Forward Error Correction in Optical Networks

P. Michael Henderson
michael.henderson@cox.net

March 27, 2001

Many different kinds of impairments affect optical communications, including optical nonlinearity, chromatic dispersion, polarization mode dispersion, amplifier induced four wave mixing, and others. As optical communications rates increase many of these impairments become more pronounced.

As we move from OC-192/STM-64 to OC-768/STM-256, for example, amplifier spacings would have to be decreased if we were limited to existing technologies. This would be extremely expensive and disruptive as fiber optic cables were dug up and additional amplifiers added. Fortunately, new techniques, including Raman amplifiers and Forward Error Correction (FEC), have been developed which will mitigate this problem. This paper examines the FEC techniques developed primarily in the ANSI T1X1.5 committee and submitted for ITU ratification as formal recommendations.

Two different approaches have been developed. The first is an in-band[1] SONET/SDH approach documented in a revision to ITU recommendation G.707 ratified in October 2000. The second is a digital wrapper approach documented in the new ITU recommendation G.709 which was ratified in February 2001. I will cover the in-band approach first but my major focus will be G.709.

While this paper is intended as a tutorial, it only presents a high level overview of these two proposed recommendations. There is a lot more in the recommendations than is covered here. Those seeking additional detail should consult the references.

In-band FEC

The concept of in-band FEC is that the line rate will not be affected – the line rates after addition of the FEC will be exactly the same OC-N or STM-N rates that we are familiar with today. For this reason, in-band FEC is sometimes called the "rate preserving" technique.

Adding FEC to SONET/SDH has a number of issues. First, the SONET/SDH frame size is fixed in time (125 μs) and not in the number of bits or octets. As the line rate increases, the number of octets in a SONET/SDH frame increases. This is a problem for FEC, because an FEC is selected to apply over a certain number of bits or symbols. Second, there are a limited number of available, unused octets in the overhead portion of SONET/SDH. In fact, at the lowest levels of SONET/SDH, there are essentially no unused octets available for use by the FEC.

For these reasons, the in-band FEC defined in the revision of G.707 only applies to SONET/SDH rates of OC-48/STM-16. This fixes the length of the frame, and allows sufficient unused overhead octets for the redundant FEC bits. Higher level SONET/SDH rates are handled by N/48 (for SONET) or N/16 (for

[1] The term "in-band" is somewhat unfortunate because both techniques have the FEC "in-band". Use of the term "in-band" somehow implies that the alternate approach is "out of band," which is just not true. However, that's the term the ANSI group used when developing this technique and that's the term I will use in this paper.

SDH) disinterleaving of the signal, taking 16 consecutive octets at a time. Rates below OC-48/STM-16 are not defined for FEC. In this discussion, I will focus on the OC-48/STM-16 signal.

For purposes of applying the FEC, the transport overhead of the SONET/SDH frame is arranged as shown in Figure 1. Note that each octet is viewed "edge on" so that we have a frame which is eight bits deep. Since we are working with an OC-48/STM-16 frame, we see 144 overhead bits in each row, and 4,176 payload bits per row. The total number of bits per row across the frame is the total of the overhead and payload bits, or 4,320 bits. Eight FEC codes are applied to each row, one for each bit of the octets.

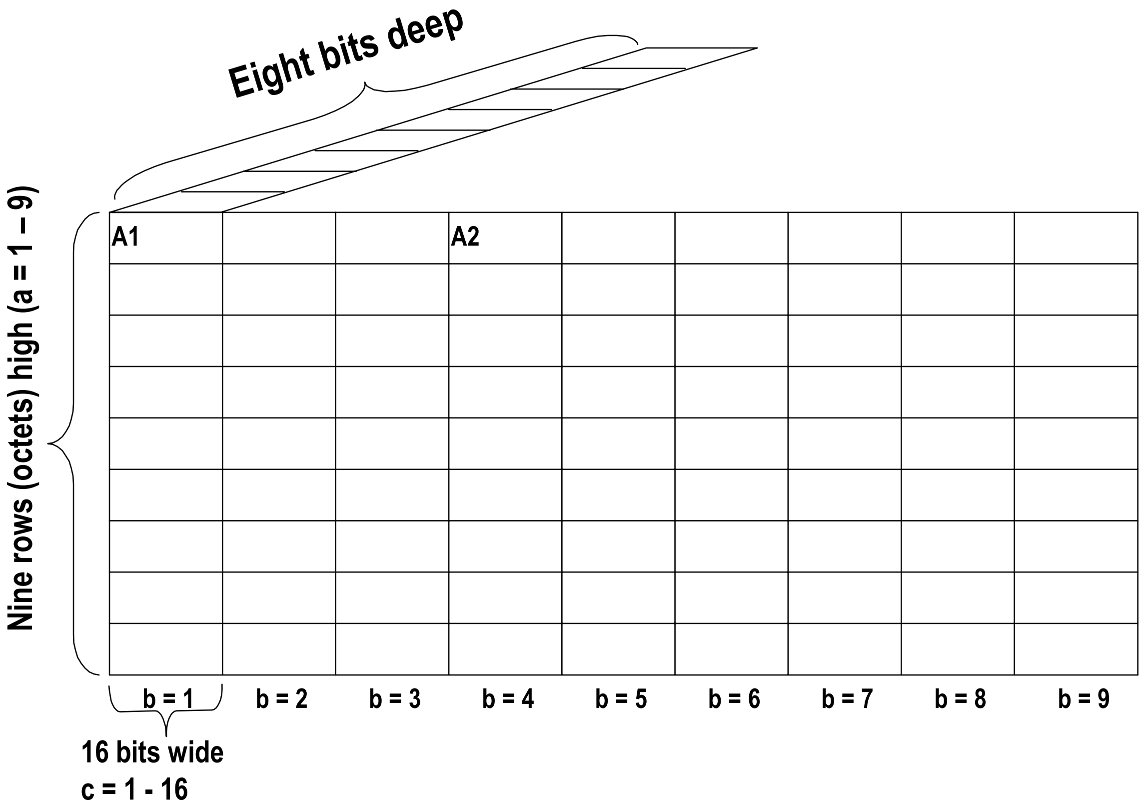

Figure 1: The SONET/SDH frame for purposes of applying the FEC. Note that each octet is arranged so that its eight bits are into the page, giving the frame a depth of eight "slices". An FEC is computed for each slice, row by row. This diagram shows ONLY the transport overhead and none of the payload.

The G.707 recommendation defines a three coordinate vector system to describe the location of bits in the first slice (all interior slices are the same as the first slice). The coordinate system is of the form **S**(a,b,c), where a (1 to 9) represents the row number, b (1 to 9) represents one of a set of 16 contiguous bits[2] in the first slice (indicated by b = 1, b = 2, etc. in Figure 1), and c (1 to 16) which represents the bits within each b. Thus, the first A2 framing character would have the coordinates **S**(1,4,1).

