information exchange between systems – Local and metropolitan area networks – Common specifications – Part 3: Media Access Control (MAC) Bridges*.

– IEEE 802.1Q-1998, *IEEE standard for local and metropolitan area networks: Virtual Bridged Local Area Networks*.

3 Definitions

This Recommendation defines the following terms:

3.1 broadcast packet: A packet with the all-ones Destination Address (FF FF FF FF FF FF_{16}).

3.2 capability and status announcement: A link-layer control protocol that is used to flood status information between stations with low overhead.

3.3 collision fragment: The fixed transmission sequence consisting of Preamble, Frame Header, DA, SA, ET and EOF.

3.4 EOF sequence: The 4-symbol sequence that is appended to the physical layer frame, consisting of the first four symbols of the TRN sequence.

3.5 ethernet: ISO/IEC 8802-type networks.

3.6 link integrity: A process that derives a user indication that the interface is attached to the phoneline and can detect at least one other station.

3.7 link level priority: The software priority class associated with the link-layer packet. This value may be mapped when converting to/from PHY Priority.

3.8 PHY priority: The 3-bit absolute priority used by the G.989.1 media access control to rank preference to frames waiting to be transmitted on the channel. Priority 7 has preference over Priority 0.

3.9 self-NEXT: NEXT from other systems of the same type.

3.10 system margin: A set of values for impairment levels at which a receiver does not exceed a specified frame error rate on a given test loop.

3.11 valid CS frame: The minimum transmitter signal that is acceptable to implementations of carrier sense and collision detection.

4 Abbreviations

This Recommendation uses the following abbreviations:

BEB Binary Exponential Backoff

CSA Capability and Status Announcement

CSMA/CD Carrier Sense Multiple Access with Collision Detection

DA Destination Address

DFPQ Distributed Fair Priority Queuing

EOF End Of Frame

FDQAM Frequency Diverse QAM

FER Frame Error Rate

HCS Header Check Sequence

JEDEC Joint Electron Device Engineering Council

LARQ Limited Automatic Repeat reQuest

LSB Least Significant Bit

MAC Medium Access Control

MII Media Independent Interface

MSB Most Significant Bit

NEXT Near-End Crosstalk

NID Network Interface Device

PAR Peak to average ratio

PDU Protocol Data Unit

PE	Payload Encoding
PHY	Physical Layer
POTS	Plain Old Telephone Service
ppm	parts per million
QAM	Quadrature Amplitude Modulation
rms	root mean square
SA	Source Address
SI	Scrambler Initialization
TCP	Transmission Control Protocol

5 Physical layer specification

A functionally oriented view of system reference model is shown in Figure 1. The MAC and G.989.1 PHY functions are defined in ITU-T Rec. G.989.1. The Payload PHY and Link Layer Protocols functions are defined in this Recommendation.

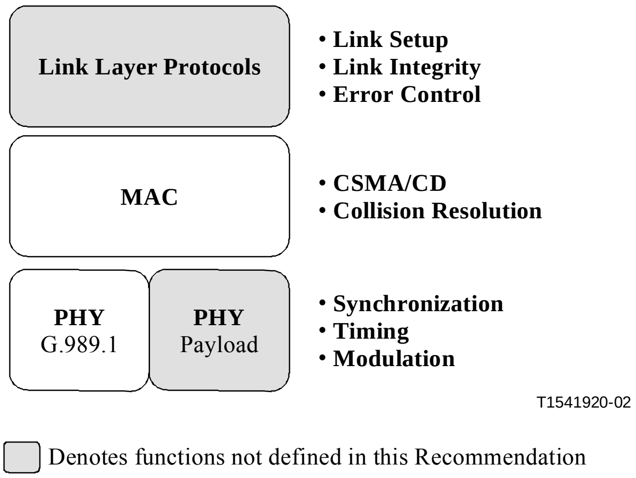

Figure 1/G.989.2 – Functional view of reference model

5.1 Overview

This Recommendation defines the contents of a PNT frame that is encapsulated by a G.989.1 header and trailer, i.e. that portion of a PNT frame following the G.989.1 FT octet, and prior to the G.989.1 EOF delimiter.

The modulation used by the PNT payload is defined in 5.2. The framing is defined in 5.3.

5.2 Modulation

The G.989.2 frame shall use Quadrature Amplitude Modulation (QAM).

The two-dimensional symbol rate shall be 4 000 000 symbols/s, with a tolerance of ±0.01%.

5.2.1 Carrier frequency and tolerance

The carrier frequency shall be 7 MHz.

The carrier clock shall be locked to the symbol clock.

5.2.1.1 Clock jitter

The rms jitter of the transmitter clock shall be less than 70 ps, averaged over a sliding 10 µs window.

5.2.2 Transmit filters

Transmit filtering sufficient to meet the PSD mask defined in ITU-T Rec. G.989.1 shall be applied. Specifically, for 4D symbol mapping (see 5.2.4.1), the transmit filter shall not be narrowed from that used for 2D symbol mapping.

NOTE 1 – Use of such a wide transmit filter for 4D symbol mapping results in redundancy in the frequency domain. The use of this "FDQAM" technique increases robustness on frequency-selective impaired channels

NOTE 2 – The scaling requirements of 5.2.5.3 result in slightly different power output levels for different constellation encodings. PE values of 1 or 2 (see 5.3.2.3) result in the highest transmit power.

5.2.2.1 Passband ripple

The ripple between 4.75 and 6.25 MHz and between 8.0 and 9.25 MHz shall be less than 2.0 dB.

5.2.3 Transmitter symbol response

The transmitter output magnitude shall be upper-bounded by the following temporal mask (see Figure 2). The response shall be measured across a 100 Ω load between tip and ring at the transmitter's W1 interface (see clause 4/G.989.1). Output before t = 0 and after t = 5.0 µs shall be less than 0.032% of the peak amplitude.

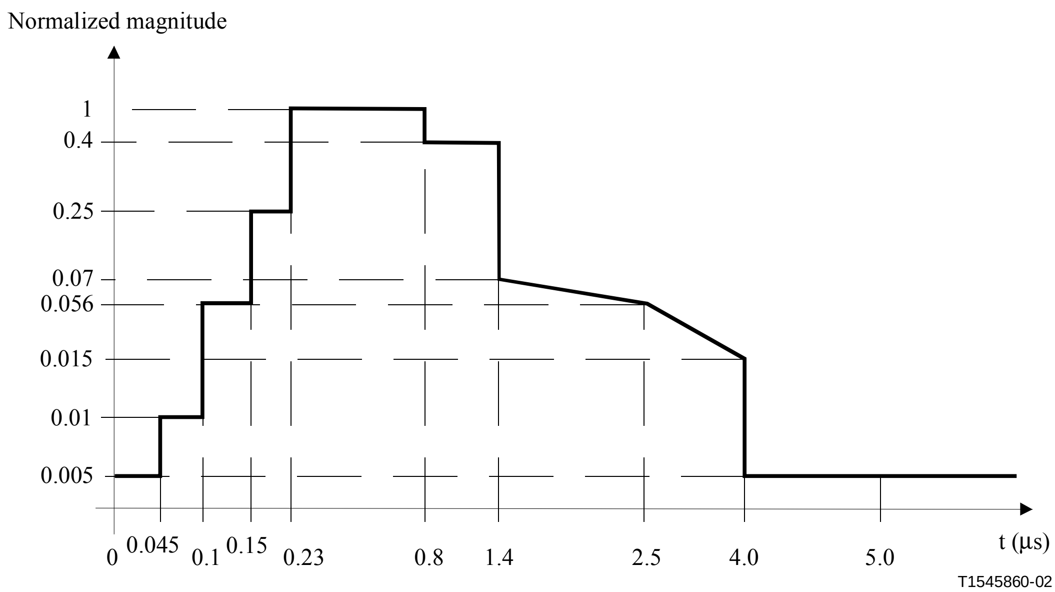

Figure 2/G.989.2 – Transmitter symbol response magnitude mask

5.2.4 Symbol mapping

All G.989.2 frames shall use the 4D symbol mapping defined in 5.2.4.1. Optionally, a portion of the G.989.2 frame may use the 2D symbol mapping defined in 5.2.4.2.

5.2.4.1 4D symbol mapping

Each 4D symbol shall consist of two consecutive two-dimensional QAM symbol intervals. The first QAM symbol shall be modulated with 2 to 8 data bits using the constellation encoding in 5.2.5. The second symbol interval shall be transmitted with zero amplitude.

5.2.4.2 2D symbol mapping

Each 2D symbol shall consist of a single QAM symbol, modulated with 2 to 8 data bits using the constellation encoding in 5.2.5.

5.2.4.3 2D to 4D transition

When switching from 2D to 4D symbol mapping within a given G.989.2 frame, an additional zero amplitude symbol shall be transmitted immediately following the last 2D symbol and immediately preceding the first 4D symbol.

5.2.5 Constellation encoder

5.2.5.1 Bit to symbol mapping

The incoming bits shall be grouped into N-bit symbol intervals, where $2 \leq N \leq 8$. Each N-bit group shall be mapped to a 2D symbol as shown in Figures 3 through 9. The symbol values are shown with bits ordered such that the right-most bit is the first bit in time and the left-most bit is the last bit in time.

All constellations except for 3 bits-per-symbol lie on a uniform square grid, and all constellations are symmetric about the real and imaginary axes. The 3 bits-per-symbol constellation points are equally spaced around an equal-amplitude circle.

