INTERNET-DRAFT G. Manhoudt
Intended Status: Proposed Standard AimValley
Expires: August 7, 2014

S. Roullot
Alcatel-Lucent

K.Y. Wong
Alcatel-Lucent

H.L. Kuo
Chunghwa Telecom

February 3, 2014

Transparent SDH/SONET over Packet
draft-manhoudt-pwe3-tsop-04

Abstract

This document describes the Transparent SDH/SONET over Packet (TSoP) mechanism to encapsulate Synchronous Digital Hierarchy (SDH) or Synchronous Optical NETwork (SONET) bit-streams in a packet format, suitable for Pseudowire (PW) transport over a packet switched network (PSN). The key property of the TSoP method is that it transports the SDH/SONET client signal in its entirety through the PW, i.e., no use is made of any specific characteristic of the SONET/SDH signal format, other than its bit rate. The TSoP transparency includes transporting the timing properties of the SDH/SONET client signal. This ensures a maximum of transparency and a minimum of complexity, both in implementation and during operation.

Table of Contents

1. Introduction

This document describes the Transparent SDH/SONET over Packet (TSoP) method for encapsulating SDH or SONET signals with bit rates of 51.84 Mbit/s or N * 155.52 Mbit/s (where N = 1, 4, 16 or 64) for Pseudowire (PW) transport over a packet switched network (PSN), using circuit emulation techniques.

The selected approach for this encapsulation scheme avoids using any particular signal characteristics of the SDH/SONET signal, other than its bit rate. This approach closely follows the SAToP method described in [RFC4553] for PW transport of E1, DS1, E3 or DS3 over a PSN.

An alternative to the TSoP method for STM-N transport over PW is known as CEP (Circuit Emulation over Packet) and is described in [RFC4842]. The key difference between the CEP approach and the TSoP approach is that within CEP an incoming STM-N is terminated and demultiplexed to its constituent VCs (Virtual Containers). Subsequently, each VC is individually circuit emulated and encapsulated into a PW and transported over the PSN to potentially different destinations, where they are reassembled into (newly constructed) STM-N signals again. The TSoP approach, on the other hand, is to encapsulate the entire STM-N in a single circuit emulating Pseudowire and transport it to a single destination over the PSN. The essential difference between both methods is that CEP offers more routing flexibility and better bandwidth efficiency than TSoP at the cost of the loss of transparency (overhead, timing, scrambling) at the STM-N layer and at the cost of added complexity associated with the inclusion of what in essence is an SDH/SONET VC cross-connect function in the PEs.

Within the context of this document, there is no difference between SONET [GR-253] signals, often denoted as OC-M, and SDH [G.707] signals, usually denoted as STM-N. For ease of reading, this document will only refer to STM-N, but any statement about an STM-N signal should be understood to apply equally to the equivalent OC-M signal, unless it is specifically mentioned otherwise. The equivalency can be described by the following relations between N and M: If N = 0 then M = 1 and if N >= 1 then M = 3 * N.

The TSoP solution presented in this document conforms to the PWE3 architecture described in [RFC3985] and satisfies the relevant general requirements put forward in [RFC3916].

As with all PWs, TSoP PWs may be manually configured or set up using a suitably expanded version of the PWE3 control protocol [RFC4447].

2. Terminology and Conventions

2.1. Conventions Used in This Document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

2.2. Acronyms and Terms

The following acronyms used in this document are defined in [RFC3985] and [RFC4197]:

AC   Attachment Circuit
CE   Customer Edge
PE   Provider Edge
PREP   Pre-Processing
PSN   Packet Switched Network
PW   Pseudowire
SDH   Synchronous Digital Hierarchy
SONET   Synchronous Optical Network

In addition, the following specific terms are used in this document:

G-AIS   Generic Alarm Indication Signal - A specific bit pattern that replaces the normal STM-N signal in the case of certain failure scenarios. The G-AIS pattern [G.709] is constructed by continuously repeating the 2047 bit pseudo random bit sequence based on the generating polynomial 1 + x^9 + x^11 according to [O.150].

IWF   Interworking Function - A functional block that segments and encapsulates a constant bit-rate signal into PW packets and that in the reverse direction decapsulates PW packets and reconstitutes the constant bit-rate signal.

LOF   Loss Of Frame - A condition of an STM-N signal in which the frame pattern cannot be detected. Criteria for raising and clearing a LOF condition can be found in [G.783].

LOPS   Loss of Packet State - A defect that indicates that the PE at the receiving end of a TSoP carrying PW experiences an interruption in the stream of received TSoP packets. See [RFC5604]

LOS   Loss Of Signal - A condition of the STM-N attachment circuit in which the incoming signal has an insufficient energy

level for reliable reception. Criteria for raising and clearing a LOS condition can be found in [G.783].

NIM   Non-Intrusive Monitor - A circuit that monitors a signal in a certain direction of transmission, without changing the binary content of it. A NIM can be used for Fault Management and Performance Monitoring purposes

PDV   Packet Delay Variation - A (statistical) measure that describes the distribution of the variation in transit times of packets in a certain flow between two reference points in the network. See [G.8260]

SF   Signal Fail - A control signal, that exists internally in a system, to convey the failed state of an incoming signal, from a server layer process to the adjacent client layer process. See [G.783]

3. Emulated STM-N Services

A TSoP emulated STM-N service over a Pseudowire makes use of a bi-directional point-to-point connection over the PSN between two TSoP-IWF blocks, located in the PE nodes that terminate the PW that interconnects them, as shown in figure 1. The TSoP-IWF blocks each consist of two half-functions, a PSN-bound IWF and a CE-bound IWF, one for each direction of transmission. As the name implies, the PSN-bound part of the TSoP-IWF performs the conversion of an STM-N bitstream to a packet flow, suitable for transport over the PSN and the CE-bound part of the TSoP-IWF performs the inverse operation.

```
       |<------------------ Emulated Service ----------------->|
       |                                                       |
       |      |<-------------- Pseudowire ------------->|      |
       |  AC  |                                         |  AC  |
       |<---->|           |<----- PSN ----->|           |<---->|
       |      |           |                 |           |      |
       |      .    PE1    .                 .    PE2    .      |
       .      +-----------+                 +-----------+      .
   +---+      |           |                 |           |      +---+
   |   |----->| PSN-bound |====> . . . ====>| CE-bound  |----->|   |
   | C |      |    IWF    |                 |   IWF     |      | C |
   | E |      | _   _   _ |                 | _   _   _ |      | E |
   |   |      |           |                 |           |      |   |
   | 1 |      |           |                 |           |      | 2 |
   |   |<-----| CE-bound  |<==== . . . <====| PSN-bound |<-----|   |
   +---+  |   |   IWF     |                 |    IWF    |   |  +---+
          |   +-----------+                 +-----------+   |
          |     TSoP-IWF                       TSoP-IWF     |
        native                                           native
        STM-N                                             STM-N
```

Figure 1. Overview of STM-N emulated service architecture

The following list provides the STM-N services, as specified in [G.707] and [GR-253], that can be supported by a TSoP PW:

1. STM-0 or OC-1 (51.84 Mbit/s)
2. STM-1 or OC-3 (155.52 Mbit/s)
3. STM-4 or OC-12 (622.08 Mbit/s)
4. STM-16 or OC-48 (2488.32 Mbit/s)
5. STM-64 or OC-192 (9953.28 Mbit/s)

The TSoP protocol used for emulation of STM-N services does not depend on the method in which the STM-N is delivered to the PE. For example, an STM-1 attachment circuit is treated in the same way regardless of whether it is a copper [G.703] or a fiber optic [G.957] link.

