5i’

Recommendation G.708

NETWORK NODE INTERFACE FOR THE SYNCHRONOUS DIGITAL HIERARCHY

(Melbourne, 1988)

The CCITT,

considering

(a) that network node interface (NNI) specifications are necessary to enable interconnection of synchronous digital network elements for transport of payloads, including digital signals of the asynchronous hierarchy defined in Recommendation G.702;

(b) that Recommendation G.707 describes the advantages offered by a synchronous digital hierarchy and multiplexing method and specifies a set of synchronous digital hierarchy bit rates;

(c) that Recommendation G.709 specifies the multiplexing structures;

(d) that Recommendations G.707, G.708 and G.709 form a coherent set of specifications for the synchronous digital hierarchy and NNI;

(e) that Recommendation G.802 specifies the interworking between networks based on different asynchronous digital hierarchies and speech encoding laws,

recommends

that the frame structure for multiplexed digital signals at the network node interface of a synchronous digital network including ISDN should be as described in this Recommendation.

1 Location of NNI

Figure 1-1/G.708 gives a possible network configuration to illustrate the location of the network node interface specified in this Recommendation.

Figure 1-1/G.708, p.

2 Basic multiplexing principle and multiplexing elements

2.1 *General*

Frame structures and overheads in this Recommendation are mainly in the context of circuit mode connection types rather than asynchronous transfer mode (ATM). ATM based multiplexing principles are under study.

Figure 2-1/G.708 shows the relationship between various multiplexing elements that are defined below, and illustrates possible multiplexing structures.

Figures 2-2/G.708, 2-3/G.708 and 2-4/G.708 are examples of how various signals are multiplexed using these multiplexing elements.

The legends used in these figures are defined in § 2.2.

Details of the multiplexing method and mappings are given in Recommendation G.709.

Note — When signals at bit rates of the various multiplexing elements of the synchronous digital hierarchy (Recommendations G.707, G.708, G.709) are different from existing hierarchy levels in Recommen dation G.702, the signals are not required to be transported via digital networks which are in line with Recommendation G.702.

2.2 *Definitions*

2.2.1 *Container , C-n (n = 1 to 4)*

This element is a defined unit of payload capacity which is dimensioned to carry any of the levels currently defined in Recommendation G.702 and may also provide capacity for transport of broadband signals which are not yet defined.

2.2.2 *Virtual container , VC -n*

Two types of virtual containers have been identified:

— *Basic virtual container, VC-n (n = 1, 2)*

This element comprises a single C-*n (n = 1, 2)* plus the basic virtual container path overhead (POH) appropriate to that level.

— *Higher order virtual container to VC-n (n = 3, 4)*

This element comprises a single C-*n (n = 3, 4)* , an assembly of tributary unit groups (TUG-2s) or an assembly of TU-3s, together with virtual container POH appropriate to that level.

Figure 2-1/G.708, p. 2 A L'ITALIENNE

Figure 2-2/G.708, p. 3

Figure 2-3/G.708, p. 4

Figure 2-4/G.708, p. 5

2.2.3 *Tributary unit TU -n (n = 1 to 3)*

This element consists of a virtual container plus a tributary unit pointer. A tributary unit pointer indicates the phase alignment of the virtual container (VC-*n*) with respect to the POH of the next higher level virtual containers in which it resides. The tributary unit pointer location is fixed with respect to this higher level POH.

In certain applications (for example, synchronous mapping providing direct observability of 64 kbit/s channels) the basic virtual container has a fixed phase-alignment with respect to the higher level virtual container. In this case, the basic virtual container (VC-1) POH and TU-1 pointer are null.

2.2.4 *Tributary unit group , TUG -2*

This element consists of a homogeneous assembly of TU-1s or a single TU-2.

2.2.5 *Administrative unit, AU -n (n = 3, 4)*

This element consists of a VC-*n* (*n* = 3, 4) plus an administrative unit pointer (AU PTR) (n = 3, 4) with respect to the STM-1 frame. The administrative unit pointer location is fixed with respect to the STM-1 frame.

2.2.6 *Synchronous transport module level 1, STM-1*

This element is the basic building block of the synchronous digital hierarchy and it comprises either one AU-4 or multiple AU-3s, together with the section overhead (SOH).

2.2.7 *Synchronous transport module level N, STM-N*

This element defines the N-th level of the synchronous-digital hierarchy and contains N synchronously multiplexed STM-1 signals.

The STM-N-signal can be obtained via single-or multiple-stage multiplexing.

Values of N correspond to the synchronous digital hierarchy levels given in Recommendation G.707.

3 Frame structure

3.1 *Level 1: 155 520 kbit/s (STM-1)*

3.1.1 *Basic frame structure*

The STM-1 frame structure is shown in Figure 3-1/G.708. The three main areas of the STM-1 frame are indicated:

— section overhead;

— AU pointers;

— STM-1 payload.

Figure 3-1/G.708, p.

3.1.2 *Section overhead*

Rows 1-3 and 5-9 of columns 1-9 of the STM-1 in Figure 3-1/G.708 are dedicated to the section overhead.

The allocation of section overhead capacity and functions is given in Figure 3-4a/G.708. An explanation of the overhead functions is given in § 5.

3.1.3 *Administrative unit (AU) pointers*

Row 4 of columns 1-9 and row 1-3 of columns 11-14 in Figure 3-1/G.708 are available for AU pointers. The positions of the pointers of the AUs for different organizations of the STM-1 payload are shown in Table 3-1/G.708. The application of pointers and their detailed specifications are given in Recommendation G.709.

H.T. [T1.708]

TABLE 3-1/G.708

Position of AU pointers

AU	Position of AU pointer
31	Areas A and B
32	Area A**s and B**
4	Area A**s and B**

Table 3-1/G.708 [T1.708], p.

3.1.4 *Administrative units in the STM-1*

The STM-1 payload can suppport the following types and numbers of administrative units:

— one AU-4; or three AU-32s; or four AU-31s.

The VC-*n* associated with each AU-*n* does not have a fixed phase with respect to the STM-1 frame. The location of the first byte of the VC-*n* is indicated by the AU-*n* pointer. The AU-*n* pointer is in a fixed location in the STM-1 frame as illustrated in Figures 2-2/G.708 to 2-4/G.708 and 3-1/G.708 to 3-3/G.708.

The AU-4 may be used to carry, via the VC-4, three TU-32s or four TU-31s. This nested arrangement is illustrated in Figures 3-2/G.708 and 3-3/G.708. The VC-3 associated with each TU-3 does not have a fixed phase relationship with respect to the start of the VC-4. The TU-3 pointer is in a fixed location in the VC-4 and the location of the first byte of the VC-3 is indicated by the TU-3 pointer (illustrated in Figures 3-2/G.708 and 3-3/G.708).

FIGURE 3-2/G.708,p .

FIGURE 3-3/G.708,p .

3.1.5 *VC-4 and VC-3 path overheads*

The allocation of the VC-4 and VC-3 path overhead capacity and functions is given in Figure 3-4/G.708. An explanation of the overhead functions is given in § 5.

The position of the VC-4 and VC-3 path overhead is specified in Recommendation G.709.

Figure 3-4/G.708, p.

3.2 *Level 4: 622 080 kbitB/Fs (STM-4)*

This level is obtained by one-byte interleaving of four STM-1s as illustrated in Figure 3-5/G.708.

The SOH of the STM-1s shall be 125 s phase aligned prior to multiplexing such that the SOH of the resulting STM-4 is contained in the first 36 columns. The AU pointer value(s) of each STM-1 is/are adjusted to indicate the start of the VC(s) with respect to this new position of the AU pointer(s) which is fixed relative to the STM-4 SOH.

FIGURE 3-5/G.708,p .

4 Interconnection of STM-1s

The synchronous digital hierarchy, specified in Recommendations G.707, G.708 and G.709, is designed to be universal, allowing transport of a large variety of signals including those specified in Recommendation G.702.

However, there are a number of options for structuring an STM-1. This section provides guidelines for the interconnection of STM-1s. Two general cases are considered:

— Case 1: STM-1s having the same structure (detailed in § 4.1);

— Case 2: STM-1s having different structures (detailed in § 4.2).

4.1 *Interconnection of STM-1s having the same structure*

The interconnection unit used between STM-1s is the VC associated with the AU. This arrangement is shown in row i) of Table 4-1/G.708.

4.2 *Interconnection of STM-1s having different structures*

In the case of STM-1s having different structures, the following guidelines should be used to facilitate interconnection by bilateral agreement or to resolve contention.

The method of interconnection between STM-1s having different structures depends on whether the type of AU is different or whether the type of TUG is different. The cases are considered in three categories:

— different types of AU-3s carrying a C-3 payload;

— different types of AU carrying the same type of TUG-2;

— different types of TUG-2s.

4.2.1 *Different types of AU-3s carrying a C-3 payload*

For the interconnection of different types of AU-3s carrying a C-3 payload, the C-3 payload is transferred from the AU-3 to a corresponding TU-3. This TU-3 is then assembled into a VC-4 using the nested approach illustrated in Figure 3-3/G.708. This arrangement is shown in row ii) of Table 4-1/G.708, and is intended to facilitate the transit of C-3 in a VC-3 across a network which cannot support the associated AU-3.

4.2.2 *Different types of AU carrying the same type of TUG*

For the interconnection of a different type of AU carrying the same type of TUG-2, the TUG-2s are transferred between the dissimilar AUs. In the absence of bilateral agreement on an AU-3 type, the AU-4 shall be used. This arrangement is shown in row iii) of Table 4-1/G.708.