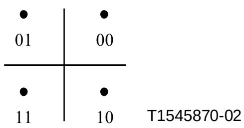

Figure 3/G.989.2 – 2 bits-per-symbol

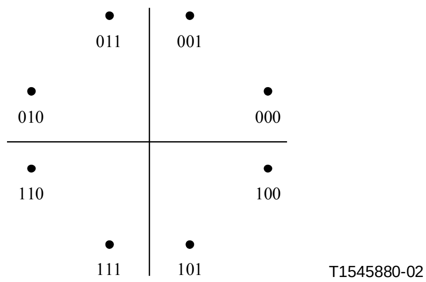

Figure 4/G.989.2 – 3 bits-per-symbol

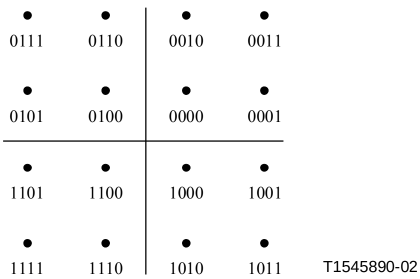

Figure 5/G.989.2 – 4 bits-per-symbol

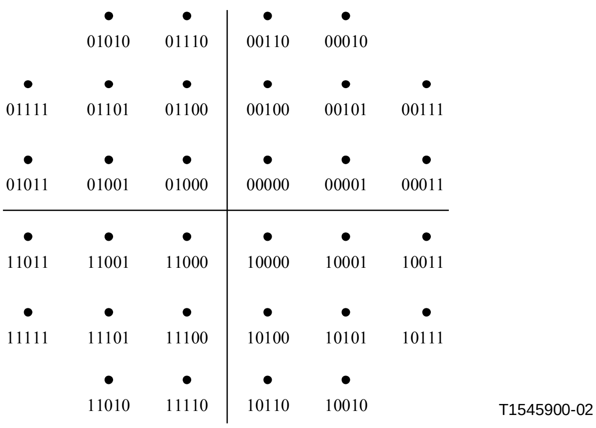

Figure 6/G.989.2 – 5 bits-per-symbol

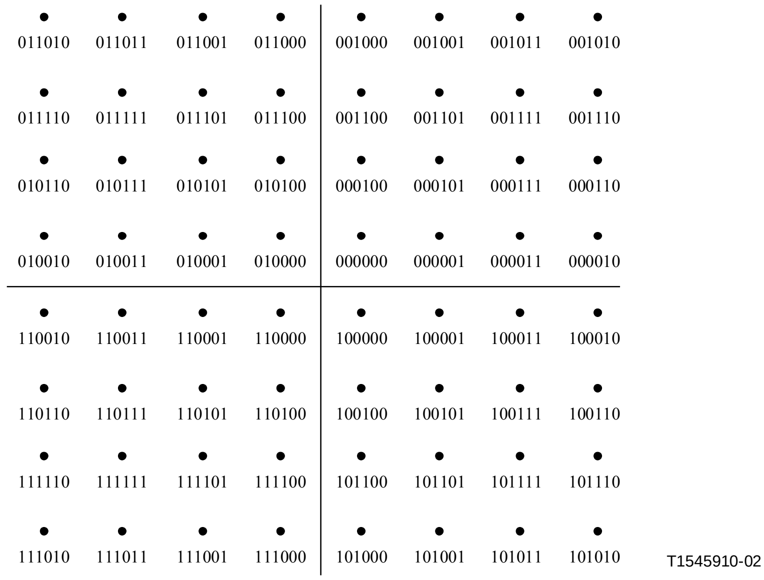

Figure 7/G.989.2 – 6 bits-per-symbol

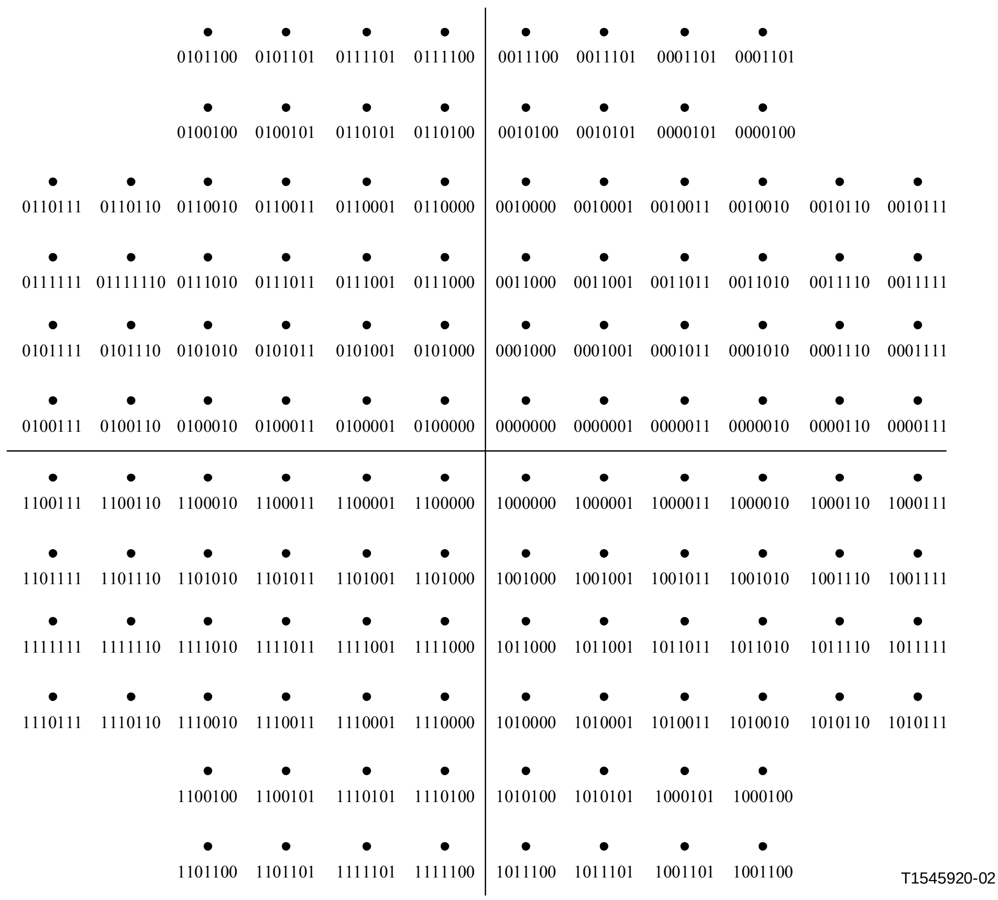

Figure 8/G.989.2 – 7 bits-per-symbol

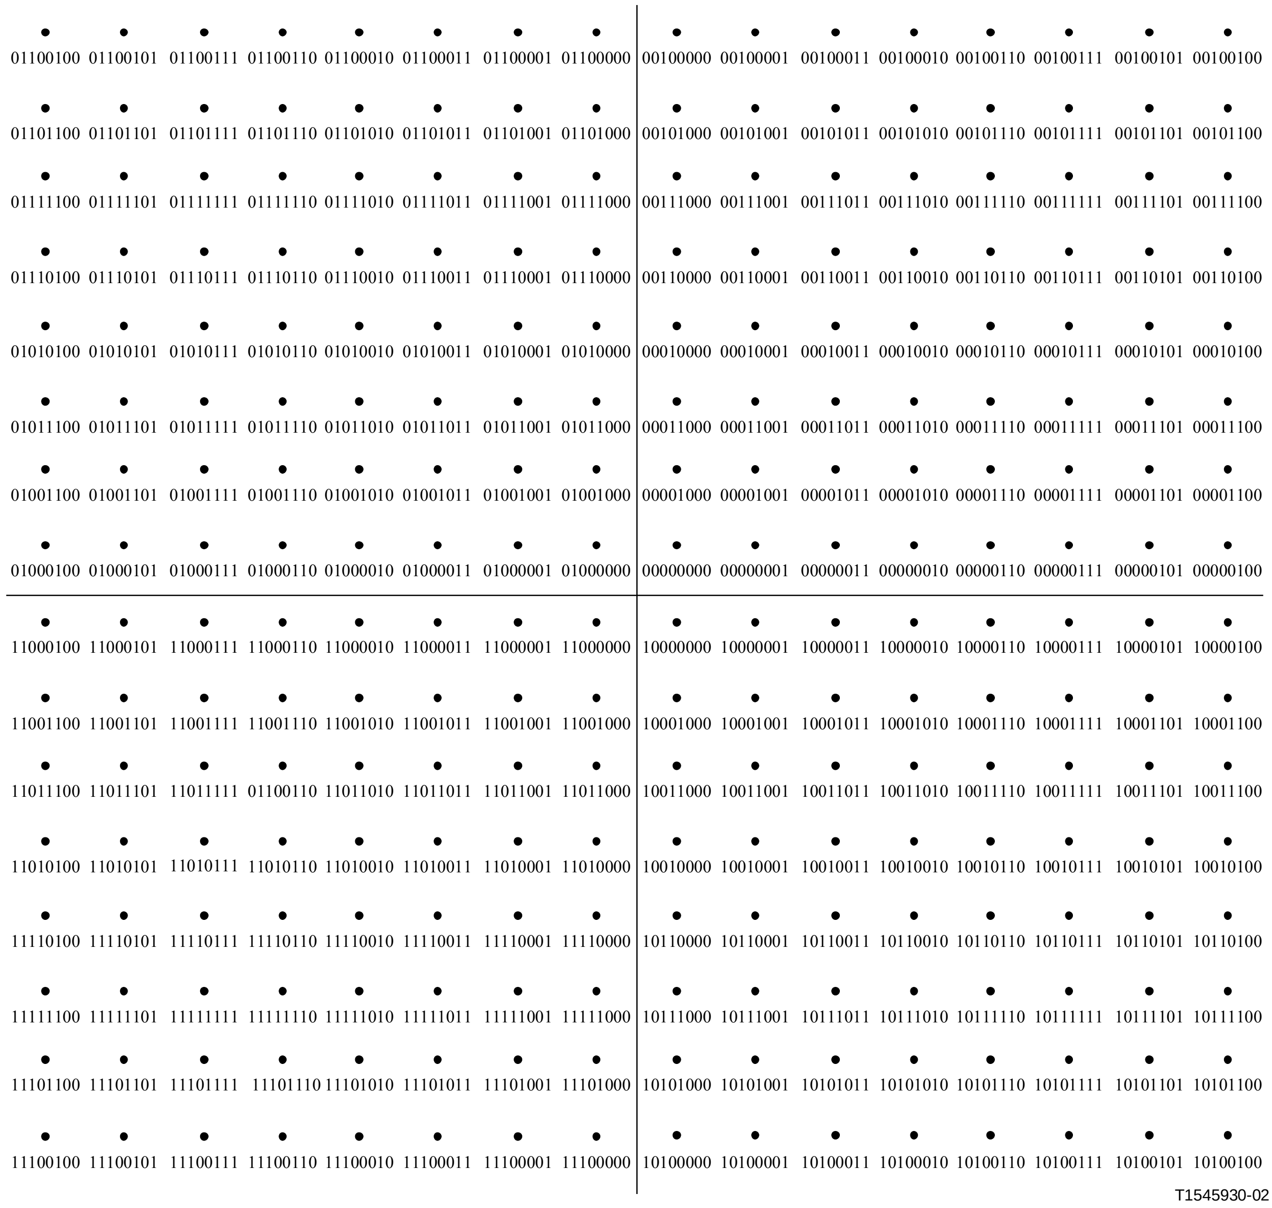

Figure 9/G.989.2 – 8 bits-per-symbol

5.2.5.2 Symbol mapping transitions

Symbol mapping transitions within a G.989.2 frame shall occur on a data bit octet boundary, for transitions between 2D and 4D symbol mapping or for transitions in bits-per-symbol mapping. If the number of data bits mapped to the symbol immediately preceding such a transition is less than the current value of N, a sufficient number of bits of value zero shall be appended to the data bits in order to complete this last symbol in the sequence.

5.2.5.3 Scaling

The relative scaling of different constellations is given by Table 1. The constellation points shall be scaled such that the reference points have the values shown, within a minimum-distance tolerance of ±4%.

Table 1/G.989.2 – Constellation reference points

Bits per symbol	Reference point	4D mapping amplitude	2D mapping amplitude
2	00	1	$\frac{1}{\sqrt{2}}$
3	000	1	$\frac{1}{\sqrt{2}}$
4	0000	$\frac{1}{3}$	$\frac{1}{3\sqrt{2}}$
5	00000	$\frac{1}{4}$	$\frac{1}{4\sqrt{2}}$
6	000000	$\frac{1}{7}$	$\frac{1}{7\sqrt{2}}$
7	0000000	$\frac{1}{9}$	$\frac{1}{9\sqrt{2}}$
8	00000000	$\frac{1}{15}$	$\frac{1}{15\sqrt{2}}$
NOTE – The constellation points are scaled such that the outermost points have approximately equal magnitude.			

5.3 Framing

The PHY payload in each physical frame is formatted as an Ethernet-type Link Level frame. In addition, the Ethernet Frame is preceded by a 3-octet Frame Control field, and followed by a 2-octet CRC-16 field and, possibly, a variable-length pad field.

The frame format is shown in Figure 10. It consists of a low bit-rate header section, a variable-rate data section, and a low bit-rate trailer. Some parts of the frame are not scrambled, as described in 5.3.6.

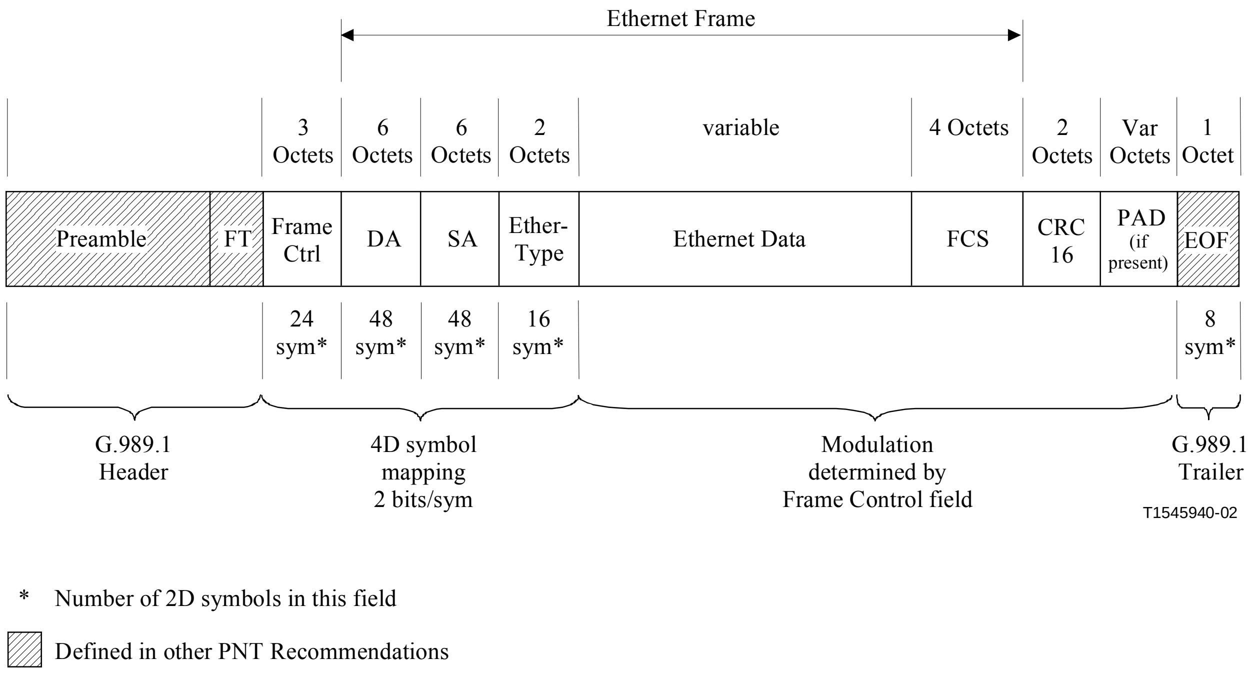

Figure 10/G.989.2 – PHY frame format

5.3.1 Bit order

Except where otherwise stated, all fields shall be transmitted most-significant octet first, with the LSB within each octet transmitted first. The LSB within a field is designated as Bit No. 0. Diagrams show most significant bits or octets to the left.

5.3.2 Frame control definition

The frame control field is a 24-bit field defined in Table 2.

Table 2/G.989.2 – Frame control field

Field	Bit No.	Bits	Description
RSVD	23	1	Reserved. This field shall be set to zero by the transmitter, and the receiver shall ignore it.
PRI	22:20	3	Priority (0-7)
SI	19:16	4	Scrambler Initialization
PE	15:8	8	Payload Encoding
HCS	7:0	8	Header Check Sequence

Hence, with the bit-ordering defined in 5.3.1, the frame control field is transmitted in the order shown in Figure 11.

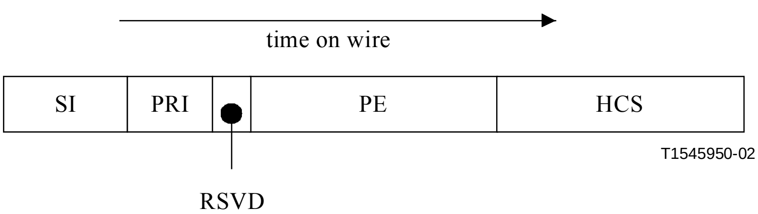

Figure 11/G.989.2 – Frame control field order

5.3.2.1 Scrambler Initialization bits

This 4-bit field shall be set to the value used to initialize the scrambler, as described in 5.3.6.

5.3.2.2 Priority

Priority refers to the MAC priority mechanism. The 3-bit PHY priority value (PRI) shall be set to the G.989.1 priority level of the frame.

For PNT stations that do not implement class-of-service the PRI field shall be ignored on receive, and shall be transmitted set to 1.

5.3.2.3 Payload encoding

This field determines the constellation encoding of the variable bit-rate portion of the PNT payload. The values are defined as follows in Table 3.

Table 3/G.989.2 – Payload encoding

Value	Interpretation
0	Non-standard mode
1	4D symbol mapping, 2 bits per symbol
2	4D symbol mapping, 3 bits per symbol
3	4D symbol mapping, 4 bits per symbol
4	4D symbol mapping, 5 bits per symbol
5	4D symbol mapping, 6 bits per symbol
6	4D symbol mapping, 7 bits per symbol
7	4D symbol mapping, 8 bits per symbol
8	Non-standard mode
9	2D symbol mapping, 2 bits per symbol
10	2D symbol mapping, 3 bits per symbol
11	2D symbol mapping, 4 bits per symbol
12	2D symbol mapping, 5 bits per symbol
13	2D symbol mapping, 6 bits per symbol
14	2D symbol mapping, 7 bits per symbol
15	2D symbol mapping, 8 bits per symbol
16-256	Reserved for use by ITU-T
NOTE – Receivers should discard frames whose PE value they do not understand.	

5.3.2.4 Header Check Sequence (HCS)

An 8-bit cyclic redundancy check (CRC) is computed as a function of the (unscrambled) 128-bit sequence in transmission order starting with the G.989.1 FT field and ending with the Ethernet SA field, with zeros substituted for the as-of-yet uncomputed HCS field. The encoding is defined by the following generating polynomial.

$$G(x) = x^8 + x^7 + x^6 + x^4 + x^2 + 1 \qquad (5\text{-}1)$$

Mathematically, the CRC value corresponding to a given frame is defined by the following procedure:

1) The first 8 bits of the input bit sequence in transmission order (i.e. the G.989.1 FT field) are complemented.
2) The 128 bits of the sequence in transmission order are then considered to be the coefficients of a polynomial M(x) of degree 127. (The first bit of the FT field corresponds to the x^{127} term and the last bit of the SA field corresponds to the x^0 term.)
3) M(x) is multiplied by x^8 and divided by G(x), producing a remainder R(x) of degree ≤ 7.
4) R(x) is multiplied by H(x) to produce N(x), where H(x) is defined as $H(x) = x^7 + x^6 + x^5 + x^4 + x^2 + x + 1$
5) N(x) is divided by G(x), producing a remainder R'(x) of degree ≤ 7.
6) The coefficients of R'(x) are considered to be an 8-bit sequence.
7) The bit sequence is complemented and the result is CRC'.

The 8 bits of CRC' are placed in the HCS field so that x^7 is the LSB of the octet and x^0 is the MSB of the octet. (The bits of the CRC' are thus transmitted in the order $x^7, x^6, \ldots x^1, x^0$.)

NOTE 1 – Although the HCS is embedded within the protected bit stream, it is calculated in such a way that the resulting 128-bit stream provides error-detection capabilities identical to those of a 120-bit stream with an 8-bit CRC appended. It should be noted that the resulting 128-bit sequence, considered as the coefficients of a polynomial of degree 127, when divided by G(x), will always produce a remainder equal to $x^7 + x^6 + x + 1$.

NOTE 2 – Because all fields covered by the HCS are transmitted at 4D symbol mapping and 2 bits per symbol, these fields should be received correctly in many cases where the payload is received in error. The HCS may be used in conjunction with soft-decision error statistics to determine with high probability whether the header was received correctly. This knowledge may be useful for optimizing the performance of ARQ and/or rate negotiation algorithms.

5.3.3 Ethernet frame

The bit fields starting with the DA field and ending with the FCS field in Figure 10 are identical to the corresponding fields described in ISO/IEC 8802-3, and are referred to as the Link-level Ethernet Frame.

NOTE 1 – An ISO/IEC 8802-3 Ethernet frame has an Ethernet preamble and start-frame-delimiter (SFD) bits prefixed to the Link-level frame; these bits are not present in PNT frames.

NOTE 2 – It is intended that assigned Ethernet MAC addresses are used for Destination Address (DA) and Source Address (SA).

The Ethernet frame consists of an integer number of octets.

Formatting and procedures for the Ethernet frame contents are further described in clause 6.

5.3.4 CRC-16

A 16-bit cyclic redundancy check (CRC) is computed as a function of the contents of the (unscrambled) Ethernet frame in transmission order, starting with the first bit of the DA field and ending with the last bit of the FCS field. The encoding is defined by the following generating polynomial.

$$G(x) = x^{16} + x^{12} + x^5 + 1 \qquad (5\text{-}2)$$

Mathematically, the CRC value corresponding to a given frame is defined by the following procedure.

8) The first 16 bits of the frame in transmission order are complemented.

9) The *n* bits of the frame in transmission order are then considered to be the coefficients of a polynomial M(x) of degree n–1. (The first bit of the Destination Address field corresponds to the $x^{(n-1)}$ term and the last bit of the FCS field corresponds to the x^0 term.)

10) M(x) is multiplied by x^{16} and divided by G(x), producing a remainder R(x) of degree ≤ 15.

11) The coefficients of R(x) are considered to be a 16-bit sequence.

12) The bit sequence is complemented and the result is the CRC.

The 16 bits of the CRC are placed in the CRC-16 field so that x^{15} is the least significant bit of the first octet, and the x^0 term is the most-significant bit of the last octet. (The bits of the CRC are thus transmitted in the order $x^{15}, x^{14}, \ldots x^1, x^0$.)

NOTE – The CRC-16, in conjunction with Ethernet's FCS, provides more protection from undetected errors than the FCS alone. This is motivated by environmental factors that will often result in a frame error rate (FER) several orders of magnitude higher than that of Ethernet, making the FCS insufficient by itself.

5.3.5 Pad

PNT payloads using 2D symbol mapping shall include a Pad field consisting of an integral number of octets.

The last octet of the Pad field is set to a value, designated as PAD_LENGTH, which is equal to the number of octets in the Pad field, minus one. The preceding Pad field octets are all set to a value of zero.

PAD_LENGTH shall be set to a value such that the transmission time of the frame, from the beginning of the G.989.1 preamble through the last symbol of the G.989.1 EOF delimiter, is at least 92.5 µs.

PNT payloads using 4D symbol mapping exclusively, shall not include a Pad field.

NOTE – The Pad field ensures that a collision fragment can discriminated from a valid frame by the transmission length detected by the carrier sense function as defined in ITU-T Rec. G.989.1. An example of a compliant formula for generating PAD_LENGTH is **max**(102–N,0), where N is the number of octets from DA to FCS, inclusive.

5.3.6 Scrambler

The PNT payload contents shall be scrambled using the frame-synchronized scrambler shown in Figure 12. Scrambling shall commence with the 9th transmitted bit of the Frame Control field (i.e., the first PE bit), and continue through to the last bit of CRC-16 (or the last bit of the Pad field if present), inclusive.

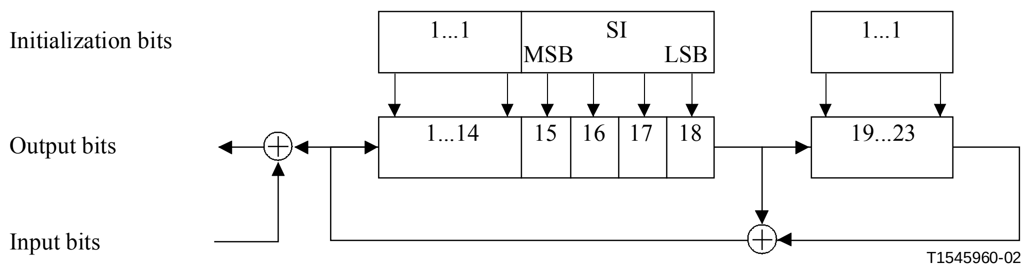

Figure 12/G.989.2 – Data scrambler

The scrambler uses the following generating polynomial.

$$G(x) = x^{23} + x^{18} + 1 \qquad (5\text{-}3)$$

Bits 15 through 18 of the shift register shall be initialized with a 4-bit random number. This value shall be placed in the SI sub-field defined in 5.3.2.1 in the order such that register position 15 is the MSB (bit 19 of Frame Control) and bit 18 is the LSB (bit 16 of Frame Control).

NOTE – The use of a random initial scrambler state results in a uniform power spectral density (PSD) measured over multiple similar frames. This eliminates the problem of tones in the PSD from highly correlated successive packets. In order to achieve this objective, the values of SI should be uniformly distributed.

5.4 Minimum device requirements

G.989.2 device at a minimum shall be capable of transmitting and receiving 4D-symbol-mapping - modulated frames.