Note that this coordinate system holds for higher levels of SONET/SDH, with only changes to c, the number of bits in each b. The maximum value of c is always N/3 for SONET and N for SDH. This is in

[2] It is 16 bits only for OC-48/STM-16. See the next paragraph in the text for the number of bits for rates above OC-48/STM-16.

line with the method for handling higher levels of SONET/SDH described earlier. The disinterleaving is done 16 octets at a time, so for OC-192/STM-64, the value of c will vary from 1 to 64.

The FEC used is a Bose-Chaudhuri-Hocquenghem-3 code[3] (BCH-3). The specific code utilized is a shortened code[4] derived from a (8191, 8152) parent code. Specifically, the code covers 4320 information bits, utilizing 39 redundant bits. In the standard, these redundant bits are referred to as a_n, where n varies from 0 to 38 (don't confuse this "a_n" with the "a" in the vector coordinates).

A BCH-3 code can correct up to three bits in error. Since the codes are eight way interleaved, a burst error of up to 24 bits can be corrected. Now, we'll investigate where those 39 a_n bits are placed in the SONET/SDH frame.

Since each c only consist of 16 bits for OC-48/STM-16, the 39 redundant bits are broken into three groups of 13 bits each, and these three groups are distributed across the b's, one group of 13 bits to each b (leaving three bits unused in each selected b).

The general philosophy taken when deciding where to place the redundant bits is that the encoding and decoding delay must be minimized. For example, even if the redundant bits are placed in the same row as the payload bits covered by the FEC, the encoder must buffer the entire row so that it can compute the FEC and place the redundant bits in the designated locations in the overhead area for that row. Since there are 4320 octets per row, it will take about 14 μs to transmit (or receive) one row at OC-48/STM-16. When decoding, the entire row must be held in memory so that corrections can be applied after the decoding computations are completed.

The existing usage of the overhead octets limit placement of the redundant bits, however. For example, Row 1 is unavailable due to the A1/A2/J0 octets. Row 4 is dedicated to pointers. For this reason, the standards committee chose to include the redundant bits either in the same row with the payload bits covered by the FEC, or in the next row.

---

[3] Named for mathematicians Raj Chandra Bose, Dijen K. Ray-Chaudhuri, and P. A. Hocquenghem.
[4] A code is "shortened" by assuming that all the unused bits are a constant, usually zero, at both encode and decode.

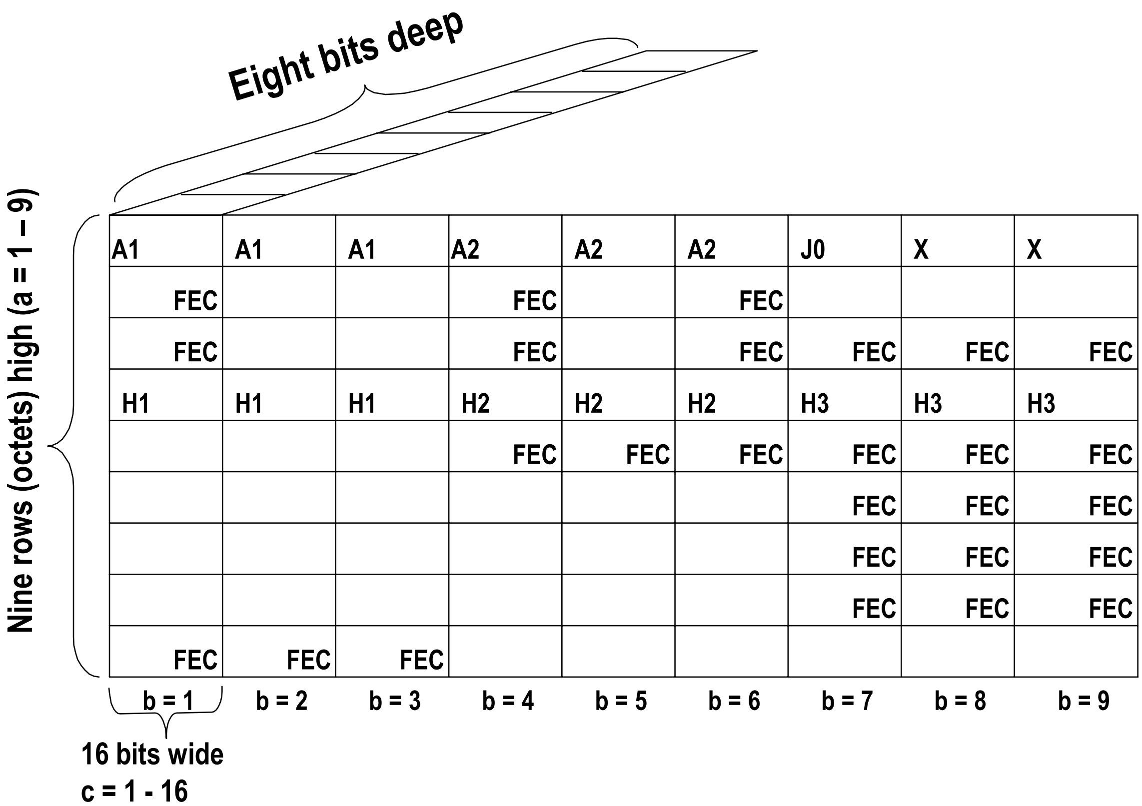

Figure 2: Location of the FEC redundant bits in the OC-48/STM-16 frame.
Each "FEC" indicator represents 13 bits.

The location of the a_n FEC redundant bits are given in Table 1. The column headed "Row" indicates the location (row) of the payload bits covered by the redundant bits. The other columns indicate the location (a,b) of the redundant bits. Although not specified in the table, each 13 redundant bits are located in bits 4 through 16 of each b, meaning that there are three unused bits preceding each set of 13 bits.

Note that the payload is defined to be the full 4320 bits of a SONET/SDH row. Obviously, when the redundant bits are in the same row, they cannot be included in the FEC calculations. Some other regenerator and section overhead octets are excluded from the FEC calculations. See the standard for more details.

