Signaling transport over IP-based networks using IETF standards

Klaus D. Gradischnig[*]
NeuStar, Inc.

Michael Tüxen[†]
Siemens AG

Abstract

This paper analyzes some of the reliability features of the transport layers of the signaling system no. 7 (SS7) and compares them to the reliability and performance achievable with the protocol stack for signaling transport over IP currently being defined in the SIGTRAN working group of the IETF. It identifies parameters which have to be adjusted and restrictions to available addressing options which have to be made in order for the SIGTRAN protocol stack to achieve the reliability and performance of SS7.

1 Introduction

The convergence of voice (i. e. circuit switched) and data (i. e. packet switched) networks for offering seamless voice and multi-media services is rapidly progressing. In order for packet based voice services to find acceptance in the marketplace, however, it is paramount that the quality of service, as perceived by the end user of such services, is at least as good as that experienced in present day circuit switched networks. While there are a number of other important aspects of convergence, such as addressing [4], this paper concentrates on the reliability and performance of signaling in a converging environment.

For the PSTN/ISDN a number of strict reliability, quality, and performance requirements can be found in the E-series of Recommendations of the ITU-T, for example. One corner stone for delivering these requirements is the reliability of the signaling network. This has been discussed extensively, for instance at the first DRCN in [1] and [3] .

Even where the convergence of circuit and data networks is not yet an issue, it is becoming more and more important to be able to combine classical SS7-based networks with IP-based networks using the latter to transport/tunnel SS7 signaling messages. Deploying such a combined architecture enables operators to make use of the advantages of IP-based equipment in an SS7-based environment, avoiding some of the problems increasingly appearing in the rapidly growing SS7 networks, such as linkset capacity and load sharing [2].

The standardization of a protocol suite for transporting SS7 signaling over IP networks is currently being developed by the Signaling Transport (SIGTRAN) working group of the Internet Engineering Task Force (IETF). In order for convergence and/or SS7 over IP becoming a success it is therefore important that the protocol suite developed by SIGTRAN can offer the same or a better performance and reliability for signaling than the current SS7 network.

This paper discusses some of the explicit and implicit performance and reliability features of the SS7 transport layers (MTP levels 2 and 3) which together with proper signaling network planning and dimensioning deliver the performance and reliability required by the signaling applications.

We will then show that using the protocol parameters given in [11] for the deployment of SCTP (the basic transport protocol of the SIGTRAN protocol suite) in the public Internet will not result in the necessary protocol behavior for signaling transport. Therefore it is necessary to use special parameter settings to be able to fulfill the SS7 performance requirements. The relation between these parameters is analyzed and values are suggested which can be used for signaling transport.

Finally we will show that, when interworking between SS7 and the SIGTRAN stack is on MTP level 3 (as opposed to on MTP level 2 or on the SCCP level), a number of restrictions to the allowed addressing options in the SIGTRAN stack are necessary in order to enable MTP's network management to fully include the elements in the IP domain and thus to guarantee the reliability required from the signaling network.

[*] NeuStar, Inc., 45980 Center Oak Plaza, Sterling, VA 20166, USA Tel: +1 571 434 5652. Fax: +1 571 434 5601. e-mail: Klaus.Gradischnig@neustar.com.

[†] Siemens AG, ICN WN CS SE 5, D-81359 München, Germany. Tel: +49 89 722 47210. Fax: +49 89 722 48212. e-mail: Michael.Tuexen@icn.siemens.de.

2 SS7 Networks

The classical PSTN/ISDN separates the networks used for the transmission of the voice information of a call and the call control messages needed for controlling the calls. The latter network is called the signaling network and normally uses the protocols defined for the signaling system no. 7 (SS7). These SS7 networks are considered in this paper.

In the SS7 terminology the physical layer is called message transfer part level 1 (MTP level 1 or MTP1), the link layer is called message transfer part level 2 (MTP level 2 or MTP2) and the network layer is called message transfer part level 3 (MTP level 3 or MTP3). The layer on top of MTP3 is called user part. In this paper only one user part will be considered: the ISDN signaling user part (ISUP). This user part handles the signaling message needed for setting up basic telephone calls.

