This text was obtained from:
Proceedings of the NATO Advanced Study Institute
on Computer Communication Networks pp 373-96
Sussex, United Kingdom

September 9-15, 1973

THE SITA NETWORK

G.J. Chretien W.M. Konig J.H. Rech

Société Internationale de Télécommunications Aéronautiques

(SITA) Neuilly France

INTRODUCTION

1.1. SITA (Société Internationale de Télécommunications Aéronautique), a cooperative company founded in 1949, embraces the majority of the international air carriers (more than 160). It provides to its members a worldwide message switching network.

1.2. Initially the network consisted of manual (torn-tape) centres, interconnected by low speed circuits (50, 75 Bauds, 60, 30, 15 words per minute, asynchronous). The Airline terminal equipment (teleprinters. Telex) was connected to the SITA manual centres, thus enabling airline messages to be exchanged via nodes of the SITA network, with consequent reduction in costs to the airlines by their sharing of communications facilities.

1.3. With the rapid development of the Air Transport Industry, the airline communications needs became increasingly important and thus the SITA network expanded very quickly, by 1963 covering the world. Network development was not, however, restricted to geographic extension; in 1963 a number of the busiest manual centres were replaced by semi-automatic systems, and three years later, due to the continuing steady increase of traffic volumes, SITA equipped the Frankfurt centre with its first computer system to perform the message switching functions. Then, in 1969, SITA began replacing the other most heavily loaded centres (Western Europe and New York) with computer systems and established a computer communication data network by interconnecting these centres with voice grade circuits (medium speed). This network, called the High Level Network, performing the task of block switching, was interfaced at that time with the rest of the network composed of manual centres. This step was soon followed by the automation of other manual centres using what are in SITA terminology called satellite processors. These stand-alone computers act as concentrators of airline teleprinter traffic and controllers of airline CRT terminals, each of them connected to one High Level Centre by medium speed circuits. By mid-1973, the SITA network comprised 150 centres including 8 high level centres and 21 satellite processors. The 29 automated centres will be referred to as the SITA medium speed network (see figure 1).

In the remaining part of the network, referred to as the low speed network, there is one other centre (Hong Kong) equipped with computer systems and connected to the overall network by means of low speed circuits. A number of other European centres are equipped with multiplexers (1200 bits per second: 24 or 12 low speed circuits) and included in the low speed network.

1.4. The traffic carried at present by the SITA network can be subdivided into two main types, not taking into account special messages necessary to guarantee safe and error-free transmission. These two main types of traffic are telegraphic and conversational.

1.4.1 Telegraphic Traffic (or Type B)

These are messages generated under a given (ATA/IATA) format, using CCITT No. 2 code. They are one-way single or multi-addressed, destined to and generated by airline teleprinters or computers.

Such messages may be sent or received also by the local Telex networks.

The security requirements imposed on these messages are stringent, i.e. the SITA network is fully responsible, (no loss or undetected duplication and capability for message retrieval).

On the other hand, the transmission time constraints are not very severe, i.e. in the order of several minutes. The average length of such messages is around 200 characters.

1.4.2 Conversational Traffic (or Type A)

This is Query/Response traffic. The Query is generated by an agent set (CRT) in an airline office connected by medium speed lines to the SITA network, then sent via the SITA medium speed network to the parent reservations computer. There the Query is processed, a Response is generated and sent back to the enquiring CRT. The response time required on such a transaction is around three seconds.

At the same time, the degree of security demanded for these messages is low compared to that for telegraphic traffic, as, due to the nature of such messages, an operator at an agent set terminal will re-enter his Query should there be no reply within a few seconds. Average message lengths are in the order of 30 characters for a Query and 120 for a Response. Used for these messages are the CCITT No. 5 and the IBM BCD codes.

1.4.3 Comparing the time constraints imposed on the two types of traffic. Type A has priority over Type B.

2. THE SITA MEDIUM SPEED NETWORK

This comprises the SITA High Level Network and its satellite processors (see Figure 1).

2.1. The High Level Centre interfaces

Each High Level Centre (HLC) will interface with the following systems, exchanging with them messages or blocks, (see Figure 2):

a) other high level centres, through medium speed lines;

b) airline reservation computers, by means of medium speed lines (2.4, 4.8 kilobits);

c) satellite processors (SP) via medium speed lines (2.4, 4.8 kilobits);

d) teleprinters, via low speed lines (50, 75 Bauds);

e) local Telex networks, through low speed lines (50, 75 Bauds);

f) ARQ, 60, 30, 15 w.p.m.