Row	(a, b) for a_n $26 \leq n \leq 38$	(a, b) for a_n $13 \leq n \leq 25$	(a, b) for a_n $0 \leq n \leq 12$
1	2,1	2,4	2,6
2	3,1	3,4	3,6
3	3,7	3,8	3,9
4	5,4	5,5	5,6
5	5,7	5,8	5,9
6	6,7	6,8	6,9
7	7,7	7,8	7,9
8	8,7	8,8	8,9
9	9,1	9,2	9,3

Table 1: x, y values for the location of each set of 13 bits of the 39 FEC redundant bits for each row. (source: ITU Study Group 15, Report R 71, p120)

Since FEC is being added to SONET/SDH, there must be a technique to signal the receiver as to whether FEC is being used by the transmitter. This is done by the FEC Status Indicator (FSI), which is located immediately prior to a_n ($0 \leq n \leq 12$) for row 3, at **S**(3,9,3). While we only see one bit in the first slice, there are 7 bits behind this bit for a total of 8 bits. It is the last two bits of this octet which are used for the FSI. A value of 01 indicates FEC "on" while a value of 00 indicates FEC "off". States 10 and 11 are invalid. To switch the FEC from on to off (or from off to on), the status bits are changed from 01 to 00 (or 00 to 01) seven frames prior to the actual state change. The state change takes place on the 8^{th} frame.

So now that we have FEC in the SONET/SDH frame, what good does it do? The answer depends upon the error rate and the characteristics of the errors. Assuming that the errors are independent and Gaussian, the output BER for a given input BER is shown in Figure 3. There are several things which should be noted from this figure. First, with an input BER of 10E-10, the output BER is approximately 10E-30. Fiber links are usually engineered for an "end-of-life" BER not to exceed 10E-10. This means that errors essentially disappear for normally operating links.

Note also that for an input BER of 10E-6 (a seriously degraded link), the output BER is about 10E-14. This allows network providers to provide acceptable error rates to their customers, even when their links seriously degrade. Finally, note that no gain is provided at 10E-3 (around this BER, the FEC could actually increase the number of errors).

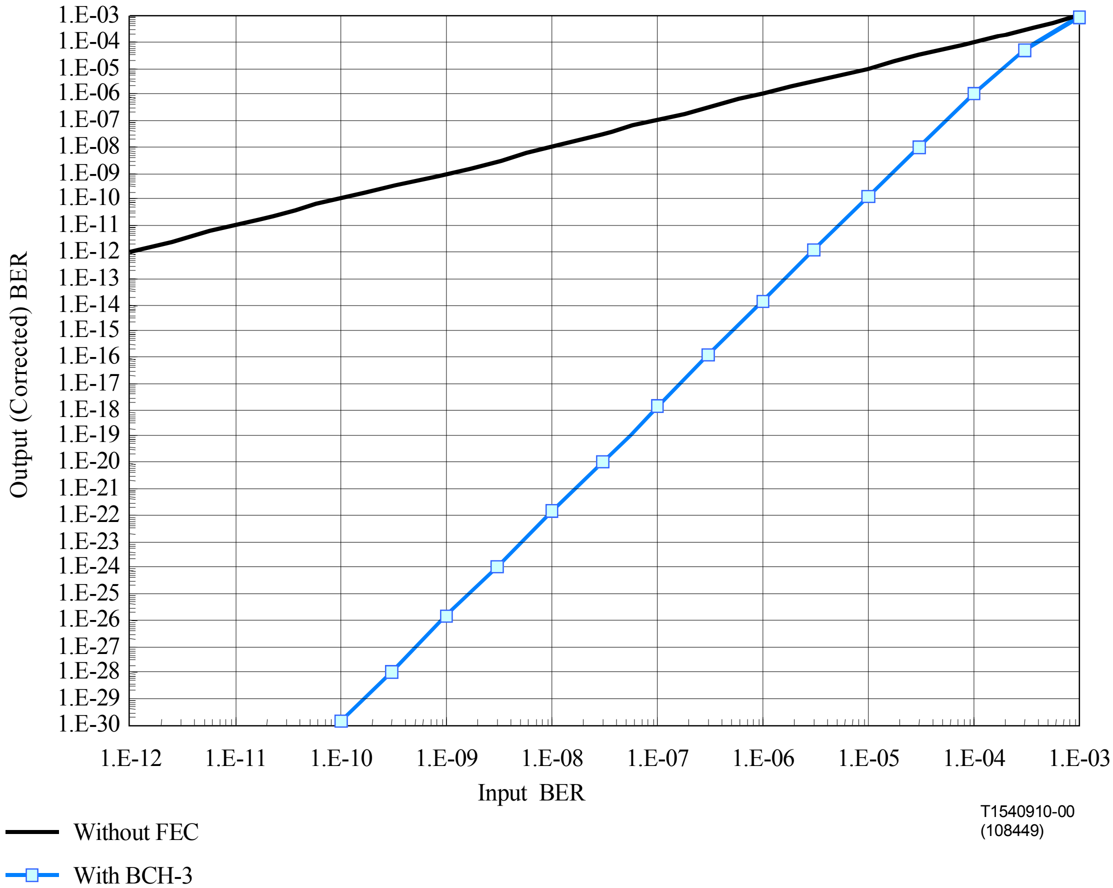

Figure 3: Input BER verses output BER for the BCH-3 code used in G.707.
(source: ITU Study Group 15, Report R 71, p163)

Digital Wrapper[5]

The digital wrapper approach defined in the new ITU recommendation G.709 is entirely different from the in-band FEC approach. First, it utilizes a fixed length frame (fixed in number of bits rather than time) and second, it utilizes an overspeed line rate so that the payload rate is the same as the SONET/SDH rates. The digital wrapper is functionally positioned between the data source (which could be a SONET/SDH framer) and the transceiver. See the Figure 4 below.

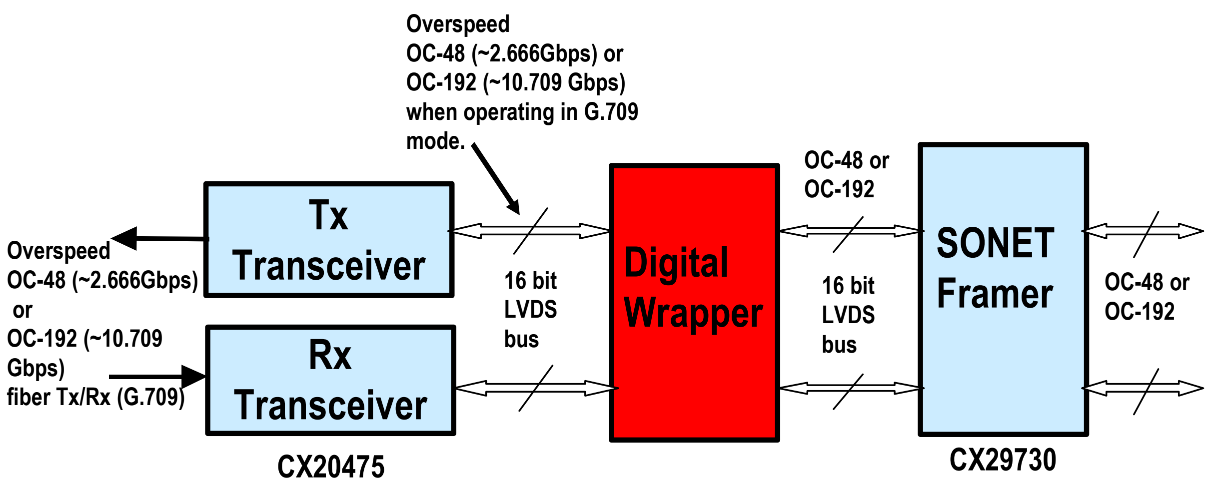

Figure 4: Placement of the digital wrapper function.

