


Simply SS7



Simply SS7

Simply

SS7



Simply SS7 is a thorough overview of the SS7 standards which begins where *From Bell to Broadband* left off. It is written and published by, and intended solely for the use of SS8 Networks. Any reproduction of any portion of this booklet without the written permission of SS8 Networks is expressly prohibited.

© SS8 Networks
2001-2002

Table of Contents

1 Page 1	What Makes a Network	
2 Page 11	SS7 Network Architecture	
3 Page 35	SS7 User Part Functionality	
4 Page 83	Signal Unit And Message Format	



Foreword

Those who are about to confront telecommunications for the first time face a frustrating experience. For one thing, the topic is broad and complex. For another thing it is often difficult to find clearly written materials in the public domain. And, finally, telecommunications is an industry obsessed with acronyms.

By the time one manages to negotiate the SS7s, STPs, CRPs, SEPs, SSPs and HLRs it is time to deal with the ISUPs, ATMs, Sonets, and BISUPs. It seems that telecommunications people would rather say to their children “GOTT!!” than they would to say “Get off the telephone!!”.

Those of us who were not born speaking telecommunications sympathize completely. All of that is really part of the reason this booklet was written. It is our hope to lead the reader through the Signalling System #7 so gently, that you may never even be aware of the fact that you now know what the acronym SS7 means until it’s too late to unlearn it.

We must confess that we also have a hidden agenda in writing this. If your vision of the Signalling System #7 is sufficiently clear, it cannot fail to lead you to easier implementation of SS8 Networks products, and a greater appreciation of their value. When both the publisher and the reader benefit, it sounds like a great reason to write a book. Let’s find out.

Al Trickey
Technical Training Manager

Acknowledgments

Nothing is written in a vacuum. I had at first thought that the best way to handle this page might be to simply print the telephone listings for all of those who worked for the company's Shelton group. Surely there has rarely been any company that could boast such a long list of helpful and cooperative fellow employees. But if I were to do that you might miss some of the most important contributors.

First, I wish to thank the many engineers who not only contributed to my education, but also gave freely of their time to handle some of very deep questions which some of our customers attending classes have, from time to time, put forth. Among these mentors have been Akif Arsoy, Cemal Dikmen, Murat Erkam, Kerem Irten and Joe Conigliaro. Indeed, Murat Erkam provided the technical editing for the first book (From Bell To Broadband), while Christopher Rogalin (of the Product Validation group) provided the standard editing.

For this book, specifically, I am highly indebted to Janet Pikulik who at one time bore overall responsibility for both training and documentation. It was she who refused to give up on me during my early (and sometimes painful) introduction to the SS7 network. Years later I sought her out whenever I struggled with any SS7 issue. She always provided the answer, or at least got me pointed in the right direction.

The thank you for the technical editing of this book belongs to Paiman Nodoushani. Imagine my surprise to have someone actually *volunteer* to perform what is often such a boring duty! But, volunteer he did. And, despite a busy schedule, he completed the task in no time.

Finally, let me say that we were fortunate to have on staff a technical writer who is a master of his craft, to whom fell the task of performing the standard editing. His name is David Belisle. And, yes, David, absolutely, although belatedly, I will, henceforth, avoid the overuse of commas. And yes, I promise in future to avoid the use of "quotes" for "emphasis".

My thanks to all of the above.

Al Trickey
Technical Training Manager


Section 1

What Makes A Network A Network

What is a network?

It might appear that the title page contains a typo. Two words are repeated. Yet, actually, that fact serves to make a point about one powerful aspect of a Telecommunications network. That aspect is the use of redundancy to ensure that a message will be delivered. If there is only one route by which you can deliver a message you had better hope that you encounter no problems on that route. Back in the early days of the Telegraph, often cutting a single wire would bring communications to a halt. Cutting a wire in today's networks is unlikely to prevent the message from getting through because other paths are almost always available.

When dealing with networks you'll often hear some reference to something in nature that resembles the structure of a network. The term "web" for example appears frequently in the graphics and text offered by those who provide access to various networks. The word "net" itself is a reference to the way in which network locations are connected together. Let's see what a "web" tells us about a network.



Of course this doesn't look like a real web. Most spiders make webs that are essentially spirals rather than the concentric circles shown here. And, spiders are generally much more graceful builders. Nevertheless the drawing serves to illustrate the point.


The structure of a web offers the spider numerous advantages. First, it covers an area hundreds or thousands of times broader than his own reach. Thus, he multiplies the chances of having an insect fly within his range.

Secondly, the vibration of an insect striking any part of the web is communicated to every part of the web thanks to the hundreds of connected intersections. The spider need not be looking in the right direction to know when he has trapped his prey.

Simply SS7


Finally, the radiating spokes provide roads to the prey which allow the spider to reach any spot on the web quickly before even a large insect has time to break free.

Most networks make use of the concentric aspects of our drawing. For our example we'll look at the network known as the PSTN (Public Switched Telephone Network). The easiest way to begin is to look at the local exchange office (often known simply as the End Office). Here we find a switch capable of making connections for the transmission lines that enter the building. For a local office most of these will be the lines that go to homes and offices and end at a telephone or, perhaps, at an office switchboard.



In the drawing above, if the network were to go no further, then this switch could connect any telephone it serves to any other telephone it serves. But, of course, there are other end offices servicing their own collection of telephones. If the drawing represents the entire network, none of the telephones shown here could be connected to any of the telephones serviced by any other end office. That could be solved by running a transmission line from each end office to each of the other end offices. That would work; but, it would be very inefficient.

A much better way would be to connect this local office to a switch whose main purpose would be the interconnection of many local offices. The local switch, then, would not need a transmission line to every other local office. Instead, transmission lines going directly to the intermediate switch would be sufficient.



The result would look something like the above drawing. Here a switch is used simply to connect switches. Note the result in the rapid expansion of the number of telephones that can now be connected. We have elected to keep the example simple by showing an intermediate switch which only connects four end offices. Now picture that switch connecting as many end offices as each end office connects telephones.


In our example, each end office could connect twenty telephones. If the intermediate switch connected twenty end offices, four hundred telephones could then be connected. In practice, of course, an end office connects a huge number of telephones. The first intermediate switch likewise connects many more end offices.

Taking our example one step further, the Tandem offices can be connected to another switch whose job is to connect switches. In this case the new switch would likely be a switch for long distance communication. Under previous FCC rules, this switch would be owned and maintained (or leased) by a long distance carrier or Interexchange Carrier (IXC).

.

Simply SS7

When this connection is made, our network begins to resemble the following drawing



As you can see, each time a switch at a higher hierarchical level is used to connect lower level switches, the network grows more rapidly. So does the potential for disaster. Since each switch connects more and more telephones, the loss of such switches would result in the loss of communications to an ever greater number of phones. There is no way to guarantee that a switch will not fail. Therefore, the network makes use of the concepts of redundancy. On occasion, this might mean putting in duplicate facilities. On many more occasions it simply means taking multiple advantage of what is already there.

For example, we showed end offices connecting into a local tandem. In reality end offices often connect to more than one local tandem. Local tandems often connect to other local tandems. It is this multiplicity of connection that characterizes a real network. In the process, determining the route by which information must go (routing) becomes more and more flexible. When enough connections exist, difficulties at any given point of connection have little effect on the network and phone calls can be connected (or messages can be delivered) even in the face of numerous disruptions within the network.

Signalling System #7 - How Come?

Before we begin to explain *what* the Signalling System #7 is, it might prove helpful to understand *why* the Signalling System #7 is. The SS7 network was not in existence 75 years ago, but the telephone network was. Actually the telephone system didn't begin as a network. It began simply as telephone lines connected to switches. The original switches were really isolated and served only the customers to whom they were directly connected. Very early in the development of the telephone system, it became apparent that the most efficient way of extending the distance over which a telephone call could be connected was to simply connect existing switches.

Almost as soon as the means of interconnecting switches became available, it also became obvious that making connections from any given switch to multiple switches vastly increased the possibilities for routing a call. The first network was born. Today that network is known as the Public Switched Telephone Network. Even if you aren't an acronym fan, you'll have to admit (at least this time) that it's easier to say PSTN.

The PSTN (Public Switched Telephone Network) began in an almost haphazard fashion. But it wasn't long before telephone people began to see the value of thinking through just how and where connections should be made. Network design had arrived on the scene. As a part of that design certain switches were connected in such a way as to do nothing more than provide access to the area of a local connection. This connection of one area to another resulted in the pairing of two more or less local areas and the switch became known as a tandem. Then came the development of regional switches which, by design, connected the tandem switches together. Nearly every network today employs a similar type of hierarchy.

To the casual user of a telephone it might appear simple to connect two telephones together for the purpose of a conversation. If there were only two telephones it would be. But of course there are hundreds of millions of telephones. Information about the call must be provided for the caller, the party being called, and the telephone company connecting the call.

Let's start with the caller. If you unplug your telephone, the first person who tries to make a call from that phone will know there is a problem as soon as they pick up the handset. Why? Because there will be no dialtone. Normally, when we remove the handset from the cradle a switch in our telephone closes to connect two wires. At the local office to which the phone is connected, the existence of this completed circuit is sensed at a subscriber interface (where the individual phone lines enter the building). To the telephone company this is a signal that you intend to make a call. This signal is known by the name "off hook." The telephone company's response to the "off hook" signal is to signal the caller by sending the humming sound we call "dialtone." To the caller this signals that the line is O.K. and that the phone company is ready to receive the instructions which the caller sends by dialing. We are all so familiar with the dialtone that few of us would expect to be able to make a call on a phone which returned no dialtone.

When the circuitry has been connected to the office serving the party being called, that office once again needs to refer to the state of the called party's line at the subscriber interface. If that party is

Simply SS7

not using the phone, the handset will be “on hook” and the circuit that allows speech to be transmitted will not be complete. In such a case the called party’s telephone office will send a voltage out on a line with the purpose of ringing the telephone. With that done the same office generates and returns an interrupted tone which travels all the way back through all the connected circuitry to the phone of the party who is making the call. That party hears this “ringback” signal and knows that the phone they have called is ringing.

Of course if that distant phone is in use (or if the handset is simply out of the cradle) that party’s telephone office will make no attempt to connect the call. Instead, they will generate and return a tone we know as a “busy” signal. The party making the call will know the call cannot be connected, and they will hang up. Or will they? If they hang up, the local office sees the circuit interruption as an “on hook” signal. But the local office cannot yet dismantle the connection to the tandem. That’s because the tandem needs to be signalled that the caller has gone “on hook”. The tandem, in turn needs to signal the regional office of the same condition. Only when the signalling is completed to the next switch can any switch release the connection.

The telephone company was aware that signalling in this way created problems. The signals were sent through the same circuits that would carry the voice during a conversation. Consider what happens when the result of the call is a busy signal. The calling party may hang up quickly. Or they may listen to the busy signal for a few seconds. All during this time no conversation can take place over this particular circuit. Nevertheless the circuit must be maintained for the sole purpose of returning the signal. The busy tone originates at the far end. If this is a coast to coast phone call, thousands of miles of wiring are involved in the circuit.

After tying up a circuit for a few seconds the caller may hang up and try again. If the result is the same, more circuitry will be tied up for seconds. How many calls will they make before they finally get to talk to someone? Multiply this by the millions of phone calls being made every hour and you’ll get some idea of the size of the problem.

By mid-century the only way the phone company could compensate for such lost conversation capability was to increase the size of the telephone infrastructure. That meant more wires, more switches, and more of everything required to provide telephone connections. Circuits weren’t being used efficiently; but, it was the only way.

Perhaps it was fortunate that all of this became obvious to the telephone companies during the booming telephone services demand of the post-war (WW II) period. The same booming economy that created the demand brought the problem into focus and began to provide the solution.

By the late Fifties electronics had already begun some very rapid advances. Solid state electronics was a reality though not yet widely in use. Machines were already talking to machines over transmission lines (Teletype). Some primitive digital networks were evolving. As usual, advances were largely of a proprietary nature with each provider of telecommunications hardware taking an approach that seemed valid to them. But, that fact created such diversity that most equipment could only communicate with other equipment of the same type. It was clear that some rules had to be established so that every equipment manufacturer would, at least, have to package and handle communications consistently. All that remained was to create standards.

The Coming of Standards

The period following World War II was one of immense optimism. Despite the fact that the Cold War had begun even before the ashes of the hot war had cooled, man had begun to dream of international cooperation. Those who dreamed that dream realized that communications was a key to bringing the dream to reality. It was necessary to provide telecommunications standards that would create the compatibility necessary to provide end-to-end communications between international networks without regard to the nation of origin. Organizations devoted to the creation of such standards began to emerge.

CCITT - Consultative Committee On International Telephone and Telegraph

The task of establishing standards was undertaken by the International Telecommunications Union, which is a United Nations Treaty organization. Using representatives of member governments versions of our own FCC, along with input from numerous large companies, the ITU assembled a group known as the CCITT (Consultative Committee On International Telephone and Telegraph). The result of their deliberations was Common Channel Interoffice Signalling System #6, first introduced in the 60s. Were there five previous versions? Of course. But only CCIOSS 6 (sometimes called SS6, but only in hindsight) survived to deployment.

Recently, the CCITT group name has been changed to the Telecommunication Standardization Sector (TS) and groups responsible for radiocommunications (RS) and telecommunications development (D) have been added. Since it is the TS group that is largely responsible for developing SS7 standards, the use of the term CCITT is rapidly being replaced by ITU-TS.

ANSI - American National Standards Institute


This group bears the responsibility for approving standards for use in the United States. The ANSI Accredited Standards Committee T1 is the ANSI committee (one of many) that is directly responsible for telecommunications standards. This committee in turn is composed of seven subcommittees, each of which bears the responsibility for a specific area of telecommunications standards. New standards are proposed by any of several accredited standards bodies including Bellcore and EIA (Electronic Industries Association).

While there are other standards for the SS7, these two (CCITT and ANSI) , along with the Japanese standard, cover most of the land based and wireline standards throughout the world. By the way, these are not really “different” standards. The basic approach is universal. It is simply that many nations decided it was desirable to modify the CCITT (ITU-TS) standards in a number of details to best serve their own needs. This is where the American National Standards Institute got involved because of its responsibilities for U.S. standards. Around the world other standards groups became involved in their own countries. The result is the existence of over 30 National Variants of the ITU-TS standard.

Nevertheless, don’t think of them as “different” standards, but rather as different “dialects” of the same standard.

A Final Word About the PSTN

It hasn't been our intention to thoroughly explore the Public Switched Telephone Network. It is a topic with a million different facets that could easily take a lifetime of study. Worse yet, it would likely not be possible to learn about every technological advance, even as they occur. Learn what has happened today, and, tomorrow all you will know is yesterday's news.



Our purpose, instead, has been twofold. First, a broad understanding of the PSTN serves, as no other example does, to illuminate the concept of network. Secondly, the fact that modern technologies can become universally available so quickly is due almost solely to the existence of the PSTN. Without the PSTN there is no Internet. Without the PSTN there is no Mobile Network. And, without the PSTN there is no Signalling System #7.

With the network concepts gained by this exposure to the PSTN, we are now ready to begin our examination of the network (SS7) designed and implemented to offer services to, and in turn receive services from, the PSTN. In all of communications there is no greater symbiotic relationship than that which exists between the PSTN and the SS7.

Section 2

SS7

Network
Architecture

SS7 Network Architecture

You've seen what a network is, why a network like the SS7 was needed, and where and how the standards were developed. Now it's time to look at the SS7 network itself. The easiest way to begin is to examine network architecture to understand the physical elements of the network and to appreciate their reason for being.


What would you say is the key element in the PSTN (Public Switched Telephone Network)? While there are a number of key elements, it really is the switching location that makes it a network. Switches are the "glue" that holds the PSTN together.

The STP (Signalling Transfer Point)

The SS7 is held together by a digital sister of the switch known as a Signalling Transfer Point. The requirements of voice switching and digital transfer are different, but they resemble each other in some ways. The PSTN requires circuit connections of voice lines.

There is no need for connection in the SS7 network. What is referred to as "circuits" in the PSTN can not carry messages until the switch makes a physical connection. Instead of circuits, the SS7 makes use of transmission lines called links. In concept, at least, these links always exist and are always available to carry messages. Instead of "connecting," the STP needs only to direct messages to the links which it selects as most appropriate to deliver the message. For example, if an STP has links heading off toward the four compass points, it might be more "appropriate" to direct a message addressed to California to a west-leading link than to an east-leading link.

Of course, geography is not always the basis on which the STP will decide where to transfer incoming messages. Other factors (such as least cost) have an impact on STP routing decisions just as they do in the PSTN. Nevertheless, the STP does not connect to these routes (since it already is connected); instead, it transfers messages to the selected route. You may see STPs illustrated in either of the two ways shown here.



Graphic Representations of STPs

Simply SS7

Both of these representations have the same intent. Since STPs are always “paired” by link connections they are always shown in pairs. The drawing on the left is more explicit because it indicates the links connecting the two. For our purposes in this text, we will illustrate STPs as shown on the left.

The reason for the pairing of STPs is redundancy. If one of the pair should be lost for any reason its “partner” STP will handle the load. For this reason, certain geographical considerations are made for any new pair to be added to the network. For example, if a company would like to establish an STP at the northern end of the San Andreas Fault and the second member of the pair at the southern end of the San Andreas Fault, network administration would be likely to turn down the request. The reason is obvious.

The Engineering design must also take into account the reason for the pairing. The pairing is designed to guarantee the continued flow of traffic even in the event of the loss of one STP. Because of this need we say that STPs are engineered to 40%. That means that under normal circumstances, the facilities required at each STP (links, cards, processors, etc.) will handle their share of the total traffic directed at the STP pair while operating at no more than 40% of capacity. Thus, if the other STP is lost, its traffic can be transferred to the remaining STP, which will now be operating at 80% of capacity (40% times 2). The formula for determining the required traffic capacity of an STP looks like this:

$$\text{Required Capacity per STP} = (.50 \times \text{Total traffic directed at pair})/.40$$


For example, if the total traffic was 1000 messages per second, each STP would be given the resources to handle 1250 messages per second. The reason for the excess capacity is to ensure that messages can be handled even when “bursts” of traffic exceeding 1000 m/s occur.

Later you will see that the same 40% engineering is recommended to establish resource requirements per link for those operating on the network with a pair of links. Links are used as the connection to the network as well as for connection between STPs. Links are named using names which suggest which two network entities are being connected. Thus, when a service provider gains access to the network by establishing links with an STP pair, the links are called “Access Links” or simply “A” links.


Network Links

In this section we’ll examine and describe all of the different types of links found in the network . Before we do there is one thing of which you should take note. From time to time we’ll use the word “type” in connection with links. It may sound as if we are referring to the kinds of transmission lines being employed. That is not the case. A variety of transmission lines are in use. The SS7 standards have no concern for the nature of the physical facilities. From a practical standpoint network entities simply take advantage of the available infrastructure, coupled with concerns about cost, reliability, etc.

We’ll look at links one type at a time, building a network as we go.




Since it's always best to begin from someplace you've already been, we'll start with our earlier drawing of paired STPs. The links that connect the pair allow messages to cross over from one to the other. They are, therefore, named "**Cross Links**" or simply "**C**" links.