4.2.3 *Different types of TUG-2s*

For the interconnection of different types of TUG-2s, the TU-1s are transferred from the TUG-22 to the TUG-21. The TUG-21 is used as the interconnection unit. In the absence of bilateral agreement on an AU-3 type, the TUG-21s are directly assembled into a VC-4. This arrangement is shown in row iv) of Table 4-1/G.708.

The method of interconnection between an AU-31 containing TUG-21s and an AU-31 containing TUG-22s is for further study.

5 Overhead functions

5.1 *Types of overhead*

Several types of overhead have been identified for application in the synchronous digital hierarchy. The types of overhead described below and their applications are shown in Figure 5-1/G.708.

5.1.1 *Section overhead (SOH)*

Section overhead capacity is added to either an AU-4 or an assembly of AU-3s to create an STM-1. The content always includes STM-1 framing. Content representing section performance monitoring and other maintenance and operational functions can be added or modified without disassembly of the STM-1, as appropriate, for various configurations of elements (e.g. intermediate regenerator monitoring, protection switching control).

H.T. [T2.708]

TABLE 4-1/G.708
{
Interconnection of STM-1s
}

	STM-1 structure A	Conversion steps	Interconnection {	Conversion steps	STM-1 structure B	Parameters
i)	{					
AU-*x*						
/C-*x*						
or						
TUG-2*p*						
}	{					
AU-*x*						
VC-*x*						
}	AU-*x*	VC-*x*	{			
VC-*x*						
← AU-*x*						
}	{					
AU-*x*						
/C-*x*						
or						
TUG-2*p*						
}	{					
x						
= 4, 32 or 31						
p						
= 1 or 2						
}						
ii)	{					
AU-3*x*						
/C-3*x*						
}	{					
AU-3*x*						
VC-3*x*						
TU-3*x*						
VC 4						
}	AU-4	VC-4	VC-4 ← AU-4	{		
AU-4/TU-3*x*						
/C-3*x*						
}	*x* = 1 or 2					
iii)	{					
AU-*x*						
/TUG-2*p*						
}	{					
AU-*x*						
VC-*x*						
TUG-2*p*						
}	{					
AU-*y*						
| ua)						
}	TUG-2*p*	{				
TUG-2*p*						
← VC-*z*						
← AU-*z*						
}	{					
AU-*z*						
/TUG-2*p*						
}	{					
x						
= 4, 32 or 31						
y						

= 4, 32 or 31 *z* = 4, 32 or 31; *z* ≠ *x* *p* = 1 or 2 }							
iv) AU-*x* /TUG-22/TU-1*p* } AU-*x* VC-*x* TUG-21 } AU-*y* | ua) } TUG-21 ← TU-1*p* ← TUG-22 ← ← VC-*z* ← AU-*z* } AU-*z* /TUG-22/TU-1*p* } *x* = 4, 32 or 31 *y* = 4, 32 or 31 *z* = 4 or 31; *z* ≠ *x* *p* = 1 or 2 (see note) }	{ { { TUG-21 { {	{					

Item to the left is carrying item to the right.

Conversion ‘‘From To’’ (to arrive at the interconnection unit).

a) In the absence of a bilateral agreement on an AU-3, an AU-4 should be used.

Note — The case where x = 31 and z = 31 is for further study.

Tableau 4-1/G.708 [T2.708], p. 12   A L’ITALIENNE

5.1.2 *Virtual container path overhead (POH)*

Virtual container path overhead provides for communication between the point of assembly of a virtual container and its point of disassembly. Two categories of virtual container path overhead have been identified:

— *Basic virtual container path overhead (VC-1, 2 POH)*

Basic virtual container POH is added to the container (C-1, 2) when the VC-1, 2 is created. Among the functions included in this overhead are virtual container path performance monitoring, signals for maintenance purposes and alarm status indications.

— *Higher order virtual container path overhead (VC-3, 4 POH)*

VC-3 POH is added to either an assembly of TUG-2s or a C-3 to form a VC-3. VC-4 POH is added to either an assembly of TU-3s or a C-4 to form a VC-4. Among the functions included within this overhead are virtual container path performance monitoring, alarm status indications, signals for maintenance purposes and multiplex structure indications (VC-3,4 composition).

Figure 5-1/G.708, p.

5.2 *Overhead descriptions*

The location of the various section and VC-3, 4 path overhead bytes in the STM-1 frame is illustrated in Figure 3-4/G.708.

5.2.1 *SOH byte descriptions*

5.2.1.1 *Framing: A1, A2*

Six bytes are dedicated to each STM-1. The pattern shall be A1A1A1A2A2A2 (A1=11110110, A2=00101000). These bytes shall be provided in all STM-1 signals within an STM-N.

5.2.1.2 *Data communication channel: D1-D12*

Twelve bytes are allocated for section data communication. These bytes are defined only for STM-1 No. 1 of an STM-N signal.

5.2.1.3 *STM identifier: C1*

This is a unique number assigned to an STM-1 prior to it being multiplexed to a higher STM-N level. Upon demultiplexing, this byte may be used to identify the position of any particular STM-1 within the incoming STM-N signal.

5.2.1.4 *Orderwire* : *E1, E2*

These two bytes provide orderwire channels for voice communication. These bytes are defined only for STM-1 No. 1 of an STM-N signal.

5.2.1.5 *User channel: F1*

This byte is reserved for user purposes (e.g. network operators). This byte is defined only for STM-1 No. 1 of an STM-N signal.

5.2.1.6 *BIP-8: B1*

One byte is allocated in each STM-1 for a bit error monitoring function of an elementary regenerator section. This function shall be a Bit Interleaved Parity 8 (BIP-8) code using even parity. The BIP-8 is computed over all bits of the previous STM-N frame after scrambling and is placed in byte B1 before scrambling. (For details of the scrambling process, see Recommendation G.709.) The B1 byte shall be monitored and recomputed at every regenerator.

Note — Bit Interleaved Parity-N (BIP-N) code is defined as a method of error monitoring. With even parity, an N bit code is generated by the transmitting equipment over a specified portion of the signal in such a manner that the first bit of the code provides even parity over the first bit of all N-bit sequences in the covered portion of the signal, the second bit provides even parity over the second bit of all N-bit sequences within the specified portion, etc. Even parity is generated by setting the BIP-N bits so that there are an even number of 1s in each of all the N-bit sequences including the BIP-N.

5.2.1.7 *BIP-24: B2* x *3*

Three bytes are allocated in each STM-1 for a bit error monitoring function of a section. This function shall be a Bit Interleaved Parity 24 (BIP-24) code using even parity. The BIP-24 is computed over all bits of the previous STM-1 frame except for the first three rows of section overhead (A1 through D3) and is placed in bytes B2 before scrambling. This parity code is not recomputed at regenerators. These bytes shall be provided in all STM-1 signals within an STM-N signal.

5.2.1.8 *APS channel: K1, K2*

Two bytes are allocated for Automatic Protection Switching (APS) signalling. These bytes are defined only for STM-1 No. 1 of an STM-N signal.

5.2.1.9 *Spare: Z1, Z2*

Six bytes are allocated for functions not yet defined. These bytes have no defined value. These bytes are reserved in all STM-1s of an STM-N.

5.2.2 *AU pointer descriptions*

5.2.2.1 *Pointer value*

Two bytes are allocated for a pointer which indicates the offset in bytes between the pointer and the first byte of the associated virtual container POH. For a complete specification and location of these bytes, see Recommendation G.709.

5.2.2.2 *Pointer action*

Three pointer action bytes are allocated in an AU-4 for frequency justification purposes. One pointer action byte is allocated for AU-3s and TU-*n* s. For a complete specification of these bytes, refer to Recommendation G.709.

In the event of a negative justification, they carry valid information.

5.2.3 *VC-n (n = 3, 4) POH byte descriptions*

5.2.3.1 *Path BIP-8: B3*

One byte is allocated in each virtual container for a path bit error monitoring function. This function shall be a BIP-8 code using even parity. The BIP-8 is computed over all bits of the previous container and is placed in the B3 byte.

5.2.3.2 *Path status: G1*

One byte is allocated to return the VC-*n* path terminating status performance information to the VC-*n* path originating point.

5.2.3.3 *Signal label: C2*

One byte is allocated to indicate the composition of the VC-*n* payload.

5.2.3.4 *VC-* n *path user channel: F2*

One byte is allocated for user communication purposes.

5.2.3.5 *VC-* n *path trace* : *J1*

This byte is used at the VC-*n* termination point to verify the VC-*n* path connection.

5.2.3.6 *Spare: Z3-Z5*

Three bytes are allocated for as yet undefined purposes.

5.2.3.7 *Multiframe indicator: H4*

This byte is allocated to provide a multiframe indication, when required.

6 Physical specification of the NNI

Specification for physical, electrical or optical characteristics of the NNI will be contained in another Recommendation which is under study.

Recommendation G.709

SYNCHRONOUS MULTIPLEXING STRUCTURE

(Melbourne, 1988)

The CCITT,

considering

(a) that Recommendation G.707 describes the advantages offered by a synchronous digital hierarchy and multiplexing method and specifies a set of synchronous digital hierarchy bit rates;

(b) that Recommendation G.708 specifies

— the general principles and frame structure of the network node interface (NNI) for the synchronous digital hierarchy;

— the overall frame size of 9 rows $\times$ 270 columns and section overhead (SOH) definition and its byte allocation;

— arrangements for international synchronous interconnection of STM-1s;

(c) that Recommendations G.707, G.708 and G.709 form a coherent set of specifications for the synchronous digital hierarchy and NNI, *recommends*

that the formats for mapping multiplexing elements into the STM-1 at the Network Node Interface (NNI) and the method of multiplexing to STM-N shall be as described in this Recommendation.