The basic building block of the digital wrapper is a Reed-Solomon code[6], specifically, an RS(255,239) code. Remember that Reed-Solomon codes operate on symbols instead of bits. In this application, the symbol is an octet so the block is 255 octets in length. The payload is 239 octets, meaning that there are 16 redundant octets in the code. An RS(255,239) code can correct up to 8 symbols in error and detect (but not correct) up to 16 symbols in error.

One of the payload octets is taken for framing and management functions, as shown in Figure 5. Having only one octet for framing and management functions is not very valuable so this basic frame is worked into a larger frame structure.

[5] The techniques described here are the most current proposals as of the date of publication of this paper. It's possible that some of the techniques, especially those in the tandem connection monitoring area, will change prior to the final ITU recommendation. Check back in late February 2001 for an update.
[6] Named for mathematicians Irving S. Reed and Gustave Solomon.

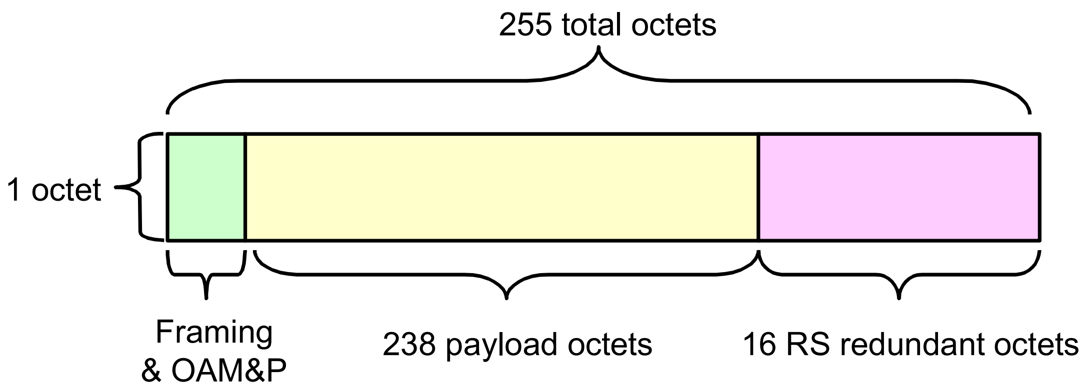

Figure 5: The basic digital wrapper Reed-Solomon code frame.

The first move into a larger frame structure is to octet interleave 16 of these basic frames into a single row of a larger frame. Then, four rows are used to create a frame. See Figure 6 which attempts to show this interleaving and larger frame structure.

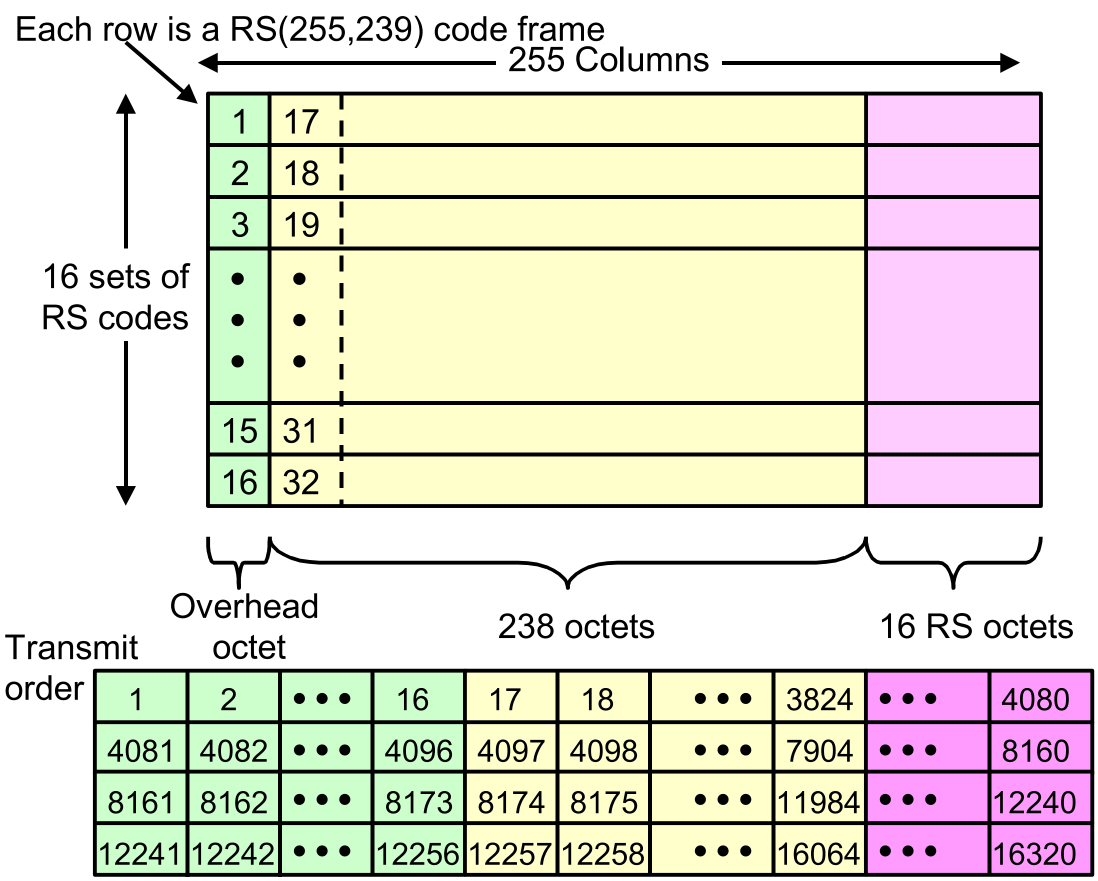

Figure 6: The individual RS code frames are 16 interleaved into a row of the larger G.709 frame. The complete frame consists of four of these 16 interleaved rows.

Note that this four-row frame structure provides an overhead area of four rows by 16 octets, for a total of 64 framing and management octets. Since there are 16 redundant octets per RS code, and 16 are interleaved per row, there are 256 redundant octets per row, and 1024 redundant octets per four-row frame. This leaves 15,232 octets for actual payload.