The circle in this drawing is another network entity. "SP" stands, simply, for Signalling Point. That name is a very broad generic title used to describe locations on the network that receive and/or send signals. This term is often used when it is not important to describe the way in which the location functions. Another broad generic term often used is "Signalling End Point" (SEP). Later we'll describe terms used to define specifically what type of activity takes place at such a "node." But, for now, we are only concerned with links.

Because the Signalling Point in the drawing above gains access to the Network using its links to the STPs, the links are called "**Access Links**" or simply "**A**" links.

Simply SS7





Just as telephones generally tie into local switches in the PSTN (Public Switched Telephone Network), network entities (or nodes) like this SP generally link to a local STP (Signalling Transfer Point). An STP, along with the nodes that connect to it, forms its own little local network. When that local network connects a bridge of links to another local network, the links are connected in every conceivable way.

“Every conceivable way” in this case means that four links are used in an arrangement that links each STP to each of the other STPs. Since the links between these local STP pairs form a bridge over which messages can be transferred from one local network to another (as well as to the network at large), these links are called “**Bridge Links**” or simply “**B**” links. The Bridge Link label applies when connecting two local pairs or when connecting two regional pairs. A different label is applied when the STP pairs are at different hierarchical levels (such as local/regional).


In the Public Switched Telephone Network there is a definite switch hierarchy. Some switches exist to connect lines while others exist to connect switches. In the SS7 network, there is often no such clear hierarchy. Nevertheless, the standards envisioned that connecting two STP pairs with essentially the same functions (like offering access for Service Providers) was slightly different than connecting this “local” STP pair to an STP pair which served a broader network.

Whether or not such a clear hierarchy really exists, in the next drawings we’ll examine links from one level of the hierarchy to another. The drawing below will help to describe the graphic representation of hierarchy. It’s really quite simple. STPs drawn at higher positions in the drawing are (presumably) at higher network levels.






An STP pair will always find a routing advantage in linking with a “higher” level STP pair. The linking arrangement is identical to the “B” link arrangement except that, as the drawing indicates, the links become *diagonal*. These links are called “**Diagonal Links**” or simply “**D**” links.



A Signalling Point may try to improve on its routing flexibility by extending its links to an STP in a remote part of the network. The links may look like A links, but because they extend the routing capabilities of the SP (and also because they need to be extended to reach farther) they are referred to as “**Extended Links**” or simply “**E**” links.

Before we proceed to the last type of link, there is one fact of which we should take note. Often looking at drawings such as these leaves the impression that every linked location is located in a different geographical location from every other linked location. In the case of paired STPs this is usually true because of the need to limit the risk to the pair.

In many other instances, however, it is not true. A Signalling Point of any type may physically reside in the same room with an STP. Indeed they may exist in the same cabinet. In some instances Signalling Points of more than one type may reside in the same physical location. Each functional entity is defined by a network address known as a Signalling Point Code (SPC). In a way it’s like having a telephone and a FAX machine on a desk. Their physical location is the same, but their addresses (telephone numbers in the PSTN) are different.



To show the final type of link, we have added a new Signalling Point into our little network. Generally the node would have its own access into the network using “A” link connections to a local STP pair. Our drawing shows the SPs using different STP pairs for access, but this may not be the case. In this instance these two nodes have some kind of a business association. They may be two nodes owned by the same company. At any rate, they need to share data; and, for whatever reasons, they prefer not to send that data through the network.

Instead, they assign a link connection for the specific purpose of fully isolating this communication from the network. Since these links connect locations which are fully associated, they are named “**Fully Associated Links**” or simply “**F**” links.

Now that we have seen all the links, and, while you have a single drawing showing all of them right in front of you, it might be a good time to review. Here is a list of each with a short description.

Access Links	Link a node (Signalling Point) to a local STP pair.
Bridge Links	Link two pairs at the same level (local/local, regional/regional).
Cross Links	Link two STPs together to form an STP pair.
Diagonal Links	Link a local STP pair to a higher level STP pair.
Extended Links	Link a node (Signalling Point) to a remote STP pair.
Fully Associated Links	Link two associated nodes together.

If you still could use a little help to remember the links, you might want to try this little poem:

Links **B** and **C** and **D**, connect the **STPs**
While **F** and **A** and **E**, link services you see.

SS7 Network Nodes


In our discussions so far we have used the generic “SP” to describe Signalling Points. With the architecture of the network in mind, it’s time to take a look at some specific signalling point types. Since the SS7 Network developed originally with the idea of improving the efficiency of the PSTN (Public Switched Telephone Network), we’ll start where the PSTN meets the SS7.

The SSP (Service Switching Point)

There are actually two types of Signalling nodes that are switch associated. The first type is called a **CCSSO** (Common Channel Signalling Switching Office). These are end or tandem offices which have the capability to use the SS7 in what is referred to as a trunk signalling mode for call set-up.

The second type (and the name you’ll hear most often) is the Service Switching Point (**SSP**). Like the CCSSO this switch can handle call set-up. Unlike the CCSSO, the **SSP** also has the ability to stop call processing, make queries of even unknown databases, and perform actions appropriate to the response. The greatest difference between the two lies with the fact that the **SSP** is equipped with whatever software is required to handle numerous feature capabilities. In a way the **CCSSO** is a more limited version of the **SSP**.

The following drawing illustrates the SSP in the network(s).



Simply SS7

The SCP (Service Control Point)

One of the first purely digital uses for the SS7 network was to provide a response to the need to translate from one form of data to another. For example, switches need to maintain tables to translate dialed digits into routing information consistent with the North American Numbering Plan (NANP). It is that plan that breaks North America down into area codes, exchanges, and finally to the lines serving individual telephones.

When telephone companies thought up a way to have customers pay for incoming calls as well as outgoing calls, the North American Numbering Plan couldn't be used to let the switch know how to route the call. Where normal numbers had an area code, the new numbers had no information at all. Where normal numbers had an exchange code, the new numbers used 800 no matter where the circuit was to terminate. What is more, people soon wanted to choose numbers such as 1-800-FLY-AWAY. Geographical significance was totally lost in such numbering.

Lost, that is, to everyone but the SS7 Network. An 800 number was still assigned to a business by activating a telephone line in an exchange, in an area code. To the switch it was just another normal number. But, from the 800 number that was dialed, the switch didn't know what the normal number was. Initially there were few 800 numbers, and they were assigned locally. That meant that local switches could maintain translation tables. But, it wasn't long before FCC rulings and 800 number demand made it impossible for local switches to store all the data, much less deal with the daily additions and deletions.

It was obvious that the only way to maintain this massive amount of data was to place it at centralized locations in the SS7 Network. For easy access and manipulation, databases were created for the storage. Now a switch which was asked to connect an 800 number could first query the database for the NANP (North American Numbering Plan) values. The switch could then route the call normally.


With the advent of the database, the switch only needed to know where to send the query. And, thus, the SCP (Service Control Point) enters the picture. The SCP acts as the "front end" of the database. It may or may not be located in the same location as the database. The important thing is that it can handle the query and send the answer back to the switch.

Toll free numbers, however, were to become far from the only database use in the network. When you use a credit card to make a call you may be entirely unknown to the local telephone company. Before the switch attempts to connect you, it needs to send a query to an SCP to locate your account in a database and verify that it is a valid (and viable) account. In the process it also learns how and where to send the bill.

In today's network you will find a database wherever a translation, verification, or simply information is required. At the doorway to that database you will find a Service Control Point. This is the node that provides the mechanisms for data to be retrieved from the database in a form that is suitable to the purposes of the node initiating the query. Since the types of services that can be offered are limited only by imagination and available data, it is likely that SCPs will continue to play a significant role in the growth of the SS7 Network.

In general, the major databases (like the 800 database) have been centralized in the network. That is not to say that a single such database exists; but, rather, that several identical databases exist throughout the network. Obviously each of these databases should contain the same information. Since new numbers are being granted every day, the task of updating is daunting. The best answer thus far is to update a database on a regular (such as daily) basis.

It may sound strange, but daily updates are not frequent enough for some companies. We'll see why in the next node description. Before we leave the SCP, the drawing below illustrates an SCP in the network.



The CRP (Customer Routing Point)

One of the primary uses of an SCP (Service Control Point) is to provide routing information for switches when that routing information is not directly provided for the switch through the digits that have been dialed. The Toll Free Numbers databases are a good example of this. Dialed 800 and 888 numbers end up being translated at a database before a voice switch can do anything further with the call. Another example is the 900 number database. These calls, of course, are not Toll Free (in fact, they cost the caller much more than the cost of a normal call to the same location). A dialed 900 number is just as unintelligible to the switch as an 800 number is.

Centralizing databases was a big step towards perfecting the efficiency of the 800 number system. Nevertheless, the act of centralizing databases created a problem or two, even while it was resolving others. Perhaps the largest of these new problems was the update requirement. As you might imagine such databases are huge. The number of additions, deletions and modifications that are necessary every day can be staggering. Continual updating would be difficult to handle. What's more, who would want to remove a database from service several times a day?

The best answer to the need for updating is to collect all the updating information received each day and pick a particularly light traffic period for the updating. During the updating, traffic can be shifted to another database. Following the updating the updated database can go back on line while the backup database gets updated. From the switch viewpoint the update is seamless. How could anyone have a problem with that?

The fact of the matter is that some do have a problem with that. Let us, for example, look at some of the concerns of a large retailer who maintains numerous order lines to handle orders by phone. This company only has three warehouses in the U.S. Orders coming from anywhere in the nation are filled at these warehouses. However, orders are processed at over 100 locations scattered around the country. The major factor in this distribution is cost. The company uses a single nationally advertised number to reach all of these locations. Routing to the specified order line is determined by the area code of the calling customer. Thus, the company is assured that a customer calling from New Jersey will result in a phone bill determined under local rates. A California facility will not be tallying up a long distance telephone bill when the call originates in New Jersey.

Normally, when this 800 number is dialed, the SSP (remember ?.....Service Switching Point) sends a query to a centralized 800 database and receives the routing information necessary to connect the call. The SCP (remember ?.....Service Control Point) correlates both the dialed digits and the caller's area code to come up with one of numerous normal telephone numbers listed in the database for this particular 800 number.


The company had a bitter experience some years ago. Their order lines were in a building in Charlotte NC which burned to the ground. The Telephone company was informed but there was no way to update the centralized database immediately. At the height of the company's busy season prospective customers in the area were unable to place orders.

The answer to their problem lay in the creation of their own routing database. When a switch makes a query to obtain switch routing information based on the dialed digits of 800 numbers assigned to this company, the query is routed to the database the company maintains for itself. Since this phone customer establishes a signalling point for the purpose of providing its own routing information, the node is called a Customer Routing Point (CRP).

The company operating the CRP has full control of the routing information being returned to the switch. When it becomes necessary to change the routing (as in the case of the burned out location), the company simply updates its database. The minute the company becomes aware of an inaccessible location, the changes can be made, and the very next call will be answered at a new location.

If you think about it you'll realize that the advantages are numerous. For example, if the incoming call load becomes excessive at one location, they can quickly be routed to others. The CRP operator knows how many queries, and thus how many phone calls, are being made to each location. The traffic statistics gathered at the CRP become an invaluable resource in numerous business decisions. And, finally, the ability to control call routing allows the company to make highly efficient use of their resources (number of phones, number of agents, etc.) at all servicing locations .

The drawing below illustrates the CRP in the network.



The IP (Intelligent Peripheral)

The SCP software programming (usually referred to as the Application or the Process) allows it to deal with the requests contained in the messages sent to it through the SS7 network. Often these messages are in the form of a query sent from a switch. The SCP is capable of interpreting the query, extracting the required data from a database, and returning it in a form understood by the switch software.

The same functionality allows the SCP to access and make available services of another kind at a different type of signalling point in the network. Sometimes this entails invoking features for which the switch is not equipped. At other times it entails utilizing an Intelligent Peripheral.

In general, the Intelligent Peripheral is home to a Process which can deal with the requests made of it through the SCP by providing the services of a variety of devices. If you are unfamiliar with the term “devices,” think of it simply as equipment.

For example, a caller may wish to make use of a service that allows the dialing of frequently used numbers simply by saying the name of the party they wish to call. If you haven’t used such a device yourself, you may have seen the commercial in which the caller gets connected by speaking into the telephone and saying something like, “Call Mom.” The call gets placed correctly because of equipment which can analyze the sound of the words and correlate it to an entry in a database. The database provides the normal telephone number which is then returned to the switch to which the caller is connected. The switch can now route the call exactly as it would have if the caller had dialed Mom’s phone number.

Providing such a voice translating device at every local switch in the world would cost far more than most phone companies could afford. The cost of the service, therefore, would be far beyond the reach of most of us. Worse yet, when someone came up with an improvement in the service, every switch in the world would have to be reprogrammed before the improvement became universally available.

For these reasons, it makes much more sense to keep the equipment (and the software necessary to make it work) at a limited number of locations around the network. Then, all the local switch needs to know is where to send the message to obtain a phone number translation of the vocal request.

The Intelligent Peripheral provides the services of many such devices. The caller may make the request vocally, by touching the digits on the phone (which most people know as “touch tone”), or even by using keyboard entries on a computer. In the future your telephone will likely be equipped with, at least, a limited version of the computer keyboard. The devices available at the IP will be able to receive and translate all of these different ways of making a service request.


Upon receiving the request sent to it by the SCP, the Intelligent Peripheral will make use of such devices which interpret “touch tone” input, voice input, computer input, or terminal input. Then, it will respond by returning data to the switch or terminal, or by sending a voice announcement to the telephone of the caller.

Some of the services provided by the IP may be services provided directly to the switch rather than to a telephone subscriber. In this way even older, less capable switching equipment may be able to respond to and/or send signalling it isn't equipped to understand. And, the switch would be able to respond to requests for services offered by new technologies simply by subscribing to the services of an IP where those technologies can be found.

Service requests are sent through the SS7. Therefore, the IP requires links to the SS7. On the other hand, the services provided (various signalling types, etc.) are related to switches and must be delivered in the voice network. This means that the IP must also have trunk connections into the PSTN.

Thanks to the Intelligent Peripheral, new services need only be installed at a limited number of centralized locations to become universally available. Service improvements which require only some reprogramming can be done quickly. Once done, such upgrades can be made available throughout the entire SS7 network.



One point should be made here. A company may equip and deploy any type of node solely for the use of its own customers, its own switches, or its own SS7 nodes. Or, it may equip and deploy for the purpose of selling such services to other companies. How "universally" available any such service becomes is, ultimately, a business decision. The Intelligent Peripheral is shown below in the network.



The Wireless Network

Judging from the title of this section it might appear that we have suddenly jumped from the discussion of the SS7 to an entirely different network. The truth is that the wireless network is truly “wireless” in only a small portion of its architecture. A mobile telephone transmits and receives to and from a transmitter/receiver. The transmitter/receiver is the first and last place in the network where the communications are wireless. The voice received from the mobile telephone must be connected into the PSTN. The voice received from the PSTN must end up connected to the transmitter for transmission to the telephone.

The wireless company needs a switch to connect a call originating or terminating at their customer in the same way that the landline phone company needs an end office switch. An easy way to illustrate this is to show it side by side with an earlier drawing we used to illustrate the End Office part of the Public Switched Telephone Network (PSTN).




Hopefully, it should be easy to see the comparison. The wires are gone. In their place is a two way radio transmission (indicated by the symbol that looks like a lightning bolt). Both the transmitter/receiver and the mobile telephones are radio transmitters. They are also both radio receivers.

From the end office, the calls have to be connected to a PSTN access switch. Multiple end offices are connected to such a switch. In the same way, the transmitter/receiver of the wireless network must be connected to a switch; and one such switch will serve to connect multiple transmitter/receiver locations. That switch needs to have trunk connections to the PSTN as well as links to the SS7.

The MSC (Mobile Switching Center)

The MSC needs “one foot” in the PSTN (Public Switched Telephone Network) and “one foot” in the SS7 (Signalling System #7) network. Landline connections must be made to connect the callers. Information pertaining to the call must be sent to the PSTN switches using the services of the SS7. And, finally, the MSC faces problems which switches in the PSTN don’t face. Every end office in the PSTN can find its customers at the other end of the wires attached to its own subscriber interface. At any given moment the MSC may not have a clue as to where its subscriber is located.

The MSC then has requirements that PSTN switches do not have. Among these is the need to keep track of its subscribers. When the subscriber is within the range of any of the MSC’s transmitter/receivers, the MSC receives an indication of the signal strength of the subscriber’s telephone at the location of the antennas.



Simply SS7

The MSC connects the antennae receiving the strongest signal. If the subscriber is “on the move” during the conversation, the signal will gradually weaken at one antennae while strengthening at another. At some point, the signal will be stronger at an antennae from which the call is *not* connected, than at the antennae from which the call *is* connected. Once the MSC determines this, a new connection is established and the call is “handed” off from one transmitter/receiver to another. “Hand offs” are not instantaneous once the signal strength balance seems to favor a new location. The MSC is smart enough to know that transient conditions (such as passing under a bridge) will have an effect on this balance.

New Standards (GSM and IS41)

The area served by a mobile provider is divided up by the ability of each antenna to receive a signal strong enough to connect. These areas are known as “cells” (like the cells in a honeycomb), and the name leads to the term “cellular telephones” or “cell phones.” Generally, an antennae is located at or near the center of each cell.

In the early days of mobile personal communications, what we have just described might well have been called “the wireless network” because those who provided such mobile communications maintained their own small network in isolation from all others. Mobile communications providers prospered in metropolitan areas where the need for such service was greatest. As a result, early subscribers found that their mobile phones were useless once they had driven far enough beyond the boundaries of the metropolitan area. The response of the providers was to create ever more distant cells. They did this by building new transmitters or by leasing the services of existing antennae that had been built for other purposes.

Inevitably, of course, the cells of a provider in one area began to overlap those of a provider from another area. That overlapping would prove to be both a boon to the subscriber and a problem for the provider.

The problem was that when a subscriber of one company’s service “roamed” into the cells of another, the new company had no idea who the “roamer” was. There was no mechanism in place that would allow a call to be “handed off” to the new provider; and the new provider had no way of gaining assurances that the call would be paid for. Astute providers knew there was money to be made by connecting “visitor” calls. As a result, many companies set out on their own to establish deals with adjoining companies. The results didn’t look very promising. In particular, the need to request the service ahead of time and the need to validate the customer resulted in having the phone call proceed in anything but a “seamless” fashion. It was time for more standards.

Europeans were among the first to recognize this problem, as well as the special needs of the MSC to handle its broadcasts and receptions. One result was a French effort to provide standards as put forth by the standards group, Groupe Service Mobile (GSM). Another came from the standards organization called Electronic Industry Association/Telecommunications Industry Association (EIA/TIA) and was called IS-41 (Interim Standard 41). There are others, some of which describe only the communication between an MSC and a Base Station system regardless of the manufacturer of either. Base Station System is a term used to describe the system used to manage the radio frequency resources of a number of Transceivers (transmitter/receivers).