The nodes in SS7-networks are called signaling points (SP) and are identified by a unique 14-bit integer called the signaling point code (SPC). SPs are connected by signaling links. The bandwidth of a signaling link is normally 64 kbit/sec. If larger bandwidths are needed, and for redundancy reasons, up to 16 links between two signaling points can be used. This group of links between two SPs is called a linkset.

The load sharing in SS7-networks is based on the signaling link selection field (SLS), a 4-bit integer provided by the userpart. The point codes of the destination (DPC) and of the origin (OPC) are also provided by the userpart and the triple consisting of the SLS, OPC and DPC is called routing label and is part of each MTP3 protocol data unit, called message signal unit (MSU).

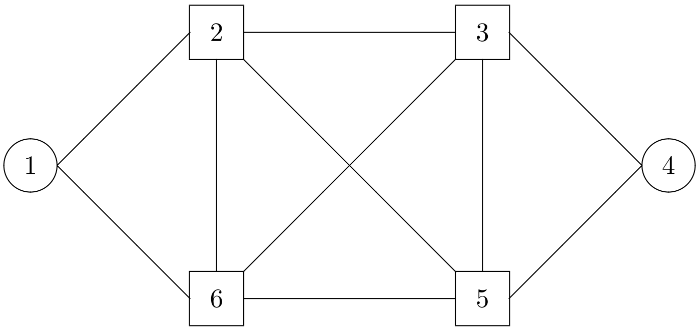

Figure 1: A simple SS7-network

The signaling points are divided into two classes:

1. A signaling end point (SEP) is the source or destination of message send by a user part. In figure 1 the SPs with point code 1 and 4 are SEPs. The protocol stack used in SEPs (with ISUP) is shown in figure 2.
2. A signaling transfer point (STP) only sends received messages of user part to other SPs. In figure 1 the SPs with point 2, 3, 5 and 6 are STPs. The protocol stack used in STPs is shown in figure 2.

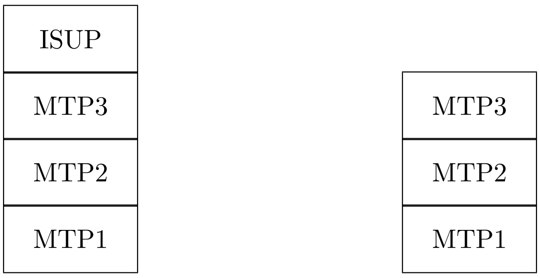

Figure 2: The protocol stack used in SEPs and STPs

3 Performance requirements

For the ISUP and signaling in general [7] and [8] list a number of reliability and performance requirements the signaling network must meet. These requirements directly or indirectly translate into the following MTP requirements (which then also fulfill the requirements of the other MTP users):

1. Not more than one in 10^{10} of all message signal units must contain an error that is undetected by the MTP.
2. Not more than one in 10^7 messages will be lost due to failure in the MTP.
3. The availability of any signaling relation (i. e. communication path between two communicating SEPs) has to be at least 0.99998 corresponding to a downtime of at most 10 minutes/year.
4. Not more than one in 10^{10} messages will be delivered out-of-sequence to the User Parts due to failure in the MTP. This value also includes duplication of messages.
5. In addition there are requirements on message transfer times in STPs, which under normal conditions are supposed to be less than 100 msec, and implicit requirements on limits for the outgoing queuing delays which must not become a dominating factor of the transfer times.

We will now address a number of these requirements in more detail and in particular describe the protocol features which make fulfilling these requirements a possibility.

Requirement 1 is a function of the quality of the underlying physical transport, the CRC function of the MTP2, and the likelihood of system internal errors of implementations.

In order to fulfill requirement 2 with unreliable hardware the MTP deploys redundant signaling links and the so called change-over procedure which allows the loss-free switching of traffic from a failed link to other links, provided the signaling link terminations on both nodes involved are still functioning and can communicate with each other via alternative links/paths.

