



An Introduction to SS7

By Guy Redmill, SS7 Product Manager

July 2001

white paper

Table of Contents

ABSTRACT	4
1 INTRODUCTION.....	4
2 BACKGROUND AND HISTORICAL PERSPECTIVE	4
Figure 2.1: In Band Signaling: CAS.....	5
Figure 2.2: Out of Band Signaling: CCS	5
3 PRINCIPLES OF SS7 SIGNALING: THE SIGNALING LINK.....	5
Figure 3.1: Overview of SS7 Stack	6
Figure 3.2: Links and Linksets	6
Figure 3.3: Routes and Routesets	7
4 PRINCIPLES OF SS7 SIGNALING: THE PROTOCOL STACK	7
4.1 THE MESSAGE TRANSFER PART	7
Figure 4.1: Layer 2 Message Structure	8
4.2 LAYER 4 CALL CONTROL PROTOCOLS, TUP AND ISUP	9
Figure 4.2.1: Example TUP Call Sequence	10
Figure 4.2.2: Example ISUP Call Sequence	11
4.3 DATA USER PART, DUP	11
4.4 SIGNALING CONNECTION CONTROL PART, SCCP	11
Figure 4.3.1: SCCP Global Title Translation:	12
4.5 TRANSACTION CAPABILITIES APPLICATION PART, TCAP	12
4.6 MAP, INAP, IS41	13
5 MESSAGE TRANSMISSION, SECURITY AND ERROR CORRECTION	13
Figure 5.1: Message Acknowledgement	13
6 NETWORK ARCHITECTURE AND ELEMENTS.....	15
6.1 NETWORK STRUCTURE	15
Figure 6.1.1: The SSP	15
Figure 6.1.2: Tandem Exchanges	15
Figure 6.1.3: Signaling Overlay Network	16
6.2 NETWORK ELEMENTS	16
6.2.1 The Local Exchange, Class 5 or Service Switching Point.....	16
6.2.2 The Tandem or Class 4 Switch	17
6.2.3 The Signal Transfer Point or STP.....	17
Figure 6.2.3.1: Combined Linksets	17
Figure 6.2.3.2 Primary and Secondary Routing.....	17
6.2.4 The Service Control Point or SCP	17
6.2.5 The Service Data Point or SDP	18
6.2.6 The Intelligent Peripheral or IP	18

6.3	SIGNALING MODES	18
6.3.1	Signaling Connectivity	18
	Figure 6.3.1.1: Fully Associated Signaling	18
	Figure 6.3.1.2: Non Associated Signaling.....	18
	Figure 6.3.1.3: Quasi Associated Signaling.....	18
6.3.2	Signaling Link Definitions.....	19
	Figure 6.3.2.1: Access Links	19
	Figure 6.3.2.2: Bridge Links	19
	Figure 6.3.2.3: Cross Links	19
	Figure 6.3.2.4: Diagonal Links	19
	Figure 6.3.2.5: Extended Links	20
	Figure 6.3.2.6: Fully Associated Links	20
7	NATIONAL IMPLEMENTATIONS AND ARCHITECTURAL VARIATIONS	20
8	SS7 AND ENHANCED SERVICES	21
	Figure 8.1: Overview of an IN Enabled Network	21
9	SS7 IN MOBILE NETWORKS	22
	Figure 9.1: Overview of a GSM Mobile Network.....	23
10	SS7 IN MOBILE NETWORKS	23
	Figure 10.1: SIGTRAN Architecture	24
11	THE FUTURE OF SS7	24
12	BROOKTROUT AND SS7	25

SS7 or Signaling System Number 7 is a set of protocols that describes a means of communication between telephone switches in public telephone networks. SS7 is a highly sophisticated and powerful form of Common Channel Signaling (CCS). The use of out-of-band signaling procedures offers considerable benefits over and above other signaling methodologies.

In common with many signaling protocols, SS7 is made up of a layered architecture. The lowest 3 layers together form the Message Transfer Part or MTP. This is responsible for the secure and reliable routing of messages, the content of which is provided by other, higher layers. MTP uses signaling links to route messages to the required destinations. Higher layers have different functions and are implemented as required by the network. Call control (i.e. the establishment and disconnection of calls) is handled by one of a series of Layer 4 Call Control Protocols, such as ISUP or TUP. Other functions are built on top of another layer called SCCP.

SS7 establishes a framework by which data is exchanged between systems in the network via dedicated signaling channels. The signaling link underpins the complete SS7 architecture. It enables communication to take place between entities within the network, permits the exchange of information and is essential for the effectiveness of the security features that make the network so resilient.

SS7 network architecture is explored in detail, as are basic call procedures. SS7 is discussed in relation to a variety of different network applications and signaling functions. The future role of SS7 is examined. The paper concludes by introducing Brooktrout's world-leading SS7 capabilities.

1

Introduction

There are two essential components to all telephone calls. The first, and most obvious, is the actual content—our voices, faxes, modem data, etc. The second is the information that instructs telephone exchanges to establish connections and route the “content” to an appropriate destination. Telephony signaling is concerned with the creation of standards for the latter to achieve the former. These standards are known as protocols. SS7 or Signaling System Number 7 is simply another set of protocols that describe a means of communication between telephone switches in public telephone networks. They have been created and controlled by various bodies around the world, which leads to some specific local variations, but the principal organization with responsibility for their administration is the International Telecommunications Union or ITU-T.

Originally designed for the purpose of conveying information relating to call establishment and teardown from exchange to exchange, the protocol architecture has been extended to cover a variety of tasks associated with collecting and reporting information necessary for the transmission of telephone calls. The SS7 standards now include specifications for a wide diversity of telephony management tasks and have proven to be extremely successful and resilient. As we have moved towards convergence between the public circuit-switched telephone network and the packet-switched IP world, SS7 has become the subject of significant attention as developers seek to integrate the two worlds and leverage the best of both. An understanding of SS7 is thus a vital component of an understanding of the current and next generation of public networks.

The purpose of this white paper is to provide a clear explanation of the role of SS7 in the telephone network today, to explore its origins and architecture and to examine its future in this rapidly changing environment.

2

Background and Historical Perspective

To understand SS7 we must first understand something of the basic inefficiency of previous signaling methods utilized in the Public Switched Telephone Network (PSTN). Until relatively recently, all telephone connections were managed by a variety of techniques centered on “in band” signaling.