Now 64 framing and management octets are really not that many. To provide for additional management information, there is a one-octet field in the overhead area which increments each frame, going from 0 to 255. This creates a "multiframe" structure, allowing management information and commands to extend over the multiframe.

Note that the digital wrapper simply carries bits (or perhaps we should say octets) – it is essentially blind to the type of payload traffic. Since the digital wrapper is designed to carry any type of traffic, including legacy SONET/SDH traffic, the line rate is adjusted so that the payload rate (not including any overhead) is equal to the standard SONET/SDH rates.

The performance of the RS(255,239) code is given in Table 2 below. Note that the performance is quite a bit better than the in-band FEC, due primarily to the higher ratio of redundant octets to payload octets. With any reasonable line BER, errors are essentially eliminated.

Input BER	Output BER
10^{-4}	5×10^{-15}
10^{-5}	6.3×10^{-24}
10^{-6}	6.4×10^{-33}

Table 2: Input verses output BER for a RS(255,239) code, as used in G.709 (source: ITU recommendation G.975, p10)

G.709 defines three rates, identified as 1, 2, and 3, corresponding to payload rates of OC-48/STM-16, OC-192/STM-64, and OC-768/STM-256. The actual line rates are higher because of the FEC and the management overhead. The exact line rates will be discussed later, after a discussion of the overhead.

The overhead, excluding the Reed-Solomon redundant octets, is divided into three sections: Transport, Data, and Payload. This division is shown in Figure 7.

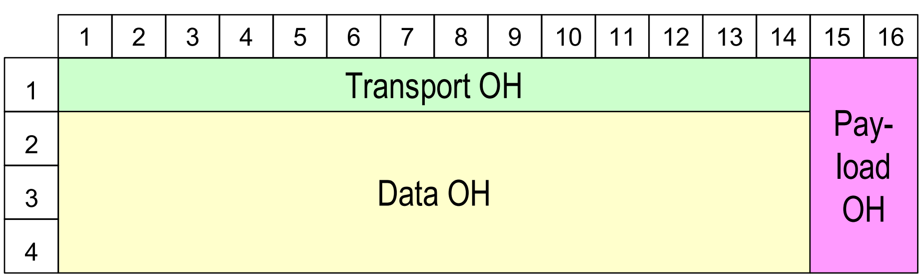

Figure 7: The division of overhead in the G.709 frame.

The transport overhead is made up of the following.

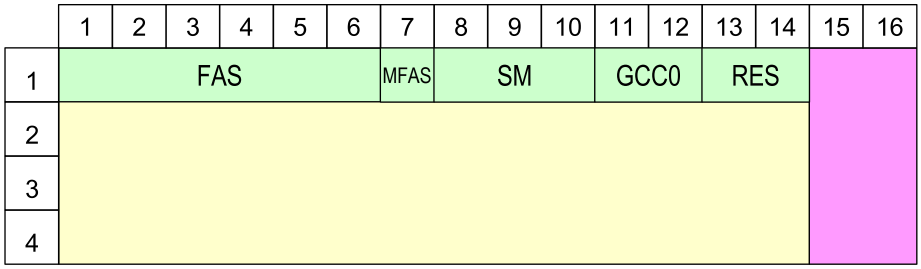

Figure 8: Details of the transport overhead.

The first six octets of the transport overhead are the frame alignment signal (FAS) octets. These are the same A1/A2 octets used in SONET/SDH (A1= 1111 0110 and A2 = 0010 1000) and are used to provide framing for the signal.

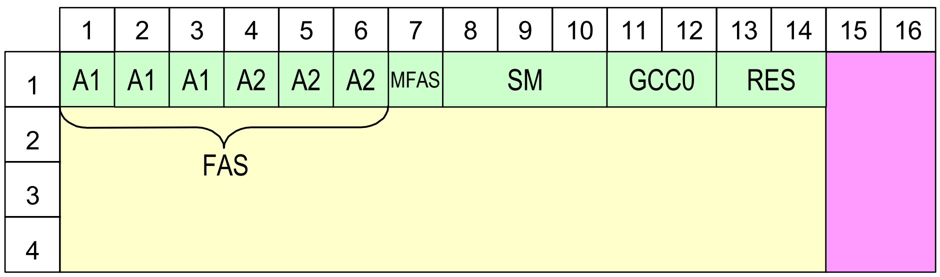

Figure 9: The framing alignment signal octets.

To provide sufficient bit timing content for clocking purposes and to reduce the probability of non-FAS octets duplicating the FAS octets, and thus causing an incorrect frame alignment, a scrambler is applied over all octets of the frame except the FAS octets. See the ITU recommendation for details of the scrambler and its initialization.

The next octet, the multi-frame alignment signal (MFAS) is a one octet field which is incremented by one each frame from zero to 255. When 255 is reached, the field is reset to zero. This provides a multi-frame structure which allows commands and management data to extend over a number of frames.

The section monitoring (SM) octets are used for the trail trace identifier (TTI), parity (BIP-8) and the backward error indicator (BEI), backward defect indicator (BDI), and incoming alignment error (IAE). See Figure 10 which indicates the position of these fields.

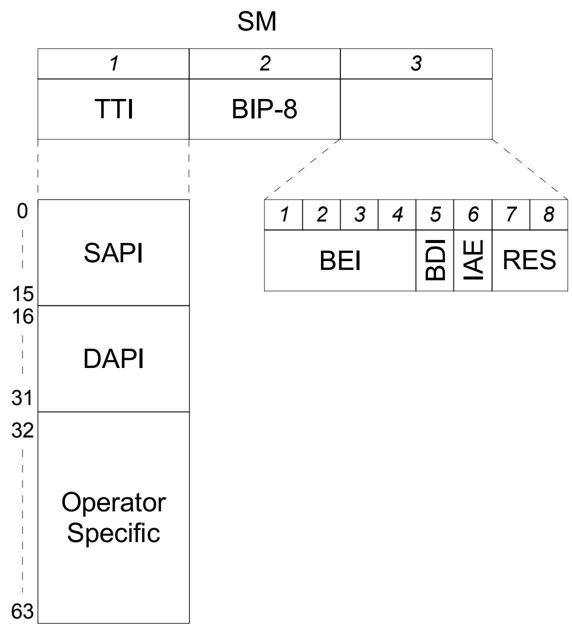

Figure 10: The section monitoring fields. (Source: T1X1.5 document 2000-246, p39)

The TTI field is a multi-frame field, with a length of 64 frames. The information is repeated four times per multiframe. The information in the multiframe is the source access point identifier (SAPI) and the destination access point identifier (DAPI).

The next two octets in the transport overhead are for the general communications channel zero (GCC0), the contents of which are undefined. The last two octets are reserved for future standardization.

Next, let's examine the data overhead.

