INTERNATIONAL TELECOMMUNICATION UNION

ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

G.709

(03/93)

GENERAL ASPECTS OF DIGITAL TRANSMISSION SYSTEM

SYNCHRONOUS MULTIPLEXING STRUCTURE

ITU-T Recommendation G.709

(Previously "CCITT Recommendation")

FOREWORD

The ITU Telecommunication Standardization Sector (ITU-T) is a permanent organ of the International Telecommunication Union. The 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 Conference (WTSC), which meets every four years, established the topics for study by the ITU-T Study Groups which, in their turn, produce Recommendations on these topics.

ITU-T Recommendation G.709 was prepared by the ITU-T Study Group XVIII (1988-1993) and was approved by the WTSC (Helsinki, March 1-12, 1993).

NOTES

1 As a consequence of a reform process within the International Telecommunication Union (ITU), the CCITT ceased to exist as of 28 February 1993. In its place, the ITU Telecommunication Standardization Sector (ITU-T) was created as of 1 March 1993. Similarly, in this reform process, the CCIR and the IFRB have been replaced by the Radiocommunication Sector.

In order not to delay publication of this Recommendation, no change has been made in the text to references containing the acronyms "CCITT, CCIR or IFRB" or their associated entities such as Plenary Assembly, Secretariat, etc. Future editions of this Recommendation will contain the proper terminology related to the new ITU structure.

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

© ITU 1993

All rights reserved. No part of this publication may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying and microfilm, without permission in writing from the ITU.

CONTENTS

Page

1 General 1
1.1 Abbreviations 1
1.2 Basic multiplexing structure 2
2 Multiplexing method 3
2.1 Multiplexing of administrative units into STM-N 3
2.2 Multiplexing of Tributary Units into VC-4 and VC-3 4
2.3 Maintenance signals 7
2.4 Timing recovery 10
3 Pointers 10
3.1 AU-n pointer 10
3.2 TU-3 pointer 17
3.3 TU-1/TU-2 pointer 20
4 Path overhead description 28
4.1 VC-3/VC-4/VC-4-Xc POH 28
4.2 VC-1/VC-2 POH 30
4.3 CRC-7 polynomial algorithm 32
5 Mapping of tributaries into VCs 33
5.1 Mapping of tributaries into VC-4 34
5.2 Mapping of tributaries into VC-3 35
5.3 Mapping of tributaries into the VC-2 38
5.4 Mapping of tributaries into VC-12 39
5.5 Mapping of tributaries into VC-11 41
5.6 VC-11 to VC-12 conversion for transport by a TU-12 43
5.7 Floating and locked mode conversion 50
5.8 Mapping of ATM cells 51

SYNCHRONOUS MULTIPLEXING STRUCTURE

(Melbourne, 1988, revised Geneva, 1991 and Helsinki, 1993)

The CCITT,

considering

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

(b) that Recommendation G.708 specifies:

- the general principles and frame structure of the network node interface (NNI) for the SDH;
- the overall frame size of 9 rows by N × 270 columns;
- the section overhead (SOH) together with its byte allocation;
- arrangements for international interconnection of synchronous transport modules (STMs);

(c) that Recommendations G.707, G.708 and G.709 form a coherent set of specifications for the SDH and NNI,

recommends

that the formats for multiplexing and mapping elements into the STM-N at the NNI shall be as described in this Recommendation.

1 General

1.1 Abbreviations

For the purpose of this Recommendation the following abbreviations are used:

AIS	Alarm indication signal
APS	Automatic protection switching
ATM	Asynchronous transfer mode
AU-*n*	Administrative unit-*n*
AUG	Administrative unit group
BIP-X	Bit interleaved parity-X
CAS	Channel associated signalling
DCC	Data communication channel
FEBE	Far end block error
FERF	Far end receive failure
HEC	Header error control
NDF	New data flag
NNI	Network node interface
NPI	Null pointer indication

POH	Path overHead
SDH	Synchronous digital hierarchy
SOH	Section overHead
STM(-N)	Synchronous transport module (-N)
TU-*n*	Tributary unit-*n*
TUG (-*n*)	Tributary unit group (-*n*)
VC-*n*	Virtual container-*n*
VC-*n*-Xc	X time concatenated VC-*n* (*n* = 2 or 4)

1.2 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. The detailed multiplexing structure is described in the following subclauses.

NOTE – The order of transmission of information in all the diagrams in this Recommendation is first from left to right and then from top to bottom. Within each byte the most significant bit is transmitted first. The most significant bit (bit 1) is illustrated at the left in all the diagrams.

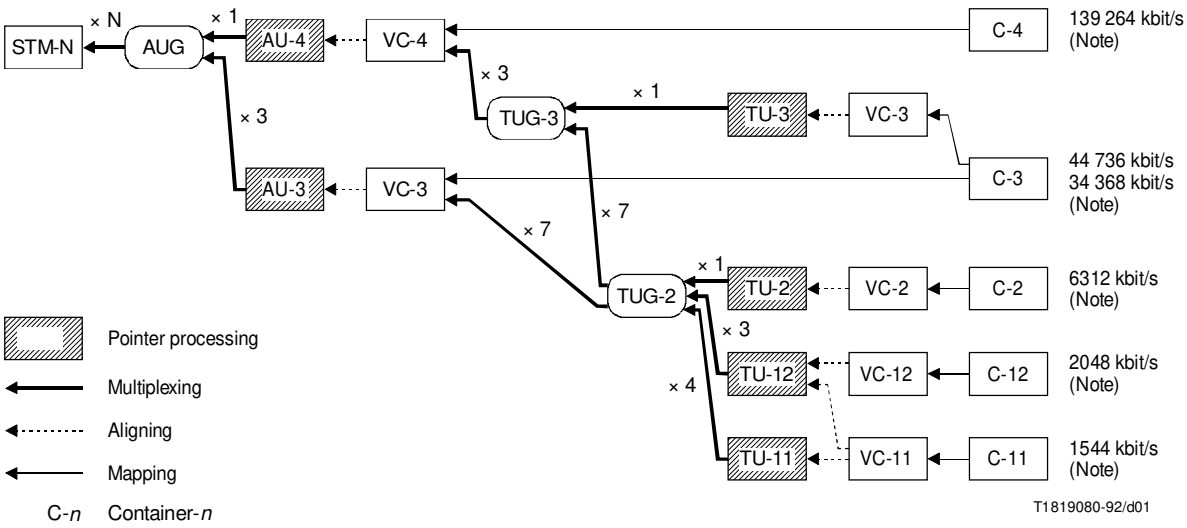

NOTE – G.702 tributaries associated with containers C- *x* are shown. Other signals, e.g. ATM, can also be accommodated.

FIGURE 1-1/G.709

Multiplexing structure

2 Multiplexing method

2.1 Multiplexing of administrative units into STM-N

2.1.1 Multiplexing of administrative unit groups (AUGs) into STM-N

The arrangement of N AUGs multiplexed into the STM-N is shown in Figure 2-1. The AUG is a structure of 9 rows by 261 columns plus 9 bytes in row 4 (for the AU-*n* pointers). The STM-N consists of an SOH as described in Recommendation G.708 and a structure of 9 rows by N × 261 columns with N × 9 bytes in row 4 (for the AU-*n* pointers). The N AUGs are one-byte interleaved into this structure and have a fixed phase relationship with respect to the STM-N.

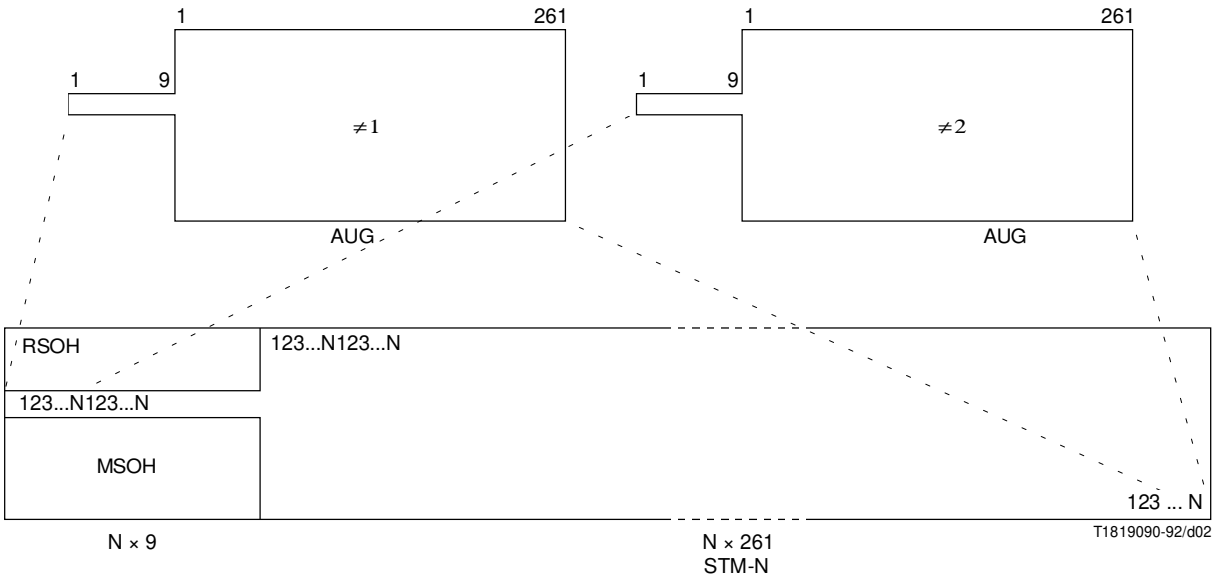

FIGURE 2-1/G.709

Multiplexing of N AUGs into STM-N

2.1.2 Multiplexing of an AU-4 via AUG

The multiplexing arrangement of a single AU-4 via the AUG is depicted in Figure 2-2. The 9 bytes at the beginning of row 4 are allocated to the AU-4 pointer. The remaining 9 rows by 261 columns is allocated to the virtual container-4 (VC-4). The phase of the VC-4 is not fixed with respect to the AU-4. The location of the first byte of the VC-4 with respect to the AU-4 pointer is given by the pointer value. The AU-4 is placed directly in the AUG.

2.1.3 Multiplexing of AU-3s via AUG

The multiplexing arrangement of three AU-3s via the AUG is depicted in Figure 2-3. The 3 bytes at the beginning of row 4 are allocated to the AU-3 pointer. The remaining 9 rows by 87 columns is allocated to the VC-3 and two columns of fixed stuff. The byte in each row of the two columns of fixed stuff of each AU-3 shall be the same. The phase of the VC-3 and the two columns of fixed stuff is not fixed with respect to the AU-3. The location of the first byte of the VC-3 with respect to the AU-3 pointer is given by the pointer value. The three AU-3s are one-byte interleaved in the AUG.

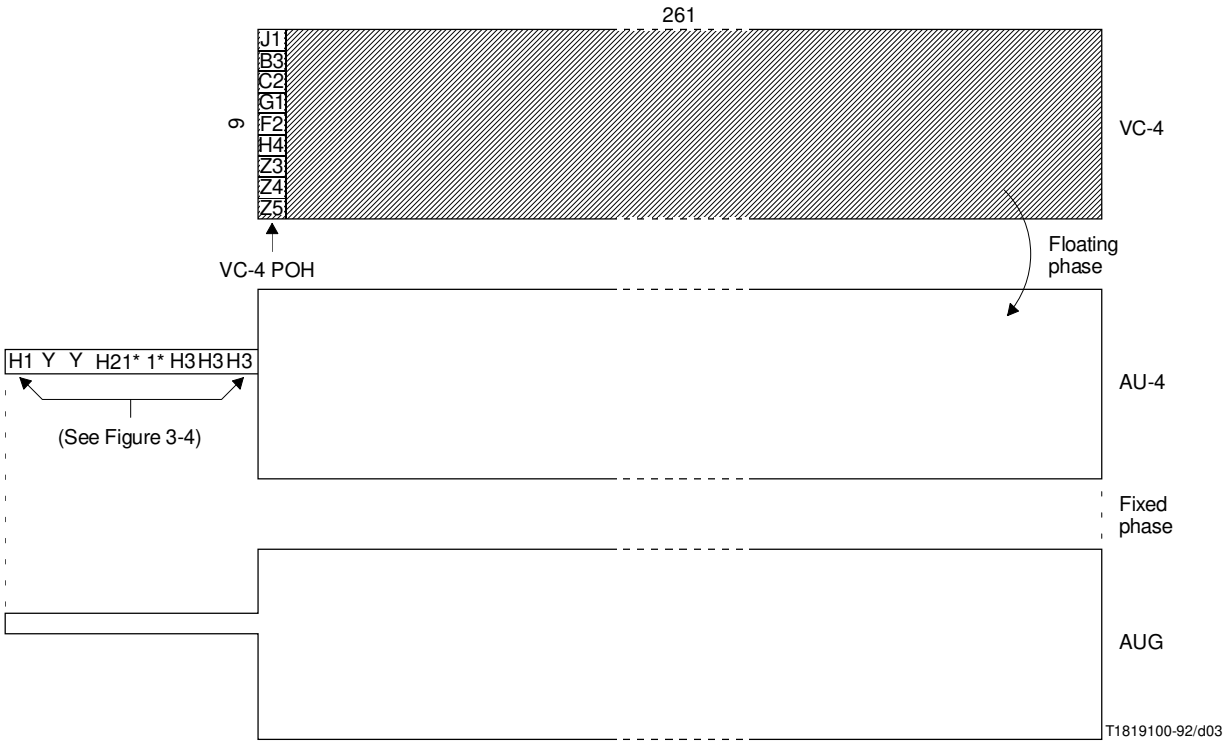

FIGURE 2-2/G.709

Multiplexing on AU-4 via AUG

2.2 Multiplexing of Tributary Units into VC-4 and VC-3

2.2.1 Multiplexing of Tributary Unit Group-3s (TUG-3s) into a VC-4

The arrangement of three TUG-3s multiplexed in the VC-4 is shown in Figure 2-4. The TUG-3 is a 9-row by 86-column structure. The VC-4 consists of one column of VC-4 POH, two columns of fixed stuff and a 258-column payload structure. The three TUG-3s are single byte interleaved into the 9-row by 258-column VC-4 payload structure and have a fixed phase with respect to the VC-4.

As described in 2.1, the phase of the VC-4 with respect to the AU-4 is given by the AU-4 pointer.

2.2.2 Multiplexing of a TU-3 via TUG-3

The multiplexing of a single TU-3 via the TUG-3 is depicted in Figure 2-5. The TU-3 consists of the VC-3 with a 9-byte VC-3 POH and the TU-3 pointer. The first column of the 9-row by 86-column TUG-3 is allocated to the TU-3 pointer (bytes H1, H2, H3) and fixed stuff. The phase of the VC-3 with respect to the TUG-3 is indicated by the TU-3 pointer.

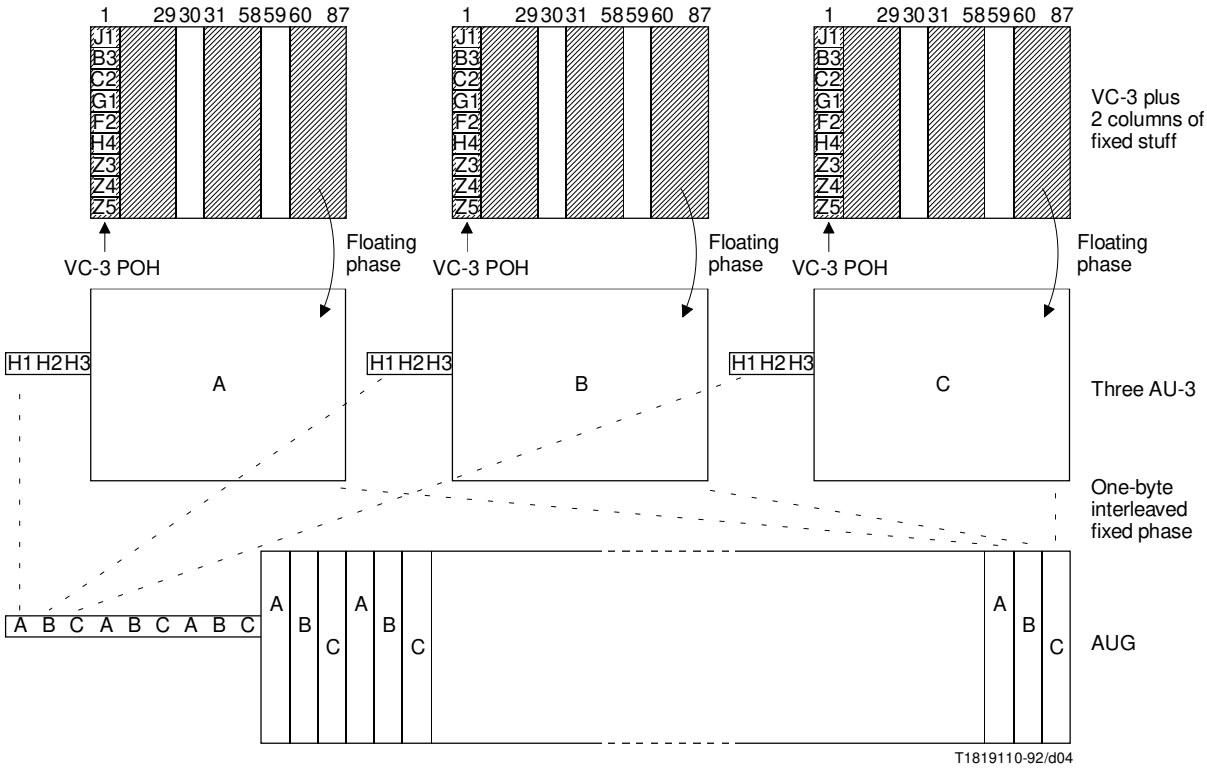

NOTE – The byte in each row of the two columns of fixed stuff of each AU-3 shall be the same.

FIGURE 2-3/G.709

Multiplexing of AU-3s via AUG

2.2.3 Multiplexing of TUG-2s via TUG-3

The multiplexing structure for the TUG-2 via the TUG-3 is depicted in Figure 2-6. The TUG-3 is a 9-row by 86-column structure with the first two columns accommodating the following :

– a Null Pointer Indication (NPI) contained in the first three bytes of the first column. This NPI can be used to distinguish between TUG-3s containing TU-3s and TUG-3s containing TUG-2s. See 3.2 for details ;

– stuffing in the other bytes of these two columns.

A group of seven TUG-2s can be multiplexed via the TUG-3.

The arrangement of seven TUG-2s multiplexed via the TUG-3 is depicted in Figure 2-7. The TUG-2s are one-byte interleaved in the TUG-3.

2.2.4 Multiplexing of TUG-2s into a VC-3

The multiplexing structure for TUG-2s into a VC-3 is depicted in Figure 2-8. The VC-3 consists of VC-3 POH and a 9-row by 84-column payload structure. A group of seven TUG-2s can be multiplexed into the VC-3.

The arrangement of seven TUG-2s multiplexed into the VC-3 is depicted in Figure 2-9. The TUG-2s are one-byte interleaved in the VC-3. An individual TUG-2 has a fixed location in the VC-3 frame.