Before we leave the subject of wireless standards we'll do our best to eliminate a common misconception. The wireless standards are not separate and distinct standards from the earlier SS7 standards. Those involved in the creation of the SS7 standards could not have foreseen all the unique requirements of mobile telephony. When it was seen that there were functional requirements and messaging requirements not covered in the original standards, other standards were written. But these standards still make use of SS7 functionality and, in practice, are used "on top" of the SS7 standards. They work, therefore, as adjuncts or "add ons" to the SS7 standards. Our discussions of SS7 functionality to be found in a later section will help you to understand this.

HLR (Home Location Register)

As you have seen, the MSC (Mobile Switching Center) is, itself, a node on the SS7 network. In the nature of its communications, it is a very busy node. It needs to communicate connection requests into the PSTN. It needs to communicate specialized information to its own Base Station System. And, finally, it needs to maintain its own subscriber database and to query the subscriber database of other providers.

Roaming was one of the first problems encountered in mobile networks that led to the development of new mobile standards (**GSM** and **IS41**). Issues of how landline connections were to be made were much more complex than those experienced in the PSTN. For example, when the roamer leaves area A and enters area B, the new area establishes a new connection. Then when the caller roams back into area A, should area A also establish a new connection? If so, you can end up with multiple criss-crossed circuits (referred to as "shoelacing")

When a subscriber to the services of one company "roams" into the cellular network of another, the provider uses the visitor's identification (Mobile Identification Number or MIN) to query the roamer's own provider for information on the subscriber. This "Home Location Register" (**HLR**) is established by each provider as a part of its own record keeping. It contains all the information necessary to validate the customer and to provide the necessary billing information. The provider who makes the query uses the data to fill in a database that will be used to correlate the charges for the call and provide the information for billing the home provider. The home provider then adds the call (based on its own charging structure) to the customer's bill, and, in turn, pays the servicing provider for the service.

With the data garnered from the Home Location Register, the new service provider is able to make an entry into a database it maintains. This database is used to help correlate this roamer with this call. We'll examine that database next.

VLR (Visitor Location Register)

Surprise! For the first time in reading this booklet you knew the definition of a network node before you even knew its name! This is the name given to the database used by mobile telephone service providers to store information on roamers (people traveling outside the area served by their own providers).

Using such techniques as subscriber databases along with the new standards, mobile providers have been able to provide services that are increasingly more “seamless”. That is, the telephone user is less and less aware that he/she has traveled from one area into another. In many areas the delays and interruptions which once distracted cell phone users are a thing of the past. The subscriber simply drives from place to place with uninterrupted, even service through entire conversations. New technologies currently emerging will bring even greater improvements.

Intelligent Network (IN) vs. Advanced Intelligent Network (AIN)

We are nearing the end of our survey of SS7 architecture. Before we go on to other things, we’ll try to clarify your understanding of two terms which you are likely to hear quite often. “Intelligence” is a term often applied to the use of computer programs (software). If we equate this with the same term as applied to human beings we will certainly miss the point. No computer program to date even begins to compete with human intelligence.

Nevertheless, a switch that can be programmed to react in numerous ways to numerous requests is certainly much more “intelligent” than a switch without such capabilities. Programmability is the key. At an earlier time no machine could perform new tasks without modifications to the machine. As often as not, offering a new service meant throwing out the old machine and replacing it with a new one. That applied to switches as much as to any other kind of machine.

The “Intelligent Network” really began when the first programmable switch was installed in the mid 60s. For the first time, the way the switch made connections could be controlled, and even modified, simply by modifying a program. Still, the network itself wasn’t intelligent until the introduction of databases and the Service Control Points which gave access to them. That came with the introduction of the centralized 800 database in the early ‘80s.


By 1988 AT&T began to see the need to standardize such things as the way features would be invoked. The problem evolved through the fact that, as the Intelligent Network grew, each vendor had begun to take his own path. As the paths in the evolution of service deployment began to diverge, AT&T approached vendors with a request to join in the development of a standard. The standard was not to define new services (although several have evolved as a result of the standard), but rather to provide the means whereby a customer could “serve up” such services from his own premises, or how the telephone network could provide the services to its own switches.

One result was the development of a Service Creation Element (SCE). The concept envisions the use of a Graphic User Interface (GUI) with “drag and drop” icons to be used for customizing services. Currently two releases of the AIN standard have occurred (AIN 0.1 and AIN 0.2). The AIN 0.2 version has been deployed on a limited basis and is undergoing testing.

Advanced Intelligent Network Architecture

The arrival of the Advanced Intelligent Network has little effect on network architecture. The reason is that the purpose of the AIN is not the creation of new services nor the redesign of the SS7 network. It is, instead, an attempt to standardize and define the best ways that new services might be developed and deployed.

One node that appeared in the Intelligent Network assumes a greater role in the Advanced Intelligent Network. This is the **SMS** (Service Management System) which provides a human/database interface. Currently employed in the Intelligent Network, the SMS utilizes a man-machine interface and command line language for service building. AIN approaches to service building will increase the importance of the SMS to the network. The drawing below illustrates AIN architecture.




Simply SS7

One Final Look

With our discussion of SS7 network architecture complete, all that remains is to illustrate the network in all of its major aspects and to review the nodes which have been discussed. The drawing below illustrates the SS7 network and its interface to the PSTN and Wireless networks.

Directly opposite the drawing is a listing of the nodes along with a brief description of how they operate within the network. Once you are comfortable with your understanding of this architecture, you'll be ready to move on to gain an understanding of how messages are sent and how they move through the network.



STP (Signal Transfer Point)

The “knots” that hold the network together. These nodes serve to provide network access to other nodes (by connection with Access Links). STPs transfer messages around the network. STPs maintain routing tables for the purposes of directing messages to their intended destinations.

SSP (Service Switching Point)

The Service Switching Point is a switch associated node which handles call set-up and has the ability to stop call processing, make queries of even unknown databases, and perform actions appropriate to the response. In general, the SS7 messages which originate or terminate here are either circuit or call routing related.

SCP (Service Control Point)

In general, Service Control Points provide access to databases. These nodes are the residences of processes which can access the database, extract the required data and return it to the node requesting the data. The database(s) to which the SCP has access may or may not reside at the same location as the SCP. The same capabilities that allow the SCP to access databases lend themselves to other uses such as providing access to an IP.

IP (Intelligent Peripheral)

The IP is the residence of processes which manage resources such as signalling sensors and voice response equipment. The resource management capabilities become available to switches on demand, thereby freeing switch locations from the need to equip with a myriad of such devices, and providing highly efficient use of both aging and up-to-date technologies.

CRP (Customer Routing Point)

The CRP provides on-premises control of the routing information requested by switches for translation of 800 type dialing (not limited to 800 numbers). The operator of the CRP is a customer who requires rapid update and control of the translation of their own numbers.

MSC (Mobile Switching Center)


The Mobile Switching Center maintains control over its own Transceiver network. Part of this control includes tracking subscribers and performing “hand offs.” The MSC also provides the landline connections into the PSTN to complete the connection of subscriber calls. Finally, the MSC makes use of the SS7 network to convey circuit related information to the PSTN and to communicate with the service providers of “roamers.”

HLR/VLR (Home Location Register/Visitor Location Register)

A database that contains customer information about local subscribers is maintained by each provider. This is the Home Location Register. Another company will access this information when a “roamer” appears, and use the data for an entry into its Visitor Location Register.

A Final Word About SS7 Architecture

Once again we'll complete a section with a word about the value of what you have just learned. It was not our intention to explore every kind of SS7 node in detail. There are other database nodes. There are other device resource nodes. There are other programming resource nodes. And, there are nodes which combine several types of services.



Our purpose, this time, has been to expose you to numerous different node concepts and help you to understand the nature of the services these nodes provide. If we have succeeded, you now have the basis for understanding any type of node to which you may be exposed in the future.

Now we can turn our attention to the tasks that hardware and software must perform to ensure that all these services can be accessed and provided with both reliability and efficiency.

Section 3

SS7

User Part

Functionality

SS7 User Part Concepts

Any discussion of functionality in the SS7 network must begin with an understanding of the basics of network communications. When human beings communicate, the conversation can be uni-directional or bi-directional. Generally, when we watch television, it is uni-directional because we rarely talk to the TV. When we do, the TV usually pays no attention.

A conversation between two people is bi-directional (if not, we don't call it a conversation.) One person listens while the other one talks, then the roles switch and the talker becomes the listener. In the SS7 network the talker/listener is a program. The program (or process) involved in the conversation is one designed to manage the services of the node.

Thus, when a voice switch location (See SSP, or Service Switching Point) receives an 800 number dialed by a subscriber, the Program-In-Charge recognizes that the switch cannot route the call and that an 800 number translation is required. As a result, it sends a request (query) to the Program-In-Charge at an SCP (Service Control Point). The job of that program is to find a normal telephone number that corresponds to the dialed 800 number, retrieve the translation, and return it to the SSP with all the necessary data, and in a form that the SSP can understand.

Each different kind of node has a different kind of Program-In-Charge designed to manage the services required of that node. These Programs-In-Charge are called "Applications." In SS7 network communications, Applications fill the role of the talkers and listeners. All communications occur between Applications.

Every node requires the services of an Application; and, the Applications are designed to perform those functions which are unique to the node. For example, locating, packaging and returning 800 database translations are requirements that are unique to an SCP.

While each node has its unique requirements, there are many requirements that are common to all nodes. For example, every node must have the ability to reliably send and receive data over links to and from the nodes at the other end of the link. Nodes directly connected to each other by links are called "adjacent" nodes. Since the transfer of messages between adjacent nodes is a common requirement of all nodes, it makes sense to create a single program to handle all of the requirements of that task. Such a program could then be provided to every node and creating the Application would be that much simpler because the Application would not need to do anything about sending messages over the links.

All nodes that deal with the connection of voice circuits likewise have common needs. In this case, each switch needs to communicate its circuit requirements, to tell the next switch who the caller is and who is being called, to communicate the information about the availability of the called party (telephone ringing?, telephone busy?), and a great deal of other call related information. Since most switches deal with these same requirements, it once again makes sense to create a single program to handle all of the requirements of the task. This program could then be supplied to every switch location in the networks.


Simply SS7

If we examine all the needs of each type of application, it soon becomes obvious that the same type of modular program can be written to serve each of the many network uses for the nodes. The modular approach results in having numerous small programs that can be used in various combinations to perform all of the repetitive requirements, leaving the Application free to perform only those things which make the service it offers unique.

Designing these programs in such a modular fashion leads to numerous efficiencies. For one, things must occur in a specific sequence. For example, if a message loses something in transmission, the Application will not be able to respond to it. Therefore, it makes sense to determine if this is the case (and, to do something about it) as soon as the message is received. Otherwise, time will be lost when the Application attempts to handle a message, only to find it cannot do so.

The modular approach leads to the creation of what is referred to as a “layered” protocol. Protocol means nothing more than a rigid set of rules which determine how communication should be handled. It covers everything from what should occur to when and how it should occur. It also prescribes exactly what the message consists of when it is sent over the links. “Layered” means each module performs its function in sequence and then hands the message off to the next module (which is “above” for incoming messages and “below” for outgoing messages).


Each of the functional program modules is termed a “user part.” The rules (protocol) dictate the sequence in which things must be done. To show this graphically, a convention has been adopted for the drawing. In this drawing, the functional modules that deal with a message just about to be transmitted over the links (or one just received from the links) are shown at the bottom. Other modules are shown “stacked” above in the sequence in which their functions are performed. The resulting picture is commonly called a “stack.” A typical SS7 stack is shown below.



Application To Application Communication

This drawing illustrates how messages are sent and received by applications using sequential layers of the layered protocol. Each layer performs several functions before passing the message to the next layer. All Applications make use of the Message Transfer Part (**MTP**).

The Applications shown here make use of the modules (user parts) labeled Signalling Connection Control Part (**SCCP**) and Transactions Capabilities Application Part (**TCAP**). These two user parts are employed when the Application deals with Database query and response.




Note that, in this drawing, the lines indicating message flow do not pass through the Integrated Services Digital Network User Part (ISUP). That part provides functions for the handling of telephone call related messaging which is sent from switch to switch. The communication shown, therefore, does not involve switching in a voice network.

Other User Parts

Another major advantage of a layered protocol lies in the efficiency that can be realized when a new user part is added. As long as a user part is designed to use the services of the existing lower layers, a new part can simply be added to the stack and thereby become available to the application.

To illustrate this, we have added a user part to the previous “stack” drawing to arrive at the drawing shown below. The part added is not a “new” part. Instead, it is a part formerly employed largely in Europe to help implement switch connections. It is called the **Telephone Users Part (TUP)**, and it has been used in Europe in ways resembling the use of the Integrated Services Digital Network in the U.S. Nevertheless, the use of ISDNUP (or at least the circuit control functionalities commonly called ISUP) is replacing TUP worldwide. Some holdouts (such as China) still use TUP.



There are other user parts as well, but we will concentrate on the parts in common use in the U.S.

SS7 User Part Functionality

Taking a deeper look at the specifics of SS7 functionality brings us right back to some of the earliest work done by the CCITT (Consultative Committee On International Telephone and Telegraph). The early work done by any standards group often consists of compiling and organizing a very long list of questions that start with “What if?” and “How?”. For example, “What if the message received was garbled in transmission?”, and “How would you detect the fact that the message was garbled” ?

Providing answers to these questions results in a long list of actions that need to be taken to guarantee that the transfer of messages is accurate and reliable, that messages can be accurately guided to their final destinations, that resources can be managed effectively, and that the Applications are delivered sufficient data of the correct types to allow them to perform their tasks.

MTP Level 1

Finally, we can begin to discuss the required functions of the user parts and to look at how software handles those functions. Unfortunately, the first level we’ll examine doesn’t lend itself to a software discussion because it isn’t software. Above the first level we will describe the actions that the (software) user parts perform. Level 1 is called the physical layer. It deals with hardware and electrical configuration.

Bear in mind that a protocol is only a set of rules. Those rules extend to what occurs in the equipment to control the links. For example, one rule for MTP level one is that a link must consist of two data channels operating in opposite directions at the same bit rate. In other words, the links must be bi-directional.

The standard also refers to the need to disable certain attachments to the link that would interfere with Full Duplex operation and might challenge bit integrity. In other words, MTP level 1 is a user part that deals with physical issues at the level of links, interface cards, multiplexors etc. It does not, therefore, concern software providers except that they need to understand these requirements in order to interface the software module layers with the physical layer.

MTP Level 2

This is a busy user part. It is the last to handle messages being transmitted and the first to handle messages being received. It monitors the links and reports on their status. It checks messages to ensure their integrity (both incoming and outgoing). It discards bad messages and requests copies of discarded messages. It acknowledges good messages so the transmitting side can get rid of superfluous copies. It places links in service, and restores to service links that have been taken out of service. It tests links before allowing their use. It provides sequence numbering for outgoing messages. And finally it reports much of the information it gathers to Level 3.

We’ll examine the functions performed by this layer (Level 2) in the paragraphs that follow.

Signal Unit Delimitation and Alignment - User part functions vary with whether the message is being received or sent. This function is performed by Level 2 when a message is being sent out on a link. The user part attaches an 8 bit code to the beginning of a message package (an MSU, or Message Signal Unit). This code, termed a “flag”, always consists of a byte with zeros at each end and six “1s” in the middle. In ANSI networks a single (lead) flag is used. In CCITT networks, flags are placed at both the lead and trail ends of the Message Signal Unit.

At the receiving end, level 2 reacts to this flag by resetting all resources used in the reading of messages. The result is that the flag indicates an appropriate position in the MSU to begin reading the message.

Quite commonly, the normal coding of other information in the MSU would result in the creation of “imitation” flags. Such imitations would be read by the receiving MTP and interpreted as flags. To prevent this, the MTP on the side sending the message analyzes the entire MSU to find sequences of “1s.” When it locates a sequence of five “1s”, it places a “0” directly after it in a process known as “bit stuffing.”

On the receiving side, the flag alerts the MTP to the beginning of a message. Thereafter, it simply removes a zero after each sequence of five ones, and the message is restored to its original form.

An overly long message is considered “misaligned” and results in the MTP changing its method of counting Signal Unit errors.

Signal Unit Error Detection - When sending a message, the MTP keeps track of the numbers of bits being transmitted in the Message Signal Unit. That number (less the 8 bits of the lead flag) is encoded using a check bits algorithm, and placed at the end of the MSU (before the final flag, if there is one).

At the receiving end, level 2 counts the number of bits seen after the lead flag and decodes the check bits to determine if they are the same. If the check bits don’t equal the count made by the MTP, the message is discarded and a copy is requested.

Signal Unit Error Correction - The standards support two methods of error correction. The type used depends on whether the transmission is land based or satellite based.

In the case of land based transmission, the transmitting side makes a copy of every message sent and retains the copy in a retransmit buffer. When the receiving MTP recognizes a corrupted message, it sends a request to the transmit side for message copies. The transmit side stops transmitting and begins delivery of the copies stored in the buffer instead. Message numbering that is applied to each transmitted message ensures the transmission of the correct messages in the correct sequence.

Positive acknowledgment from the receiving MTP eventually indicates that the receiving side has caught up. At that point the transmitting side stops transmitting copies and returns to its normal copy and send procedures.

For satellite transmission, the procedure is similar with one major exception. A limit is set on held copies. Once the limit is reached, transmission is stopped and the stored copies are transmitted repetitively until acknowledgments are received. This continues until the number of copies which remain unacknowledged reaches some pre-set lower threshold.

Signalling Link Alignment - This functionality is invoked when a link is initially placed into service, and again when a link is being restored to service following a link failure.

The procedure passes through four states, beginning with an Out-Of-Alignment state and ending with an error monitoring proving period. The standards provide for a choice of two proving procedures which can be imposed by the node.

For either proving procedure, the MTP sends the simplest form of message format (the form known as Fill In Signal Unit, FISU) while maintaining a count of FISUs in error.

For the procedure known as Normal Alignment, the proving period lasts for 2.0 seconds and a maximum of four errors are allowed. For the procedure known as Emergency Alignment, the proving period lasts for .5 second and a maximum of one error is allowed. The time period is dependent on the transmission rate. The values given here are those of a standard 64Kbps.

Alignment Error Rate Monitor (AERM) - During the alignment proving period, the MTP watches the FISUs to ensure that each contains an even multiple of 8 bits and is six octets long including the lead flag. The AERM is incremented for each time that this is not so.

Signal Unit Error Rate Monitor (SUERM) - During normal transmission, the MTP watches Signal Units to ensure that each is an even multiple of 8 bits and is at least six octets long. “Ones density” is also checked (imitation flags) and the check bits are tallied against the actual Signal Unit length.

Errors cause the SUERM to be incremented. When the error count reaches 64 the link is removed from service and the alignment procedure begins anew. The SUERM is an increment/decrement counter designed to determine error rate rather than error count. Therefore, each time a series of 256 valid messages has been monitored, the SUERM is decremented. This method of determining an error rate is generally known as the “leaky bucket” technique.

Flow Control - Level 2 monitors congestion on the links. How the sensing is to be accomplished is not specified in the standards. Generally, this is a matter of detecting message buildup at link-associated message queues. When a threshold is passed, the MTP sends a a Signal Unit back to the sending side to indicate SIB (Status Indicator Busy). The receiving side stops returning acknowledgments (both positive and negative).