Also, in case the STM-N is carried in an OTN signal [G.709], the functionality in the TSoP-IWF operates in the same way, but a PWE3 Preprocessing (PREP) functional block will be present between the AC and the PE to perform the OTN (de)multiplexing functions.

The TSoP-IWF function in figure 1 is further broken down in functional blocks in figure 2. These individual functional blocks are described in the next two sections.

```
     AC                         TSoP-IWF                           PSN
   ------>|<------------------------------------------------->|<--------

          +---------------------------------------------------+
          |                 +-------+                         |
          |  +-------+  SF  | PSN-  |  CE-side_               |
          |  |       |----->| bound |  defect     +--------+  |
   STM-N  |  | STM-N |   +->| NIM   |------------>|        |  |
   =========>|  Rx   |   |  +-------+             |        |  |
          |  |       |===0=======================>| PSN-   |  |   Packet
          |  +-------+      +--------+            | bound  |===========>
          |                 | Subst. |            | IWF    |  |
          |                 | Signal |===========>|        |  |
          |                 | Gen.   |        +-->|        |  |
          |                 +--------+        |   +--------+  |
          |                                   |               |
          | PSN-bound direction               |               |
   - - - -|- - - - - - - - - - - - - - - - - -|- - - - - - - -|- - - - -
          | CE-bound direction                |               |
          |                                   |               |
          |                 +--------+        | PSN-side_     |
          |                 | G-AIS  |        | defect        |
          |                 | Gen.   |====+   |               |
          |                 +--------+    |   |               |
          |                               |   |   +--------+  |
          |                 +--------+    |   +---|        |  |
          |  +-------+      |        |<===+       |        |  |
          |  |       |      | STM-N/ |  No_Packet | CE-    |  |
   <=========| STM-N |<=====| G-AIS  |<-----------| bound  |<===========
   STM-N  |  |  Tx   |      | Switch |            | IWF    |  |   Packet
          |  |       |  +-->|        |<========0==|        |  |
          |  +-------+  |   +--------+         |  |        |  |
          |             |      +----------+    |  +--------+  |
          |             |      | optional |    |              |
          |             +------| CE-bound |<---+              |
          |                    |   NIM    |                   |
          |                    +----------+                   |
          +---------------------------------------------------+
```

Figure 2. TSoP functional block diagram

3.1. PSN-bound Direction

In the PSN-bound direction the STM-N signal is received from the CE via an AC by the STM-N Rx function. This function recovers the optical or electrical signal and converts it to a suitable internal format. In addition, it detects the LOS condition and it asserts the SF signal whenever this is the case. The STM-N Rx block is equivalent to the OSn_TT_Sk & OSn/RSn_A_Sk (in the case of an optical STM-N) or the ESn_TT_Sk & ESn/RSn_A_Sk (in the case of an electrical STM-N interface) function pairs defined in [G.783].

The PSN-bound IWF segments the STM-N ingress bitstream, which it receives from the STM-N Rx function, in blocks of equal length. Each block of bits is supplied with the appropriate TSoP Encapsulation Headers and then delivered to the PSN Multiplexing layer to add the required headers for transport over the PSN.

The PSN-bound NIM function controls the state of the CE-side_defect signal. It will assert this signal in case the SF signal is asserted or in case another defect is detected in the incoming STM-N signal. The inclusion of other defects than LOS in the CE-side_defect signal is OPTIONAL.

When the CE-side_defect signal is asserted, the PSN-bound IWF will set the corresponding flag (L-bit) in the overhead of the affected packets. Packets in which the L-bit is set MUST have a substitution payload (created by the Substitution Signal Generator function) of the same length as the regular TSoP payload. This substitution payload is RECOMMENDED to be the G-AIS pattern or a fixed "all ones" pattern.

Lastly, when the PSN-side_defect state is asserted, the PSN-bound IWF will set the corresponding flag (R-bit) in the overhead of all packets that are transmitted while this signal is in the asserted state.

3.2. CE-bound Direction

In the CE-bound direction, the CE-bound IWF receives the PW packets from the PSN and strips off the PSN, PW, and TSoP encapsulation headers and writes the payload data in a buffer. The output data stream towards the CE is created by playing out this buffer with a suitable clock signal. The thus reconstructed STM-N signal is forwarded to the STM-N/G-AIS Switch function.

The No_Packet signal is asserted by the CE-bound IWF in case the internal packet buffer empties due to lack of input packets from the PSN or in case a packet is missing or invalid.

The PSN-side_defect signal is asserted by the CE-bound IWF in case the LOPS condition is detected by the CE-bound IWF (see section 6.2.2). The state of this signal controls the value of the R-bit in the overhead of the packets returned towards the far-end TSoP-IWF.

The G-AIS Generator generates a G-AIS signal at the nominal frequency of the recovered STM-N signal, +/- 20 ppm.

The STM-N/G-AIS Switch normally takes its input from the CE-bound IWF and forwards the recovered STM-N signal towards the STM-N Tx function, but during the time that the No_Packet signal is asserted, it will select the G-AIS Generator as its active input and forward a G-AIS signal towards the STM-N Tx function.

The CE-bound NIM function is an OPTIONAL function that can be used to detect additional defects in the recovered CE-bound STM-N signal. The presence of such defects (e.g. STM-N LOF) MAY be used as an additional reason for the STM-N/G-AIS Switch function to select the G-AIS signal as its active input.

Lastly, the STM-N Tx function converts the internal signal that is output by the STM-N/G-AIS Switch block into a regular STM-N signal towards the CE via the AC. The STM-N Tx block is equivalent to the OSn_TT_So & OSn/RSn_A_So (in the case of an optical STM-N) or the ESn_TT_So & ESn/RSn_A_So (in the case of an electrical STM-N interface) function pairs defined in [G.783].

4. TSoP Encapsulation Layer

4.1. TSoP Packet Format

The general format of TSoP packets during transport over the PSN is shown in Figure 3.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                              ...                              |
|                        PSN/PW Headers                         |
|                              ...                              |
+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+
|                              ...                              |
|                  TSoP Encapsulation Headers                   |
|                              ...                              |
+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+
|                                                               |
|                              ...                              |
|                                                               |
|                      TSoP Payload Field                       |
|                                                               |
|                              ...                              |
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 3. Generic TSoP Packet Format

4.2. PSN/PW Headers

A TSoP PW can be transported over different types of PSNs based on different switching technology. Below the transmission over MPLS is described, but other methods are not precluded. The selected method will determine the format of the PSN/PW Headers part and influence the order of the fields in the TSoP Encapsulation Headers part.

4.2.1 Transport over an MPLS(-TP) PSN

In case a TSoP PW is forwarded over an MPLS(-TP) PSN, a standard "bottom of stack" PW label as shown in figure 4 is prepended before the TSoP Encapsulation Headers. Subsequently, one or more MPLS(-TP) labels need to be pushed according to the standard MPLS transport methods outlined in [RFC3031] and [RFC3032].