1 Basic multiplexing structure

Descriptions of the various multiplexing elements are given in Recommendation G.708.

The relationships between the various multiplexing elements are shown in Figure 1-1/G.709. The detailed multiplexing structure is described in the following sections.

Figure 1-1/G.709, p.

2 Mapping formats and multiplexing method

2.1 *Mapping and multiplexing up to STM-1*

2.1.1 *Mapping of VC-4 into AU-4*

The STM-1 mapping format for transporting one VC-4 in an AU-4 is shown in Figure 2-1/G.709. The VC-4 consists of a 9-row by 261-column payload structure; the first column of the VC-4 is devoted to path overhead (POH). The payload of the VC-4 shown in Figure 2-1/G.709 is a single C-4. Other possible VC-4 payloads include a single 139 | 64 kbit/s signal in a C-4, four VC-31s (shown in Figure 2-2/G.709 and carried in four TU-31s), three VC-32s (shown in Figure 2-3/G.709 and carried in three TU-32s), and a group of either 21 TUG-21s or 16 TUG-22s (shown in Figure 2-4/G.709).

The STM-1 format shown in Figure 2-1/G.709 consists of an AU-4 plus section overhead (SOH). The VC-4 does not have a fixed phase with respect to the AU-4 (and the STM-1); therefore, the location of the first byte of the VC-4 with respect to the AU-4 frame is given by the AU-4 pointer. Note that the AU-4, including the AU-4 pointer, has a fixed location in the STM-1 frame.

Figure 2-1/G.709, p.

2.1.2 *Mapping of four VC-31s into AU-4*

The STM-1 mapping format for transporting four VC-31s in an AU-4 is shown in Figure 2-2/G.709. Each TU-31 consists of a 9-row by 64-column payload structure plus six bytes of POH plus a three-byte TU-31 pointer. The payload of the VC-31 shown in Figure 2-22/G.709 is a single C-31. Other possible VC-31 payloads include a single 34 | 68 kbit/s signal in a C-31 (shown in Figure 5-10/G.709) or a group of either five TUG-21s or four TUG-22s (shown in Figure 2-5/G.709).

The four VC-31s are carried independently in the 261-column VC-4. Each of the VC-31s does not have a fixed phase with respect to the start of the VC-4. Therefore, the location of the first byte of each VC-31 with respect to the VC-4 POH is given by a 3-byte TU-31 pointer (H1, H2, H3). These four TU-31 pointers reside in a fixed location in the VC-4 as shown in Figure 2-2/G.709.

As described in § 2.1.1, the phase of the VC-4 with respect to the AU-4 is given by the AU-4 pointer.

2.1.3 *Mapping of three VC-32s into AU-4*

The STM-1 mapping format for transporting three VC-32s in an AU-4 is shown in Figure 2-3/G.709. Each TU-32 consists of a 9-row by 84-column payload structure plus one column of POH and one 3-byte TU-32 pointer. The payload of the VC-32 shown in Figure 2-3/G.709 is a single C-32. Other possible VC-32 payloads include a single 44 | 36 kbitB/Fs signal in a C-32 or a group of seven TUG-21s (shown in Figure 2-5/G.709).

The three VC-32s are carried independently in the 261-column VC-4. Each of the VC-32s does not have a fixed phase with respect to the start of the VC-4. Therefore, the location of the first byte of each VC-32 with respect to the VC-4 POH is given by a 3-byte TU-32 pointer (H1, H2, H3). These three TU-32 pointers reside in a fixed location in the VC-4 as shown in Figure 2-3/G.709; 36 fixed stuff bytes are also required in the VC-4.

As described in § 2.1.1, the phase of the VC-4 with respect to the AU-4 is given by the AU-4 pointer.

Figure 2-2/G.709, p.

Figure 2-3/G.709, p.

2.1.4 *Mapping of TUG-2s into AU-4*

The STM-1 mapping format transporting TUG-21s and TUG-22s into an AU-4 is shown in Figure 2-4/G.709. The AU-4 can carry a group of either 21 TUG-21s or 16 TUG-22s.

The TUG-21 payload structure has 9 rows and 12 columns. When used to transport TUG-21s, the VC-4 consists of one column of VC-4 POH, eight columns of fixed stuff, and a remaining 252-column payload structure. The 21 TUG-21s are mapped into this 9-row by 252-column structure using a fixed phase with respect to the VC-4. The TUG-21s are single byte interleaved as shown in Figure 2-4/G.709.

The TUG-22 payload structure has 9 rows and 16 columns. The VC-4 consists of one column of VC-4 POH, four columns of fixed stuff and 256 payload columns when used to carry the 16 TUG-22s. The TUG-22s are single byte interleaved into the 9-row by 256-column structure.

As described in § 2.1.1, the phase of the VC-4 with respect to the AU-4 is given by the AU-4 pointer.

Figure 2-4/G.709, p.

2.1.5 *Mapping of four AU-31s into STM-1*

The STM-1 mapping format for transporting four VC-31s within four AU-31s is shown in Figure 2-6/G.709. A VC-31 is defined to be a 9-row by 64-column payload structure, plus six bytes of POH, located in row 1 to 6 of the first column, according to the figure.

Each AU-31 pointer has a fixed phase with respect to the STM-1 frame. As shown in Figure 2-6/G.709, the four AU-31 pointers are located in columns 11 to 14, rows 1 to 3 of the STM-1, one pointer in each column. Columns 11 to 270 of the STM-1 are divided between each of the AU-31s; thus, each AU-31 occupies alternately every fourth column.

The phase of each VC-31 is not fixed with respect to its AU-31. Therefore, the location of the first byte of each VC-31 with respect to the AU-31 frame is given by AU-31 pointer (H1, H2, H3). The payload of the VC-31 shown in Figure 2-6/G.709 is a single C-31. Other possible VC-31 payloads include a single 34 368 kbit/s signal in a C-31 and a group of five TUG-21s or four TUG-22s (shown in Figure 2-5/G.709).

Figure 2-5/G.709, p.

Figure 2-6/G.709, p.

2.1.6 *Mapping of three AU-32s into STM-1*

The STM-1 mapping format for transporting three VC-32s within three AU-32s is shown in Figure 2-7/G.709. A VC-32 is defined to be a 9-row by 85-column payload structure, with the first column consisting of VC-32 POH. When mapped into its AU-32, two columns of fixed stuff are added to each VC-32 payload to make it equal the AU-32 payload capacity. These two fixed stuff columns are fixed with respect to the VC-32 POH and are inserted between columns 29 and 30, and between columns 57 and 58 of the VC-32.

Each AU-32 pointer has a fixed phase with respect to the STM-1 frame. As shown in Figure 2-7/G.709, the three AU-32 pointers are located in the fourth row of the first nine columns of the STM-1 frame, between the bytes of the SOH. The remaining 261 columns of the STM-1 are divided between each of the AU-32s; thus, each AU-32 occupies alternately every third column of the 261. AU-32 number one consists of three bytes of AU-32 pointer, plus STM-1 columns 10, 13, 16, | | | where columns 1 through 9 contain the SOH and the AU-32 pointers.

The phase of each VC-32 (plus fixed stuff columns) is not fixed with respect to its AU-32. Therefore, the locations of the first byte of each VC-32 with respect to the AU-32 frame is given by the AU-32 pointer (H1, H2, H3). The payload of the VC-32 shown in Figure 2-7/G.709 is a single C-32. Other possible VC-32 payloads include a single 44 | 36 kbitB/Fs signal into a C-32 (shown in Figure 5-8/G.709) and a group of seven TUG-21s (shown in Figure 2-5/G.709).

Figure 2-7/G.709, p.

2.1.7 *Mapping of TUGs into a VC*

Figure 2-5/G.709 shows the schematic mapping of TUG-2s into a VC-3. The details of these mappings are given in § 5; this section presents the general multiplexing principles involved.

The VC-31 consists of six bytes of VC-31 POH plus a 9-row by 64-column payload structure. This payload structure can be used to carry five TUG-21s or four TUG-22s. The individual TUG-2 has a fixed location in the VC-31 frame; this is shown schematically in Figure 2-5/G.709.

The VC-32 consists of nine bytes of VC-32 POH plus a 9-row by 84-column payload structure. This payload structure can be used to carry seven TUG-21s. Again, the individual TUG-21 has a fixed location in the VC-32 frame.

Each TUG-21 can carry a single VC-21 or four VC-11s or three VC-12s. Each TUG-22 can carry a single VC-22 or four VC-12s or five VC-11s. The VCs do not have a fixed phase with respect to the VC-3 POH; TU pointers are used to indicate the position of the VCs in the TUG frame.

2.2 *STM-N multiplexing*

2.2.1 *STM-N frame format*

The STM-N signal is formed by single byte interleaving N STM-1 signals. The STM-N frame structure is depicted in Figure 2-8/G.709.

The first byte of the STM-N signal shall be the first A1 framing byte from STM-1 No. 1 followed sequentially by the first A1 byte from STM-1 No. 2 through No. N. The first bit to be transmitted shall be the most significant bit of the first A1 framing byte from STM-1 No. 1.

