AEÜ International Journal of Electronics and Communications

http://www.urbanfischer.de/journals/aeue

SCTP – A Multi-link End-to-end Protocol for IP-based Networks

Andreas Jungmaier, Erwin P. Rathgeb, Michael Schopp and Michael Tüxen

Dedicated to Professor Paul J. Kühn on the occasion of his 60th birthday

Abstract **IP-based networks grow rapidly and coexist with other universal basic telecommunication infrastructures, namely with the ISDN, mobile networks and ATM networks. Since all infrastructures are capable of carrying all applications and services, there is a growing need for interworking. To be able to transfer signaling traffic originating in the other networks via IP networks, a family of protocols is being defined based on SCTP, a new end-to-end transport protocol. This paper describes some features of SCTP in detail and highlights some advantages over TCP. In addition, a modification of SCTP is proposed and evaluated which significantly enhances the performance in long delay environments.**

Keywords **End-to-end transport, Internet, Multi-link, Multi-homing, SCTP, Signaling, SS7, TCP/IP**

1. Trends in wide area networking

Traditionally, the focus in wide area networks (WAN) was purely on voice communication. Therefore, classical telecommunication networks were optimized for voice traffic. Data communication evolved independently with the introduction of mini-computers and was mainly used in local area networks (LAN). The need for wide area data networking – mainly for LAN interconnection – only developed over the last 20 years. Specialized data networks based on X.25, Frame Relay or leased lines were used. Single end users were connected by using the ubiquitous voice infrastructure, e.g. via modems.

Only with the increasing penetration of PCs in private homes and with the introduction of the World Wide Web (WWW), the need for a basic infrastructure for computer based communication became apparent. This led to the development of the integrated services digital network (ISDN [1]) based on the well proven concepts from the "classical" telecommunication networks.

The broadband-ISDN (BISDN) was intended to be the next evolutionary step in that direction allowing to overcome the bandwidth restrictions of the "narrowband-ISDN" and thus to provide a truly universal infrastructure for all services and applications. Since the circuit switched concept of the ISDN did not provide the required flexibility, a packet oriented concept, the asynchronous transfer mode (ATM [2, 3]), was defined as a basis for the BISDN. ATM networks have since been introduced by many public network providers.

With the introduction of digital systems, e.g. GSM (Global System for Mobile communication), cellular mobile networks have grown substantially. While GSM was mainly designed for voice traffic, the introduction of GPRS (General Packet Radio Service) [4] and UMTS (Universal Mobile Telecommunications System) [5] will provide enhanced support also for data and multimedia.

In parallel – and in competition – the LAN technologies and the TCP/IP [6, 7] based wide area networking concepts have evolved significantly. The deployment of public IP networks was supported by the liberalization and deregulation of the telecommunication markets. Due to the fact that the volume of computer based traffic has surpassed that of voice traffic already, there is a trend towards purely IP based networks [8]. These networks will be capable to also support real time services (voice, multimedia) with sufficient and guaranteed quality.

As a consequence, several basic infrastructures – with individual strengths and weaknesses – will coexist, which will all be able to support voice and data as well as video and multimedia. For economically sound overall solutions, cooperation among these infrastructures is a key issue.

One area of significant effort and progress during the definition of ISDN, BISDN and GSM was signaling and call control which led to the standardization and development of systems for user-network (Digital Subscriber Signaling System, DSS1, DSS2 [9, 10]) and network internal (Signaling System No. 7, SS7 [11]) signaling. The SS7 protocols [12–15] provide not only a sound basis for reliable control of ISDN and BISDN networks, but also for the "Intelligent Network" (IN) functions and for mobility management in cellular networks [16].

A crucial issue for the evolution of the heterogeneous networks is that signaling transport is also provided via pure IP networks with the required reliability and performance. Therefore, a corresponding protocol suite is currently under definition in the Signaling Transport (SIGTRAN) group of the Internet Engineering Task Force (IETF). The basis of this suite is the new end-to-end protocol for reliable data transport via IP networks called SCTP

Received September 26, 2000. Revised November 3, 2000.

A. Jungmaier, E. P. Rathgeb, Computer Networking Technology Group, Institute for Experimental Mathematics, University of Essen, Ellernstrasse 29, 45326 Essen, Germany.
E-mail: ajung@exp-math.uni-essen.de
M. Schopp, Siemens AG, ICM N MR ST 10, Hofmannstr. 51, 81359 Munich, Germany.
M. Tüxen, Siemens AG, ICN WN CS SE 51, Hofmannstr. 51, 81359 Munich, Germany.

1434-8411/01/55/1-46 $15.00/0

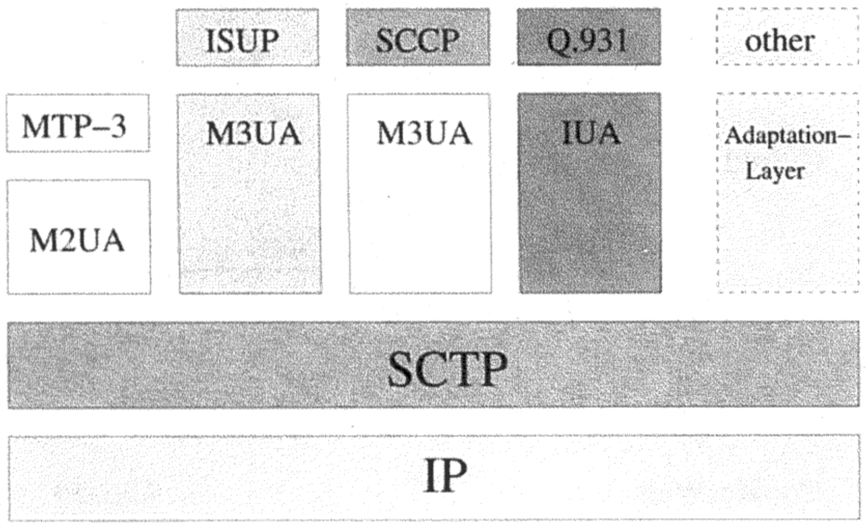

M2UA — MTP–2 User Adaptation Layer
M3UA — MTP–3 User Adaptation Layer
IUA — ISDN Q.931 User Adaptation Layer
SCCP — Signaling Connection Control Part

Fig. 1. SIGTRAN protocol architecture.