All telephone conversations require a bearer trunk to transport them from origin to destination. In the early days of telephony, this was simply a single wire dedicated to a customer. Switching meant connecting customers together via intermediary pieces of wire. In addition to carrying the conversation or bearer “content,” all telephone bearer trunks also carried the signaling information necessary to control the telephone call concerned. This is known as “Channel Associated Signaling,” sometimes abbreviated to CAS. As one might imagine, this is fundamentally inefficient, as it means that even if the destination phone is unable to accept an incoming call, a complete bearer channel is fully occupied, from the point of origin to the point of destination, in the attempt to connect to it. Far better would be a way to signal to the destination phone without using the

valuable bearer circuit until absolutely necessary. Although this wouldn't eliminate the need for every customer to connect directly to the phone exchange, it would enable lines between phone exchanges, a valuable resource, to be used more effectively and indeed not to be used at all if the far end was found to be busy. This mode of signaling, where the information is carried separately from the bearer channels is known as "Common Channel Signaling," or CCS. CCS can also result in the allocation of a single, dedicated resource to signaling and allow it to be responsible for the control of large numbers of individual voice circuits.

Figure 2.1: In Band Signaling: CAS





Figure 2.2: Out of Band Signaling: CCS



SS7 is simply a highly sophisticated and powerful form of CCS. Another important benefit of CCS systems to consider is that they are essentially independent of switch and transmission technology. This means that evolution of SS7 standards is independent of the evolution of the underlying equipment. Standards can develop that are constant and can be used across different networks, and new functionality can be added without reference to the transmission equipment.

Since its inception, SS7 has assumed a role of growing importance, as most PSTNs have migrated from less efficient signaling methodologies towards a full implementation of SS7, and the majority of international networks have followed suit. Today, SS7 is responsible for routing calls across countries, between countries, and has a central role in mobile networks.

3 Principles of SS7 Signaling: The Signaling Link

As we have seen, the primary rationale for the development of SS7 was to find a way of signaling across networks without wasting valuable resources. Once this was achieved, a whole slew of services became possible. In order to understand the power and effectiveness of SS7, we need, therefore, to discover the fundamental elements that form the SS7 architecture.

In common with many signaling protocols, SS7 is made up of a layered architecture. Each layer has a specific role and responsibility. The lowest 3 layers together form the Message Transfer Part or MTP. This is responsible for the secure and reliable routing of messages, the content of which is provided by other, higher layers. MTP uses signaling links to route messages to the required destinations. Higher layers have different functions and are implemented as required by the network. Call control (i.e. the establishment and disconnection of calls) is handled by one of a series of Layer 4 Call Control Protocols, such as ISUP or TUP. Other functions are built on top of another layer called SCCP. We will focus on these layers in more detail in section 4.

SS7 establishes a framework by which data is exchanged between systems in the network via dedicated signaling channels. These channels are known as Signaling Data Links (SDLs) or simply as Links. Each SS7 system within the network acts as a Signaling Point (SP), and communicates with other SPs via dedicated links. SPs can be classified according to their precise function, but we shall turn to this subject in section 6.

Links connect SPs to their neighbors and form communication paths or routes between them. Within a SS7 network, all SPs are identified by a unique address. This address is called a Point Code (PC). All SS7 messages have a point of origin and a destination and hence are assigned an Originating Point Code (OPC) and a Destination Point Code (DPC). These identities allow them to be routed by the network to the appropriate destination. Point Codes are contained within messages built by the Message Transfer Part Layers. The length of a point code is determined by the specification that is deployed and consists of a series of bits. Because there is a finite length to this field, there is a limit to the number of point codes that can be active in any one network.

Links are responsible for the transportation of messages, directed by higher-level software, to their destination. The MTP interprets and directs messages to the appropriate destination and attaches the correct address labels (or instructions) to the message before directing them to the required signaling link.

The signaling link underpins the complete SS7 architecture. It enables communication to take place between entities within the network, permits the exchange of information and is essential for the effectiveness of the security features that make the network so resilient.

Links are usually organized into groups known as linksets. A linkset is a collection of links that share the same destination and are usually, but not always, established directly between SPs. When links are collected in linksets, the total load of messages is typically shared between the active links. This is one of the ways in which SS7 is inherently secure, for if one of the links in the linkset should fail, the remaining links can assume responsibility for the load formerly carried by the now inactive link. There can be up to 16 links in a linkset and a single SP may support a number of linksets that are established between itself and other SPs.

Figure 3.2: Links and Linksets






Figure 3.1 Overview of a SS7 Stack



When one SP is in communication with another, there is said to be a route between the two. A route is the path that exists between any two SPs. The route may include a single linkset, or multiple linksets; the term simply refers to the existence of a direct communications path between the two. SPs can also route messages to other SPs to which they are not directly connected. In these cases, the route describes the path taken by the signaling between the two end points. Where alternative paths exist from one SP to another, a collection of routes can be said to exist. A collection of routes to the same destination is known as a Routeset.

In figure 3.2 above, we can see that three SPs (A, B and C) are connected to each other via signaling links. In each case, we see two links, united in a linkset. Each SP node can be said to have a route to each other. However, the arrangement also allows for the possibility that any two points can be reached indirectly via a third, as well as directly. Because signaling travels separately from bearer traffic, we can see how complex routing strategies can be developed in SS7 networks.

Figure 3.3: Routes and Routesets



to ensure that messages reach their destination. For example, if the connections between A and B were to fail, messages from A could still reach B via C.

The simple point code address appended to each node simply serves to illustrate how each node can be uniquely identified by a numeric value, which is easily translated into binary form.

In Figure 3.3 above, we can see that with the addition of an additional node, the number of potential paths between, A and D for example, has been significantly increased. A has 2 routes to B and 4 to D, yielding 2 routesets with 2 and 4 routes respectively. SS7 signaling allows messages to be routed to SPs to which the initial node is not directly connected. This concept will be explored in more detail in section 6.

In the signaling hierarchy, we can see that the lowest common denominator, and hence most fundamental unit, is the link. Links are grouped in linksets, which make up routes, which make up routesets. In the same way that links within a linkset can work together to ensure secure transmission of information, the possibility of supporting alternative routes to the same destination serves to promote network security.

Understanding this allows us to approach the SS7 network from a sound architectural basis. Everything else that follows about SS7 can be understood in the context of moving information around the network across the most convenient available route

4 Principles of SS7 Signaling: The Protocol Stack

SS7 is really a suite of protocols that use a common transport mechanism for the distribution of messages between functional entities. The common transport layer is, as we saw earlier, the Message Transfer Part or MTP Layers. (NB, this is of course separate from the underlying transmission infrastructure, which is usually a 64 or 56 Kb/s channel in a multi channel E1 or T1, or a dedicated single channel link such as V.35). Above this layer, there are a number of alternative functional blocks (ISUP, SCCP, MAP, INAP, etc) that perform functions like call control management (for the establishment and tear down of conventional telephone calls), or search for and retrieve information relating to subscribers and the users of the call control services.

4.1) The Message Transfer Part

MTP is in turn divided into 3 distinct functional layers that perform specific tasks. Overall, it is concerned with the safe routing of messages and the management of SS7 links. The OSI reference model provides a framework for a layered architecture of functional entities, each of which may offer functionality to other layers. The SS7 protocol stack provides an architecture which is, in many respects, similar, but which departs from the OSI model particularly in the higher layers. However, at the lower 3 layers, the functionality follows the OSI model and will therefore be familiar to anyone with knowledge of this.