In each HLC there are High Level functions and Low Level functions. The High level functions allow communication between any pair of HLCs via the interconnecting links. By definition, the High Level Network (HLN) consists of the HLCs with only their High Level functions and the Interconnecting links. The remainder of the SITA network is referred to as the Low Level Network (LLN).

2.2. The Satellite Processor Interfaces

Each Satellite Processor (SP), (see Figure 2), interfaces with the following systems:

a) its parent HLC by means of medium speed line (2.4, 4.8 kbps);

b) teleprinters, via low speed lines (50, 75 bauds);

c) local TLX network;

d) Agent sets via medium speed lines.

3. THE HIGH LEVEL NETWORK

Physical description

The network is at present composed of eight high level centres interconnected by eleven full duplex medium speed lines. The centres are located in Amsterdam, Brussels, Frankfurt, London, Madrid, New York, Paris and Rome, (see Figure 1). All circuits are operated at 4.8 kilobits, except for those connected to the Brussels Centre, but this is a transitional situation.

3.1 Principles of the High Level Network

The SITA High Level Network (HLN) follows a block switching principle.

Each message which has to be transmitted over the HLN is received first by an HLC called HLC of entry. The message is then subdivided into blocks with an appropriate header. Each of them is now sent independently towards the HLC of exit which will reassemble the message before passing it to the Low Level Network. An HLC of transit, that is a High Level Centre which receives a block while not being the HLC of exit, will perform only switching functions in accordance with the address in the block header.

3.2 High Level User Requirements

The requirements of the SITA users on the network are two-fold:

i) Low transit time, i.e. the time to transmit a block through the network.

ii) High availability, i.e. the monthly average downtime of the paths between any two high level centres must be as small as possible. For the HLN, the mean transit time is in the order of 700 milliseconds.

As far as the availability is concerned, continuous service on a 24-hour basis is expected and any downtime in a high level centre should not exceed 2 hours monthly. For this reason:

a) each high level centre is equipped with two or three computer systems, one of them being in a standby mode while the other(s) is (are) performing the necessary switching functions. A typical configuration of an HLC is given by figure 3.

b) between any two high level centres there are at least two independent paths (i.e. paths having no links or centre in common). Thus each HLC is connected to at least two other HLCs.

c) The maximum distance between any two centres is 3 links.

3.3 Organisation of the Network

Messages received by an HLC of entry for transmission on the HLN are segmented into blocks of a maximum length of 240 characters (7+ parity bits). The upper limit for the number of blocks per message is 16. In order to guarantee high transmission performances across the HLN one needs to provide:

a) secure transmission of blocks between two adjacent HLCs by a link control procedure;

b) addressing information for each block;

c) message reassembly for multi-block messages;

d) high security for type B messages.

3.4 Link Control Procedure

A synchronous full duplex link control procedure over one circuit (point-to-point symmetric for each direction) on which the two centres send their blocks continuously is used. Each data block has an envelope which comprises a sequence number and a block check character (longitudinal, transversal parity). A block, to be accepted as valid by the receiving centre has to be in the proper sequence and must have a correct block check. In this case an acknowledgement is sent, otherwise a negative acknowledgement is generated. Furthermore, the sending centre initiates a timer each time it sends a block. This timer is killed when the acknowledgement of the block is received.

Otherwise at its expiration, the sending centre repeats all blocks following the first unacknowledged block. Type A blocks have priority over Type B blocks, realised by a priority indication in the envelope. Link Control blocks (e.g. acknowledgement, negative-acknowledgement, etc.). have the highest priority at the output queues in each HLC.

Note: The subdivision of messages into blocks serves a dual purpose:

a) to increase the line efficiency, i.e. to increase the ratio of the total number of data characters to be transmitted over the total number of characters effectively transmitted (data, envelope, link control messages, repetition due to erroneous transmission) in the case of a line with a high error rate, e.g. 10^{-4} bit.

b) to avoid long waiting times in an output queue for Type A blocks.

3.5 Addressing

In order to be routed through the HLN, each block needs the address of the HLC of Exit. The blocks are switched through core memory in each transit HLC and transmitted in a store and forward mode according to paragraph 3.4.