The receiving side periodically retransmits the Busy status indication and the transmitting side resets the timer for Excessive Delay of Acknowledgment (T_7). This means the transmit side continues to retain the copies in its retransmission buffer. Eventually, either the receiving side will stop sending “Busy” or a second timer (T_6) on the transmit side will time out and a “link failure indication” will be generated.

MTP Level 2 Functionality

As the monitor of both link status and message integrity, the MTP Level 2 is required to report the results of its monitoring both to the adjacent node (in the form of negative and positive acknowledgments along with requests for retransmission; and also signal units containing status information about the links) and to the MTP Level 3 (in the form of status information necessary for level 3 to make decisions about traffic diversions). MTP level 2 functions are illustrated in this drawing.

MTPL3

MTPL2



MTP Level 2 Timers

In general, the user part timers are “countdown” timers whose settings determine either a fixed period of time for an activity (as with alignment proving periods) or a time limitation which is applied to an activity, after which the activity is deemed to have failed (as with “excessive delay of acknowledgment”).


Some of these timers have fixed values, while for others, the standard provides for a range of values. Such ranges provide some flexibility to network nodes to allow node operating efficiency to be optimized. A table of the MTP Level 2 timers follows:

Timer #	Significance	Value or Range
T1	Aligned/ready timeout Started at end of link alignment to allow remote end to complete 4 proving periods before “in-service” declaration.	13 sec. + wide error margin
T2	Not aligned Started/stopped on entry/exit of “not aligned” alignment state. While running, local node sends status indication “O.”	11.5 sec.
T3	Aligned Started/stopped on entry/exit of “aligned” alignment state. While running, local node sends status indication “N” or “E.”	11.5 sec.
T4	Proving Started on entry of proving state. Timeout indicates successful proving unless proving aborts (up to 4 times).	Rate Dependent 2 sec. “N” .5 sec. “E” @ 64kbits
T5	“Busy” Redelivery Interval Starts when “busy” is sent to remote transmitting end. At timeout “busy” is sent again and T5 is reset until “busy” is no longer valid.	80 - 120 ms
T6	Link Failure Remote terminal sets this long timer upon receipt of a “busy”. If congestion does not abate by timeout, “link failure indication” is generated.	3 - 6 sec. @ 64 kbits
T7	Excessive Delay of Acknowledgment Set when messages enter the retransmit buffer and reset when acknowledgment is received. If no acknowledgment is received before timeout, a “link failure indication” is sent to level 3.	0.5 - 2 sec. @ 64 kbits

MTP Level 3

The functions of level 3 are divided into two major categories. One of these is Message Routing (or Signalling Message Handling). The other is Signalling Network Management.


Network Management can be further broken down into four sub-categories. We'll start with Traffic Management and examine each of these sub-categories in turn.



Traffic Management

Understanding the functionality of level 3 becomes easier if you first understand the following three things. First, level 2 monitors the links and the messages. That monitoring provides the basis for reporting status. Status reports, in turn, provide the data necessary for managing the links; and, this is the job of level 3.

The second thing to understand is that links are used both individually and collectively in a number of ways. For example, a group of links extending from a given signalling point to another can be treated as a **linkset**. In concept, messages can be sent over any of these links because all messages will arrive at the same place. This grouping of links together provides choices for the MTP. Messages can be sent on a single link. Or, messages can be distributed across all the links in the linkset. Or, messages can be distributed over most, but not all of the links in the linkset. In the latter scenario, the MTP can take a link out of service and still transfer the messages, because it can simply choose another link in the linkset for the traffic. Then, when the link removed for service is restored, traffic can be returned from the alternate link to the original link. The drawing below illustrates the linkset concept.




In this drawing the links are shown as having been assigned priorities (0 to 3). MTP level 3 provides a code (known as **Signalling Link Selection**) which is normally rotated with each message. Normally, each new code assigns the traffic to a different link in order of priority. Thus, messages are distributed evenly across all links in the set. However, the signalling point can also direct (usually through a configuration) that only two or three of the links be used. As long as those two to three links can handle the traffic, the entire linkset is not used. However, if a link fails (because the traffic is too heavy or due to some physical failure) the unused link(s) can be used as alternates

to which to transfer the traffic. Then, when the original link can once again handle the traffic, it returns to service and the traffic is placed on it once more. In some cases, a higher level user part (the SCCP) is asked to provide a service in which the messages are guaranteed to be in sequence. That guarantee can be made only if all messages are delivered on the same (single) link. To accommodate this request, MTP level 3 **stops** rotating the Signalling Link Selection (SLS) for the duration of the service request; and, all messages are delivered in sequence on a single link.

Links are treated collectively as a part of a linkset. Linksets in turn are treated collectively as a part of a route set. The route set concept is used by the MTP to select from among the available linksets when sending a message to some distant location. As with voice switches this routing can be “best way,” “lowest cost,” etc. From the standpoint of signalling nodes, route sets are generally of little practical significance. The reason is that often a node will maintain only two linksets, which is sufficient to give the node access to a mated pair of STPs. Having only two “routes” (which is what a linkset is called when seen as part of a “route set”) available, there is little or no advantage for the node to be able to select a linkset over which to direct traffic. In such an instance the node will likely declare (through a configuration) that both linksets are part of a route set. Having provided such data, the requirements of the MTP will be satisfied.

An STP, on the other hand, may have numerous link sets going in many directions. In this case, the STP managers will pay a good deal of attention to setting up route sets which offer the greatest advantage. The drawing below illustrates the routeset concept.



The drawing shows routing from the perspective of the Message Transfer Part at the STP with the darker shading (15-15-00). It should be obvious that, when a message is to be sent to the signalling point at 20-20-20, linkset #3 (LS3) and linkset #4 (LS4) are not likely to be the most desirable linksets for the STP to use. There are, of course, many more connections in the network than those shown here. A message addressed to 20-20-20 would (eventually) arrive there even if sent on LS3 and LS4. It is even possible that there might be reason (such as lower cost) to direct the messages there. However, it seems more likely that the managers of this node would use LS1 or LS2 to send messages in this direction.

The routeset might include one or both linksets. Priorities would be assigned to each route (linkset) in the routeset although in this instance it would appear that there would be little to recommend either route in preference to the other.

Simply SS7

Back where we started this discussion we said there were *three* things that needed to be understood about MTP level 3. The first is the nature of a linkset. The second is the nature of a routeset. The third is to understand that the MTP level 3 both reports adverse network situations and responds to those situations by redirecting traffic.

The question is where is the MTP that does each of these? A service provider's node (such as an SSP or a CRP) generally has links only to an STP pair. It therefore has little way of collecting information about the network beyond its links. The STP, on the other hand, maintains links to multiple service provider nodes as well as to other STPs. In general, then, the MTP at the service provider's node will receive Route Management messages from the STP and will respond by diverting its outgoing traffic using Traffic and Link Management functionality. If the service provider's node maintains enough links with enough access to the network, it may succeed in getting its messages around network trouble spots.

Let's see how the MTP at a location which is not an STP might attempt to do this.

Forced Rerouting

When an STP is unable to route to a specific signalling point in the network, it sends a **Transfer Prohibited (TFP)** signal to adjacent nodes. When an STP is unable to route to a specific signalling cluster in the network, it sends a **Transfer Cluster Prohibited (TCP)** signal to adjacent nodes. The level 3 User Part at the node responds by diverting traffic to another available linkset, then returning traffic to the original linkset when normal routing can be re-established

Controlled Rerouting

When an STP encounters difficulty in routing to a specific signalling point in the network (but, transfer is possible), it sends a **Transfer Restricted (TFR)** signal to adjacent nodes. When an STP encounters difficulty in routing to a specific signalling cluster in the network (but, transfer is possible), it sends a **Transfer Cluster Restricted (TCR)** signal to adjacent nodes. The level 3 User Part at the node responds by diverting traffic to a more efficient linkset, then returning traffic to the original linkset when normal routing can be re-established.

MTP Restart

Before returning to the network, a node which has been isolated by the unavailability of its linksets needs time to determine any network routing changes which have occurred while it was not available. Before full startup it sends a **TRW (Traffic Restart Wait)** signal which lets adjacent nodes know not to send traffic even though the links may appear to have resumed service. When the restarting node is satisfied that "enough" links (usually 50%) are available, level 3 sends **TRA (Traffic Restart Allowed)**.

Management Inhibiting

This is a link labeling procedure which does not prevent link usage, but rather, reserves a link for the purpose of sending test messages. A link under congestion cannot be inhibited. If the link is the only one available for traffic, it cannot be inhibited. Finally, an inhibition request requires positive confirmation from the MTP at the other end.

Link Management

MTP level 3 directs the activities of placing links in and out of service. Instantaneous removal or replacement of links would leave no time for the adjacent (directly connected by link) nodes to react. The result would be lost, duplicated, or corrupted messages. Therefore the link management functionality enforces an orderly withdrawal and replacement of links through a mandatory message exchange.

Signalling Link Activation

MTP level 3 directs the activation of the links (at the request of an application), deactivates links in response to changing link status, and reactivates links removed from service once those links can resume service.

Signalling Link Changeover

When traffic must be removed from a link taken from service and brought to another link, MTP level 3 controls the process to eliminate lost messages, duplication, or mis-sequencing. It does so by warning the adjacent node through a **ChangeOver Order (COO)** signal and receives a **ChangeOver Acknowledgment (COA)** in response.

Signalling Link Changeback

Traffic must be returned to a link from which it was diverted when that link is successfully realigned. MTP level 3 controls the process to eliminate lost messages, duplication, or mis-sequencing. It does so by warning the adjacent node through a **ChangeBack Declaration (CBD)** signal and receives a **ChangeBack Acknowledgment (CBA)** in response.

Signalling Link Test

A **Signalling Link Test Message (SLTM)** and a **Signalling Link Test Acknowledgment (SLTA)** are (optionally) exchanged immediately after alignment and periodically (30-90 sec.) while the link is in service. The message data ensures unified agreement about the Signalling Link Code used by both adjacent nodes to identify the link

Routing Management

Routing Management consists of procedures which are really the “flip side” of the Traffic Management procedures. Most of these are procedures performed by the MTP at an STP. The reason this is so is that it is the STP which has routing responsibilities in the SS7 network. As a result, it is the STP which receives most of the information about unavailable signalling points or signalling points experiencing difficulties (such as numerous congested links).

Transfer-Prohibited

Transfer-Cluster-Prohibited

An STP sends a **Transfer Prohibited (TFP)** signal to adjacent nodes when it is unable to route to a specific signalling point in the network.

An STP sends a **Transfer Cluster Prohibited (TCP)** signal to adjacent nodes when it is unable to route to a specific signalling cluster in the network.

Transfer-Restricted

Transfer-Cluster-Restricted

An STP sends a **Transfer Restricted (TFR)** signal to adjacent nodes when it encounters difficulty in routing to a specific signalling point in the network (but, transfer is possible).

An STP sends a **Transfer Cluster Restricted (TCR)** signal to adjacent nodes when it encounters difficulty in routing to a specific signalling cluster in the network (but, transfer is possible).

Transfer-Allowed

Transfer-Cluster-Allowed

These messages (**TFA** & **TCA**) are sent by the STP to the adjacent nodes when the condition which caused it to send the restricted or prohibited messages has been cleared.

Signalling-Route-Set-Test

Nodes generally respond to Route Management Messages sent by STPs by invoking local link management signals like **ChangeOver Order (COO)**.

In addition, the affected node will send a periodic **Signalling-Route-Set-Test (SRST)** to check the routing status. The typical sending interval is 30 seconds.

Congestion (Flow) Control

At some point in time you may be reading through the standards and end up with some confusion about whether flow control is a level 2 or level 3 function. You can limit this confusion by remembering that level 2 is responsible for monitoring and reporting the results of the monitoring. Level 3, on the other hand, uses this information to manage developing problems by causing link, link-set, or routeset diversions and also by reporting to the upper levels.

In the case of flow control, the MTP level 2 must first determine whether congestion exists; and, if it does, to report how bad the congestion is. When congestion occurs at the receiving end, MTP level 2 sends a congestion indication over the link in question to the MTP at the transmitting side. The receiving MTP also stops sending the acknowledgments that would allow the transmitting node to eliminate successfully transmitted messages from its retransmit buffer. Also, at the transmitting node, a long timer is started, which, when it expires, will provide a “link failure indication.”

One of the clearest indications of congestion is the buildup of messages in the transmit and retransmit buffers of the sending node. The standards provide for threshold values to be imposed upon the amount of messaging in the transmit and/or retransmit buffers. Whether the event sensing is placed at the transmit buffer, the retransmit buffer, or both is based on the relative sizes of both. Three identical sets of event values are set for three overlapping message levels. Each level has a congestion onset value, a congestion abatement value, and a discard onset value. When the first congestion onset is reached, the MTP level 2 will inform MTP level 3. Level 3 responds by informing the user part sending the messages (SCCP, ISUP, TUP etc.). If the buildup of messages continues, it will reach the discard onset threshold. When this occurs level three begins to discard messages.

When we examine the message formats in a later section of this book we will see that a message can be assigned a priority. Messages containing only information of a non-critical nature are assigned the lowest priority. This priority value corresponds to the level of congestion indicated when the lowest congestion onset threshold is reached. Because this is so, the messages that are discarded when the lowest discard onset threshold is reached are the lowest priority messages. If the message buildup falls below a threshold which has been set lower than that of congestion onset, the MTP will stop reporting congestion. This is called the congestion abatement threshold.

If the congestion buildup continues despite discard of low priority messages, it will eventually reach congestion onset at the next highest level. The MTP will indicate a higher level of congestion. Then, if the discard onset threshold is reached, the MTP will discard the messages of the next higher priority. At the same time it will still be discarding low priority messages. This escalation of congestion level along with the discard of higher priority messages continues for three priority levels.


The standards support the assignment of four priorities. However, only three levels of Flow Control message discard are assigned because the highest priority is given to network management messages with the understanding that such messages are not subject to discard.

Simply SS7

The title of the previous section connects MTP level 2 Flow Control with MTP level 3 Congestion Control. In reading the standards you may have difficulty making such a connection. However, bear in mind that the job of level 2 is to monitor and report (both to level 3 and to the MTP at the other end of the links). The job of level 3 is to manage (as with links, link sets and routes) and to report to higher level user parts. Congestion status, traffic control, route and link management, status indicating and flow control, are therefore all interrelated.

The distinction between level 2 and level 3 congestion issues is largely a matter of whether congestion is detected on incoming messages queues or outgoing transmit/retransmit buffers. To be correct, then “flow control” (level 2) procedures are invoked when congestion is detected on the receiving end. “Congestion status” procedures are invoked when congestion is detected at the sending end.

The drawing below illustrates level 3 congestion concepts. The levels at which to place CO, CA and CD are dependent on buffer sizes, traffic densities, and other variables which are specific to each side. Where, and at what level to set the thresholds, then, depend on implementation.



CO - Congestion Onset CA - Congestion Abatement CD - Congestion Discard


Message Discrimination and Distribution

MTP level 3 has two more important functions to fulfill. These both fall into the category of message routing and consist of message discrimination and message distribution. Level 3 must examine each message to determine two things. The first question to be answered is, “Is this message intended for this SP?” The second question is, “If intended for this signalling point, for which user part is it intended?” The answer to the first question is obtained by invoking the discrimination functionality while the answer to the second invokes the distribution functionality.

All messages handed off to level 3 (by either a higher or lower level) are either en route to the application via level 4 user parts (SCCP, ISDNUP, TUP, etc.) or they are en route to the adjacent node via level 2 and the SS7 signalling links.

For most nodes, an incoming message not intended for that node was probably incorrectly addressed. MTP level 3 will examine the Destination Point Code (**DPC**) and discard any messages whose **DPC** does not match the Level 3 configuration. Only STPs have the routing responsibility to take such a message, and after consulting routing data, send the message on its way into the network. No message should arrive at an end point node to which it was not addressed.

As for message distribution, the MTP relies on a field in the message to indicate the user part for which the message was directed. Later we’ll take a look at that field along with many others. The flow chart below illustrates message discrimination and distribution.



MTP Level 3 Timers

In general, the user part timers are “countdown” timers whose settings determine either a fixed period of time for an activity (as with alignment proving periods) or a time limitation which is applied to an activity, after which the activity is deemed to have failed (as with “excessive delay of acknowledgment”)

Some of these timers have fixed values, while for others, the standard provides for a range of values. Such ranges provide some flexibility to network nodes to allow node operating efficiency to be optimized. The following is a table of the MTP Level 3 timers:

Timer #	Significance	Value or Range
T1	Delay to avoid message mis-sequencing on changeover	.5 to 1.2 sec.
T2	Wait for changeover Ack	.7 to 2.0 sec
T3	Delay to avoid message mis-sequencing on changeback	.5 to 1.2 sec.
T4	Wait for changeback Ack (1st attempt)	5 to 1.2 sec.
T5	Wait for changeback Ack (2nd attempt)	5 to 1.2 sec.
T6	Delay to avoid mis-sequencing on controlled routing	5 to 1.2 sec.
T7	Wait For Signalling Data Link Connection Ack	1 to 2 sec.
T8	Transfer-prohibited inhibited timer (transient solution)	0.8 to 1.2 sec.
T10	Wait to repeat signalling-route-set-test message	30 to 60 sec.
T11	Transfer-restricted timer	30 to 90 sec.
T12	Wait for uninhibit Ack	0.8 to 1.5 sec.
T13	Wait for forced inhibit	0.8 to 1.5 sec.
T14	Wait for inhibition Ack	2 to 3 sec.
T15	Wait to repeat signalling-route-set-congestion test	2 to 3 sec.
T16	Wait for route-set-congestion status update	1.4 to 2.0 sec.
T17	Delay to avoid oscillation of initial alignment failure and link restart	0.8 to 1.5 sec.
T19	Failed link craft referral timer	480 to 600 sec.
T20	Wait to repeat local inhibit test	90 to 120 sec.
T21	Wait to repeat remote inhibit test	90 to 120 sec.

MTP Level 3 Timers (continued)

Timer #	Significance	Value or Range
T22	Restarting SP waiting for links to become available	network dependent
T23	Restarting SP waiting to receive traffic restart allowed	network dependent
T24	Restarting SP with transfer function waiting to broadcast traffic restart allowed messages	network dependent
T25	SP adjacent to restarting SP waiting for traffic restart allowed message	30 to 35 sec.
T26	Restarting SP with transfer function waiting to repeat traffic restart waiting messages	12 to 15 sec.
T27	Minimum duration of availability for full restart	2 to 5sec.
T28	SP adjacent to restarting SP waiting for traffic restart waiting message	3 to 35 sec.
T29	Timer started when TRA sent in response to unexpected TRA or TRW	60 to 65 sec.
T30	Timer to limit sending of TFPs and TFRs sent in response to unexpected TRA or TRW	30 to 35 sec.

Notes: Timers T9 and T18 are not shown because they are not used.


Timers T22, T23, and T24 are started in sequence by a restarting SP.

SCCP

It may sound strange, but we'll begin our discussion of the Signalling Connection Control Part with information about what it can do, but doesn't do. With those things out of the way, we'll be able to focus on what it *does* do.