(Stream Control Transmission Protocol[1]) [17] presented in this paper.

2. SCTP – An overview

2.1 Signaling transport using SCTP

SCTP was originally designed for the transport of message-based signaling information over Internet Protocol (IP) networks. Within the framework [17–21] defined by SIGTRAN, the services of SCTP are made available to existing signaling protocols via adaptation layers. Figure 1 illustrates the protocol architecture showing several examples of adaptation layers and signaling application protocols.

Within SIGTRAN, special attention has been given to a scenario where telecommuncations signaling is transported between a Signaling Gateway (SG) and a Media Gateway Controller (MGC) [22]. A SG supports interworking in the control plane between SS7-based PSTNs (Public Switched Telephone Network) and IP networks (e.g. for Voice over IP, VoIP). Whereas call control related protocols like the ISDN User Part (ISUP) reside in the MGC, the transport-oriented SS7 protocols (Message Transfer Part Level 1-3, MTP [23]) are terminated in the SG. As shown in Figure 2, MTP-3 primitives are relayed between the SG and the MGC via the adaptation layer M3UA on top of SCTP.

However, this is only one example usage of SIGTRAN protocols. With the appropriate adaptation layers, SCTP can be applied whenever messages of classical telecommunications signaling application protocols (like e.g. SS7 user parts or Q.931 [9]) have to be transported over IP. Examples are:

- interconnection of PSTN islands via IP networks (virtual trunking)

[1] Initially called Simple Control Transmission Protocol.

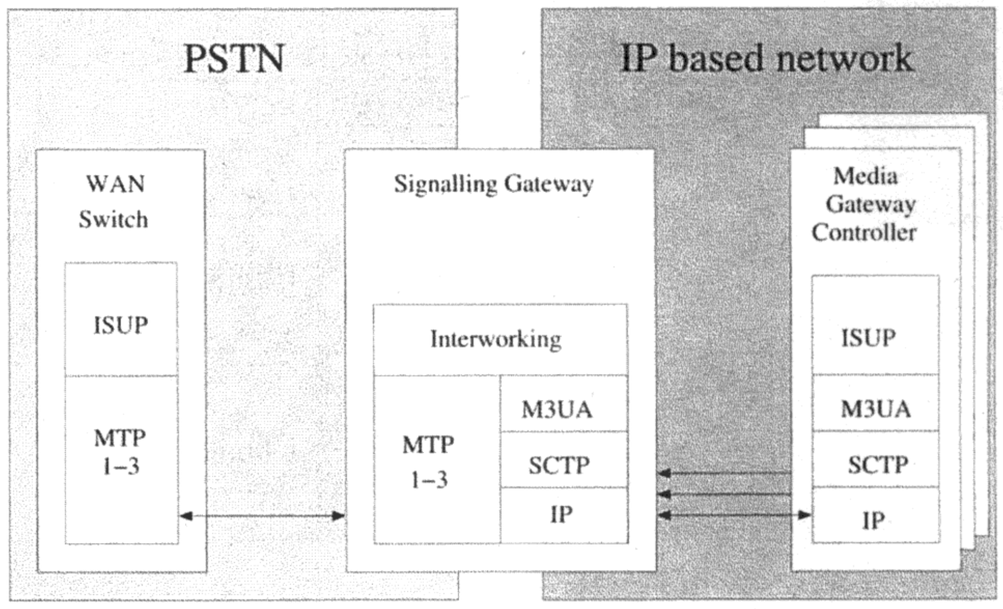

Fig. 2. Interworking between PSTN und IP-based Network with respect to ISUP-based call control signaling.

- access from the PSTN to IP-based Service Control Points (SCP) for Intelligent Network services
- access to IP-based Location Registers (HLRs or VLRs) for mobility management in mobile networks
- signaling transport in the UMTS Radio Access Network

Besides that, SCTP is also a new IETF transport protocol which can be used directly in any IP-based network. Due to its properties described in the next sections, it has the potential to even become an alternative to the established transport protocols TCP and UDP (User Datagram Protocol) for a variety of applications.

2.2 Basic properties of SCTP

SCTP is a reliable transport protocol operating on top of a potentially unreliable connectionless packet service such as IP. It offers acknowledged error-free non-duplicated transfer of SCTP packets (messages). Detection of data corruption and loss or duplication of data is achieved by using checksums and sequence numbers. A selective retransmission mechanism is applied to correct loss and errors.

While TCP works byte stream oriented, SCTP transfers data units which are called chunks. SCTP performs Path-MTU discovery, i.e. it determines the maximum possible size at which IP packets can be sent to a destination without being segmented. SCTP segments large messages into chunks which fit into such an IP packet. Several small chunks resulting from small messages may be multiplexed into one IP packet (bundling).