3.6 Reassembly of messages

Since blocks of any one message are sent independently and possibly via different paths, each message needs to be reassembled in the HLC of exit. To achieve this, each block envelope contains a sequence number per Entry/Exit pair which means that the address of the entry centre is also needed in the envelope. To properly control the sequencing of blocks for reassembly, special "end to end control messages" are used.

3.7 Security Requirements

Type B messages have very high security requirements. Therefore, to avoid any loss in case of HLC CPU failure, these messages will be stored on duplicated mass memory units (drums) in the HLC of Entry and reassembly will be performed on drums in the HLC of Exit. For Type A multi-block messages, the reassembly will also be performed on drums although this is not essential for security reasons (see 1.4.2.)

Single block type A messages, since no reassembly is needed, are not stored on drums. Each time a certain number of messages has been reassembled, an End to End acknowledgement (see 3.6) is sent to the HLC of Entry. As a result, this centre no longer has responsibility for these messages and therefore may drop them from drums (see Figure 4).

3.8 Retrieval

Type B messages are stored on magnetic tapes for retrieval purposes in the HLC of Entry.

3.9 Transparency

Messages received by an HLC from the Low Level side in 5-bit per character code, will be padded up in 7+ parity bit characters.

3.10 Routing

The routing scheme adopted for the HLN can be described as follows:

a) for a given configuration of the HLN i.e.(no change in conditions of link or centre), there will be always just one route selected between any two centres.

b) Each centre will inform all others with status messages every time it observes a change in the condition of the links to which it is connected i.e. each centre in the HLN will have a map of the whole network configuration.

c) Each HLC A has listed all possible routes between any centre and itself. Routes corresponding to one particular HLC B are listed in Centre A according to pre-established priorities.

Obviously to each route there corresponds one outgoing link in Centre A.

d) Each time Centre A has to send a block to Centre B, it will, according to paragraph 3.10.C, select the highest priority route available and so determine the corresponding outgoing link.

This routing scheme is, of course, not the optimum with regard to the actual traffic pattern but it is simple. The updating of the routing tables is done three times a year.

3.11 Traffic flow control on the HLN

Should any centre be in overload situation, it will stop accepting blocks from other HLCs and will inform them

about its situation. If a centre receives such information, it will

a) drop all received blocks to be transmitted to the overload centre

b) stop transmitting blocks of messages received from outside sources.

4. THE SITA LOW LEVEL NETWORK

On this network there is just Type B traffic. Each time a message arrives in a centre, it is stored on paper-tape (or on mass memory units in the Hong-Kong Centre) before being sent to the next SITA centre or delivered to its addressees. For message protection, each sending centre (SITA centre or airline office) will number sequentially all messages it sends and each receiving centre will make sure that there is no break in the sequence. Since Type B messages can be multi-addressed, each SITA centre will have a responsibility table. For illustration, consider the example in Figure 5.

Station A has a multi-address message to send to stations B and C, thus A sends the message to SITA Centre 1, which then sends to B one copy and another to SITA Centre 2;

Centre 2 in turn relays the message to Station C. Centre 2, not being responsible for Station B when the message is sent over link 1-2, and having received the message from Centre 1, will not send it back to Centre 1 for onward transmission to Station B. This responsibility table avoids infinite duplication of multi-address messages.

5. THE HIGH LEVEL CENTRE

In this paragraph only the low level functions of such a centre are considered (see paragraph 2.1).

5.1. Airline Reservation Computers

An HLC exchanges type A and type B messages with a connected airline reservations computer. The type A messages received by the HLC are sent through the SITA medium speed network (see figure 6) to the enquiring airline display terminal.

The exchange of messages is effected by using a synchronous full duplex link control procedure (point to point, symmetric for each direction) over a multi-circuit link. This procedure follows the ATA/IATA standard. For each full duplex circuit, functionally it is comparable to the procedure briefly outlined in paragraph 3.4. All blocks of a given message are sent via one single circuit. In the envelope of each block there is a sequence number which allows message reassembly.

Type B messages and multiblock type A messages are reassembled and stored on duplicated drums in the HLC. Single block type A messages are just core switched. Each time any message is correctly reassembled and stored on drums the HLC generates a message acknowledgement comparable to that described in paragraph 3.7.