Transmitters at a minimum shall be capable of transmitting all constellations from 2 bits-per-symbol to 8 bits-per symbol (PE values 1-7). Receivers at a minimum shall be capable of receiving all constellations from 2 bits per symbol to 6 bits per symbol (PE values 1-5).

NOTE – G.989.2 devices can also be made to function in a mode that is interoperable with existing non-standard, lower-speed PNT devices in the field. Such operation is outside the scope of this Recommendation.

5.5 Transceiver input impedance

5.5.1 Passband return loss

The average return loss of the transceiver with respect to a 100-ohm resistive load shall exceed 12 dB between 4.75 and 9.25 MHz. This requirement applies to the transceiver powered on or in low-power mode (transmitter powered off). The average return loss with respect to a 100-ohm resistive load shall exceed 6 dB between 4.75 and 9.25 MHz with the transceiver unpowered.

5.5.2 Stopband input impedance

The magnitude of the input impedance shall be greater than 10 ohms from 0-30 MHz and shall conform to the following lower-bound mask (see Table 4).

Table 4/G.989.2 – Input impedance lower-bound mask

Frequency range (kHz)	Min impedance (Ohms)
0 < f ≤ 0.285	1 M
0.285 < f ≤ 2.85	100 k
2.85 < f ≤ 28.5	10 k
28.5 < f ≤ 95	4.0 k
95 < f ≤ 190	2.0 k
190 < f ≤ 285	1.4 k
285 < f ≤ 380	1.0 k
380 < f ≤ 475	850
475 < f ≤ 570	700
570 < f ≤ 665	600
665 < f ≤ 760	525
760 < f ≤ 855	450

Table 4/G.989.2 – Input impedance lower-bound mask

Frequency range (kHz)	Min impedance (Ohms)
855 < f ≤ 950	400
950 < f ≤ 1000	350
1000 < f ≤ 1400	175
1400 < f ≤ 2300	100
2300 < f ≤ 2850	50
2850 < f ≤ 3085	25
3085 < f ≤ 3725	10
3725 < f ≤ 3935	25
3935 < f ≤ 4000	50
10 000 < f ≤ 10 450	40
10 450 < f ≤ 10 925	25
10 925 < f ≤ 13 125	10
13 125 < f ≤ 14 175	25
14 175 < f ≤ 16 800	50
16 800 < f ≤ 21 000	100
21 000 < f ≤ 30000	50

This requirement applies to the transceiver powered on, in low-power mode (transmitter powered off), or unpowered.

NOTE – Figure 13 shows an example of the input impedance of a compliant transceiver.

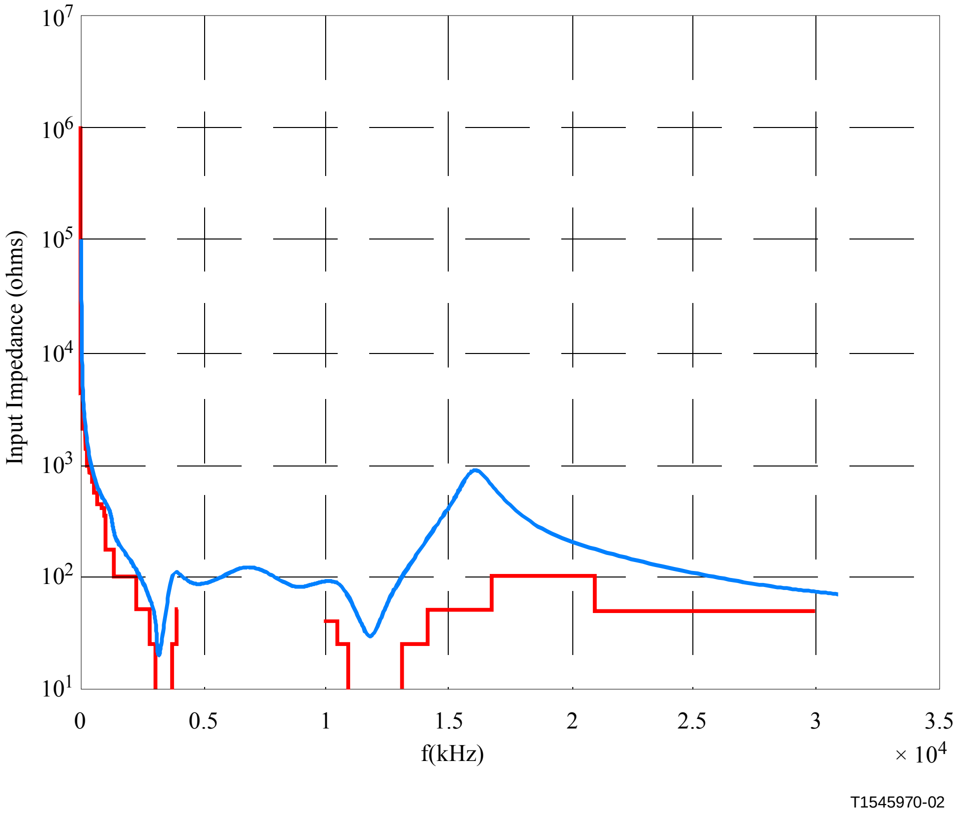

Figure 13/G.989.2 – Example input impedance (Informative)

6 Link Layer Protocol Specification

6.1 Overview

This clause defines procedures for implementing the following link control functions:

- Rate Negotiation.
- Link Integrity.
- Capability Announcement.
- Limited Automatic Repeat reQuest (LARQ).

These link functions use control frames to carry protocol messages between stations. This Recommendation includes a standardized mechanism for Link Layer network control and encapsulation. Control frames are data link layer frames that are identified by the Ethertype value $886C_{16}$ in the Type/Length field of the frame, and further distinguished by individual sub-types. Link Control frames are not seen by layer 3 (IP) of the network stack, and are not bridged between PNT network segments.

6.2 Basic Link Layer Frame Format

There are two basic formats for a Link Control Frame: a short subtype and a long subtype. The short subtype format may be used for control frames where the amount of control information is less than 256 octets. The long subtype format is used for control frames where the amount of control information exceeds 255 octets.

NOTE – The control and encapsulation frames described in this Recommendation use the short subtype format.

6.2.1 Short format

Table 5/G.989.2 – Short Format Link Control Frame

Field	Length	Explanation
DA	6 octets	Ethernet Destination Address
SA	6 octets	Ethernet Source Address
Ethertype	2 octets	$886C_{16}$
SSType	1 octet	0 Non-standard 1 Rate Request Control Frame 2 Link Integrity Short Frame 3 Capabilities Announcement 4 LARQ 5 Vendor-specific short format type 6-127 Reserved for future use by the ITU-T (See Note)
SSLength	1 octet	Number of control information octets, starting with the SSVersion field (or the first octet following SSLength if it is not defined as SSVersion) and ending with the second (last) octet of the Next Ethertype field. Min is 2 and max is 255.
SSVersion	1 octet	Version number of the control information
Data	0-252 octets	Control information
Next Ethertype	2 octets	Ethertype/length of next layer protocol; 0 if none
Pad	41-0 octets	Padding required to meet minimum if data < 41 octets
FCS	4 octets	Frame Check Sequence
NOTE – SSType values 128-255 correspond to the long subtype.		

The SSVersion field specifies which format version of the control information is used. This allows for the future extension of each SSType.

Receivers shall check SSLength to ensure that sufficient control information is present. Newer, backwards compatible, frame format versions may contain additional fixed data fields, but for backwards compatibility should contain the fixed fields specified in earlier formats.

Receivers shall interpret all supported SSType frames using the latest supported SSVersion that is less than or equal to the SSVersion indicated in the received frame. Unknown fields shall be ignored. Encapsulated data from unsupported (newer) SSVersions of supported encapsulating SSType frames shall be passed to the layer above.

The Next Ethertype field is required for all Short Format Link Control Frame headers. Among other things, it supports backwards compatibility by enabling receivers to always strip short format link layer headers. If the Next Ethertype field is zero, then the frame is a basic control frame and should be dropped after processing the control information it contains. The Next Ethertype field is always the last two octets of the control header. The position of the Next Ethertype field in the frame shall be determined using the SSLength field in order to ensure forward compatibility.

If the Next Ethertype field is non-zero, then the frame is an *encapsulating* control frame. An encapsulated data frame is an encapsulating control frame with any Next Ethertype field not matching 0000_{16} or $886C_{16}$. G.989.2 receivers shall be capable of removing at least one encapsulating Short Format Link Control Frame header from any received encapsulated data frame. Future versions of this Recommendation may require the processing of multiple headers on encapsulated data frames, such as might occur if a Rate Request Control Frame were inserted into (i.e. piggybacked-on) a regular data frame with a LARQ header. When Next Ethertype is restricted by the specification to the value 0000_{16} for a specific Link Layer control frame SSType or LSType, then encapsulation of data frames is not allowed when using that Link Layer control frame type. The only Link Layer frame type which supports encapsulation of data frames is the LARQ frame.

If the SSType is not understood by the receiver (a fact possibly announced via future CSA options), then the frame shall be ignored. Note that all nodes are required to understand the LARQ SSTYPE (although they are not required to implement LARQ). Because the LARQ SSTYPE is the only SSTYPE allowed to encapsulate data payloads, G.989.2 transmissions will never result in dropped payloads due to unfamiliar SSTYPE values.

The header and trailer for standard Ethernet frames are shaded with gray, in order to highlight the formats of the control information frames.

6.2.2 Long format

Table 6/G.989.2 – Long Format Link Protocol Frame

Field	Length	Explanation
DA	6 octets	Ethernet Destination Address
SA	6 octets	Ethernet Source Address
Ethertype	2 octets	$0886C_{16}$
LSType	2 octets	32768 Reserved for future use by the ITU-T 32769 Vendor-specific long-format 32770 to 65535 Reserved for future use by the ITU-T
LSLength	2 octets	Number of control information octets, starting with the SSVersion field (or the first octet following SSLength if it is not defined as SSVersion) and ending with the second (last) octet of the Next Ethertype field. Min is 2 and max is 65535.
LSVersion	1 octet	Version number of the following protocol information
Data	LSLength – 3 octets	LSType protocol dependent data
Next Ethertype	2 octets	Ethertype/length of next layer protocol; 0 if none
Pad	42-0 octets	Pad to minimum size if needed
FCS	4 octets	Frame Check Sequence

An LSVersion, similar to SSVersion, is used for Long Format subtypes. A Next Ethertype field is required for all Long Format subtypes. If Long Format subtypes (LSType values) are not understood by the receiver (a fact possibly announced via future CSA options), then the frame shall be ignored. Processing requirements with respect to forwards compatibility, dropping of unknown frame types with Next Ethertype = 0000_{16}, and removal of Long Format headers with Next_Ethertype $\neq 0000_{16}$, are identical to those for Short Format Control Frame headers.

6.2.3 Order of transmission

The transmission order of the frame fields is from the top to the bottom of each Table.

Within a field, the most-significant octet of the field is the first octet of the field to be transmitted, with the LSB of each octet transmitted first. Subsequent octets within a field are transmitted in decreasing order of significance.

When subfields are pictured in any table, the ordering shown is decreasing significance from the top to the bottom of the table.

6.3 Rate Negotiation control function

The Rate Negotiation function in a destination station uses Rate Request Control Frames (RRCF) to provide information to a source station as to the payload encoding that the source station should use to encode future frames sent to this destination, and to generate test frames to assist a receiver in selecting the most appropriate band to use.

NOTE – The policy which the destination station uses to select the desired payload encoding, and the policy it uses to decide when to transmit Rate Request Control Frames, are outside the scope of this Recommendation. In general, the usable encoding is a function of the channel quality between source and destination; this generally differs between each pair of stations depending on the wiring topology and specific channel impairments.

6.3.1 Rate Request Control Frame Format (RRCF)

The RRCF specifies a maximum constellation (bits per symbol) that the receiver (ReqDA) wishes to be used in a given mode, and indicates which modes are supported. The payload modulation is defined in 5.2. (See Table 7.)

Table 7/G.989.2 – Rate Request Control Frame definition

Field	Length	Meaning
DA	6 octets	Ethernet Destination Address
SA	6 octets	Ethernet Source Address
Ethertype	2 octets	$886C_{16}$
SSType	1 octet	= 1
SSLength	1 octet	Number of control information octets, starting with the SSVersion field and ending with the second (last) octet of the Next Ethertype field. The minimum value of SSLength is 8 for SSVersion 0.
SSVersion	1 octet	= 0
OpCode	1 octet	Operation code for this control message. See Table 9 for definitions.
ModeNum	1 octet	The mode number is specified in this control field. Each mode has a two octet descriptor. – Mode 1 indicates 4D symbol mapping; – Mode 2 indicates 2D symbol mapping. G.989.2 stations shall ignore mode values other than 1 and 2.
NumAddr	1 octet	Number of addresses specified in the payload of this control message. NumAddr may be zero. The SA in the Ethernet header is always used, and is referred to in the following clauses as RefAddr0.

Table 7/G.989.2 – Rate Request Control Frame definition

Field	Length	Meaning
Mode1_PE	1 octet	4D symbol mapping, 7 MHz carrier: The PE value that should be used to send data when the 4D symbol mapping mode is selected. Values 1 through 7, inclusive, are the only ones valid.
Mode1_rank	1 octet	The rank order of the ReqDAs' preference for this mode. 1 is highest preference, and the other modes are assigned successively larger rank values; no two modes shall have the same rank.
Mode2_PE	1 octet	Optional, only present if ModeNum = 2. 2D symbol mapping, 7 MHz carrier: If included, this field is the PE value that should be used to send data when the 2D symbol mapping mode is selected. (0, 9..15) are the only valid values.
Mode2_rank	1 octet	Optional, only present if ModeNum = 2. Rank order of ReqDAs' preference for this mode
RefAddr1	6 octets	Optional. Present if NumAddr ≥ 1. The second MAC Address for which the rates are being specified; only broadcast and multicast address types are allowed.
RefAddr2	6 octets	Optional. Present if NumAddr ≥ 2. The third MAC Address for which the rates are being specified; only broadcast and multicast address types are allowed.
• • •		[additional instances of RefAddr, until the number of RefAddr fields equals NumAddr]
Next Ethertype	2 octets	= 0
Pad		To reach minFrameSize if required
FCS	4 octets	Frame Check Sequence

G.989.2 stations shall ignore mode values beyond ModeNum = 2. If a receiver does not specify a mode in an RRCF, or specifies a PE of 0 for a mode, then transmitters shall not use that mode. Modes may only be unspecified if no other mode information follows.

The ModeNum and NumAddr fields are placed next to each other so that all the fixed fields can be referenced at known offsets in the frame.

Table 8 describes the assigned values that may appear in the mode description entries in the Rate Request Control Frame.

Table 8/G.989.2 – PE values for Rate Request Control Frames

PE	Data rate	Meaning
0	N/A	Non-standard
1	4 Mbit/s	4D symbol mapping, QAM, 2 bits per symbol
2	6 Mbit/s	4D symbol mapping, QAM, 3 bits per symbol
3	8 Mbit/s	4D symbol mapping, QAM, 4 bits per symbol
4	10 Mbit/s	4D symbol mapping, QAM, 5 bits per symbol
5	12 Mbit/s	4D symbol mapping, QAM, 6 bits per symbol
6	14 Mbit/s	4D symbol mapping, QAM, 7 bits per symbol
7	16 Mbit/s	4D symbol mapping, QAM, 8 bits per symbol
8	N/A	Non-standard
9	8 Mbit/s	2D symbol mapping, QAM, 2 bits per symbol
10	12 Mbit/s	2D symbol mapping, QAM, 3 bits per symbol
11	16 Mbit/s	2D symbol mapping, QAM, 4 bits per symbol
12	20 Mbit/s	2D symbol mapping, QAM, 5 bits per symbol
13	24 Mbit/s	2D symbol mapping, QAM, 6 bits per symbol
14	28 Mbit/s	2D symbol mapping, QAM, 7 bits per symbol
15	32 Mbit/s	2D symbol mapping, QAM, 8 bits per symbol

Table 9 describes the values that may appear in the OpCode entry in the Rate Request Control Frame.

Table 9/G.989.2 – OpCode values for Rate Request Control Frames

OpCode	Meaning
0	Rate Change Request
1	Rate Test Request
2	Rate Test Reply
3-255	Reserved

6.3.2 TableRate negotiation overview

The following terms are used in describing the rate negotiation procedures:

mode specification A Payload Encoding (PE) and Rank associated with a given mode. A **mode** is a single combination of symbol mapping and symbol rate. Two modes are defined in this Recommendation.

Logical channel, channel A flow of frames from a sender to one or more receivers on a single network segment, consisting of all the frames with a single combination of DA and SA.

Receiver A station that receives frames sent on a particular channel. If the destination is a unicast address there is at most one **receiver**. If the destination is a group address (including broadcast), there may be many **receivers**.

Receiver PE The preferred PE to be used on this channel, as determined by the **receiver**.

RRCF Rate Request Control Frame. Sent from the **receiver** to the **sender** to effect a change in PE.

RefAddr0 The SA in the Ethernet header of the RRCF frame. This is the DA of the **receiver** (for the channel), and is always used by the channel **sender** as the first RefAddr processed.

RefAddr1..RefAddr<n> Other addresses including broadcast and multicast addresses for which the **receiver** is indicating rate information to the **sender**. The channel **receiver's** station address (RefAddr0) should not be put in the list of additional RefAddrs.

NOTE – At least one RefAddr field is necessary to support rate negotiation for broadcast and multicast addresses since these cannot be used as the source address in the Ethernet header.

Sender The sending station for a channel, usually the station owning the source MAC address.

Sender PE The preferred PE associated with a channel, as noted by the **sender**.

6.3.2.1 Channels

Rate Negotiation is defined over simplex logical channels. A separate channel is defined for each combination of Ethernet DA and SA. There is no explicit channel setup procedure. A new channel is implicitly defined when a packet is received from a new SA or sent to a new DA. Each channel has a single **sender** but can have multiple **receivers**. **Receivers** operate independently.