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                 Label                 | TC  |S|     TTL       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 4. PW Label (S = 1)

4.3. TSoP Encapsulation Headers

4.3.1. Location and Order of TSoP Encapsulation Headers

The TSoP Encapsulation Headers MUST contain the TSoP Control Word (figure 6) and MUST contain a Minimum length RTP Header [RFC3550] (figure 7). The TSoP Encapsulation Headers must immediately follow the PSN/PW header, as shown in figure 5.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                              ...                              |
|                        PSN/PW Headers                         |
|                              ...                              |
+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+
|                       TSoP Control Word                       |
+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+
|                                                               |
+--                                                           --+
|           Minimum length RTP Header (see [RFC3550])           |
+--                                                           --+
|                                                               |
+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+
|                              ...                              |
|                      STM-N data (Payload)                     |
|                              ...                              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 5. General TSoP Packet Format

4.3.2. Usage and Structure of the TSoP Control Word

The purpose of the TSoP control word is to allow:

1. Detection of packet loss or misordering
2. Differentiation between PSN and attachment circuit problems as a cause for outage of the emulated service
3. Signaling of faults detected at the PW egress to the PW ingress

The structure of the TSoP Control Word is in accordance with the general PW Control Word format specified in [RFC4385]. The TSoP CW format is shown in Figure 6 below. This TSoP Control Word MUST be present in each TSoP PW packet.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|0 0 0 0|L|R|RSV|FRG|    LEN    |       Sequence number         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 6. Structure of the TSoP Control Word

The use of Bits 0 to 3 is described in [RFC4385]. These bits MUST be set to zero unless they are being used to indicate the start of an Associated Channel Header (ACH). An ACH is needed if the state of the TSoP PW is being monitored using Virtual Circuit Connectivity Verification [RFC5085] or in case OAM functionality according to [RFC6371] is added.

L (bit 4) - If this bit is set, it indicates that the STM-N data ingressing in the PSN-bound IWF is currently experiencing a fault condition. Once set, if the fault is rectified, the L-bit MUST be cleared. For each frame that is transmitted with L-bit = 1, the PSN-bound IWF MUST insert such an amount of substitution data in the TSoP payload field that the TSoP frame length, as it is during normal operation, is maintained. The CE-bound IWF MUST play out an amount of G-AIS data corresponding to the original TSoP Payload Field for each received packet with the L-bit set.

Note: This document does not prescribe exactly which STM-N fault conditions are to be treated as invalidating the payload carried in the TSoP packets. An example of such a fault condition would be LOS.

R (bit 5) - If this bit is set by the PSN-bound IWF, it indicates that its local CE-bound IWF is in the LOPS state, i.e., it has lost a preconfigured number of consecutive packets. The R-bit MUST be cleared by the PSN-bound IWF once its local CE-bound IWF has exited the LOPS state, i.e., has received a preconfigured number of consecutive packets. See also section 6.2.2.

RSV (bits 6 to 7) - This field MUST be set to 0 by the PSN-bound IWF and MUST be ignored by the CE-bound IWF. RSV is reserved.

FRG (bits 8 to 9) - This field MUST be set to 0 by the PSN-bound IWF and MUST be ignored by the CE-bound IWF. FRG is fragmentation; see [RFC4623].

LEN (bits 10 to 15) - This field MAY be used to carry the length of the TSoP packet (defined as the length of the TSoP Encapsulation Header + TSoP Payload Field) if it is less than 64 octets, and MUST be set to zero otherwise. When the LEN field is set to 0, the preconfigured size of the TSoP packet payload MUST be assumed to be as described in Section 5, and if the actual packet size is inconsistent with this length, the packet MUST be considered malformed.

Sequence number (bits 16 to 31) - This field is used to enable the common PW sequencing function as well as detection of lost packets. It MUST be generated in accordance with the rules defined in Section 5.1 of [RFC3550] for the RTP sequence number:

- o Its space is a 16-bit unsigned circular space
- o Its initial value SHOULD be random (unpredictable).

It MUST be incremented with each TSoP data packet sent in the specific PW.

4.3.3. Usage of the RTP Header

A minimum length RTP Header as specified in [RFC3550] MUST be included in the TSoP Encapsulation Header. The reason for mandating the insertion of an RTP Header by the PSN-bound IWF is that it is expected that in most cases the CE-bound IWF will need to use the contained timestamps to be able to recover a clock signal of sufficient quality. By avoiding to make the presence of RTP Headers subject to configuration, the design of the CE-bound IWF can be simplified and another potential source of errors during commissioning is eliminated.

The RTP Header fields in the list below (see also figure 7) MUST have the following specific values:

V (version) = 2
P (padding) = 0
X (header extension) = 0
CC (CSRC count) = 0
M (marker) = 0

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| V |P|X|   CC  |M|     PT      |       Sequence number         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                           Timestamp                           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                              SSRC                             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 7. Structure of the RTP Header field

The PT (payload type) field is used as follows:

1. One PT value SHOULD be allocated from the range of dynamic values (see [RTP-TYPE]) for each direction of the PW. The same PT value MAY be reused for both directions of the PW and also reused between different PWs.
2. The PSN-bound IWF MUST set the PT field in the RTP header to the allocated value.
3. The CE-bound IWF MAY use the received value to detect malformed packets.

The sequence number MUST be the same as the sequence number in the TSoP control word.

The RTP timestamps are used for carrying timing information over the network. Their values MUST be generated in accordance with the rules established in [RFC3550].

A TSoP implementation MUST support RTP timestamping at the PW ingress with a nominal clock frequency of 25 MHz. This is also the default value. Other clock frequencies MAY be supported to generate the RTP Timestamps. Selection of the applicable clock frequency is done during commissioning of the PW that carries the emulated STM-N service.

The SSRC (synchronization source) value in the RTP header MAY be used for detection of misconnections, i.e., incorrect interconnection of attachment circuits. In case this option is not used, this field should contain an all zero pattern.

The usage of the options associated with the RTP Header (the timestamping clock frequency, selected PT and SSRC values) MUST be aligned between the two TSoP IWFs during Pseudowire commissioning.

5. TSoP Payload Field

In order to facilitate handling of packet loss in the PSN, all packets belonging to a given TSoP PW are REQUIRED to carry a fixed number of octets in its TSoP Payload Field.

The TSoP Payload Field length MUST be defined during PW commissioning, MUST be the same for both directions of the PW, and MUST remain unchanged for the lifetime of the PW.

All TSoP implementations MUST be capable of supporting the following TSoP Payload Field length:

- STM-N (for N = 0, 1, 4, 16 and 64) - 810 octets

Notes:

1. Whatever the selected payload size, TSoP does not assume alignment to any underlying structure imposed by STM-N framing (octet, frame, or multiframe alignment). The STM-N signal remains scrambled through the TSoP encapsulation and decapsulation processes.

2. With a payload size of 810 octets, the STM-N emulated service over the PSN will have a nominal packet rate of 8000 packets/s when N = 0 and a nominal packet rate of 24000 * N packets/s for N >= 1.

TSoP uses the following ordering for packetization of the STM-N data:

- The order of the payload octets corresponds to their order on the attachment circuit.

- Consecutive bits coming from the attachment circuit fill each payload octet starting from most significant bit to least significant.

6. TSoP Operation

6.1. Common Considerations

Edge-to-edge emulation of an STM-N service using TSoP is only possible when the two PW attachment circuits are of the same type, i.e., both are STM-N with equal N.