In order to enable the design of signaling networks fulfilling requirement 3 the MTP provides several procedures supporting redundancy in the network. On 64kbit/s links link failures are discovered within 128 msec by the error rate monitor of MTP2. If an alternative link or path exists, MTP3 initiates the changeover procedure. Furthermore, [6] sets a limit of 800 msec (plus transmission times of changeover control messages) for this procedure. Thus, changeover from a failed link to an alternate one is effected within about one second or less. Should no alternate link or path exist, the MTP3 provides other procedures (transfer prohibited and forced rerouting) so that upstream nodes can re-route the traffic, if possible. Q.706 does not specify any performance requirements for these procedures. They are, however, of the same complexity as the changeover procedure and thus also expected to complete within 1 second or less. Within the same time frame a user of the MTP is informed that a destination is not reachable anymore and thus could initiate re-routing on the user level (call-routing could, e. g., route the call via a different transit switch). It should be noted that routing tables in an MTP network (normally) consist of prioritized lists of routes to the various destinations. When a route fails the next one in the list is used. The routing tables in the whole network are configured such that, together with the management procedures of the MTP, no matter which linksets fail, usage of alternate routes will not cause loops in the network. This allows a rapid, localized recovery from linkset failures without the need for a network wide co-ordination. This can, however, in principle also lead to situations where, although physically possible, no routes might be available between two nodes.[1]

To enable the fulfillment of requirement 4, MTP3 performs explicit or timer based sequence control procedures wherever possible when rerouting traffic via alternate links or routes or when reverting traffic back to the original routes.

While STP transfer times are an implementation and not a protocol issue the MTP provides several mechanisms to limit outgoing queues (requirement 5) and thus overall signaling transfer times. The error rate monitor of MTP2 not only rapidly discovers failed links but will also take a link out of service when the signal unit error rate approaches $4 \cdot 10^3$.[2] If outgoing congestion occurs on links, MTP management takes action and informs traffic sources to reduce traffic, either based on message priorities and/or on a sampling basis. If congestion is of a lasting nature (e. g. caused by too many link failures) a conditional rerouting procedure (transfer restricted procedure) can optionally be deployed.

4 The SIGTRAN Architecture

The SIGTRAN protocol suite was developed to allow an interworking between SS7 network elements and IP-based elements. Interworking units called signaling gateways (SGs) interconnect the SS7 network and the IP network. The IP-based application services are provided by making use of a cluster architecture for redundancy. The general architecture is described in [10].

[1] The same, though, can be said for policy based routing protocols.

[2] See [9] for a discussion of properties of an error rate monitor suitable for high speed signaling links.

The protocol stack running at a SEP is shown in figure 2. A corresponding architecture for an IP-based SEP, like a media gateway controller (MGC) or softswitch, is shown in figure 3.

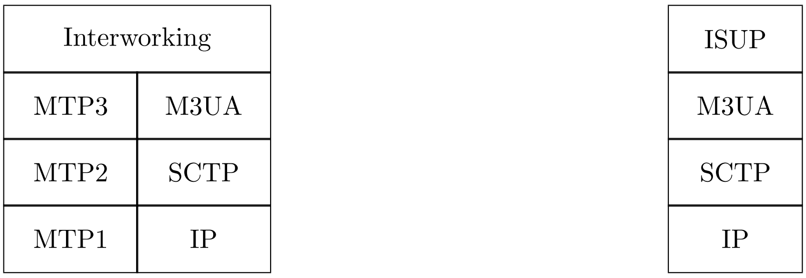

Figure 3: Protocol stack at the SG and the MGC using M3UA

The protocol stack handled at the SEP is divided into a lower, for instance MTP1, MTP2 and MTP3, and an upper, for instance ISUP, part. The lower part runs on the SG and the upper part runs on the IP-based cluster. The primitives defined between the lowest layer of the upper part, such as ISUP, and the top layer of the lower part, such as MTP3, are transported using an adaptation layer, in our case M3UA. M3UA is discussed in some more detail in section 6. In our example the cut in the SS7 protocol stack is made between MTP3 and ISUP but it would also be possible to do it at the MTP2/MTP3 boundary. The adaptation layer is different for different splits of the protocol stack and provides mechanisms for the transport of the primitives and the cluster management. A number of different adaptation layers are being developed by the SIGTRAN group:

1. IUA: The boundary is Q.921/Q.931.
2. M2UA: The boundary is MTP2/MTP3.
3. M2PA: The boundary is MTP2/MTP3, but for a symmetric scenario.
4. M3UA: The boundary is MTP3/user part
5. SUA: The boundary is SCCP/SCCP user.