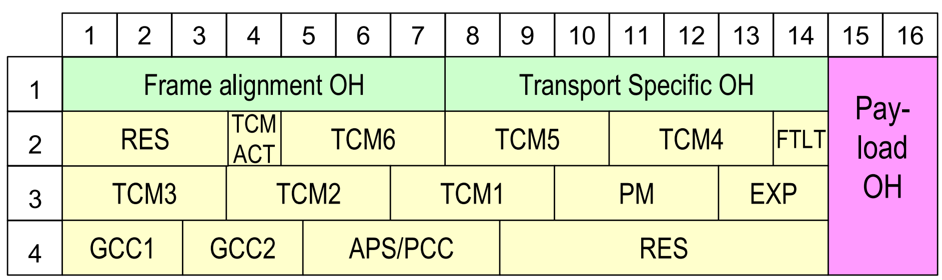

Figure 11: The data overhead.

First, look at row 3, columns 10, 11, & 12. This is the path monitoring field (PM) which contains information similar to the section monitoring field. See Figure 12 for details of the PM fields. Path monitoring is end-to-end.

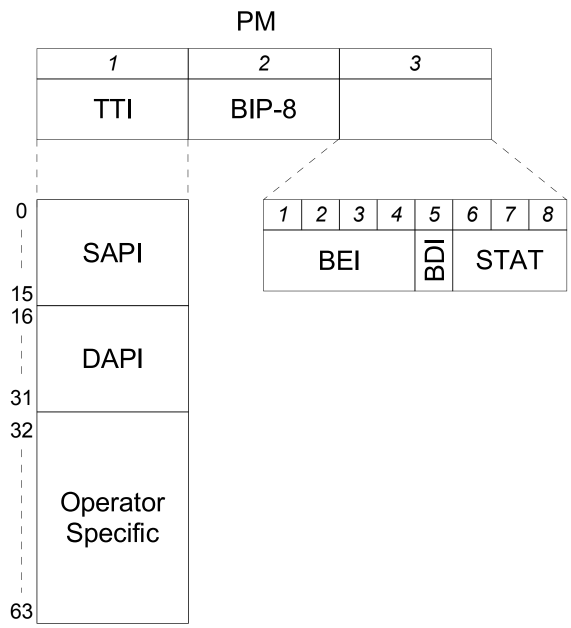

Figure 12: The path monitoring fields. (source: Draft G.709)

Note the addition of the STAT (status) field when compared to the SM fields. This field indicates the presence (or absence) of a maintenance signal. A value of 001 indicates a normal signal. A value of 101 indicates a locked (LCK) signal, while 110 indicates an open connection indication (OCI), and 111 indicates an alarm indication signal (AIS). All other values are reserved.

Jumping back to row 2, the first field is reserved. The next field is the tandem connection monitoring activation/deactivation (TCM ACT) field. This definition of this field is for further study. For now, TCM fields are activated by agreement between the network providers.

The next sequence of fields after the TCM ACT field relates to tandem connection monitoring (TCM), which is similar to line overhead in SONET/SDH. The idea is to provide fields which allow each segment to monitor the defects in that segment and report its status back to its peer. See Figure 13 which illustrates tandem connection operation. Note that cascaded and overlapping segments supported.

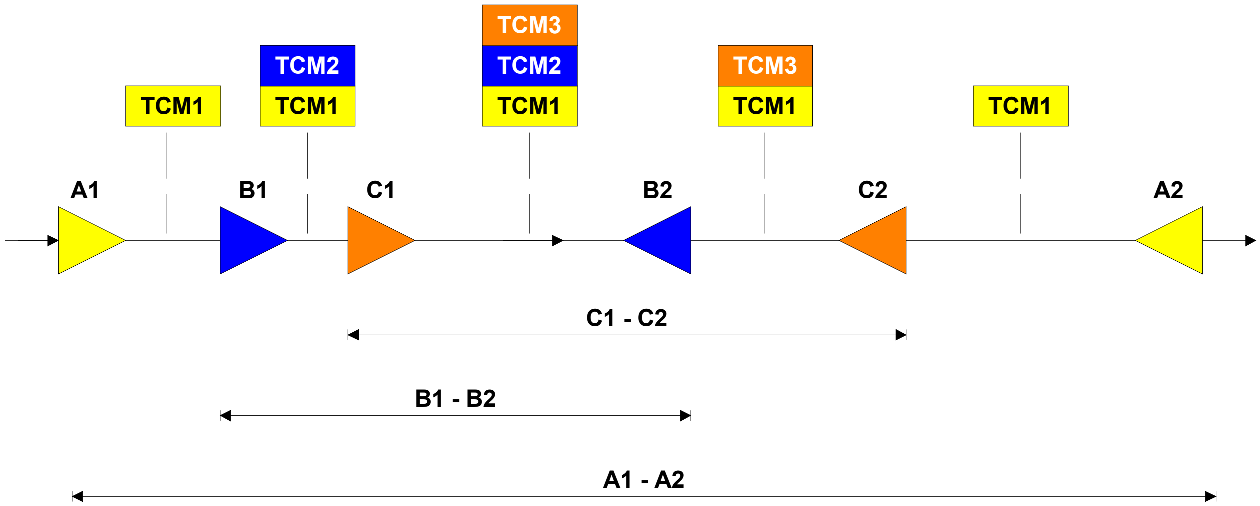

Figure 13: Tandem connections and the TCMi field. (Source: Draft G.709)

The contents of each TCMi field is similar to the PM field described earlier. See Figure 14 for details.

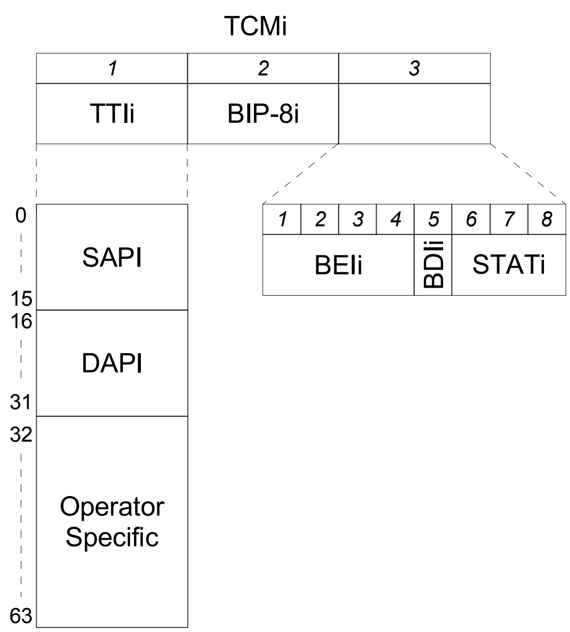

Figure 14: The path monitoring and tandem connection monitoring fields. (Source: Draft G.709)

Here, the values of the STATi field are as follows. A value of 001 indicates the TCM field is in use, without incoming alignment error (IAE). A value of 010 indicates in use, with IAE. A value of 101 indicates a locked (LCK) signal, while 110 indicates an open connection indication (OCI), and 111 indicates an alarm indication signal (AIS). All other values are reserved.