While TCP performs a strict in-sequence delivery of data per connection, SCTP has a more flexible delivery scheme. SCTP distinguishes different "streams" of messages within one SCTP "association" (see below). This enables a delivery scheme where the message sequence is only maintained per stream (partial in-sequence delivery) to reduce unnecessary head-of-line blocking among independent streams. Furthermore, SCTP provides a bypassing mechanism, such that messages are delivered to the SCTP user as soon as they have been completely received (order-of-arrival delivery).

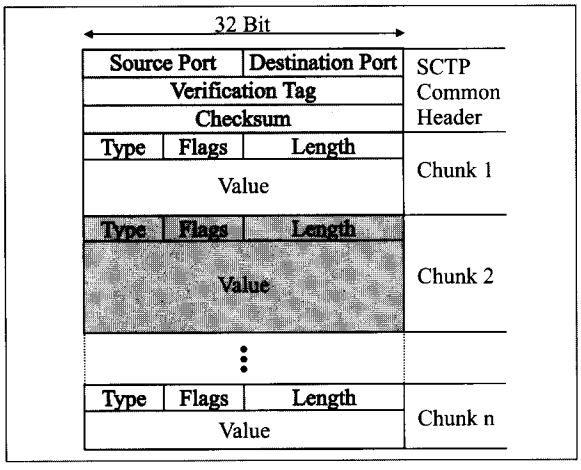

Fig. 3. SCTP packet format.

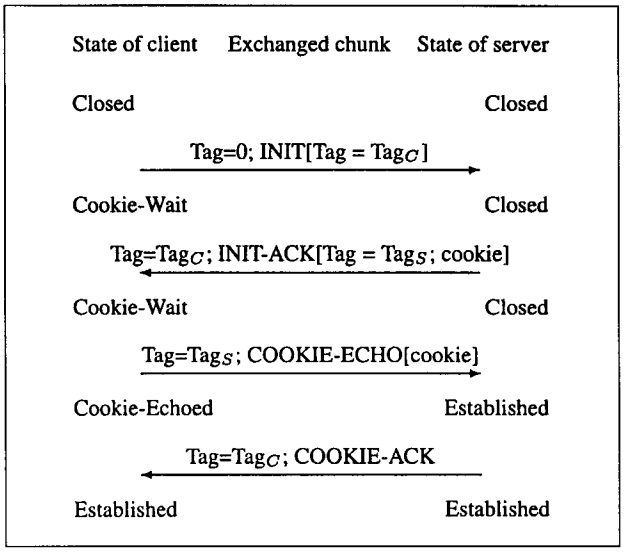

Fig. 4. Simple SCTP association setup.

Flow and congestion control in SCTP have been designed to assure that SCTP traffic behaves in the same way as TCP traffic does. This allows for seamless introduction of SCTP services into existing IP networks [24].

An advantage of SCTP over TCP is the support of so-called multi-homed nodes (i.e. nodes which can be reached under several IP addresses). A TCP connection is defined by a pair of transport addresses (IP address + port number). In SCTP however, each side of an association provides the other side with a list of multiple IP addresses (or other network addresses) combined with a single SCTP port number. An SCTP association spans over all possible source/destination combinations between the two involved nodes. Thus, each multi-homed node can be reached from another node via several paths (under several transport addresses). The status of each path is monitored by SCTP with respect to reachability, delay and number of consecutive retransmissions. Path monitoring, the usage of an alternate path for retransmissions and status dependent path selection make SCTP more robust against partial network failures than TCP.

As described below, specific mechanisms make SCTP also more resistant against blind attacks than TCP.

2.3 SCTP packet format

The protocol data units (PDU) of SCTP are called SCTP packets. If SCTP runs over IP (as described in [17]), an SCTP packet forms the payload of an IP packet.

As shown in Figure 3, an SCTP packet is composed of a common header and chunks. Multiple chunks may be multiplexed into one packet up to the Path-MTU size. A chunk contains either control information or user data.

The common header consists of 12 bytes. For the identification of an association, SCTP uses the same port concept as TCP and UDP. For the detection of transmission errors, each SCTP packet is protected by a 32 bit checksum (Adler-32 algorithm), which is more robust than the 16 bit checksum of TCP and UDP. SCTP packets with an invalid checksum are silently discarded. The common header contains a verification tag whose purpose will be described later.

Each chunk begins with a chunk type field, which is used to distinguish data chunks and different types of control chunks, followed by chunk specific flags and a chunk length field needed because chunks have a variable length. The value field contains the actual payload of the chunk.

3. Protocol behavior of SCTP

In the following, we will illustrate the behavior of SCTP in a few examples. These were derived from a formal SDL (Specification and Description Language) protocol description which we generated from the IETF documents and used as basis for our SCTP implementation. A description of the SCTP state machine can be found in [24]. The SCTP entity initiating association setup will be called "client", the peer entity will be called "server".

3.1 Setup of an SCTP association

Figure 4 illustrates the basic association setup (without error cases or collisions). Only the verification tag of the SCTP packet, the chunk type and the relevant chunk parameters (in brackets) are shown.

First, each SCTP entity is in the state "Closed". The client requests the setup of an association by sending an INIT control chunk containing a randomly generated tag value (Tag_C) as parameter. The SCTP packet itself contains a verification tag with value 0. The server handles the request and creates a transmission control block (TCB) containing all data related to the association. By using a hash function and a secret key, a message authentication

code (MAC) for the TCB is created. TCB and MAC are combined and sent to the client in the INIT-ACK chunk as so-called "cookie". The SCTP packet is marked with the Tag_C. A randomly generated tag value (Tag_S) is sent as parameter in the INIT-ACK chunk. The server remains in the state "Closed" without keeping any association related data or reserving any resources after having sent the INIT-ACK.