5.2 Satellite Processor

Message exchange between HLC and SP is described in paragraph 6.1.

5.3 Teleprinters

An HLC exchanges type B messages with low speed teleprinters. Messages, if received correctly by the HLC (e.g. proper format, serial number in sequence etc...), are stored on duplicated drums for security reasons. The HLC will then analyse the address(es) contained in the message before sending it on the appropriate circuit(s).

5.4 Telex

HLCs can also be connected to local telex networks and exchange messages with telex subscribers. As soon as the connection is established messages are processed as teleprinter message.

5.5 HLC computer systems

At present, SITA uses 3 types of computer systems. Philips DS 714 MK II, UNIVAC 418 II, UNIVAC 418 III. The Philips DS 714 MK II and UNIVAC 418 II are briefly described in the following paragraphs. The UNIVAC 418 III machine belongs to the same family as the 418 II but is faster and more efficient.

5.5.1 Philips Centres

Two DS 714 processors equipped with 2 drums, 7 tape units, 1 console, 1 papertape reader, 1 printer.

Each DS 714 processor has: 32 bit words, 2.2 us memory cycle (memory is byte addressable), 96 K words memory and direct memory access.

Two types of communication multiplexers are used:

a) up to 125 asynchronous lines with speeds from 45 up to 100 bauds;

b) up to 30 synchronous or asynchronous lines with speeds from 200 up to 9600 bps.

At present, for example, the Paris HLC is equipped with 3 multiplexers of type a for 300 lines and one of type b for 12 lines.

5.5.2 Univac 418 II Centres

Each processor has: 18 bit words, 2 μs memory cycle, a 64 K word memory, direct memory access.

The "two 418 II" system has one card reader, one printer, 4 drums of 260 K words each. The line controllers (CTMC) can control up to 32 low speed asynchronous lines or up to 16 synchronous lines. The Rome HLC, for instance, is at present equipped with 2 CTMSs (5 medium speed lines, 50 low speed lines plus 2 remote multiplexer lines).

5.6 Example

As an example this paragraph describes the handling of a type B message received by a UNIVAC 418 II centre.

5.6.1 The UNIVAC 418 II design and program structure

The executive of the 418 II provides 4 levels of priority (1. for real time program and the executive; 2. for time dependent jobs; 3. for batch; 4. for computations).

The main routine called "the switcher" gives control to the programs according to their priority.

When an interrupt occurs the interrupt-answering routine takes control of the processor to toggle the buffers and log the generated status. Then according to the priority of the interrupted program, the control will be returned to the interrupted program or to the main routine.

Within each priority level an application program (the sequencer routine) defines the order in which the application programs are scheduled.

5.6.2 Processing of Type B messages coming from the Low Level Network (see Figure 7).

Each time 5 characters have been received, the interrupt is processed by the interrupt-answering routine which performs the toggling of two 7-word buffers (5 words are used for the input characters and 2 words to chain the buffers).

The pre-edit routine performs:

a) the character packing in an input staging buffer (3 characters per word);

b) the checking of message delimiters and errors which can occur (stuck tape, over length, cancel signal, line open...).

The input staging buffer is taken in a list of free chained buffers. When such a buffer is full or the message completely received (for the last buffer) the input staging routine writes this buffer on drums.

"EDIT" gets control from the "sequencer" routine and is scheduled by the Input staging routine. The main function of EDIT is to process the input messages and to build the output messages.

The "output staging" routine and "the real time interrupt answering service output" perform the following functions:

- reading of message on drum;

- logging on tape unit;

- unpacking of characters in the two alternate buffers used for output.

5.7 HLC system recovery

In each HLC one (or two) processor(s) perform the traffic handling functions while another is in a standby mode. This standby processor may perform off-line functions (e.g. batch processing). Then should the on-line processor fail, the standby system takes over.

5.7.1 <u>For the UNIVAC system</u>

Switchover is effected manually. An alarm warns the system control room when the on-line processor fails. The supervisor then switches manually the drums and line controllers to the standby processor.

5.7.2 <u>For the Philips system</u>

If the on-line system falls the system control is automatically transferred to the standby processor which aborts any off-line operation in progress and immediately starts the traffic handling function. A failure in the on-line system may be detected:

a) by the hardware or software traps of the on-line system which then ceases to operate.

b) by the alarm and switcher unit which should receive periodically a signal from the on-line processor.