Before byte interleaving STM-1 signals to form an STM-N signal, all of the SOH and the AU-*n* (*n* = 3 or 4) pointers in the signals to be interleaved must be 125 µs frame aligned. The alignment is accomplished by adjusting the values of the AU-*n* pointers to reflect the new relative positions of the VC-*n* s.

Note that is is permitted to mix STM-1s containing AU-3s and STM-1s containing AU-4s in the same STM-N.

Figure 2-8/G.709, p.

2.2.2 *STM-N interleaving*

If an STM-N level signal is input to a byte interleaver with STM-M level output (M > N), N bytes of each STM-N are consecutively placed on the output STM-M signal. This method of interleaving is illustrated in Figure 2-9/G.709 where STM-X, STM-Y and STM-Z (X + Y + Z = M) inputs are sequentially interleaved to form an STM-M output.

Figure 2-9/G.709, p.

2.2.3 *Concatenated STM-1s*

STM-1 signals can be concatenated together to form an STM-Nc which can transport payloads requiring greater than one C-4 capacity. A concatenation indication, used to show that this multi-C-4 payload carried in a single VC-4-Nc should be kept together, is contained in the AU-4 pointer. See § 3.4 for details.

2.3 *Maintenance signals*

2.3.1 *Section maintenance signals*

The section alarm indication signal (AIS) is detected as an all 1s in bits 6, 7, 8 of byte K2 after descrambling.

Far end receive failure (FERF) is to return an indication to the transmitting STM-N MUX that the receiving STM-N MUX has detected an incoming section failure or is receiving section AIS.

FERF is detected by a 110 code in bit positions 6, 7 and 8 of the K2 APS byte after descrambling.

2.3.2 *Path maintenance signals*

The VC-*n* (*n* = 3, 4) unequipped indication is an all 0s VC-*n* path signal label after descrambling. This code indicates to VC-*n* terminating equipment that the VC-n is intentionally unoccupied so that alarms can be inhibited. This code is generated as an all 0s VC-*n* path signal label and a valid VC-n path BIP-8 (byte B3); the VC-*n* payload is unspecified.

An alarm indication signal (AIS) is a signal sent downstream as an indication that an upstream failure has been detected and an alarm generated. The TU-*n* (*n* = 1, 2, 3) path AIS is specified as all 1s in the entire TU-*n* , including the TU-*n* pointer. Similarly, the AU-*n* (*n* = 3, 4) path AIS is specified as all 1s in the entire AU-*n* , including the AU-*n* pointer. All path AISs are carried within STM-N signals having valid SOH.

The path status byte (G1) is used to convey the terminating path status and performance to the originator of a VC-*n* (*n* = 3 or 4). Bits 1 through 4 convey the count of errors detected using the path BIP-8 code. This code has nine legal values, 0-8. The remaining seven possible values should be interpreted as zero errors.

2.4 *Timing recovery*

The STM-N (N ≥" 1) signal must have sufficient bit timing content at the NNI. A suitable bit pattern, which prevents a long sequence of 1s and 0s, is provided by using a scrambler. Its operation shall be functionally identical to that of a frame synchronous scrambler of sequence length 127 operating at the line rate.

The generating polynomial shall be $1 + x^6 + x^7$. Figure 2-10/G.709 gives a functional diagram of the frame synchronous scrambler.

The scrambler shall be reset to 1111111 on the most significant bit of the byte following the last byte of the first row of the STM-N SOH. (This is the most significant bit of the $9 \times N + 1$ transmitted byte of the STM-N; see Figure 2-8/G.709.) This bit, and all subsequent bits to be scrambled, shall be added modulo 2 to the output from the x^7 position of the scrambler. The scrambler shall run continuously throughout the complete STM-N frame.

The first row of the STM-N SOH ($9 \times N$ bytes, including the A1 and A2 framing bytes) shall not be scrambled.

Note — Care should be taken in selecting the binary content of the bytes reserved for national use and which are excluded from the scrambling process of the STM-N signal, to ensure that long sequences of 1s or 0s do not occur.

Figure 2-10/G.709, p.

2.5 *Conceptual steps for STM-N assembly*

For a better understanding of the detailed structure of the STM-N frame shown in Figure 2-8/G.709, the conceptual steps required to assemble the STM-N frames in the direct (non-nested) arrangement are listed:

1) Each VC-*n* (*n* = 3 or 4) has either six or nine bytes devoted to path overhead (POH) functions. Of these, the BIP-8 error check byte (B3) is calculated over the entire contents of the VC-*n* and the result is placed in the B3 byte of the following frame.

If it is appropriate, the VC-*n* unequipped signal consisting of an all 0s pattern for the VC-*n* is inserted. (See § 2.3.)

2) After all of the required VC-*n* s have been assembled, AU-*n* pointer values are calculated so as to frame align all of the AU-*n* s within a single STM-N frame.

If the contents of the VC-*n* are lost due to an equipment or other failure, the AU-*n* path AIS signal is inserted into the AU-*n* AU-*n* path AIS is defined in § 2.3.

3) The SOH bytes are then added to the STM-N frame. It is convenient to consider the last five rows of the SOH first. Of the $N \times 45$ such SOH bytes, $N \times 9$ are allocated to the $N \times 3$ B2, $N \times 3$ Z1 and $N \times 3$ Z2 bytes. Thus, each STM-1 has a full complement (3) of these bytes in the STM-N. The remaining STM-N SOH bytes in the last five rows (K1 and K2, D4-D12 and E2) are limited to the first STM-1 in any STM-N signal. The content of the unused SOH bytes of STM-1 No. 2 through No. N are for national use.

4) The $N \times 3$ B2 bytes of an STM-N contain a bit interleaved parity $N \times 24$ (BIP-$N \times 24$) code using even parity which is calculated across the entire previous STM-N frame excluding the first three rows of SOH.

5) A line signal failure would result in the insertion of a section AIS at this point in the assembly of an STM-N (see § 2.3).

6) The remaining bytes of SOH contained in the first three rows ($27 \times$ N bytes) of the STM-N are added next. Of these, the B1, E1, F1, D1-D3 bytes are present only in STM-1 No. 1 of any STM-N signal. The content of the unused SOH bytes of STM-1 No. 2 through No. N are for national use.

7) The STM-1s are then byte interleaved to form an STM-N, as described in § 2.2.2, and subsequently serialized and scrambled as described in § 2.4.

8) The final operation is the calculation of a BIP-8 code over the entire STM-N bit stream on a frame-by-frame basis. The result is loaded into byte B1 of STM-1 No. 1 in the following frame when the SOH is loaded.

3 Pointer

3.1 *AU pointer*

The AU pointer provides a method of allowing flexible and dynamic alignment of the VC within the AU frame.

Dynamic alignment means that the VC is allowed to ‘‘float’’ within the AU frame. Thus the pointer is able to accommodate differences not only in the phases of the VC and SOH, but in the frame rates as well.

3.1.1 *AU pointer location*

The AU-4 pointer is contained in bytes H1, H2 and H3 as shown in Figure 3-1/G.709. The three individual AU-32 pointers are contained in three separate H1, H2 and H3 bytes as shown in Figure 3-2/G.709. Likewise the four individual AU-31 pointers are contained in four separate H1, H2 and H3 bytes as shown in Figure 3-3/G.709.

Figure 3-1/G.709, p.

Figure 3-2/G.709, p.

Figure 3-3/G.709, p.

3.1.2 *AU pointer value*

The pointer contained in H1 and H2 designates the location of the bytes where the VC begins. The two bytes allocated to the pointer function can be viewed as one word as shown in Figure 3-4/G.709. The last 10 bits (bits 7-16) of the pointer word carry the pointer value. The two S bits (bits 5 and 6) indicate the AU type.

As illustrated in Figure 3-4/G.709, the AU-4 pointer value is a binary number with a range of 0 to 782 which indicates the offset between the pointer and the first byte of the VC. As shown in Figure 3-1/G.709, the H1 and H2 bytes contain the pointer value while the position which the pointer indicates is the very first byte of the consecutive three bytes. Figure 3-4/G.709 also indicates two additional valid pointers: the concatenation indication (CI) and the null pointer indication (NPI). The CI is indicated by 1001 in bits 1-4, with bits 5-6 unspecified, and ten 1s in bits 7-16. The NPI is indicated by 1001 in bits 1-4, with bits 5-6 unspecified, and five 1s in bits 7-11 followed by five 0s in bits 12-16.

As illustrated in Figure 3-4/G.709, the AU-32 pointer value is also a binary number with a range of 0 to 782. Since there are three AU-32s in the STM-1, each AU-32 has its own associated H1, H2 and H3 bytes. In Figure 3-2/G.709, the H bytes are shown in sequence. The first H1, H2, H3 set refers to the first AU-32, and the second set to the second AU-32, and so on. The same is true for the information bytes. For the AU-32s, each pointer operates independently.

Likewise, as illustrated in Figure 3-4/G.709, the AU-31 pointer value is a binary number with a range of 0 to 581. Since there are four AU-31s in the STM-1, each AU-31 has its own associated H1, H2 and H3 bytes. In Figure 3-3/G.709, the H bytes are shown in sequence. The first H1, H2, H3 set refers to the first AU-31, the second set to the second AU-31, and so on. The same is true for the information bytes. For the AU-31s, each pointer operates independently.