After receiving the INIT-ACK, the client returns the cookie to the server in the COOKIE-ECHO chunk using Tag_S in the SCTP packet. The TCB validated with the MAC is used to verify that the server had created the cookie initially. Only then the server assigns resources to the association and changes to the state "Established". The client is informed about the transition by means of a COOKIE-ACK chunk and transfer of data chunks can begin. An association is released by exchanging a sequence of acknowledged shutdown chunks.

3.2 Security issues of SCTP

While TCP uses a three-way handshake to establish a connection, SCTP uses the four-way handshake described above to set up an association. This additional effort enables security mechanisms which make SCTP robust against blind attacks, i.e. attacks where the attacker is not able to intercept PDUs but tries to interfere by sending malicious PDUs to one or more SCTP nodes.

The cookie concept and the fact that resources are only reserved in the server after cookie validation make SCTP robust against blind denial-of-service attacks, where an attacker sends setup requests using a fake source IP address. A prominent example of such an attack is TCP SYN flooding. To avoid such attacks, a simple cookie concept had also been proposed earlier as TCP enhancement. However, since it is not an integral part of TCP, it is not available in every implementation.

The verification tag, which is chosen during association setup and included in every SCTP packet, makes it difficult to insert extraneous content into an established association without intercepting SCTP packets, and guessing the 32 bit value is practically unfeasible if good random generators are used. SCTP packets with an incorrect verification tag are silently discarded.

3.3 Reliable transfer and flexible delivery

While TCP couples the reliable user data transfer with a strict order-of-transmission delivery, SCTP separates the reliable transfer from the delivery mechanism. This makes it possible to adapt delivery to the needs of the applications using SCTP. Some applications may only need partial ordering of packets while others might even be satisfied with a reliable transfer without resequencing.

SCTP operates on two levels:

- Within an association, reliable packet transfer is assured by using a checksum, a sequence number and a selective retransmission mechanism. Without reconstructing the initial sequence, every correctly received data chunk is delivered to a second, separate level.
- The second level provides a flexible delivery mechanism based on several independent "streams" of packets within an association.

3.3.1 Reliable transfer

Detection of loss and duplication of data chunks is supported by numbering all data chunks with a Transport Sequence Number (TSN) in the sender. The acknowledgements returned to the sender are based on the TSNs.

Retransmissions can be timer controlled, or initiated by the fast retransmission algorithm. For the former, the timer value is derived from continuous measurements of the round trip delay. Whenever the retransmission timer expires, unacknowledged data chunks are resent and the timer is started again after doubling its current value (as in TCP).

When the receiver detects one or more gaps in the sequence of data chunks, each received SCTP packet is acknowledged by sending a specific "Selective Acknowledgement" (SACK) control chunk reporting all gaps. Whenever the sender receives four consecutive gap reports for the same data chunk, this data chunk is immediately retransmitted (Fast Retransmit). Most up-to-date operating systems already support a similar optional extension to TCP [25].

3.3.2 Flexible packet delivery

The user of SCTP assigns each packet to one of several "streams" within an association. When an association is set up, the number of available streams per direction is exchanged between the peer entities. Within each stream, SCTP assigns independent Stream Sequence Numbers (SSN) to the user packets. These numbers are used at the receiver to determine the sequence of delivery. SCTP performs in-sequence delivery per stream to avoid head-of-line blocking among independent streams of packets within one association. With TCP, this could only be achieved by setting up several connections (one per stream) which would lead to additional cost and overhead.

SCTP also allows to mark packets for out-of-order (i.e. order-of-arrival) delivery. This can be used for important messages which may bypass others like, e.g., transaction abort messages of an application. If no resequencing is required, all packets could be marked accordingly. This is a feature which can not be efficiently provided by TCP.

3.4 The concept of transmission paths

An essential property of SCTP is the support of multihomed nodes, i.e. nodes which can be reached under several IP addresses. If the SCTP nodes and the (IP-)network

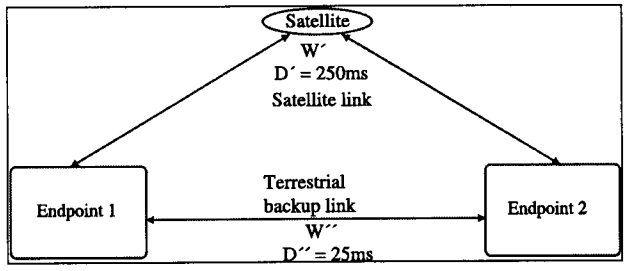

Fig. 5. Satellite link with terrestrial backup.

are configured such that traffic from one node to another follows physically different paths for different destination IP addresses, associations become tolerant against physical network failures and similar problems.

3.4.1 IP address management at association setup

If a client is multi-homed, it informs the server about all its IP addresses in the INIT chunk. Initially, the client only needs to know one IP address of the server, because the server provides all of its IP addresses to the client in the INIT-ACK chunk. SCTP is able to handle IP version 4 and IP version 6 addresses (even mixed). An SCTP entity considers each IP address of its peer to identify the endpoint of one "transmission path" towards this node.

If no explicit IP addresses are contained in the INIT or INIT-ACK chunk, the source IP address of the IP packet carrying the SCTP packet is used. This facilitates the application of SCTP when Network Address Translation (NAT) is used, e.g. at the edge of private IP networks. An additional optional feature has been introduced into the SCTP specification [17] allowing to use a host name instead of one or more IP addresses.