6.2. IWF Operation

6.2.1. PSN-Bound Direction

Once the PW is commissioned, the PSN-bound TSoP IWF operates as follows:

The ingressing STM-N bit-stream is segmented, such that each segment contains the configured number of payload octets per packet. This forms the TSoP Payload Field. The STM-N bit-stream MUST NOT be descrambled before segmentation and packetization for PW transport.

Subsequently, the TSoP Encapsulation Headers are prepended according to the rules in section 4.3.

Lastly, the PSN/PW Headers are added to the packetized service data, and, depending on the applicable layer 1 technology, additional overhead is added. The resulting packets are transmitted over the PSN.

6.2.2. CE-Bound Direction

Once the PW is commissioned, the CE-bound TSoP IWF operates as follows:

Each time a valid TSoP packet is received from the PSN, its sequence number is checked to determine its relative position in the stream of received packets. Packets that are received out-of-order MAY be reordered. Next, the data in the fixed length TSoP payload field of each packet is written into a (jitter) buffer in the order indicated by its sequence number. In case data is missing due to a lost packet or a packet that could not be re-ordered, an equivalent amount of dummy data (G-AIS pattern) is substituted.

Subsequently, the STM-N stream towards the CE is reconstructed by playing out the buffer content with a clock that is reconstructed to have the same average frequency as the STM-N clock at the PW ingress. In addition, this clock signal must have such properties that the following requirements can be met:

- o A reconstructed SDH-type STM-N signal delivered to an Attachment Circuit MUST meet [G.825] and [G.823] jitter and wander requirements (for synchronization interfaces), or,

- o A reconstructed SONET-type OC-M signal delivered to an Attachment Circuit MUST meet [GR-253] jitter and wander requirements.

The size of the buffer in the CE-bound TSoP IWF SHOULD be configurable to allow accommodation to the PSN specific packet delay variation (see appendix B).

The CE-bound TSoP IWF MUST use the sequence number in either the TSoP Control Word or in the RTP header for detection of lost and misordered packets. The use of the sequence number in the TSoP Control Word for this purpose is RECOMMENDED.

The CE-bound TSoP IWF MAY reorder misordered packets. Misordered packets that can not be reordered MUST be discarded and treated the same way as lost packets.

The payload of received TSoP packets marked with the L-bit set MUST be replaced by the equivalent number of bits from the G-AIS pattern. Likewise, the payload of each lost or malformed (see section 6.3) TSoP packet MUST be replaced with the equivalent number of bits from the G-AIS pattern.

Before a TSoP PW has been commissioned and after a TSoP PW has been decommissioned, the IWF MUST play out the G-AIS pattern to its STM-N attachment circuit.

Once a TSoP PW has been commissioned, the CE-bound IWF begins to receive TSoP packets and to store their payload in the buffer, but continues to play out the G-AIS pattern to its STM-N attachment circuit. This intermediate state persists until a preconfigured degree of filling (for example half of the CE-bound IWF buffer) has been reached by writing consecutive TSoP packets or until a preconfigured intermediate state timer (started when the TSoP commissioning is complete) expires. See appendix B for considerations regarding the configuration of the initial degree of filling of this buffer.

Each time an STM-N signal is replaced by a G-AIS signal at the same nominal bitrate, this signal may start at an arbitrary point in its repeating 2047-bit sequence. Once the starting point is selected, the G-AIS signal is sent uninterrupted until the condition that invoked it has been removed. The frequency of the clock that is used to generate this G-AIS signal MUST have an accuracy that is better than +/- 20 ppm relative to the nominal STM-N frequency. Appendix D describes the effect of G-AIS insertion on downstream SDH equipment.

Once the preconfigured amount of the STM-N data has been received, the CE-bound TSoP IWF enters its normal operational state where it continues to receive TSoP packets and to store their payload in the buffer while playing out the contents of the jitter buffer in accordance with the required clock. In this state, the CE-bound IWF

performs clock recovery, MAY monitor PW defects, and MAY collect PW performance monitoring data.

The CE-bound IWF enters the LOPS defect state in case it detects the loss of a preconfigured number of consecutive packets or if the intermediate state timer expires before the required amount of STM-N data has been received. While in this state, the local PSN-bound TSoP IWF SHOULD mark every packet it transmits with the R-bit set. The CE-bound IWF leaves the LOPS defect state and transits to the normal state once a preconfigured number of consecutive valid TSoP packets have been received (successfully reordered packets contribute to the count of consecutive packets).

The RTP timestamps inserted in each TSoP packet at the PW ingress allow operation in differential mode, provided that both PW ingress and PW egress IWFs have a local clock that is traceable to a common timing source.

The use of adaptive mode clocking mode, i.e., recovering the STM-N clock in the CE-bound IWF by essentially averaging the arrival times of the TSoP packets from the PSN without using RTP information, is not recommended for TSoP-based circuit emulation. Appendix C provides some considerations regarding the implementation and configuration of the CE-bound IWF.

6.3. TSoP Defects

In addition to the LOPS state defined above, the CE-bound TSoP IWF MAY detect the following defects:

- o Stray packets
- o Malformed packets
- o Excessive packet loss rate
- o Buffer overrun
- o Buffer underrun
- o Remote packet loss

Corresponding to each defect is a defect state of the IWF, a detection criterion that triggers transition from the normal operation state to the appropriate defect state, and an alarm that MAY be reported to the management system and thereafter cleared. Alarms are only reported when the defect state persists for a preconfigured amount of time (typically 2.5 seconds) and MUST be cleared after the corresponding defect is undetected for a second preconfigured amount of time (typically 10 seconds). The trigger and release times for the various alarms may be independent.

Stray packets MAY be detected by the PSN and PW demultiplexing

layers. The SSRC field in the RTP header MAY be used for this purpose as well. Stray packets MUST be discarded by the CE-bound IWF, and their detection MUST NOT affect mechanisms for detection of packet loss.

Malformed packets are detected by mismatch between the expected packet size and the actual packet size inferred from the PSN and PW demultiplexing layers (taking the value of the L-bit into account). Differences between the received PT value and the PT value allocated for this direction of the PW MAY also be used for this purpose. Malformed, in-order packets MUST be discarded by the CE-bound IWF and replacement data generated as with lost packets.

Excessive packet loss rate is detected by computing the average packet loss rate over a configurable amount of time and comparing it with preconfigured raise and clear thresholds.

Buffer overrun is detected in normal operational state when the (jitter) buffer of the CE-bound IWF cannot accommodate newly arrived TSoP packets.

Buffer underrun can detected in normal operational state when the (jitter) buffer of the CE-bound IWF has insufficient data to maintain playing out the STM-N signal towards the CE at the recovered clock rate. In this situation G-AIS MUST be substituted until the buffer fill has reached its preconfigured degree of filling again.

Remote packet loss is indicated by reception of packets with their R-bit set.

6.4. TSoP Performance Monitoring

Performance monitoring (PM) parameters are routinely collected for STM-N services and provide an important maintenance mechanism in SDH networks. However, STM-N level PM data provides the information over the performance of the end-to-end STM-N connection, which may extend well beyond the part in which it is carried over a TSoP Pseudowire.

It may be important to be able to measure the performance of a TSoP Pseudowire section, which forms a part of the STM-N end-to-end connection, in isolation. For that reason a set of packet level counters are specified that can be used to assess the performance of the TSoP Pseudowire section. Collection of the TSoP PW performance monitoring data is OPTIONAL and, if implemented, is only performed after the CE-bound IWF has exited its intermediate state.