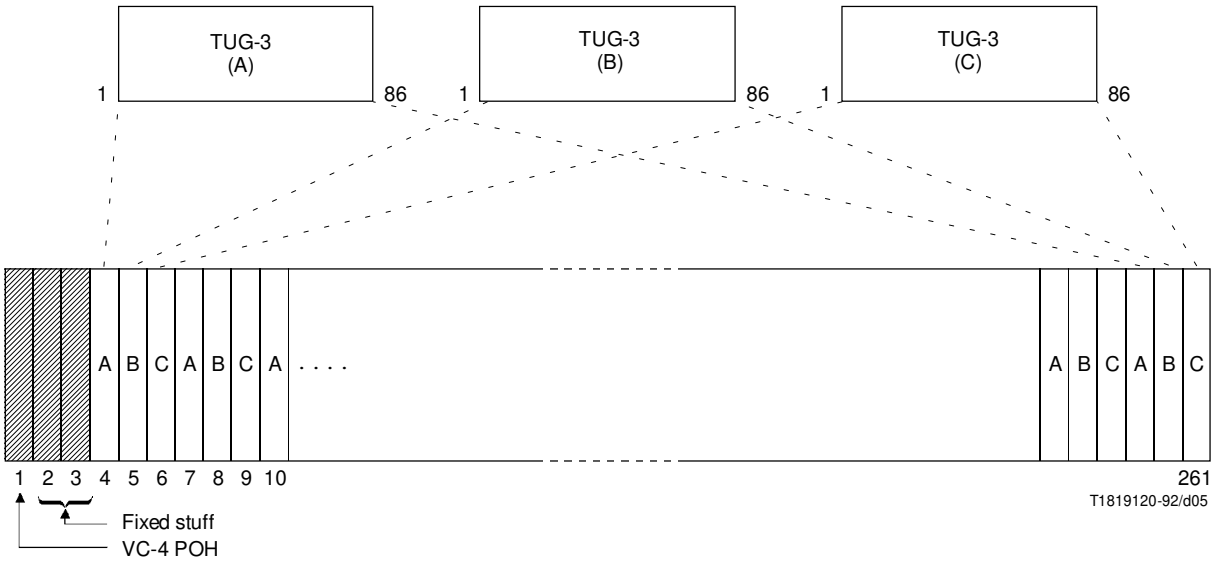

FIGURE 2-4/G.709

Multiplexing of three TUG-3s into a VC-4

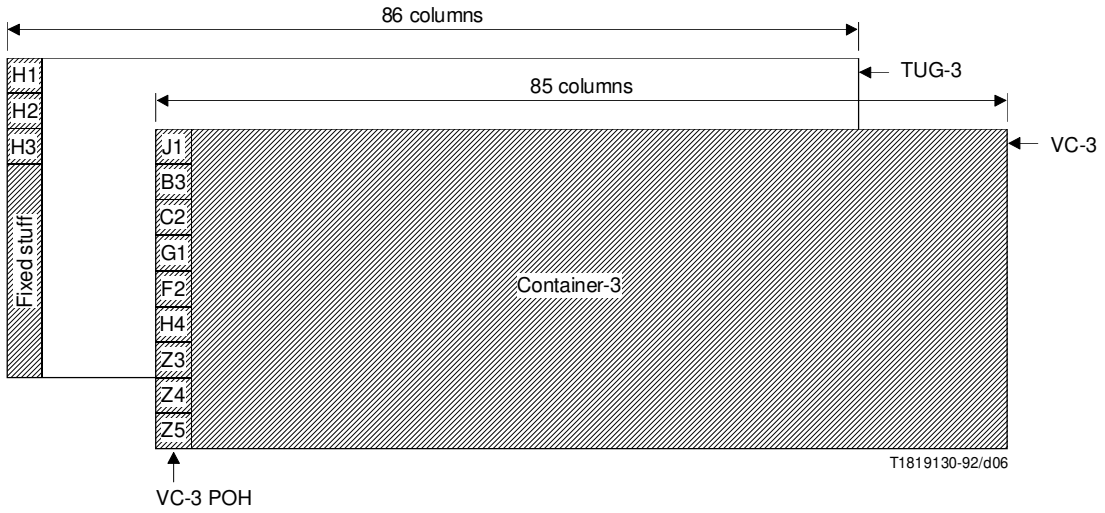

FIGURE 2-5/G.709

Multiplexing of a TU-3 via a TUG-3

2.2.5 Multiplexing of a TU-2 via TUG-2s

The multiplexing arrangement of a single TU-2 via the TUG-2 is depicted in Figure 2-9.

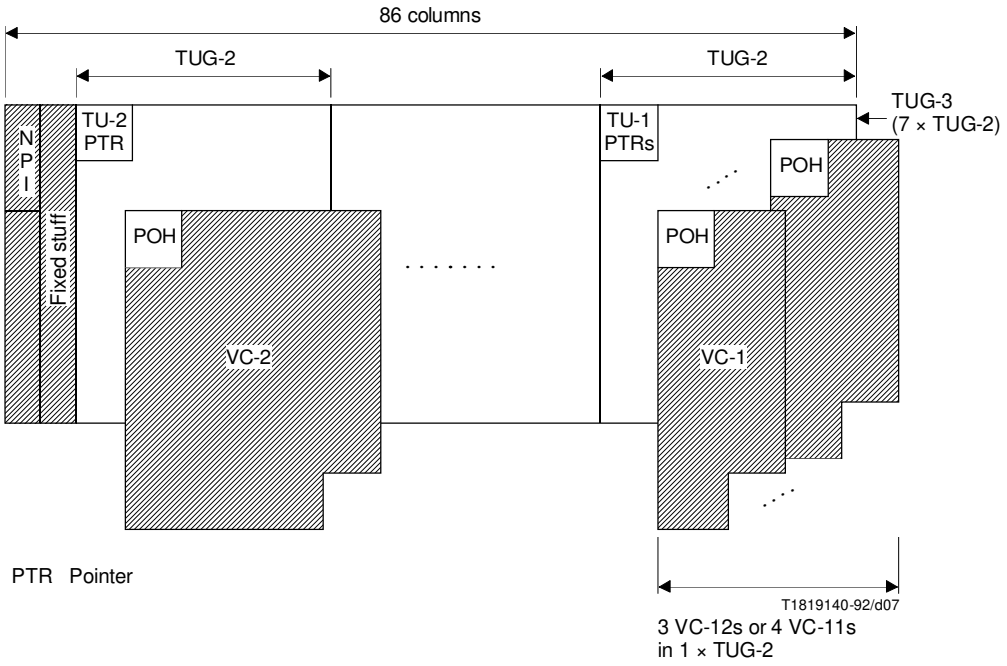

FIGURE 2-6/G.709

Multiplexing of seven TUG-2s via a TUG-3

2.2.6 Multiplexing of TU-1s via TUG-2s

The multiplexing arrangements of four TU-11s or three TU-12s via the TUG-2 are depicted in Figure 2-9. The TU-1s are one-byte interleaved in the TUG-2.

2.3 Maintenance signals

2.3.1 Section maintenance signals

An alarm indication signal (AIS) is a signal sent downstream as an indication that an upstream failure has been detected and alarmed.

The multiplex section AIS (MS-AIS) is detected as an all "1" in bits 6, 7 and 8 of the K2 byte after descrambling.

The multiplex section far end receive failure (MS-FERF) is used to return an indication to the transit end that the received end has detected an incoming section failure or is receiving section AIS.

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

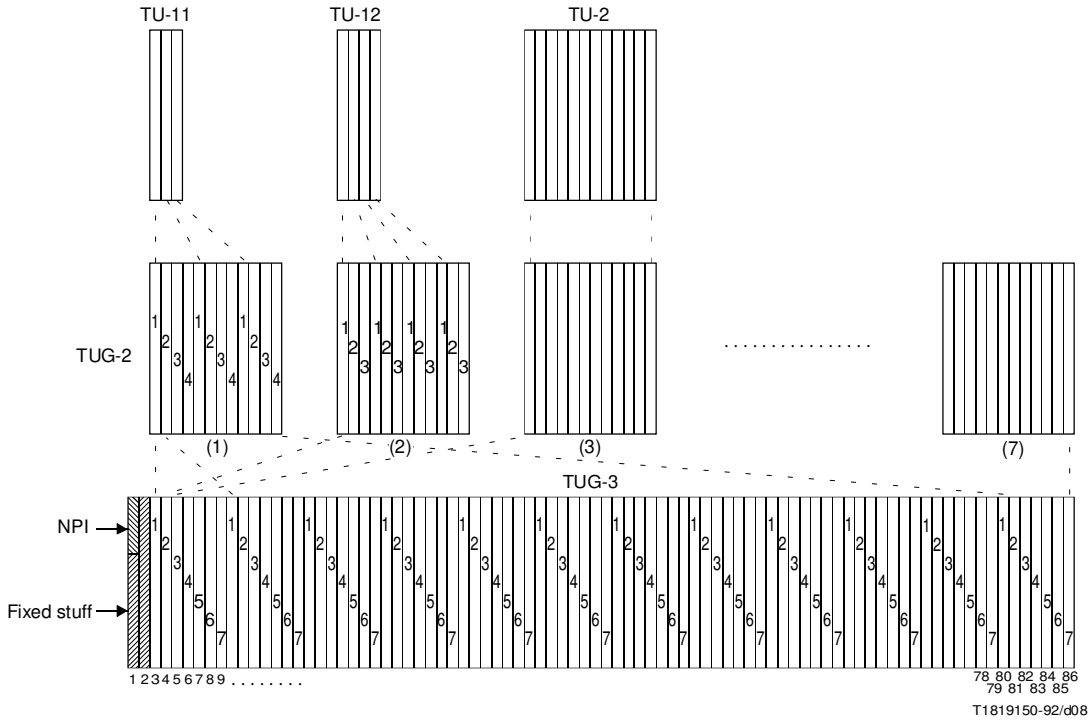

FIGURE 2-7/G.709

Multiplexing of seven TUG-2s via a TUG-3

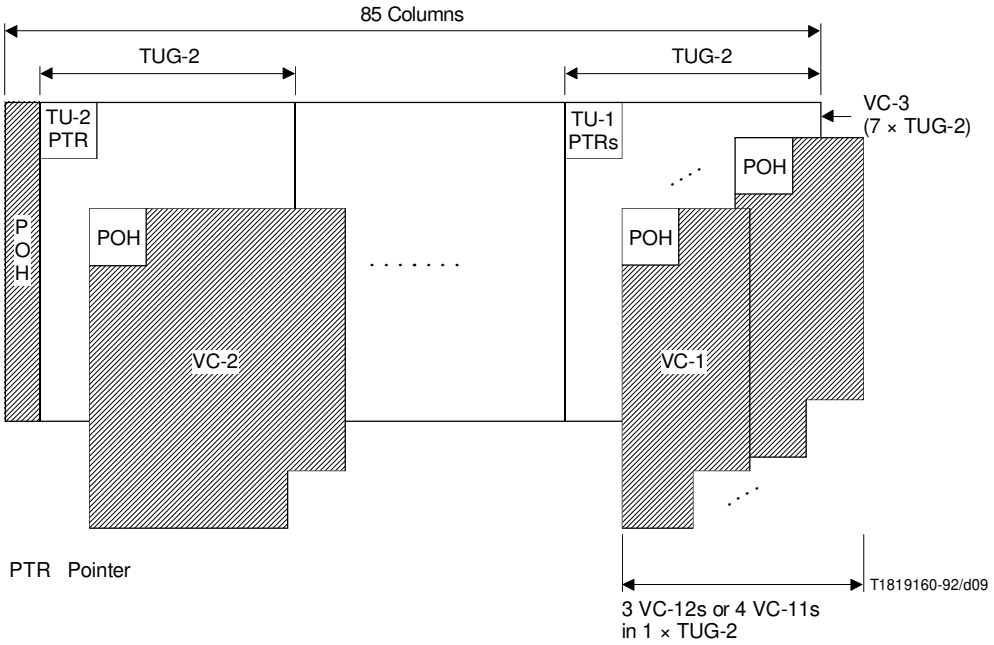

FIGURE 2-8/G.709

Multiplexing of seven TUG-2s into a VC-3

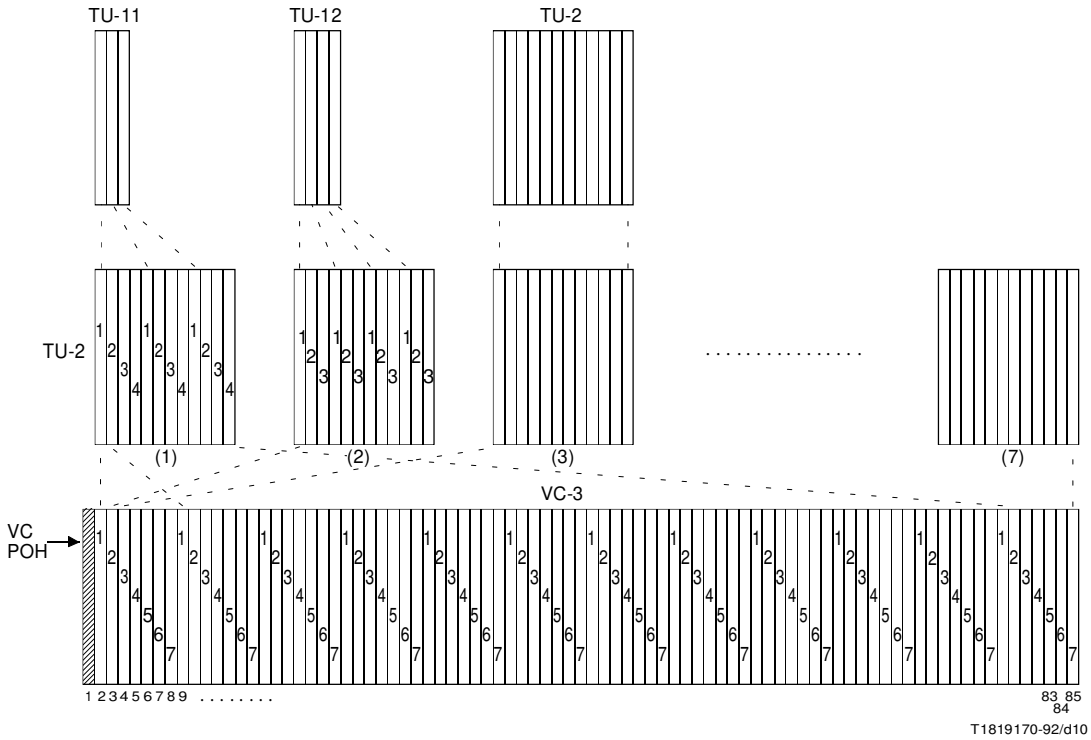

FIGURE 2-9/G.709

Multiplexing of seven TUG-2s into a VC-3

2.3.2 Path maintenance signals

The VC-*n* (*n* = 3, 4) or VC-4-Xc unequipped indication is an all "0"s in the virtual container path signal label (byte C2) after descrambling. The VC-*n* (*n* = 1, 2) unequipped indication is an all "0s" in the lower path signal label (bits 5-7 of the V5 byte). This code indicates to hte virtual container terminating equipment that the virtual container is intentionally unoccupied so that alarms can be inhibited. This code is generated as an all "0"s in the virtual container path signal label and a valid virtual container path BIP-8 (B3); the virtual container payload is unspecified.

The TU-*n* (*n* = 1, 2, 3) path AIS is specified as all "1"s in the entire TU-*n*, including the TU-*n* pointer. Similarly, the AU-*n* (*n* = 3, 4) path AIS is specified as all "1"s in the entire AU-*n*, including the AU-*n* pointer. All path AIS's are carried within STM-N signals having valid SOH.

The path status byte (G1) is used to convey to the originator of a VC-*n* (*n* = 3 or 4) or to the VC-4-Xc the terminating path status and performance. Bits 1 through 4 convey the count of errored parity blocks called far end block error (FEBE), 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 error. Bit 5 is a path FERF.

2.4 Timing recovery

The STM-N (N = 1, 4, 16) signal must have sufficient bit timing content at the NNI. A suitable bit pattern, which prevents a long sequence of "1"s and "0"s is provided by using a scrambler.

The operation of the scrambler 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 gives a functional diagram of the frame synchronous scrambler.

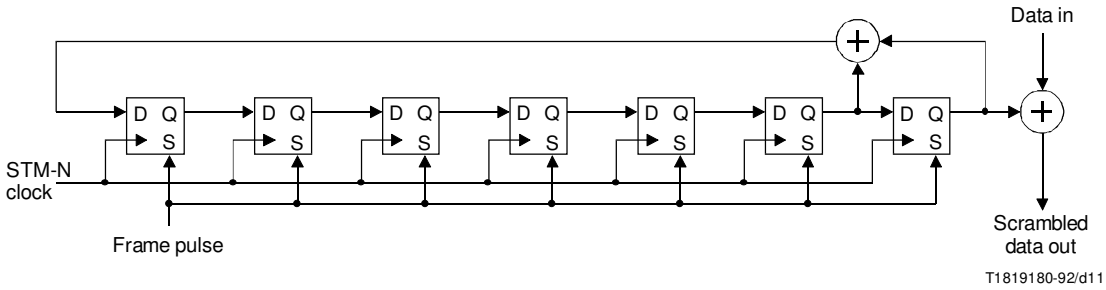

FIGURE 2-10/G.709

Frame synchronous scrambler (functional diagram)

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 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 "1"s or "0"s do not occur.

3 Pointers

3.1 AU-*n* pointer

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

Dynamic alignment means that the VC-*n* is allowed to "float" within the AU-*n* frame. Thus, the pointer is able to accommodate differences, not only in the phases of the VC-*n* and the SOH, but also in the frame rates.

3.1.1 AU-n pointer location

The AU-4 pointer is contained in bytes H1, H2 and H3 as shown in Figure 3-1. The three individual AU-3 pointers are contained in three separate H1, H2 and H3 bytes as shown in Figure 3-2.

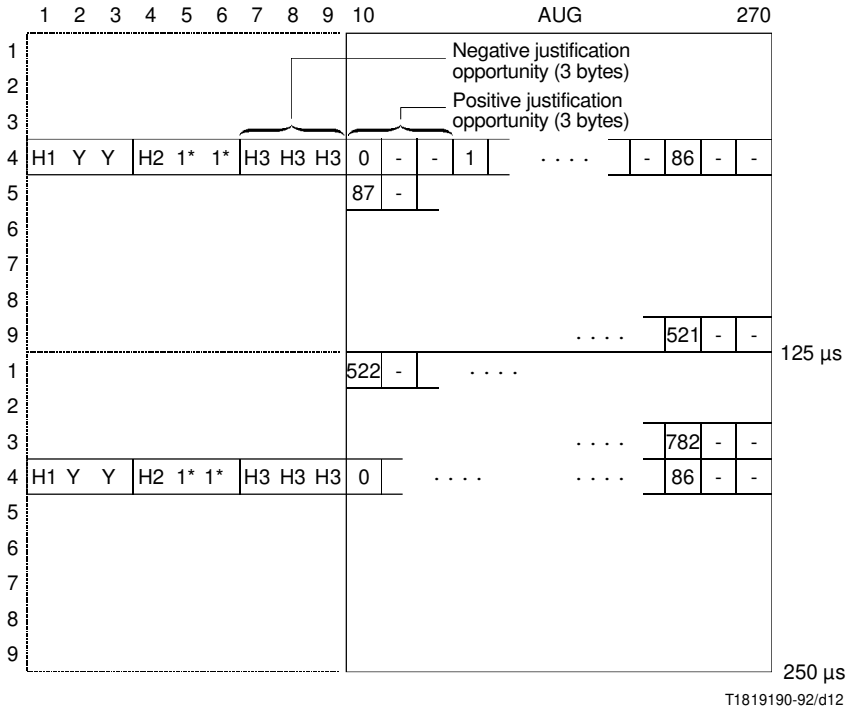

FIGURE 3-1/G.709

AU-4 pointer offset numbering

3.1.2 AU-n pointer value

The pointer contained in H1 and H2 designates the location of the byte where the VC-*n* begins. The two bytes allocated to the pointer function can be viewed as one word as shown in Figure 3-3. The last ten bits (bits 7-16) of the pointer word carry the pointer value.