MTP Layer 1 represents the physical layer. That is, the layer that is responsible for the connection of SS7 Signaling Points into the transmission network over which they communicate with each other. Primarily, this involves the conversion of messaging into electrical signal and the maintenance of the physical links through which these pass. In this way, it is analogous to the Layer 1 of ISDN or other, perhaps more familiar, protocols.


MTP Layer 2 is designed to provide reliable transfer of signaling information between SPs. In this role, it examines transmitted data to check for errors and to correct them when they are discovered, if possible. With potentially large amounts of information being transmitted, MTP Layer 2 must also monitor message flow control, sorting messages based on queues and buffers. Precisely how security is achieved will be examined in section 5.

All SS7 messages are transmitted across SS7 signaling links. Flow control is of vital importance as the bandwidth available in a signaling link is usually either 64,000 bits per second or 56,000 bits per second. Link monitoring is therefore essential to the smooth operation of a SS7 system.

MTP Layer 2 is also responsible for the assembly of outgoing messages into packets known as signaling units, of which there are three types. Message Signal Units (MSUs) are the packets that carry the actual SS7 messages from a higher layer in the stack. Link Status Signal Units (LSSUs) are used to transmit information about the links themselves. The third kind of packet is the Fill In Signaling Unit or FISU. This final category carries no specific information itself, but is used as packing, as one of the features of a SS7 link is that it is never idle. FISUs are transmitted continuously across the link when there are no other messages to transmit. Gaps in transmission would be detected as errors, so it is essential that this background noise be maintained. Each message consists of an ordered sequence of bytes, separated by delineating flags or identifiers. A typical message can be broken down into key components. Figure 4.1.1 illustrates the structure of the 3 principal Layer 2 messages. Finally, it should be noted that MTP Layer 2 does not actually take any actions itself—it works under the direction of MTP Layer 3. This mechanism will be discussed in section 5.

Figure 4.1.1: Layer 2 Message Structure

Fill-In Signal Unit (FISU)



Link Status Signal Unit (LSSU)



Message Signal Unit (MSU)



Key:

BSN: Backward Sequence Number

BIB: Backward Indicator Bit

FSN: Forward Sequence Number

FIB: Forward Indicator Bit

LI: Length Indicator

SI: Status Indicator

SIO: Service Indicator Octet

SIF: Service Information Field

FCS: Frame Check Sequence Field

MTP Layer 3 has two basic functions. Message routing, which relates to the sending of received messages to appropriate destinations, either above or below it in the stack; and Network Management which concerns the control of traffic routing, the links which bear the traffic, and dealing with errors.

Each of the layers that sit above MTP Layer 3 can be considered as its users. They rely upon MTP Layer 3 for the safe delivery of messages they send to it and for the safe reception and onward routing of messages that are bound for them. Since there are several layers that can sit above MTP Layer 3 (SCCP, ISUP, TUP, etc), there must be some way of indicating which layer should receive a particular message. Message Discrimination, as it is known, is an important function of MTP Layer 3 and data carried in the MSU message allows this to take place.

Each message type is divided into a number of sub fields that contain information that is important in message routing and ensuring transmission integrity, as well as the actual data that is sent from end to end itself. MSUs have a sub field known as the Service Information Octet (or SIO) that allows message discrimination to take place. Each of the possible protocols that run at Layer 4 (SCCP, ISUP, TUP, DUP) has been allocated a particular value. MTP Layer 3 inspects the value and then ensures that the data part of the message is passed to the correct receiving layer.

Routing of messages to the appropriate external destination is equally important. The Signaling Information Field (or SIF) contains information that allows routing to take place. MTP L3 will complete this field by adding to it the OPC of the SP that generates the message and the DPC of the SP for which it is bound. Another important part of the SIF field is the Signaling Link Selection value (SLS). This establishes the correct link (out of those available) down which the message should be transmitted. As we have seen, links are grouped into linksets and part of MTP L3's task is to balance the distribution of messages that are transmitted between two points across the number of links that are available in a particular linkset.

SS7 signaling links are able to carry information about many thousands of bearer circuits, so load balancing or "loadsharing" is essential to the efficient operation of the network. If all of the information were to pass down a single link, the damage caused by failure of that link could be catastrophic. By distributing messages down available links in a linkset, MTP L3 is able to lessen the chances of a total breakdown in message transmission and attempts to avoid congestion on a single link. MTP L3 is also able to perform corrective actions in the event of link failure. These shall be explored in section 5.

4.2) Layer 4 Call Control Protocols, TUP and ISUP

One or more of the available Layer 4 protocols provides the data that is transmitted by MTP in the form of MSUs. These can be placed into two basic categories: circuit-related call control protocols (e.g. ISUP and TUP) and non-circuit-related protocols (e.g. SCCP). Circuit-related call control protocols are involved in the management of information relating to the connection and disconnection of telephone calls and are said to be circuit-related as each call requires the availability of a dedicated circuit (although, as we have seen, this is not actually established until the network has determined that the dialed party is available).

TUP or the Telephone User Part was the first SS7 call control protocol to be designed by the ITU-T and establishes a framework protocol for the exchange of messages designed to describe call set-up and teardown. These messages are used in the correct sequence to ensure that the relevant information is sent from the originating SP to the required destination and that the destination can signal back information regarding its availability to accept the call. Finally, when the parties involved complete their dialogue, information regarding the termination of the call is transmitted from both terminating SPs to ensure that the circuit is cleared effectively and that each party is available for further message transmission. If a call is not cleared effectively, the network can become confused about the status of the elements involved and problems can occur.

Specific messages have been defined for each of the relevant operations and the protocol stipulates how these may be used and the information that they should convey. In any given sequence of events, there are conditions that have to be met before the next stage can begin. If these conditions fail to be met, then the protocol defines alternative procedures to ensure that the system continues to operate effectively. Note that the protocol itself may not explicitly refer to the MTP layers. The security methods employed at MTP to ensure continuous operation are independent of the layer 4 procedures. This stratified approach further promotes reliability in SS7.

All of the initial information regarding a particular call is transmitted in either the IAM (Initial Address Message) or IAI (Initial Address Message with Additional Information). The IAI contains more detailed information regarding the call. The originating SP generates the IAM and IAI when an attached subscriber picks up the phone and dials their desired number.

There are two different modes of signaling that can be employed. The SP may begin signaling to the destination before all of the digits have been collected, if it is able to do so. In this case, correctly described as "Overlap" dialing, the first message it transmits will be incomplete and will be followed by additional messages, known as SAMs (Subsequent Address Messages) or SAOs (Subsequent Address One digit Message), which contain the remaining digits until the dialing is complete.