The following counters are defined:

ENCAP_TXTOTAL_PKTS (counter size: 32 bits) - The total number of TSoP packets that is transmitted towards the PSN by the PSN-bound IWF function. This includes packets with the L-bit set.

DECAP_RXTOTAL_PKTS (counter size: 32 bits) - The total number of TSoP packets that is received from the PSN by the CE-bound IWF function. This includes malformed packets, out-of-order packets and packets with the L-bit set.

DECAP_REORDERED_PKTS (counter size: 32 bits) - The number of out-of-order TSoP packets that is received from the PSN by the CE-bound IWF, based on the received sequence numbers, for which the ordering could be corrected by the CE-bound IWF.

DECAP_MISSING_PKTS (counter size: 32 bits) - The number of TSoP packets that did not arrive at the CE-bound IWF from the PSN, based on the received sequence numbers.

DECAP_MALFORMED_PKTS (counter size: 32 bits) - The number of TSoP packets that is received from the PSN by the CE-bound IWF function which contains one of the following RTP related errors: TSoP Payload Field length mismatch, PT-value mismatch (if checked) and/or SSRC mismatch (if checked).

DECAP_OUTOFORDER_PKTS (counter size: 32 bits) - The number of out-of-order TSoP packets that is received from the PSN by the CE-bound IWF, based on the received sequence numbers, for which the ordering could not be corrected by the CE-bound IWF.

DECAP_OVERRUN_BITS (counter size: 64 bits) - The number of bits of TSoP Payload that is received from the PSN but dropped by the CE-bound IWF due to the fact that the (jitter) buffer has insufficient capacity available to store the complete TSoP Payload Field content.

DECAP_UNDERRUN_BITS (counter size: 64 bits) - The number of bits that is not played out towards the CE by the CE-bound IWF because the (jitter) buffer is empty at the moment they need to be played out.

DECAP_PLAYEDOUT_PKTS (counter size: 32 bits) - The number of packets that has been successfully played out towards the CE by the CE-bound IWF containing valid STM-N payload, including the packets that have been received with the L-bit set, containing substituted data. Packets which are lost in transmission over the PSN or packets which are (partially) discarded by the CE-bound IWF due to some error condition are not counted.

Note that packets with the L-bit set are considered normal data from the perspective of TSoP Pseudowire Performance Monitoring, since in such cases the location of the fault is in the STM-N path, before it ingresses the PSN-bound IWF, so outside the scope of the TSoP PW.

7. Quality of Service (QoS) Issues

TSoP SHOULD employ existing QoS capabilities of the underlying PSN.

If the PSN providing connectivity between PE devices is Diffserv-enabled and provides a PDB [RFC3086] that guarantees low jitter and low loss, the TSoP PW SHOULD use this PDB in compliance with the admission and allocation rules the PSN has put in place for that PDB (e.g., marking packets as directed by the PSN).

If the PSN is Intserv-enabled, then GS (Guaranteed Service) [RFC2212] with the appropriate bandwidth reservation SHOULD be used in order to provide a bandwidth guarantee equal or greater than that of the encapsulated STM-N traffic.

8. Congestion Control

As explained in [RFC3985], the PSN carrying the PW may be subject to congestion. TSoP PWs represent inelastic constant bit-rate flows and cannot respond to congestion in a TCP-friendly manner prescribed by [RFC2914], although the percentage of total bandwidth they consume remains constant.

Unless appropriate precautions are taken, undiminished demand of bandwidth by TSoP PWs can contribute to network congestion that may impact network control protocols.

Whenever possible, TSoP PWs SHOULD be carried across traffic-engineered PSNs that provide either bandwidth reservation and admission control or forwarding prioritization and boundary traffic conditioning mechanisms. IntServ-enabled domains supporting Guaranteed Service (GS) [RFC2212] and DiffServ-enabled domains [RFC2475] supporting Expedited Forwarding (EF) [RFC3246] provide examples of such PSNs. Such mechanisms will negate, to some degree, the effect of the TSoP PWs on the neighboring streams.

If TSoP PWs run over a PSN providing best-effort service, they SHOULD monitor packet loss in order to detect "severe congestion". If such a condition is detected, a TSoP PW SHOULD shut down bi-directionally for some period of time as described in Section 6.5 of [RFC3985].

Note that:

1. The TSoP IWF can inherently provide packet loss measurement since the expected rate of arrival of TSoP packets is fixed and known

2. The results of the TSoP packet loss measurement may not be a reliable indication of presence or absence of severe congestion if the PSN provides enhanced delivery. For example:

   a) If TSoP traffic takes precedence over non-TSoP traffic, severe congestion can develop without significant TSoP packet loss.

   b) If non-TSoP traffic takes precedence over TSoP traffic, TSoP may experience substantial packet loss due to a short-term burst of high-priority traffic.

3. The availability objectives for the digital paths that are supported by an STM-N signal (see [G.827]) MUST be taken into account when deciding on temporary shutdown of TSoP PWs.

This specification does not define the exact criteria for detecting "severe congestion" using the TSoP packet loss rate or the specific methods for bi-directional shutdown the TSoP PWs (when such severe congestion has been detected) and their subsequent re-start after a suitable delay. This is left for further study. However, the following considerations may be used as guidelines for implementing the TSoP severe congestion shutdown mechanism:

1. If the TSoP PW has been set up using either PWE3 control protocol [RFC4447], the regular PW teardown procedures of these protocols SHOULD be used.

2. If one of the TSoP PW end points stops transmission of packets for a sufficiently long period, its peer (observing 100% packet loss) will necessarily detect "severe congestion" and also stop transmission, thus achieving bi-directional PW shutdown.

9. Security Considerations

TSoP does not enhance or detract from the security performance of the underlying PSN; rather, it relies upon the PSN mechanisms for encryption, integrity, and authentication whenever required.

TSoP PWs share susceptibility to a number of Pseudowire layer attacks and will use whatever mechanisms for confidentiality, integrity, and authentication are developed for general PWs. These methods are beyond the scope of this document.

Although TSoP PWs MUST employ an RTP header to achieve an explicit transfer of timing information, SRTP (see [RFC3711]) mechanisms are NOT RECOMMENDED as a substitute for PW layer security.

Misconnection detection capabilities of TSoP increase its resilience to misconfiguration.

Random initialization of sequence numbers, in both the control word and the optional RTP header, makes known plaintext attacks on encrypted TSoP PWs more difficult. Encryption of PWs is beyond the scope of this document.

10. Applicability Statements

TSoP is an encapsulation layer intended for carrying SDH STM-N circuits over the PSN in a structure-agnostic and fully transparent fashion.

TSoP fully complies with the principle of minimal intervention, minimizing overhead and computational power required for encapsulation.

TSoP provides sequencing and synchronization functions needed for emulation of STM-N bit-streams, including detection of lost or misordered packets and perform the appropriate compensation. Furthermore, explicit timing information is provided by the presence of an RTP timestamp in each TSoP packet.

STM-N bit-streams carried over TSoP PWs may experience delays exceeding those typical of native SDH networks. These delays include the TSoP packetization delay, edge-to-edge delay of the underlying PSN, and the delay added by the jitter buffer. It is recommended to estimate both delay and delay variation prior to setup of a TSoP PW. See appendix B for more information on jitter buffer configuration.