As illustrated in Figure 3-3, the AU-4 pointer value is a binary number with a range of 0 to 782 which indicates the offset, in three byte increments, between the pointer and the first byte of the VC-4 (see Figure 3-1). Figure 3-3 also indicates one additional valid pointer, the concatenation indication. The concatenation indication is indicated by "1001" in bits 1-4, bits 5-6 unspecified, and ten "1"s in bits 7-16. The AU-4 pointer is set to concatenation indication for AU-4 concatenation (see 3.1.7).

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

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

3.1.3 Frequency justification

If there is a frequency offset between the frame rate of the AUG and that of the VC-*n*, 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.

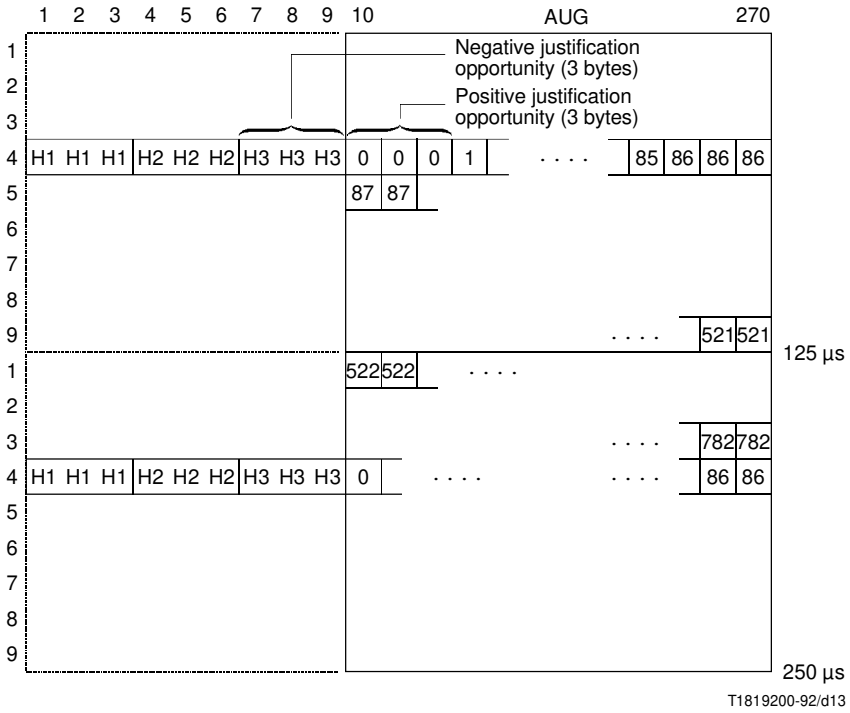

FIGURE 3-2/G.709

AU-3 pointer offset numbering

If the frame rate of the VC-*n* is too slow with respect to that of the AUG, then the alignment of the VC-*n* 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 3-4.

For AU-3 frames, a positive justification byte appears immediately after the individual H3 byte of the AU-3 frame containing inverted I-bits. Subsequent pointers will contain the new offset. This is illustrated in Figure 3-5.

If the frame rate of the VC-*n* is too fast with respect to that of the AUG, then the alignment of the VC-*n* 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 3-6.

For AU-3 frames, a negative justification byte appears in the individual H3 byte of the AU-3 frame containing inverted D-bits. Subsequent pointers will contain the new offset. This is illustrated in Figure 3-7.

3.1.4 New data flag (NDF)

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

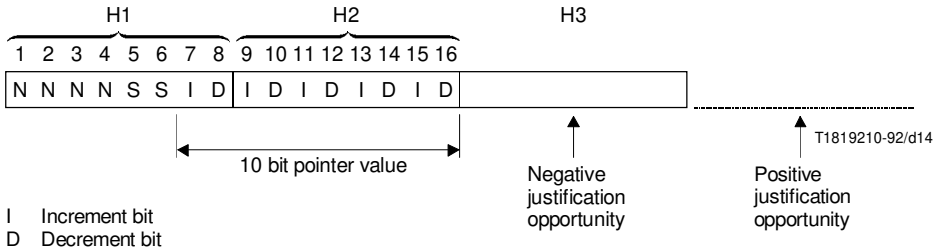

I Increment bit
D Decrement bit
N New data flag bit

New data flag
– Enabled "1001"
– Disabled "0110"

Negative justification
– Invert 5 D-bits
– Accept majority vote

SS values	AU-*n*/TU-*n* type
10	AU-4, AU-3, TU-3

Positive justification
– Invert 5 I-bits
– Accept majority vote

Pointer value (bits 7-16)
– Normal range
AU-4, AU-3: 0-782 decimal
TU-3: 0-764 decimal

Concatenation indication
– 1001SS1111111111 (S bits are unspecified)

Null pointer indication (NPI)
– 1001SS1111100000 (S bits are unspecified)

NOTES

1 NPI value applies only to TU-3 pointers.

2 The pointer is set to all "1"s when an AIS occurs.

FIGURE 3-3/G.709

AU-*n*/TU-3 pointer (H1, H2, H3) coding

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.

3.1.5 Pointer generation

The following summarizes the rules for generating the AU-*n* pointers.

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

2) The pointer value can only be changed by operation 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. If the previous pointer is at its maximum value, the subsequent pointer is set to zero. 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. If the previous pointer value is zero, the subsequent pointer is set to its maximum value. No subsequent increment or decrement operation is allowed for at least three frames following this operation.

5) If the alignment of the VC-*n* 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-*n* 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.

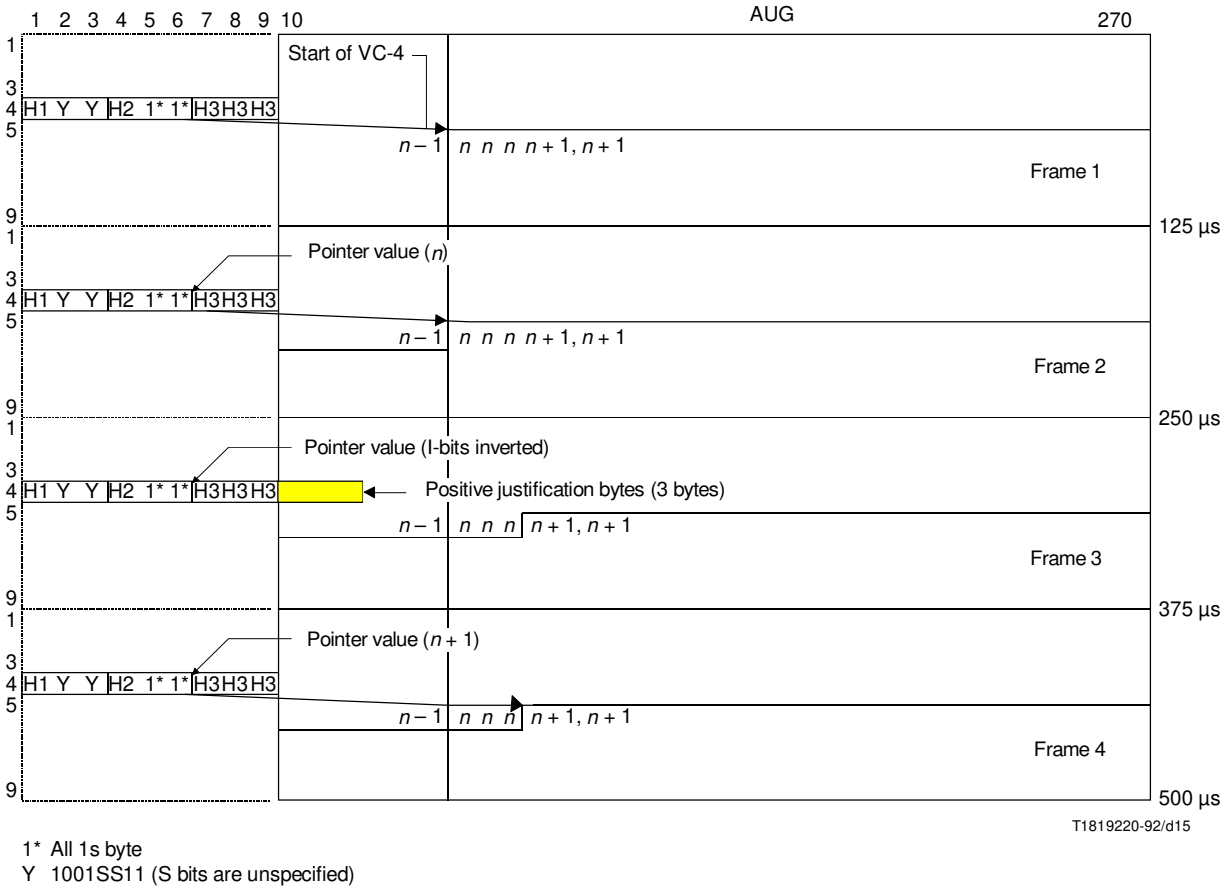

1* All 1s byte
Y 1001SS11 (S bits are unspecified)

FIGURE 3-4/G.709

AU-4 pointer adjustment operation – positive justification

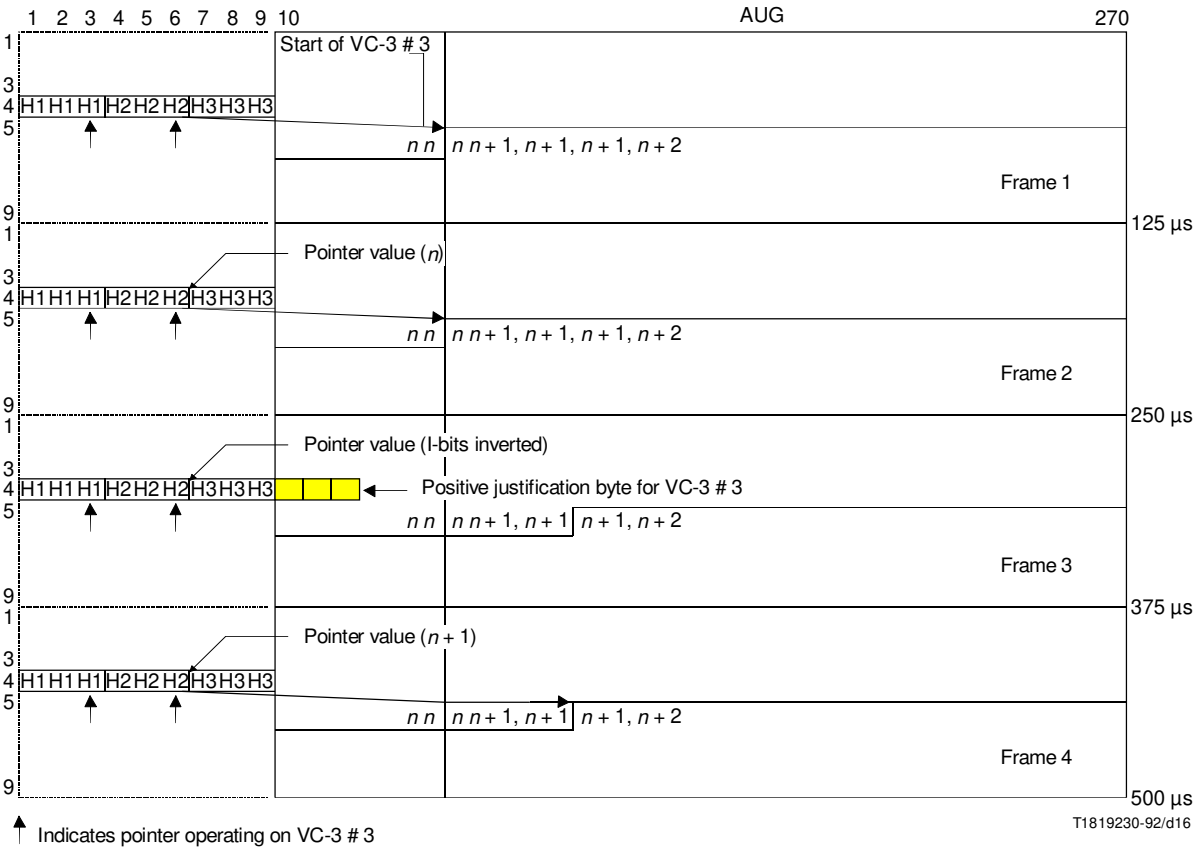

FIGURE 3-5/G.709

AU-3 pointer adjustment operation – Positive justification

3.1.6 Pointer interpretation

The following summarizes the rules for interpreting the AU-*n* pointers.

1) During normal operation, the pointer locates the start of the VC-*n* within the AU-*n* 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 the rules 3, 4 or 5. Any consistent new value received three time consecutively overrides (i.e. takes priority over) rules 3 or 4.

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 unless the receiver is in a state that corresponds to a loss of pointer.

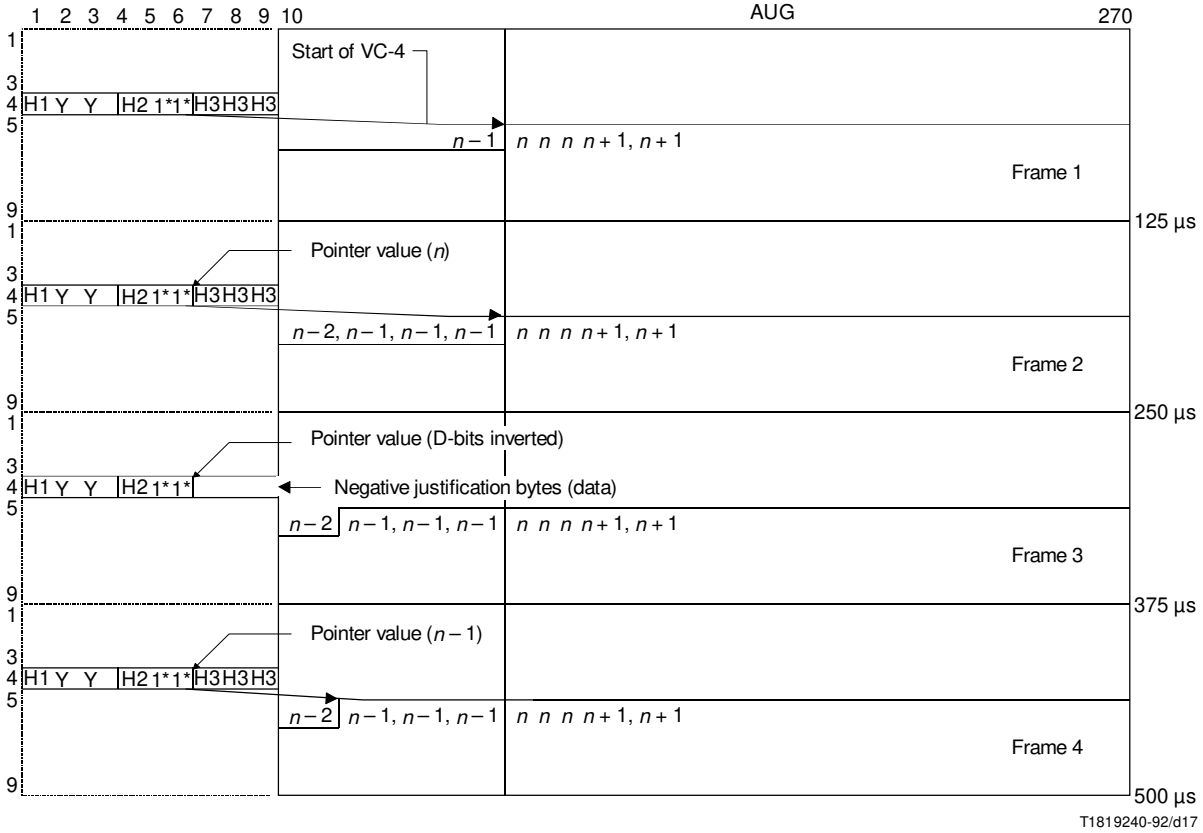

FIGURE 3-6/G.709

AU-4 pointer adjustment operation – Negative justification

3.1.7 AU-4 concatenation

AU-4s can be concatenated together to form an AU-4-Xc which can transport payloads requiring greater capacity than one container-4 capacity. A concatenation indication, used to show that this multi container-4 payload carried in a single VC-4-Xc should be kept together, is contained in the AU-4 pointer.The capacity available for the mapping, the multi container-4, is X times the capacity of the container-4 (e.g. 599 040 Mbit/s for $X = 4$ and 2 396 160 kbit/s for $X = 16$). Columns 2 to X of the VC-4-Xc are specified as fixed stuff. The first column of the VC-4-Xc is used for the POH. The POH is assigned to the VC-4-Xc (e.g. the BIP-8 covers 261 X columns of the VC-4-Xc). The VC-4-Xc is illustrated in Figure 3-8.

The first AU-4 of an AU-4-Xc shall have a normal range of pointer values. All subsequent AU-4s within the AU-4-Xc shall have their pointer set to concatenation indication "1001" in bits 1-4, bits 5-6 unspecified, and ten "1"s in bits 7-16. The concatenation indication determines that the pointer processors shall perform the same operations as performed on the first AU-4 of the AU-4-Xc.

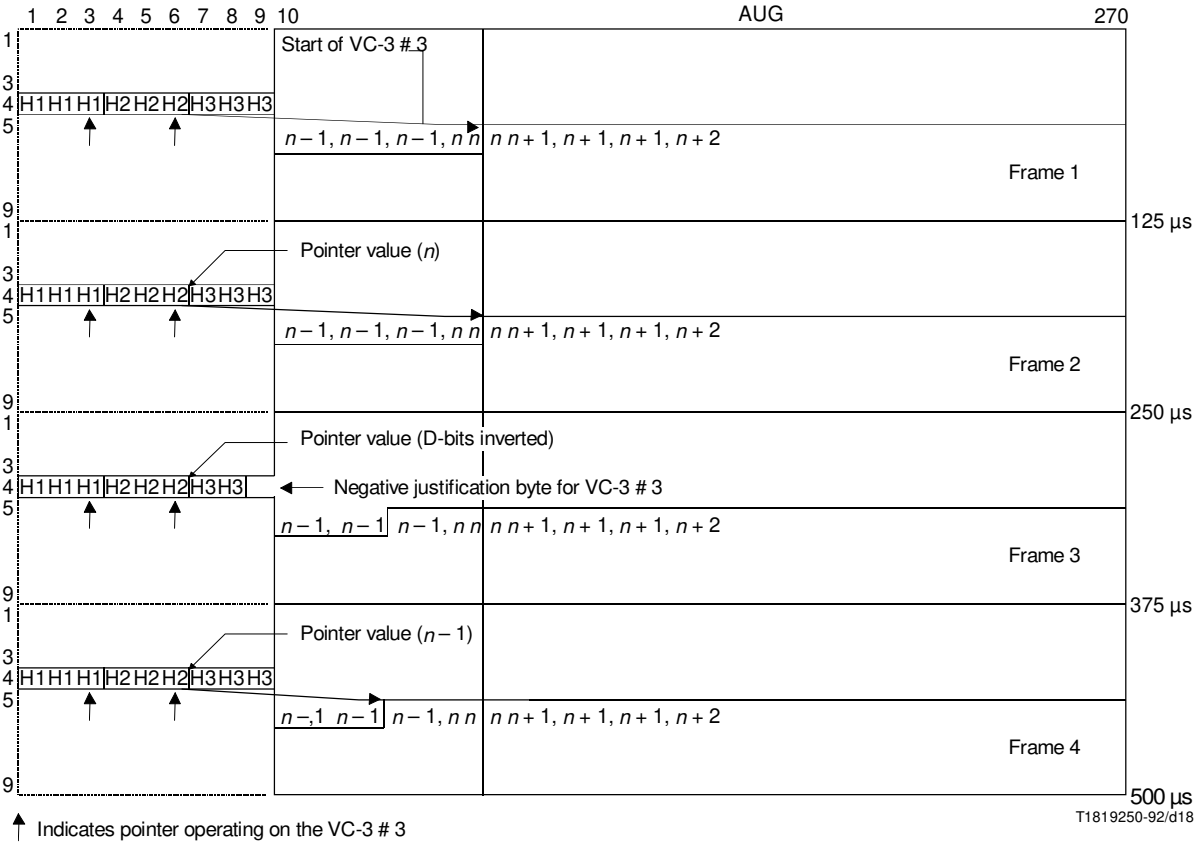

FIGURE 3-7/G.709

AU-3 pointer adjustment operation – Negative justification

3.1.7.1 Pointer generation

The following additional pointer generation rule shall apply for AU-4 pointers. If an AU-4-Xc signal is being transmitted, a pointer is generated for the first AU-4 only. The Concatenation Indication is generated in place of the other AU-4 pointers in the AU-4-Xc. All operations indicated by the AU-4 pointer in the first AU-4 apply to each AU-4 in the AU-4-Xc.