You may recall from an earlier drawing that we can trace a message through the user parts simply by drawing a straight line from level 1 to the Application or from the Application to level 1. As the drawing below illustrates, there is one area where, if we did that, the line would be drawn through the MTP layers, then through the SCCP, and finally through ISDNUP. This would indicate that the SCCP can be used with ISDNUP for sending messages in the circuit switched PSTN (Public Switched Telephone Network). And indeed, the standards did anticipate such usage for the purpose of end-to-end phone call routing. ISUP routing, while referring to the final destination of the call, actually passes from switch to switch following the same path as the voice connection.


As the SS7 network was deployed, telephone companies and other service providers found adequate functionality for their purposes with the ISUP (circuit concerned) portion of the ISDNUP. As a result there are really no deployed uses for SCCP used in conjunction with ISDNUP. As a result, we needn't spend any time looking into the functional aspects of such use.



Another subject which we can dispose of rather quickly is the matter of connectionless vs. connection oriented services. The standards were written for the SCCP to support both connection oriented and connectionless services. The concept of a “connection” is not the same for the SS7 network as it is for the voice switching network. In general, the physical “connections” or links always exist in the SS7. The term “connection oriented” refers to a “virtual” connection rather than a physical one. In a “virtual” connection a conversation takes place before the conversation takes place. That is, the two sides first agree to establish the connection before they begin the exchange of messages for which the connection is intended.

This may seem like overkill. The idea is to make the rules of communication more rigid in order to better guarantee the integrity of the communication. The phases of such a connection resemble those of a telephone call. The phone call consists of a setup phase followed, by a voice conversation, followed by a release phase. A “virtual connection” consists of a connection establishment phase followed by a data transfer phase, followed by a connection release phase. Connection oriented services tend to use more resources and create greater “overhead.” The trade-off for increased integrity is lower “speed.” In other words, even though the data transfer rates are the same for connectionless service, less significant data is passed per unit of time. Service providers have generally found connectionless services to be highly reliable with little tendency toward corruption; and, have, therefore opted *not* to make use of connection oriented services.

Now we can concentrate on the actual uses of SCCP. The service types offered by the SCCP are shown below along with the functions of the user part.



SCCP Service Types

Both services shown here are connectionless. Essentially connectionless means that messages are simply delivered to their destination and returned to the point of origination without any additional special handling. No virtual connection is established.

For **class 0** usage, messages are transported without reference to other messages. Delivery of messages is not guaranteed to be sequential. On the other hand, for **class 1**, the SCCP calls on the services of MTP level 3 to modify its normal handling of links in link sets. MTP level 3 normally provides a rotating code value (Signalling Link Selection or **SLS**) to share the load in a link set. When asked to do so by SCCP, the MTP stops rotating this code. The result is that the **SLS** stays the same, message after message. Messages delivered on the same link remain in sequence.

SCCP Specialized Routing Functions

Many of the chief benefits of the use of the SCCP lie in the specialized routing functions. The addressing capabilities here are what allows the locating of database information or the invoking of features at a switch. To this point we have seen only the addressing capabilities of the MTP. The MTP, of course, is concerned only with transferring messages to the other end of links. For this reason its addressing is limited to the use of the point code of the location to which the link goes. The MTP deals only with the signalling point code of its own location in the network (which becomes its Origination Point Code), and the signalling point code of the node at the other end of the link (which becomes its Destination Point Code).

Destination Point Code (**DPC**) is important to the SCCP as well. When supplied by the originator of a query it indicates a network location where there is a process to retrieve data from a database (or where a service or feature can be invoked). However, it may be that data can be retrieved from more than one database at that location. For this reason, SCCP needs to provide addressing which can be used to differentiate between databases or between various features or services that can be invoked at a node. The value used is called a subsystem number.

The values provided for database identification range from 0 to 255. As of 1996, some of these were assigned to the most common usages. The “0” value, if used, indicates an “unknown” database. The first seven values are assigned for uses of the databases required for such things as the ISDN/Telphony Numbering Plan (1), The Data Numbering Plan (3), the Land Mobile Numbering Plan (6), etc. With the exception of some reserved (for future use) values, the remainder can be used. The SCCP uses **DPC** and **SSN** (subsystem number) to route to the appropriate network location and to the appropriate database which can be accessed at that location.


Still another addressing mechanism is available to SCCP users. This one is known as Global Title Translation (**GTT**). To understand the need for such a mechanism, let’s examine what happens at an SSP when someone dials an 800 number. Ordinarily, the dialed digits received at the switch would contain an area code, an exchange code, and a line number. The switch could then consult its routing table and determine the next switch through which to route the call. An 800 number, on the other hand, does not comply with the North American Numbering Plan which the switch would normally refer to in choosing the routing. The telephone being called is connected by voice

line to an end office somewhere. That end office has an address in the North American Numbering Plan and the line to the telephone has a line number. Therefore, the telephone has a regular telephone number to which the switch (at the SSP) could route the call. Of course, every switch in the country could maintain an 800 number database for the purpose of translating the dialed digits of an 800 number into a normal telephone number. In this case, possible and practical are not the same thing. If it were done this way, every switch in the country would find it necessary to update its routing table every day. There *is* a better way. If the switch simply maintains a database which provides it with the Destination Point Code and Subsystem Number of a centralized 800 database, it can send a query there and receive a normal telephone number in return. Now it can handle the call as if a normal telephone number were dialed in the first place.

With the rate at which databases proliferate, it is only a matter of time before each switch will be maintaining a database of databases that will need frequent updating. This is one of the problems solved by GTT. Using GTT, the SSP need not know where the database is located. Instead, it makes a request for a Global Title translation and passes it on to the STP. The STP may have the DPC and SSN of the database in its own tables. If so, it redirects the request and later returns the data to the Point Code making the original request. If the first STP does not have the necessary address, it will usually have the location of another STP (perhaps at a different level of hierarchy) to which it can pass the Global Title request. When an STP which has access to the data is reached, it can retrieve the data and send it back to the location which was the one that sent the request to its final STP destination. Node by node, the data gets returned to the requestor. The originator of the request can then make use of the data without ever knowing where it came from.

There are at least two benefits to be derived from Global Title Translations. The first is that SPs can have access to data of all types without having to maintain ever more cumbersome tables. New data can become universally available very quickly. The second is that companies can have better control over the data kept within their own networks. For example, the STP serving as the gateway for network can use GTT to hide the location of databases from outsiders. Access to the databases is controlled and can either be restricted, or provided at a fee.

The drawing below illustrates Global Title usage.



SCCP Subsystem Management

To understand SCCP Subsystem Management, it pays to start with an understanding of how reliability is gained in any network. As was mentioned earlier, the development of standards is filled with “what ifs.” In the case of the SCCP, one of the pertinent what ifs is, “What if a query arrives at the SCP when the database is down for maintenance?” The answer once again lies in that major ally of all networksredundancy. The database front end can still fulfill the request as long as there is a database available containing the same data as the one currently out of service.

The timing of the switch over from an out-of-service database to an in-service database is critical. If the query is received at the same moment that a database is withdrawn from service, it may be too late for the front end process to stop its attempt to get data from the now out-of-service database. Likewise, a precipitous shut down of the database could result in queries receiving only partial answers. The obvious answer is to plan a shutdown period in such a way that the database goes out-of-service only after those who need the data have been informed of the impending lack of accessibility. Likewise, those who are normally accessing the data need time to switch to an alternate database.

Generally, when a replicated subsystem needs to leave service, the SCCP database is examined for “Concerned Points”. Such a point is a location which has to be informed of subsystem status because it is a location normally accessing data from the database in question. There are a number of ways in which replicated (multiply copied) databases can be implemented. Whichever way is used, the “Concerned Point” will generally have access to the data from a redundant source. The withdrawing database will make an effort to acquire agreement for the withdrawal from the redundant database by sending an **N-Coord request** primitive to the local SCCP. The local SCCP, in turn, sends a Subsystem-Out-of-Service-Request to the SCCP at the location of the redundant database. At the same time it sets a timer to wait for permission to be granted. If the timer times out before permission is granted, the local SCCP sends an **N-Coord confirmation** primitive back to the subsystem which indicates “denied”.


At the redundant database, the availability of resources (replicated database, traffic handling capabilities, etc.) is checked. If the resources are sufficient for the abandoned query/response load to be handled, an **N-Coord response** primitive is sent to the SCCP at the location of the database which requested permission to withdraw; and the request is granted.

The SCCP broadcasts messages to the **Concerned Points** found in its configuration information. When a database is withdrawn from service “Subsystem-Prohibited” is broadcast to its **CPs**. When that same database returns to service, the SCCP broadcasts “Subsystem-Allowed.” The same messages should be sent to the mated database to keep the replicate informed of the status of the database whose load it is now handling.

When the “Subsystem-Prohibited” is received at the CP, the local SCCP sets a timer. When the timer times out a “Subsystem-Status-Test” is sent to the SCCP peer at the Out-of-Service location. The timer is reset when the message is sent and the cycle begins anew. “Subsystem-Status-Test” is no longer sent after the database returns to service. “Subsystem-Status-Test” (**SST**) is a request for the prohibited database

You may wonder, “Why send an **SST** (Subsystem-Status-Test)?” After all, the SCCP at the database which is withdrawn from service will broadcast **SSA** (Subsystem-Allowed) when the database is returned to service. The answer is, “What if somehow the SSA is not received?” It is, after all, a single broadcast on a link transmitting 64Kbps. The SST will repeat until an “allowed” response is received. This guarantees that resources can be returned to their original state as soon as possible.

The drawing below illustrates the broadcast of Subsystem-Allowed and Subsystem-Prohibited:



SCCP Timers

In general, SCCP timers are used to coordinate Subsystem Management and to coordinate connection establishment and release.

Some of these timers have fixed values, while for others, the standard provides for a range of values. Such ranges provide some flexibility to network nodes to allow node operating efficiency to be optimized. A table of the SCCP timers follows:


Timer #	Significance	Value or Range
T(stat.info)	Interval between SST requests	30 sec.
T(coord.chg)	Wait for Subsystem-Out-of-Service-Grant	30 sec.
T(ignore'SST)	Delay from receipt of Subsystem-Out-of-Service-Grant to actually ending service	30 sec.
T(rtg.stat.info)	Interval between requests for subsystem routing status	30 sec.
T(conn'est)	Wait to receive Connection Confirm message	provisional 3 to 6 min.
T(rel)	Wait to receive Release Complete message	provisional 10 to 20 sec.
T(reset)	Wait to receive Reset Confirm message	provisional 20 to 40 sec.
T(ias)	Delay to send message over connection	provisional 5 to 10 min.
T(iar)	Wait to receive message over connection	provisional 11 to 22 min.
T(fr)	Wait to reuse local reference number	30 sec.
T(guard)	Wait to resume normal procedures for temporary connections	provisional 8 to 16 min.
T(int)	Wait for Release Complete message; or to release resources after T(rel) expiry	<= 1 min.
T(repeat_rel)	Wait for Release Complete message; or to repeat sending Released after T(rel) expiry	10 to 20 sec.

TCAP

As we have moved upward in the stack, each layer has had functionality which can generally be described as “offering services” to the layer above. The Transaction Capabilities Application Part offers its services to user designed applications as well as to OMAP (Operations, Maintenance and Administration Part) and to IS41-D (Interim Standard 41, revision D) and GSM MAP (Global Systems Mobile). Many drawings will show OMAP, IS41-D, and GSM MAP as separate layers above TCAP. Actually, they are constructed within the TCAP procedures. An examination of the messaging used for each would reveal what appears to be variations on TCAP. It would be accurate to say that the purpose of TCAP is to allow applications to exchange information using signalling that is not circuit related. It would also be accurate to say that TCAP is used largely by switching locations to obtain data from databases (SSP from 800 Db, MSC from HLR, etc.), or to invoke features at another switch (like Automatic Callback or Automatic Recall). Using the mobile message transports (GSM, IS41-D) also provides the procedures for database updating by reference to another database. In the Mobile network, for example, those who “roam” out of the network maintained by the company they pay for the service, are tracked by the company into whose area they wander by transferring data from database to database. In this case the “Home Location Register” (HLR) is accessed by the new company and information is transferred into the new company’s “Visitor Location Register” (VLR).

Unlike circuit related messages, TCAP messages need not be sent from switch to switch following a developing voice circuit. Instead, the messages are generally sent to distant locations using the end-to-end message routing provided by the services of the SCCP. Indeed, TCAP messages end up being attached to SCCP messages. The two user parts go together, with SCCP providing the specialized routing and TCAP providing the appropriate message structuring and parameters to acquire and package the data.

Remember that TCAP uses the services of the SCCP which in turn uses the services of the MTP. The result is that a TCAP message ends up being sequenced in the message packet as shown in the drawing below. The arrow indicates the direction of the message flow, so the portion at the right of the drawing will be the first part received.




Simply SS7

As illustrated in the drawing, the TCAP message consists of two major portions. The first is the Transaction Portion. The Transaction Portion contains none of the data of the message. Instead, it consists largely of protocol control. It begins with an octet (TCAP Package Type) that identifies the type of dialog that is to take place.

Following the package type, the transaction portion indicates how many octets of data are contained in the entire message. It then deals (in housekeeping terms) with a Transaction ID. The transaction id is a value which is assigned to a complete query/response dialog and remains the same for the duration of the dialog. The reason this is necessary is that the application which is the originator of the query may have numerous queries out at one time. Some of those queries may be brief, single package transactions. Others may be multi-package transactions segmented to fit by the SCCP. Either way, the originator of the request needs to be able to correlate the query with the response information. That correlation is provided by an identifying value generated with the query and returned with all data associated with the response.

The drawing below illustrates the fields in the Transaction Portion:



A description of the information conveyed in each of these fields follows.

Transaction Portion Fields

TCAP Package Type - This field is filled in by the signalling point making the request to let the receiver know the nature of the request. It also contains the data necessary to relate this message to other messages which may be part of the same continuing transaction. Here are the types used.

Unidirectional - A one way message with no reply expected.

Query With Permission - The signalling point making the request does not expect to use the same transaction to send other messages. It therefore grants the receiving application permission to release allocated resources after responding.

Query Without Permission - The signalling point making the request *does* expect to use the same transaction to send other messages. It therefore denies the receiving application permission to release allocated resources after responding.

Response - This is the answer that the receiving application sends to a *Query With Permission*. The message indicates the termination of the transaction.

Conversation With Permission - The signalling point receiving the request for a query (with or without permission), tells the originating signalling point to continue the transaction and grants permission to terminate when the transaction is complete.

Conversation Without Permission - The signalling point receiving the request for a query (with or without permission), tells the originating signalling point to continue the transaction but denies permission to terminate when the transaction is complete.

Abort - While the standards intended the Abort as a means for an originating signalling point to end a transaction, it generally becomes useful only when protocol errors are received.

Total TCAP Message Length - This value indicates the number of octets there are from this point in the stream of data to the end of the TCAP message including all components and all parameters.

Transaction ID Identifier - This might better be called Transaction ID Indicator. It is not the actual identifier, but only an indication that an identifier is present and will follow.

Transaction ID Length - With unidirectional messages, there is no Transaction ID and this value will be zero. For messages other than unidirectional, the requesting node provides an ID for the purpose of correlating the query with the response. Such an ID will be four octets long. Sometimes a component will be sent to the originating node by the responding node and a reply is expected. In such instances, the responding node fills in both the original (requesting) Transaction ID and one provided by its own application. The dual IDs are, obviously, longer than a single ID.


Transaction ID - This is the actual value assigned to the transaction for the purpose of providing the correlation between queries sent and responses received. Generally, the ID is provided by the originator of the message. However, a second ID can be provided by the receiver when it finds it necessary to send a message back to the originator and anticipates a response.

Component Sequence Identifier - This value indicates the sequence of components to follow without indicating the actual number of components. While in the Transaction Portion, this code is generally considered to be part of a header for the component portion.

Component Sequence Length - This is the final field before the component portion. It indicates the length of the message beginning with the first field of the component portion and ending with the last parameter of the last component.

Component Portion Fields

The drawing below illustrates the fields of the Component Portions. Invoke Component Fields are shown. The Operation Code Fields are replaced by other fields in different components.



Component Type Identifier - TCAP uses four component types. Two of these make use of two subtypes (“Last” or “Not Last”). The result is that this field may contain any one of six type indications. These are as follows:

Invoke Component - This component is used to petition for a specific operation. These generally take the form of invoking a feature at a switch or of making a query at a database (e.g. 800 number translation). There may be more than a single Invoke in a transaction. If it is a final (or only) Invoke, the code generally includes an “L” (last). If more Invoke Components are to come (in this transaction) the code generally includes an “NL” (not last).

Return Response Component - This component is the one in which the requested results are returned. For example, this would be where the requesting SP would find a telephone number which can be routed, resulting from a requested 800 number translation. As with Invoke, there may be more than a single Return Result in a transaction. If it is a final (or only) Return Result, the code generally includes an “L” (last). If more Return Result components are to come (in this transaction) the code generally includes an “NL” (not last).

Return Error Component - As the name implies, this component is generally used to report the reasons for the failure of an operation. The errors reported have generally to do with errors made by the application. This includes such things as inconsistent use of IDs, unexpected results, or invalid parameters.

Reject Component - This component is also used for error reporting. However, in this instance errors are attached to an indication of which part of the message contained the error. They may be reported as having occurred in the Transaction portion or in any of the other three components.

Component Length - This field indicates the length (in octets) of the component in which it is found. The component portion may consist of multiple components, each of which will have its own length indicator.

Component ID Identifier - The presence of this identifier indicates that this component has an Invoke ID or that another ID is added for the purpose of correlation of Invokes with Return Results. If sent by the originator, the value will be used in the Return Response (or any of the other components) returned. Note that this is separate from Transaction IDs because Transactions can involve more than one Invoke.

Component ID Length - The length indicated here is that of the ID field. This is used because the ID field length will vary from 0 for a Unidirectional message, to 4 or 8 depending on whether the ID represents an Invoke ID only, or a combination of an Invoke ID and a correlation ID.

Component IDs - In this field there may be an (optional) ID assigned to an Invoke component. This is not a network requirement and, when used, is significant only to the sender of the Invoke. When an Invoke ID is sent, a correlation ID becomes mandatory for the SP returning components to use the correlation value with any component returned which is in response to that Invoke.

The Fields Between Component IDs and Parameter Identifiers - In the next section we will discuss the Operation Code fields which were illustrated in the previous drawing. In the Return Result Component, these fields are missing entirely. In the Return Error Component, these fields are replaced by Error Code Identifier, Error Code Length and Error Code Fields. In the Reject Component these fields are replaced by Problem Code Identifier, Problem Code Length and Problem Code Fields.

Operation Code Identifier - This field carries an identifier for the National or for Private TCAP Networks. In the U.S., ANSI and Bellcore standards are implemented in the National network. Private networks may use their own coding internally, but any communications with other networks (e.g. the PSTN) must be ANSI compatible.

Operation Code Length - This field indicates the length of the operation code only. For the National TCAP network the value is always 2 (octets). No such limitation applies to Private networks.

Operation Code - Operation codes are not specified in world wide standards. They are, instead, considered to be implementation dependent. Generally, they are used to specify the operation and how it is to be carried out. Most often, they are separated into categories such that one portion of the code will indicate a family of operations while another provides specifics.

Parameter Set Identifier - This is a single octet field which identifies the individual types of parameters. For example, this might be a timestamp, digits, or a network identifier.