These adaptation layers all use one common transport layer protocol running on top of IP. This transport protocol used for the signaling transport is the Stream Control Transport Protocol (SCTP) and is described in the following section.

Almost all the adaptation layers have a similar cluster management for handling host failures. The logical unit providing a service is called an application server (AS). For example, an AS running M3UA could handle all messages with a specific DPC. An AS consists of one ore more application server processes (ASPs) which one can think of as processes in the UNIX sense. The ASPs belonging to one AS can run on different physical nodes to provide redundancy. The redundancy concept is quite general supporting the (n+k) model. A special case is the (1+1) case with one ASP being active and the other being in stand-by mode. The necessary state sharing between the ASPs of an AS is out of scope of the standardization process. This state sharing is necessary because in case of a failure stable calls should not be affected.

It should be noted that the cluster management for the different adaptation layers is not identical and has to be implemented differently. A new working group, the reliable server pooling (RSerPool) working group, started to define a common solution for this problem. Like the development of SCTP, which started as a transport protocol for signaling transport, but is now a general purpose transport protocol, applications of the RSerPool protocol suite will not be limited to SIGTRAN applications. The usage of the RSerPool protocol suite will take some complexity away from the adaptation layers resulting in simpler solutions.

5 SCTP

The stream control transmission protocol (SCTP) is defined in [11]. SCTP is a connection oriented protocol and an SCTP connection is called association. An association is identified by the two SCTP endpoints. The following features of SCTP can be used to fulfill the performance requirements given in section 3.

SCTP supports a streams concept. An SCTP association provides multiple uni-directional streams. All user data is transmitted using one of these streams and SCTP guarantees in-sequence delivery within each stream. As different streams are independent, message loss of user data transported on one stream does not delay messages using other streams. This is avoids head-of-line blocking which could otherwise become an issue. Since in SS7 networks normally multiple links interconnect signaling points with the MTP distributing the traffic over these (independent) links using the SLS field in the routing label, head-of-line blocking is less of an issue in SS7 networks.[3] Independent SS7 messages, such as messages with different SLS value, can thus be transmitted on independent streams. The number of streams is negotiated during the initialization phase of the association.

SCTP also provides multihoming. This means that an SCTP endpoint can have multiple IP addresses. Each IP address of the peer is considered a path. In the association shown in figure 4 endpoint A has 3 paths towards endpoint Z and endpoint Z has 2 paths towards endpoint A.

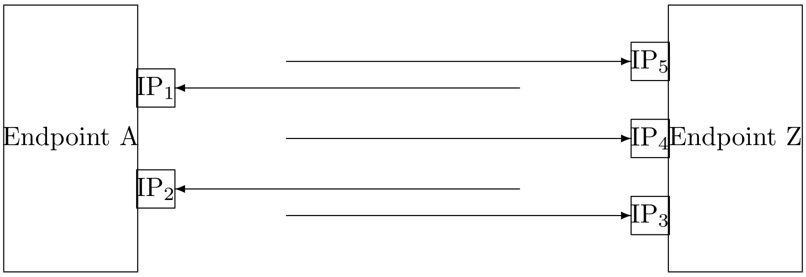

Figure 4: Path concept

Each SCTP endpoint chooses one of the paths as primary path. It is used to transmit all user data in the first attempt. In case of retransmissions an alternate path is used, if available. All paths are supervised using so-called heartbeat messages so that an SCTP endpoint knows if a path is active or inactive. After a limited number of retransmissions a path is considered inactive and a different one is chosen as primary path.