Alternatively, the SP will wait until all of the dialed digits have been collected before transmitting an IAM. This mode of signaling is known as "Enbloc".

On receipt of an IAM or IAI, the receiving SP will verify that the target subscriber is available. If it is, then the SP transmits an acknowledgement to the original SP in the form of an ACM (Address Complete Message). It is at this point in the course of a normal call that the bearer circuit is actually seized between the two SPs, although it is allocated originally by the initiating side.

The two parties will be in a position to communicate once the far end has answered the call. Again, this is signalled via a specific message, the AN or Answer Message. Additional information can be carried by any of these messages. For example, the AN can contain information relating to the charging strategy that should be applied to the call.

Once both parties have completed their dialogue, the call can be terminated via a specific clearing sequence that varies according to who hangs up, or “releases” the call first.

Figure 4.2.1: Example TUP Call Sequence



Each message can carry a defined set of information. The parameters that can be transmitted are described in the relevant standards (ITU-T Q721 – 725) and the software that controls the SP must ensure that they are set to the correct values before transmission and also interpret those received.

TUP does not encompass the transmission of ISDN data services and, although still widely deployed is largely being superseded by ISUP, which does support this feature. Although the ITU-T defined a core set of TUP standards, many national authorities implemented their own interpretation of the available parameters. Hence, it is possible to find considerable local variation in a supposedly common standard, although these variations are often not significant. Some particularly modified versions are to be found in France and the UK, where the TUP standard was evolved to support ISDN services prior to the launch of ISUP by the ITU-T. These versions are often referred to as NUPs (National User Parts). The NUP deployed in the UK is specifically known as BT-NUP or IUP; that in France as SSUTR2. China also has a unique variant that has many significant revisions from the standard TUP specifications. This is known as Chinese TUP.

ISUP, or ISDN User Part describes signaling functions relating to the management and set-up of telephone calls that can include voice, non-voice and data transmission. The ISUP procedures and protocol are described in ITU-T documents Q761-7. It builds on the principles of the TUP and extends its capabilities to include full integration with the ISDN access network. This means that calls that require end-to-end data transmission can be established across the network. ISUP also supports many additional services, such as those found in private networks (PBX networks). Due to the many enhancements, many countries are migrating towards an ISUP-based network. However, local differences continue to be implemented in order to meet specific requirements. This leads to some interesting problems at the boundaries between networks and we shall return to this topic in section 8.


ISUP messages can convey a great deal more information than those specified in the TUP protocol and the rich set of messages and parameters is also subject to considerable national variations. Calls can be overlap or enbloc and signaling procedures exist to implement a wide diversity of services. Many of these services are described under the umbrella term “CLASS,” or Customer Local

¹ The IAM will only contain the called number. IAI can also have calling, original, additional called, closed user group and charging information.

² For Overlap dialing, SAO only has 1 digit.

Area Signaling Services. These are services that are made available to subscribers that are connected to a particular switch. ISUP signaling can be used to deploy CLASS services effectively. Such services include “Ring Back When Free,” “Call Waiting,” “Number Display” and “Call Barring”.

Figure 4.2.2: Example ISUP Call Sequence



Key:

IAM: Initial Address Message

ACM: Address Complete Message

ANM: Answer Message

REL: Release Message

RLC: Release Complete Message

Note that, once again, it is essential for both legs of the telephone call to clear correctly and that strict procedures apply.

ISUP and TUP also include procedures for the maintenance of the circuits they control. There are specific messages that can be initiated to change the state of particular circuits. These include circuit resets and circuit blocking actions. When a system is brought into service, it will typically start with all associated circuits in a blocked state. By unblocking them in conjunction with the SP to which they are connected, a SP signals its availability and readiness to accept and place calls.

4.3) Data User Part, DUP

Data User Part, or DUP defines a protocol to control inter-exchange circuits that are used for Data Calls. This functionality has been superseded to some extent by the evolution of ISUP, and the protocol is now little, if at all, deployed.

4.4) Signaling Connection Control Part, SCCP

The SCCP provides enhanced functions to support the transfer of circuit-related and non-circuit-related signaling information. MTP only uses point codes routing to determine the destination of messages. SCCP however, uses a number of more discriminatory addressing methods to ensure that data reaches its destination.

SCCP uses two principal modes of data transfer: connection-oriented, in which a session must be initiated before data transfer can begin and connectionless, in which data transfer can take place without prior negotiation. These modes of transfer can be described as classes and there are 4 possible classes:

- Class 0: Connectionless – data is transferred without a fixed connection or session
- Class 1: Connectionless – data is transferred without a fixed connection or session, but in sequence.
- Class 2: Connection oriented – a session is initiated prior to the transfer of data
- Class 3: Connection oriented – a session is initiated prior to the transfer of data, the data is transferred in sequence

Although an individual signaling point can be identified by its point code, there may exist several components within the system to which different kinds of messages could be directed. MTP lacks the capability to distinguish between these, but SCCP can perform discrimination to ensure that messages reach their correct destination.


For example, many SPs contain databases that can be queried by other SS7 systems. There may be more than one database in a given SP. In this case, each database can be termed a “sub-system”. SCCP is able to reach any destination in the network by using a combination of the sub-system number (SSN) and the point code address. Since the SSN is represented by a single byte in the MSU structure, there can be a maximum of 256 SSNs at a particular SP. Any given SCCP layer maintains a record of available SSNs that it can access and much of the messaging associated with this layer concerns the monitoring of this data.

SCCP supports another extremely important means of routing data: Global Title Translation (GTT). A Global Title is a form of message used when the SCCP is unable to produce a PC and SSN to route data effectively. Instead, it produces a Global Title message, which contains the data it has received and an indication of the kind of data it requires. The SCCP concerned does not actually know where to find the required information, but it does know of a place that will be able to locate it. The SCCP system sends this global title message to a destination that is pre-configured in its own routing table as being able to provide an answer to the query.

In a given network, there may be only one location capable of dealing with Global Titles and this will connect to all other SCCP systems in the network. GTT is particularly useful as it allows systems to gain information from outside their own network without knowing where to look for it. For example, roaming subscribers in a mobile network will not be able to register initially, as their Mobile Subscriber Number will not be recognized by any of the databases in the visited network. Global Title allows the visited network to interrogate a GTT location and determine the origin of the visitors so that they can be registered with their home networks. If a particular database is unavailable, the Global Title location can redirect the query to another without the knowledge of the interrogating SCCP.

Global Title also limits the information that given SPs need to retain and monitor. If each switch needed to know where to locate all of the information that it might require, the volume of data each would have to store would be unmanageable. Instead, there can be nominated points within the network to which such queries can be directed. This is another example of the way in which SS7 can provide powerful call routing and control management, as the protocols have a built-in capacity for enhancements. Once services such as these are abstracted from the telephone exchange itself, they can grow without reference to the original apparatus of call delivery, as they depend only upon the transfer of information.