In all cases, the STM-1 SOH and AU pointer bytes are not counted in the offset. For example, in an AU-4, the pointer value of 0 indicates that the VC starts in the byte location that immediately follows the last H3 byte, whereas an offset of 87 indicates that the VC starts three bytes after the K2 byte.

Figure 3-4/G.709, p. 28

Figure 3-4/G.709 [T2.709], p. 28

3.1.3 *Frequency justification*

If there is a frequency offset between the frame rate of the SOH and that of the VC, then the pointer value will be incremented or decremented as needed, accompanied by a corresponding positive or negative justification byte or bytes. Consecutive pointer operations must be separated by at least three frames (i.e. every fourth frame) in which the pointer value remains constant.

If the frame rate of the VC is too slow with respect to that of the SOH, then the alignment of the VC must periodically slip back in time and the pointer value must be incremented by one. This operation is indicated by inverting bits 7, 9, 11, 13 and 15 (I-bits) of the pointer word to allow 5-bit majority voting at the receiver. Three positive justification bytes appear immediately after the last H3 byte in the AU-4 frame containing inverted I-bits. Subsequent pointers will contain the new offset. This is illustrated in Figure 3-5/G.709.

For AU-32 frames, a positive justification byte appears immediately after the associated H3 byte of the individual AU-32 frame containing inverted I-bits. Subsequent pointers will contain the new offset. This is illustrated in Figure 3-6/G.709. The same is true for AU-31 as shown in Figure 3-7/G.709.

If the frame rate of the VC is too fast with respect to that of the SOH, then the alignment of the VC must periodically be advanced in time and the pointer value must be decremented by one. This operation is indicated by inverting bits 8, 10, 12, 14 and 16 (D-bits) of the pointer word to allow 5-bit majority voting at the receiver. Three negative justification bytes appear in the H3 bytes in the AU-4 frame containing inverted D-bits. Subsequent pointers will contain the new offset. This is illustrated in Figure 3-8/G.709.

For AU-32 frames, a negative justification byte appears in the H3 byte of the individual AU-32 frame containing inverted D-bits. Subsequent pointers will contain the new offset. This is illustrated in Figure 3-9/G.709. The same is true for AU-31 as shown in Figure 3-10/G.709.

Figure 3-5/G.709, p.

Figure 3-6/G.709, p.

Figure 3-7/G.709, p.

Figure 3-8/G.709, p.

Figure 3-9/G.709, p.

Figure 3-10/G.709, p.

3.1.4 *New data flag*

Bits 1-4 (N-bits) of the pointer word carry a new data flag (NDF) which allows an arbitrary change of the pointer value if that change is due to a change in the payload.

Four bits are allocated to the flag to allow error correction. The decoding may be performed by accepting NDF enabled if at least three bits match. Normal operation is indicated by a 0110 code in the N-bits. NDF is indicated by inversion of the N-bits to 1001. The new alignment is indicated by the pointer value accompanying the NDF and takes effect at the offset indicated. NDF should be enabled when the pointer value transits between its normal value and the CI or NPI.

3.1.5 *Pointer generation*

The following summarizes the rules for generating the AU pointers.

1) During normal operation, the pointer locates the start of the VC within the AU frame. The NDF is set to 0110.

2) The pointer value can only be changed by rules 3, 4 or 5.

3) If a positive justification is required, the current pointer value is sent with the I-bits inverted and the subsequent positive justification opportunity is filled with dummy information. Subsequent pointers contain the previous pointer value incremented by one. No subsequent increment or decrement operation is allowed for at least three frames following this operation.

4) If a negative justification is required, the current pointer value is sent with the D-bits inverted and the subsequent negative justification opportunity is overwritten with actual data. Subsequent pointers contain the previous pointer value decremented by one. No subsequent increment or decrement operation is allowed for at least three frames following this operation.

5) If the alignment of the VC changes for any reason other than rules 3 or 4, the new pointer value shall be sent accompanied by NDF set to 1001. The NDF only appears in the first frame that contains the new values. The new location of the VC begins at the first occurrence of the offset indicated by the new pointer. No subsequent increment or decrement operation is allowed for at least three frames following this operation.

3.1.6 *Pointer interpretation*

The following summarizes the rules for interpreting the AU pointers.

1) During normal operation, the pointer locates the start of the VC within the AU frame.

2) Any variation from the current pointer value is ignored unless a consistent new value is received three times consecutively or it is preceded by one of rules 3, 4 or 5.

3) If the majority of the I-bits of the pointer word are inverted, a positive justification operation is indicated. Subsequent pointer values shall be incremented by one.

4) If the majority of the D-bits of the pointer word are inverted, a negative justification operation is indicated. Subsequent pointer values shall be decremented by one.

5) If the NDF is set to 1001, then the coincident pointer value shall replace the current one at the offset indicated by the new pointer value regardless of the state of the receiver.

3.2 *TU-3 pointers*

There are two types of TU-3 pointers: TU-31 and TU-32. The TU-3 pointer provides a method of allowing flexible and dynamic alignment of VC-3 within the TU-3 frame, independent of the actual contents of the VC. Dynamic alignment means that the VC-3 is allowed to ‘‘float’’ within the TU-3 frame.

3.2.1 *TU-3 pointer location*

Three individual TU-32 pointers are contained in the three separate H1, H2 and H3 bytes as shown in Figure 3-11/G.709. Four individual TU-31 pointers are contained in the four separate H1, H2 and H3 bytes as shown in Figure 3-12/G.709.

3.2.2 *TU-3 pointer value*

The TU-3 pointer value contained in H1 and H2 designates the location of the byte where the VC-3 begins. The two bytes allocated to the pointer function can be viewed as one word as shown in Figure 3-4/G.709. The last ten bits (bits 7-16) of the pointer word carry the pointer value. The two S bits (bits 5 and 6) indicate the TU type.

The TU-32 pointer value is a binary number with a range of 0-764 which indicates the offset between the pointer and the first byte of the VC-32 as shown in Figure 3-11/G.709.

The TU-31 pointer value is a binary number with a range of 0-581 which indicates the offset between the pointer and the first byte of the VC-31 as shown in Figure 3-12/G.709.

Figure 3-11/G.709, p.

Figure 3-12/G.709, p.

3.2.3 *Frequency justification*

If there is a frequency offset between the TU-3 frame rate and that of the VC-3, then the pointer value will be incremented or decremented as needed accompanied by a corresponding positive or negative justification byte. Consecutive pointer operations must be separated by at least three frames in which the pointer value remains constant.

If the frame rate of the VC-3 is too slow with respect to that of the TU-3 frame rate, then the alignment of the VC must periodically slip back in time and the pointer must be incremented by one. This operation is indicated by inverting bits 7, 9, 11, 13 and 15 (I-bits) of the pointer word to allow 5-bit majority voting at the receiver. A positive justification byte appears immediately after the individual H3 byte in the TU-3 frame containing inverted I-bits. Subsequent TU-3 pointers will contain the new offset.

If the frame rate of the VC-3 is too fast with respect to that of the TU-3 frame rate, then the alignment of the VC must be periodically advanced in time and the pointer must be decremented by one. This operation is indicated by inverting bits 8, 10, 12, 14 and 16 (D-bits) of the pointer word to allow 5-bit majority voting at the receiver. A negative justification byte appears in the individual H3 byte in the TU-3 frame containing inverted D-bits. Subsequent TU-3 pointers will contain the new offset.

3.2.4 *New data flag*

Bits 1-4 (N-bits) of the pointer word carry a NDF which allows an arbitrary change of the value of the pointer if that change is due to a change in the VC-3.

Four bits are allocated to the flag to allow for error correction. The decoding may be performed by accepting NDF enabled if at least three bits match. Normal operation is indicted by a 0110 code in the N-bits; NDF is indicated by inversion of the N-bits to 1001. The new alignment is indicated by the pointer value accompanying the NDF and takes effect at the offset indicated.

3.2.5 *Pointer generation*

The following summarizes the rules for generating the TU-3 pointers.

1) During normal operation, the pointer locates the start of the VC-3 within the TU-3 frame. The NDF is set to 0110.

2) The pointer value can only be changed by rules 3, 4 or 5.

3) If a positive justification is required, the current pointer value is sent with the I-bits inverted and the subsequent positive justification opportunity is filled with dummy information. Subsequent pointers contain the previous pointer value incremented by one. No subsequent increment or decrement operation is allowed for at least three frames following this operation.

4) If a negative justification is required, the current pointer value is sent with the D-bits inverted and the subsequent negative justification opportunity is overwritten with actual data. Subsequent pointers contain the previous pointer value decremented by one. No subsequent increment or decrement operation is allowed for at least three frames following this operation.

5) If the alignment of the VC changes for any reason other than rules 3 or 4, the new pointer value shall be sent accompanied by the NDF set to 1001. The NDF only appears in the first frame that contains the new value. The new VC location begins at the first occurrence of the offset indicated by the new pointer. No subsequent increment or decrement operation is allowed for at least three frames following this operation.

3.2.6 *Pointer interpretation*

The following summarizes the rules for interpreting the TU-3 pointers.

1) During normal operation the pointer locates the start of the VC-3 within the TU-3 frame.

2) Any variation from the current pointer value is ignored unless a consistent new value is received three times consecutively or it is preceded by one of rules 3, 4 or 5.

3) If the majority of the I-bits of the pointer word are inverted, a positive justification is indicated. Subsequent pointer values shall be incremented by one.