5.7.3 Since the on-line processor records the current status of the dynamic tables on both drums, the standby processor will use this "snap-shot" of the current work to take over the task whenever necessary. A drum failure has no effect on either system as two drums in the system are always maintained as images of each other.

5.8 <u>HLC Saturation</u>

All sources which send traffic to an RLC have free access to that HLC.

5.8.1 In case of memory congestion the on-line processor in the HLC stops receiving traffic from other HLCs in accordance with paragraph 5.3.1.

5.8.2 In case of drum overload, telegrams are automatically sent to stop teleprinters from generating traffic and stored messages are drained on to magnetic tapes.

5.9 Statistics

Each HLC records statistics of the incoming traffic loads, type A and type B, of the High Level circuits, for performance measurements and billing purposes.

6. SATELLITE PROCESSORS (SP)

These computer systems are stand-alone processors with simply a communication controller and no peripherals. They each depend on only one HLC.

SPs are built to operate without manual intervention. As previously mentioned, they act as traffic concentrators and solicit/send traffic from/to airline CRT terminals.

6.1 HLC - SP

Data is exchanged between an SP and its parent HLC via a multi-circuit link using a full duplex synchronous link control procedure with message control (point to point, symmetric for each direction of each circuit), almost identical to that described in paragraph 5.1, Type B message reassembly is not performed in the SP. Blocks of one message are always sent via the same circuit unless the considered circuit is down, in which case the remaining blocks are transmitted via another circuit. Each message the SP receives from TTY, TLX or agent set is, after having been divided into blocks, sent to its parent HLC. The SP does not perform local switching functions.

The message, whether Type B or multi-block Type A, is reassembled on drum in the HLC of entry; if single block Type A, it is just core switched. Each data block exchanged between an HLC and an SP carries an identification number of the TTY/Agent set circuit over which the message has been received (input) or over which the message has to be sent (output). Furthermore, SP and HLC exchange service messages, allowing the HLC to have a remote control of the stations connected to the SP.

6.2 Teleprinter Lines

SPs accept messages from teleprinters using the ATA/IATA format. As soon as one block of TTY data (here maximum length 240 characters) is received, it is immediately sent with an envelope to the HLC. To reduce memory occupancy, the SP does not wait for receipt of the complete message.

Should an abnormal situation occur, (e.g. incomplete messages, line open, etc.) the SP Informs its HLC by sending an appropriate service message. HLCs send messages to SPs block-by-block, using the above mentioned link control procedure. However, Type B messages are not transmitted continuously. As soon as the first block is sent, the HLC does not send the following block of the same message until it is informed by the SP via a service message that the block (which at the link level had already been acknowledged) is completely transmitted to the outstation. Here again, this method of sending Type B messages reduces memory occupancy in the SP.

6.3 Telex

An SP has capability of carrying out automatic connection or response in conjunction with the local telex network. The SP utilises service messages to inform its HLC of all the steps of connection and disconnection. The messages received/sent from/via the Tlx network are processed according to paragraph 6.2.

6.4 Agent Set Lines

SPs control lines with Agent sets, using three different polling procedures: IBM 1006; Raytheon DIDs 400; Uniscope 100. Each Agent set enquiry received by the SP is sent block-by-block to the parent HLC, where it is reassembled on drums, if multi-block, and sent via the HLN towards the parent reservation computer (see Figure 6). Type A response messages generated by a reservations computer are sent by the HLC to the SP, block-by-block, according to the link control procedure. The message, if multi-block, is reassembled in the SP before being sent on the Agent Set line.

6.5 Satellite Processor Implementation

Two types of machines are used to perform satellite processor functions:

a) Raytheon 706: 16 bit words; 1μs memory cycle; 32 K words memory; the communication controller works with one control word group - its limits are 8 full duplex synchronous and 72 full duplex asynchronous lines; Direct Memory Access; special characters are recorded in main memory.

b) Thomson Houston (GE) 4020: 24 bit words;1.6 μs memory cycle; 32 K words memory; the communication controller works with two groups of control words -its limits are 64 synchronous or asynchronous full duplex lines; special characters are wired.