Figure 4.3.1: SCCP Global Title Translation



Interconnecting SCCP systems connect to a common Global Title location to permit access to databases that remain hidden to them.

4.5) Transaction Capabilities Application Part, TCAP

TCAP is primarily designed to be used for the querying and retrieval of information from databases. It formats data that can be presented using SCCP transport to a number of different databases. It can request that operations be carried out and await the result. It can also control the flow of information to the operation and the presentation of results to one or other of the higher layers that utilize TCAP's services. TCAP initiates queries and receives responses. In order to ensure that responses and queries can be correlated and sorted into the correct order, a numeric value is inserted into each query. The responding SP simply copies this number into its response so that the two can be cross-referenced.

There are two kinds of operations that are supported by TCAP. "Dialogues" take place between TCAP and one of the higher layers that use its services. Within a given dialogue, many operations may be active. Each operation may yield a result known as a component. Components can be stored by TCAP until such time as it is notified by a dialogue handling indicator that it is appropriate to dispatch them as a single TCAP message. When TCAP receives a message, it unbundles all of the separate components and sends

them individually to the appropriate higher layer. Since multiple dialogues can take place at any one time, each separate dialogue is given an individual identity that is conserved in the components. A TCAP transaction occurs when all stored components relating to a dialogue are presented to SCCP for routing to the relevant TCAP.

4.6) MAP, INAP, IS41


The Mobile Application Part (MAP), Intelligent Networking Application Part (INAP) and IS41 are layers of the SS7 stack that use TCAP to process their operations and procedures, and SCCP to transport their data to the relevant destinations. These layers allow the delivery of a wide range of key services above and beyond standard telephony and will be discussed in sections 8 and 9.

5 Message Transmission, Security and Error Correction

The security function of MTP L2 is concerned solely with the safe delivery and reception of signaling units. Each of the message types that we saw in section 4.1 is “read” by MTP L2 in order to ensure that they have arrived correctly. Since each message type begins with a fixed flag, it is possible to separate each individual message from every other. Each field within a message has an expected length. Variable length fields are usually divisible by 8. Hence, it is a straightforward matter to check that fields within messages of an appropriate length have been received. The next issue is to verify that all expected messages have been received. This can be accomplished by using sequence numbers.

The most important error correction methods concern the MSUs, which carry the real data from the application layers above MTP L3. Every transmitted message contains a Forward Sequence Number (FSN), as detailed in figure 4.1, that uniquely identifies it. This number can have a value between 0 and 127. When a message is received by either SP involved in the transmission, the receiving end must notify the transmitting side that the message has been received and that it is valid. This is achieved by the insertion of another field, the Backward Sequence Number (BSN). This contains the value of the last FSN to be received in a good message. Not every message requires acknowledgement however. For instance, if a series of messages are received, all in good order, the receiving end need only acknowledge the most recent.

Figure 5.1: Message Acknowledgement



The transmitting end retains all unacknowledged messages in a buffer until it receives an acknowledgement. The requirement to only acknowledge the most recently received successful message aids efficiency. The same principle applies in reverse. Each message has both a FSN and BSN, which allows acknowledgement to proceed on a continual basis. Under normal circumstances the Backward Indicator Bit (BIB) and Forward Indicator Bit (FIB) are set to equal each other. This means that they will usually be set to either 0 or 1. If a corrupt message is received, the receiving side will change the BIB to be different from the FIB. Once the point of origin receives this message, it will recognize that a message needs to be retransmitted and understand which one it should be (from the BSN). When 128 messages have been sent and acknowledged, the counter is reset and the FSN sequence begins again from 0. The procedure described here is known as Basic Error Correction. Essentially, MTP L2 performs a series of checks and cross-references on all transmitted and received messages to ensure that it has access to viable information. It will store everything until it has been notified that it is safe to discard it. Of course, this mechanism depends upon the link remaining in an active state. SS7 includes a number of other error

checking measures to ensure that messages are received successfully. For example, there is a general timer that is established once a message is transmitted and switched off when it has been acknowledged. If it fails to be turned off, it will expire and may cause more drastic action to be taken, such as the cessation of all traffic across a link.

When a link has been brought into service there is a suite of mechanisms that can be used to ensure both message integrity and transmission. During the initial alignment phase, the two ends of the link engage in an exchange of information that verifies that the link is capable of exchanging data. This process involves LSSU messages.

In addition to the transmission and re-transmission of individual messages, MTP L2 monitors the overall state of the link. The Signal Unit Error Monitor (SUERM) exists to count the number of errors that occur. These include errors that lead to re-transmission. When the number of errors reaches a critical point (64), then the link can be re-started. However, in order to avoid continual re-starting, the reception of a given number of good messages causes the “corruption counter” to decrement. In normal operating conditions there should be a balance between increments of errors and decrements of successes.

There are two other error correction methods of note: CRC and PCR. CRC stands for Cyclic Redundancy Check and this is a simple counter that calculates the total length of bits in a message that is transmitted and then performs an operation on this number. The resulting value is inserted into a message prior to transmission. Once received, the receiving end performs the calculation in reverse to discover the original value, that is, the length of the transmitted message. PCR or Preventative Cyclic Retransmission is a special mode of error correction that is reserved exclusively for use in systems that utilize satellite transmission. The problem with a satellite is that it moves relative to the earth. This means that normal timers that are established on message transmission are at risk of being exceeded even though there may not be any problems with the link. PCR does not use the BIB and FIB fields to register the receipt of a corrupt message. Instead, all messages that are unacknowledged are retransmitted during idle periods. Only 127 messages can be sent without receiving the appropriate acknowledgement. Once the buffer is full, no new messages can be transmitted and the system must enter a phase known as forced retransmission. PCR has several other requirements and is fairly processor intensive, but it does provide additional assurances over and above the standard error correction methods.

Other procedures exist to ensure secure transmission of LSSU and FISU messages. It is important to note that MTP L2 has procedures to cater for both transmission and acknowledgement of messages and their inspection to ensure that if they are found to be corrupt, they are retransmitted. However, it is MTP L3 that takes corrective actions to ensure continuous operation of SS7 systems.

It is the task of MTP L3 to route messages to active links. If a link fails and the failure is detected by MTP L2, it has to re-route messages towards a remaining active link, if one is available. Typically, the routing will take place to a link within the same linkset, but it is possible for the routing to take place to an alternative route. Indeed, it is here that we see the full capabilities of SS7 for routing that we described in section 3 displayed. Routing tables are established which provide MTP L3 with the data necessary to make decisions to route messages to the required destination. It manages signaling links, linksets, routes and routesets and can take appropriate actions to ensure safe delivery of all messages, if possible.