6.3.2.2 Sending RRCFs

Rate Request Control Frames (RRCFs) (all OpCodes) should be sent with a priority corresponding to Link Layer priority 7. RRCFs shall never be sent with a Link Layer priority of 6. RRCFs may be sent with a lower Link Layer priority, from the set [5, 4, 3, 0]. However, the Link Layer priority of an RRCF shall never be lower than the highest Link Layer priority received in the last 2 s from the station to which the RRCF is being sent. Rate Change Requests (OpCode = 0) shall always be sent with 4D symbol mapping QAM at 2 bits per symbol (PE = 1). Selection of the encoding for Rate Test Request frames and Rate Test Reply frames is described below.

6.3.2.3 Interval timer

Each station should maintain a timer with a period of 128 s. There should be no attempt to synchronize this timer between stations. The timer should not be modified by receipt or transmission of any frames. The timer interval is used when determining which nodes have been actively sending to multicast and broadcast addresses (see 6.3.3.2) and when sending reminder RRCFs in reference to multicast and broadcast addresses (see 6.3.4.1).

6.3.3 Sender operation

6.3.3.1 Sender – Transmit Data Frame

The sender shall maintain a table of logical channel state information. Logical channel state information includes the node type (e.g. G.989.2 or unknown), the sender PE and the receiver PE for each mode for which this information has been specified. When sending a frame, the sender shall determine the sender PE from the table. If necessary, a new channel entry may be made in the table with a default sender PE of PE = 1.

6.3.3.2 Sender – Receive Rate Change Request (RRCF OpCode 0)

For each of the RefAddrs in the RRCF, starting with RefAddr0, the SA of the RRCF frame, the station shall update the sender PE according to the mode specification in the RRCF. If no logical channel state information exists for RefAddr0, the station should create a new logical channel state entry and initialize the sender PE according to the mode specification in the RRCF. If no logical channel state information exists for additional RefAddrs, the station may either ignore those addresses or create new logical channel state entries and initialize the sender PE according to the mode specification in the RRCF.

For multicast addresses and the broadcast address, senders should use a rate that is receivable by all nodes actively listening to that address. Sender stations may enforce a minimum PE which they will use to transmit to a given multicast channel, based on application-level information about QoS. It is desirable to send at the highest rate supported by the channel. Hence, if a RefAddr is a multicast address or the broadcast address, the sender should use the PE value which yields the highest raw bit-rate, but which is not greater than any of the band specifications provided by the nodes actively listening to that address. Active multicast listeners shall be defined as any stations that have, in either of the last two 128-s intervals, either:

- sent any frame to the multicast address; or
- sent a RRCF to this station with the multicast address listed in the RefAddr list.

Active broadcast listeners shall be defined as any stations that have, in either of the last two 128-s intervals, either:

- sent any frame to the broadcast address; or
- sent a RRCF to this station with the broadcast address listed in the RefAddr list.

6.3.3.3 Sender – Receive Rate Test Request frame (RRCF OpCode 1)

For each supported mode, the sender shall generate a Rate Test Reply frame (RRCF OpCode 2) to the requestor encoded using the specified payload encoding. The contents of the RRCF shall be the current logical channel state info.

Support for Rate Test Request frames is only required in stations that implement both modes 1 and 2. Stations which only implement mode 1 may discard received Rate Test Request frames and send no response.

6.3.3.4 Sender – Active G.989.2 nodes

An active G.989.2 node is any station from which a valid frame has been received in either of the last two 128-s intervals.

6.3.4 Receiver operation

6.3.4.1 Receiver – Receive a frame

Receivers shall employ means to limit the numbers of RRCFs generated. The procedures defined in this clause, or an alternative procedure which generates no more RRCFs than the procedure described here, shall be used.

Stations that are interested in receiving frames of a specific multicast address or of the broadcast address shall provide a mechanism to ensure that all sources of frames sent to that multicast address (or the broadcast address, as appropriate) are reminded of this node's desire to receive frames directed to that address at least once every 128 s (see 6.3.3.2).

6.3.4.2 Receiver – Send Rate Test Request frame (RRCF OpCode 1)

Periodically, but at a rate not to exceed once every 128 s (except as described below), a receiver may send a Rate Test Request frame to a sender to test if the channel can support a different band. The band encodings represent the encodings for which the receiver would like the sender to generate test frames. NumAddr shall be set to 0 in Rate Test Request frames.

Rate Test Request frames should be sent encoded at the current negotiated rate for the channel from the receiver to the sender.

Support for Rate Test Request frames is only required in stations that implement additional bands beyond Band 1. Stations that only implement Band 1 need not provide a mechanism for generating Rate Test Request frames.

6.3.4.3 Receiver – Receive Rate Test Reply frame (RRCF OpCode 2)

Upon receipt of a Rate Test Reply frame, the receiver should use the demodulation statistics for this frame, and any previously received Rate Test Reply frames using this encoding, to make a decision as to the channel's capability to support the tested band encoding. If the decision is that the channel is not capable of supporting the tested band encoding, the receiver shall not generate another Rate Test Request frame for at least 128 s. If the decision is that the channel is capable of supporting the tested band encoding, the receiver may repeat the test to collect more data, at a maximum rate of one RateTest Request frame every second, with a maximum of 16 additional tests. At this point, the receiver should generate a Rate Change Request to the sender specifying the new band encoding.

Support for Rate Test Reply frames is only required in stations that implement additional bands beyond Band 1. Stations that only implement Band 1 may silently discard received Rate Test Reply frames.

6.4 Link Integrity Function

The purpose of the Link Integrity Function is to provide a means for the station to determine whether it is able to receive frames from at least one other station on the network. In the absence of other traffic, a station periodically transmits a Link Integrity Control Frame (LICF) to the Broadcast MAC address, with the interval between such transmissions governed by the method described below.

All stations shall implement the following functions to ensure that, with high probability, within any 1 s interval there is either:

– at least one LICF sent to the Broadcast MAC address from this station; or
– at least one packet addressed to the Broadcast MAC address received from each of at least two other stations.

Additionally, all stations shall send at least one LICF every 64 s.

The method is described below:

– Stations should be capable of generating Link Integrity frames even in inactive or sleep mode. While in sleep or inactive mode, stations that do not want to be or cannot be awakened should not send Link Integrity frames.
– A Link Packet may be any broadcast frame received with a valid header FCS. It is recommended that only LICF frames be treated as G.989.2 Link Packets.
– Each station maintains a free-running timer with a period of 1 s. There should be no attempt to synchronize this timer between stations. The timer should not be modified by any link state transitions or by the reception of any frames. This timer is the source of the time-out event used in the link integrity state table below.

– Each station maintains a 6-bit FORCE_SEND counter that is initialized to a random value between 30 and 63. This initialization value may be selected once at node startup and used for each re-initialization of the FORCE_SEND counter, or a new random value may be selected for each re-initialization of the FORCE_SEND counter.

– Each station has a register (SA1) that can be set from the SA of a received Link Packet.

– An LICF should be sent with a priority corresponding to Link Layer priority 7.

– The PE for an LICF shall be determined by accessing the RRCF logical channel information for the broadcast channel. An exception to this criterion is if Link Integrity frames are not sent with the currently negotiated broadcast PE value, then they shall be sent with PE = 1. This allows stations in sleep or inactive mode to maintain active status on the network. Receipt of a Link Integrity frame with PE = 1 shall not cause a transmission of an RRCF.

– Each station shall send a Link Integrity Control Frame (LICF) with the format shown in Table 11, according to the state diagram in Figure 14.

NOTE – While in sleep or inactive mode, stations should perform Link Integrity and wake-up processing on all receive packets. No further processing of receive packets is necessary. Any relevant power-management processing will be done on LARQ and non-LARQ data frames and the understanding is that non-WoLAN (Wake-on-LAN) frames should be discarded.

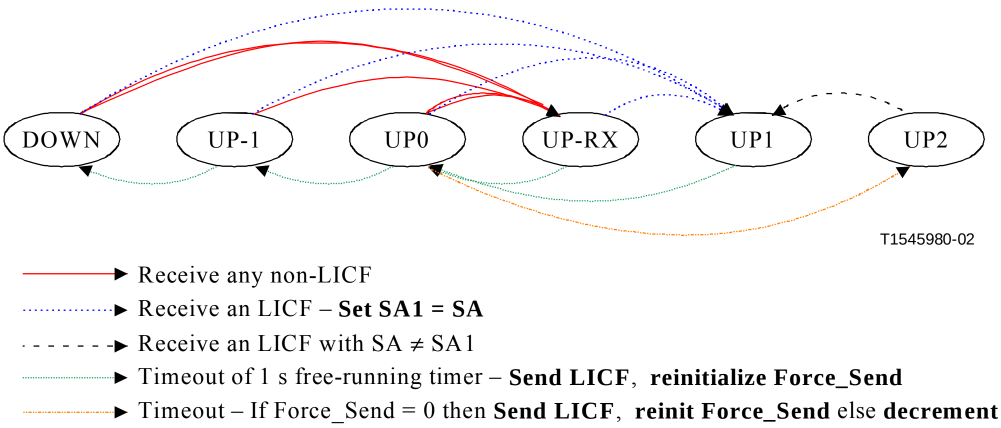

Figure 14/G.989.2 – Link Integrity state diagram

Figure 14 gives a pictorial view of the state transitions, with some minor loss of detail, including omission of events that do not cause state transitions (and have no associated actions), and the collapsing of multiple events into a single transition with a more complex description of the action.

Table 10 is a complete state table, with associated actions. The time-out event is the periodic expiration of a one-second free-running timer.

Table 10/G.989.2 – Link Integrity Finite State Machine

	DOWN	UP-1	UP0	UP-RX	UP1	UP2
Receive any non-LICF	UP-RX (none)	UP-RX (none)	UP-RX (none)	UP-RX (none)	UP1 (none)	UP2 (none)
Receive LICF with SA = SA1	UP1 Set SA1← SA	UP1 Set SA1← SA	UP1 Set SA1← SA	UP1 Set SA1← SA	UP1 (none)	UP2 (none)
Receive LICF with SA ≠ SA1	UP1 Set SA1← SA	UP1 Set SA1← SA	UP1 Set SA1← SA	UP1 Set SA1← SA	UP2 (none)	UP2 (none)
Time-out and Force_Send = 0	DOWN Send LICF [a)], reinit Force_Send	DOWN Send LICF [a)], reinit Force_Send	UP-1 Send LICF [a)], reinit Force_Send	UP0 Send LICF, reinit Force_Send	UP0 Send LICF, reinit Force_Send	UP0 Send LICF, reinit Force_Send
Time-out and Force_Send > 0	DOWN Send LICF [a)], reinit Force_Send	DOWN Send LICF [a)], reinit Force_Send	UP-1 Send LICF [a)], reinit Force_Send	UP0 Send LICF, reinit Force_Send	UP0 Send LICF, reinit Force_Send	UP0 decrement Force_Send

a) Devices which can transmit using more than one MAC source address (e.g. a bridge) should send a CSA request frame to the broadcast address instead of sending an LICF for the cases indicated in the table.

NOTE 1 – Initial State: DOWN, Force_Send initialized: 30 ≤ Force_Send ≤ 63.

NOTE 2 – Link Integrity Status is indicated when in any state but DOWN.

Table 11/G.989.2 – Link Integrity Short Frame

Field	Length	Meaning
DA	6 octets	Ethernet Destination Address = FF FF FF FF FF FF_{16}
SA	6 octets	Ethernet Source Address
Ethertype	2 octet	$886C_{16}$
SSType	1 octet	= 2
SSLength	1 octet	Number of control information octets, starting with the SSVersion field and ending with the second (last) octet of the Next Ethertype field. Minimum is 4 for SSVersion 0
SSVersion	1 octet	= 0
LI_pad	1 octet	Ignored on reception
Next Ethertype	2 octets	= 0
Pad	40 octets	Any value octet
FCS	4 octets	

6.5 Capability and Status Announcement

This clause defines procedures for network-wide negotiation, capability discovery and status announcement. It is based on periodic broadcast announcements, called Capabilities and Status Announcements (CSA) sent in CSA Control Frames (CSACFs). The defined status flags allow determination of a station's G.989.2 version, optional feature capability, and link-layer priority usage, as well as communication of network configuration commands.

The purpose of the protocol is to distribute to all stations the complete set of status flags in use on the network, so that stations can make operational decisions based on those flags with no further interaction.

Stations shall use the CSA Control Frame as described in Table 12 and the CSA Flag definitions in Table 13. Stations shall send a CSA Control Frame once per minute or when a change in the station's current status requires the announcement of new (or deleted) flags.

Table 12/G.989.2 – Capability and Status Announcement frame

Field	Length	Meaning
DA	6 octets	Ethernet Destination Address = FF FF FF FF FF FF_{16}
SA	6 octets	Ethernet Source Address, not necessarily corresponding to the MAC address to which the frame contents are applicable (see CSA_SA).
Ethertype	2 octet	$886C_{16}$
SStype	1 octet	= 3
SSLength	1 octet	Number of control information octets, starting with the SSVersion field and ending with the second (last) octet of the Next Ethertype field. Minimum is 32 for SSVersion 0.
SSVersion	1 octet	= 0
CSA_ID_Space	1 octet	Identifies the registration space of CSA_MFR_ID 0 Unspecified 1 JEDEC 2 PCI 3-255 Reserved for use by ITU-T
CSA_MFR_ID	2 octets	HW manufacturer ID: Identifies the manufacturer of the PHY controller implementation, as assigned by the entity identified in CSA_ID_Space. The purpose of this field plus the part number and revision is to identify specific PHY implementations. This is not a board or assembly-level identifier.
CSA_Part_No	2 octets	HW Manufacturer Part Number: The part number of the PHY controller chip
CSA_Rev	1 octet	HW Revision
CSA_Opcode	1 octet	0 Announce 1 Request
CSA_MTU	2 octets	Maximum size link-level PDU. This receiver accepts in octets; the default value is 1526 octets. 1526 is the minimum value that shall be advertised by a G.989.2 station.
CSA_SA	6 octets	MAC address of the station to which the capabilities and status are applicable
CSA_pad	2 octets	Reserved for version 0. Shall be sent as 0; ignored on reception. Creates field alignment to 32-bit WORD boundaries
CSA_CurrentTxSet	4 octets	Configuration flags, plus all current in-use status for this station. Flag definitions are specified in Table 13.

Table 12/G.989.2 – Capability and Status Announcement frame

Field	Length	Meaning
CSA_OldestTxSet	4 octets	A copy of the "oldest" TX flags for this stations, from the period ending at least one period (minute) earlier. Flag definitions are specified in Table 13.
CSA_CurrentRxSet	4 octets	The union of recent flags received from other stations. Flag definitions are specified in Table 13.
Next Ethertype	2 octets	= 0
Pad		Pad to reach minFrameSize if necessary
FCS	4 octets	

A station sending a CSA Control Frame announcing a status change shall send a second copy of the most recent CSACF a short interval after the first, since it is always possible to lose a frame due to temporary changes in the channel, impulse noise, etc. The interval should be randomly selected (not simply fixed), and chosen from the range 1 to 1000 μs, inclusive.

CSA Control Frames are sent with a priority corresponding to Link Layer priority 7.

CSA Control Frames are always sent to the broadcast address (FF FF FF FF FF FF_{16}).

The PE for a CSA control frame shall be determined by accessing the RRCF logical channel information for the broadcast channel.

A Request OpCode is defined to allow a station to quickly gather complete information about all stations. Upon receiving a CSA control frame with the Request OpCode, a station shall transmit a current CSA message after a delay of a short interval, using the same mechanism (and parameters) that delays the second copy of CSA announcements, described above.

6.5.1 CSA control frame

Table 12 defines the format of a Capabilities and Status Announcement control frame. The first three fields beyond the Ethernet header comprise the standard header for short format control frames.

6.5.2 Status, Configuration, Option and Priority Flags

Table 13 contains the flags used for CSA_CurrentTxSet, CSA_OldestTxSet, and CSA_CurrentRxSet in Capabilities and Status Announcement control frames.

Table 13/G.989.2 – CSA Flags

Octet	Field	Length	Description
Flags0	TxPriority7	1	Station is (was) transmitting frames with LL priority 7. (always set)
	TxPriority6	1	Station is (was) transmitting frames with LL priority 6.
	TxPriority5	1	Station is (was) transmitting frames with LL priority 5.
	TxPriority4	1	Station is (was) transmitting frames with LL priority 4.
	TxPriority3	1	Station is (was) transmitting frames with LL priority 3.
	TxPriority2	1	Station is (was) transmitting frames with LL priority 2.
	TxPriority1	1	Station is (was) transmitting frames with LL priority 1.
	TxPriority0	1	Station is (was) transmitting frames with LL priority 0. (always set)

Table 13/G.989.2 – CSA Flags

Octet	Field	Length	Description
Flags1	Reserved	7	Shall be sent as 0 and ignored by receiving stations
	Supports 2D symbol mapping	1	This station supports 2D symbol mapping payload encodings.
	Reserved	8	Shall be sent as 0 and ignored by G.989.2 stations when received
Flags2	G.989.2	1	G.989.2 mode
Flags3	Reserved	4	Shall be sent as 0 and ignored by receiving stations
	G.989.2 version	3	000 Reserved 001 Reserved 010 G.989.2 011-111 Reserved

Thirty-two bits of flags are supported for announcing status and configuration information. The flags are divided into three basic groups: mode selection flags, supported options, and in-use TX link layer priority announcements. These flags are added to the global state as soon as announced, and removed when no longer announced by any station, either through explicit deletion or by timing them out. An in-use TX link layer priority will be announced for a period of one to two minutes after the last frame actually sent with the priority, until the aging mechanism causes it to be deleted from CurrentTxSet.

The **default set** of status flags, used to initialize the **NewTxSet** (defined below), is defined to be the priorities 0 and 7, the station's G.989.2 version, and any supported options.