Parameter Set Length - The length of the Parameters field is variable. Its actual length is indicated by this field.

Parameters - This field contains the actual parameter values. For example, if the Identifier indicated a timestamp, these fields would contain grouping of octets which would give the year (2 octets), the month (2 octets), the day (2 octets) the hour (2 octets), the minutes (2 octets) all in binary. Following this there might also be fields expressing the difference between local time and Greenwich Meridian Time.

The term “Parameter Set” used here refers to the fact that Parameters are grouped by specific data types. The examples here were all of the use of the “timestamp” set. Others (such as the “digits” set) contain data grouped by different data categories. Not all parameters in a set must be used. Operation codes can indicate which values are to be selected, thereby creating a “Parameter Set” which is specific to this component.

ISUP


As mentioned previously, the Integrated Services Digital Network User Part (ISDNUP) is usually thought of as being composed of standards referring to switch-to-switch circuit related messaging (ISUP),

ISUP is used throughout the PSTN (Public Switched Telephone Network) to provide the messaging necessary for the set up and teardown of all circuits, both voice and digital. Wireless networks also make use of ISUP to establish the necessary switch connections into the PSTN. In the telephone network ISUP messages follow the path of the voice circuits. That is, ISUP messages are sent from each switch to the switch where the next circuit connection is required.

ISUP offers two types of services, known as **Basic** and **Supplementary**. Basic Services consist of those services employed in the process of setting up and tearing down a call. Supplementary Services consist of those services employed in passing all messages that may be necessary to maintain and/or modify the call.

ISUP functionality can be further broken down into three procedural categories. The first of these is **Signalling Procedure Control** (SPRC) which directly interfaces with the services of the MTP. The SPRC, in turn, provides support for **Circuit Supervision Control** (CSC) and for **Call Processing Control** (CPC). The application, which deals with the circuit connection requirements of the switch, and simultaneously with SS7 signalling, is usually referred to as a Call Control application.

The drawing below illustrates the architecture of ISUP procedures.



Simply SS7

Call Setup

ISUP protocol is employed in the process of setting up and tearing down a call (referred to as Basic Services) and in passing all of those messages which may be necessary to maintain and/or modify the call (referred to as Supplementary Services). For our purposes here we will assume analog lines in use between the connected parties and the telephone company.

One of the easiest ways to understand ISUP messaging is to follow a scenario of the process involved in making a simple telephone call. At the beginning of the process the calling party lifts the phone off the receiver. This places a current on the subscriber line interface of the local exchange. This is DC signalling. The local company acknowledges the presence of that current by sending a dialtone to the calling party.

Having heard the dialtone, the calling party dials the phone, thereby sending the address (telephone number) of the called party to the local exchange. The local exchange waits until all digits have been dialed, and then examines the digits to see if the called party is local (no area code) or whether the call must be routed through a long distance carrier. If the call is long distance, the call will be routed to the long distance carrier through a point-of-presence (**POP**) in the Local Access Transport Area (LATA) of the calling party. The prefix and subscriber number (last four digits) will be used in the first message (**IAM**) routed to the distant exchange.

The call setup is sent using the ISUP protocol through the SS7 network. The STP is used to route this message, and beyond that, plays no significant role in setting up the voice circuits.

Once all of this information has been collected, the originating exchange creates an initial address message (**IAM**) and sends it to the intermediate tandem. All the information necessary for the tandem exchange to establish a connection is carried in this **IAM**.

The exchange thus addressed may not be the final destination of the call. It may be a tandem being used as an intermediate switch to reach the final destination. The local exchange decides how to route the call by reference to its trunk routing tables. These tables define the voice circuits to use in the establishment of an end-to-end circuit with the least number of hops. The local exchange uses the circuit information to create a call setup message which is sent to this first voice connection exchange.

The tandem exchange acknowledges receipt of the **IAM** by sending an address complete message (**ACM**) back to the originating exchange. This indicates that the tandem has reserved a circuit designated for reservation in the **IAM**. Receipt of the **ACM** triggers the originating exchange to send the “phone ringing” (ringback) tone to the calling party.

While the intermediate tandem is sending the ACM back to the originating exchange, it can begin setting up the next circuit between itself and the destination exchange. This is accomplished, once again, through the use of an IAM sent to the next destination (in this example, the final destination). This IAM contains the called and calling party addresses that the tandem received from the originating exchange.

The **IAM** also specifies the signalling method to be used for this call. For example, if the IAM specifies the use of ISUP protocol from **end-to-end**, then the call will be set up using the ISUP protocol. In the unlikely event that the exchange does not support ISUP to this destination, or that there are no facilities available that use ISUP, the call will be rejected and a reason for the rejection will be returned to the originating exchange.

The IAM may specify that ISUP is **preferred**, but not mandatory. In such a case the call will be set up using ISUP (if available) or some other method such as TUP or multifrequency signalling. The **IAM** may also specify that ISUP is required **where available**, but that it need not be available “all the way.” In such a case, other methods may be employed at intermediate exchanges which cannot provide ISUP.

When the **IAM** is received, the destination exchange examines the **IAM** to see if it contains an indication of further information in subsequent messages. When a determination is made that all information is present, the exchange checks the line of the called party to determine its availability. If that line is busy, the destination exchange returns a REL (release) message to the originator, the originating exchange places a “busy” tone on the calling party’s line and all voice circuits are released. A small time lag in the generation of the ringback to the calling party’s phone ensures that ringback is never heard before a busy tone.

When the destination exchange finds that the called party’s number is not busy, it sends an address complete message (**ACM**) back to the intermediate exchange. The intermediate exchange has already acknowledged the originating exchange with an **ACM**, and, thus, needs only to maintain its part of the voice circuit. The destination exchange sends an alerting signal to the called party’s phone and rings the phone.

No further messaging occurs until the ringing phone is lifted from the cradle. When that happens, the destination exchange senses DC loop current on its subscriber interface. The destination exchange then sends an answer message (**ANM**) back to the intermediate exchange. The final leg of the voice circuit is immediately cut through when the intermediate exchange receives the **ANM**. The intermediate exchange now sends an ANM to the originating exchange which begins the cut through of the entire voice circuit to its destination and the call is connected.


Once the call has been connected, no further SS7 messaging is necessary until either party goes “on hook.”

When the phone is hung up, the exchange local to that phone sends a **REL** (Release) to the next exchange. That exchange releases its circuitry and sends a REL to the next switch and an **RLC** (release Complete) to the previous switch to acknowledge that the circuitry has been released.

Simply SS7

The messages involved in setting up the call (both primitives and SS7 messages) are shown in the next drawing. The standards leave room for variation. For example, as you will see in the drawing, intermediate exchanges generally await the return of an ACM from the destination exchange before returning an ACM to the previous exchange. But, in fact, some intermediate exchanges will return an ACM as soon as they have received an IAM from the previous exchange.

ISUP Normal Call Scenario



Here is a repeat of some of the message definitions to assist you in understanding the drawing.

IAM - Initial Address Message

This is an ISUP message containing all the information necessary for a switch to establish the connection.

ACM - Address Complete Message

This message serves as the acknowledgment of an IAM. The ACM indicates that the switch sending it has reserved the circuit designated for reservation in the IAM. Receipt of the **ACM** triggers the originating exchange to send the “phone ringing” (ringback) tone to the calling party.

ANM - Answer Message

When the called party picks up the phone, the destination exchange senses DC loop current on its subscriber interface. As a result, that exchange sends an answer message (**ANM**) back to the intermediate exchange. Each switch in the circuit completes its portion of the circuit and returns an ANM to the next switch closer to the calling party. When the ANM reaches the originating exchange, the final leg of the voice circuit is immediately cut through and the call is connected.

REL - Release

This message is sent first by the exchange sensing that the phone was hung up. Each subsequent exchange sends its own REL to the next exchange and initiates release of the circuitry.

RLC - Release Complete


Each exchange receiving an **REL** sends an **RLC** message back to acknowledge receipt of the REL and to indicate that circuit release has been initiated.

ISUP Message Structures


Compared to TCAP messages, ISUP messages are really quite simple. The reason is that ISUP messaging deals with the relatively rigid requirements of circuit connection and control. This means that the data is easily structured in the form of parameters and that these parameters can then be put together in a myriad of ways to achieve any result necessary. ISUP messages, then, do not require the “tight” protocol control of a TCAP message.

As was mentioned earlier, the standards do provide for ISUP to use the services of SCCP for end-to-end routing. However, to date those who work in the PSTN (Public Switched Telephone Network) have not seen fit to employ ISUP in this way. They are, instead, content (even happy) to use ISUP to communicate from switch to switch following the voice path. The wealth of information about circuitry that can be gathered this way makes end-to-end routing advantages appear rather small.

The drawing shown here illustrates the ISUP message in a message packet. This time, the packet is borne along using the services (and the protocol control) of only the MTP.



The next drawing illustrates the fields of the Header Portion.




Header Portion Fields

Destination Point Code - This field indicates the signalling point that is the destination of the message. Thus far we have referred to point codes without any real description of their nature. The time has come. Signalling Point code resembles telephone numbers to some extent. The ANSI standard breaks the zoning of a national network into three components. In order, they are called Network, Network Cluster, and Network Cluster Member. As with telephone numbers, each successive group represents a smaller hierarchical (not necessarily geographical) area with the final number representing an individual node. The ANSI standards provide that each group is represented by a byte. The CCITT standards provide for a code of three groups (named differently) with the first and third group being represented by 3 bits; and the middle group being represented by a byte. We'll take a look at both types when we later examine message packets.

Origination Point Code - This field indicates the signalling point that is the originator of the message.

SLS Code - In earlier portions of this book, the Signalling Link Selection code was shown to be the mechanism by which the links in link sets are assigned traffic. If the SLS is rotated (and it normally is) messages are directed to successive links with the result that loads are shared across the links of a linkset, or even across the linksets of a combined linkset. However, in those instances in which sequential delivery of messages must be guaranteed, the rotation of the SLS can be stopped for the duration of the message transfer and all messages will be delivered on the same link.

Trunk CIC Code - The **Circuit Identification Code** is a two octet field used to provide identification of the specific trunk circuit used to establish the voice (or data) connection path. The standards do not specify how such identifiers should be allocated and this is generally done by mapping internal values with actual trunk configurations from switch to switch. Circuit Identification Codes are mapped to actual voice/data channels. The mapping results in both sides of the circuit agreeing on a common code (the CIC) to identify the shared circuitry.



Parameters Portion Fields

Message Type - This field indicates one of more than sixty ISUP message types. Each Message Type is associated with its own sets of parameters. This means that some message types make use of all three types of parameters while others (such as **Unequipped CIC**) use no parameters at all. Still others use only Optional Parameters.


Mandatory Fixed Parameters - As the name implies, this group consists of parameters which are of a fixed length for each specific message type. The message type dictates whether these are used, and if used, the specific parameters that apply. The message type defines the parameter and, therefore, no indicators of length are required. For example the IAM (Initial Address Message) uses Nature of connection indicators, Forward call indicators, and Calling Party's category in this portion of the message.

Mandatory Variable Parameters - These parameters are of variable length. Therefore, throughout the variable part there are single octet pointers to show where a parameter begins. There are also length indicators to indicate the length of the parameters. The combination of pointing to the beginning of a parameter and indicating its length makes the parameter readable to the receiving side. Once again, the actual parameters used are dependent on message type.

Optional Parameters - These parameters are also related to message type, but the sender may decide whether or not they are used. In many instances, the Fixed and Variable length types also can be included in the Optional grouping. To make these readable, pointers to the beginning of the parameter are provided along with length indicators. In addition, optional parameters are provided with a name field.

ISUP Message Fields

The following drawing illustrates the sequence of ISUP message fields.



ISUP Timers

Timer #	Initiation Cause	Termination Cause	Time-Out Value
T1	Sending of RE lease Message	Receipt of Re lease Complete	4-15 sec.
T2	Not specified for U.S. Networks		
T3	Not specified for U.S. Networks		
T4	Not specified for U.S. Networks		
T5	Sending of initial RE lease	Receipt of Re lease Complete	1 min.
T6	Receipt of SUS pend by controlling exchange	Receipt of RES ume	10-32 sec.
T7	Sending of I nitial A ddress M essage	Receipt of A ddress C omplete M essage or A Nswer M essage	20-30 sec.
T8	Receipt of I nitial A ddress M essage with previous or incoming continuity check	Receipt of C OnTinuity message	10-15 sec.
T9	Receipt of A ddress C omplete M essage at outgoing international exchange	Receipt of A Nswer M essage	2-4 min.
T10	Not specified for U.S. Networks		
T11	Receipt of I nitial A ddress M essage in interworking situations	Sending of A ddress C omplete M essage	15-20 sec.
T12	Sending of B LOcking Message	Receipt of B Locking A cknowledgment Message	4-15 sec.
T13	Sending of initial B LOcking Message	Receipt of B Locking A cknowledgment Message	1 min.
T14	Sending of Un B Locking Message	Receipt of Un B Locking A cknowledgment Message	4-15 sec.
T15	Sending of initial Un B Locking Message	Receipt of Un B Locking A cknowledgment Message	1 min.
T16	Sending of Re Set Circuit Message (not due to expiry of T5)	Receipt of Re lease C omplete M essage	4-15 sec.
T17	Sending of initial Re Set Circuit Message	Receipt of Re lease C omplete M essage	1 min.
T18	Sending of C ircuit G roup B lock-ing	Receipt of C ircuit G roup B locking A cknowledgment Message	4-15 sec.

ISUP Timers (Continued)

Timer #	Initiation Cause	Termination Cause	Time-Out Value
T19	Sending of initial Circuit Group Blocking	Receipt of Circuit Group Blocking Acknowledgment Message	1 min.
T20	Sending of Circuit Group Unblocking	Receipt of Receipt of Circuit Group Unblocking Acknowledgment Message	4-15 sec.
T21	Sending of initial Circuit Group Unblocking	Receipt of Receipt of Receipt of Circuit Group Unblocking Acknowledgment Message	1 min.
T22	Sending of Circuit Group ReSet	Receipt of Circuit Group ReSet acknowledgment	4-15 sec.
T23	Sending of initial Circuit Group ReSet	Receipt of Receipt of Circuit Group ReSet Acknowledgment	1 min.
T24	Sending of check tone	Receipt of backward check tone	<2 sec.
T25	Detection of initial continuity check failure	On expiry	1-10 sec.
T26	Detection of second continuity check failure	Detection of continuity	1-3 min.
T27	Receipt of CONtinuity Message indicating failure of repeat continuity check	Receipt of Continuity Check Request Message	>3 min.
T28	Sending of Circuit Query Message	Receipt of Circuit Query Response Message	10 sec.
T29	Not specified for U.S. Networks		
T30	Not specified for U.S. Networks		
T31	Release of ISDN User Part end-to-end signalling (Connectionless)	On expiry	>6 min.
T32	Sending of response to end-to-end connection request	Receipt of first end-to-end message from remote end	3-5 sec.
T33	Sending of INformation Request Message	Receipt of INformation Message	12-15 sec.

ISUP Timers (Continued)

Timer #	Initiation Cause	Termination Cause	Time-Out Value
T34	Sending of LoopBack Acknowledgment in response to receipt of Continuity Check Request	Receipt of COntinuity Message or RELease Message	10-15 sec.
T35	Not specified for U.S. Networks		
T36	Receipt of Initial Address Message indicating another segment to follow	Receipt of Receipt of unsolicited INformation Message	2-4 sec.
T37	When ISDN User Part availability test is started	Receipt of a message from the affected ISDN User Part	30 sec.
T _{ACC,r}	Receipt of ACC indicator	On expiry	5 sec.
T _{CCR}	Sending of Continuity Check Request Message	Receipt of LoopBack Acknowledgment Message	2 sec.
T _{CCR,r}	Receipt of initial COntinuity Message indicating failure	Receipt of Continuity Check Request Message	20 sec.
T _{CGB}	Receipt of Circuit Group Blocking Message	Receipt of Circuit Group Blocking or Circuit Group UNblocking Message	5 sec.
T _{CRA}	Sending of Circuit Reservation Acknowledgment Message	Receipt of Initial Address or RELease Message	10 sec.
T _{CRM}	Sending of Circuit Reservation Message	Receipt of Circuit Reservation Acknowledgment Message	3-4 sec.
T _{CVT}	Sending of Circuit Validation Test Message	Receipt of Circuit Validation Response Message	10 sec.
T _{EXM,d}	Sending of Initial Address Message to succeeding network	On expiry	Network dependent
T _{GRS}	Receipt of Circuit Group ReSet Message	Receipt of Circuit Group ReSet	5 sec.
T _{HGA}	Carrier Loss	Carrier Restoration	0-5 min.
T _{SCGA}	On failure of initial demand continuity check in SCGA group	On success of demand continuity check in SCGA group	0-2 min.
T _{SCGA,d}	On failure of demand continuity check in SCGA group	On expiry	5-120 sec.

Section 4


Signal Unit And Message Format

Signal Unit Formats

The SS7 uses only three packets (signal units) in transmission. The majority of the fields are identical in each of these units. We'll examine each unit, and hopefully, end up with an understanding of the ways in which the protocol manages to deliver information from node to node and to the final destination of the message. We'll also examine those things that are done to make this happen in a reliable and efficient fashion.

The Message Signal Unit (MSU)

We'll start with the most complex of the signal units. While it may seem strange to work from the complex to the simple, we do so with a purpose. Most of the fields that appear in the Message Signal Unit appear in the other two signal units also. Once you understand the MSU, you'll need only some brief explanation to understand the other two as well.



Message Signal Unit Fields

We'll examine these fields to determine the purpose of each and what kind of data belongs there. Notice the arrow indicating the direction of the data flow. This indicates the sequence in which the unit is assembled by the transmitting MTP and also the sequence in which the receiving MTP sees the data.

Simply SS7

The Flag - When dealing with a digital protocol, knowing where to start reading is of paramount importance. In the SS7 protocol the key to making the correct assumptions about the data lies in knowing exactly what data is being read. That, in turn, lies in knowing exactly where the data lies within the signal unit. For those reasons, there can be no confusion about where to start reading the message.

This is so important that the MTP can take no chances that it might start to read the data at some point other than the actual start of the signal unit. To avoid this, a unique eight bit code is placed at the beginning of the unit. This code is a byte with zeros at either end and six ones in the middle. There may be considerable data in the remainder of the signal unit. It is very likely that a zero followed by six ones and a zero will occur elsewhere in the signal unit. To ensure that there are no false flags, the sending MTP reads through the signal unit. Each time it reads five ones in a row, it inserts a zero. This is a procedure known as “bit stuffing.” The MTP then attaches the flag and transmits the message.

At the receiving end, the MTP sees the flag and begins its reading of the signal unit. Every time it sees five ones, it removes the following zero. In this way, the confusion of multiple flags is eliminated, and the signal unit ends up restored to its original form.

The original standards (CCITT) provided for the use of a second flag to be used at the end of the signal unit. The later ANSI standard saw no value in this, and instead, supports the use of a single flag at the beginning. In this way a single flag becomes both the beginning of one unit and the ending of the previous one. The result is a shorter signal unit and a higher rate (signal units per time period) of transmission.