3.1.7.2 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 AU-4 are identical to those performed on the first AU-4 within the AU-4-Xc. Any variation from the concatenation indication is ignored unless a consistent new pointer value is received three times consecutively.

3.2 TU-3 pointer

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 content of the VC-3.

3.2.1 TU-3 pointer location

Three individual TU-3 pointers are contained in the three separate H1, H2 and H3 bytes as shown in Figure 3-8.

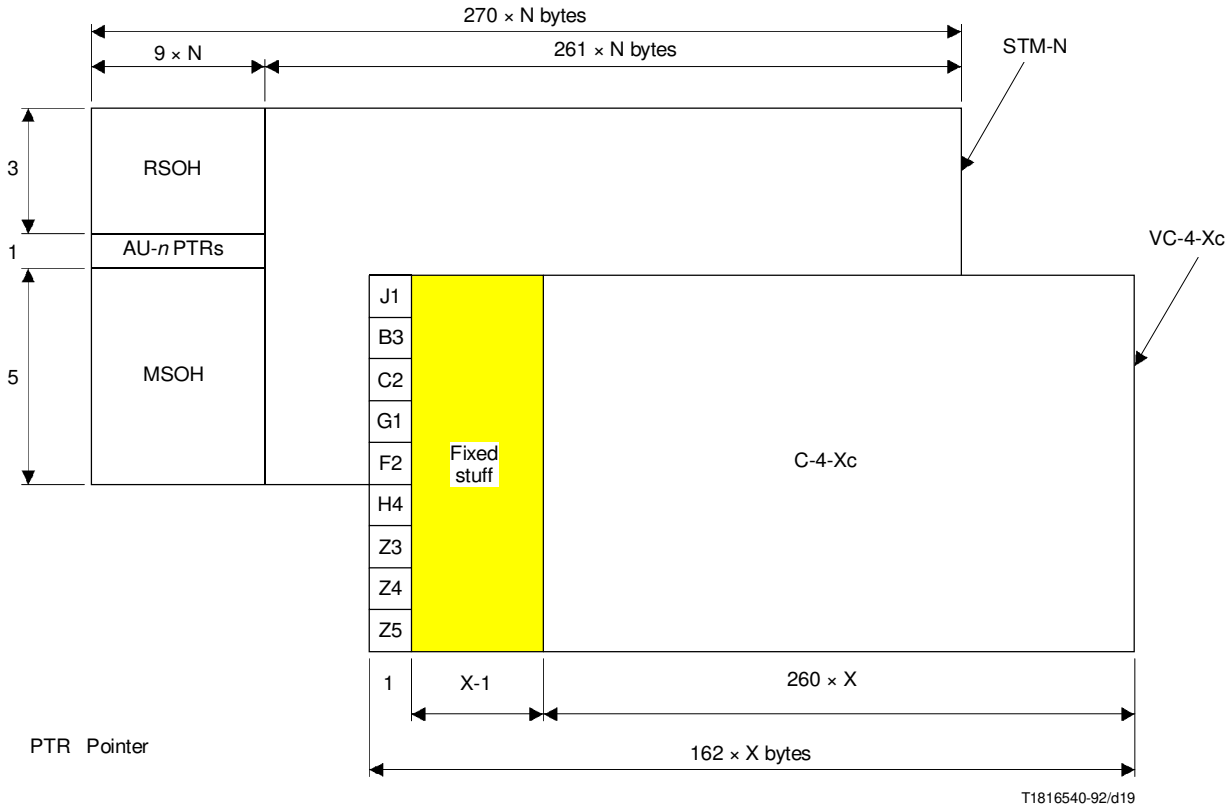

FIGURE 3-8/G.709

VC-4-Xc structure

When TUG-2s are multiplexed into a VC-4, the TU-3 pointer location is set to a Null Pointer Indication (NPI). NPI is indicated by "1001" in bits 1-4, bits 5- 6 unspecified, five "1"s in bits 7-11 followed by five "0"s in bits 12-16.

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 3-3. The last ten bits (bits 7-16) of the pointer word carry the pointer value.

The TU-3 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-3 as shown in Figure 3-9.

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.

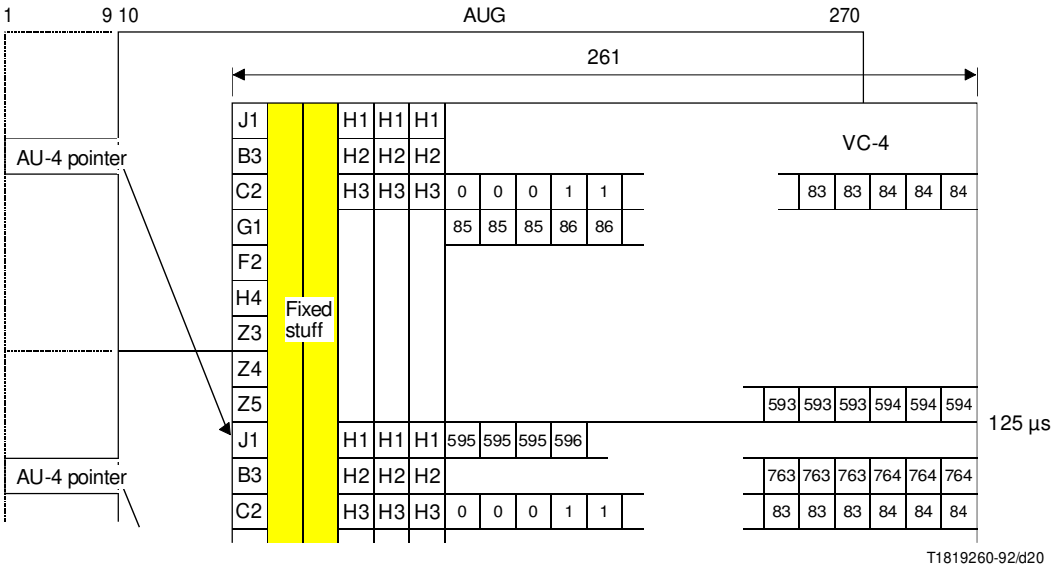

FIGURE 3-9/G.709

TU-3 pointer offset numbering

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-3 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-3 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 (NDF)

Bits 1-4 (N-bits) of the pointer word carry an 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 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 indicated offset.

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 operation 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. If the previous pointer is at its maximum value, the subsequent pointer is set to zero. 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. If the previous pointer value is zero, the subsequent pointer is set to its maximum value. No subsequent increment or decrement operation is allowed for at least three frames following this operation.

5) If the alignment of the VC-3 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-3 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. Any consistent new value received three times consecutively overrides (i.e. takes priority over) rules 3 or 4.

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 unless the receiver is in a state that corresponds to a loss of pointer.

6) If the TU-3 pointer contains the NPI, then any variation is ignored unless a consistent new pointer value is received three times consecutively.

3.3 TU-1/TU-2 pointer

The TU-1 pointer is only used with floating mapping. Floating and locked modes of operation are described in 5.7.

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-1/VC-2.

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 3-10.

3.3.2 TU-1/TU-2 pointer value

The Tributary Unit pointer word is shown in Figure 3-11. The two S bits (bits 5 and 6) indicate the Tributary Unit type.

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 Tributary Unit sizes as illustrated in Figure 3-12. The pointer bytes are not counted in the offset calculation.

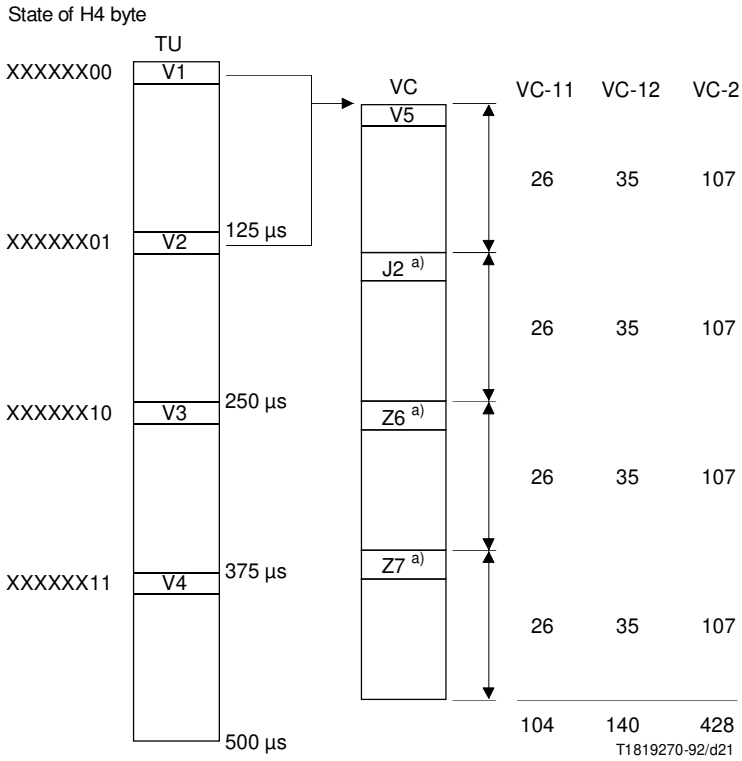

a) Allocation of these bytes is provisional.

TU Tributary unit
VC Virtual container
V1 VC Pointer 1
V2 VC Pointer 2
V3 VC Pointer 3 (action)
V4 Reserved

NOTE – V1, V2, V3 and V4 bytes are part of the TU and are terminated at the pointer processor.

FIGURE 3-10/G.709

Virtual container mapping in multiframed tributary unit

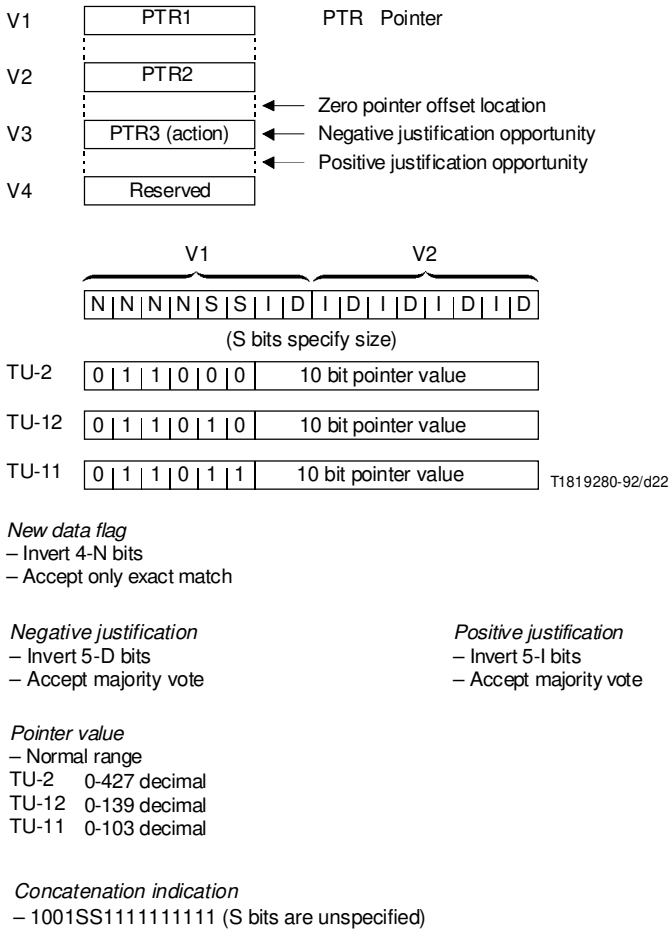

FIGURE 3-11/G.709

TU-1/TU-2 pointer coding

3.3.3 TU-1/TU-2 multiframe indication byte

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

– 500 microseconds (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 channel associated signalling for 2048 kbit/s payloads in the locked TU-1 mode;

– 3 ms (24-frame) multiframe for byte synchronous channel associated signalling for 1544 kbit/s payloads in the locked TU-1 mode.

The value of the H4 byte, read from the VC-3/VC-4 POH, identifies the frame phase of the next VC-3/VC-4 payload as shown in Figure 3-13. The coding of the H4 byte is illustrated in Figures 3-14, 3-15 and 3-16.

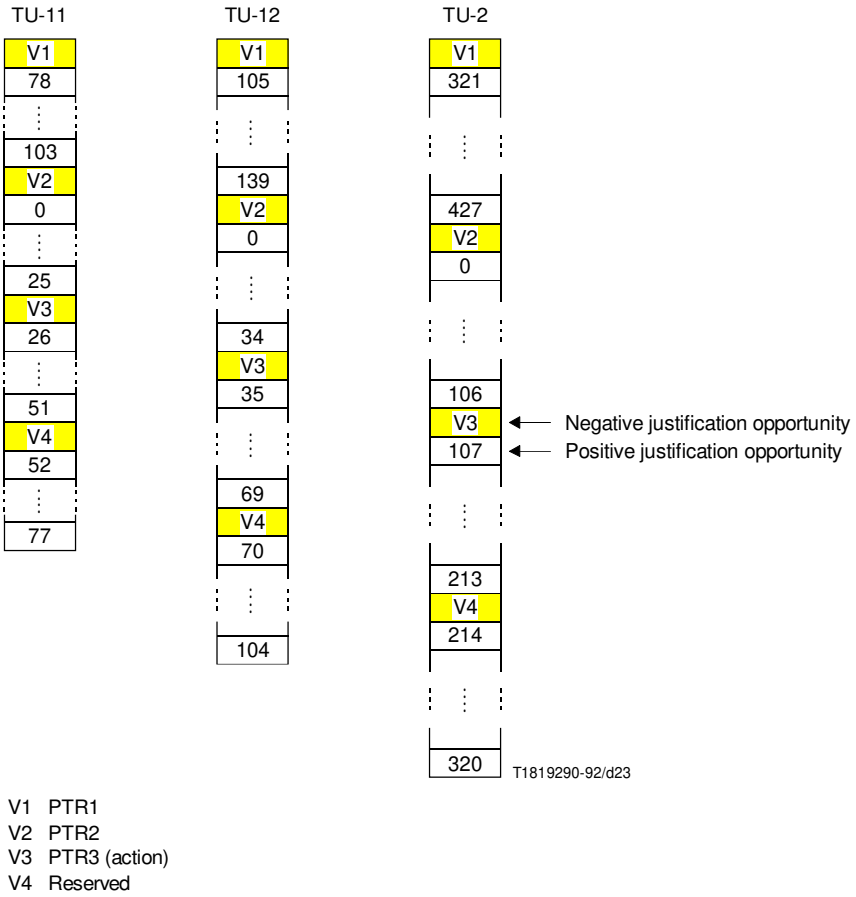

FIGURE 3-12/G.709

TU-1/TU-2 pointer offsets

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 microseconds multiframe. The 2 or 3 ms multiframe of any signalling within floating TU-1s is indicated by per-tributary unit multiframe indicators carried within the TU-1. Figure 3-10 shows the VC-1/VC-2 mapping in the multiframed TU-1/TU-2.

A converter from locked to floating Tributary Units is permitted to pass H4 through transparently. A converter from floating to locked Tributary Units must recover and align the multiframes from all of the floating tributary units and so can transmit any convenient full multiframe on the locked tributary unit 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 in 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 3-12. 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 tributary unit 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 for negative justification.

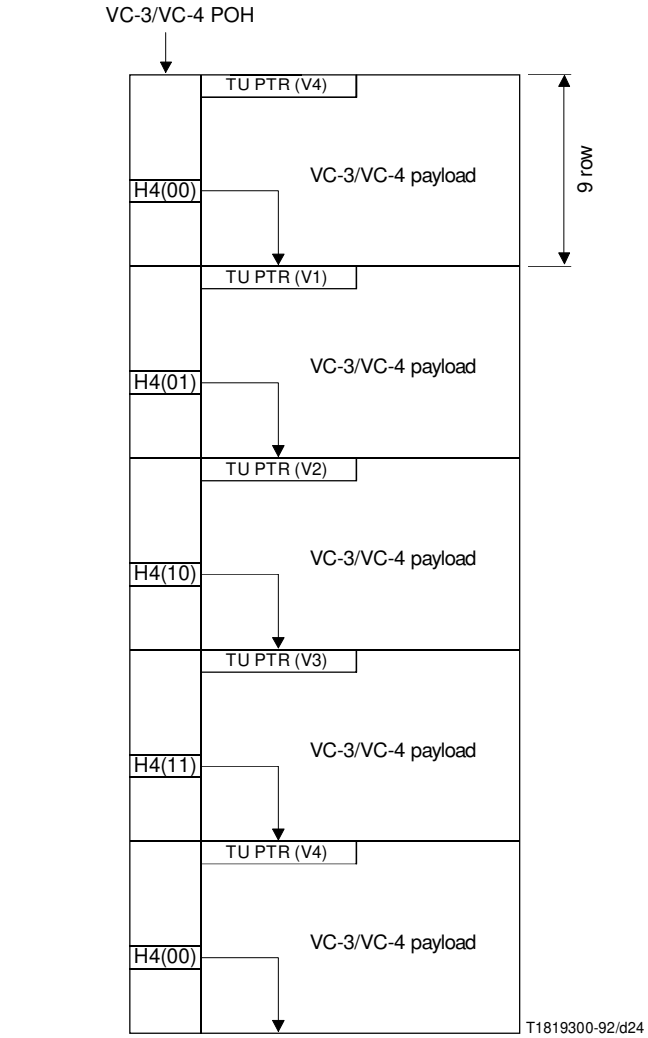

FIGURE 3-13/G.709

An example of TU-1/2 multiframe indication using H4 byte (the case of 500 μs multiframe)

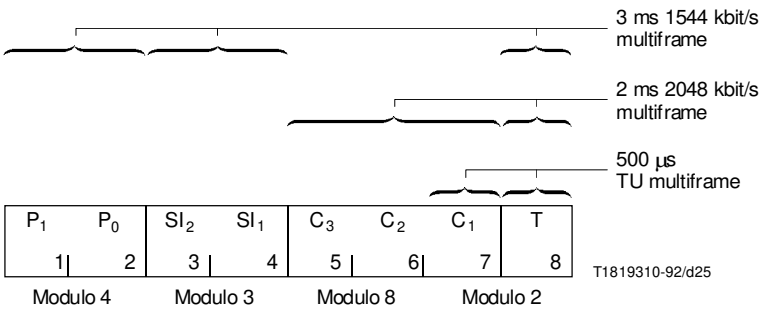

FIGURE 3-14/G.709

Tributary Unit multiframe indicator byte (H4)

Bit				Frame	Time
1 2	3 4	5 6 7	8		
0 0	0 0	0 0 0	0	0	0
0 0	0 0	0 0 0	1	1	
0 0	0 1	0 0 1	0	2	
0 0	0 1	0 0 1	1	3	500 µs TU Multiframe
0 0	1 0	0 1 0	0	4	
0 0	1 0	0 1 0	1	5	
0 1	0 0	0 1 1	0	6	
0 1	0 0	0 1 1	1	7	
0 1	0 1	1 0 0	0	8	
0 1	0 1	1 0 0	1	9	
0 1	1 0	1 0 1	0	10	
0 1	1 0	1 0 1	1	11	
1 0	0 0	1 1 0	0	12	
1 0	0 0	1 1 0	1	13	
1 0	0 1	1 1 1	0	14	
1 0	0 1	1 1 1	1	15	2 ms 2048 kbit/s signalling cycle
1 0	1 0	0 0 0	0	16	
1 0	1 0	0 0 0	1	17	
1 1	0 0	0 0 1	0	18	
1 1	0 0	0 0 1	1	19	
1 1	0 1	0 1 0	0	20	
1 1	0 1	0 1 0	1	21	
1 1	1 0	0 1 1	0	22	
1 1	1 0	0 1 1	1	23	3 ms 1544 kbit/s signalling cycle
0 0	0 0	1 0 0	0	24	
0 0	0 0	1 0 0	1	25	
0 0	0 1	1 0 1	0	26	
0 0	0 1	1 0 1	1	27	
0 0	1 0	1 1 0	0	28	
0 0	1 0	1 1 0	1	29	
0 1	0 0	1 1 1	0	30	
0 1	0 0	1 1 1	1	31	
0 1	0 1	0 0 0	0	32	
0 1	0 1	0 0 0	1	33	
0 1	1 0	0 0 1	0	34	
0 1	1 0	0 0 1	1	35	
1 0	0 0	0 1 0	0	36	
1 0	0 0	0 1 0	1	37	
1 0	0 1	0 1 1	0	38	
1 0	0 1	0 1 1	1	39	
1 0	1 0	1 0 0	0	40	
1 0	1 0	1 0 0	1	41	
1 1	0 0	1 0 1	0	42	
1 1	0 0	1 0 1	1	43	
1 1	0 1	1 1 0	0	44	
1 1	0 1	1 1 0	1	45	
1 1	1 0	1 1 1	0	46	
1 1	1 0	1 1 1	1	47	6 ms = cycle repeat time

TU Tributary unitt

NOTE – Full H4 coding sequence is mandatory in locked TU mode and optional in floating TU mode.