4) If the majority of the D-bits of the pointer word are inverted, a negative justification is indicated. Subsequent pointer values shall be decremented by one.

5) If the NDF is set to 1001, then the coincident pointer value shall replace the current one at the offset indicated by the new pointer value regardless of the state of the receiver.

3.3 *TU-1/TU-2 pointers*

The TU-1 pointer is only used with floating mapping. Floating and locked modes of operation are described in § 5.2.

The TU-1 and TU-2 pointers provide a method of allowing flexible and dynamic alignment of the VC-1/VC-2 within the TU-1 and TU-2 multiframes, independent of the actual contents of the VC.

3.3.1 *TU-1/TU-2 pointer location*

The TU-1/TU-2 pointers are contained in the V1 and V2 bytes as illustrated in Figure 3-13/G.709.

Figure 3-13/G.709, p.

3.3.2 *TU-1/TU-2 pointer value*

The TU pointer word is shown in Figure 3-14/G.709.

The pointer value (bits 7-16) is a binary number which indicates the offset from V2 to the first byte of the VC-1/VC-2. The range of the offset is different for each of the TU sizes as illustrated in Figure-3-15/G.709. Note that the pointer bytes are not counted in the offset calculation.

Figure 3-14/G.709, p.

Figure 2-15/G.709, p.

3.3.3 *TU-1/TU-2 multiframe indication byte*

TU-1/TU-2 multiframe indication byte (H4) relates to the lowest level of multiplexing structure and indicates a variety of different multiframes for use by certain payloads. Specifically it provides:

— 500 µs (4-frame) multiframe identifying frames containing TU-1/TU-2 pointers in the floating TU-1/TU-2 mode, and reserved byte locations in the locked TU-1 mode;

— 2 ms (16-frame) multiframe for byte synchronous out-slot-signalling for 2048 kbitB/Fs payloads in the locked TU-1 mode;

— 3 ms (24-frame) multiframe for byte synchronous out-slot-signalling for 1544 kbitB/Fs payloads in the locked TU-1 mode.

The coding of the H4 byte is illustrated in Figures 3-16/G.709 to 3-18/G.709.

Figure 3-16/G.709, p. 40

Figure 3-17/G.709 [T3.709], p. 41

Figure 3-18/G.709 [T4.709], p. 42

For network elements that operate only in the floating TU-1/TU-2 mode, a simplified multiframe alignment byte may be used. The simplified version provides only the 500 µs multiframe. The 2 or 3 ms multiframe of any signalling within floating TU-1s is indicated by per-TU multiframe indicators carried within the TU-1. Figure 3-13/G.709 shows the VC-1/VC-2 mapping in the multiframed TU-1/TU-2.

A converter from locked to floating TUs is permitted to pass H4 through transparently. A converter from floating to locked TUs must recover and align the multiframes from all of the floating TUs, and thus can transmit any convenient full multiframe on the locked TU side.

3.3.4 *TU-1/TU-2 frequency justification*

The TU-1/TU-2 pointer is used to frequency justify the VC-1/VC-2 exactly the same way that the TU-3 pointer is used to frequency justify the VC-3. A positive justification opportunity immediately follows the V3 byte.

Additionally, V3 serves as the negative justification opportunity such that when the opportunity is taken, V3 is overwritten by data. This is also shown in Figure 3-15/G.709. The indication of whether or not a justification opportunity has been taken is provided by the I- and D-bits of the pointer in the current TU multiframe. The value contained in V3 when not being used for a negative justification is not defined. The receiver is required to ignore the value contained in V3 whenever it is not used as negative justification.

3.3.5 *TU-1/TU-2 sizes*

Bits 5 and 6 of TU-1/TU-2 pointer indicate the size of the TU. Four sizes are currently provided as shown in Table 3-1/G.709.

H.T. [T1.709]

TABLE 3-1/G.709

Size (binary) TU pointer range (in 500 µs) }	Designation	{
01	TU-22	0 à 571
00	TU-21	0 à 427
10	TU-12	0 à 139
11	TU-11	0 à 103

Table 3-1/G.709 [T1.709], p.

Note that this technique is only used at the TU-1/TU-2 levels.

3.3.6 *New data flag (NDF)*

Bits 1-4 (N-bits) of the pointer word carry an NDF. It is the mechanism which allows an arbitrary change of the value of a pointer, and possibly also the size of the TU, if that change is due to a change in the payload. If the change includes a change in size then, implicitly, there must be a simultaneous new data transition in all of the TUs in the TUG-21.

As with the TU-3 pointer NDF, the normal value is 0110 (transmitted), and the value 1001 (received exactly) indicates a new alignment for the VC, and possibly a new size. If a new size is indicated, then all TU pointers (1 to 4) in the TUG-21 must simultaneously indicate NDF with the same new size. The new alignment, and possibly size, is indicated by the pointer value and size value accompanying the NDF and takes effect at the offset indicated. The NDF should be enabled when the pointer value transits between its normal value and the concatenation indication (CI).

3.3.7 *TU concatenation*

TU-2s may be concatenated to form a TU-2-mc (concatenated m × TU-2s) to carry payloads requiring a capacity of more than a C-21 (for the TU-21 case) or C-22 (for the TU-22 case). A CI (1001 in bits 1-4, bits 5-6 unspecified, and all 1s in bits 7-16 of the TU-2 pointer) is used to show that this multi-C-2 payload, carried in a single VC-2-mc (concatenated m × VC-2), must be kept together.

Note that the TU-2 is carried in a TUG-2 as shown in Figure 5-4/G.709 and Figure 5-5/G.709.

If a TU-2 pointer contains the concatenation indication, then the pointer processor determines that this TU-2 is concatenated to the previous TU-2, and all operations indicated by the previous TU-2 pointer are to be performed on this TU-2 as well.

3.3.8 *TU pointer generation and interpretation*

The rules for generating and interpreting the TU-1/TU-2 pointer for the VC-1/VC-2 are an extension to the rules provided in §§ 3.2.5 and 3.2.6 for the TU-3 pointer with the following modifications:

1) The term TU-3 is replaced with TU-1/TU-2 and the term VC-3 is replaced with VC-1/VC-2.

2) Additional pointer generation rule 6: If the size of the TU within a TUG-21 is to change, then an NDF, as described in rule 5, is to be sent in all TUs of the new size in the group simultaneously.

3) Additional pointer interpretation rule 6: If an NDF of 1001 and an arbitrary new size of TU are received simultaneously in all of the TUs within a TUG-21, then the coincident pointers and sizes shall replace the current ones immediately.

3.4 *Pointer operation for STM-1 concatenation*

A concatenation indication contained in the AU-4 pointer is used to show that the STM-1 is part of an STM-Nc.

The AU-4 within the first STM-1 of an STM-Nc shall have a normal range of pointer values. All subsequent AU-4s within the grouped STM-Nc shall have their pointer values set to 1001 in bits 1-4, bits 5-6 unspecified, and all 1s in bits 7-16. Since this value does not indicate a valid offset, the pointer processors shall interpret this value to mean that they shall perform the same operations as performed on the first AU-4 of the grouped STM-Nc. The NDF must be set when changing a pointer to/from the concatenation value.

3.4.1 *Pointer generation*

The following additional pointer generation rule shall apply for AU-4 pointers:

If an STM-Nc signal is being transmitted, a pointer is generated for the AU-4 within the first STM-1 only. The concatenation indication is generated in place of the other pointers; all operations indicated by the AU-4 pointer in the first STM-1 apply to each STM-1 in the STM-Nc.

3.4.2 *Pointer interpretation*

The following additional pointer interpretation rule shall apply for AU-4 pointers:

If the pointer contains the concatenation indication, then the operations performed on the STM-1 are identical to those performed on the first STM-1 within the STM-Nc. Rules 3 and 4 of § 3.1.6 do not apply to this pointer.

4 Path overhead

4.1 *VC-1/VC-2 path overhead*

The first byte in the VC-1/VC-2 pointed to by the TU-1/TU-2 pointer is the VC-1/VC-2 path overhead byte. This byte is designated as V5.

This byte provides the functions of error checking, signal, label and path status of the VC-1/VC-2 paths. The bit assignments of the VC-1/VC-2 POH are specified in the following paragraphs and are illustrated in Figure 4-1/G.709.

V5 is used only in floating mode VC-1/VC-2s and is designated as an R-byte in locked mode VC-1/VC-2s. Floating mode and locked mode operation is described in § 5.8.

Bits 1 and 2 are used for error performance monitoring. A bit interleaved parity (BIP) scheme is specified. Bit 1 is set such that parity of all odd number bits (1, 3, 5 and 7) in all bytes in the previous VC-1/VC-2 is even and bit 2 is set similarly for the even number bits (2, 4, 6 and 8). Note that the calculation of the BIP-2 includes the VC-1/VC-2 POH bytes but excludes the TU-1/TU-2 pointers.

Bit 3 is a VC-1B/FVC-2 path far-end-block-error (FEBE) indication that is set to 1 and sent back towards a VC-1/VC-2 path originator if one or more errors were detected by the BIP-2, and is otherwise set to 0.

Bit 4 is unused (X). The receiver is required to ignore the value of this bit.

Bits 5 through 7 provide a VC-1/VC-2 signal label. Eight binary values are possible in these three bits. Value 0 indicates ‘‘VC-1/VC-2 path unequipped’’, and value 1 indicates ‘‘VC-1/VC-2 path equipped — non-specific payload’’. The remaining six values are reserved to be defined as required in specific VC-1/VC-2 mappings. Any value received, other than 0, indicates an equipped VC-1/VC-2 path.