TSoP carries STM-N streams over PSN in their entirety, including any control plane data contained within the data. Consequently, the emulated STM-N services are sensitive to the PSN packet loss. Appropriate generation of replacement data can be used to prevent LOF defects and declaration of severely errored seconds (SES) due to occasional packet loss. Other effects of packet loss on this interface (e.g., errored blocks) cannot be prevented. See appendix D for more information.

TSoP provides for effective fault isolation by forwarding the local attachment circuit failure indications to the remote attachment circuit.

TSoP provides for a carrier independent ability to detect misconnections and malformed packets via the PT and SSRC fields in the RTP Header. This feature increases resilience of the emulated service to misconfiguration.

Being a constant bit rate (CBR) service, TSoP cannot provide TCP friendly behavior under network congestion.

Faithfulness of a TSoP PW may be increased by exploiting QoS features of the underlying PSN.

TSoP does not provide any mechanisms for protection against PSN outages, and hence its resilience to such outages is limited. However, lost packet replacement and packet reordering mechanisms increase resilience of the emulated service to fast PSN rerouting events.

A key requirement for TSoP to achieve transparent transport of the timing information of an STM-N signal, is that the recovered STM-N signal meets all relevant SDH and SONET jitter/wander requirements (see section 6.2.2). It will depend on the synchronization situation of the PSN whether or not a given CE-bound TSoP implementation can meet this requirement. In appendix C a number of network synchronization situations are listed, in which it is possible to meet this requirement with a reasonable CE-bound IWF design. In other network synchronization scenarios, the application of TSoP is not generically recommended.

11. IANA Considerations

IANA is requested to assign a new MPLS Pseudowire (PW) type for the following TSoP encapsulated services:

PW type	Description	Reference
0x0020	STM-0 or OC-1	RFC XXXX
0x0021	STM-1 or OC-3	RFC XXXX
0x0022	STM-4 or OC-12	RFC XXXX
0x0023	STM-16 or OC-48	RFC XXXX
0x0024	STM-64 or OC-192	RFC XXXX

The above value is suggested as the next available value and has been reserved for this purpose by IANA.

RFC Editor: Please replace RFC XXXX above with the RFC number of this document and remove this note.

12. Acknowledgements

The authors of this document are much indebted to the authors of [RFC4553]. This latter RFC has been used as a template and example for the current document. Moreover, many paragraphs and sentences have been copied from this RFC without alteration or with only slight modification into the current document.

Furthermore, we thank Zhu Hao, Jeff Towne, Willem van den Bosch, Peter Roberts and Matthew Bocci for their valuable feedback.

13. References

13.1. Normative References

[G.707] ITU-T Recommendation G.707/Y.1322 (01/2007) - Network node interface for the synchronous digital hierarchy (SDH)

[G.783] ITU-T Recommendation G.783 (03/2006) - Characteristics of synchronous digital hierarchy (SDH) equipment functional blocks

[O.150] ITU-T Recommendation O.150 (05/1996) - General requirements for instrumentation for performance measurement on digital transmission equipment

[G.823] ITU-T Recommendation G.823 (03/2000) - The control of jitter and wander within digital networks which are based on the 2048 kbit/s hierarchy

[G.825] ITU-T Recommendation G.825 (03/2000) - The control of jitter and wander within digital networks which are based on the synchronous digital hierarchy (SDH)

[GR-253] Telcordia GR-253-CORE - Synchronous Optical Network (SONET) Transport Systems: Common Generic Criteria, Issue 5, October 2009

13.2. Informative References

[G.703] ITU-T Recommendation G.703 (11/2001) - Physical/electrical characteristics of hierarchical digital interfaces

[G.709] ITU-T Recommendation G.709/Y.1331 (12/2009) - Interfaces for the Optical Transport Network (OTN)

[G.781] ITU-T Recommendation G.781 (09/2008) - Synchronization layer functions

[G.811] ITU-T Recommendation G.811 (09/1997) - Timing characteristics of primary reference clocks

[G.827] ITU-T Recommendation G.827 (09/2003) - Availability performance parameters and objectives for end-to-end international constant bit-rate digital paths

[G.828] ITU-T Recommendation G.828 (03/2000) - Error performance parameters and objectives for international, constant bit rate synchronous digital paths

[G.957] ITU-T Recommendation G.957 (06/1999) - Optical interfaces for equipments and systems relating to the synchronous digital hierarchy

[G.8260] ITU-T Recommendation G.8260 (02/2012) - Definitions and terminology for synchronization in packet networks

[G.8262] ITU-T Recommendation G.8262/Y.1362 (07/2010) - Timing characteristics of a synchronous Ethernet equipment slave clock

[G.8263] ITU-T Recommendation G.8263/Y.1363 (02/2012) - Timing characteristics of packet-based equipment clocks

[RTP-TYPE] RTP PARAMETERS, <http://www.iana.org/assignments/rtp-parameters>

Appendix A. Parameter Configuration for TSoP PW Set-up

The following parameters of the TSoP IWF MUST be agreed upon between the peer IWFs during the PW setup. Such an agreement can be reached via manual configuration or via one of the PW set-up protocols:

1. Type of attachment circuit, i.e., the value of N of the STM-N signal, which corresponds to a bit rate as mentioned in section 3.

2. Payload size, i.e., the (constant) number of octets that is transmitted in the TSoP Payload Field of each TSoP packet. The default value is 810 octets.

3. Timestamping clock frequency: 25 MHz (default) or an alternative value.

4. The configurability of the following parameters (see section 6.2.2) governing the behavior of the CE-bound IWF buffer is optional:

   a) The maximum amount of payload data that may be stored in the CE-bound IWF payload buffer

   b) The desired degree of filling of the CE-bound IWF buffer in steady state (see appendix B)

   c) The "intermediate state" timer, i.e., the maximum amount of time that the CE-bound IWF waits before it starts to play out data towards the CE

5. The content of the following RTP header fields must be provided by the user:

   a) The 7-bit RTP Payload Type (PT) value; any value can be assigned to be used with TSoP PWs. Default is an all zero pattern.

   b) The 32-bit Synchronization Source (SSRC) value. Default is an all zero pattern.

6. The number of TSoP packets that must be missed consecutively before the CE-bound IWF enters the LOPS defect state (default: 10) and the number of TSoP packets that must be received consecutively to clear the LOPS defect state (default: 2). See section 4.3.2 and [RFC5604]

7. To support the optional excessive packet loss event by the CE-bound IWF, the following parameters must be configured:

   a) The length of the observation period for detecting excessive packet loss. Default value is 10 s.

   b) The minimum number of lost packets that is to be detected in the observation interval before an excessive packet loss alarm is raised. Default value is 30% of the expected packets.

   c) The maximum number of lost packets that is to be detected in the observation interval to clear an excessive packet loss alarm. Default value is 1% of the expected packets.

Appendix B. Buffer Configuration in the CE-bound IWF

The buffer in the CE-bound IWF (often called the "jitter buffer") is used to compensate the differences in transit time that each bit of the STM-N signal experiences between the moment it ingresses the PSN-bound IWF and the moment it ingresses the CE-bound IWF. There are two mechanisms that cause the transit times of individual bits to be different:

1. The packetization delay (Tpkt(n)) in the PSN-bound IWF: After arrival in the PSN-bound IWF, STM-N bit #n has to wait until sufficient bits have been received to fill the complete Payload Field of a TSoP packet. Clearly if two STM-N bits end up in the same TSoP packet, the bit that arrives earlier has to wait longer than the bit that arrives later. The (variable part of the) packetization delay, Tpkt(n), varies between zero and the time between the transmission of two subsequent TSoP packets.