In addition to the routing of messages across available routes to their appropriate destinations, MTP L3 is also responsible for the even distribution of the load of SS7 messages. Where more than one link is available in a linkset, messages will be shared between them according to some kind of algorithm. One commonly deployed method is to send information that relates to even-numbered circuits across one link, odd-numbered across another. If it discovers that a particular link is unavailable, MTP L3 forces retransmission of the unacknowledged messages down an alternative link in the linkset (if possible). This is accomplished by use of a “changeover” message. Because the status of links is continually monitored, MTP L3 is able to detect when the link returns to service and can then force a “changeback”. This ensures that the load continues to be distributed evenly and keeps the system working efficiently.

Whatever the circumstances, MTP L3 tries to take remedial action to ensure transmission success. Furthermore, it actively monitors situations to ensure that operation is as efficient as possible and will always seek to restore the initial configuration. The resilience of SS7 networks is owed to the adaptability of MTP L3 in selecting alternative routes from those configured by the user and attempting to make best efforts at even distribution of messages.

6 Network Architecture and Elements


SS7 networks are composed of a number of different kinds of signaling point, each designed to meet different requirements. In most networks a hierarchy exists to further differentiate the network components. Additionally, there are a number of ways in which signaling can be passed through the network and the use of a particular mode relates to the specific function of the SPs concerned.

6.1) Network Structure

Entry to a SS7 network usually occurs at the point where normal subscriber lines connect to a local telephone exchange. This exchange is known variously as the "Local Exchange," "Central Office" or "Class 5 Switch". However, because it has an interface from what may be a wide range of access devices to the SS7 network, it forms a special kind of SP: the Service Switching Point or SSP.

Figure 6.1.1 shows one SSP linked in direct communication with another. This situation is not uncommon and typically reflects the simplest level of network design: it makes perfect sense for neighboring exchanges to connect directly to each other and transfer traffic via the most direct path. We can assume that a signaling linkset exists between the two SSPs and that normal measures of redundancy are implemented. The configuration also serves to illustrate one of the principal modes of signaling: Fully Associated Links. This and other modes shall be described in section 6.3

Figure 6.1.1: The SSP




If the network were to consist of a series of exchanges, each interconnected with each other, the resulting architecture would be extremely complex. However, a hierarchical approach alleviates this problem as a higher level of exchange can be introduced, the "junction" or "tandem" exchange. This allows local exchange SSPs to communicate with each other indirectly and removes the necessity for complex interconnection arrangements. Figure 6.1.2 illustrates this principle. The tandem exchange is also known as a "Class 4" switch.

Figure 6.1.2: Tandem Exchanges

Key:


LE: Local Exchange (SSP)

TE: Tandem Exchange



In some networks, particularly those in the United States, the signaling path is completely separated from the bearer paths, giving rise to a completely separate signaling network at a national level. This network exists solely for the purposes of routing SS7 signaling information to the required destination. Bearer traffic is routed via a separate transmission network under the control of local and tandem exchanges. Such a network is illustrated in figure 6.1.3. A special kind of switch is used for the exclusive control of signaling information: the Signal Transfer Point or STP. As network traffic grows, they are increasingly likely to be deployed within networks, particularly in an effort to separate different kinds of traffic.

Figure 6.1.3:
Signaling Overlay Network



In this case, the dashed lines represent some of the bearer channel transmission paths. Not every network has this level of complexity, but the diagram serves to illustrate how a network can be constructed to provide an efficient method of interconnecting disparate exchanges whilst offering multiple routing possibilities.

The final level of hierarchy is the point of interconnection to other, neighboring, networks. These interfaces are at special exchanges that may have to support protocol conversion, if the adjacent network supports a different user part variant. These exchanges are called “Gateway Exchanges” or “International Switching / Gateway Centers,” depending upon whether the other network lays in the same country or another.

6.2) Network Elements

All nodes in a SS7 network can be described as signaling points. However, there is considerable differentiation and specialization, which leads to a degree of classification being necessary.

6.2.1) The Local Exchange, Class 5 or Service Switching Point

This switch is the point entry into the SS7 network. It serves to connect various forms of access devices (analogue telephone handsets, digital subscriber lines, primary and basic rate ISDN, etc) to the SS7 network. Its other important role is to offer services to subscribers, as it may connect to external databases in the intelligent network (see section 8). Some of these services may be deployed locally, others at a national level.

6.2.2) The Tandem, Trunk Exchange or Class 4 Switch


This device exists to concentrate traffic from a number of local exchanges to simplify routing across the network. It can also act as a service switching point, in that it can also offer services to calls that it receives.

6.2.3) The Signal Transfer Point or STP

The STP exists solely to route SS7 messages to the appropriate destination. It does not offer termination services and it does not typically deploy a user part. It uses the capabilities of MTP to ensure that messages reach the correct destination. It refers to a database or routing table to determine the correct destination for a particular message. STPs are usually deployed in pairs, as they are linked to ensure redundancy. A device connecting to a STP will connect to both in the pair to achieve routing resilience.


One key advantage of STP functionality is that it enables a SP to signal to other SPs without being directly connected to them. This capability means that either dedicated STPs are deployed at the heart of a SS7 network, or STP-like functionality is implemented in trunk exchange types. In this case, the SP connecting to a device with STP connectivity may be able to specify to which distant (or non-adjacent) SP it wishes to direct traffic. Often, of course, this knowledge is not available, so the STP or trunk exchange may consult routing tables or other databases. Special kinds of linksets are used to connect to STPs to take advantage of these routing possibilities.

Figure 6.2.3.1:
Combined Linksets



A combined linkset is deployed when a SP connects to a pair of STPs. A linkset is configured between the SP and each STP within the pair. These two, separate linksets are then configured in such a way that the SP is able to loadshare signaling traffic between them, effectively creating a combined linkset.

Figure 6.2.3.2:
Primary and Secondary Routing



Primary and secondary routing procedures adopt the same principle as a combined linkset, but, whereas the SP with a combined linkset will loadshare between the linksets, the SP with a primary and secondary routing algorithm will nominate one linkset as the primary and the other as a secondary. Under normal conditions, all signaling traffic will be directed across the primary linkset. In failure conditions, the secondary or backup linkset will be utilized.

6.2.4) The Service Control Point or SCP

The Service Control Point or SCP acts as a service node within the SS7 network and allows services to be deployed in one location that can be accessed via special signaling messages by all switches within the network. This concept is known as the "Intelligent Network" or "IN," and is discussed in more detail in section 8.

6.2.5) The Service Data Point or SDP

The Service Data Point or SDP is a dedicated database used to store subscriber data within the IN.

6.2.6) The Intelligent Peripheral or IP

The Intelligent Peripheral or IP is a device that is usually associated with a Class 4 or Class 5 switch in IN networks, and offers voice interaction to subscribers, under the control of the SCP. It can only be deployed at a location that also terminates bearer channels and offers switching capabilities.