3.4.2 Path monitoring

An SCTP entity monitors all transmission paths to the peer entity of an association. HEARTBEAT chunks are sent over all paths which are currently not used for the transmission of data chunks. Each HEARTBEAT chunk has to be acknowledged by a HEARTBEAT-ACK chunk.

Each path is assigned a state which is either "active" or "inactive". A path is "active" if it has been used in the recent past to transmit an (arbitrary) SCTP packet which has been acknowledged by the peer. If transmissions on a certain path fail repeatedly, this path is considered "inactive".

3.4.3 Path selection

During setup of an SCTP association, one of the IP addresses from the returned list is selected as initial "primary path". By default, data chunks are transmitted via this primary transmission path. For retransmissions however, another active path may be selected. To support the measurement of round trip delays, SACK chunks should be sent to the source address of the IP packet which carried the data chunk that had triggered the SACK.

The users of SCTP are informed about the status (state and measurements) of a transmission path on request or when a transmission path changes its state. They may then instruct the local SCTP entity to use a new primary path.

4. SCTP in a multi-link scenario

4.1 A satellite link with backup

We will consider the simple network shown in Figure 5 consisting of two nodes connected by two paths (W' and W'') to examine the behavior of SCTP. W' is assumed to be a high bandwidth link with a high transmission delay D', whereas W'' provides a low bandwidth connection with low delay D'. This is a scenario where a satellite connection is supported by a terrestrial backup link. The satellite connection is assumed to carry the main traffic load for economic reasons. As mentioned above, an SCTP association can use both of these links. We assume that the SCTP entities on both sides use W' as primary path which carries a load high enough to cause fast retransmission in case of packet loss. According to the specification [17], SACK chunks will also be sent via W' in case of packet loss. However, the backup link W' will be used for retransmission.

4.2 Experimental network

Figure 6 explains the setup used for deriving the results presented below. Both endpoints run an SCTP implementation that has been developed in cooperation between the University of Essen and Siemens [22]. It implements the current specification [17] and was successfully tested during two interoperability meetings held by members of the IETF SIGTRAN working group.

The implementation is available under the GNU Public License from http://www.sctp.de and runs on common Unix operating systems (Linux, FreeBSD) as user-mode program. In the laboratory setup, both endpoints are

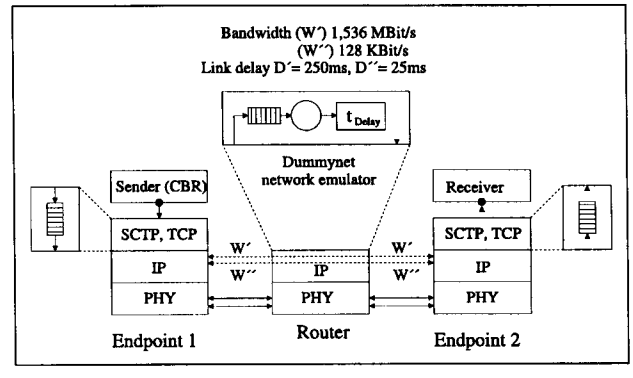

Fig. 6. Network structure – lab setup.

multi-homed Linux workstations that have routes to different IP sub-networks. They are connected by a FreeBSD workstation configured as router. The router uses the Dummynet network emulator [26] allowing to configure links with limited bandwidth, limited queues, configurable delays and random loss.

We emulated a T1 link via a geo-stationary satellite (bandwidth 1.536 MBit/s, $D' = 250$ ms) supported by a dual ISDN channel backup link (bandwidth 128 KBit/s, $D'' = 25$ ms). In this environment, both TCP and SCTP achieve a maximum throughput of 52 KByte/s in an error-free case, because of limitations through the bandwidth-delay product [24].

We assume that for a significant period of time 1% of all packets will experience randomly distributed transmission errors, e.g. due to atmospheric phenomena.

We observe a unidirectional data transmission where a sender transmits data units (1 KByte length) with a constant rate of 30 KByte/s, which both TCP and SCTP can easily support in the error-free case. For deriving the TCP results we used a modified version of the freeware program ttcp [27] that is commonly used to measure maximum throughput of TCP connections.

User packets include a time stamp allowing to determine the delivery delay. Delivery delay is defined as interval starting when the sender (cf. Figure 6) calls the SEND primitive of TCP or SCTP, respectively, and ending when the data is delivered to the receiver. It is determined by:

- the waiting time in the sender queue (operating system dependent parameter for TCP),
- the delay in the access queue of the link used,
- the transmission time on this link,
- the simulated link delay,
- the time until (lost) data is correctly reordered.

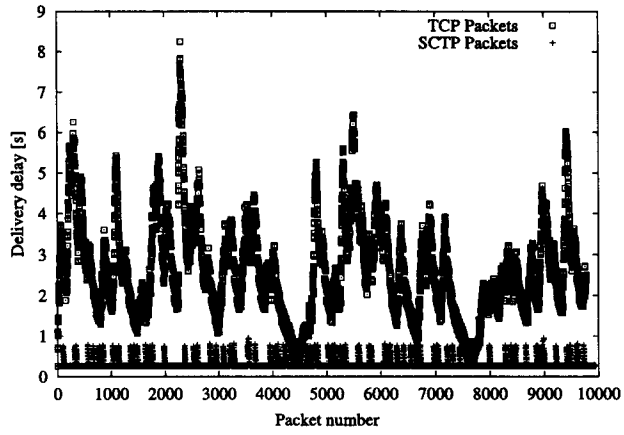

Fig. 7. Comparison of delivery delay for TCP and SCTP packets.