FIGURE 3-15/G.709

Tributary Unit multiframe indicator byte (H4) full coding sequence

Bit				Frames	Time
1 2	3 4	5 6 7	8		
1 1	1 1	1 1 0	0	0	0
1 1	1 1	1 1 0	1	1	
1 1	1 1	1 1 1	0	2	
1 1	1 1	1 1 1	1	3	Multitrame d'unité TU de 500 µs

TU Tributary unit

NOTES

1 Use of reduced mode can be detected by bits 3 and 4 = "1".

2 Reduced H4 coding sequence is optional in floating TU mode.

FIGURE 3-16/G.709

Tributary Unit multiframe indicator byte (H4)
reduced coding sequence

3.3.5 TU-1/TU-2 sizes

Bits 5 and 6 of TU-1/TU-2 pointer indicate the size of the TU. Three sizes are currently provided as indicated in Table 1.

TABLE 1/G.709

Size	Designation	TU-*n* pointer range (in 500 µs)
00	TU-2	0 – 427
10	TU-12	0 – 139
11	TU-11	0 – 103

NOTE – 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 Tributary Unit, 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 Tributary Units in the TUG-2.

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-*n*, and possibly new size. If a new size is indicated, then all Tributary Unit pointers in the TUG-2 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.

3.3.7 TU-2 concatenation

TU-2s may be concatenated to form a TU-2-mc (concatenated m × TU-2s) when a payload is required of more than a Container-2. This forms a multi Container-2 payload which is carried in a single VC-2-mc (concatenated m x VC-2). The rules by which TU-2s can be concatenated are separated into three categories:

- concatenation of contiguous TU-2s in the higher order VC-3;
- sequential concatenation of TU-2s in the higher order VC-4;
- virtual concatenation of TU-2s in the higher order VC-4.

The details and extensibility of the concept of virtual concatenation of TUs are for further study.

3.3.7.1 Concatenation of contiguous in the higher ordern VC-3

TU-2s which are contiguous in time in the higher order VC-3 in which the are carried are concatenated together by the use of the concatenation indication ("1001" in bits 1-4, bits 5-6 unspecified, and all ones in bits 7-16 of the TU-2 pointer). The Concatenation Indication determines that the TU-2 pointer processor performs all the operations as indicated by the first TU-2 pointer in the TU-2-mc.

With this type of concatenation the VC-2-mc contains a single Virtual container POH which appears in VC-2 # 1 of the VC-2-mc.

3.3.7.2 Sequential concatenation of TU-2s in the higher order VC-4

This type of concatenation which allows the simultaneous transport of TU-2-mcs and TU-3s in hte same VC-4 is under study.

3.3.7.3 Virtual concatenation of TU-2s in the higher order VC-4

This method of cancatenation allows for the transport of a single VC-2-mc in m × TU-2 without the use of concatenation indication in the pointer bytes. The method only requires the path termination equipment to provide concatenation functionality.

Virtual concatenation requires the concatenated Tributary Unit signals at the origin of the path to be launched with the same pointer value. The so formed Tributary Units at each inteface shall be kept in a single higher order VC-4

When the higher order VC-4 is terminated, the restrictions that apply in passing the concatenated Tributary Units from one interface to another is that all of the concatenated Tributary Units are connected to a single higher order VC-4 and that the time sequencing of the concatenated Tributary Uunits in not altered.

Differences in delay of the individual concatenated VC-2 signals may occur due to pointer processing at intermediate equipment. The maximum difference in pointer value within a concatenated group at any interface is for further study. At the path termination the VC-2-mc can be reconstructed by using the pointer values for alignment.

Each concatenated VC-2 signal will carry its own POH. At the VC-2-mc path termination, the individual BIP-2s are aggregated to give a single BIP error monitor.

The details and the extensibility of the virtual concatenation method within the VC-4 are under study.

NOTE – With virtual concatenation, the available capacity of the VC-2-mc is lower than that with contiguous concatenation due to the fact that with the virtual concatenation each VC-2 carries its own POH contrary to contiguous concatenation where only the VC-2 # 1 of the VC-2-mc carries its own POH. In order to be able to interconnect VC-2-mc using different types of concatenation, the mapping of signals in VC-2-mcs should be based on the lower available capacity, namely the capacity of VC-2-mc based on virtual concatenation. Stuffing bytes should be inserted in the VC-2-mc payload based upon contiguous concatenation to accommodate the difference in capacity.

3.3.8 TU-1/TU-2 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 tributary unit within a TUG-2 is to change, then an NDF, as described in rule 5, is to be sent in all Tributary Units of the new size in the group simultaneously.

3) Additional pointer interpretation rule 7: If an NDF of "1001" and an arbitrary new size of tributary unit are received simultaneously in all of the tributary units within a TUG-2, then the coincident pointers and sizes shall replace the current ones immediately.

4) If the TU-2 pointer contains the concatenation indication, then any variation is ignored unless a consistent new pointer value is received three times consecutively.

4 Path overhead description

4.1 VC-3/VC-4/VC-4-Xc POH

The VC-3 POH is located in the first column of the 9-row by 85-column VC-3 structure.

The VC-4 POH is located in the first column of the 9-row by 261-column VC-4 structure.

The VC-4-Xc POH is located in the first column of the 9-row by 261 X-column VC-4-Xc structure.

The VC-3/VC-4/VC-4-Xc POH consists of nine bytes denoted J1, B3, C2, G1, F2, H4, Z3-Z5 (see Figures 2-2, 2-3 and 3-8). These bytes are classified as follows:

– Bytes used for end-to-end communication with independent payload function: J1, B3, C2, G1.

– Payload type specific bytes: H4, F2, Z3.

– Byte reserved for future international standardization: Z4.

– Byte which can be overwritten in an operator domain (without affecting the end-to-end performance monitoring facility of the byte B3): Z5.

VC-3/VC-4/VC-4-Xc path trace (J1) – This is the first byte in the virtual container; its location is indicated by the associated AU-*n* or TU-*n* pointer. This byte is used to repetitively transmit a high order path access point identifier so that a path receiving terminal can verify its continued connection to the intended transmitter. This path access point identifier may use a 64-byte free format string or the 16-byte E.164 format as described below within a national network. At an international boundary, only the E.164 format shall be used. Where the 16-byte format is transferred in the 64-byte field it shall be repeated four times.

Path trace format: A 16-byte frame is defined for the transmission of this E.164 numbering. The first byte of the string is a frame start marker and includes the result of a CRC-7 calculation over the previous frame. The following 15 bytes are used for the transport of 15 ASCII characters required for the E.164 numbering format. The 16-byte frame is given below:

1CCC	CCCC	Frame start marker
0XXX	XXXX	byte 2
.	.	
.	.	
0XXX	XXXX	byte 16

0XXXXXXX: ASCII character of the E.164 string.
ASCII: American Standard Code for Information Interchange

CCCCCCC: Result of the CRC-7 calculation over the previous frame. The description of this CRC-7 calculation is given in 4.3.

Path BIP-8 (B3) – One byte is allocated in each VC-3, VC-4 or VC-4-Xc 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, VC-4 or VC-4-Xc before scrambling. The computed BIP-8 is placed in the B3 byte of the current VC-3, VC-4 or VC-4-Xc before scrambling.

Signal label (C2) – One byte is allocated to indicate the composition of the VC-3/VC-4/VC-4-Xc. Table 2 which is based on Hex code, provides codes for this byte.

TABLE 2/G.709

C2 byte mapping code

MSB 1234	LSB 5678	Hex code	
0000	0000	00	unequipped
0000	0001	01	equipped-non-specific
0000	0010	02	TUG structure
0000	0011	03	locked TU
0000	0100	04	Asynchronous mapping of 34 368 and 44 736 kbit/s into the container-3
0001	0010	12	Asynchronous mapping of 139 264 kbit/s into the container-4
0001	0011	13	ATM
0001	0100	14	MAN (DQDB)
0001	0101	15	FDDI

MAN Metropolitan area network

DQDB Distributed queue dual bus

FDDI Fiber distributed data interface

NOTES

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

2 Value "1" is only to be used in cases where a mapping code is not defined in the above table. For interworking with old equipment (i.e. designed to transmit only the values "0" and "1"), the following conditions apply:

– for backward compatibility, old equipment shall interpret any value received other than value "0" as an equipped condition;

– for forward compatibility, when receiving value "1" from old equipment, new equipment shall not generate a signal label mismatch alarm.

3 Identification of the payload as whether it is an AU-4 or AU-3 structured can be indicated by checking for the Y bytes in the AU-*m* pointer area.

4 There are 247 spare codes left for future use.

5 Mapping for MAN (DQDB) and FDDI are for further study.

Path status (G1) – One byte is allocated to convey back to a VC-3/VC-4/VC-4-Xc 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 4-1, 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/VC-4-Xc path remote alarm indication is sent back by VC-3/VC-4/VC-4-Xc assembler whenever the VC-3/VC-4/VC-4-Xc assembler is not receiving a valid signal. The VC-3/VC-4/VC-4-Xc path FERF is bit 5, which is set to a "1" to indicate VC-3/VC-4/VC-4-Xc path FERF, and is otherwise set to zero. The specific received conditions under which VC-3/VC-4/VC-4-Xc path FERF is initiated are path AIS, signal failure conditions or path trace mismatch. Bits 6, 7 and 8 are not used.

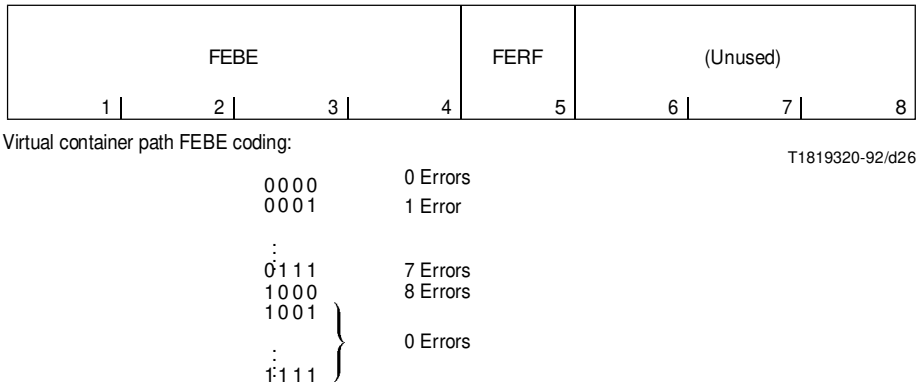

FIGURE 4-1/G.709

VC-3/VC-4 path status (G1)

Path user channel (F2, Z3) – These bytes are allocated for user communication purposes between path elements and are payload dependant.

Position indicator (H4) – This byte provides a generalized position indicator for payloads and can be payload specific (e.g. H4 can be used as a multiframe position indicator for the VC-1/VC-2.

Spare (Z4) – This byte is allocated for future, as yet undefined purposes. This byte has no defined value. The receiver is required to ignore the value contained in this byte.

Network operator byte (Z5) – This byte is allocated for specific management purposes. For tandem connection maintenance, byte Z5 is used in the following manner: bits 1 to 4 are used as an incoming error count (with the MSB of the IEC in bit 1) and bits 5 to 8 are used as a communication channel. The exact implementation of the tandem connection features are for further study.

4.2 VC-1/VC-2 POH

The byte V5 is allocated (bit 4 provisionally allocated) and bytes J2, Z6 and Z7 are provisionally allocated to the VC-1/VC-2 POH. The V5 byte is the first byte of the multiframe and its position is indicated by the TU-1/TU-2 pointer. POH bytes are only used in floating mode. The position of these bytes in the multiframe is given in Figure 3-10.

V5 byte: The byte V5 provides the functions of error checking, signal label and path status of the VC-1/VC-2 paths. The bit assignments of the V5 byte are specified in the following subclauses and are illustrated in Figure 4-2.

Virtual container path signal label coding:

L1	L2	L3	Meaning
0	0	0	Unequipped
0	0	1	Equipped – non-specific
0	1	0	Asynchronous, floating
0	1	1	Bit synchronous, floating
1	0	0	Byte synchronous, floating
1	0	1	Equipped – unused
1	1	0	
1	1	1	

a) Provisional allocation.

T1819740-93/d27

Virtual container path FEBE coding: 0 No errors

1 One or more errors

NOTES

1 Virtual container path overhead is defined only in VC-2 #1 of VC-2-mc.

2 Value "1" is only to be used in cases where a mapping code is not defined in the above table. For interworking with old equipment (i.e. designed to transmit only the values "0" and "1"), the following conditions apply:

– for backward compatibility, old equipment shall interpret any value received other than value "0" as an equipped condition;

– for forward compatibility, when receiving value "1" from old equipment, new equipment shall not generate a signal label mismatch alarm.

FIGURE 4-2/G.709

VC-1/VC-2 POH V5 byte

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 bytes V1, V2, V3 (except when used for negative justification) and V4.

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

Bit 4[1)] is a VC-1/VC-2 path remote failure indication (RFI). This bit is set to one if a failure is declared, otherwise it is set to zero. The VC-1/VC-2 path RFI is sent back by the VC-1/VC-2 assembler.

A failure is a defect that persists beyond the maximum time allocated to the transmission system protection mechanisms.

Bit 5 through 7 provide a VC-l/VC-2 signal label. Eight binary values are possible in these three bits. Value 000 indicates "VC-l/VC-2 path unequipped", and value 001 indicates "VC-l/VC-2 path equipped-non-specific payload". Three values are defined to indicate specific mappings as shown in Figure 4-2. The use of these three values is optional although they are not to be used for any other purpose. The remaining three values are reserved to be defined in other specific VC-l/VC-2 mappings. Any value received, other than 000, indicates an equipped VC-1 /VC-2 path.

1) Allocation of these bit and bytes is provisional.

Bit 8 is a VC-l /VC-2 path FERF. This bit is set to one if either a TU-l /TU-2 path AIS or a signal failure condition is being received, otherwise it is set to zero. The VC-l/VC-2 path FERF is sent back by VC-l /VC-2 assembler.

J2 byte[1)]: J2 is used to repetitively transmit a low order path access point identifier so that a path receiving terminal can verify its continued connection to the intended transmitter. This path access point identifier uses the E.164 numbering format. A 16-byte frame is defined for the transmission of this E.164 numbering. This 16-byte frame is identical to the 16-byte frame defined in 4.1 for the description of the byte Jl .

Z6[1)] byte: Z6 is under consideration to provide a tandem connection monitoring function in the same way as the byte Z5 in the higher order POH. This shall not affect the end to end performance monitoring facility of the BIP-2 of the V5 byte.

Z7[1)] byte: Z7 byte is reserved for further study.

4.3 CRC-7 polynomial algorithm

4.3.1 Multiplication/division process

A particular CRC-7 word is the remainder after multiplication by X^7 and then division (modulo 2) by the generator polynomial $X^7 + X^3 + 1$, of the polynomial representation of the previous path trace identifier multiframe (PTI).

When representing the contents of the block as a polynomial, the first bit in the block, i.e. byte 1 bit 1 should be taken as being the most significant bit. Similarly, C_l is defined to be the most significant bit of the remainder and C_7 the least significant bit of the remainder.

4.3.2 Encoding procedure

Contrary to e.g. CRC-4 procedure in 2 Mbit/s signals, the CRC-7 word is static because the data is static (the PTI represents the source address). This means that the CRC-7 checksum can be calculated *a priori* over the PTI multiframe. For consistency with existing Recommendations, the CRC-7 checksum is to be calculated over the previous multiframe. In theory this means that the 16 byte string that is loaded in a device for repetition transmission should have the checksum as the last byte although in practice it does not really matter because the PTI is static.

The encoding procedure is as follows:

i) The CRC-7 bits in the PTI are replaced by binary 0s.

ii) The PTI is then acted upon by the multiplication/division process referred to in 4.3.l.

iii) The remainder resulting from the multiplication/division process is inserted into the CRC-7 location.

The CRC-7 bits generated do not affect the result of the multiplication/division process because, as indicated in i) above, the CRC-7 bit positions are initially set to 0 during the multiplication/division process.

4.3.3 Decoding procedure

The decoding procedure is as follows:

i) A received PTI is acted upon by the multiplication/division process referred to in 4.3.1 after having its CRC-7 bits extracted and replaced by 0s

ii) The remainder resulting from the division process is then compared on a bit-by-bit basis with the CRC-7 bits received.

iii) If the remainder calculated in the decoder exactly corresponds to the CRC-7 bits received, it is assumed that the checked PTI is error free.

[1)] Allocation of these bit and bytes is provisional.

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 shows TU-1 and TU-2 sizes and formats.