A properly designed SS7 network as described in section 2 always provides at least two physically separated ways to transmit user data. To provide the same level of redundancy using the IP based solution this multihoming feature can be used. Each SCTP endpoint should use at least two IP addresses and interfaces connected to two different routers. The forwarding in the IP network has to be engineered in a way that packets sent on different paths (in the SCTP sense) take two different ways through the IP network not using any common resources. This avoids single points of failure. This multihoming support of SCTP provides redundancy against network failures. Combined with the handling of host failures provided by the adaptation layers a communication with complete redundancy is possible.

It should be noted that this multihoming feature of the SCTP provides much of the functionality of the change-over procedures described in section 3. As a consequence it is arguably superfluous for the SCTP supports, like MTP2, the retrieval of unacknowledged messages in case of the failure of an association. While requirement 2 of section 3 can be fulfilled without this feature when using the SCTP, any user adaptation layers (e. g. M2PA) actually making use of this SCTP feature must then include additional functionality to fulfill requirement 3.

To fulfill the timing requirements given in section 3 one has to make sure that lost packets are retransmitted in a timely manner. SCTP has two mechanisms which trigger a retransmission:

1. Timer based retransmission: a timer is used to supervise successful transmission by analyzing acknowledgement messages. The timer value is based on measurements of the round trip delay. Exponential back-off is used on repeated retransmission attempts but the timer is modified to remain in the interval $[\mathrm{RTO}_{\mathrm{min}}, \mathrm{RTO}_{\mathrm{max}}]$. The suggested values are $\mathrm{RTO}_{\mathrm{min}} = 1\,\mathrm{sec}$ and $\mathrm{RTO}_{\mathrm{max}} = 60\,\mathrm{sec}$.

2. Fast retransmission: If a sender gets 4 reports about missing data these data chunks are retransmitted immediately without waiting for the expiry of the retransmission timer.

Using the parameter settings suggested in [11] it takes up to a minute to detect the failure of a path. Individual messages will be delayed for half a minute. These times violate the requirements given in section 3.

To fulfill the performance requirements one can do the following:

1. Limit the round trip delay of the IP network to 100 msec or less.

2. Make sure that sufficient bandwidth is available, preferably using QoS engineering to guarantee the required bandwidth.

[3] Multiple links, however, can cause load sharing problems [2].

3. Lower $\mathrm{RTO}_{\mathrm{min}}$ to the order of the round trip delay and lower $\mathrm{RTO}_{\mathrm{max}} \geq \mathrm{RTO}_{\mathrm{min}}$ to 500 msec to absolutely limit the retransmission timeout.

4. SCTP provides the ability to delay the sending of acknowledgements. This feature should be disabled to provide more accurate measurements of the round trip delay.

In normal load situations this will result in retransmissions after one round trip delay. On a failed primary path all packets sent to it before it is considered inactive will have to be retransmitted at least once. This time can be limited by choosing a small number of retransmissions forcing the path to become inactive. On the other hand, this number should not be too small to avoid oscillation. Also the number of gap reports which trigger a fast retransmission could be lowered. Using these mechanisms it is possible to lower the time to detect a path failure to a value less than 1 second.

One critical issue, which needs more research, are changes to the flow control algorithm used by SCTP. Currently it is compatible with the one used by TCP. While the behavior of SCTP is better than that of TCP in case of network failures due to the support of multihoming, it would be helpful to use a more aggressive retransmission algorithm. But then fairness with other traffic in the network is no longer guaranteed.

6 M3UA

M3UA (currently) defines a routing key as "a set of SS7 parameters and parameter values that uniquely define the range of signaling traffic to be handled by a particular Application Server. Parameters within the routing key cannot extend across more than a single SS7 Destination Point Code." In particular, a routing key could be based not only on the DPC of an MSU but, for instance, also on OPC or fields in the data of the user part, like ISUP's circuit identification code (CIC). In addition IP based network elements may share a point code with the SG. This general addressing architecture, however, can make a seamless interworking with SS7's network management difficult or impossible. Depending on the network configuration (mated/multiple SGs or not) a seamless MTP network management puts restrictions on the routing key granularity and/or imposes special requirements on the SG.