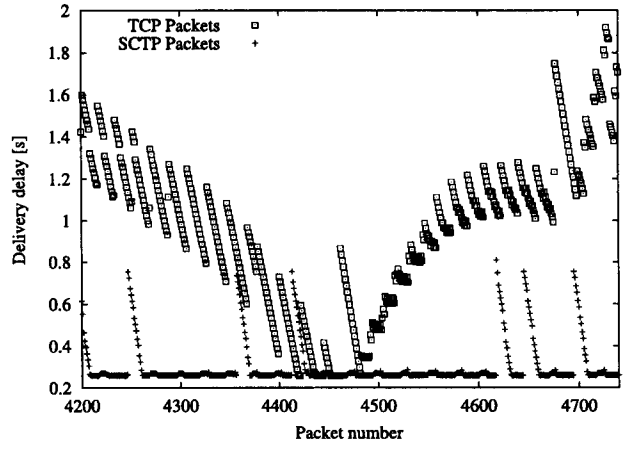

Fig. 8. Delivery delay for TCP and SCTP packets (detail).

4.3 Comparison of TCP and SCTP

In the scenario according to Figure 5, the TCP connection can only use the main satellite link. Modern versions of TCP, which implement the selective acknowledgement [25] will initiate packet retransmission after four successive acknowledgements have indicated a lost packet. Therefore, it is expected that the delivery delay of a lost and resent packet amounts to at least

$$t_{lost,TCP} = 3 \cdot D' + 4 \cdot \Delta t_{msg} \qquad (1)$$

taking into account only the additional link delays and the constant interval between two SEND calls of the protocol user (Δt_{msg}).

Standard SCTP will use the same SACK mechanism as TCP. However, SCTP will also use the secondary path and will retransmit via this path (cf. Section 3.4.3). In the described scenario, SCTP will benefit from the lower delay on the backup link. A lower bound for the delivery delay of a retransmitted SCTP packet is

$$t_{lost,SCTP} = 2 \cdot D' + D'' + 4 \cdot \Delta t_{msg} + t_{et''} \qquad (2)$$

where $t_{et''}$ is the emission time on the secondary link W''. The emission time of the packet on the primary path is neglected due to the high bandwidth as is the transmission time of the SACK chunks since these are very short.

Figure 7 shows the packet delivery delays for TCP and SCTP. TCP is strongly affected by the random errors. Buffering of data in the receiver's reordering queue while waiting for retransmission of a previously lost packet causes typical peaks of the delivery delay. Adaptation, e.g. after learning about a loss, and subsequent exhaustion of the congestion window result in additional buffering in the sender queue. Multiple losses during a short period cause a significant fluctuation of the delivery delay.

SCTP shows a more stable behavior, since by retransmitting (and quickly acknowledging) data via the secondary path, loss can be recovered independently from the properties of the primary path and without significant effects on the congestion window. Therefore, SCTP reaches a mean throughput of 29 KByte/s, while the mean throughput of TCP decreases to 18 KByte/s.

Figure 8 shows that the delivery delay for TCP packets reduces to approx. $D' = 250$ ms, when there is no loss. As soon as packets are lost (e.g. around packet number

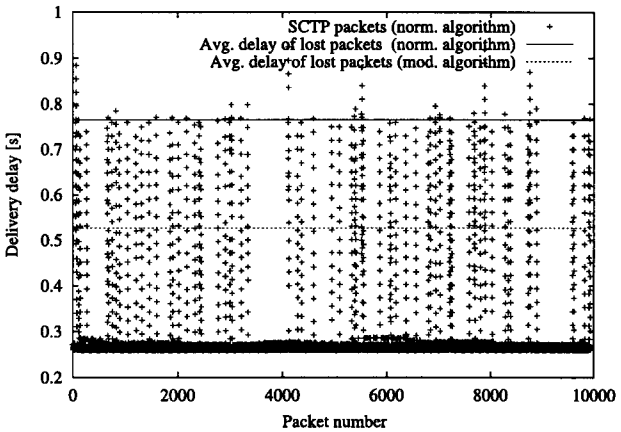

Fig. 9. Delivery times of SCTP packets in the satellite scenario.

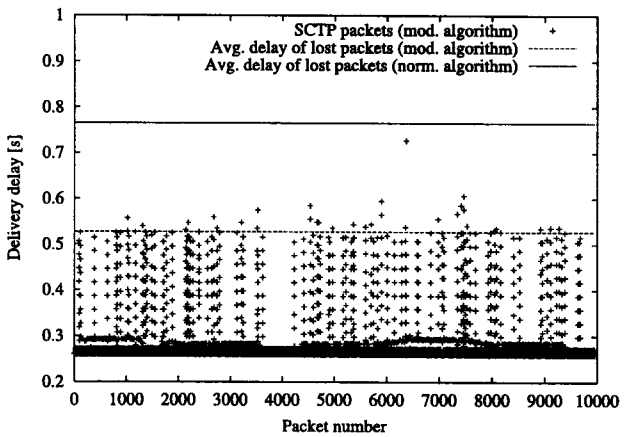

Fig. 10. Delivery times of SCTP packets (modified).

4470), the protocol has to wait for retransmission resulting in a delivery delay of at least $t_{lost,TCP} = 870$ ms.

4.4 Modified retransmission algorithm for SCTP

4.4.1 Standard SCTP

A standard SCTP implementation sends lost packets via the secondary path after four SACKs have indicated a possible packet loss. The SACKs themselves are sent via the primary path. According to Equation (2), this causes a delivery delay of at least $t_{lost,SCTP} = 708$ ms. From our experiments (see Figure 9) we measured an average delivery delay for retransmitted packets of 765 ms. This value is slightly higher than $t_{lost,SCTP}$ because of delays caused by the access queue to the secondary link.