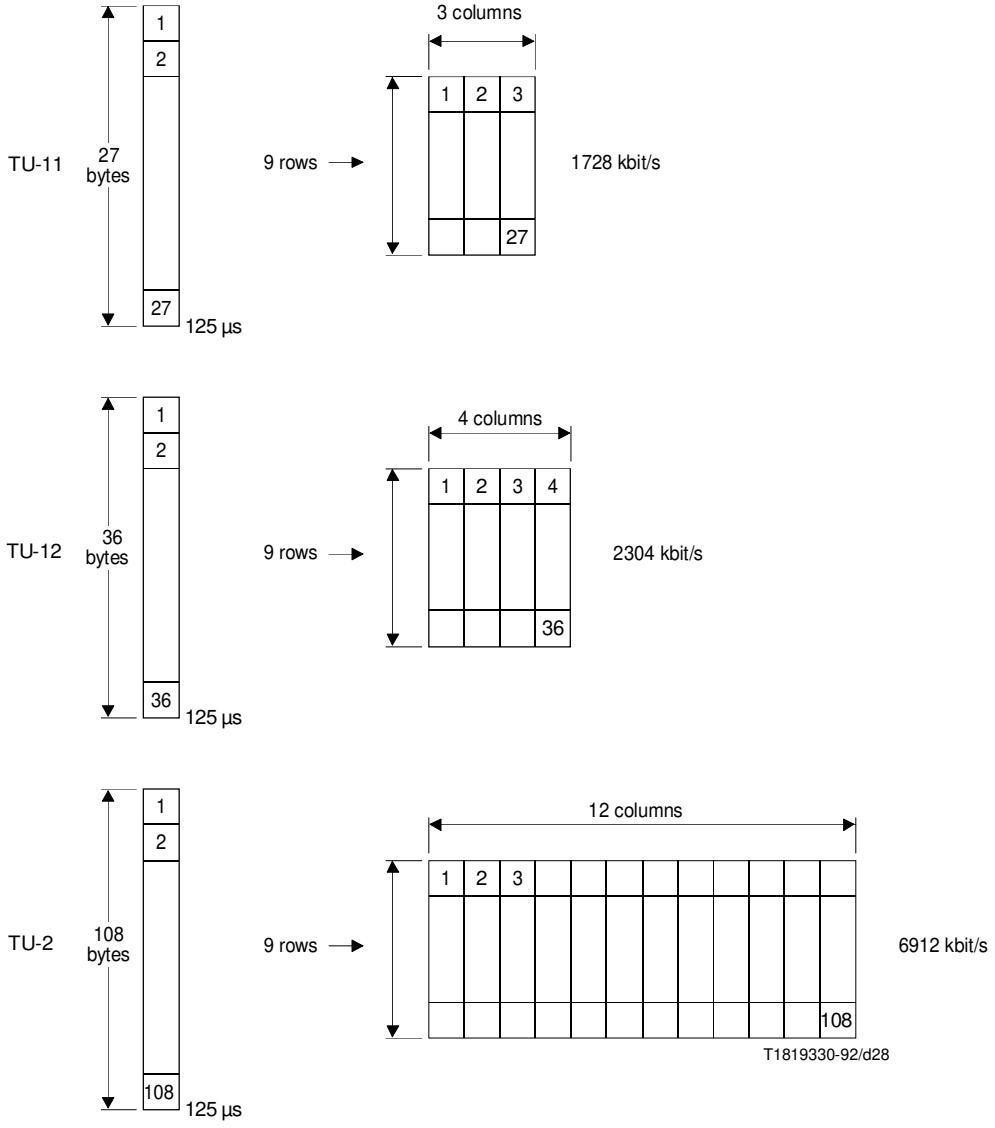

NOTE – The tributary unit pointer bytes (V1-V4) are located in byte 1 (using a four frame multiframe).

FIGURE 5-1/G.709

TU-1 and TU-2 sizes and formats

5.1 Mapping of tributaries into VC-4

5.1.1 Asynchronous mapping of 139 264 kbit/s

One 139 264 kbit/s signal can be mapped into a VC-4 of an STM-1 frame as shown in Figures 5-2 and 5-3.

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

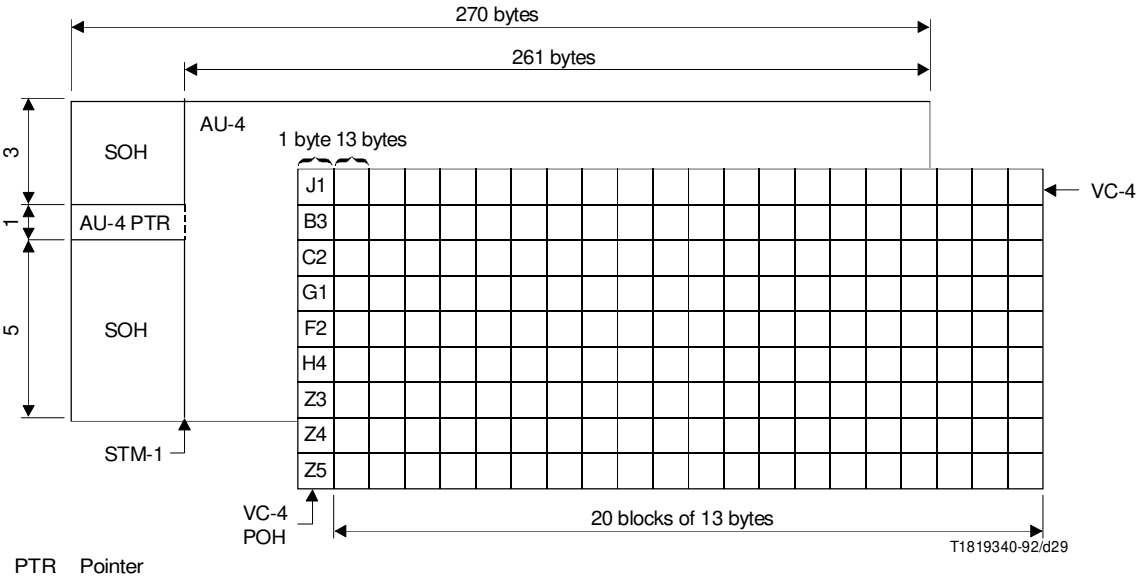

FIGURE 5-2/G.709

Mapping of VC-4 into STM-1 and block structure of VC-4 for asynchronous mapping of 139 264 kbit/s

This payload can be used to carry one 139 264 kbit/s signal:

- Each of the nine rows is partitioned into 20 blocks, consisting of 13 bytes each (Figure 5-2).
- In each row, one justification opportunity bit (S) and five justification control bits (C) are provided (Figure 5-3).
- The first byte of each block consists of:

  a) either eight information bits (I) (byte W); or

  b) eight fixed stuff bits (R) (byte Y); or

  c) one justification control bit (C) plus five fixed stuff bits (R) plus two overhead bits (O) (byte X); or

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

- The last 12 bytes of one block consist of information bits (I).

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

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

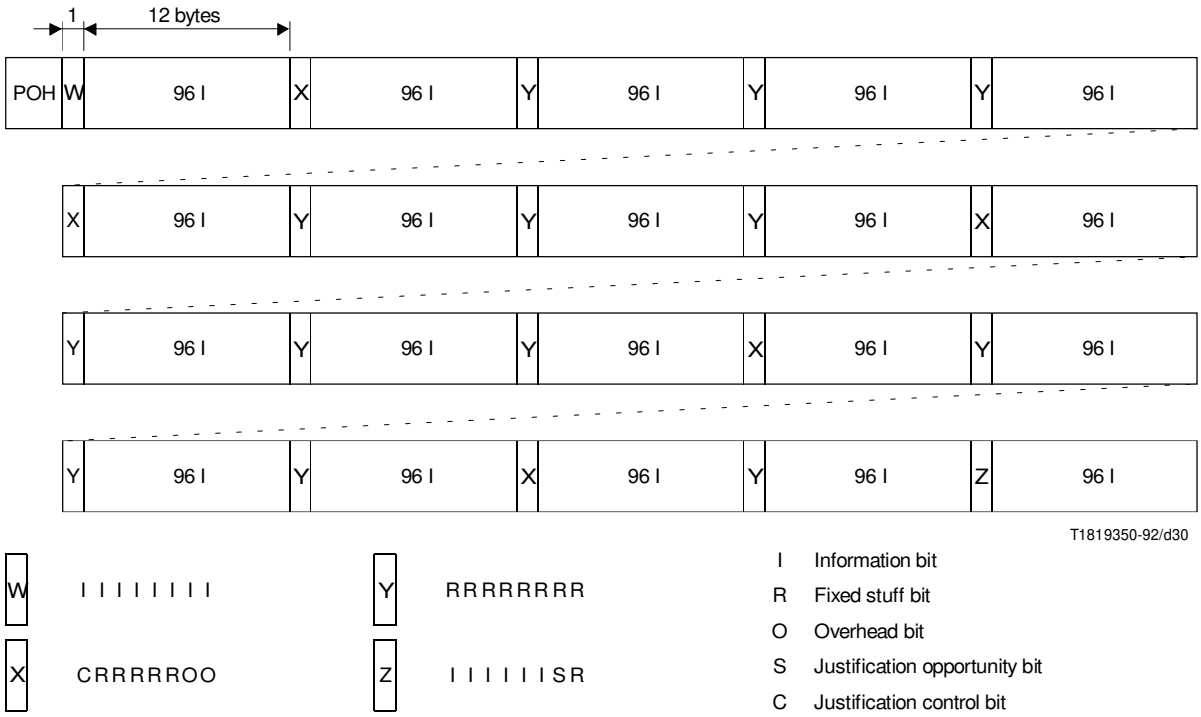

NOTE – This figure shows one row of the nine-row VC-4 container structure.

FIGURE 5-3/G.709

Asynchronous mapping of 139 264 kbit/s tributary into VC-4

The set of five justification control bits (C) in every row is used to control the corresponding justification opportunity bit (S). CCCCC = 00000 indicates that the S bit is an information bit, whereas CCCCC = 11111 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.

5.2 Mapping of tributaries into VC-3

5.2.1 Asynchronous mapping of 44 736 kbit/s

One 44 736 kbit/s signal can be mapped into a VC-3 as shown in Figure 5-4.

The VC-3 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 bits is used to control the justification opportunity (S) bit. CCCCC = 00000 indicates that the S bit is a data bit, whereas CCCCC = 11111 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.

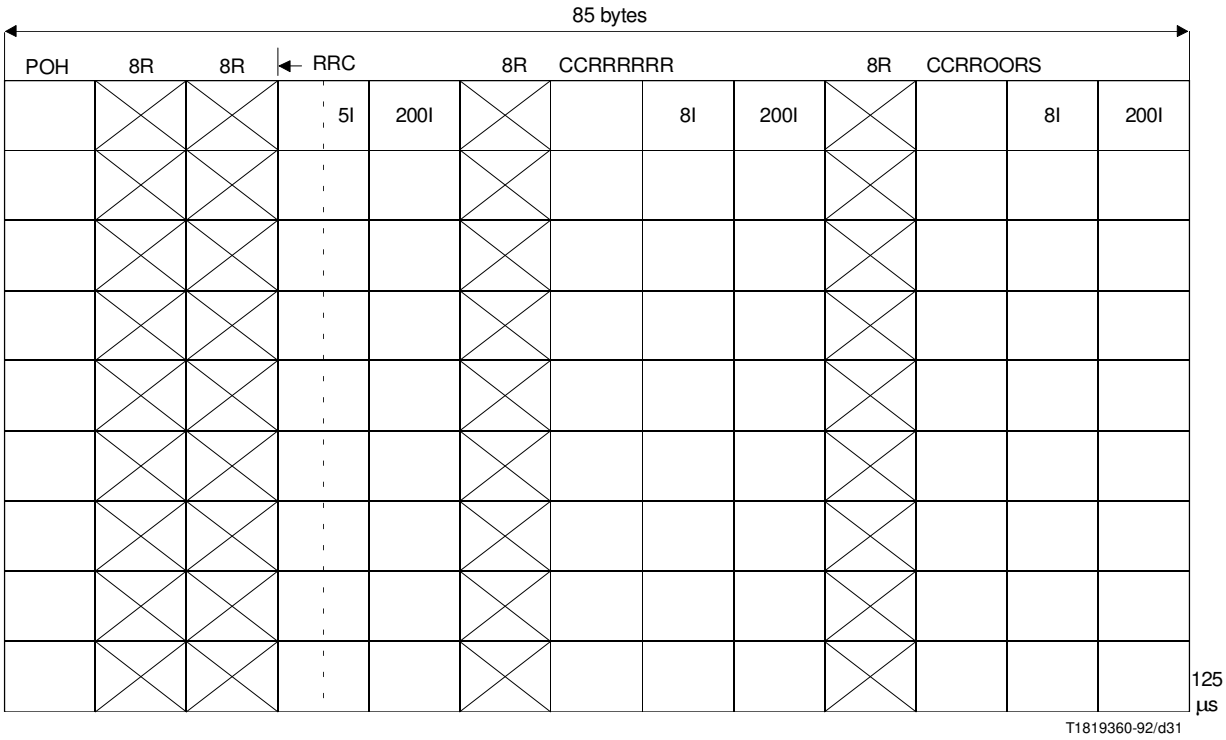

R Fixed stuff bit

C Justification control bit

S Justification opportunity bit

I Information bit

O Overhead bit

FIGURE 5-4/G.709

Asynchronous mapping of 44 736 kbit/s tributary into VC-3

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.

5.2.2 Asynchronous mapping of 34 368 kbit/s

One 34 368 kbit/s signal can be mapped into a VC-3 as shown in Figure 5-5.

In addition to the VC-3 POH, the VC-3 consists of a payload of 9×84 bytes every 125 µs. This payload is divided in three subframes, each subframe consisting of:

- 1431 information bits (I);
- two sets of five justification control bits (C_1, C_2);
- two justification opportunity bits (S_1, S_2);
- 573 fixed stuff bits (R).

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

$C_1C_1C_1C_1C_1$ = 00000 indicates that S_1 is a data bit while $C_1C_1C_1C_1C_1$ = 11111 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.

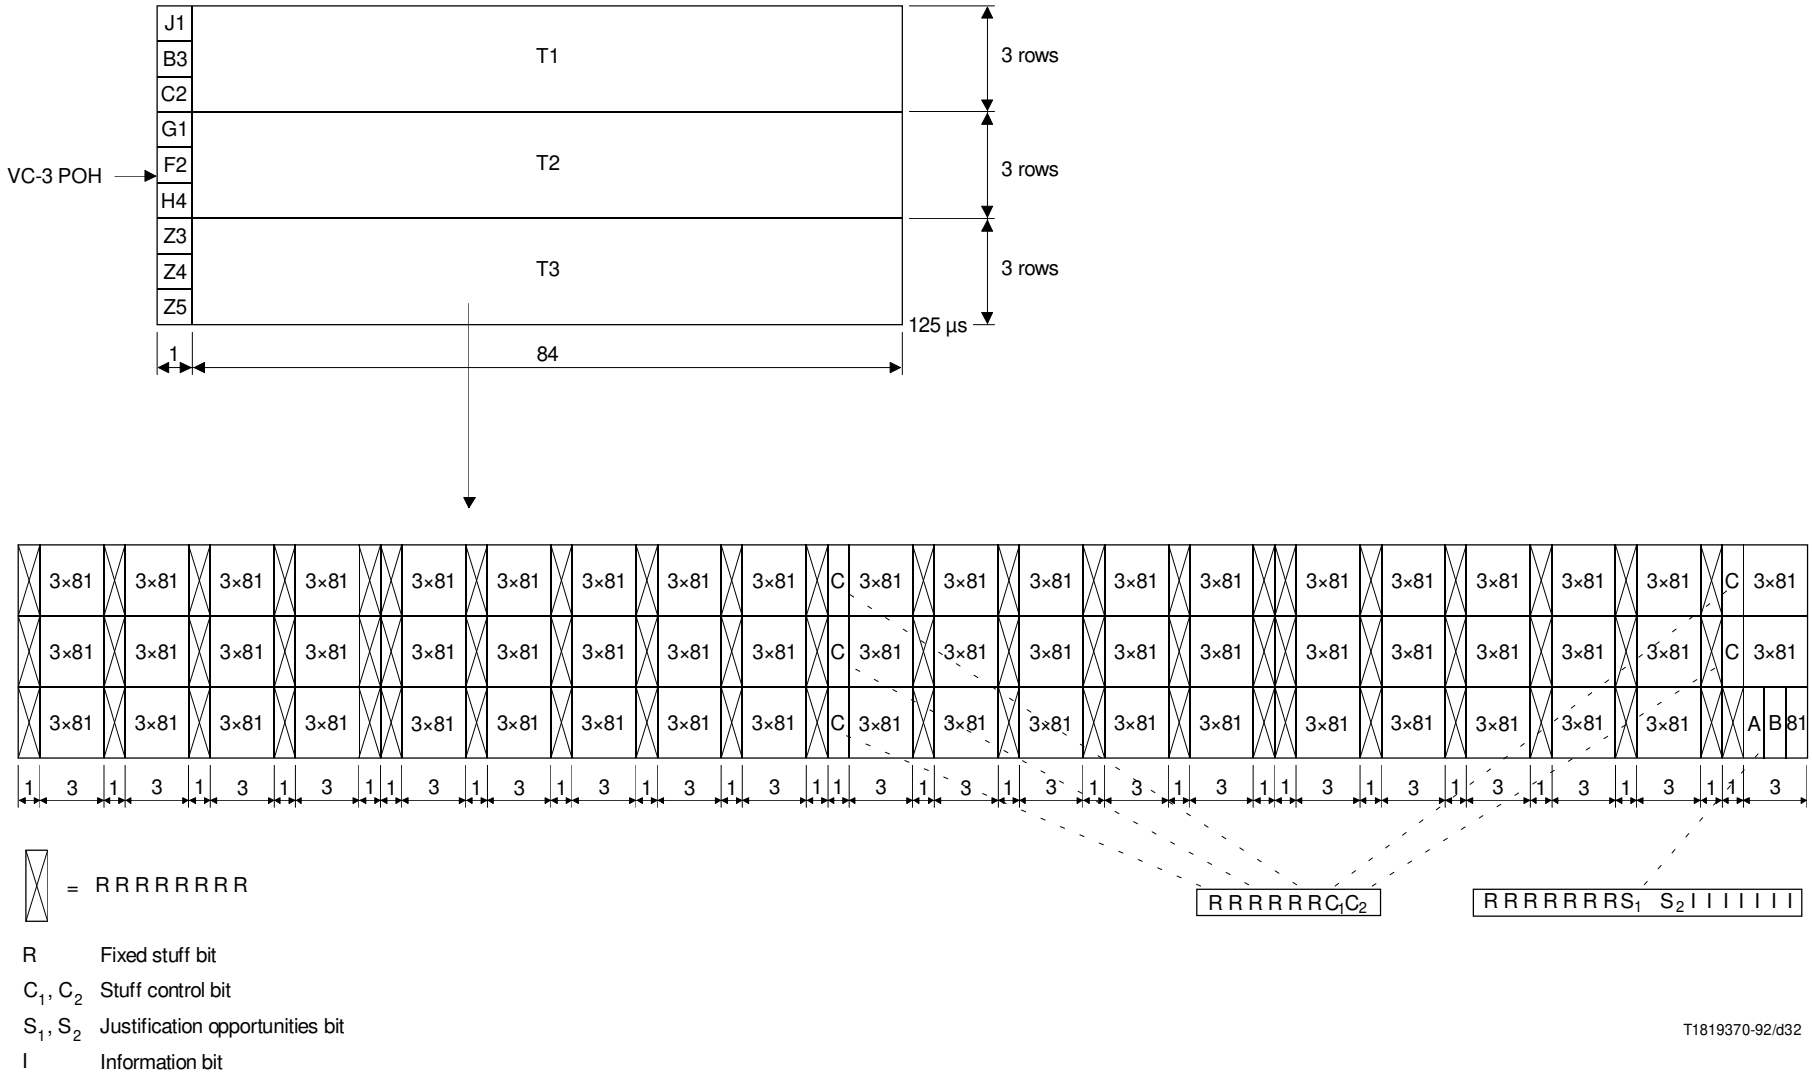

FIGURE 5-5/G.709

Asynchronous mapping of 34 368 kbit/s tributary into a VC-3

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 34 368 kbit/s. In these cases, S_1 bit should be a fixed stuff and 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 34 368 kbit/s mappings.