Backward Sequence Number - From this point on, it will be helpful if you equate the term “backward” with “receiving node” and the term “forward” with “transmitting node.”. It is the receiving node which make changes to this value (the BSN). It does so when it is returning a signal unit to positively acknowledge the receipt of a unit or to make a negative acknowledgment of a unit. In the latter case, the MTP will usually also request that the message be retransmitted. This will become more clear when we examine the Forward Sequence Number.

Backward Indicator Bit - Once again, it is the receiving node that will make changes to this value (the BIB). It will change this bit to the opposite of Forward Indicator Bit in the same signal unit being used to send a negative acknowledgment back to the transmitting side. The transmitting side reads this changed bit state as a request for retransmission.

Forward Sequence Number - This time it is the transmitting node which makes changes to the value. The transmitting MTP maintains a numbering resource which provides cyclical and sequential values in the range of zero to one hundred twenty seven (0 - 127). It places the value into this field and then simultaneously transmits the signal unit and copies it into a retransmit buffer. This provides the receiving side with a value by which to refer to the signal unit.

Forward Indicator Bit - Once again, it is the transmitting node that deals with this value. On transmission it ensures that the Forward Indicator Bit matches the Backward Indicator Bit.

Error Correction - The standards support two methods of correcting errors. In one (the **Preventive Cyclic Retransmission Error Correction** Method) the messages are retained on the transmitting side until acknowledged. During every break in transmission (no messages to be sent) the transmitting side simply retransmits all messages that have not yet been acknowledged. This method is generally used only for satellite transmission.

The other method is the **Basic Error Correction** Method. Now that we have seen the first five Message Signal Unit (MSU) fields, we'll follow this method through its sequence to see what both the transmitting MTP and the receiving MTP need to do to ensure the delivery of good messages.

First the transmitting side uses its numbering resource to provide a value for the Forward Sequence Number. Then it transmits the MSU and sends a copy to its retransmit buffer. The receiving MTP, of course, monitors the incoming message (see Signal Unit Error Rate Monitor in Section 3). If no error occurs, it will send an acknowledgment before it has seen the entire series of 128 signal units (0 - 127) applied by the transmitting side. This needs to be done because the transmitting side will not apply a value to a new message if that number matches a value in its retransmit buffer. If this occurs, the transmit node simply stops transmitting and the MTP indicates a "link failure."

For the **acknowledgment**, the MTP uses whichever signal unit it would normally be returning to the transmitting node. Since the MTP normally reports link status periodically, this would commonly be a Link Status Signal Unit (LSSU). To make the acknowledgment, the receiving MTP takes the Forward Sequence Number for the last valid signal unit and copies it to the Backward Sequence Number field of the signal unit it is returning. It leaves the Backward Indicator Bit alone so that the Forward and Backward Indicator Bits are returned as received (both the same). When the signal unit arrives at the transmitting side, it is recognized as an acknowledgment without request for retransmission. The transmitting side now simply deletes from its retransmit buffer all signal units having the value of the Backward Sequence Number and all prior Sequence Numbers.

When the receiving MTP detects an error, it once again uses the next planned return signal unit and copies the Forward Sequence Number of the last valid signal unit into the Backward Sequence number. This time, it toggles (from 1 to 0 or from 0 to 1) the Backward Indicator Bit so that it is no longer the same as the Forward Indicator Bit. When this unit is received at the transmitting node, it recognizes the unequal Forward and Backward Indicator Bits as a request for retransmission. It deletes all signal units with a value equal to or less than that of the Backward Sequence Number and begins retransmission of all signal units beginning with the one that is one higher than that of the Backward Sequence Number. Transmission is halted until the retransmission is complete.

We have completed our discussion of the first five fields of the Message Signal Unit. We have also looked at the ways in which the MTP uses these fields for the purpose of error correction. Now we are ready to move on to the remaining fields in the Signal Unit.

Simply SS7

Length Indicator - This field may seem a little strange. For one thing, it doesn't do what it was originally intended to do. The original intention was for this field to indicate the number of octets of significant data which followed it. That data was located in the Service Information Octet and the Signalling Information Field.

All along you may have wondered why this system in current use is called Signalling System #7. Were there six predecessors? The answer is, that at least on paper, there were. The only deployed version was Common Channel Inter-Office Signalling System #6. Its limited deployment bore no resemblance to the nearly universal deployment of the SS7. In **CCIOSS #6**, MSUs were handled differently. For one thing, there was an attempt to limit the size of a signal unit. The Length Indicator was assigned six bits to indicate the amount of following data (up to 64 octets). For **CCIOSS #6** that was adequate.

Signalling System #7, however, allows up to 272 octets of data in the Signalling Information Field alone. The Signalling Information Field is where the significant data that represents the actual message is placed. Of what use is a length indicator that can only represent values from 0 to 63? The answer is that as a length indicator it is of little use, but as an identifier of Signal Unit type it is very helpful. Signalling System #7 uses three types of signal units.

The least complex signal unit is the Fill In Signal Unit (FISU) which contains no message or service data beyond the Length Indicator. Therefore, a length indication value of 0 identifies the Fill In Signal Unit.

Another, slightly more complex signal unit is the Link Status Signal Unit (LSSU). It has a Link Status Field which can contain one or two octets. Therefore, a length indication value of 1 or 2 identifies the Fill In Signal Unit.


Finally, the Message Signal Unit, as we have said, may contain much more data in the fields following the Length Indicator than can be counted using six bits. In that case, the counter simply goes as far as it can. If more than 64 octets are found, the value stays at 63 (0 to 63 range). The Signalling Information Field will never have less than 2 octets of data. Taken with the Service Information Octet, that means that the length indication will always be more than 2 if the signal unit is an MSU. The meaning then, of the Length Indicator value is as follows:

Length Indicator = 0	Fill In Signal Unit
Length Indicator = 1 or 2	Link Status Signal Unit
Length Indicator > 2	Message Signal Unit

Any missing data is easily assessed using the sophisticated algorithm found in the Check Bits, so there is now no loss resulting from the fact that the Length Indicator cannot accurately indicate the amount of data in message fields.

Spare - From the drawing you can see that the spare is simply a two bit field. Its only purpose, generally, is to keep the entire signal unit to an even number of octets. Each field which is not an octet (FSN, BSN) has its counterpart which brings the number to eight (FIB, BIB). The spare becomes the counterpart for the six bit Length Indicator. This is particularly helpful at specific times. For example, when the MTP is aligning (or restoring) a link, it sends Fill In Signal Units and looks for errors. With no significant data, the total data must be evenly divisible by eight. A single bit (or 2,3,4,5,6 or 7) lost or gained will change this no-remainder division and the unit can be adjudged invalid.

Signalling Information Octet - The Signalling Information Octet is apportioned into sub-fields of four bits, two bits and two bits. The drawing below illustrates this field. We will examine each of the sub-fields in greater detail.



Service ID - The Service ID field provides information about the type of message being sent. This provides level 3 with the necessary information for message distribution. Codes are provided for the user parts with the exception of TCAP. The reason for this is that TCAP messages are always appended to SCCP messages. An SCCP code along with a UNITDATA message type indicate a TCAP message. Other codes are provided when the message is a signalling network message dealing with management messages or with testing and maintenance messages.

Simply SS7

Sub-Service Field - This field contains bits for the indication of network in which the message sender is deployed. It also provides a spare for network dependent usage

Priority Field - The standards provide for two bits to be used as a spare. When operating within a national network, these bits are provided for national use. We have referred to it here as a priority field, but that is only one possible use. The two bits of the field can allow messages to be categorized in order of importance. The standards allow for this prioritizing to be network dependent. The categories shown in the drawing are simply those in use at some SSPs (Service Switching Points). Categories for other types of nodes would of course, be different. In the previous section you learned that one of the implementation dependent functions of the MTP was the indication of multi-level congestion, along with thresholds which could be set for the deliberate discard of messages under conditions of congestion.

If message priorities are implemented, they can be used to provide the MTP with an identification of the messages that can be discarded at each congestion level. Thus, at the lowest congestion level, the MTP can be implemented to discard the least significant messages. At each successively higher level of congestion, the MTP can discard increasingly important messages. The standards, however, do not allow for the discard of network management messages.

Network Indicator Field - This indicator is used by signalling message handling functions to determine which version of the User Parts to employ. Where a node operates significantly within the confines of a national standard (such as ANSI) this indicator will be set for National.


Signalling Information - Every field that we have looked at so far is designed to control the delivery and provide further specific information about the data contained in this field. This is the message field of the Message Signal Unit. In the U.S. this field never contains less than 2 octets of data, nor does it ever contain more than 272 octets. This allows the transmission of 265 octets of information along with a label. There may also be additional housekeeping information. Such information may be used at level 4 to link information blocks together (among other possible uses).

The drawing that follows will begin our examination of this field.

Routing Label - In the drawing there are two routing labels depicted. It should be made clear that both of these routing labels will not be found in the same MSU. The attempt made in the drawing is to illustrate an ANSI routing label and a CCITT (ITU-TS) routing label. The one to be found in an MSU will depend, of course, on which standard is being employed by the node sending the message. In North America, this will generally be ANSI.

For ANSI networks the routing label is identified as being the first 56 bits in the SIF (Signalling Information Field). ANSI provides for a network location to be identified by a 24 bit code. The first eight bits of this code is given the name of Network ID. Signalling Point Codes, in general, follow the general schema introduced by the telephone numbering scheme known as the North American Numbering Plan. In that plan, each succeeding value generally represents a smaller geographical area (area code, exchange) until the last number represents an individual phone line.

For the SS7, the distinction is hierarchical rather than geographical. That is the network identifier identifies a network to which a point code belongs. This can be a code indicating a broad general network, or it can be a reserved code used to indicate that the identity is assigned to a group of



commonly administered nodes which, together, do not qualify for full network status. With one of these reserved codes used as the network identifier, an indication is given that the next value (Network Cluster) will be used to identify the network to which the node belongs.

Thus the code can be used to identify smaller networks within a larger network and, finally, narrow down the addressing (Network Cluster Member) to address a single signalling point.

Simply SS7

The only reserved Network Cluster Member is the “empty” byte (eight zeros) which is used to identify an STP. The standards reserve this value for STP usage, but they don’t compel its use. What that means is that when you see a node with a Network Cluster Member value of zero, you know that the node is an STP. However, when you see a node with a Network Cluster Member code of other than zero, there is no guarantee that the node is *not* an STP.

The remaining field is the **Signalling Link Selection Code**. You may recall that different SCCP services require different handling of this code. SCCP services break down into four categories. These services are connectionless or connection oriented, each of which is further categorized as requiring or not requiring in-sequence delivery. The MTP level 3 normally rotates the SLS and each new code will direct the message to a new available link. When the SCCP requires it, the MTP stops rotating the code and each new message is directed to the same link, thereby guaranteeing in-sequence delivery.

For the ANSI routing label, the only thing remaining to be identified is the field labeled “User Data.” This is the actual message. The letter “n” represents the total amount of data in the SIF. Since the routing label is always 56 bits (7 octets), “n-7” represents the size of the message (total octets in the field, less the routing label).

The differences in the CCITT (ITU-TS) routing label are relatively minor. The names of the signalling point code fields are changed (*Zone* instead of Network ID, *Network* instead of Network Cluster, and **Signalling Point ID** instead of Network Cluster Member). The ITU-TS standard also elected to allot a different number of bits to two of the three portions of the code (3-8-3 instead of 8-8-8). Fewer bits were also allotted to the SLS (4 instead of 5, and the ANSI 1996 standard supports 8 bits for the SLS). And, finally, the computation of the message size based on the total data in the SIF changes because only four octets need to be subtracted from the total data.

One remaining field in the MSU is the **Check Bits** field. This is an algorithm that allows the MTP to determine whether the number of bits transmitted is the same as the number of bits received.


The last field is the trail flag. The trail flag is not used in the U.S. Therefore, in ANSI networks the MSU ends with the Check Bits.

Before we go on to other Signal Units (which should seem simple now) we’ll take a brief look at the format of some of the most common SS7 messages. It is well beyond the intended scope of this book to examine all of the message types, all of the parameters, etc., that can be sent on the SS7. We *do* need to look at some of the data that might appear in the message portion of the SIF.

SS7 Message Examples

Initial Address Message

Certainly among the most common messages to be seen on the SS7 network is the ISUP *Initial Address Message (IAM)*. In the next drawing we have illustrated an ANSI **IAM** including the coding of the parameter name and the number of octets (or the range of octet numbers) in the complete coding of the parameter.




This drawing also illustrates a field that has not been discussed elsewhere. This field is known as the **Circuit Identification Code (CIC)**. It follows the Routing Label and is used with ISUP messages as a way of identifying the chosen voice circuit to the next switch, thereby ensuring accurate circuit connections.

TCAP Message with Invoke Component

This drawing illustrates A TCAP message to Invoke an operation. In the drawing the Transaction, Component and Parameter Set portions are separated by various shadings. The component shown at the bottom is there simply to indicate that the message may contain numerous component portions. Package Type Identifiers are coded as follows:

Unidirectional	1 1 1 0 0 0 0 1	Response	1 1 1 0 0 1 0 0
Query With Permission	1 1 1 0 0 0 1 0	Query W/O Permission	1 1 1 0 0 0 1 1
Conversation With Permission	1 1 1 0 0 1 0 1	Conversation W/O Permission	1 1 1 0 0 1 1 0
Abort	1 1 0 1 0 1 1 0		



The Link Status Signal Unit (LSSU)


As the name implies, the Link Status Signal Unit is used by the MTP to provide status information to the signalling points at either end of the link. These indications are placed in the one or two octet field called the Link Status field. Status indications relate to the current mode under which the MTP is monitoring the link. Recall that our earlier discussions about functionality mentioned two error rate monitors. The first was the AERM (Alignment Error Rate Monitor) employed during link alignment.

At the beginning of the alignment process (while in the **Out-of-alignment** state) the MTP provides a status indication of “O.” When the alignment has proceeded to the proving period, the status of “N” (for **N**ormal proving period) or the status of “E” (for **E**mergency proving period) is sent.

By the way, as long as we are on the subject of Normal vs. Emergency alignment you may have wondered how the selection is made. The node itself does the selection, either as required by the application or by pre-configuration.

The format of the Link Status Signal Unit is illustrated below.

*Trail flag not used in U.S.



When a link is first powered up, and before the “O” status indication, the “OS” (**Out-of-Service**) status is sent. The same status is sent any time the link can neither receive nor send MSUs for any reason other than a Processor Outage.

The sending of these status indications is somewhat of an interactive process on the part of the MTPs at either end of the link. For example, if the MTP at one end of the link has an “N” proving status it will send “N”. If the other end sends an “E” status, the side with the Normal status will not change that status. However, it will not cause a delay in the link realignment by performing the longer proving period. It will, instead, perform the shorter “E” proving period to be consistent with the realignment being performed by the other side. Even as it does so, it still sends “N”.

While the MTP is monitoring normal transmission (using the SUERM or Signal Unit Error Rate Monitor) it will send a status of “B” (link **B**usy). That same MTP stops sending either positive or negative acknowledgments. The MTP at the other end of the link sets a long timer to await the clearance of the congestion. It also resets the “excessive delay of acknowledgment” timer every time it receives (periodically during the congestion period) a new indication of “B.”


A status of “PO” is sent when the MTP detects a problem in delivering messages to levels 3 or 4. The same MTP then begins message discard. When the opposite MTP receives “PO”, it informs level 3 and places Fill In Signal Units on the link.

The Fill In Signal Unit (FISU)

If you pull a link out of the link port, the MTP will detect a “failed link.” This is due to the fact that the MTP continues to monitor Signal Units and now has no data with which to deal. While nodes are communicating, there are many times when no messages are being sent. If there were no data to read, the MTP would have to conclude that it had detected a “failed link”. For this reason, the MTP cannot allow any time period in which the link carries no data. The transmitting MTP fills all such blanks with **Fill In Signal Units**.

The format of the Fill In Signal Unit is illustrated below.

*Trail flag not used in U.S.



As you can see, the format just became even more simple. The MTP can continue to monitor for valid flags, octet integrity and package size. It simply finds no data to send to level 3. On the transmitting side, no changes are made. The FSN applied to the last message simply reappears in each FISU until a new message is sent.

There is another time when the FISU becomes handy. During the alignment procedure the MTP uses the AERM to monitor the link. Until the link is placed in an In Service state, there are no messages being sent. Therefore, the MTP puts FISUs on the link and monitors them during the proving period.

A Backward Glance

We have reached the end of our look at Signalling System #7. The scope of this book was never intended to cover every conceivable aspect of the standards. That can be done only by studying the standards.

Our hope here was to give you a broad exposure to numerous aspects of the standards while, at the same time, clarifying some of the mystery. If, at this point, you feel you have profited from reading this book, our goal has been reached.

Where should you go next? We'll mention only one source. The **McGraw-Hill** publication entitled **Signaling System #7** by **Travis Russell** is, in our opinion, the best available in the public domain. While most of Simply SS7 resulted from grueling hours with the standards themselves, we must confess that we kept the Russell book handy if only to get his “slant” on the information. If your next step takes you beyond the scope of Simply SS7, having your own copy of Signaling System #7 is a wonderful way to go. Beyond that, we can only recommend that you get your own copy of the standards.