In column 14 of row 2, there is a one-octet field for a fault type and fault location reporting communication channel (FTFL). This octet is used to create a 256 octet multiframe message which is divided into a 128 octet forward field and a 128 octet backward field. Three fields are defined in each subfield. The first octet is the fault indication field which indicates either (1) no fault, (2) signal fail, or (3) signal degrade.

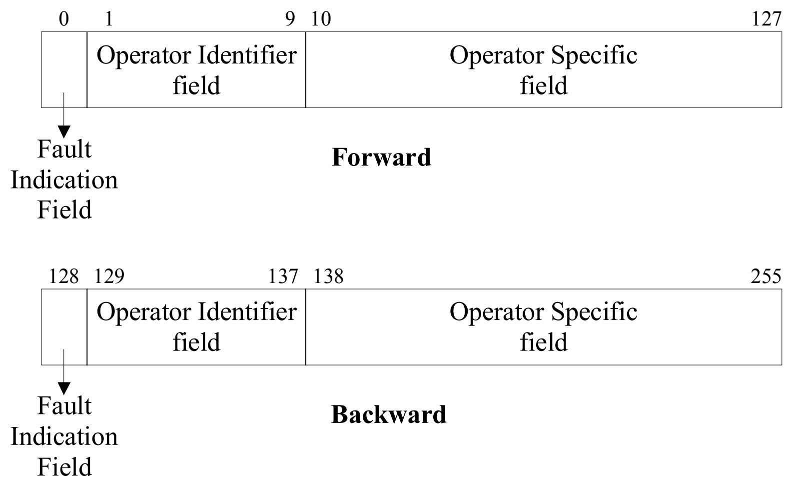

Figure 15: The fault type and fault location communication channel fields. (Source: T1X1.5 document 2000-246, p58)

The experimental field (EXP) is in row 2, columns 13 and 14. This field is not subject to standardization but is available to network operators within their own (sub)networks to support applications which require data overhead.

GCC1 and GCC2 are two general communication channels for use by any equipment which has access to the data overhead. The contents are not subject to standardization.

The automatic protection switching and protection control channel field (APS/PCC) is for further study. This is a serious omission because these fields are needed to communicate failures and coordinate the switchover to the protection fiber. The final field is reserved (RES) for future standardization.

The last area of overhead is the payload overhead. See Figure 16.

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
1	Frame alignment OH							Transport OH							RES	JC	D	D
2	Data OH														RES	JC	D	D
3															RES	JC	D	D
4															PSI	NJO	PJO	D

Figure 16: Payload overhead. Note that two columns of payload are included (the D indicates "data").

The primary field in the payload area is the payload structure identifier (PSI), located in column 15, row 4. This is a 256 octet multiframe field, although only the first octet is defined as the payload type (PT). The rest of the octets are reserved. The values for the PT field are defined in the recommendation. The rest of the octets in column 15 are reserved.

The octets in column 16 depend upon whether the payload is plesiochronous or synchronous. If the payload is completely synchronous, all of the octets are reserved. For plesiochronous traffic, there can be differences between the input and output clocks, necessitating some mechanism for handing overruns and under runs. This is achieved by use of the negative justification opportunity (NJO) and positive justification opportunity (PJO) octets. The justification control (JC) octets control the justification and are duplicates. Majority voting applies when interpreting the JC octets.

As mentioned earlier, there are three data rates defined in G.709, corresponding to payload rates of OC-48/STM-16 (OTU1), OC-192/STM-64 (OTU2) and OC-768/STM-256 (OTU3). As described above, an OTU1 signal has 238 payload octets out of 255 total octets so we have to multiply the OC-48/STM-16 line rate by 255/238, giving a line rate of approximately 2.666 057 143 Gbps. The rates for OTU2 and OTU3, however, are bit different.

It is important to be able to multiplex lower rate signal into a higher rate signal. Since the FEC in the higher rate signal covers the entire frame, it is not necessary to transport the FEC of the lower rate signal. It is, however, necessary to transport the overhead of the lower rate signal. Since there is one octet of overhead per RS frame in an OTU1 signal, it is necessary to allow an extra overhead octet per RS frame, or 16 extra octets per row, in the OTU2 frame so that this lower rate overhead can be transported. This means that only 237 octets per RS frame can be used for payload in the OTU2 frame, requiring that the line rate be 255/237 times the OC-192/STM-64 line rate, or about 10.709 225 316 Gbps.

Likewise, for the OTU3 signal, two octets per RS frame, or 32 octets per row, must be allowed for multiplexing. This requires that the line rate be 255/236 times the OC-768/STM-256 line rate, or about 43.018 413 559 Gbps. The frame formats for OTU2 and OTU3 are shown in Figure 17 and Figure 18 and the line rates are summarized in Table 3.

Level of Signal	**Approximate Line Rate (Gbps)**	**Time to transmit one row (μs)**	**Time to transmit one frame (μs)**
1	2.666 057 143	12.242	48.971
2	10.709 225 316	3.047	12.191
3	43.018 413 559	0.759	3.034

Table 3: Characteristics of the three levels of signals defined in G.709: line rate, time to transmit one row (buffering delay), and time to transmit one frame (four rows). Times are rounded to three decimals.

Multiplexing is not covered in this version of G.709 but is intended for a future revision of the recommendation. The subject is covered in an ITU contribution available at http://ties.itu.int/u/tsg15/sg15/wp3/q11/0009/cd/mmb03b.pdf (unfortunately not free public access).