2. The packet delay variation (Tpdv(n)), i.e. the difference in transit time that the TSoP packet containing bit #n experiences relative to some reference (minimum) transit time, due to the presence of non-empty shapers and queues (or any other cause for variable delay) on its path through the PSN.

```
                  ----- direction of transmission ----->

   +------+    +-------------+               +-------------+    +------+
   | Clk1 |    |     Clk2    |               |    Clk3     |    | Clk4 |
   |------|    |-------------|               |-------------|    |------|
   |      |    |             |               |             |    |      |
   |      |    | +---------+ |  +---------+  | +---------+ |    |      |
   | SDH1 |----|-| Tpkt(n) |=|==| Tpdv(n) |==|=| Tjbf(n) |-|----| SDH2 |
   |      |    | +---------+ |  +---------+  | +---------+ |    |      |
   |      |    |packetization|    packet     |   jitter    |    |      |
   |      |    |    delay    |     delay     |    buffer   |    |      |
   +------+    +-------------+   variation   +-------------+    +------+
     CE1             PE1                           PE2             CE2

          |----|             |===============|             |----|
            AC1                     PSN                      AC2
          (STM-N)                                          (STM-N)

                   +-----------------------------------+
                   |  Clk1 is the system clock of CE1  |
                   |  Clk2 is the system clock of PE1  |
                   |  Clk3 is the system clock of PE2  |
                   |  Clk4 is the system clock of CE2  |
                   +-----------------------------------+
```

Figure 8. Delay components in a uni-directional TSoP Pseudowire (from PE1 to PE2).

Figure 8 schematically shows three contributors to the over-all transit delay of STM-N bits between ingressing PE1 and egressing PE2: The paketization delay in PE1 and the packet delay variation over the PSN, as mentioned above, which cause delay differences between the bits and the jitter buffer delay in PE2, Tjbf(n), which is intended to equalize these differences.

When the CE-bound IWF in PE2 starts up, it writes the bits in the TSoP Payload Field of incoming packets into the jitter buffer until a preconfigured degree of filling is reached. From this moment onward, bits are played out from this buffer under control of an STM-N clock (derived from Clk3) that is locally synthesized in PE2. A necessary condition to avoid overflow or underflow of the jitter buffer is that this clock must have the same average frequency as the clock in PE1 (derived from Clk2) that governs the encapsulation process.

The dwelling time of an individual bit in the jitter buffer, Tjbf(n), is determined by the actual delay time that this bit experienced in transiting through PE1 and the PSN. Bits traveling fast reside long in the jitter buffer, while slow bits reside only a short while in

the buffer, such that the total time is equal for all bits.

For a given bit #n, this behavior can be expressed by the relation Tpkt(n) + Tpdv(n) + Tjbf(n) = constant (within the allowed jitter and wander limits of an STM-N signal). Since especially Tpdv(n) can vary significantly, the initial degree of filling of the jitter buffer, Tjbf(0), should be configured large enough to allow lowering the buffer fill when at a later time bits arrive that happen to be slower. This compensating effect of the jitter buffer fill on the overall transit time requires the frequency of Clk3, and so the frequency of the outgoing STM-N signal, to change slow compared to the PDV variations.

Note that the Tjbf(n) term adds to the total delay that an STM-N bit experiences crossing the TSoP PW section. For this reason the initial degree of filling of the jitter buffer should not be configured larger than necessary, i.e. it is a compromise between the delay permitted by the application and the probability of buffer underrun due to large PDV excursions. Depending on the requirements for a particular STM-N client signal, a small probability of buffer underrun may be acceptable in order to meet the delay specifications. See appendix D for a description of the effects of jitter buffer underrun on CE2.

Apart from the initial fill level of the jitter buffer, its total size can also be configurable (see section 6.2.2). The fill level increases when a number of faster packets arrives and the buffer needs sufficient "headroom" to avoid overflow under such conditions. However, it is not necessary to reserve more "headroom" than is needed to accommodate the fastest packets, corresponding to the minimum delay of the PW.

Appendix C. Synchronization Considerations for the CE-bound IWF

In section 6.2.2 the requirements for reconstruction of the STM-N client signal from the incoming TSoP packets in the CE-bound IWF function are given, without reference to the quality of the CE-bound IWF system clock or the method of STM-N clock recovery. This is done on purpose, to avoid as much as possible restrictions on the implementation, as these factors depend on the (network) synchronization situation. This appendix provides information on this dependency.

Figure 9 shows the reference network that is used to analyze the dependency of the CE-bound IWF requirements on the synchronization situation in the network. Since network synchronization operates uni-directional, only the corresponding direction of transmission is depicted. The return path will never be used at the same time as a

synchronization reference signal: If Ref4 is derived from AC2, the S1-byte of the returning STM-N will contain the message "Don't Use for Synchronization".

```
               ----- direction of transmission ----->

                     Ref2                   Ref3
                       |                     |
  Ref1                 |                     |                 Ref4
   |                   V                     V                   |
   |              +--------+             +--------+              |
   V              |  Clk2  |             |  Clk3  |              V
+------+          |--------|             |--------|          +------+
| Clk1 |          | STM-N /|             | TSoP  /|          | Clk4 |
|------|          |      / |             | -PW  / |          |------|
|      |          |     /  |             |     /  |          |      |
| SDH1 |--------->|    /   |============>|    /   |--------->| SDH2 |
|      |  STM-N   |   /    |    TSoP     |   /    |   STM-N  |      |
|      |          |  /     |   Packets   |  /     |          |      |
+------+          | / TSoP |             | /      |          +------+
  CE1             |/   -PW |             |/ STM-N |             CE2
       |          +--------+             +--------+          |
       |             PE1                     PE2             |
       |          |        |             |        |          |
       |-- AC1 -->|        |==== PSN ===>|        |-- AC2 -->|
       |          |                               |          |
       |          |--------- Pseudowire --------->|          |
Up-    |                                                     |  Down-
stream |------------- STM-N (multiplex) section ------------>| stream


        +---------------------------------------------------+
        |  Clk1 is the system clock of CE1, locked to Ref1  |
        |  Clk2 is the system clock of PE1, locked to Ref2  |
        |  Clk3 is the system clock of PE2, locked to Ref3  |
        |  Clk4 is the system clock of CE2, locked to Ref4  |
        +---------------------------------------------------+
```

Figure 9. Reference network for analysis of TSoP synchronization requirements

The intention of the requirements in section 6.2.2 is to ascertain that the reconstruction of the STM-N client signal in PE2 is sufficiently faithful that not only its binary content is recovered but also its timing properties are maintained. This latter aspect implies that the recovered STM-N signal can still be used as a

synchronization reference signal in the downstream STM-N network as is required per [G.825] and [GR-253] and that normal processing rules (see [G.781]) can be applied to the S1-byte.

The requirements regarding maintaining STM-N timing properties of the reconstructed STM-N are based on two rules:

1. The reconstructed STM-N signal at the TSoP PW egress must have the same average frequency as the STM-N signal at the TSoP PW ingress.