4.4.2 Modified SCTP

In the following, we investigate a modification of the SCTP retransmission algorithm where all SACK chunks that contain gap reports and thus possibly indicate losses shall be sent via the secondary path W″ with the lower delay in order to minimize the delivery delay further. Thus it is possible to keep the delivery delay of packets relatively independent of the error rate of the primary path. This modification is particularly advantageous for applications that react critical to strong variations of delivery delay.

Acknowledgements are also used to compute the round trip delay of the path used. Therefore, only those may be used for that purpose, which do not contain gap reports, and are in fact transmitted via the primary path W′.

The delivery delay of retransmitted packets is now at least

$$t_{lost,SCTPmod} = D' + 2 \cdot D'' + 4 \cdot \Delta t_{msg} + t_{et''} \qquad (3)$$

where the emission times for the primary link and SACK chunks are assumed to be negligible, again.

As expected, the time needed for a retransmission decision decreases for the modified algorithm. Figure 10 shows delivery delays per packet as well as the average delivery delays of lost packets for both standard SCTP and the modified version.

For the described scenario a value of $t_{lost,SCTPmod} = 483$ ms can be computed. In the experiment we determined the average delivery delays of lost packets to be 527 ms which again is slightly higher than the estimated value due to the queueing delay in the access queue of the low bandwidth secondary link.

According to Equations (2) and (3) the average delivery delay for retransmitted packets decreases by

$$t_{im} = t_{lost,SCTP} - t_{lost,SCTPmod} = 225\,\text{ms}\,. \qquad (4)$$

The measured value is $t_{im,exp} = 238$ ms, and thus verifies the expected effects.

5. Conclusion and future work

SCTP is a new end-to-end transport protocol for IP based networks. Initially developed for transfer of signaling data originating from classical telecommunication networks (ISDN, etc.) via IP networks, it may also find widespread use for other applications in the Internet.

Some of the key advantages of SCTP are security features that prevent blind attacks, as well as flow- and congestion control features known from TCP which have been enhanced for multi-homing. Thus, the robustness of this protocol is significantly improved compared to TCP.

The measurements for a scenario with a satellite link protected by a terrestrial secondary path presented in this paper verify the obvious advantages compared to TCP. Moreover, with the proposed modification of the SCTP retransmission algorithm, a further optimization of the end-to-end delay can be achieved without compromising protocol conformance.

Currently further work is being done to improve SCTP security against more elaborate attacks. In addition, simu-

lation studies are performed to evaluate the protocol behavior of SCTP across a wider range of parameters and with more than two alternative paths.

References

[1] Bocker, P.; Arndt, G.; Frantzen, V.; Hagenhaus, L.; Huber, M.; Maegerl, G.; Rothamel, H.; Schweizer, L.: ISDN. Digitale Netze für Sprach-, Text-, Daten-, Video- und Multimediakommunikation. Springer, 1997.

[2] Huber, M.; Händel, R.; Schröder, S.: Atm networks – concepts, protocols, applications. Addison-Wesley, 1998.

[3] Rathgeb, E.; Wallmeier, E.: Atm – Infrastruktur für die Hochleistungskommunikation. Springer, 1997.

[4] Digital cellular telecommunications system (phase 2+); general packet radio service (gprs); service description; stage 1 (gsm 02.60). ETSI, November 1998. ETSI EN 301 113 (1998-11).

[5] Technical specification group services and system aspects 3rd generation mobile system release 1999 specifications (release 1999). 3rd Generation Partnership Project (3GPP), July 2000. 3G TS 21.101 V3.1.0 (2000-07).

[6] Rfc 793 – transmission control protocol – darpa internet program protocol specification. IETF, Network Working Group, September 1981.

[7] Rfc 791 – internet protocol – darpa internet program protocol specification. IETF, Network Working Group, September 1981.

[8] Hoffmann, G.: Aufbau eines gbit/s-Netzes für die Wissenschaft in Deutschland. Elektrotechnik und Informationstechnik **Heft 6** (2000), pp. 399–404.

[9] Digital subscriber signaling system no. 1 (dss 1) – ISDN user-network interface layer 3 – general aspects. International Telecommunication Union, May 1998. ITU-T Recommendation Q.931.

[10] Broadband integrated services digital network (b-ISDN) – digital subscriber signaling system no. 2 (dss2) – user network interface (uni) layer 3 specification for basic call/connection control. International Telecommunication Union, February 1995. ITU-T Recommendation Q.2931.

[11] Introduction to ccitt signaling system no. 7. International Telecommunication Union, March 1993. ITU-T Recommendation Q.700.

[12] Bafutto, M.; Kühn, P.; Willmann, G.: Modelling and performance analysis of common channel signaling networks. AEÜ Archiv für Elektronik und Übertragungstechnik **Vol. 47** (1993), pp. 411–419.

[13] Bafutto, M.; Kühn, P.; Willmann, G.: Capacity and performance analysis of signaling networks in multivendor environments. IEEE JSAC **Vol. 12** (1994), pp. 490–500.

[14] Kühn, P.; Pack, C.; Skoog, R.: Common channel signaling networks: Past, present, future. IEEE JSAC **Vol. 12** (1994).

[15] Willmann, G.; Kühn, P.: Performance modeling of signaling system no. 7. IEEE Comm. Magazine **7** (1990), pp. 44–56.

[16] Kühn, P.; Schopp, M.: Signalling networks for ISDN, in and mobile networks – modelling, analysis and overload control. Proc. of the 10th ITC Specialists Seminar on Control in Communications, 1996. Ed. Koerner, U. (ed.). Lund University.