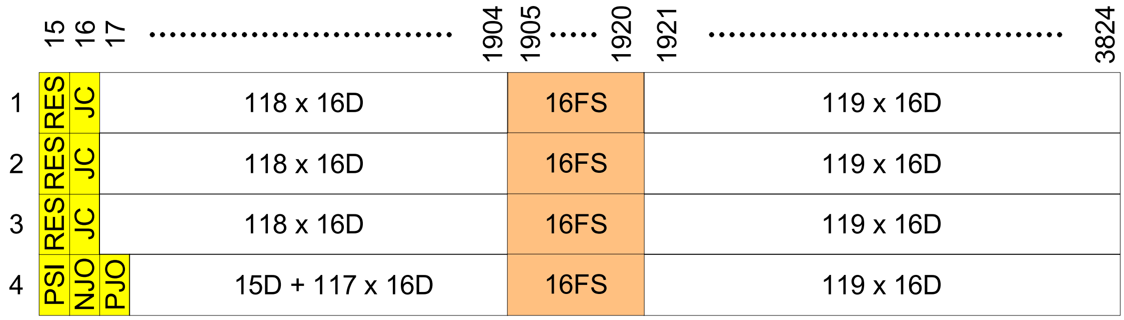

Figure 17: Payload mapping into an OTU2 frame. The PJO is only used when justification is required. Otherwise, it carries payload data. (source: T1X1.5 document 2000-246, p68)

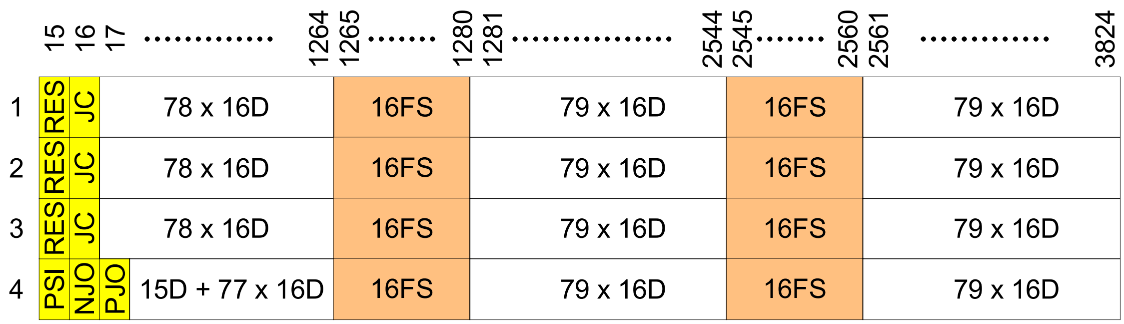

Figure 18: Payload mapping into an OTU3 frame. The PJO is only used when justification is required. Otherwise, it carries payload data. (source: T1X1.5 document 2000-246, p69)

Another area which is not fully specified is a method of framing data traffic within the G.709 frames. A technique, known as Generic Framing Procedure (GFP), is under development. Although it is not complete, enough work has been done to provide us with an outline of the technology.

GFP is an unframed technique, similar to ATM. That is, there is no framing character, as exists in packet over SONET (POS), because any framing character must be shielded inside the frame which leads to non-deterministic lengthening of the data frame. ATM has a five-octet header, with a header error control (HEC) applied over the header. The HEC octets are a part of the header. The find the header, the receiver simply selects five octets and computes the HEC. If it matches, the receiver jumps 48 octets to the beginning of the presumed next header. If that presumed header checks for HEC, and some number off additional presumed headers check for HEC, the receiver assumes it has attained frame alignment.

GFP is similar except that the frame is variable, with a maximum length of 64K octets. The header is four octets and consist of two octets which give the length of the frame, and two octets of CRC-16 applied to the two octet length field. See Figure 19. Since the length field is two octets, the maximum length frame is 65,535 octets.

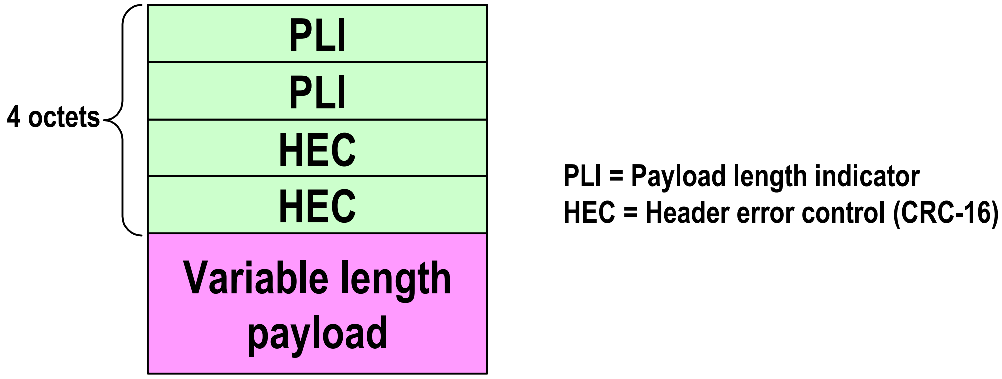

Figure 19: The Generic Framing Procedure frame, showing the fixed header.

It's important to note that the length in the PLI fields includes the length of the header. Lengths of 00 to 03 are reserved, with 00 being used for an idle frame (with HEC of 00, also). Lengths of 01, 02, and 03 are used for OAM&P and have an implied payload length of 8 octets (total frame length of 12 octets).

A fixed scrambler of 0xB6AB31E0 is exclusive-OR'ed with the header to provide DC balance and sufficient one's density during idle transmission periods. The payload is scrambled with the same scrambler as is used in ATM ($1 + x^{43}$) to reduce the probability of the payload replicating the header.

There can be little doubt that G.709 is intended as a replacement for SONET/SDH. Essentially all the functions included in SONET/SDH are included in G.709, plus a few additional (especially in the area of tandem connections). A logical question at this point is whether G.709 can actually replace SONET/SDH. First, let's examine overhead.

SONET/SDH has 4 columns of overhead for every 90 columns, an overhead rate of 4.44% of the line rate. G.709 has 17 octets of overhead for every 255 octets, giving an overhead rate of 6.67% of the line rate at the OTU1 rate. However, G.709 operates at a higher line rate than SONET/SDH which means that the G.709 payload actually runs at the full SONET/SDH rate. Thus, G.709 will provide a higher true payload rate than SONET/SDH.

To put data (IP traffic, for example) into the G.709 payload, some type of framing must be used but this type of data framing must also be used to carry traffic in SONET/SDH. And since the user does not see the line rate and doesn't care about it, if it can be accomplished by the network provider with little effort G.709 can provide higher user data rates.

G.709 also provides a significantly stronger FEC than SONET/SDH. The additional gain provided by this FEC may provide an important advantage to network engineers.

So will G.709 replace SONET/SDH in the network, especially for data traffic? SONET/SDH provides management overhead and is well understood. Network providers have used it for years and chips are readily available. Like most new technologies, G.709 will probably not replace SONET/SDH but will co-exist with it. The strength and size of G.709's eventual position in the network will depend upon whether network providers gain a significant advantage by using it. Only time will tell.

References

G.707 Study Group 15, "*Report of Working Party 3/15, Multiplexing and Switching (Geneva Meeting, 3-14 April 2000), Part II B.5 – Draft Determined Revised Recommendation G.707/Y.1322*". Document COM15-R71, available at http://www.itu.int/itudoc/itu-t/com15/reports/

G.709 "*Draft ITU-T G.709*". Document number T1X1.5/2000-246, available through http://www.t1.org/t1x1/_x1-grid.htm (Check T1X1.5 year 2000 docs and new docs. A more current version may be available from the ITU by the time you read this.)