6.3) Signaling Modes


Signaling connectivity within a SS7 network can be managed in a number of ways; similarly, there are a number of different kinds of signaling links. The basic principle we have already described: Fully Associated Mode. Alternatives to this are Non Associated Signaling and Quasi Associated Signaling. Figures 6.3.1 – 6.3.6 illustrate these principles.

6.3.1) Signaling Connectivity

Figure 6.3.1.1:

Fully Associated Signaling

Signaling path shown as a dashed line, bearer paths as a solid line




With fully associated signaling, the signaling follows the same direct path as the bearer channels.

Figure 6.3.1.2:

Non Associated Signaling

Signaling path shown as a dashed line, bearer paths as a solid line




With non-associated signaling, the signaling is routed from a SP to a distant SP via at least two intermediary STP nodes.

Figure 6.3.1.3:

Quasi Associated Signaling

Signaling path shown as a dashed line, bearer paths as a solid line




Quasi Associated Signaling describes a signaling path between two SPs connected to the same pair of STPs. Some references make no distinction between quasi and non-associated signaling, using quasi to describe both.

6.3.2) Signaling Link Definitions

Six different kinds of links are defined in SS7 standards. They are described using an alphabetical notation.


A or Access links are the normal links deployed between a signaling point and a signal transfer point. They are highlighted in the figure below.

Figure 6.3.2.1:
Access Links




B or Bridge Links are used to interconnect two pairs of STPs.

Figure 6.3.2.2:
Bridge Links




C or Cross Links connect an individual STP to its pair in what is known as a “mated pair.”

Figure 6.3.2.3:
Cross Links




D or Diagonal Links are less commonly encountered, as they connect mated STP pairs to other such pairs at a different hierarchical level in the network. This may be to gain access to a resource connected to the distant STP pair.

Figure 6.3.2.4:
Diagonal Links




E or Extended Links connect signaling points to additional STP mated pairs. This provides additional routing possibilities.

Figure 6.3.2.5:
Extended Links



F or Fully Associated Links are established directly between two signaling points without passing through an intermediary signaling-only network. The signaling traffic follows the same path as the bearer traffic.

Figure 6.3.2.6:
Fully Associated Links



7 National Implementations and Architectural Variations

As we discussed in section 4, although there are global standards for SS7, there are many regional differences that have led to widespread variation in standards. The principal bodies responsible for the evolution of SS7 standards remain the ITU-T and ANSI organizations, but ETSI, Bellcore and a number of national PTTs continue to develop their own versions. While it is not important to know each standard, it is important to recognize the key differences and to be aware of the fact that SS7 is not a generic standard, but a collection of standards that follow the same core model.

Differences in the user parts are common. For example, the UK user part (now known as IUP) uses a message called an IFAM instead of the IAM used in standard TUP. ANSI ISUP, used within the United States does not support overlap dialing (described in section 4), whereas most other standards do. The result of this kind of differentiation is that there is often a requirement for a protocol converter at the boundary of two SS7 networks. Similarly, many manufacturers do not wish to invest in numerous different SS7 variants, so rely on suppliers to offer conversion solutions. Of course, vendors wishing to deploy the most complete solution should partner with a SS7 supplier that can offer a comprehensive range of SS7 variants.

Variation occurs at all layers of the SS7 protocol stack. MTP implementations, for instance, tend to fall into two camps—those which use the standard ITU-T 14 bit point code address scheme and those that use a derivative of the ANSI 24 bit point code scheme (Japan offers an unusual 16 bit version that has features of both). Expanding the length of the point code field allows many more addresses to be included within a particular network and allows for greater differentiation within the addressing scheme.

Despite efforts to reconcile national standards, it is likely that this widespread variation will continue for the foreseeable future—we are already seeing the release of new country specific versions of ISUP, such as UK ISUP and Chinese ISUP—so consideration will have to be given to these issues in any SS7 project.

Variation also occurs on an architectural basis. For example, almost all North American reference sources assume that the SS7 network exists as a packet-based, signaling-only overlay to the PSTN. The PSTN itself manages bearer connections under the control of the SS7 network. However, most networks outside of North America utilize 64 Kb/s channels inside an E1 as the principal means of transmitting SS7 signaling. This means that a SS7 channel can often be accompanied by bearer channels on the remaining capacity of the E1. This implies less reliance on STP networks outside of the US and this is partially true. In fact, there are many ways in which SS7 signaling can reach its destination. It is vital to remember that, for a supposedly global standard, a great deal of effort has been expended in adapting it to local circumstances in all respects.

8 SS7 and Enhanced Services


The Intelligent Network or IN defines a series of standards that allow new services to be added to existing networks with minimal upgrade costs and interference. Until recently, if a new service were to be deployed in a network, it would need to be implemented at every local exchange. This obviously places limitations on the upgrade process—validation and development of the new service is critically dependent upon the supplier of the local exchange equipment.

An intelligent network divorces service logic from the switching plane and concentrates services into dedicated network resources. SSPs and other Signaling Points communicate with these Service Control Points or SCPs via a special set of signaling commands (such as INAP, the Intelligent Networking Application Part), transported over TCAP messaging. Each new service is defined as a series of operations—the retrieval of data and the connection of calls to network resources, such as Intelligent Peripherals (IPs) that can collect user information such as DTMF tones and play messages to the user. This means that new services can be created rapidly, as they need only affect the SCP and IP, not the complete network.

Most IN services are initiated by triggers, such as dialed digits. A local exchange or SSP can detect when an IN service is invoked by matching dialed digits to those held in its routing tables. As soon as a call requiring IN services is recognized, it will hold the call and signal to its nominated SCP using IN signaling. The SCP will refer to its database to check both the caller (identified by the calling party number in the IAM field) and the particular service profile activated. Once this is identified, the SCP can negotiate with the SSP to provide information about the new destination of the call and make requests to connect it to a voice resource, such as an IP so that instructions can be relayed to the dialer.

A typical IN-enabled network is set out in figure 8.1 (below).

Figure 8.1:
Overview of an IN Enabled Network



In the diagram, a single SCP is available to multiple SSPs, indicating how services can be concentrated in one point.

Services that are enabled by IN include 800 number translation, calling cards, mass calling and distribution, tele-voting and follow-me. In North America, the term AIN (or advanced Intelligent Network) is usually synonymous with IN.

A more detailed explanation of IN architectures is available in another Brooktrout White Paper, *SS7 API and IN Architecture*.

9 SS7 in Mobile Networks

Although much of the technology in mobile networks uses specific radio interfaces, the signaling used across the fixed communications part is often based upon SS7. For example, the Mobile Application Part (MAP) defines operations and procedures for the control and presentation of subscriber information to ensure that roaming and messaging services can take place. MAP uses TCAP, SCCP and MTP as transport layers. The call control signaling is typically carried out using an ISUP derivative—this is important because there are points of interconnection between the mobile network and the fixed network.