5.3 Mapping of tributaries into the VC-2

5.3.1 Byte synchronous mapping of 8448 kbit/s

Under study.

5.3.2 Asynchronous mapping of 6312 kbit/s

One 6312 kbit/s signal can be mapped into a VC-2. Figure 5-6 shows this over a period of 500 microseconds.

V5	I I I I I I I R	(24 × 8) I	R
R	C_1C_2OOOOI R	(24 × 8) I	R
I I I I I I I I	C_1C_2OOOOI R	(24 × 8) I	R
R	C_1C_2 I I I S_1S_2R	(24 × 8) I	
J2 a)	I I I I I I I R	(24 × 8) I	R
R	C_1C_2OOOOI R	(24 × 8) I	R
I I I I I I I I	C_1C_2OOOOI R	(24 × 8) I	R
R	C_1C_2 I I I S_1S_2R	(24 × 8) I	
Z6 a)	I I I I I I I R	(24 × 8) I	R
R	C_1C_2OOOOI R	(24 × 8) I	R
I I I I I I I I	C_1C_2OOOOI R	(24 × 8) I	R
R	C_1C_2 I I I S_1S_2R	(24 × 8) I	
Z7 a)	I I I I I I I R	(24 × 8) I	R
R	C_1C_2OOOOI R	(24 × 8) I	R
I I I I I I I I	C_1C_2OOOOI R	(24 × 8) I	R
R	C_1C_2 I I I S_1S_2R	(24 × 8) I	

125 µs
250 µs
375 µs
500 µs

T1819380-92/d33

a) Provisional allocation.

R Fixed stuff
C Justification control bit
S Justification opportunity bit
I Information bit
O Overhead bit

FIGURE 5-6/G.709

Asynchronous mapping of 6312 kbit/s tributary

In addition to the VC-2 POH, the VC-2 consists of 3152 data bits, 24 justification control bits, eight justification opportunity bits and 32 overhead communication channel bits. The remaining are Fixed Stuff bits (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.

$C_1C_1C_1$ = 000 indicates that S_1 is a data bit while $C_1C_1C_1$ = 111 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 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.

5.3.3 Bit synchronous mapping of 6312 kbit/s

The bit synchronous mapping for 6312 kbit/s tributaries is shown in Figure 5-7.

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

5.3.4 Byte synchronous mapping of 6312 kbit/s

Under study.

V5	I I I I I I I R	(24 × 8) I	R
R	10 OOOO I R	(24 × 8) I	R
I I I I I I I I	10 OOOO I R	(24 × 8) I	R
R	10 I I I R I R	(24 × 8) I	
J2 [a)]	I I I I I I I R	(24 × 8) I	R
R	10 OOOO I R	(24 × 8) I	R
I I I I I I I I	10 OOOO I R	(24 × 8) I	R
R	10 I I I R I R	(24 × 8) I	
Z6 [a)]	I I I I I I I R	(24 × 8) I	R
R	10 OOOO I R	(24 × 8) I	R
I I I I I I I I	10 OOOO I R	(24 × 8) I	R
R	10 I I I R I R	(24 × 8) I	
Z7 [a)]	I I I I I I I R	(24 × 8) I	R
R	10 OOOO I R	(24 × 8) I	R
I I I I I I I I	10 OOOO I R	(24 × 8) I	R
R	10 I I I R I R	(24 × 8) I	

125 µs
250 µs
375 µs
500 µs

T1819390-92/d34

[a)] Provisional allocation.

R Fixed stuff
I Information bit
O Overhead bit

FIGURE 5-7/G.709

Bit synchronous mapping of 6312 kbit/s tributary

5.4 Mapping of tributaries into VC-12

5.4.1 Asynchronous mapping of 2048 kbit/s

One 2048 kbit/s signal can be mapped into a VC-12. Figure 5-8 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 opportunity bits and eight overhead communication channel bits. The remaining are fixed stuff bits (R). The O bits are reserved for future overhead communication purposes.

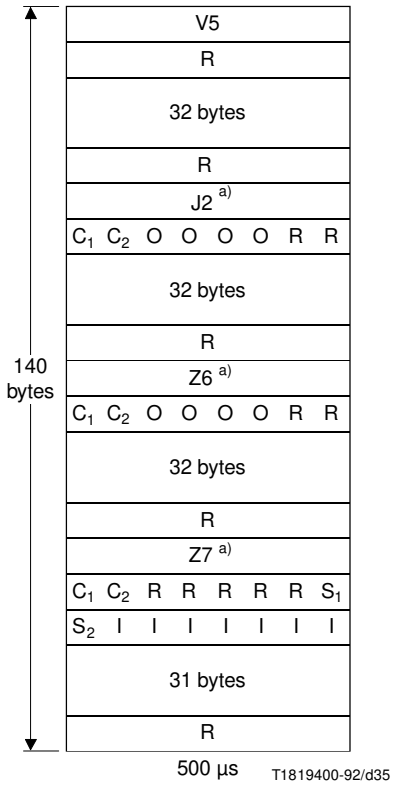

a) Provisional allocation.

I Information bit
O Overhead
C Justification control
S Justification opportunity
R Fixed stuff

FIGURE 5-8/G.709

Asynchronous mapping of 2048 kbit/s tributary

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. $C_1C_1C_1$ = 000 indicates that S_1 is a data bit while $C_1C_1C_1$ = 111 indicates that S1 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.

5.4.2 Bit synchronous mapping of 2048 kbit/s

The bit synchronous mapping for 2048 kbit/s tributaries is shown in Figure 5-9.

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

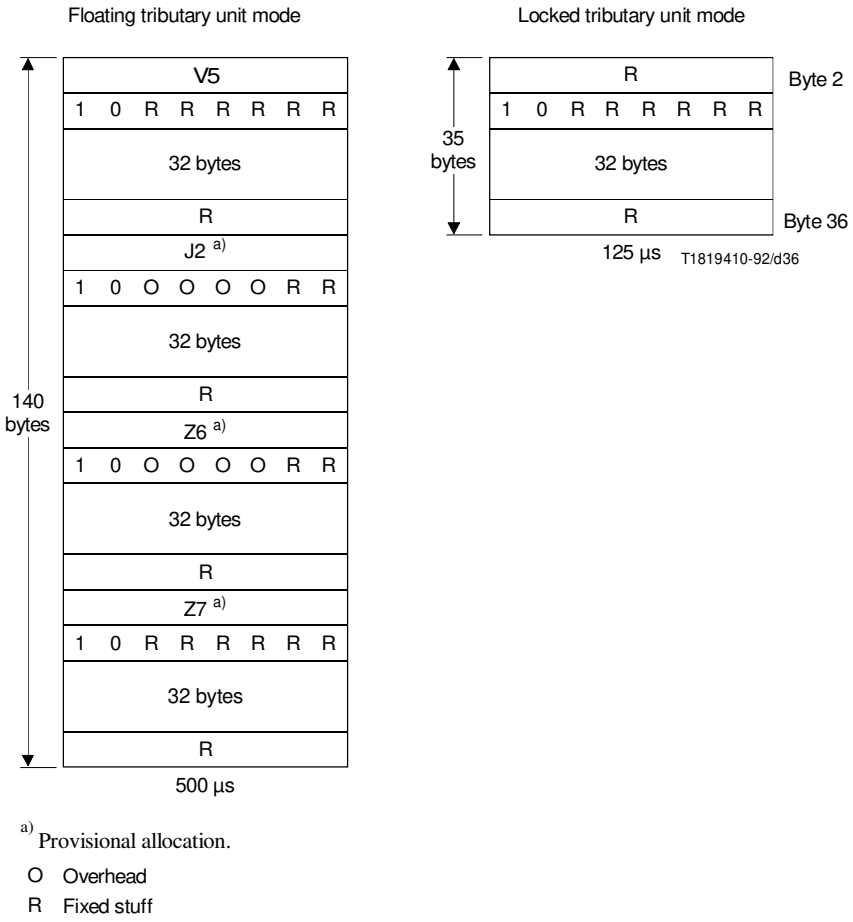

FIGURE 5-9/G.709

Bit synchronous mapping for 2048 kbit/s tributary

5.4.3 Byte synchronous mapping for 2048 kbit/s

Figure 5-10 shows byte synchronous mapping for 30-channel 2048 kbit/s tributaries Employing Channel Associated signalling (CAS). Signalling is carried in bytes 19, 54, 89 and 124 in floating mode and in byte 19 in locked mode.

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 by the position indicator byte (H4) in locked tributary unit mode. This is illustrated in Figure 5-11.

Byte synchronous mapping of 31 channel tributaries is shown in Figure 5-12. Byte 19 carries tributary channel 16.

5.5 Mapping of tributaries into VC-11

5.5.1 Asynchronous mapping of 1544 kbit/s

One 1544 kbit/s signal can be mapped into a VC-11. Figure 5-13 shows this over a period of 500 µs.

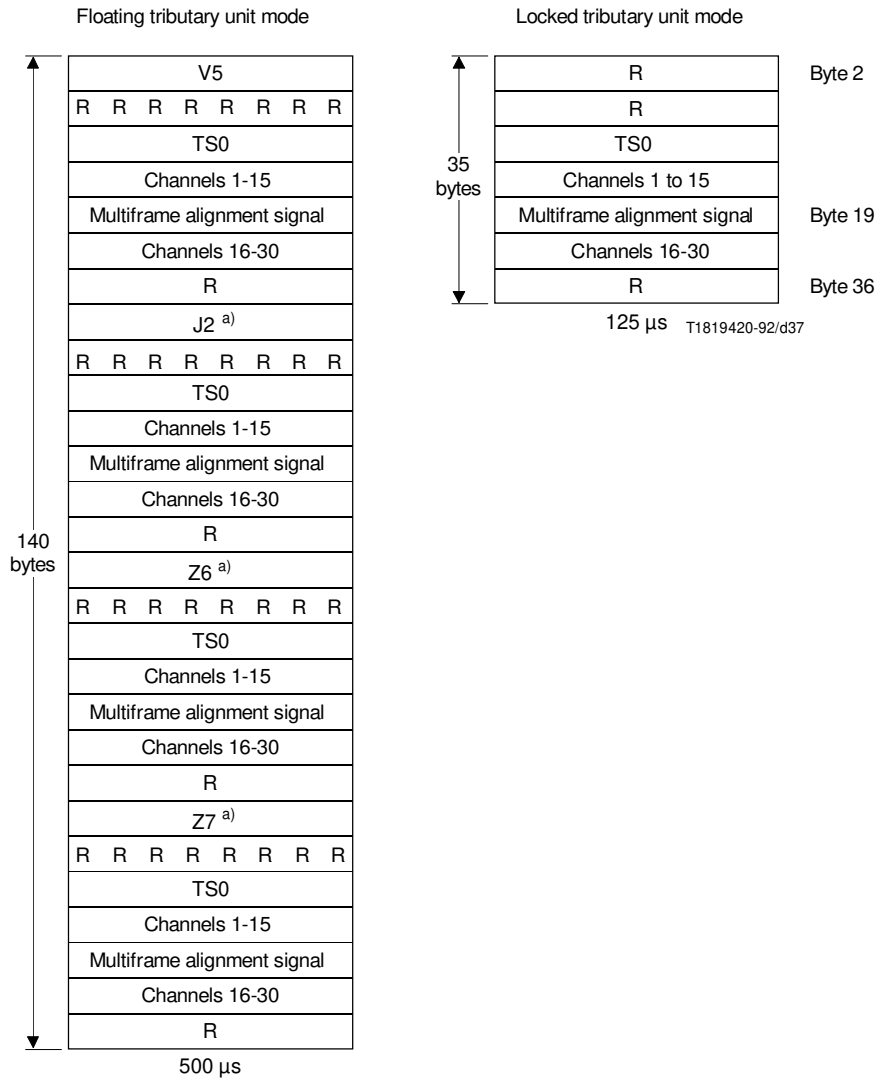

a) Provisional allocation.

FIGURE 5-10/G.709

Byte synchronous mapping for 2048 kbit/s tributary (30 channels with channel associated signalling)

Locked

H4 value				CAS format								Channel
C_3	C_2	C_1	T	S_1	S_2	S_3	S_4	S_1	S_2	S_3	S_4	
0	0	0	0	0	0	0	0	x	y	x	x	None
0	0	0	1	a	b	c	d	a	b	c	d	1/16
0	0	1	0	a	b	c	d	a	b	c	d	2/17
1	1	1	1	a	b	c	d	a	b	c	d	15/30

T1819430-92/d38

FIGURE 5-11/G.709

Out slot signalling assignments (30-channel signalling operations)

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 are fixed stuff bits (R). The 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.

$C_1C_1C_1 = 000$ indicates that S_1 is a data bit while $C_1C_1C_1 = 111$ indicates that S1 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 mapping of 1544 kbit/s

The bit synchronous mapping for 1544 kbit/s tributaries is shown in Figure 5-14.

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

5.5.3 Byte synchronous mapping for 1544 kbit/s

The byte synchronous mapping for 1544 kbit/s tributaries is shown in Figure 5-15.

The S_1, S_2, S_3 and S_4 bits contain the signalling for the 24 x 64 kbit/s channels. The phase of the signalling bits can be indicated in the P_1 and P_0 bits in floating tributary unit mode, and by the position indicator byte (H4) in locked mode. This is illustrated in Figure 5-16. The usage of the P bits is optional, since the common channel signalling methods and some of the channel associated signalling methods (Recommendation G.704) do not need the P bits. The out slot signalling assignments for one of the channel associated signalling methods is shown in Figure 5-17.

5.6 VC-11 to VC-12 conversion for transport by a TU-12

When transporting a VC-11 in a TU-12, the VC-11 is adapted by adding fixed stuff with even parity as shown in Figure 5-18. Thus the resulting TU-12 payload can be monitored and cross-connected in the network as though it were a VC-12 with its BIP value unchanged while preserving end-to-end integrity of the real VC-11 path.

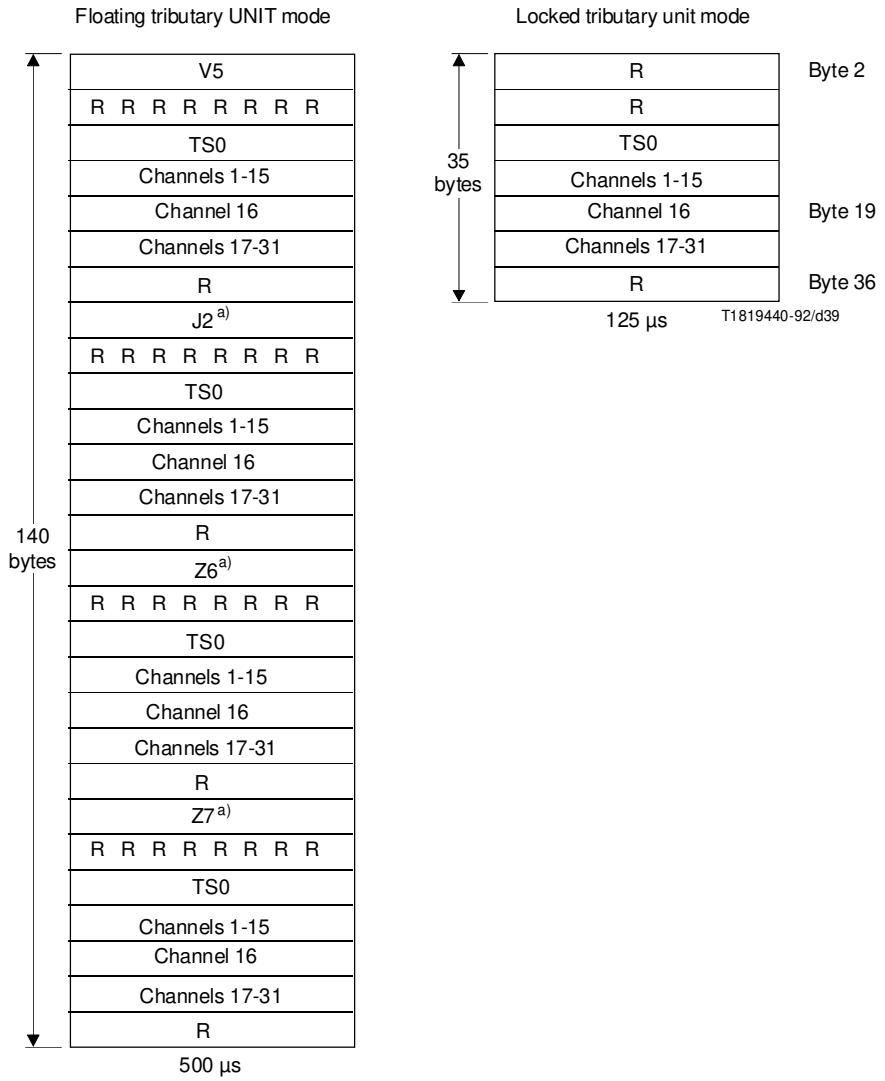

a) Provisional allocation.

FIGURE 5-12/G.709

Byte synchronous mapping for 2048 kbit/s tributary (31 channels with common channel signalling)

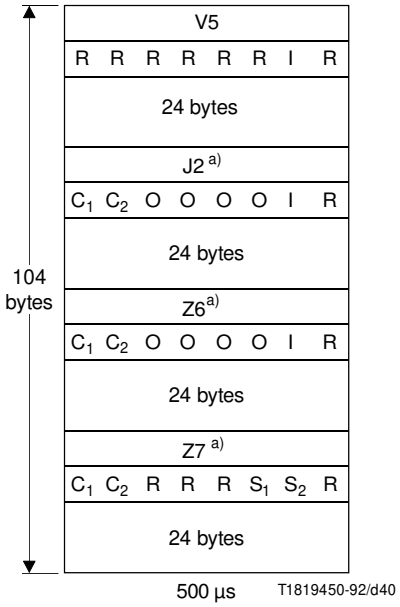

a) Provisional allocation.

I Information bit
O Overhead
C Justification control
S Justification opportunity
R Fixed stuff

FIGURE 5-13/G.709

Asynchronous mapping of 1544 kbit/s tributary

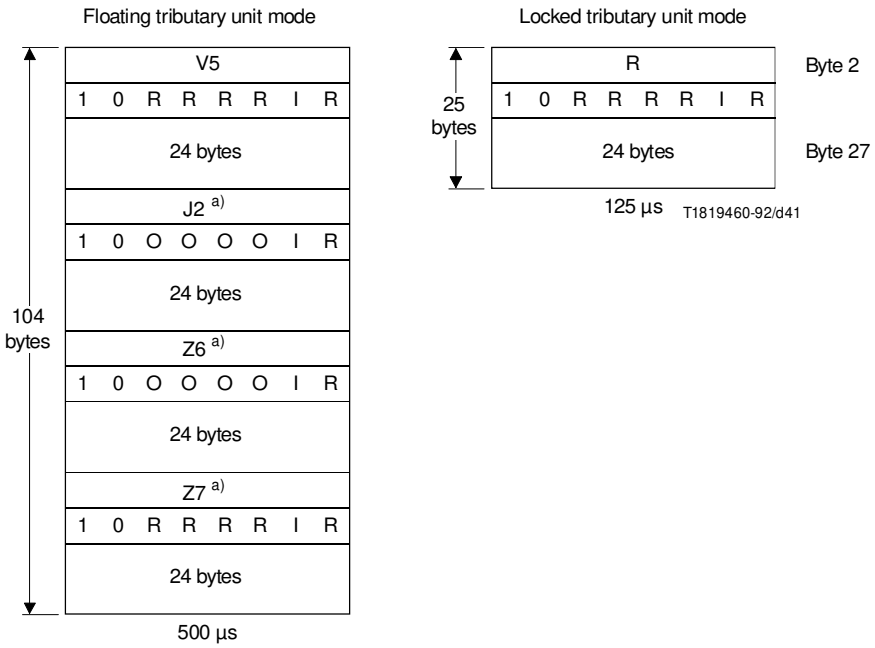

a) Provisional allocation.