Bit 8 is a VC-1/VC-2 path remote alarm indication. This bit is set to a 1 if either a TU-1/TU-2 path alarm indication signal (AIS) or a signal failure condition is being received, otherwise it is set to 0. The VC-1/VC-2 path remote alarm indication is sent back by the VC-1/VC-2 assembler.

Figure 4-1/G.709, p.

4.2 *VC-3/VC-4 path overhead*

The VC-3/VC-4 POH will be assigned to and remain with the payload until the payload is demultiplexed and will be used for functions that are necessary in transporting all VC-3/VC-4. Note that this does not preclude the allocation of other overhead in specific mappings (such as justification control for mapping asynchronous 44 | 36 kbit/s signals). That type of overhead is payload specific whereas the POH defined in this section is payload independent.

The VC-4/VC-32 POH consists of nine bytes denoted J1, B2, C2, G1, F2, H4, Z1-Z3. The VC-31 POH consists of six bytes denoted J1, B3, C2, G1, G2, H4.

4.2.1 *VC-3/VC-4 path trace (J1)*

This is the first byte in the VC; its location is indicated by the associated AU/TU pointer. This byte is used to repetitively transmit a 64 byte, fixed length string so that a path receiving terminal can verify its continued connection to the intended transmitter. The content of the message is not constrained by this standard since it is assumed to be user programmable at both transmit and receive ends.

4.2.2 *Path BIP-8 (B3)*

One byte is allocated in each VC-3 or VC-4 for a path error monitoring function. This function shall be a BIP-8 code using even parity. The path BIP-8 is calculated over all bits of the previous VC-3 or VC-4 before scrambling. The computed BIP-8 is placed in the B3 byte of the VC-3 or VC-4 before scrambling.

4.2.3 *Signal label (C2)*

One byte is allocated to indicate the composition of the VC-3/VC-4. Of the 256 possible binary values, two are defined here and the remaining 254 values are reserved to be defined as required in specific VC-3/VC-4 mappings.

— Value 0 indicates ‘‘VC-3/VC-4 path unequipped’’. This value shall be originated if the section is complete but there is no VC-3/VC-4 path originating equipment.

— Value 1 indicates ‘‘VC-3/VC-4 path equipped — non-specific payload’’. This value can be used for all payloads that need no further differentiation, or that achieve differentiation by other means such as messages from an operations system.

Note that any value received, other than value 0, constitutes an ’’equipped’’ condition.

4.2.4 *Path status (G1)*

One byte is allocated to convey back to a VC-3/VC-4 path originator the path terminating status and performance. This feature permits the status and performance of the complete duplex path to be monitored at either end, or at any point along that path. As illustrated in Figure 4-2/G.709, bits 1 through 4 convey the count of interleaved-bit blocks that have been detected in error by the path BIP-8 code (B3). This count has nine legal values, namely 0-8 errors. The remaining seven possible values represented by these four bits can only result from some unrelated condition and shall be interpreted as zero errors. VC-3/VC-4 path remote alarm indication is sent back by the VC-3/VC-4 assembler whenever the VC-3/VC-4 assembler is not receiving a valid signal. The VC-3/VC-4 path remote alarm indication is bit 5, which is set to one to indicate VC-3/VC-4 path remote alarm and is otherwise set to zero. The specific received conditions under which VC-3/VC-4 path remote alarm is initiated are path AIS, signal failure conditions or path tracer mismatch. Bits 6, 7 and 8 are not used.

4.2.5 *Path user channel (F2)*

One byte is allocated for user communication purposes between path elements.

4.2.6 *Multiframe indicator (H4)*

This byte provides a generalized multiframe indicator for payloads. Currently, this indicator is only used for TUG-structured payloads as described in § 3.3.3.

4.2.7 *Spare (Z3-Z5)*

Three bytes are allocated for future, as yet undefined, purposes. These bytes have no defined value. The receiver is required to ignore the value contained in these bytes.

Figure 4-2/G.709, p.

5 Mapping of tributaries into VCs

Accommodation of asynchronous and synchronous tributaries presently defined in Recommendation G.702 shall be possible. At the TU-1/TU-2 level, asynchronous accommodation utilizes only the floating mode, whereas synchronous accommodation utilizes both the locked and the floating mode.

Figure 5-1/G.709 shows TU-1 and TU-2 sizes and formats.

5.1 *Mapping of tributaries into VC-4*

5.1.1 *Asynchronous 139 | 64 kbit/s*

One 139 | 64 kbit/s signal can be mapped into a VC-4 container of an STM-1 frame as shown in Figures 5-2/G.709 and 5-3/G.709.

The VC-4 container consists of nine bytes (1 column) path overhead (POH) plus a 9 row by 260 column payload structure as shown in Figure 5-2/G.709.

This payload can be used to carry one 139 | 64 kbit/s signal:

— Each of the nine rows is partitioned into 20 blocks, consisting of 13 bytes each (Figure 5-2/G.709).

— In each row one justification opportunity (S) bit and five justification control (C) bits are provided (Figure 5-3/G.709).

— The first byte of one block consists of:

i) eight information (I) bits (byte W), or

ii) eight fixed stuff (R) bits (byte Y), or

iii) one justification control (C) bits, plus five fixed fixed stuff (R) bits, plus two overhead (O) bits (byte X), or

iv) six information (I) bits, plus one justification opportunity (S) bit, plus one fixed stuff (R) bit, (byte Z).

— The last 12 bytes of one block consists of information bits (I).

The sequence of all these bytes is shown in Figure 5-3/G.709.

The overhead (O) bits are reserved for further overhead communication purposes.

The set of five justification control (C) bits in every row is used to control the corresponding justification opportunity (S) bit. C C C C C = 0 0 0 0 0 indicates that the S bit is an information bit, whereas C C C C C = 1 1 1 1 1 indicates that the S bit is a justification bit. Majority vote should be used to make the justification decision in the desynchronizer for protection against single and double bit errors in the C bits.

The value contained in the S bit when used as justification bit is not defined. The receiver is required to ignore the value contained in this bit whenever it is used as a justification bit.

Figura 5-1/G.709, p.

Figura 5-2/G.709, p.

Figura 5-3/G.709, p.

5.1.2 *TUG-22*

Sixteen TUG-22s can be mapped into a VC-4. This is illustrated in three-dimensional form in a) of Figure 5-4/G.709 and in linear form in b) of Figure 5-4/G.709.

Figure 5-4/G.709, p.

5.1.3 *TUG-21*

Twenty-one TUG-21s can be mapped into a VC-4. This is shown in three-dimensional form in a) of Figure 5-5/G.709 and in linear form in b) of Figure 5-5/G.709.

Figure 5-5/G.709, p. 50

5.1.4 *TU-32*

Three TU-32s can be mapped into a VC-4. This is illustrated in Figure 5-6/G.709.

Figure 5-6/G.709, p.

5.1.5 *TU-31*

Four TU-31s can be mapped into a VC-4. This is illustrated in Figure 5-7/G.709.

Figure 5-7/G.709, p.

5.2 *Mapping of tributaries into VC-32*

5.2.1 *Asynchronous 44 | 36 kbit/s*

One 44 | 36 kbit/s signal can be mapped into a VC-32, as shown in Figure 5-8/G.709.

The VC-32 consists of nine subframes every 125 µs. Each subframe consists of one byte of VC-3 POH, 621 data bits, a set of five justification control bits, one justification opportunity bit and two overhead communication channel bits. The remaining bits are fixed stuff (R) bits. The O bits are reserved for future overhead communication purposes.

The set of five justification control (C) bits is used to control the justification opportunity (S) bit. C C C C C = 0 0 0 0 0 indicates that the S bit is a data bit, whereas C C C C C = 1 1 1 1 1 indicates that S bit is a justification bit. Majority vote should be used to make the justification decision in the desynchronizer for protection against single and double bit errors in the C bits.

The value contained in the S bit when used as justification bits is not defined. The receiver is required to ignore the value contained in this bit whenever it is used as a justification bit.

Figure 5-8/709, p.

5.2.2 *TUG-21*

Seven TUG-21s can be mapped into a VC-32. This is illustrated in Figure 5-9/G.709. The figure also illustrates the formation of the TUG-21 from TU-11, TU-12 and TU-21.

Figure 5-9/G.709, p.

5.3 *Mapping of tributaries into VC-31*

5.3.1 *Asynchronous 34 | 68 kbit/s*

One 34 | 68 kbit/s signal can be mapped into a VC-31 as shown in Figure 5-10/G.709.

In addition to the VC-31 POH, the VC-31 consists of a payload of 9 × 64 bytes every 125 µs. This payload is divided in three subframes, each subframe divided in 12 sectors and consisting of:

— 1431 information (I) bits.

— two sets of five justification control bits (C_1, C_2).

— two justification opportunity bits (S_1, S_2).

— 93 fixed stuff bits (R).

Two sets (C_1, C_2) of five justification control bits are used to control the two justification opportunity bits S_1 and S_2 respectively.

$C_1C_1C_1C_1C_1$ = 0 0 0 0 0 indicates that S_1 is a data bit while $C_1C_1C_1C_1C_1$ = 1 1 1 1 1 indicates that S_1 is a justification bit. C_2 bits control S_2 in the same way. Majority vote should be used to make the justification decision in the desynchronizer for protection against single and double bit errors in the C bits.