6.5.3 Capability and Status Announcement – Terms and parameters

6.5.3.1 Capabilities and Status Period (CS period)

Non-persistent status information expires after a basic time interval of one minute. Each station has a repeating timer set to this interval. The timers in different stations are not synchronized, and synchronization should be avoided, in general. The description below refers to the time between one expiration of this timer and the next as a "period". The "current" period refers to the time since the most recent expiration of the timer.

A CSA frame is sent at the end of each interval.

6.5.3.2 Variables

The descriptions in this clause make use of the following variables:

DeleteSet A computed value used to detect newly removed status information.

NewRxFlags, ReallyNewRxFlags Computed values used to detect new status flags.

CSP_Timer A free-running timer with a period of 60 s.

RetransmitTimer A one-shot timer, set to a random interval in the range 1 ms to 1000 ms, inclusive, after sending a CSA in which CSA_CurrentTxSet and CSA_OldestTxSet are different, or when a CSA is received with the CSA_Opcode set to 1 (Request). This timer is cancelled if a second CSA is sent as a result of the CSP_Timer expiring.

6.5.4 Status and Priority Set State Parameters

Each station maintains five basic sets of status and priority information. In addition, three more composite sets are defined as the union of two or more of the basic sets.

NewTxSet The set of flags announced during the current CS period, updated immediately when a new link layer priority is used or new volatile status is set. When the CSP_Timer expires, CurrentTxSet is given the value of NewTxSet, and NewTxSet is reset to the **default** set.

PreviousTxSet The set of flags that were announced during the previous CS period (the ending value of NewTxSet from the previous CS period).

OldestTxSet The set of flags rolled over from PreviousTxSet at the end of the previous CS period (the value of PreviousTxSet from the previous CS period). Flags that are present in OldestTxSet and missing from PreviousTxSet were not actively used or detected (by the sender) for an entire CS period, and will be deleted. This set is sent in CSA frames as CSA_OldestTxSet.

NewRxSet The union of all CSA_CurrentTxSet flags received in CSAs from other stations during the current CS period. This is rolled over into PreviousRxSet at the expiration of the CSP_Timer, then reset to the empty set (0).

A volatile status flag (one of the priority flags) in this set may subsequently be deleted if the only station previously announcing that flag stops using it. The deletion from that station's CurrentTxSet is noted by the difference from its OldestTxSet. The fact that it was the only sender is noted by the absence of the flag in that station's CurrentRxSet, indicating that it has received the flag from no other stations.

If deleted from NewRxSet, a flag shall also be deleted from PreviousRxSet.

PreviousRxSet The set of announced flags received during the previous CS period (the ending value of NewRxSet from the previous CS period). A flag may be deleted from this set, as described under NewRxSet above.

CurrentTxSet The set of flags that were announced during the previous CS period plus any new status and priority flags (or changed configuration/options flags) used during the current CS period, i.e. the union of PreviousTxSet and NewTxSet. This set is sent in CSA frames as CSA_CurrentTxSet.

CurrentRxSet The union of NewRxSet and PreviousRxSet. This set is sent in CSA frames as CSA_CurrentRxSet.

CurrentInUseSet The union of CurrentTxSet and CurrentRxSet. This set is used to determine the operational mode of the station and to modify the mapping between the LL priority of the frame and the actual PHY priority usage.

6.5.5 Capabilities and Status Announcement Protocol Operation

6.5.5.1 New transmit frame – Priority detection

The CSA Protocol does not directly process transmit frames. When the LARQ protocol is in use (Full Profile stations), CSA looks at the **LL priority** of the frame as it would normally be sent to the driver.

a) If the **LL priority** is not already in NewTxSet, add it to NewTxSet.

b) If the **LL priority** was not already in NewTxSet and it is not in PreviousTxSet, then send a new CSA control frame with the CSA_Opcode set to 0 (Announce), and start the RetransmitTimer. If the timer was already running, then cancel and restart it. Update the current PHYpriority mapping function for the driver.

6.5.5.2 Receive CSA Control Frame

The receiver may want to save a copy of some or all of the most recent CSA from each other station as a simple way of tracking other station's capabilities and status.

a) Record (optionally) the status and options flags from the CSA_CurrentTxSet in a table indexed by the CSA_SA address. The options flags are used to select use of optional functions between pairs of stations that implement the same options.

b) If the CSA_Opcode in the frame is 1 (Request), then start the RetransmitTimer. If the timer is already running it is recommended that it be left running, although this is not required and cancellation followed by restart is allowed.

c) If CSA_CurrentTxSet has a flag not already in NewRxSet, then add the flag to NewRxSet, and check to determine if this flag is not present in the PreviousRxSet. The corresponding Boolean expressions are as follows:

$$\text{NewRxFlags} = \text{CSA\_CurrentTxSet} \wedge \overline{\text{NewRxSet}} \quad (6\text{-}1)$$

$$\text{NewRxSet} = \text{NewRxSet} \vee \text{NewRxFlags} \quad (6\text{-}2)$$

$$\text{ReallyNewFlags} = \text{NewRxFlags} \wedge \overline{\text{PreviousRxSet} \vee \text{CurrentRxSet}} \quad (6\text{-}3)$$

d) Compare CSA_OldestTxSet with CSA_CurrentTxSet. If a flag has been deleted, and if that flag is also missing from CSA_CurrentRxSet, then delete the flag from NewRxSet, and PreviousRxSet. The corresponding Boolean expressions are as follows:

$$\text{DeleteSet} = \text{CSA\_OldestTxSet} \wedge \overline{\text{CSA\_CurrentTxSet}} \wedge \overline{\text{CSA\_CurrentRxSet}} \quad (6\text{-}4)$$

$$\text{NewRxSet} = \text{NewRxSet} \wedge \overline{\text{DeleteSet}} \quad (6\text{-}5)$$

$$\text{PreviousRxSet} = \text{PreviousRxSet} \wedge \overline{\text{DeleteSet}} \quad (6\text{-}6)$$

$$\text{CurrentRxSet} = \text{NewRxSet} \vee \text{PreviousRxSet} \quad (6\text{-}7)$$

e) If either ReallyNewFlags or DeleteSet are non-zero, then update the network mode and priority mapping, as necessary.

NOTE – The symbols "$\vee$" and "$\wedge$" specify bit-wise logical "OR" and "AND" operations, respectively.

6.5.5.3 CSP_Timer time-out

When a CSP_Timer time-out occurs a new CS period has begun. Roll over the various status sets, re-compute the composite sets, and send a CSA. Set the RetransmitTimer, if needed.

– OldInUseSet = CurrentInUseSet

– Move NewRxSet to PreviousRxSet.

– Set NewRxSet to 0 (empty set).

– Move PreviousTxSet to OldestTxSet.

– Move NewTxSet to PreviousTxSet.

– Set NewTxSet to the default set, consisting of this station's G.989.2 version, current configuration flags if any (normally none), currently supported options, and the default priority set {0,7}.

– Update CurrentTxSet, CurrentRxSet, and CurrentInUseSet.

$$\text{CurrentRxSet} = \text{NewRxSet} \vee \text{PreviousRxSet} \tag{6-8}$$

$$\text{CurrentTxSet} = \text{NewTxSet} \vee \text{PreviousTxSet} \tag{6-9}$$

$$\text{CurrentInUseSet} = \text{CurrentRxSet} \vee \text{CurrentTxSet} \tag{6-10}$$

– Send a CSA frame with the CSA_Opcode set to 0 (Announce), including the updated flags.

– If CSA_CurrentTxSet and CSA_OldestTxSet in the CSA frame just sent were different, start the RetransmitTimer. If the timer was previously running, then cancel it and restart it.

– If one or more status flags have been deleted, then recompute the network operating mode and/or priority mapping function due to changed status flags. The mode/mapping recomputation should be performed if CurrentInUseSet is not equal to OldInUseSet.

6.5.5.4 Retransmit time-out

If the RetransmitTimer expires, send a current CSA frame for this station with the CSA_Opcode set to 0 (Announce). The timer is not restarted.

6.5.6 Network mode selection based on CurrentInUseSet

The mode selection flags of the CSA protocol are intended to facilitate compatibility with proprietary modes, and for future versions of ITU-T Rec. G.989.3.

6.5.7 Priorities

There is a cost of slightly lower maximum attainable bandwidth associated with the lower PHY-priorities in the G.989.2 MAC protocol if a default mapping scheme of link-layer to PHY-layer priorities is employed. This cost can become especially burdensome if only lower-priority traffic is being carried on the network. Therefore, the CSA protocol includes procedures for the re-mapping of lower link-layer priorities to higher PHY PHY-layer priorities when no station on the network is sending traffic marked for those higher priorities.

The choice of Physical Layer (PHY) priority for a given frame is based on its assigned Link Layer (LL) priority. The default mapping from LL priority to PHY priority is specified in 6.5.7.3. The LL priority of a frame at the sender must be conveyed to the receiving station in order to allow proper recovery of link layer protocol at the receiver. This requires either a fixed, one-to-one, mapping of LL to PHY priorities, or some mechanism for carrying the LL priority within each frame. The LARQ protocol, defined in 6.6, carries the assigned LL priority from a sending station to a receiving station, providing the required mechanism, and thereby creating the opportunity to apply non-default LL to PHY priority mappings, which in turn, allows for higher maximum attainable bandwidth.

NOTE – A station may optionally use an IEEE 802.1Q header to convey the LL priority. However, stations that do not support IEEE 802.1Q headers are unlikely to comprehend frames that include such headers. Such frames may not be properly received by all intended recipients unless a determination is made that all receivers of the frame support the use of IEEE 802.1Q headers.

6.5.7.1 Transmit frames – Choice of physical priority

When the assignment of a Physical layer priority to the frame occurs, any changes to the PHY priority re-mapping function due to the use of a new priority should already have been made. The station should use the re-mapped PHY priority to transmit the frame (including placing this value in the Frame Control Header) unless the frame has no LARQ header, in which case the default LL-to-PHY mapping shall be used.

6.5.7.2 Received frame priorities

The LL priority of received frames indicated to higher-layer protocols (before any reassignment due to a LARQ or IEEE 802.1Q header) shall be determined using the default PHY-to-LL priority map, except that Minimal Profile stations should indicate the priority of all frames as LL 0. The mechanism that guarantees correct LL priority for received frames is the restoration of LL priority from the LARQ (or optionally, IEEE 802.1Q) header. LARQ header processing is always performed after the default LL priority has been assigned in the receive path.

6.5.7.3 Default Link Layer to Physical Layer Map

The IEEE 802.1p specification (incorporated in ISO/IEC 15802-3:1998) places the default (unassigned/best-effort) priority above both priorities 1 and 2, when an 8-level priority system is in use. Therefore, Link Layer priority 0 will be mapped above both LL 1 and LL 2 for default Physical Layer priority assignment. IEEE 802.1p designates priority level 7 for Network Control and priority level 6 for traffic requiring latency of <10 µs (typically characterized as voice-like traffic). However, on G.989.3 networks, PHY priority level 7 shall be reserved for traffic requiring latency of less than 10 µs, and Network Control traffic is redirected to PHY priority level 6. Link Layer priority 5 shall be reserved for traffic requiring latency of less than 100 µs. Thus, the default mapping for LL to PHY priorities includes the swapping of priorities 6 and 7.

For transmitted frames, the set of LL priorities [0, 1, 2, 3, 4, 5, 6, 7] are by default mapped in order to the following set of PHY priorities [2, 0, 1, 3, 4, 5, 7, 6].

For received frames, PHY priorities [0, 1, 2, 3, 4, 5, 6, 7] are, by default, mapped to LL priorities [1, 2, 0, 3, 4, 5, 7, 6].

6.5.8 Priority mapping and LARQ

The PHY priority re-mapping is performed below LARQ in the protocol stack, and is not applied to the priority field in the LARQ (or optionally, IEEE 802.1Q) header. PHY priority re-mapping shall not be performed on data frames (those that are not Link Control frames) unless a LARQ (or optionally, IEEE 802.1Q) header has been added with the original LL priority. PHY priority re-mapping is performed on Link Control Frames.

6.5.9 Priority re-mapping based on CurrentInUseSet

Without priority mapping, a station would pass the original **LL priority** to the transmitter, where that value would be used to select the associated PHY priority from the default map. With priority re-mapping, the **default-assigned PHY priorities** are increased to make use of higher PHY priorities that would otherwise be unused. The re-mapping function is simple. For each PHY priority **P** that corresponds to an in-use LL priority, the new priority **P'** to use is that priority increased by the number of higher unused priorities. For example, if [1, 3, 4, 7] are in use, then priority 4 will be increased by 2 to 6, since there are two higher unused priorities (5, 6). Tables 14 and 15 contain a few more examples that should make this clear (including the default LL-to-PHY translation). The columns in the tables represent **LL priorities** before mapping. The left-hand section shows some sets of in-use priorities, with the right-hand section showing the new PHY priority that the driver should use in each case.

Table 14/G.989.2 – Default LL to PHY TX priority mapping

									TX LL priority							
									0	1	2	3	4	5	6	7
CurrentInUse Priorities (any)									**Default TX PHY Priorities**							
a	n	y	t	x	s	e	t		2	0	1	3	4	5	7	6

Table 15/G.989.2 – Direct LL to PHY TX priority re-mapping

									TX LL priority							
									0	1	2	3	4	5	6	7
CurrentInUse Priorities (LL)									**Remapped TX PHY Priorities**							
0							7		6	5	5	6	6	6	7	7
0						6	7		5	4	4	5	5	5	7	6
0	1			4			7		5	4	4	5	6	6	7	7
0			3		5	6	7		3	2	2	4	4	5	7	6

The shaded entries show mappings that no sender should be using. However, if there is any possibility of an implementation sending with an out-of-date mapping, or sending a priority that has not been included in the mapping, then it should always use the priority of the next lower valid mapping.

Here is one example in detail. If the CurrentInUse are [0, 1, 4, 7], then the corresponding set of in-use PHY priorities is [2, 0, 4, 6]. Then increase each by the number of missing higher priorities: 2→5, 0→4, 4→6 and 6→7. Just to be safe, the any unused PHY priorities are also remapped to the new value of the next lower in-use priority, giving: 1→4, 3→5, 5→6, 7→7.

So the in-use LL priorities [0, 1, 4, 7] result in transmitting PHY priorities [5, 4, 6, 7]. A complete map for all the LL priorities adds the remaining remapped values for the default priorities corresponding to the unused LL priorities: LL [0, 1, 2, 3, 4, 5, 6, 7] gives PHY [5, 4, 4, 5, 6, 7, 7].

6.6 LARQ: Limited Automatic Repeat Request Protocol

Limited Automatic Repeat reQuest (LARQ) is a protocol that reduces the effective error rate when frame errors occur. Its primary distinction from similar, sequence number-based protocols is that it does not guarantee reliable delivery of every frame, but instead conceals errors in the physical layer through fast retransmission of frames. The goal is to significantly enhance the usability of networks that may, at least occasionally, have frame error rates (FER) of 1 in 10^{-2} or worse. Protocols such as TCP are known to perform poorly when FER gets high enough, and other applications, such as multimedia over streaming transport layers, are also susceptible to poor performance due to high FER conditions.

The protocol provides a negative acknowledgment (NACK) mechanism for receivers to request the retransmission of frames that were missed or received with errors. There is no positive acknowledgment mechanism. There is no explicit connection set-up or tear-down mechanism. A reminder mechanism gives receivers a second chance to detect missing frames when relatively long gaps (in time) occur between frames.

LARQ functions as an adaptation layer between the Link Layer (layer 2) and the IP network layer (layer 3).

Stations implement LARQ per "LARQ channel", where a LARQ channel is identified by the triple {source address, destination address, priority}. Stations may enable or disable LARQ processing on a channel dynamically, based on information about network frame error rates. However, it is recommended that LARQ be left enabled at all times, since the per-packet processing overhead is quite low, and the complexity associated with enabling and disabling the protocol (including determination of appropriate parameters) probably outweighs any likely performance gains.

Stations should implement LARQ, and if they do so, they shall use the specified control frame formats and should use the recommended procedures defined below.

Stations not adding LARQ (or optionally, IEEE 802.1Q) headers shall not re-map PHY priorities, and shall treat all received traffic as "best effort", that is, all traffic shall be assigned to Link Layer Priority 0.

Stations may choose to add LARQ headers on transmitted frames with the LARQ_NoRtx flag set to 1. This flag indicates that the station does not retransmit frames for this channel, but adding the LARQ header allows the station to use PHY priority re-mapping since the LL priority of successfully received frames will be restored from the LARQ header.

All stations shall be capable of removing LARQ headers from received frames (de-encapsulating the original payloads). Furthermore, if the implementation supports multiple LL priorities in its receive protocol processing, then it shall restore the LL priority from the LARQ header, if one is present. If a station does not implement LARQ, then it shall drop LARQ control frames and it shall discard frames marked as retransmissions in the LARQ header.

6.6.1 Frame formats – Encapsulating headers

The text below uses the terms "insert" and "remove" when discussing LARQ headers. The formal definition of the LARQ frame format provides a Next Ethertype field that contains the original frame's Ethertype value. In practice, it will generally be the case that LARQ frames will be created by inserting the 8 octets starting with the Ethertype $886C_{16}$ into the original frame between the Ethernet header's source address and the original frame's Ethertype. The original frame's Ethertype becomes relabeled as the Next Ethertype field of the final frame.

NOTE – The LARQ header carries LLC priority across the network. The use of IEEE 802.1Q headers is not required for this function, and G.989.2 drivers are not required to support the use of IEEE 802.1Q headers for conveying priority.