GSM MAP has been specified by ETSI (ETS GSM 09.02) and is widely deployed around the world. One of the principal reasons for its success has been the innovation that separates the mobile handset from a smart card that uniquely identifies the subscriber and their subscription profile: the SIM card. The SIM card contains the International Mobile Subscriber Identity, a value that can be transmitted by the handset and used for authentication in any GSM network. This requirement to authorize and register users, particularly as they move between radio reception points, means that a GSM network requires sophisticated databases to hold all of the subscriber data. Global Title Translation has an important role to play in this process, because networks need to locate information outside of their own database infrastructure. IS41 is the standard used in the US. This is specified by ANSI.

The architecture of a GSM mobile network is illustrated in figure 9.1.

The architecture is divided into 3 planes: the Network Subsystem, the Base Station Sub System and the Mobile Station. The Network Subsystem contains the Mobile Switching Centre and associated databases. This is essentially similar to any other kind of SS7 signaling point, in that it controls the switching and routing of calls between mobile subscribers, and between mobile subscribers and fixed networks. It utilizes databases such as the Home Location Register (HLR) and Visitor Location Register (VLR) to monitor the presence of subscribers in the network it controls and to maintain data regarding their location. The MAP protocol defines interfaces between the MSC and other components of the architecture. These interfaces, identified by an alphabetical notation, ensure that the MSC is able to maintain all relevant information in real time, as subscribers move between base station cells and into other networks.

The Network Subsystem uses a dedicated protocol known as the “A” Interface (GSM ETS 08.08) to communicate with the Base Station Subsystem. Here, the Base Station Controller communicates with and directs the Base Station Transceivers that are the point of contact with the mobile terminals. The Mobile Station is composed of the combination of a mobile handset and a SIM card. Other interfaces include the “B” Interface, between the MSC and the VLR; the “C” interface, between the MSC and the VLR, the “D” interface, between the VLR and HLR; the “E” interface between MSCs; the “F” interface between the MSC and the EIR (Equipment Identification Register) the “G” interface between VLRs; and the “H” interface between MSCs and SMS Service Gateways.

Figure 9.1:
Overview of a GSM Mobile Network




10 SS7 and IP

IP based transport is rapidly gaining ground in the telephony industry. Long distance routing of telephone calls over an IP network is more cost effective than comparable routing over more conventional methods. IP access standards are emerging that offer great flexibility. Nowhere amongst the new generation of IP protocols, however, has emerged one that offers all of the signaling capabilities of SS7. This has lead to a movement to leverage the capabilities of SS7 in the IP world and to find a role for SS7 in softswitch architectures. This has been achieved by the Signaling Transport Group of the Internet Engineering Task Force (IETF).

The SIGTRAN standards, as they are known, describe a way of presenting SS7 signaling information over an IP transport in such a way that all of the benefits of SS7 are maintained. The standards allow next generation IP-based networks to interface with existing SS7 networks and to exchange information with no loss of service capability. SIGTRAN decomposes the SS7 stack and allows different layers to communicate using an IP transport layer. The architecture is shown in figure 10.1. Instead of using MTP as a transport protocol, SIGTRAN separates this from the user parts and simply transports the information that would be passed to each layer of an IP infrastructure. A new IP transport protocol, Simple Control Transmission Protocol (SCTP), has been defined that improves upon previous ones to ensure reliable transfer of information in a way that meets the requirements of SS7 systems.

Currently, SIGTRAN is primarily intended to be used at the interface between PSTN and IP networks, transferring information from a device known as a Signaling Gateway to another called a Media Gateway Controller, but it is likely that it will become more widely deployed as its benefits become more apparent, for example, to allow communication between different softswitches. In the current architecture, the Signaling Gateway (SG) terminates conventional SS7 links and performs MTP actions, but the signaling information for either MTP L3 or other layers is sent over an IP connection to the Media Gateway Controller (MGC), which then directs routing and connections in a similar manner to a conventional switch.

Figure 10.1:
SIGTRAN Architecture



Variations on this architecture define user adaptation layers for MTP L2, SCCP, ISUP and other protocol layers, essentially allowing them to be deployed remotely from the control device, yet for the control device to communicate with them over a common, non-proprietary, public interface.

11 The Future of SS7

SS7 is here to stay. Despite the revolutionary talk of all IP networks, such are the capabilities of SS7 that it will become an integral part of this infrastructure. More importantly, however, there is a tremendous investment in the conventional network that has not stopped. SS7 requirements continue to grow and access to SS7 signaling remains as essential as ever. Indeed, with a growing set of service providers and operators emerging, its importance is almost certainly increasing. SS7 allows a service provider to maximize connectivity options and it remains the best way to access the large subscriber bases that connect to the PSTN. Indeed, for most operators it is the preferred means of network connectivity and is not a matter of choice. New standards that built on its proven capabilities are emerging, such as WIN and CAMEL, and it is destined to play a huge role in 3G Mobile networks. In the network core it is likely to remain unchallenged for some considerable time and with the advent of new IP-based SS7 networks, it is likely to play a key role in the next generation of public networks. SS7 is not only here to stay, it is going to become more important than ever before.

12

Brooktrout and SS7

Brooktrout offers a comprehensive range of development tools for SS7, giving access to all layers of the protocol stack and supporting one of the broadest ranges of International variants available today. Our proven SS7 products are deployed in over 60 countries worldwide and in all major networks. Brooktrout development kits provide high-performance software with enhanced features to maximize the capabilities of your SS7 systems. Our evolving software libraries will enable you to take advantage of new opportunities in the next generation network. The Brooktrout advantage is clear:

- Widely deployed—60 countries
- Asynchronous, comprehensive API
- Generic API for all variants
- Global support
- Single Point Code redundancy
- High link density
- Flexible configuration options
- Portable
- Non-intrusive mode
- Full stack support
- Evolving roadmap

If you require SS7, there is no better partner than Brooktrout Technology. Contact us to see how we can help you.



U.S. Corporate Headquarters

Brooktrout, Inc.
250 First Avenue
Needham, MA 02494-2814
U.S.A.
Phone: +1 781 449-4100
Fax: +1 781 449-9009

European Headquarters

Brooktrout Technology Europe, Ltd.
Hoeilaart Office Park
Vandammestraat 5, Box 2
1560 Hoeilaart, Belgium
Phone: +32 2 658-0170
Fax: +32 2 658-0180

Sales

Needham, MA
+1 877 842-3944

Salem, NH
+1 603 898-1800

Los Gatos, CA
+1 408 370-0881

Miami, FL (Latin America)
+1 305 347-5113

Toronto
+1 416 860-6240

U.K.
+44 1344 380 280

Germany
+49 89 74120 133

Australia
+61 2 8221 8811