The value contained in S_1 and S_2 when they are justification bits is not defined. The receiver is required to ignore the value contained in these bits whenever they are used as justification bits.

Note — The same mapping could be used for bit or byte synchronous 34 | 68 kbit/s. In these cases, S_1 bit should be a fixed stuff and the S_2 bit an information bit. By setting the C_1 bits to 1 and the C_2 bits to 0, a common desynchronizer could be used for both asynchronous and synchronous 34 | 68 kbit/s.

Figure 5-10/G.709, p.

5.3.2 *TUG-22*

Four TUG-22s can be mapped into a VC-31. This is illustrated in Figure 5-11/G.709. The figure also illustrates the formation of the TUG-22 from TU-11, TU-12 and TU-22.

5.3.3 *TUG-21*

Five TUG-21s can be mapped into a VC-31. This is illustrated in Figure 5-12/G.709. The figure also illustrates the formation of the TUG-21 from TU-11, TU-12 and TU-21.

5.4 *Mapping of tributaries into VC-22*

5.4.1 *Asynchronous 8448 kbit/s*

One 8448 kbit/s signal can be mapped into a VC-22. Figure 5-13/G.709 shows this over a period of 500 µs.

In addition to the VC-22 POH, the VC-22 consists of:

— 4220 information (I) bits.

— 24 justification control bits (C_1, C_2).

— eight justification opportunity bits (S_1, S_2).

— 316 fixed stuff (R) bits.

Figure 5-11/G.709, p. 56

Figure 5-12/G.709, p. 57

Two sets (C_1, C_2) of three justification control bits are used to control the two justification opportunity bits S_1 and S_2 respectively.

$C_1C_1C_1$ = 0 0 0 indicates that S_1 is a data bit while $C_1C_1C_1$ = 1 1 1 indicates that S1 is a justification bit. C_2 bits control S_2 in the same way. Majority vote should be used to make the justification decision in the desynchronizer for protection against single bit error in the C bits.

The value contained in S_1 and S_2 when they are justification bits is not defined. The receiver is required to ignore the value contained in these bits whenever they are used as justification bits.

Figure 5-13/G.709, p.

5.4.2 *Synchronous 8448 kbit/s*

One bit or byte synchronous 8448 kbitB/Fs signal can be mapped into a VC-22. Figure 5-14/G.709 shows this over a period of 500 µs.

Note — A common desynchronizer can be used for both asynchronous and synchronous mappings.

Figure 5-14/G.709, p.

5.5 *Mapping of tributaries into VC-21*

5.5.1 *Asynchronous 6312 kbit/s*

One 6312 kbit/s signal can be mapped into a VC-21. Figure 5-15/G.709 shows this over a period of 500 µs.

In addition to the VC-2 POH, the VC-21 consists of 3152 data bits, 24 justification control bits, eight justification opportunity bits and 32 overhead communication channel bits. The remaining bits are fixed stuff (R). The O bits are reserved for future overhead communication purposes.

Two sets (C_1, C_2) of three justification control bits are used to control the two justification opportunities S_1 and S_2 respectively. $C_1C_1C_1$ = 0 0 0 indicates that S_1 is a data bit while

$C_1C_1C_1$ = 1 1 1 indicates that S_1 is a justification bit. C_2 controls S_2 in the same way. Majority vote should be used to make the justification decision in the desynchronizer for protection against single bit errors in the C bits.

The value contained in S_1 and S_2 when they are justification bits is not defined. The receiver is required to ignore the value contained in these bits whenever they are used as justification bits.

5.5.2 *Bit synchronous 6312 kbit/s*

The bit synchronous mapping for 6312 kbit/s tributary is shown in Figure 5-16/G.709.

Note that a common desynchronizer can be used for both asynchronous and bit synchronous mapping.

Figure 5-15/G.709, p.

Figure 5-16/G.709, p.

5.5.3 *Byte synchronous 6312 kbit/s*

Under study.

5.6 *Mapping of tributaries into VC-12*

5.6.1 *Asynchronous 2048 kbit/s*

One 2048 kbit/s signal can be mapped into a VC-12. Figure 5-17/G.709 shows this over a period of 500 µs.

In addition to the VC-1 POH, the VC-12 consists of 1023 data bits, six justification control bits, two justification bits and eight overhead communication channel bits. The remaining bits are fixed stuff (R) bits. The O bits are reserved for future overhead communication purposes.

Two sets (C_1, C_2) of three justification control bits are used to control the two justification opportunities S_1 and S_2 respectively. $C_1C_1C_1$ = 0 0 0 indicates that S_1 is a data bit while $C_1C_1C_1$ = 1 1 1 indicates that S_1 is a justification bit. C_2 controls S_2 in the same way. Majority vote should be used to make the justification decision in the desynchronizer for protection against single bit errors in the C bits.

The value contained in S_1 and S_2 when they are justification bits is not defined. The receiver is required to ignore the value contained in these bits whenever they are used as justification bits.

Figure 5-17/G.709, p.

5.6.2 *Bit synchronous 2048 kbit/s*

The bit synchronous mapping for 2048 kbit/s tributaries is shown in Figure 5-18/G.709.

Note that a common desynchronizer can be used for both asynchronous and bit synchronous mappings.

Figure 5-18/G.709, p.

5.6.3 *Byte synchronous mapping for 2048 kbit/s*

Figure 5-19/G.709 shows byte synchronous mapping for 30-channel 2048 kbit/s tributaries employing Channel Associated Signalling (CAS). Signalling is carried in byte 19. The signalling assignments are shown in Figure 5-20/G.709.

The S_1, S_2, S_3 and S_4 bits contain the signalling for the 30×64 kbit/s channels. The phase of the signalling bits is indicated in the P_1 and P_0 bits in floating TU mode, and in the multiframe indicator byte (H4) in locked TU mode. This is illustrated in Figure 5-20/G.709.

Byte synchronous mapping of 31 channel tributaries is shown in Figure 5-21/G.709. Byte 19 carries tributary channel 16.

Figure 5-19/G.709, p.

Figure 5-20/G.709, p.

Figure 5-21/G.709, p.

5.7 *Mapping of tributaries into VC-11*

5.7.1 *Asynchronous 1544 kbit/s*

One 1544 kbit/s signal can be mapped into a VC-11. Figure 5-22/G.709 shows this over a period of 500 µs.

In addition to the VC-1 POH, the VC-11 consists of 771 data bits, six justification control bits, two justification opportunity bits and eight overhead communication channel bits. The remaining bits are fixed stuff (R) bits. The eight O bits are reserved for future communication purposes.

Two sets (C_1, C_2) of three justification control bits are used to control the two justification opportunities, S_1 and S_2 respectively. $C_1C_1C_1$ = 0 0 0 indicates that S_1 is a data bit while $C_1C_1C_1$ = 1 1 1 indicates that S_1 is a justification bit. C_2 controls S_2 in the same way. Majority vote should be used to make the justification decision in the desynchronizer for protection against single bit errors in the C bits.

The value contained in S_1 and S_2 when they are justification bits is not defined. The receiver is required to ignore the value contained in these bits whenever they are used as justification bits.

Figure 5-22/G.709, p.

5.7.2 *Bit synchronous 1544 kbit/s*

The bit synchronous mapping for 1544 kbit/s tributaries is shown in Figure 5-23/G.709.

Note that a common desynchronizer can be used for both asynchronous and bit synchronous mappings.

Figure 5-23/G.709, p.

5.7.3 *Byte synchronous mapping for 1544 kbit/s*

The byte synchronous mapping for 1544 kbit/s is depicted in Figure 5-24/G.709.

The S_1, S_2, S_3 and S_4 bits contain the signalling for the 24 × 64 kbit/s channels. The phase of the signalling bits can be indicated in the P_1 and P_O bits in floating TU mode, and in the multiframe indicator byte (H4) in locked mode. This is illustrated in Figure 5-25/G.709. The usage of the PP bits has options, because the common signalling method and another channel associated signalling method (e.g. Recommendation G.704, §§ 3.1.3 and 3.2.3) do not need the PP bits. The operations of the alternative channel associated signalling method is shown in Figure 5-26/G.709.

Figure 5-24/G.709, p.

Figure 5-25/G.709, p.

Figure 5-26/G.709 [T5.709], p.

5.8 *Floating and locked mode conversion*

There are two possible multiplexing modes of the TU structures: floating and locked.

In the floating TU mode four consecutive 125 µs VC-n frames are organized into a 500 µs multiframe, the phase of which is indicated by the multiframe indicator byte (H4) in the VC-*n* POH. This 500 µs TU multiframe is shown in Figure 3-13/G.709.

Locked TU mode of transport is a fixed mapping of synchronous structured payloads into a VC-*n* . This provides a direct correspondence between subtending tributary information and the location of that information within the VC-*n* . Since the tributary information is fixed and immediately identifiable with respect to the TU-*n* or AU-*n* pointer associated with the VC-*n* , no TU pointers are available for payload usage.

Figure 5-27/G.709 illustrates the conversion between floating and locked TU modes for each of the four TU sizes. Note that certain bytes (R) in the current set of mapping are not used in the floating mode in order that those mappings can be used in both floating and locked modes. Since the V1-V4 and V5 bytes are reserved, the 500 µs TU multiframe is unnecessary. Therefore the role of the multiframe indicator byte (H4) in locked mode is to define 2 and 3 ms signalling frames for byte synchronous mappings.

Figure 5-27/G.709, p.