Table 16/G.989.2 – LARQ Reminder Control Frame

Field	Length	Meaning
DA	6 octets	Ethernet Destination Address
SA	6 octets	Ethernet Source Address
Ethertype	2 octets	$886C_{16}$
SSType	1 octet	= 4
SSLength	1 octet	Number of control information octets, starting with the SSVersion field and ending with the second (last) octet of the Next Ethertype field. SSLength is 6 for SSVersion 0.
SSVersion	1 octet	= 0
LARQ_hdr data	3 octets	LARQ Control Header data with LARQ_Ctl bit = 1, LARQ_NACK = 0
Next Ethertype	2 octets	= 0
Pad	38 octets	
FCS	4 octets	Frame Check Sequence

Table 17/G.989.2 – LARQ NACK Control Frame

Field	Length	Meaning
DA	6 octets	Ethernet Destination Address
SA	6 octets	Ethernet Source Address
Ethertype	2 octets	$886C_{16}$
SSType	1 octet	= 4
SSLength	1 octet	Number of control information octets, starting with the SSVersion field and ending with the second (last) octet of the Next Ethertype field. SSLength is 12 for Nack frames with SSVersion 0.
SSVersion	1 octet	= 0
LARQ_hdr data	3 octets	LARQ Control Header data with LARQ_Ctl bit = 1, LARQ_NACK = 1..7
NACK_DA	6 octets	Original Destination Address
Next Ethertype	2 octets	= 0
Pad	32 octets	
FCS	4 octets	Frame Check Sequence

Table 18/G.989.2 – LARQ Encapsulation Frame

Field	Length	Meaning
DA	6 octets	Destination Address (from original Ethernet PDU)
SA	6 octets	Source Address (from original Ethernet PDU)
Ethertype	2 octets	$886C_{16}$
SStype	1 octet	= 4
SSLength	1 octet	Number of control information octets, starting with the SSVersion field and ending with the second (last) octet of the Next Ethertype field. SSLength is 6 for SSVersion 0 = 6
SSVersion	1 octet	= 0
LARQ_hdr data	3 octets	LARQ Encapsulation header data (with LARQ_Ctl bit = 0)
Next Ethertype	2 octets	From original Ethernet PDU
Payload	Min 46 octets	From original Ethernet PDU payload
FCS	4 octets	Frame Check Sequence

Table 19/G.989.2 – LARQ_Encapsulation Header Data

Octet	Field	Length	Meaning
Flags0	LARQ_Mult	1 bit	Multiple Retransmission Flag. 0 in the original transmission of a data frame. For retransmitted frames (LARQ_Rtx = 1), set to the value of LARQ_Mult in the NACK frame that caused the retransmission. This flag can be used by receivers to measure the round-trip times associated with the miss/nack/receive-rtx process.
	LARQ_Rtx	1 bit	0 for first transmission of a frame, 1 if frame is retransmitted. Stations not implementing LARQ shall drop any data frame if this bit is 1.
	LARQ_NewSeq	1 bit	1 if the sequence number space for the channel has been reset, and older sequence numbers should not be nacked; 0 otherwise
	LARQ_NoRtx	1 bit	0 if implementation supports retransmission, 1 if only priority is meaningful. May be used on a per-channel basis
	LARQ_Ctl	1 bit	"0" when in Encapsulation Format
	Priority	3 bits	Link Layer Priority of this frame
Flags1_Seq0	Reserved	4 bits	Reserved, shall be 0
	LARQ_seq_high	4 bits	High 4 bits of Sequence number
Seq1	LARQ_seq_low	8 bits	Low 8 bits of Sequence number

The exact application of the LARQ_Rtx, LARQ_NewSeq and LARQ_NoRtx bits requires further explanation as found in Table 20.

Table 20/G.989.2 – LARQ_Rtx, LARQ_NewSeq, LARQ_NoRtx bits

LARQ_ Rtx	LARQ_ NewSeq	LARQ_ NoRtx	Interpretation
0	0	0	Normal transmission on an active channel. This combination is used for the first transmission of a frame on an active LARQ channel. The receiver of this frame should send NACKs for earlier sequence numbers that are determined to be missing when this frame is received, or for this frame, if this frame has a CRC error but the LARQ header appears to be in sequence for the channel.
0	0	1	Used for the first transmission of a frame which will not be retransmitted in response to a NACK. The sender should use this combination when it does not save the frame for retransmission in response to receiving a NACK. If a receiver is keeping state, then it should send this frame up when it has either received frames for all previous sequence numbers, or given up attempts to receive frames for all previous sequence numbers.

Table 20/G.989.2 – LARQ_Rtx, LARQ_NewSeq, LARQ_NoRtx bits

LARQ_ Rtx	LARQ_ NewSeq	LARQ_ NoRtx	Interpretation
0	1	0	Used for the first transmission of a frame with a new sequence number space. The sender uses this combination when there are no saved frames for the channel, excepting this frame. The receiver should send all frames for this channel up to the next layer, since there is no longer the possibility of receiving any frames with previous sequence numbers. The receiver of this frame should send a NACK for this frame, if this frame has a CRC error but the LARQ header appears to be in sequence for the channel.
0	1	1	Used for the first transmission of a frame with a new sequence number space which will not be retransmitted in response to a NACK. The sender uses this combination when there are no saved frames for the channel. The receiver should send all frames for this channel up to the next layer, since there is no longer the possibility of receiving any frames with previous sequence numbers.
1	0	0	Retransmission of a frame for this channel. Sender uses this combination to send a frame which has been transmitted before, and for which a NACK will cause an additional retransmission. The receiver shall accept this frame if it is not a duplicate. If the receiver is not maintaining state for the channel, then this frame must be discarded because it would be impossible to determine the duplicate status for the frame. The receiver of this frame should send a NACK for this frame, if this frame has a CRC error but the LARQ header appears to be in sequence for the channel.
1	0	1	Retransmission of a frame for this channel. Sender uses this combination to send a frame which has been transmitted before, but has not been saved for retransmission in response to receiving a NACK. The receiver shall accept this frame if it is not a duplicate. If the receiver is not maintaining state for the channel, then this frame must be discarded because it would be impossible to determine the duplicate status for the frame.
1	1	0	Retransmission of a frame for this channel. The sender uses this combination when there are no older saved frames for the channel, excepting this frame. The receiver shall accept this frame if it is not a duplicate. If the receiver is not maintaining state for the channel, then this frame must be discarded because it would be impossible to determine the duplicate status for the frame. The receiver should send this frame and all older frames for this channel up to the next layer, since there is no longer the possibility of receiving any frames with previous sequence numbers. The receiver of this frame should send a NACK for this frame, if this frame has a CRC error but the LARQ header appears to be in sequence for the channel.

Table 20/G.989.2 – LARQ_Rtx, LARQ_NewSeq, LARQ_NoRtx bits

LARQ_ Rtx	LARQ_ NewSeq	LARQ_ NoRtx	Interpretation
1	1	1	Retransmission of a frame for this channel. The sender uses this combination when there are no older saved frames for the channel. The receiver shall accept this frame if it is not a duplicate. If the receiver is not maintaining state for the channel, then this frame must be discarded because it would be impossible to determine the duplicate status for the frame. The receiver should send this frame and all older frames for this channel up to the next layer, since there is no longer the possibility of receiving any frames with previous sequence numbers.

Table 21/G.989.2 – LARQ Control Header Data

Octet	Field	Length	Meaning
Flags0	LARQ_Mult	1 bit	Multiple Retransmission Flag. 0 in the first Nack sent for a given sequence number, 1 in all retransmitted Nacks.
	LARQ_NACK	3 bits	NACK Count If 0 in a LARQ Control Frame, then this is a Reminder.
	LARQ_Ctl	1 bit	Set to 1 for LARQ Control Header data format
	Priority	3 bits	Link Layer Priority of this frame
Flags1_Seq0	Reserved	4 bits	Reserved, shall be 0
	LARQ_seq_high	4 bits	High 4 bits of Sequence number
Seq1	LARQ_seq_low	8 bits	Low 8 bits of Sequence number

6.6.2 LARQ – Overview

The following terms are used in describing the LARQ procedures:

control frame A frame generated by a LARQ protocol module that contains only a LARQ protocol header as its payload.

Current sequence number The most recently received new sequence number for a channel.

Data frame Any standard Ethernet frame from higher (than LARQ) protocol layers. A LARQ-enabled station encapsulates the original payload of an Ethernet frame by inserting a LARQ header (short form control header with LARQ_hdr data) between the source address and the remainder of the frame before the frame is passed down to the driver for transmission on the network.

Forget timer An implementation-dependent mechanism to allow a receiver to reset the sequence number space of a channel when a received sequence number is not the next expected (**Current Sequence Number** + 1). One second is a suggested default value.

hold timer, lost timer	An implementation-dependent timing mechanism that limits the time a receiver will hold onto a received frame while waiting for a missing frame to be retransmitted. Conceptually, there is one such timer per missing sequence number. The timer interval is **Maximum Hold Interval**.
logical channel, channel	A flow of frames from a sender to one or more receivers on a single network segment consisting of all the frames with a single combination of destination address, source address, and link layer priority.
NACK, Nack, nack	An indication from a receiver to a sender requesting retransmission of one or more frames. Also, the action of providing such an indication. E.g. "to nack a sequence number" meaning to send a NACK indication.
NACK timer	An implementation-dependent timing mechanism used by a receiver to retransmit NACKs for missing sequence numbers. Conceptually, there is one such timer per missing sequence number per **logical channel**. The timer is reset each time a NACK is sent for a sequence number. The timer interval is **NACK Retransmission Interval**.
new	A new sequence number is one whose difference from the **current sequence number** for the channel, modulo the size of the sequence number space and considered as a signed integer, is greater than 0. In particular, the numbers (current + 1) through (current + 2047).
old	An old sequence number is one whose difference from the **current sequence number** for the channel, modulo the size of the sequence number space and considered as a signed integer, is less than or equal to 0. In particular, the numbers (current – 2048) through (current) are old. Note, however, that most of the old sequence numbers are also out of sequence.
out of sequence	Any sequence number that falls outside a reasonable range, old or new, of the current sequence number for a logical channel is considered out of sequence. It is recommended that plus or minus twice the value of MaximumSaveLimit (defined below) be used as the "reasonable range" when checking for out of sequence.
receiver	A station that receives frames sent on a particular channel. If the destination address is a unicast address, there is at most one receiver. If the destination address is a group address (including broadcast), then there may be many receivers.
reminder	A control frame sent by the channel **sender** with the most recently used sequence number for a channel which has been inactive for **Reminder Interval** after its most recent data frame.
reminder timer	An implementation-dependent timing mechanism used by a sender to generate a reminder frame after a period of inactivity for a channel. The timer is reset each time a new data frame is transmitted. Conceptually, there is one such timer per channel. The timer interval is **Reminder Interval**.

save timer	An implementation-dependent timing mechanism that limits the time a sender will save a frame waiting for retransmission requests. The timer interval is **Maximum Save Interval**.
sender	The sending station for a channel, usually the station owning the source MAC address.
sequence numbers	Sequence numbers are maintained separately for each logical channel by the sender.

6.6.2.1 Channels

LARQ is defined for operation on simplex logical channels. A separate logical channel is defined for each combination of Ethernet Destination Address, Ethernet Source Address and Link Layer priority. There is no explicit channel setup procedure. A new channel is implicitly defined when a station chooses to send LARQ encapsulated frames for a new combination of DA, SA and Link Layer priority. The station that sends such frames (usually the owner of the SA, except in the case of a bridge masquerading as SA) is the **sender** for the channel. Each channel has a single **sender**. Any station that receives the frames and processes the LARQ headers is a **receiver**. There may be any number of **receivers**. **Receivers** operate independently.

6.6.3 Sender operation

6.6.3.1 Variables and parameters

The clauses that follow make use of the following parameters:

Send Sequence Number	The sequence number of the most recently transmitted data frame.
Reminder Timer Interval	A fixed interval. The default is 50 µs. Lower values will increase the overhead of reminders on network load, while higher values increase the latency for end-of-sequence frames requiring retransmission. Implementations should not use values outside of the range 25-75 µs, based on 150 µs maximum save and hold times.
Minimum Retransmission Interval	An interval used to prevent too-frequent retransmissions of a single frame. Most important for multicast channels. The default is 10 µs
Maximum Save Limit	The maximum number of frames that will be saved for a single logical channel. This is implementation-dependent, and varies with the maximum frame rate the sender is expected to support. Values of 100 or more can be useful for high-speed applications such as video.
Maximum Save Interval	The maximum time that the sender will normally save a frame for possible retransmission. The default is 150 µs.

6.6.3.2 Sender – New Channel

Select implementation-dependent parameters, if necessary.

Select an initial value for **Send Sequence Number**.

6.6.3.3 Sender – Transmit New Data Frame

Access the logical channel state information for the DA, SA and Link Layer priority of the frame.

Increment **Send Sequence Number**, modulo 4096 (the size of the sequence number space).

Build the LARQ header with the new value of **Send Sequence Number**, and the Multiple Retransmission flag set to 0. The Priority field in the LARQ header is set to the Link Layer priority value specified for the frame. If no priority is specified, then the priority shall be set to 0. The method of specifying priority and the choice of value are implementation-dependent and outside the scope of this Recommendation.

Insert a LARQ header (short form control frame format with LARQ_hdr data) between the SA and the Ethertype/Length field of the original frame. The new frame is eight bytes longer than the original.

Save a copy of the frame.

Send the frame.

Restart the **reminder timer** for the channel.

Start a **save timer** for the sequence number. When no other resource limitations apply, a sending station should normally save a frame for **Maximum Save Interval**, which corresponds to **Maximum Hold Interval** used by LARQ receivers.

NOTE – The Link Layer priority for the frame is determined in an implementation-dependent manner, for instance, by examining the IEEE 802.1p (ISO/IEC 15802-3) priority passed along with packets in certain driver implementations.

6.6.3.4 Sender – Process a NACK Control Frame

The priority and Original Destination Address (NACK_DA) are read from the LARQ NACK header.

Access the logical channel state information for the Sender channel, where the channel DA is the NACK_DA and the channel SA is the Ethernet DA from the Nack control frame.

The NACK Count in the LARQ header indicates the number of sequence numbers requested for retransmission. The first indicated sequence number is the value Sequence Number in the NACK header, followed by the next (NACK Count – 1) sequence numbers. For each indicated sequence number starting with the first:

– If a copy of the original frame is no longer available, go to the next sequence number.
– If the most recent retransmission of the frame is within **Minimum Retransmission Interval** of the current time, go to the next sequence number.
– Prepare a copy of the original frame with its original LARQ header for retransmission.
– Copy the value of the Multiple Retransmission Flag from the NACK header into the LARQ header of the frame to be retransmitted.
– Set the LARQ_Rtx flag to 1.
– Send the retransmitted frame.

Do not send a retransmission if a received Nack control frame has an error.

6.6.3.5 Sender – Reminder Timer expiration

If the **reminder timer** expires, create a Reminder control frame, with the Sequence Number set to the current value of Send Sequence Number for the channel. The priority for the Reminder control frame is the same as the priority for the channel.

Send the frame.

Do not restart the **reminder timer** for the channel.

6.6.3.6 Sender – Save Timer expiration

The **save timer** is implementation-dependent. Its purpose is to set an upper bound on how long frames will be saved by a sender for possible retransmission. If set too long, host resources may be wasted saving frames that will never be retransmitted.

This timer is conceptually implemented per sequence number. Release any resources associated with the saved frame.

6.6.3.7 Sender – Resource management

A LARQ implementation requires careful attention to resource management. The resources include the buffers used for saving copies of data for retransmission, the buffers and other resources used to manage the re-ordering of frames to incorporate retransmissions, and the various timers used to govern proper behavior and efficient protocol operation. Resource management is implementation-dependent. However, the following guidelines are recommended.

Saved copies of frames should be kept for **Maximum Save Interval** (default is 150 μs), other considerations notwithstanding.

Maximum Save Limit, the maximum number of saved frames for any channel, should be a function of the maximum rate that new frames may be generated. Very slow devices might usefully save only a couple of frames for retransmission. A high-speed device serving video streams might save 100 or more frames for a single channel.

Senders that save relatively few frames are more likely to receive NACK control frames for sequence numbers that can no longer be retransmitted. Such behavior is inefficient, but causes no other problems.

6.6.4 Receiver operation

6.6.4.1 Channel variables and parameters

The description below of correct protocol operation uses the following variables. The actual implementation may vary so long as the behavior remains unchanged.

Current Sequence Number	The most recent sequence number received in a LARQ header for the channel, whether in a data frame or a reminder control frame.
Oldest missing sequence number	The oldest sequence number for a frame not yet received which has not been declared lost.
Maximum Hold Interval	The longest interval that a frame will be held awaiting an earlier missing frame. The default is to use the same value as **Maximum Save Interval**, which has a default of 150 μs.
Maximum Receive Limit	The maximum number of frames that a receiver will buffer while awaiting an earlier missing frame. The default should normally be the same as the **Maximum Save Limit**.
NACK Retransmission Interval	The interval after which a receiver will retransmit a Nack control frame for a missing sequence number, with the expectation that earlier Nack control frames or data frame retransmissions were lost. The default for fixed implementations is 20 μs.

6.6.4.2 Receiver – New channel

When a data frame with a LARQ header or a LARQ Reminder control frame is received with a new combination of DA, SA and Link Layer priority, the receiver shall initialize state information for a new channel.

The primary piece of state information is the **Current Sequence Number** for the channel. **Current Sequence Number** is initialized to the sequence number immediately preceding that found in the LARQ header of the received frame. This assignment takes place prior to processing the received frame and results in the frame either appearing to be the next expected data frame, or the reminder for the next expected data frame.

6.6.4.3 Receiver – LARQ Data or Reminder frame

Look up the channel state information based on the Ethernet DA and SA in the received frame plus the Link Layer priority from the LARQ header. (Set up a new channel if necessary.)

If the received sequence number of the received frame is out of sequence, the channel state may be reset. If the sequence number (before resetting) is old, and the Forget timer has expired, then the sequence space may be reset to the value of the received frame's sequence number.

If the received sequence number is newer than the **Current Sequence Number** (after any reset of the sequence number space), then perform new sequence number processing steps below; otherwise perform the old sequence number processing steps.