I Information bit

O Overhead

R Fixed stuff

NOTE – O bits are currently not defined in the locked tributary unit mode.

FIGURE 5-14/G.709

Bit synchronous mapping for 1544 kbit/s tributary

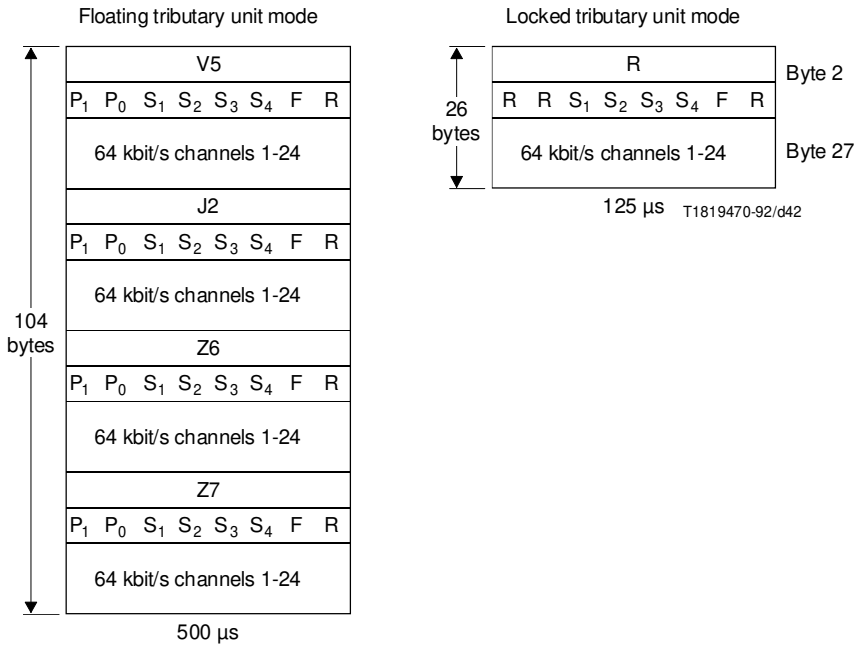

FIGURE 5-15/G.709

Byte synchronous mapping for 1544 kbit/s tributary

Locked																	
					Floating												
					Signalling												
H4 value					2 state				4 state				16 state				
P_1	P_0	S_2	S_1	T	S_1	S_2	S_3	S_4	S_1	S_2	S_3	S_4	S_1	S_2	S_3	S_4	P_1	P_0
0	0	0	0	0	A_1	A_2	A_3	A_4	A_1	A_2	A_3	A_4	A_1	A_2	A_3	A_4	0	0
0	0	0	0	1	A_5	A_6	A_7	A_8	A_5	A_6	A_7	A_8	A_5	A_6	A_7	A_8	0	0
0	0	0	1	0	A_9	A_{10}	A_{11}	A_{12}	A_9	A_{10}	A_{11}	A_{12}	A_9	A_{10}	A_{11}	A_{12}	0	0
0	0	0	1	1	A_{13}	A_{14}	A_{15}	A_{16}	A_{13}	A_{14}	A_{15}	A_{16}	A_{13}	A_{14}	A_{15}	A_{16}	0	0
0	0	1	0	0	A_{17}	A_{18}	A_{19}	A_{20}	A_{17}	A_{18}	A_{19}	A_{20}	A_{17}	A_{18}	A_{19}	A_{20}	0	0
0	0	1	0	1	A_{21}	A_{22}	A_{23}	A_{24}	A_{21}	A_{22}	A_{23}	A_{24}	A_{21}	A_{22}	A_{23}	A_{24}	0	0
0	1	0	0	0	A_1	A_2	A_3	A_4	B_1	B_2	B_3	B_4	B_1	B_2	B_3	B_4	0	1
0	1	0	0	1	A_5	A_6	A_7	A_8	B_5	B_6	B_7	B_8	B_5	B_6	B_7	B_8	0	1
0	1	0	1	0	A_9	A_{10}	A_{11}	A_{12}	B_9	B_{10}	B_{11}	B_{12}	B_9	B_{10}	B_{11}	B_{12}	0	1
0	1	0	1	1	A_{13}	A_{14}	A_{15}	A_{16}	B_{13}	B_{14}	B_{15}	B_{16}	B_{13}	B_{14}	B_{15}	B_{16}	0	1
0	1	1	0	0	A_{17}	A_{18}	A_{19}	A_{20}	B_{17}	B_{18}	B_{19}	B_{20}	B_{17}	B_{18}	B_{19}	B_{20}	0	1
0	1	1	0	1	A_{21}	A_{22}	A_{23}	A_{24}	B_{21}	B_{22}	B_{23}	B_{24}	B_{21}	B_{22}	B_{23}	B_{24}	0	1
1	0	0	0	0	A_1	A_2	A_3	A_4	A_1	A_2	A_3	A_4	C_1	C_2	C_3	C_4	1	0
1	0	0	0	1	A_5	A_6	A_7	A_8	A_5	A_6	A_7	A_8	C_5	C_6	C_7	C_8	1	0
1	0	0	1	0	A_9	A_{10}	A_{11}	A_{12}	A_9	A_{10}	A_{11}	A_{12}	C_9	C_{10}	C_{11}	C_{12}	1	0
1	0	0	1	1	A_{13}	A_{14}	A_{15}	A_{16}	A_{13}	A_{14}	A_{15}	A_{16}	C_{13}	C_{14}	C_{15}	C_{16}	1	0
1	0	1	0	0	A_{17}	A_{18}	A_{19}	A_{20}	A_{17}	A_{18}	A_{19}	A_{20}	C_{17}	C_{18}	C_{19}	C_{20}	1	0
1	0	1	0	1	A_{21}	A_{22}	A_{23}	A_{24}	A_{21}	A_{22}	A_{23}	A_{24}	C_{21}	C_{22}	C_{23}	C_{24}	1	0
1	1	0	0	0	A_1	A_2	A_3	A_4	B_1	B_2	B_3	B_4	D_1	D_2	D_3	D_4	1	1
1	1	0	0	1	A_5	A_6	A_7	A_8	B_5	B_6	B_7	B_8	D_5	D_6	D_7	D_8	1	1
1	1	0	1	0	A_9	A_{10}	A_{11}	A_{12}	B_9	B_{10}	B_{11}	B_{12}	D_9	D_{10}	D_{11}	D_{12}	1	1
1	1	0	1	1	A_{13}	A_{14}	A_{15}	A_{16}	B_{13}	B_{14}	B_{15}	B_{16}	D_{13}	D_{14}	D_{15}	D_{16}	1	1
1	1	1	0	0	A_{17}	A_{18}	A_{19}	A_{20}	B_{17}	B_{18}	B_{19}	B_{20}	D_{17}	D_{18}	D_{19}	D_{20}	1	1
1	1	1	0	1	A_{21}	A_{22}	A_{23}	A_{24}	B_{21}	B_{22}	B_{23}	B_{24}	D_{21}	D_{22}	D_{23}	D_{24}	1	1

T1819480-92/d43

FIGURE 5-16/G.709

Out slot signalling assignments
(24-channel signalling operations)

Frame number	n	$n+1$	$n+2$	$n+3$	$n+4$	$n+5$	$n+6$	$n+7$
Use of S_i bit (i = 1, 2, 3, 4)	Fs	Y_1	Y_2	Y_3	Y_4	Y_5	Y_6	X
(Note 1)	(Note 2)	(Note 3)						(Note 5)

NOTES

1 Each S_i (i = 1, 2, 3, 4) constitutes an independent signalling multiframe over eight frames. S_i includes the phase indicator in itself, so that the PP-bits cannot be used for the phase indicator.

2 The Fs bit is either alternate 0, 1 or the following 48 bit digital pattern:

A101011011 0000011001 1010100111 0011110110 10000101

For the 48-bit digital pattern, the "A" bit is usually fixed to state 1 and is reserved for optional use. The pattern is generated according to the following primitive polynomial (refer to Recommendation X.50):

$$x^7 + x^4 + 1$$

3 Y_j bit (j = 1 to 6) carries channel associated signalling or maintenance information. When the 48 bit pattern is adopted as Fs frame alignment signal, each Y_j bit (j = 1 to 6) can be multiframed, as follows.

$$Y_{j1}, Y_{j2}, \ldots, Y_{j12}$$

Y_{j1} bit carries the following 16-bit frame alignment pattern generated according to the same primitive polynomial as for the 48-bit pattern.

A011101011011000

The "A" bit is usually fixed to 1 and is reserved for optional use. Each Y_{ji} (i = 2 to 12) bit carries channel associated signalling for sub-rate circuits and/or maintenance information.

4 S_i bits (Fs, Y_1, . . ., Y_6 and X), all at state 1 indicate alarm indication signal (AIS) for six 64 kbit/s channels.

5 The X-bit is usually fixed to state 1. When backward AIS for six 64 kbit/s channels is required to be sent, the X-bit is set to state 0.

FIGURE 5-17/G.709

Out slot signalling assignments (24-channel signalling operations)

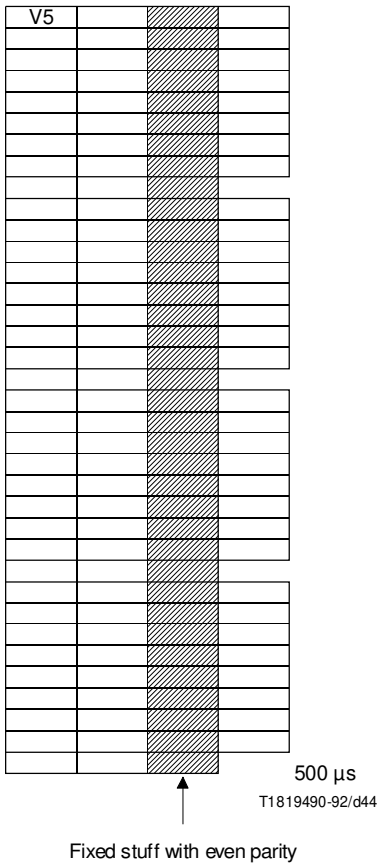

FIGURE 5-18/G.709

Conversion of VC-11 to VC-12 for transport by TU-12

5.7 Floating and locked mode conversion

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

In the floating tributary unit mode four consecutive 125 µs VC-*n* frames (*n* = 11, 12, 2) are organized into a 500 µs multiframe, the phase of which is indicated by the position indicator byte (H4) in the VC-*m* POH (*m* = 3, 4). This 500 µs tributary unit multiframe is shown in Figure 5-19.

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

Figure 5-19 illustrates the conversion between floating and locked TU modes for each of the three tributary unit 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 virtual container 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.

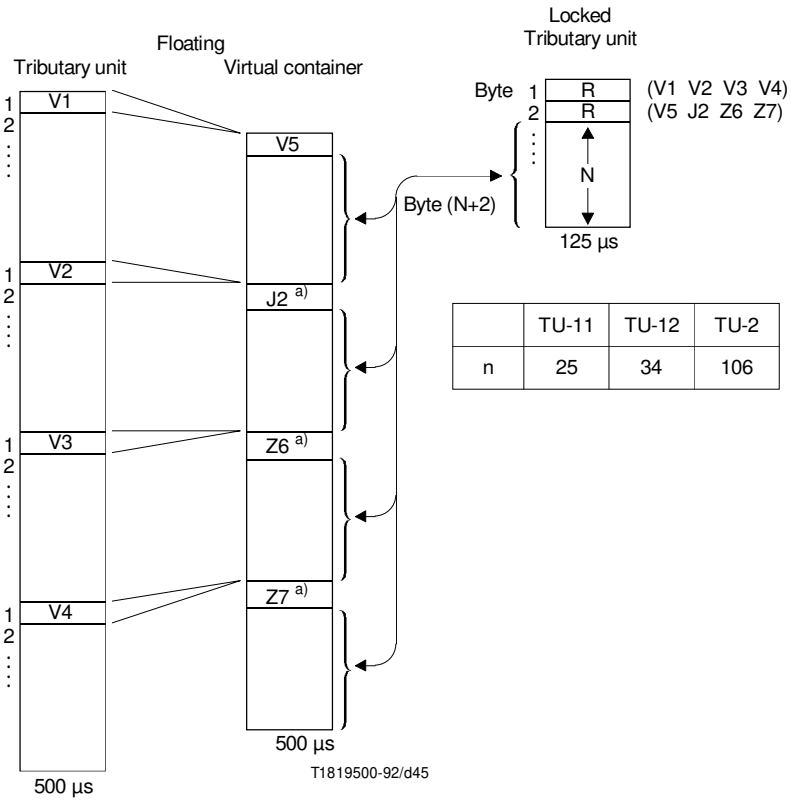

	TU-11	TU-12	TU-2
n	25	34	106

a) Provisional allocation.

V1 PTR-1
V2 PTR-2
V3 PTR-3 (action)
V4 Reserved
V5 VC-1/VC-2 POH
PTR Pointer

FIGURE 5-19/G.709

Conversion between floating and locked Tributary Unit modes

5.8 Mapping of ATM cells

The mapping of ATM cells is performed by aligning the byte structure of every cell with the byte structure of the virtual container used including the concatenated structure (VC-*x* or VC-*x*-mc, $x \geq 1$). Since the relevant container-*x* or container-*x*-mc capacity may not be an integer multiple of the ATM cell length (53 octets), a cell is allowed to cross the container-*x* boundary.

The ATM cell information field (48 bytes) shall be scrambled before mapping into the VC-*x* or VC-*x*-mc. In the reverse operation, following termination of the VC-*x* or VC-*x*-mc signal, the ATM cell information field will be descrambled before being passed to the ATM layer. A self-synchronising scrambler with generator polynomial $x^{43} + 1$ shall be used. The scrambler operates for the duration of the cell information field. During the 5-octet header the scrambler operation is suspended and the scrambler state retained. During the start-up procedure the scrambler seed has an all 1s value. Cell information field scrambling is required to provide security against false cell delineation and cell information field replicating the STM-N frame alignment word.

When the VC-*x* or VC-*x*-mc is terminated, the cell must be recovered. The ATM cell header contains a header error control (HEC) field which is used in a similar way to a frame alignment word to achieve cell delineation. This HEC method uses the correlation between the header bits to be protected by the HEC (32 bits) and the control bit of the HEC (8 bits) introduced in the header after computation with a shortened cyclic code with generating polynomial $g(x) = x^8 + x^2 + x + 1$.

The remainder from this polynomial is then added to the fixed pattern "01010101" in order to improve the cell delineation performance. This method is similar to conventional frame alignment recovery where the alignment word is not fixed but varies from cell to cell.

More information on HEC cell delineation is given in Recommendation I.432.

5.8.1 Mapping of ATM cells into VC-4

The ATM cell stream is mapped into container-4 with its octet boundaries aligned with the container-4 byte boundaries. The container-4 is then mapped into VC-4 together with the VC-4 POH (see Figure 5-20). The ATM cell boundaries are thus aligned with the VC-4 byte boundaries. Since the container-4 capacity (2340 octets) is not an integer multiple of the cell length (53 octets), a cell may cross a container-4 boundary.

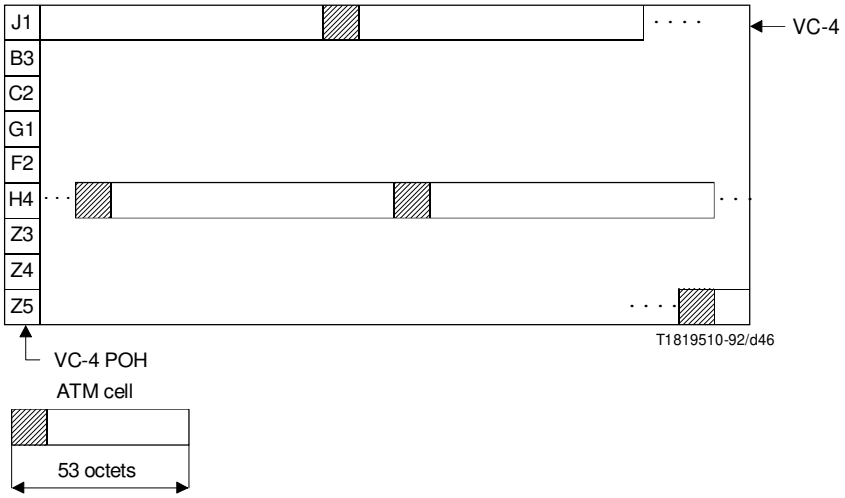

FIGURE 5-20/G.709

Mapping of ATM cells in the VC-4

5.8.2 Mapping of ATM cells into the VC-4-Xc

The ATM cell stream is mapped into a container-4-Xc with its octet boundaries aligned with the container-4-Xc byte boundaries. The container-4-Xc is then mapped into VC-4-Xc together with the VC 4-Xc POH and (X-1) columns of fixed stuff (see Figure 5-21). The ATM cell boundaries are thus aligned with the VC-4-Xc byte boundaries. Since the container-4-Xc capacity (2340.X octets) is not an integer multiple of the cell length (53 octets), a cell may cross a container-4-Xc boundary.

The H4 byte is not used for indicating cell offset and is reserved for future use.

5.8.3 Mapping of ATM cells into other VCs

Detailed mappings are under study.

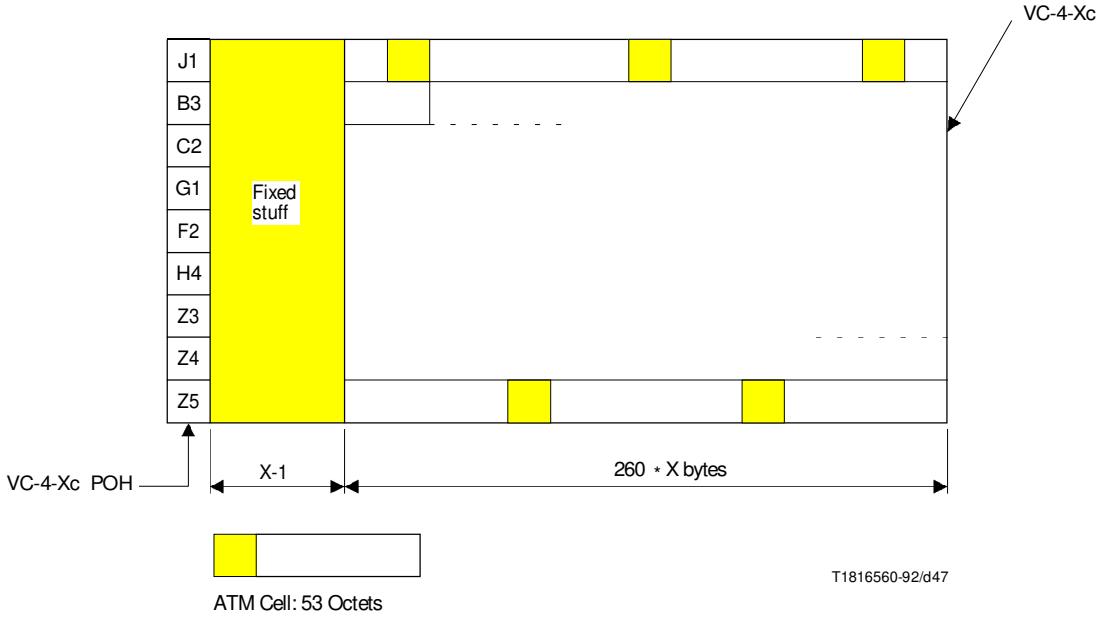

FIGURE 5-21/G.709

Mapping of ATM cells into a VC-4-Xc