Case 1: IP network elements are only connected to a single gateway

In this case the RK granularity must not be finer than can be handled by the relevant user/application management (e. g. for ISUP: OPC, DPC, CIC). The M3UA implementation in the MGC can be fairly simple and several MGCs can share a point code with each other and/or the SG. If, however, one user part is distributed over more than one IP network element, the SG has no MTP management procedures available to inform the network should one such IP element become unavailable. In this case the SG will have to perform the appropriate user part management functions (e. g. circuit blocking/unblocking) and may even have to track call state. Such an approach, however, negates the advantages achievable by the separation of SG and MGC, for example. In addition, relying on a single SG for SS7 connectivity is also problematic from an availability point of view.

Case 2: Multiple/mated SGs

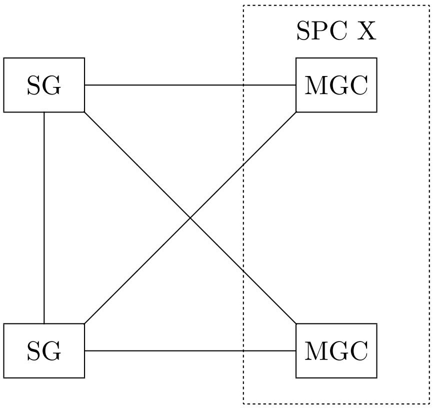

Figure 5:

In this case the RK granularity must not be finer than a destination point code, i. e., the SGs and the IP based elements must each have their own point code. Looking at figure 5, if the two MGCs share a point code and one MGC looses its connectivity to one of the SGs, that SG has no means to force the SS7 network to reroute the signaling traffic to the other SG without causing potential problems as MTP network management works on a point code basis only. In addition, as the M3UA in MGC now has more than one possible route to SS7 destinations it must now have full MTP3 endpoint functionality.

Although case 1 could be attractive when beginning to off-load traffic from a switch to a soft-switch (so as not to have to modify routing in the existing network) providing IP based elements with their own point codes and interconnecting them to mated SGs is the better solution with regard to the overall network reliability. Where point code shortages exist, appropriate MTP network partitioning can provide a solution.

On the other hand, from an MTP network management point of view there is no requirement that an RK must not "extend across more than a single SS7 Destination Point Code". A routing key might very well span more than one destination point code. MTP network management would only have to separately perform any management actions for all point codes concerned.

7 Summary

The protocol stack for the transport of signaling information over IP which is being defined by the SIGTRAN working group of IETF provides an important element for the convergence of voice and data networks.

Using appropriately engineered IP networks as underlying network layer and suitably modified parameters for the transport protocol (SCTP) forms the basis for achieving the performance and fault detection capabilities needed for signaling applications. Assigning IP based elements like MGCs their own point codes allows seamless network management in an SS7 network crossing the MTP/IP boundary. In combination this results in a converged signaling network architecture which can deliver the reliability and performance end users of the PSTN/ISDN have become accustomed to.

References

[1] K. D. Gradischnig, *Reliable SS7 Networks – The SS7 Challenge*, DRCN 1998.

[2] K. D. Gradischnig, St. Krämer, M. Tüxen, *Loadsharing – A key to the reliability of SS7-networks*, DRCN 2000.

[3] J.-C. Luengo et al, *Several Approaches to Robust Signalling Network Planning*, DRNC 1998.

[4] R. Shockey, *ENUM: Phone numbers meet the net*, Communications Convergence, July 2001.

[5] G. Sidebottom, et al., *SS7 MTP3-User Adaptation Layer (M3UA)*, Internet draft, work in progress, `draft-ietf-sigtran-m3ua-07.txt`, July 2001.

[6] ITU-T Recommendation Q.706, *Message transfer part signalling performance*, 03/93.

[7] ITU-T Recommendation Q.709, *Hypothetical signalling reference connection*, 03/93.

[8] ITU-T Recommendation Q.766, *Performance objectives in the integrated services digital network application*, 03/93.

[9] ITU-T Recommendation Q.2144, *B-ISDN Signalling ATM adaptation layer – Layer management for the SAAL at the network node interface*, 10/95.

[10] L. Ong, et al., *Framework Architecture for Signaling Transport*, RFC 2719, October 1999.

[11] R. Stewart, et al., *Stream Control Transmission Protocol*, RFC 2960, October 2000.