6.6.4.4 Receiver – LARQ frames with CRC or other errors

For best performance, implementations should allow the LARQ protocol module to process errored frames, such as those with payload CRC errors. This will allow Nack indications to be sent quickly since the receiver will not have to wait for the next frame to detect the loss. At the same time, it provides a second opportunity for detecting lost frames at the end of a sequence, when a later Reminder would be the only protection.

If errored frames are used, they shall be used only to detect a very small set of missing sequence numbers for an existing channel (one missed frame is recommended). In particular, if the errored frame appears to have a valid LARQ header, and the frame's source MAC address, destination MAC address, and LARQ header priority match an existing logical channel, and if the sequence number is (**Current Sequence Number** + 1), then treat this frame as a Reminder control frame for the purposes of processing. Note that Reminder control frames are always dropped after processing.

In all other cases, drop the errored frame with no further processing. Do not set up a new channel if the frame has an error. Do not send a retransmission if a Nack control frame has an error. Do not reset a channel (for sequence numbering purposes) for an errored frame.

6.6.4.5 Receiver – New Sequence Number

If the frame has an error indicated by a lower layer driver, such as a CRC error, and the sequence number of the frame is anything other than (**Current Sequence Number** + 1), then drop the frame with no further processing. Otherwise, process the frame as a Reminder control frame.

If the difference between the new sequence number of the received frame and the oldest missing sequence number is greater than (**Maximum Receive Limit** – 1), then repeat the following steps until the acceptable limit is reached.

– Cancel the **Nack retransmission timer** and the **lost frame timer** for the oldest missing sequence number.
– If there is a saved frame for the next sequence number, then deliver in-sequence frames to the next layer above until the next sequence number with a missing frame is reached (which may be the next expected sequence number for the channel, (**Current Sequence Number** + 1)). The value from the Priority field from the LARQ header for each frame is delivered to the next layer along with each associated frame. The method of specifying priority to the next layer is implementation dependent and outside the scope of this Recommendation.

If the sequence number is the next expected sequence number (**Current Sequence Number** + 1) and the frame is a good data frame and there are no older missing sequence numbers, then send the frame up to the next layer.

If the sequence number is newer than (**Current Sequence Number** + 1), or is a reminder for (**Current Sequence Number** + 1), then send one or more Nack control frames requesting retransmission of the missing frame(s).

– The destination address for the Nack is the source address of the received frame. The source address is this station's MAC address. The destination address of the received frame is placed in the original destination address field (NACK_DA) in the LARQ Nack control frame header. The Multiple Retransmission flag is set to 0. The [first] missing sequence number is placed in the sequence number field. The priority for the NACK control frame is the same as the priority for the channel.
– If multiple Nack control frames shall be sent, the earliest sequence number shall be sent first.
– For each missing sequence number a **Nack retransmission timer** is started, set to expire at the current time plus **Nack Retransmission Interval**.
– For each missing sequence number, a **lost frame timer** is started, set to expire at the current time plus **Maximum Hold Interval**.

If the frame is a good data frame and was not delivered to the next layer, then save it.

If the frame is a reminder frame (or an errored data frame) then drop it.

Advance the **Current Sequence Number** to the sequence number in the received frame.

6.6.4.6 Receiver – Old Sequence Number

If the sequence number is the same or older than **Current Sequence Number**, then it will generate no control frames, although it may itself be dropped, held, or sent up to the next higher layer, possibly causing other held frames to be sent up as well. It may cause the cancellation of a **Nack retransmission timer** or **lost frame timer** associated with that sequence number.

– If the frame is not a good (e.g. bad CRC) data frame, or its sequence number is older than the oldest missing frame, or it has already been received (this is a duplicate retransmission), or it is a Reminder frame, then drop the frame and skip further processing for this frame.
– Cancel the **Nack retransmission timer** and the **lost frame timer** for the sequence number.
– If the sequence is not the oldest missing sequence number, then save the frame.
– If the sequence number is the oldest missing sequence number, then deliver the frame up to the next higher layer. If there is a saved frame for the next sequence number, then deliver in-sequence frames to the layer above until the next sequence number with a missing frame is reached (which may be the next expected sequence number for the channel). The value from the Priority field from the LARQ header for each frame is delivered to the next layer along with each associated frame. The method of specifying priority to the next layer is implementation-dependent and outside the scope of this Recommendation.

6.6.4.7 Receiver – Nack retransmission timer expires

If a **Nack retransmission timer** expires, then send another Nack control frame for the associated sequence number. The priority for the Nack control frame is the same as the priority for the channel. Multiple sequence numbers may be nacked at the same time, if their timers expire at similar times.

The Multiple Retransmission flag is set to 1 for Nack control frames sent as a result of retransmission timer expiration.

While there is no explicit limit on the number of Nack control frames sent for a particular sequence number, note that the Nack timer is canceled if the frame is received or if the sequence number is declared lost.

6.6.4.8 Receiver – Lost frame timer expires

The lost frame timer is implementation-dependent. Its purpose is to set an upper bound on how long frames will be held before they are sent up when a frame is really lost. If set too long, network resources may be wasted on NACK control frames sent for frames that the sender on the channel will never retransmit. Further, higher layer transport timers may also become involved. The default value of 150 ms is strongly suggested as an upper bound.

Upon expiration, the sequence number is declared lost, resulting in the cancellation of the **Nack retransmission timer** and the **lost frame timer** for the sequence number. If there is a saved frame for the next sequence number, then send up in-sequence frames until the next sequence number with a missing frame is reached (which may be the next expected sequence number for the channel).

If the **lost frame timers** for multiple sequence numbers expire at the same time, then the timers are processed in sequence from oldest to newest.

6.6.4.9 Receiver – Forget timer

The forget timer is an implementation-dependent mechanism to allow a receiver to reset the sequence number space of a channel when a received sequence number is not the next expected (**Current Sequence Number** + 1) and a relatively long interval has expired since the last frame received on the channel. Once expired, a receiver will accept any unusual sequence number as the next expected sequence number, allowing for undetected resets of other stations, disconnection from the network, etc. The definition of "unusual sequence number" is implementation-dependent, but generally means any old sequence number or any new sequence number that is not close to the current sequence number, where "close" is 1 or some other small integer. A one-second default is suggested.

6.6.4.10 Receiver – Resource management

In general, the receiver will want to set upper bounds on the number held frames per channel and the number of held frames across channels. The bounds may vary based on the priority of the channel.

Timer intervals may vary based on factors such as the priority of the channel, or measured intervals for successful retransmissions.

The description above suggests per-sequence number timers. This is for descriptive purposes only, and does not imply any implementation mechanism.

6.7 Vendor-specific formats

The following two frame types allow vendor-specific extensions (Tables 22 and 23). The short format vendor-specific subtype allows short control messages and encapsulation headers, while the Long Format subtype allows extensions that require longer messages.

Table 22/G.989.2 – Vendor-specific short frame

Field	Length	Meaning
DA	6 octets	Ethernet Destination Address
SA	6 octets	Ethernet Source Address
Ethertype	2 octet	$886C_{16}$
SSType	1 octet	= 5
SSLength	1 octet	Number of control information octets, starting with the SSVersion field and ending with the second(last) octet of the Next Ethertype field. SSLength shall be ≥ 6 for SSVersion 0.
SSVersion	1 octet	= 0
Vendor OUI	3 octets	An ISO/IEC (IEEE) assigned Organizationally Unique Identifier
Control data	0-249 octets	Vendor-specific control data
Next Ethertype	2 octets	= next Ethertype if an encapsulation format, or 0 if no encapsulated frame
Pad	0-38 octets	Any value octet
FCS	4 octets	

Table 23/G.989.2 – Vendor-specific long frame

Field	Length	Meaning
DA	6 octets	Destination Address
SA	6 octets	Source Address
Ethertype	2 octet	$886C_{16}$
LSType	2 octets	= 32769
LSLength	2 octets	Number of control information octets, starting with the LSVersion field and ending with the second(last) octet of the Next Ethertype field. LSLength shall be > 6 for LSVersion 0.
LSVersion	1 octet	= 0
Vendor OUI	3 octets	An ISO/IEC (IEEE) assigned Organizationally Unique Identifier
Control data	1-65531 octets	Vendor-specific data
Next Ethertype	2 octets	= next Ethertype if an encapsulation format, or 0 if no encapsulated frame
Pad	40-0 octets	If needed to make minimum size frame. Should be zero
FCS	4 octets	

6.8 Minimal Link Protocol Support Profile

The Minimal Link Protocol Support Profile allows for less complex implementations of this Recommendation. While each of the component protocols serves an important function in the operation of the network, it is possible to implement minimal support for some of the more complex protocols while maintaining compatibility with fully functional implementations, without detracting from the overall performance of other stations. The shorter name, Minimal Profile, will be used in the following description. The alternative is full support of all the link protocols, called the Full Link Protocol Support Profile, or Full Profile for short.

6.8.1 Limitations of Minimal Profile stations

A Minimal Profile station can send only best-effort data traffic, and shall treat all received traffic as best-effort. A Minimal Profile station cannot advertise or use optional features that may be defined in the future. Due to the lack of support for LARQ, a Minimal Profile station may see dramatically reduced network throughput.

6.8.2 Full forward-compatible support for both short and long format Link Protocol frames

A Minimal Profile station shall be able to handle all G.989.2 Link Protocol frames, that is, those marked with the Ethertype $886C_{16}$ in the Ethernet header of the received frame. This includes dropping control frames with unknown subtypes and de-encapsulating data frames with unknown subtypes. The length field shall be used to locate the Next_Ethertype field in order to determine whether frames are control or data (encapsulated) frames.

6.8.3 Full support for rate-selection and link integrity

A Minimal Profile station shall implement the standard G.989.2 Link Integrity function, including suppression of LICFs.

A Minimal Profile station shall implement the full set of rate-selection functions.

6.8.4 Minimum support for LARQ

A Minimal Profile station shall properly handle frames with LARQ headers. It shall drop received control frames. It shall properly remove LARQ headers from data frames. In addition, if the LARQ header on a data frame has the retransmission flag set, then the frame shall be dropped in order to prevent duplicate and out-of-order frames.

A Minimal Profile station should add LARQ headers to data frames being transmitted, setting the priority to 0 and the LARQ_NORtx flag to 1 in the LARQ headers.

If LARQ headers are added, the minimal station shall use default priority mapping.

6.8.5 Receive-only support for CSA

A Minimal Profile station listens to CSA Control Frames and performs mode selection based on the configuration flags received. In particular, it uses the union of the CSA_CurrentTxSet and CSA_CurrentRxSet as the set of in-use flags. A Minimal Profile station does not send CSA Control Frames, and can therefore never advertise optional features, or use non-default priorities.

6.8.6 Priority for transmitted frames

In addition to control frames, a Minimal Profile station shall only send normal data frames using the default priority assigned to best-effort/unspecified QOS. The Link Layer priority value for this QOS is 0. If the station is not adding LARQ headers, then data frames shall be sent using the default physical layer priority for link layer priority 0. (I.e. it shall use physical layer priority 2.) If LARQ headers are being added as specified above, then the LARQ header priority field is set to 0, and the station again uses the default re-mapping function for link layer priorities to determine the actual Physical layer priority to use for Link Layer priority 0. (i.e. it shall use Physical Layer priority 2.)

6.8.7 Priority for Received frames

A Minimal Profile station should only indicate LL priority 0, if any priority is indicated, for received frames, regardless of the Physical Layer priority or priority value in a LARQ header.

6.8.8 Additional requirements for Full Profile stations

In support of Minimal Profile stations, a minor addition is also required to the CSA specification for Full Profile stations. Any station that is not sending CSA frames, but which is determined to be a G.989.2 station as a result of traffic received from that station, is treated as if it advertised a default set of status flags, including no supported options, only LL priority 0 in use, and G.989.2 operation.

Annex A

Receiver performance

A.1 Sensitivity

The receiver shall meet the requirements of this clause for received frames over a clock frequency tolerance of ±100 ppm.

A.1.1 Maximum signal

The receiver shall detect frames with peak voltage up to –6 dBV across tip and ring at a frame error rate of no greater than 10^{-4} with additive white Gaussian noise at a PSD of less than –140 dBm/Hz, measured at the receiver.

A.1.2 Minimum sensitivity

The receiver shall detect 1518-octet frames encoded with PE = 1 with rms voltage as low as 2.5 mV at no greater than 10^{-4} frame error rate. The rms voltage is computed only over time during which the signal is present.

The receiver shall detect no more than 1 in 10^4 1518-octet, PE = 1 frames with rms voltage less than 1.0 mV.

Both criteria assume additive white Gaussian noise at a PSD of less than –140 dBm/Hz, measured at the receiver, and assume a flat channel.

A.2 Immunity to narrowband interference

The receiver shall demodulate frames with payload encoded at PE = 5 and at PE = 13 (if implemented), and differential rms voltage as low as 20 mV (measured over the header), with additive white Gaussian noise with PSD up to –130 dBm/Hz at the receiver input, at a frame error rate less than 10^{-4}, in the presence of interferers listed in this clause.

A.2.1 Differential-mode interference

The requirements of this clause shall be met in the presence of any single-tone interferer with frequency and level characteristics that fall within the ranges listed in Table A.1:

Table A.1/G.989.2 – Metallic interferer amplitudes

Frequency range (MHz)	Maximum peak-to-peak interferer level (Volts)
0.01-0.1	6.0
0.1-0.6	3.3
0.6-1.7	1.0
1.7-4.0	0.1
7.0-7.3	0.1
10.0-10.15	0.1
14.0-14.35	0.28
18.068-18.168	0.5
21.0-21.45	0.5
24.89-24.99	0.5
28.0-29.7	0.5

The applied voltage is specified across tip and ring at the input to the transceiver.

A.2.2 Common-mode interference

The requirements of this clause shall be met in the presence of any single-tone interferer with frequency and level characteristics that fall within the ranges listed in Table A.2:

Table A.2/G.989.2 – Longitudinal interferer amplitudes

Frequency range (MHz)	Maximum peak-to-peak interferer level (Volts)
0.01-0.1	20.0
0.1-0.6	20.0
0.6-1.7	10.0
1.7-4.0	2.5
7.0-7.3	2.5
10.0-10.15	2.5
14.0-14.35	5.0
18.068-18.168	5.0
21.0-21.45	5.0
24.89-24.99	5.0
28.0-29.7	5.0

The applied voltage is specified between the center tap of a test transformer and ground, at the input to the transceiver.

NOTE – The common mode rejection of the test transformer used to insert the signal should exceed 60 dB up to 100 MHz.

A.3 System margin

Ten test loops, provided in Annex B, shall be used to verify receiver performance. The following impairments will be applied in each loop test: additional (flat) attenuation, additive white Gaussian noise, narrowband interferers, and 120 Hz impulse noise ("light dimmer noise").

The Frame Error Rate (FER) specified for each table entry shall be achieved for the specified payload encoding at an impairment level greater than that that shown in the table entry.

A system margin requirement for a single time-varying channel is also defined.

An entry of "–" in a table indicates that there is no requirement under the specified conditions.

A.3.1 Attenuation impairment

The attenuator setting described in the following Table A.3 is the additional attenuation applied in series with the specified wire loop.

Table A.3/G.989.2 – Additional attenuation settings

Test	Required impairment attenuator setting (dB) Loop number					
	1	**4**	**5**	**6**	**8**	**9**
PE = 1 FER = 10^{-2}	34	16	22	11	12	18
PE = 1 FER = 10^{-3}	33	15	21	10	11	17
PE = 5 FER = 10^{-2}	30	9	18	6	8	–
PE = 5 FER = 10^{-3}	29	8	17	5	7	–

A.3.2 Additive white noise impairment

White noise power at 0 dB attenuator setting (Table A.4): –70 dBm/Hz. The output of the noise attenuator shall be added at the receiver. For loop 1, 20 dB of flat-channel attenuation shall be placed in series with the loop.

Table A.4/G.989.2 – White noise attenuation settings

Test	Required impairment attenuator setting (dB) Loop number					
	1	**4**	**5**	**6**	**8**	**9**
PE = 1 FER = 10^{-2}	42	40	36	46	43	39
PE = 1 FER = 10^{-3}	43	41	37	47	44	40
PE = 5 FER = 10^{-2}	58	57	53	63	60	–
PE = 5 FER = 10^{-3}	59	58	54	64	61	–

A.3.3 Narrowband interference impairment

Narrowband interference peak-to-peak amplitude at 0 dB attenuator setting (Table A.5): 2.0 volts at 7.0 and 7.3 MHz. White Gaussian noise is simultaneously applied at a level of –135 dBm/Hz.

Table A.5/G.989.2 – Narrowband interference attenuation settings

Test	Required impairment attenuator setting (dB) Loop number					
	1	**4**	**5**	**6**	**8**	**9**
PE = 1 FER = 10^{-2}	26	26	26	26	26	26
PE = 1 FER = 10^{-3}	26	26	26	26	26	26
PE = 5 FER = 10^{-2}	26	30	26	32	30	–
PE = 5 FER = 10^{-3}	26	31	26	33	31	–

A.3.4 Impulse noise requirements

Impulse noise peak-to-peak amplitude at 0 dB attenuator setting (Table A.6): 3.0 volts. White Gaussian noise is simultaneously applied at a level of –135 dBm/Hz. The impulse shall be defined as two cycles of a 5.0 MHz square wave summed with four cycles of a 7.0 MHz square wave.

Table A.6/G.989.2 – Impulse noise attenuation settings

Test	Required impairment attenuator setting (dB) Loop number	
	2	**9**
PE = 1 FER = 10^{-2}	3	3
PE = 1 FER = 10^{-3}	3	3
PE = 5 FER = 10^{-2}	3	–
PE = 5 FER = 10^{-3}	3	–

A.3.5 Dynamic channel system margin

The system shall receive no more than 5 1518-octet frames in error out 3000 when sent at a rate of 5 frames per 10 µs period over loop #2 under the following conditions:

a) During this test, the 330 pF capacitor terminating one of the stubs shall be switched in and out of the loop once per second; i.e. an open-circuit termination shall be used for a period of 1 s every 2 s.

b) White noise at a level of –140 dBm/Hz shall be added at the receiver.

c) The PE shall be 5 (6 bits/symbol).