   Since all TSoP packets carry the same number of payload bits and a sequence numbering mechanism is applied to the TSoP packets, the TSoP decapsulation function can regenerate an STM-N signal that has, averaged over time, the same number of bits, as long as the jitter buffer does not overrun or underrun.

2. The jitter and wander properties of the recovered STM-N signal must meet the applicable SDH/SONET standards.

   The sub-sections below provide the synchronization situations in the network in which this requirement can be met with a reasonable design of the CE-bound IWF. The implication is that for other network synchronization situations such is, in general, not possible.

The considerations in this appendix are applicable during "normal" operation. As soon as the input STM-N signal to PE1 is lost or the TSoP PW itself is failing, PE2 forwards a G-AIS signal towards CE2 at a frequency that may deviate 20 ppm from the nominal value. A sustained G-AIS pattern will cause a Loss of Frame condition in CE2, which will consequently stop reading the contained S1 information and look for an alternative synchronization reference signal or revert to hold-over mode.

C.1. Layer 1 Synchronized PEs

The PEs that terminate the TSoP PW operate synchronously in case Ref2 and Ref3 are traceable over the physical layer to the same clock.

In such cases the timing characteristics of the STM-N signal can be transferred over the PW by applying differential mode, i.e. use the RTP timestamps in the CE-bound IWF to determine the rate at which the STM-N bits need to be played out. This will ensure that the average frequency is maintained over the PW section.

To satisfy the STM-N wander requirements, Clk3 must filter the phase noise on Ref3 down to the levels specified in [G.825] or [GR-253]. In case the phase noise on Ref3 stays below the network limit

specified for Synchronous Ethernet [G.8262], it is sufficient that Clk3 meets the jitter transfer specifications of [G.8262] to achieve this goal.

C.2. Synchronous CEs

The CEs that terminate the STM-N multiplex section operate synchronously in case Ref1 and Ref4 are traceable over the physical layer to the same clock.

In such cases, Ref2 can be derived from Clk1, via AC1 and Ref3 can be derived from Clk4 via AC2. The timing characteristics of the STM-N signal can be transferred over the PW by applying differential mode, i.e. use the RTP timestamps in the CE-bound IWF to determine the rate at which the STM-N bits need to be played out. This will ensure that the average frequency is maintained over the PW section.

To satisfy the STM-N wander requirements, Clk3 must filter the phase noise on Ref3 down to the levels specified in [G.825] or [GR-253]. In case the phase noise on Ref3 stays below the network limit specified for STM-N [G.825], it is sufficient if Clk3 meets the jitter transfer specifications of [G.813] to achieve this goal.

C.3. Pleisiochronous CEs

In case Clk1 and Clk4 are both traceable to different [G.811] type clocks, the STM-N link operates pleisiochronously, i.e. the long term frequency difference between both clocks is less than 2E-11.

In such cases, Ref2 can be derived from Clk1, via AC1 and Ref3 can be derived from Clk4 via AC2. Again, differential mode clock recovery is assumed in the CE-bound IWF. The very small frequency difference will cause a re-center action of the jitter buffer in PE2, each time it overflows or underflows. The time interval between such events depends on the depth of the buffer and the actual frequency difference between Ref1 and Ref4.

As an example, if the frequency difference is 2E-11 (this represents the worst case) and if an overflow or underflow event shifts the buffer fill by 125 microseconds, a re-center event happens once every 72 days. In the downstream SDH node, a re-center event causes 1 errored second (ES) in all VC paths that are carried over this STM-N signal (see appendix D). A level of one ES per 72 days is negligible compared to the ES limits formulated in [G.828] for VC-paths.

Appendix D. Effect of G-AIS Insertion on a Downstream SDH Node

There are a number of network events that force the CE-bound IWF to replace the content of the Payload Field of a TSoP packet by the same number of bits from the local G-AIS pattern generator (see figure 2). This is the case when:

a) A TSoP packet is lost, or,

b) A TSoP packet arrives out-of-order and its position can't be restored, or,

c) A TSoP packet arrives with its L-bit set.

In case the STM-N payload of K consecutive TSoP packets is replaced by G-AIS, the effect is that parity violations are detected in the downstream SDH Node (CE) for a period T = K / (N * 8000) seconds. As long as K is small, this will typically lead to counting 1 Errored Second (ES) event (or occasionally 2 in case the event straddles two adjacent 1 second monitoring periods) for the STM-N signal itself and all contained VC-path layers. As long the number of ES that is induced in a given VC-path by the effects above, remains small compared to the applicable limit defined in [G.828], this effect of G-AIS substitution can be tolerated.

A second consequence of G-AIS substitution is that the G-AIS bits may overwrite the STM-N alignment word (A1 and A2 bytes) in the recovered STM-N signal. This A1-A2 alignment word repeats every 0.125 ms. The length of the period T determines the number of framing patterns that is affected by the inserted G-AIS bits.

Each time an STM-N framing word is changed by G-AIS bits overwriting it, this will cause a framing anomaly in the downstream SDH Node. Such anomalies have no effect on the STM-N framer as long as a next A1-A2 word is correct and is still in its expected location. Only in case the G-AIS streak lasts for a 3 ms or longer period, these persistent anomalies will cause a LOF defect (see [G.783]) to be raised and consequently the declaration of a Severely Errored Second (SES). Note that the limits for SES events are much stricter than for ES event (see [G.828]). On the other hand, the probability of losing 3 ms worth of TSoP packets (e.g. 72 consecutive TSoP packets for STM-1/OC-3) is much lower than losing one or a few packets in a row.

In case the jitter buffer is re-centered, for instance due to a buffer overrun or underrun, the position of the A1-A2 framing pattern will in general shift once by an amount of time that is not an integral multiple of 0.125 ms. Also such an event will not cause a

LOF defect in the downstream STM-N node, because an STM-N framer that conforms to [G.783] must be able to detect the out-of-frame condition within 0.625 ms and find the new frame position within 0.250 ms, so the entire operation is completed well within 3 ms (the minimum time to declare a LOF defect).

Note that in case of buffer underrun, G-AIS is transmitted while the CE-bound IWF waits for the buffer to reach its configured initial fill level. This waiting time has to be added to the STM-N out-of-frame detection time and frame recovery time mentioned above, to assess the overall impact on the downstream SDH node. This implies that if the initial fill level is configured somewhat larger than 2 ms (actually, 3 - 0.625 - 0.250 = 2.125 ms), an underrun event can trigger a LOF defect and consequently a SES event in the downstream SDH network. So while a larger jitter buffer diminishes the probability of an underrun event, the consequences of such an event are more severe once this ~2 ms threshold is crossed. This must be weighed carefully.

Authors' Addresses

Gert Manhoudt
AimValley B.V.
Utrechtseweg 38
1213 TV Hilversum
The Netherlands
E-mail: gmanhoudt@aimvalley.nl

Stephan Roullot
Alcatel-Lucent Centre de Villarceaux
Route de Villejust
91620 Nozay
France
E-mail: stephan.roullot@alcatel-lucent.com

Kin Yee Wong
Alcatel-Lucent
600 March Road
Kanata, Ontario, K2K 2E6
Canada
E-mail: kin-yee.wong@alcatel-lucent.com

Kuo, Huei-Long
Chunghwa Telecom Co. Ltd.
No.21-3, Sec. 1, Xinyi Rd.
Zhongzheng Dist., Taipei City 100
Taiwan
E-mail: hlkuo@cht.com.tw