[17] Stewart, R.; Xie, Q. et al.: Rfc 2960 – stream control transmission protocol. IETF, Network Working Group, October 2000.

[18] Ong, L.; Rytina, I.; Garcia, M. et al.: Rfc 2719 – framework architecture for signaling transport. IETF, Network Working Group, October 1999.

[19] Sidebottom, G. et al.: Ss7 mtp3-user adaptation layer (m3ua). IETF, Signaling Transport Working Group, September 2000. draft-ietf-sigtran-m3ua-04.txt, work in progress.

[20] Morneault, K. et al.: Ss7-mtp2 user adaptation layer (m2ua). IETF, Signaling Transport Working Group, March 2000. draft-ietf-sigtran-m2ua-04.txt, work in progress.

[21] Morneault, K. et al.: ISDN q.921-user adaptation layer. IETF, Signaling Transport Working Group, September 2000. draft-ietf-sigtran-iua-06.txt, work in progress.

[22] Jungmaier, A.; Schopp, M.; Tüxen, M.: Das simple control transmission protocol (sctp) – Ein neues Protokoll zum Transport von Signalisierungsmeldungen über ip-basierte Netze. Elektrotechnik und Informationstechnik **Heft 6** (2000), pp. 381–388.

[23] Functional description of the message transfer part (mtp) of signaling system no. 7. International Telecommunication Union, March 1993. ITU-T Recommendation Q.701 (03/93).

[24] Jungmaier, A.; Schopp, M.; Tüxen, M.: Performance evaluation of the stream control transmission protocol. ATM 2000 – Proceedings of the IEEE Conference on High Performance Switching and Routing, 2000. pp. 141–148.

[25] Floyd, S. et al.: Rfc 2018 – tcp selective acknowledgment options. IETF, Network Working Group, October 1996.

[26] Rizzo, L.: Dummynet: a simple approach to the evaluation of network protocols. ACM Computer Communication Review **27** (Jan. 1997), pp. 31–41.

[27] Parker, S.; Schmechel, C.: Rfc 2398 – some testing tools for tcp implementors. IETF, Network Working Group, August 1998.

Andreas Jungmaier was born in 1972 in Oberhausen and studied Electrical Engineering at the University of Duisburg from 1992 to 1997. He wrote his diploma thesis on parallel distributed algorithms for analog circuit simulation at the E.N.S.P. Strasbourg as ERASMUS student. In 1998 he worked as research assistant at the Data Processing Department of the University of Duisburg. Within a European research project, he spent 3 months at the Universidad de Cantabria, Santander (Spain) working on formal specification techniques for transport protocols as well as software/hardware co-design. Since April 1999 he is with the Computer Network Technology Group in the IEM, University of Essen. There he is doing research on SCTP and network security in public wide area networks.

Erwin P. Rathgeb was born in Ulm in 1958. He got his diploma and PhD degrees in Electrical Engineering from the University of Stuttgart in 1985 and 1991, respectively. From 1985 to 1990 he was member of the scientific staff at the Institute of Communication Networks and Computer Engineering (Prof. Paul J. Kühn) at the University of Stuttgart. From 1990 to 1991 he was a member of staff at Bellcore, Morristown, NJ, then at Bosch Telekom in Backnang. In 1993 he joined Siemens in Munich. In various positions in systems engineering and product planning he contributed to concepts for commercial ATM nodes and

ATM-based multiservice networks. Since January 1999 he holds the Alfried Krupp von Bohlen und Halbach-Chair for "Computer Networking Technology" at the Institute for Experimental Mathematics, University of Essen. His current research interests include concepts and protocols for next generation internets and network security.

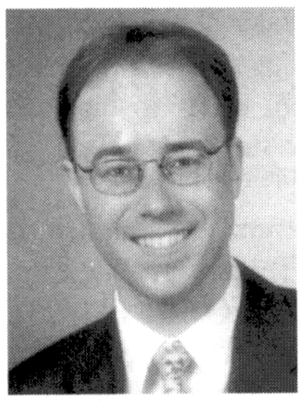

Michael Schopp was born in 1966 in Stuttgart, Germany. He received the Dipl.-Ing. degree in Electrical Engineering from the University of Stuttgart in 1992. He then joined Philips Communications Industries (PKI) in Nuremberg where he worked in a teletraffic group. From 1993 to 1998 he was a member of the scientific staff of the Institute of Communication Networks and Computer Engineering (Prof. Paul Kühn) at the University of Stuttgart where he was working towards a Dr.-Ing. degree with research in the area of modeling and performance evaluation of signaling and IN concepts for advanced mobility management in mobile communication systems. In 1998, Michael Schopp joined Siemens, Munich. Since then, he has been working for Siemens Information and Communications on standardization of the UMTS radio access network and on IP-based concepts for signaling and third generation mobile communication systems.

Michael Tüxen was born in 1966 in Oldenburg, Germany. He studied mathematics at the University of Göttingen and received the Dipl. Math. degree in 1993. From 1993 to 1996 he was a member of the scientific staff at the Sonderforschungsbereich "Geometrie und Analysis" (SFB 170) in Göttingen. He received the Dr. rer. nat. degree in 1996. In 1997 he joined the Systems Engineering group of ICN WN CS within the Siemens AG in Munich. He is working in the Study Group 2 of ITU on performance analysis of signaling protocols. At the IETF he is working in the Working Groups Signaling Transport (SIGTRAN) and Reliable Server Pooling (RSerPOOL) on design of protocols for signaling transport over IP networks and architectures for reliable distributed processing.