NOTE – Switching a capacitor in and out of the loop simulates a switch-hook transition on a common telephone.

A.3.6 Telephony ringing signal immunity

G.989.2 devices are intended to operate in the presence of a telephony ringing signal event from a telephone central office

When subjected to the simulated telephony ringing signal as shown in Figure A.1, the device Frame Error Rate for 1518-octet frames encoded with PE = 1 shall not exceed 0.1%.

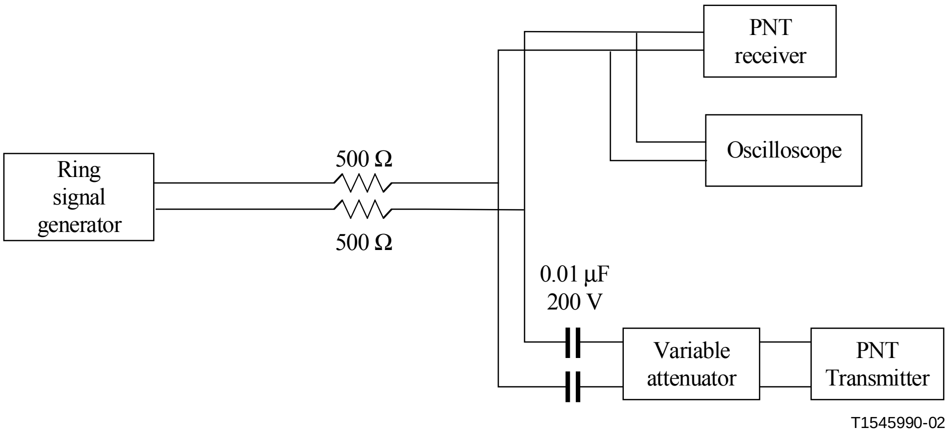

Figure A.1/G.989.2 – Telephony ringing signal conditions

The simulated ringing signal shall consist of a 20 Hz sinusoid with a level of 90 V_{rms} superimposed on a d.c. bias level of –52 V. The ringing signal shall be continuously cycled on and off with an on-time of 2 s and an off-time of 4 s.

NOTE – The two 0.01 µF capacitors in the Figure A.1 are required to d.c. isolation, as many attenuators have low impedance at d.c. that could significantly reduce the ringing voltage.

Annex B

Network test loops

Ten test loops are defined for evaluating the performance of PNT receivers. This annex includes specification of the wire types and the topologies.

B.1 Wire model

The wires labelled "quad", "flat 4-wire" and "UTP-5" are representative of typical in-premises telephony wiring.

For simulations, the following model is used to generate primary parameters R, L, G, and C *vs.* frequency:

$$R(f) = \sqrt[4]{R_o^4 + a \cdot f^2}$$

$$L(f) = \frac{l_0 + l_\infty \cdot \left(\frac{f}{f_m}\right)^b}{1 + \left(\frac{f}{f_m}\right)^b}$$

$$G(f) = g_0 \cdot f^{g_e}$$

$$C(f) = c_\infty + \frac{c_0}{f^{c_e}}$$

The parameter set for each of the wire types used in B.2 is given in Table B.1. The assumption is that R(f) is in units of ohms/km, L(f) is in units of mH/km, G(f) is in units of µMhos/km, and C(f) is in units of µF/km.

Table B.1/G.989.2 – Model parameters for wires

Model parameter	Quad cable	Flat 4-wire	UTP-5
r_0	252.58	399.6273	172.2
a	0.164	0.470	0.173
l_0	0.763	0.789	0.613
b	0.493	0.406	0.516
l_∞	0.576	0.592	0.446
f_m	239.8×10^3	432.9×10^3	310.6×10^3
g_0	0.0268	0.322	0.000175
g_e	0.547	0.467	0.540
c_0	0.075	0.0248	0
c_∞	0.044	0.0427	0.0516
c_e	0.152	0.0758	0

B.2 Test loops

In the following diagrams (Figures B.1 to B.10), unterminated stubs are marked "open". Stubs terminated by resistors are marked with the resistance value in ohms. Stubs terminated by capacitors are marked with the capacitance value in picofarads.

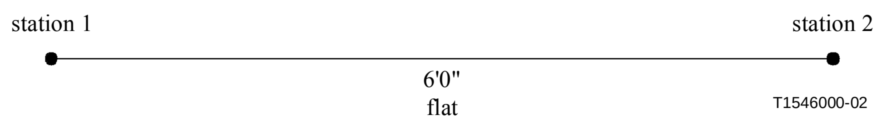

Figure B.1/G.989.2 – Test Loop No. 1

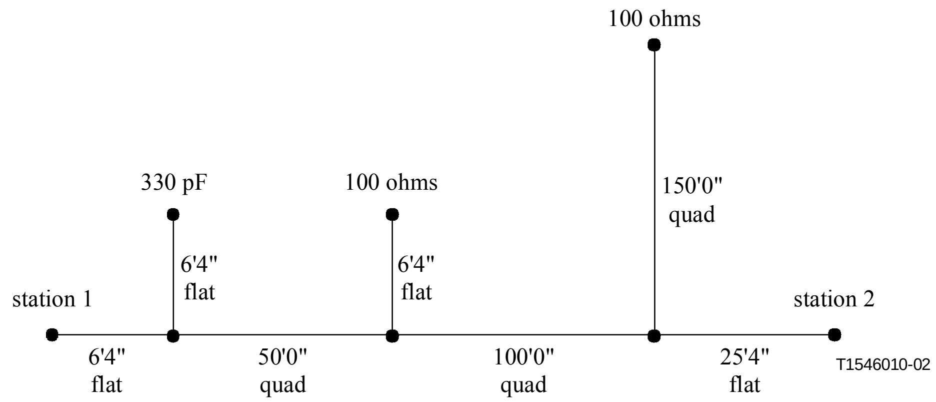

Figure B.2/G.989.2 – Test Loop No. 2

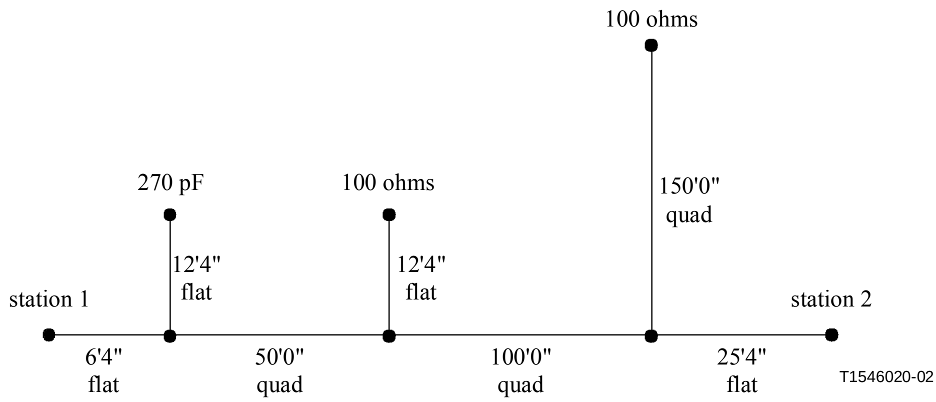

Figure B.3/G.989.2 – Test Loop No. 3

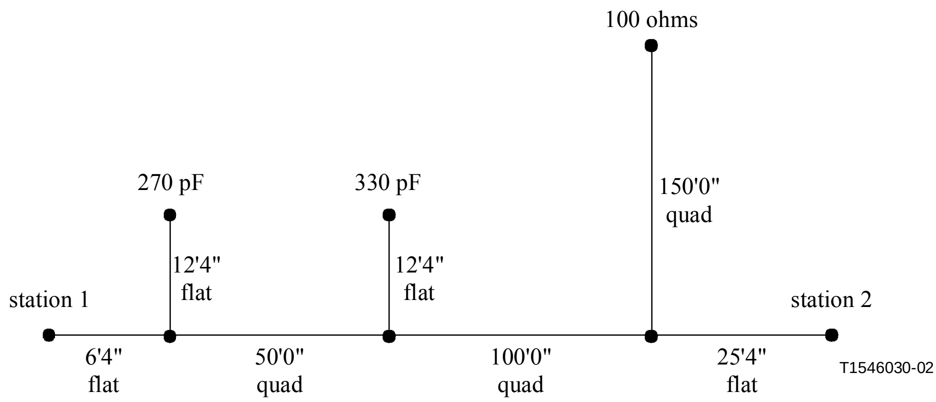

Figure B.4/G.989.2 – Test Loop No. 4

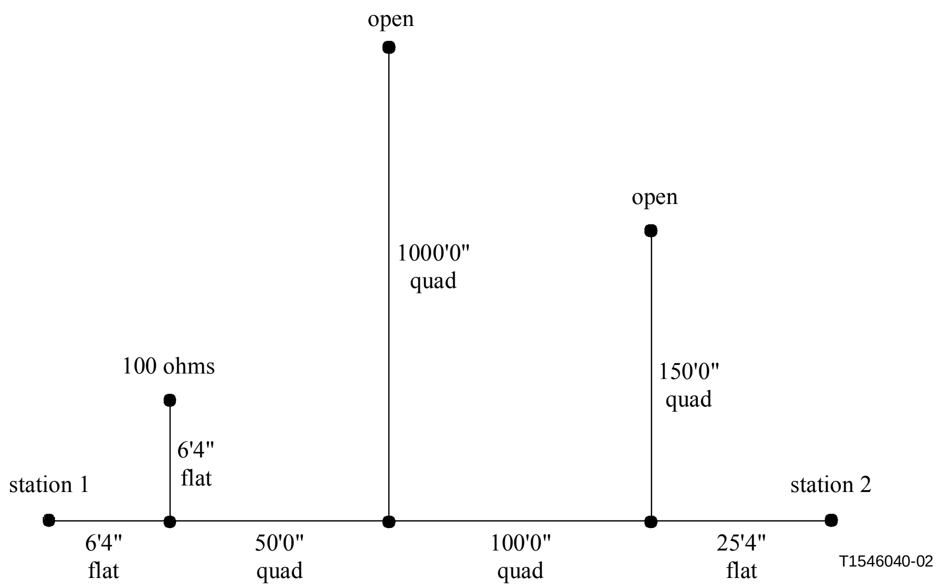

Figure B.5/G.989.2 – Test Loop No. 5

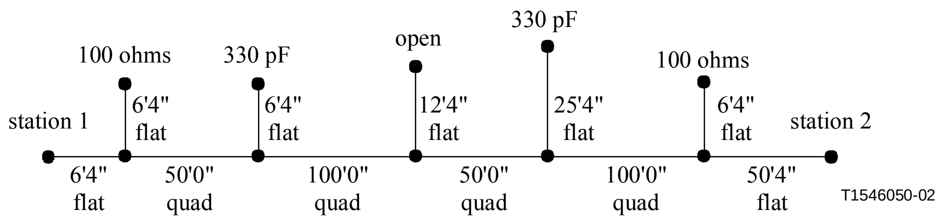

Figure B.6/G.989.2 – Test Loop No. 6

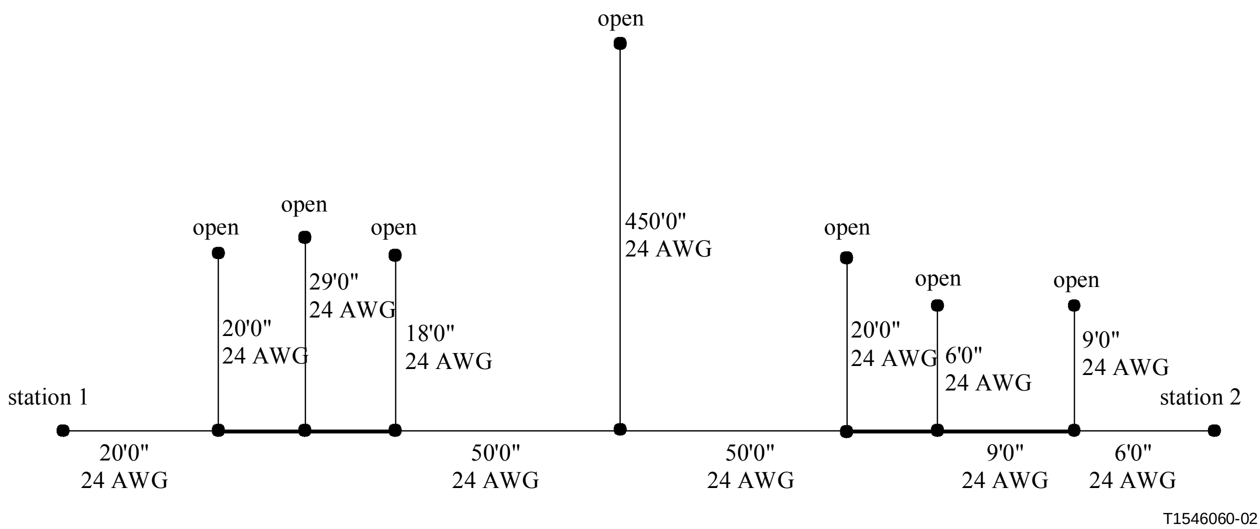

Figure B.7/G.989.2 – Test Loop No. 7

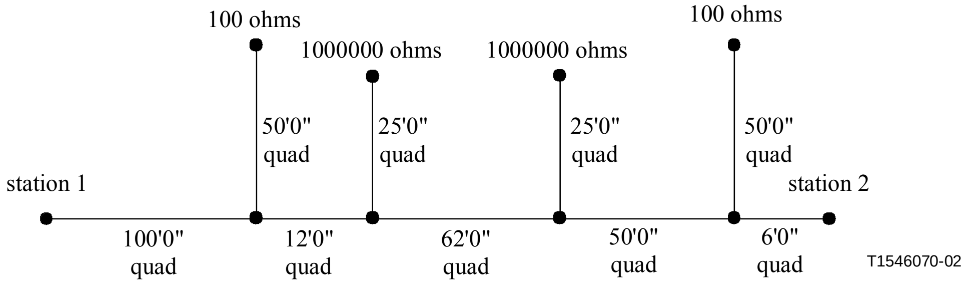

Figure B.8/G.989.2 – Test Loop No. 8

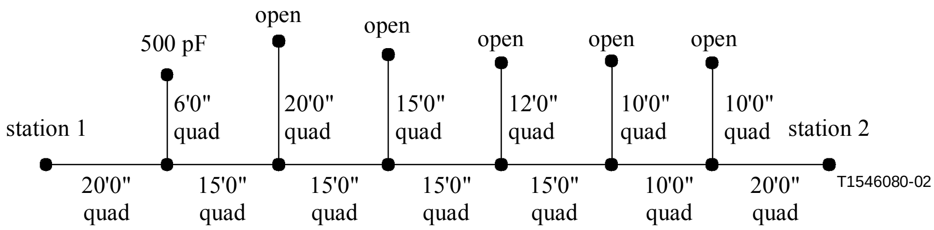

Figure B.9/G.989.2 – Test Loop No. 9

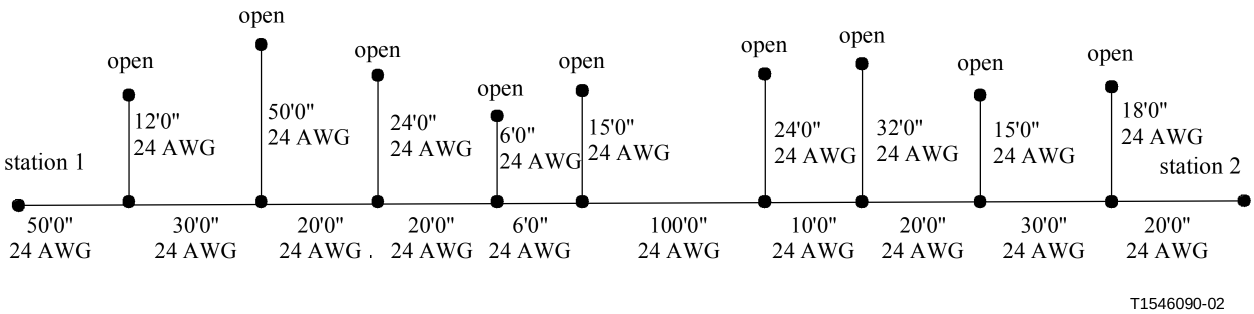

Figure B.10/G.989.2 – Test Loop No. 10

SERIES OF ITU-T RECOMMENDATIONS

Series A	Organization of the work of ITU-T
Series B	Means of expression: definitions, symbols, classification
Series C	General telecommunication statistics
Series D	General tariff principles
Series E	Overall network operation, telephone service, service operation and human factors
Series F	Non-telephone telecommunication services
Series G	**Transmission systems and media, digital systems and networks**
Series H	Audiovisual and multimedia systems
Series I	Integrated services digital network
Series J	Cable networks and transmission of television, sound programme and other multimedia signals
Series K	Protection against interference
Series L	Construction, installation and protection of cables and other elements of outside plant
Series M	TMN and network maintenance: international transmission systems, telephone circuits, telegraphy, facsimile and leased circuits
Series N	Maintenance: international sound programme and television transmission circuits
Series O	Specifications of measuring equipment
Series P	Telephone transmission quality, telephone installations, local line networks
Series Q	Switching and signalling
Series R	Telegraph transmission
Series S	Telegraph services terminal equipment
Series T	Terminals for telematic services
Series U	Telegraph switching
Series V	Data communication over the telephone network
Series X	Data networks and open system communications
Series Y	Global information infrastructure and Internet protocol aspects
Series Z	Languages and general software aspects for telecommunication systems

Printed in Switzerland
Geneva, 2002