Acronyms

A Listing of Common SS7 And Other Telecommunications Acronyms

Acronym	Definition
ACD	Automatic Call Distributor
ACG	Automatic Code Gapping
ACM	Address Complete Message
AFR	Automatic Flexible Routing
AHT	Average Handle Time
AIN	Advanced Intelligent Network
AIOD	Automatic Identified Outward Calling
AMA	Automatic Message Accounting
AMATPS	AMA Teleprocessing System
AMP	AIN Maintenance Parameter
ANI	Automatic Number Identification
ANM	Answer Message
ANSI	American National Standards Institute
API	Application Programming Interface
ARP	Address Resolution Protocol
ASA	Average Speed of Answer
ASN.1	Abstract Syntax Notation 1
ASE	Application Service Element
ATB	All Trunks Busy
ATP	Acceptance Test Procedure
AUI	Attachment Unit Interface
AW	Admin Workstation
BAF	Bellcore AMA Format
BBG	Basic Business Group
BCC	Bellcore Client Company
BCD	Binary Coded Decimal

Acronym	Definition
BCI	Backward Call Indicators
BCLID	Bulk Calling Line Identification
BCM	Basic Call Model
BER	Basic Encoding Rules
BG	Business Group
BGID	Business Group Identification
BRI	Basic Rate Interface
BSN	Backward Sequence Number
CAC	Carrier Access Code
CAP	Competitive Access Provider
CC	Call Control
CCA	Call Control Adjunct
CCITT	Consultative Committee on International Telephone & Telegraph
CCS	Common Channel Signalling
CDAR	Customer Dialed Account Recording
CDP	Customized Dialing Plan
CDPD	Cellular Digital Packet Data
CED	Call Entered Digits
CGB	Circuit Group Blocking Message
CGU	Circuit Group Unblocking Message
CIC	Carrier Identification Code
CIDS	Calling Identity Delivery & Suppression
CLID	Calling Line ID
CLLI	Common Language Location Identification
CMC	Cellular Mobile Carrier
CMS	(AT&T'S) Call Management System
CNAB	Call Name Delivery Blocking
CO	Central Office
COT	Continuity Test Message
CPC	Call Processing Control

Acronym	Definition
CPE	Customer Premises Equipment
CPG	Call Progress Message
CR	Conditional Requirement
CRA	Circuit Reservation Acknowledgment Message
CRM	Circuit Reservation Message
CS-1	Capability Set 1
CSC	Circuit Supervision Control
CSU	Channel Service Unit
CT	Call Type
CVR	Circuit Validation Response Message
CVT	Circuit Validation Test Message
DACS	Digital Access Cross-Connect System
DCE	Data Circuit Equipment
DMP	Device Management Protocol
DN	Dialed Number
DNIS	Dialed Number Identification Service
DP	Dial Pulse
DPC	Destination Point Code
DSVD	Digital Simultaneous Voice and Data
DTE	Data Terminal Equipment
DTMF	Dial Tone Multifrequency
DUP	Data User Part
DXI	Data Exchange Interface
EA	Equal Access
EADAS	Engineering & Administration Data Acquisition System
EADASN	EADAS Network Administration
EAEO	Equal Access End Office
EAMF	Equal Access Multifrequency
EBCDIC	Extended Binary Coded Decimal Interchange Code
EDP	Event Detection Point

Acronym	Definition
EIA	Electronic Industries Association
EIR	Equipment Identification Register
EKTS	Electronic Key Telephone Service
EMS	Event Management Service
EO	End Office
ESN	Electronic Serial Number
EXM	Exit Message
FCS	Frame Check Sequence
FISU	Fill-in Signal Unit
FRAD	Frame Relay Access Device
FRL	Facility Restriction Level
FUNI	Frame User Network Interface
FSD	Feature Specific Document
FSN	Forward Sequence Number
FSS	Facility Selective Service
FTE	Full Time Equivalent
FTP	File Transfer Protocol
FX	Foreign Exchange
GN	Generic Name
GRS	Group Reset Message
GSC	Gateway Switching Center
GSM	Group Special Mobile
GTT	Global Title Translations
GTV	Global Title Value
GUI	Graphical User Interface
HDLC	High Level Data Link Control
HFC	Hybrid Fiber Coaxial Cable
HLR	Home Location Register
IAM	Initial Address Message
IC	Interexchange Carrier

Acronym	Definition
ICP	Intelligent Call Processing
ICR	Intelligent Call Router
IDLC	Integrated Digital Loop Carrier
ISP	Intermediate Service Part
IDT	Integrated Digital Terminal
INR	Information Request Message
IP	Intelligent Peripheral or Internet Protocol
IPC	Interprocess Communication
IPI	Intelligent Peripheral Interface
ISP	Intermediate Service Part
ISPC	International Signalling Point Code
ISDN	Integrated Services Digital Network (Used with CPE)
ISDNUP	ISDN User Part
ISUP	ISDN User Part (Used with Circuit Oriented)
IWX	Interworking Function
IXC	Interexchange Carrier
LAA	Longest Available Agent
LAN	Local Area Network
LATA	Local Access & Transport Area
LI	Length Indicator
LSSU	Link Status Signal Unit
LSSGR	LATA Switching & Signalling Generic Requirements
LOCREQ	Location Request
MAP	Mobility Application Part
MBG	Multi-switch Business Group
MCC	Mobile Country Code
MIN	Mobile Identification Number
MGW	Mini-Gateway Prototype
MLHG	Multi-line Hunt Group
MMI	Man-Machine Interface

Acronym	Definition
MSC	Mobile Switching Center
MSISDN	Mobile Station ISDN Number
MSU	Message Signal Unit
MUX	Multiplexor
MTP	Message Transfer Part
NAA	Next Available Agent
NCA	Non-Call Associated
NCP	Network Control Point
NDC	National Destination Code
NIC	Network Interface Controller
NNI	Network Node Interface
NPA	Numbering Plan Area
NSP	Network Services Part
ODBC	Open Database Connectivity
OE	Office Equipment
OMAP	Operations & Maintenance Application Part
OPC	Origination Point Code
OPI	Open Peripheral Interface
OS	Operations System
OSI	Open Systems Interface
OTGR	Operations Technology Generic Requirement
PBX	Private Branch Exchange
PCS	Personal Communications Services
PG	Peripheral Gateway
PIC	Point In Call
PIM	Peripheral Interface Manager
PPP	Point-to-Point Protocol
PRI	Primary Rate Interface
PROFREQ	Profile Request
PSN	Alternative to PSTN (Public Switched Telephone Network)

Acronym	Definition
PSTN	Public Switched Telephone Network
REGNOT	Registration Notification
RISC	Reduced Instruction Set Computing
ROUTREQ	Routing Request
SANC	Signalling Area Network Code
SCCP	Signalling Connection Control Part
SCP	Service Control Point
SDLC	Synchronous Data Link Control
SEP	Signalling Endpoint
SF	Status Field
SI	Service Indicator
SIF	Signalling Information Field
SIO	Signalling Information Octet
SLC	Signalling Link Code
SLIP	Serial Line Internet Protocol
SLS	Signalling Link Selection
SLP	Service Logic Program
SMDS	Switched Multimegabit Digital Service
SMS	Service Management System
SN	Services Node
SNA	Systems Network Architecture
SNM	Signalling Network Management
SNMP	Simple Network Management Protocol
SNT	Signalling Network Testing
SP	Signalling Point
SPC	Signalling Point Code
SPID	Service Provider Identifier
SPM	Signalling Point Manager
SQL	Structured Query Language
SPR	Signalling Point w/SCCP Relay

Acronym	Definition
SPRC	Signalling Procedure Control
SRTC	Subrate Channel
SS7	Signalling System 7
SSF	Sub-Service Field
SSN	SubSystem Number
SSP	Service Switching Point
STP	Signalling Transfer Point
SU	Signal Units
TA	Technical Advisory
TC	Transaction Capabilities
TCAP	Transaction Capabilities Application Part
TCM	Traveling Class Mark
TCP/IP	Transmission Control Protocol/Internet Protocol
TDM	Time Division Multiplexor
TDP	Trigger Detection Point
TLDN	Temporary Local Directory Number
TR	Technical Reference
TUP	Telephone Users Part
UDP	User Datagram Protocol
UDT	Unitdata
UDTS	Unitdata Service
VAD	Voice Activated Dialing
VANC	Voice Activated Network Control
VLR	Visitor Location Register
VPN	Virtual Private Network
VRU	Voice Response Unit
WAN	Wide Area Network
WATS	Wide Area Telephone Service
XUDT	Extended Unitdata
XUDTS	Extended Unitdata Service

Index

Numerics

40% engineering rule	12
800 numbers	18
900 numbers	20

A

Abort (TCAP)	65
Access Links	12, 13
ACM	70
ACM - Address Complete Message	72
address complete message	71
adjacent nodes	35
Advanced Intelligent Network	28
Aligned	44
Aligned/ready time-out	44
Alignment Error Rate Monitor (AERM)	42
ANM	71
ANM - Answer Message	73
ANSI	7
answer message	71
Applications	35

B

Backward Indicator Bit	84
Backward Sequence Number	84
Base Station System	26
Basic Error Correction	41, 85
Basic Services	69
bit stuffing	40
Bridge Links	14
busy signal	6
Busy Redelivery Interval	44

C

Call Processing Control (CPC)	69
call setup	70
CCITT	7
CCSSO (Common Channel Signalling Switching Office)	17
ChangeBack Acknowledgment (CBA)	49
ChangeBack Declaration (CBD)	49
ChangeOver Acknowledgment (COA)	49

ChangeOver Order (COO)	49, 50
Check Bits	90
Circuit Supervision Control (CSC)	69
Component ID	67
Component ID Identifier	67
Component ID Length	67
Component Length	67
Component Portion Fields	66
Component Sequence Identifier	66
Component Sequence Length	66
Component Type Identifier	66
Concerned Points	60
congestion abatement	51
congestion onset	51
Controlled Rerouting	48
Conversation With out Permission	65
Conversation With Permission	65
Cross Links	13
CRP	31
CRP (Customer Routing Point)	20

D

Destination Point Code	59, 75
Diagonal Links	15
discard onset	51

E

E link status	93
End Office	2
Error Code Field	67
Error Code Identifier	67
Error Code Length	67
Return Error Component	67
Excessive Delay of Acknowledgment	44
Extended Links	15

F

The Fill In Signal Unit (FISU)	94
Flag	84
flags	40
Flow Control	42
flow control	51
Forced Rerouting	48

Index	Simply SS7
Forward Indicator Bit	84
Forward Sequence Number	84
four priorities	51
Fully Associated Links	16
G	
Global Title Translations	59
GSM	26
H	
Header Portion Fields	75
HLR	31
Home Location Register	27
I	
IAM	70
IAM - Initial Address Message	72
Initial Address Message	90
Intelligent Network	28
Interexchange Carrier	3
International	7
International Telecommunications Union	7
Invoke Component	66
IP	31
IP (Intelligent Peripheral)	22
IS41	26
ISUP	56, 69
ISUP Message Structures	74
ISUP Timers	78
ITU-TS	7
L	
layered protocol	36
Length Indicator	86
Link Failure	44
link management	49
linkset	46
The Link Status Signal Unit (LSSU)	93

M

Management Inhibiting	49
Mandatory Fixed Parameters	76
Mandatory Variable Parameters	76
message discrimination	53
message distribution	53
The Message Signal Unit (MSU)	83
Message Signal Unit Fields	83
Message Transfer Part	37
Message Type	76
Mobile Identification Number	27
MSC	31
MSC (Mobile Switching Center)	25
MTP Level 1	39
MTP Level 2	39
MTP Level 2 Functionality	43, 45
MTP Level 2 Timers	44
MTP Level 3 Timers	54
MTP Restart	48

N

N link status	93
N-Coord Confirmation	60
N-Coord Request	60
N-Coord Response	60
Network	90
Network Cluster	89
Network Cluster Member	89
Network ID	89
Network Indicator Field	88
Normal Alignment	41
North American Numbering Plan	18
Not aligned	44

O

O link status	93
off hook	5
OMAP (Operations, Maintenance and Administration Part)	63
on hook	6
Operation Code	68
Operation Code Identifier	68
Operation Code Length	68
Optional Parameters	76

Origination Point Code	75
------------------------	----

P

Parameter Set Identifier	68
Parameter Set Length	68
Parameters	68
Parameters Portion Fields	76
PO	93
Preventive Cyclic Retransmission Error Correction	41
Preventive Cyclic Retransmission Error Correction Method	85
Priority Field	88
Problem Code Field	67
Problem Code Identifier	67
Problem Code Length	67
Proving	44
PSTN	2

Q

Query With Permission	65
Query Without Permission	65

R

Reject Component	67
REL	71
REL - Release	73
Release	71
Release Complete	71
Response	65
Return Response Component	67
RLC	71
RLC - Release Complete	73
Route Management	48
route set	47
Routing Label	89
Routing Management	50

S

SCCP Service Types	58
SCCP Specialized Routing Functions	58
SCCP Subsystem Management	60
SCCP Timers	62
SCP	31

SCP (Service Control Point)	18
Service Creation Element	28
Service ID	87
Service Management System	29
Signal Unit Delimitation and Alignment	40
Signal Unit Error Correction	41
Signal Unit Error Detection	40
Signal Unit Error Rate Monitor (SUERM)	42
Signalling Connection Control Part	37
Signalling Information	88
Signalling Information Octet	87
Signalling Link Activation	49
Signalling Link Alignment	41
Signalling Link Changeback	49
Signalling Link Changeover	49
Signalling Link Selection	46
Signalling Link Selection Code	90
Signalling Link Test	49
Signalling Link Test Acknowledgment (SLTA)	49
Signalling Link Test Message (SLTM)	49
Signalling Point ID	90
Signalling Procedure Control (SPRC)	69
Signalling-Route-Set-Test	50
Signalling-Route-Set-Test (SRST)	50
SLS Code	75
SP	13
SS7 stack	36
SSA	61
SSN	58
SSP	31
SSP (Service Switching Point)	17
SST	61
STP (Signalling Transfer Point)	11
Sub-Service Field	88
subsystem number	58
Subsystem-Allowed	60
Subsystem-Prohibited	60
Subsystem-Status-Test	60
Supplementary Services	69

T

TCA	50
TCAP	63
TCAP Message with Invoke Component	92
TCAP Package Type	65

Telephone Users Part	38
TFA	50
Total TCAP Message Length	65
TRA (Traffic Restart Allowed)	48
Traffic Management	45
Transaction ID	66
Transaction ID Identifier	65
Transaction ID Length	65
Transaction Portion	64
Transactions Capabilities Application Part	37
Transfer Cluster Prohibited (TCP)	50
Transfer Cluster Restricted (TCR	50
Transfer Cluster Restricted (TCR)	48
Transfer Prohibited (TFP)	48, 50
Transfer Restricted (TFR)	48, 50
Transfer-Allowed	50
Transfer-Cluster-Allowed	50
Trunk CIC Code	75
TRW (Traffic Restart Wait)	48
 U	
 Unidirectional	 65
 V	
 Visitor Location Register	 28
VLR	31
 W	
 wireless network	 24
 Z	
Zone	90

SS8 Networks Software Products

SS8 offers a broad range of solutions for Telecom manufacturers and Service Providers. Current products fall into three categories. First there are three **Development Platforms**, each offering a range of different and overlapping strengths for the development of Telecom applications.

Next there are plug and play wireless **Short Messaging** and **Over The Air** servicing products.

Then there are numerous solutions which allow providers to reach compliance with the Federal mandates of the **CALEA** Act.

Finally, a **Signalling Gateway** will be making its way to market in the spring of 2002.

On the following pages you will find some information on each of these. For additional information write, call, FAX, surf, or e-mail us at any of these addresses:

SS8 Networks East

2 Enterprise Dr., Shelton, CT USA 06484
203.925.6121 (tel) • 203.926.2664 (fax)
<http://www.SS8.com>
<http://www.SS7.com>
info-adapts@SS8.com (e-mail)

SS8 Networks Software Products

Development Platforms

Distributed7™ - Take years of experience and evolution of our famed AccessMANAGER™ SS7 development platform, add even more development, administration and maintenance tools, and place it all in a multi-host distributed SS7 User Part environment and you have Distributed7™. Standalone or Distributed, Front End or Back End, Distributed7™ takes advantage of User Part or application redundancy and makes efficient use of SS7 links by allowing multiple processes to appear on the SS7 network as a single Signalling Point.

Connect7™ - With products like AccessMANAGER™ and Distributed7™ why add a 3rd development platform? The answer is that many equipment manufacturers need a low-cost SS7 network interface product to create applications embedded within their own products. While low cost, such interfaces also need to provide high performance and high reliability. Connect7™ is a “stack on a card” approach that offers portability to any operating system, high performance and redundant interface card reliability. Utilities aid the development process and compilable source code is included to provide Application templates.

SS8 Networks Software Products

Wireless Products

SMserver™ - When you look for a Short Message Center, look for one that will deliver messages from *any* alphanumeric source. Make sure it delivers VM waiting messages, paging messages, e-mail messages, web delivered messages, operator messages and subscriber originated messages. Then be sure it can deliver schedule information, weather reports, stock market quotes, horoscopes or anything else you can provided in alphanumeric format. Next, be sure that you have enough development control that you can interface with any possible future alphanumeric origination system.

After that be certain it supports GSM 900, 1800, 1900, and iDEN as well as IS-41 networks including Is-136 TDMA and IS-637 CDMA. Then ensure that it will deliver mail in batches or by broadcast or to subscriber selected groups. Also see to it that it makes periodic delivery attempts and that the period is variable by the number of attempts made and even by the errors that have been returned as the cause of delivery failures.

Finally, check to be certain it prepares billing records and traffic records. When you have found all this and more, you'll be looking at SMserver™ from EDS.

OTAserver™ - Want to cut fraud loss? Simple. Just change the A-key of mobile stations instantly. Want to eliminate the cost of using a third party service bureau for the activation and support processes? Just as simple. Control Over-The-Air service provisioning and Over-The-Air activation yourself. OTAserver™ supports both CDMA and TDMA networks simultaneously and it's scalable, supports real time billing, and collects detailed operational statistics. Did we mention that it is easy to deploy and administer? Or that it offers scalable fault tolerance?

SS8 Networks Software Products

CALEA Compliant Intercept Products

As the scramble to meet CALEA requirements intensifies, every involved manufacturer has tried to make certain that its product will be ready. But for the Service Provider most of these have missed one important factor.....centralization. SS8 products provide the user friendly administrative platforms that accept the warrant information, send provisioning information backward into the network, control the initiation and termination of the surveillance and control the distribution of the Call Data and/or Call Content to the warrant specified Law Enforcement Agencies using approved J-Standard protocols. Even in networks where all switches are capable of making the required deliveries, SS8 products can cut costs and vastly improve control. Various combinations of the basic products shown below are also possible.

CALEAserver™ - Our basic intercept product centralizes delivery of both the Call Data and the Call Content. An administrative platform provides easy input of data through data entry screens, sends provisioning data backward into the network, controls an integrated intermediate voice switch for delivery of Call Content and delivers properly formatted Call Data to the LEAs specified in the warrant.

CDCmanager™ -For those who prefer to let their own network switches deliver the Call Content, this SS8 product is an ideal way to deal with the question of managing the Call Data Channels. All of the capabilities of CALEAserver™ are available in this product except for those which manage and control an intermediate surveillance content delivery switch. Warrant Administration is still centralized and the platform delivers properly formatted data on the **Call Data Channels**.

EINmanager™ - In some networks the Call Data comes from an IP network and the Call Content is VOIP. This product collects and distributes this data in the same way that the CALEAserver™ collects, converts and delivers SS7 data and PSTN voice content. All of the Administrative aspects of CALEAserver™ are here as well.

SS8 Networks Software Products

Signaling Gateway / Signaling Gateway Client

Our latest entrant combines our years of experience with SS7 networks and a clear view of the inherent savings of operating signaling points in a IP network.

Signaling Gateway™ - We began with our best development platform and added the SIGTRAN capabilities needed to translate, send and receive SS7 messaging to/from the SS7 network and to/from an IP network. The platform allows us to operate in a distributed environment with all the built in administrative, management, maintenance, and statistical tools of the world's best SS7 development platform with all the work of developing a gateway already done.

Signaling Gateway Client™ - For those who would like to develop their own SS7 signaling point to operate within an IP network, we offer the natural partner to our Signaling Gateway. This product allows the newly developed application to use the necessary SIGTRAN tools to communicate with the Signaling Gateway for the purpose of sending, receiving and generating SS7 signaling

Link Concentration

The problem with SS7 links is that they are costly to acquire and costly to maintain. Worse yet, redundancy requirements often mean that networks end up paying for far more link capacity than they need.

Easy7™ - Our answer to this problem is a link concentrator, once again operating on a platform with all the administrative, management, maintenance, and statistical tools of our best development platform. But all the development has been done. The TSP brings all of its own SS7 links to Easy7, and Easy7 connects into the network at large with fewer links. That means fewer links need to be paid for, and payback times are often astonishingly short. Full redundancy is maintained and lowered operating costs result. We'll be happy to help you analyze the link needs of your network and tell you how much you'll save.