Each system provides a protected memory area for remote load of the program sent by the parent HLC, an auto restart against power failure and an interrupt stall. This is to avoid as far as possible any manual intervention. In each system the communication controller generates interrupts when buffers are full (receipt) or empty (sending), or when special characters are detected and when error conditions occur (data lost, parity error, carrier lost, etc.). Programs are arranged in order of priority.

6.6 Example

One of the two systems is described hereunder in more detail. This example is based on the Thomson Houston system.

6.6.1 Each program has a priority level and the priority is re-considered after each unit of the system timer (100 ms). Each procedure (see 2.2a,b,c,d) involves one program for input and one for output, associated with their specific interrupt sub routine (see Figure 8) which performs special character recognition and buffer switching up to the complete reception or transmission of a data block. Input and output programs analyse the validity of the block according to the procedure rules. Blocks correctly received are transferred with their identifiers in a general chained queue and the input program disregards them completely. A "Message Analyser and Message Transfer" program scans the above-mentioned queue. In each block found there, it checks the following points:

a) is the output circuit operational;

b) is the code of the characters in the block consistent with the receiving end of the output circuit;

c) is the block of the message In sequence with the preceding block of the same message;

d) is the output program ready to process the block (output queue).

If all the above mentioned conditions are met, the block is transferred to the corresponding output program, Otherwise it may be dropped or remains In the queue. If it is dropped the SP informs the HLC by a service message.

6.6.2 The main memory is partitioned in fixed length buffers which are taken one after the other in a free chained list. When the buffers again become available they are returned to the free chained list. A block correctly received is stored in chained buffers. After each buffer pickup the SP checks, if memory is still available. In case of saturation, the SP first informs teleprinters to stop generating traffic, then, if necessary, reduces the poll rate on the agent set circuits and finally stops receiving traffic from HLCs. Return to normal working conditions is always done progressively to avoid burst effects. The parameters of each data block received (address of the first buffer, block chaining number, input line, priority etc.) are placed in a special buffer called the "Block

Identifier" which is of 3 character length and taken from a pool of chained buffers. All the blocks are then represented by these special buffers.

The number of special buffers is computed so that it corresponds to the maximum number of blocks that can be received in the main memory.

6.6.3 Each time an SP restarts after a failure (power failure, watchdog timer, interrupt stall) it informs its HLC that it is ready for reload by means of a simplified version of the Link Control procedure via the HLC-SP link. The HLC can at will,

a) load the SP

b) reconfigure the SP circuits

c) dump its memory

d) reinitialise it.

7. RESULTS - PERFORMANCES

At present the majority of traffic handled by the SITA network is still type B. However, a rapid increase of the type A is expected within the next few years.

7.1 High Level Centres

As an example consider the High Level Centre in Paris: the traffic switched there from the Low Level Network to the High Level Network or vice versa, is in the order of 1.2 messages per second during peak conditions.

If the messages switched from Low Level to Low Level are also taken into account, 4 messages per second are received and 5.5 are sent. The difference between the last two figures is due to multi-addressing.

Average switching times for type B messages entering or leaving the HLN are around 150 ms.

As far as blocks which are just core switched are concerned, the mean value is in the order of 20 ms.

Downtimes per month of HLCs are around 2 hours, including all scheduled stops for service reasons (configuration changes, new program versions, etc.)

7.2 HLN

The probability that any pair of HLCs is completely isolated i.e. no route between them is available, is negligible. The corresponding downtimes are in the order of minutes a month.

The average number of messages leaving the HLN during peak conditions is about 7 per second.

Usual line loads during normal peak conditions are between 0.30 and 0.40 Erlang.

7.3 Response Time

The average response time for Type A messages, i.e. the time lapse between the Instant an operator presses the transmit key of his terminal to send his query and the instant the first character of the reply appears on the screen (see Figure 6) ranges from 1.4 seconds to 3 seconds, depending on the number of links involved in route. A typical response time distribution is shown in figure 9.

7.4 Satellite Processors

The downtime of SPs is around 16 hours per month including scheduled and unscheduled stops (e.g. preventive maintenance, configuration changes, etc.)

Presently, certain SPs switch up to 6 message blocks per second during peak conditions of traffic.

Switching times are in the order of 5 ms per block for both systems.

* * *

Note:

The Authors have restricted their article to the SITA operating network in 1973. However, it is clear that this network is in continuous expansion and that new automated centres will be operational within a few months.

* * *

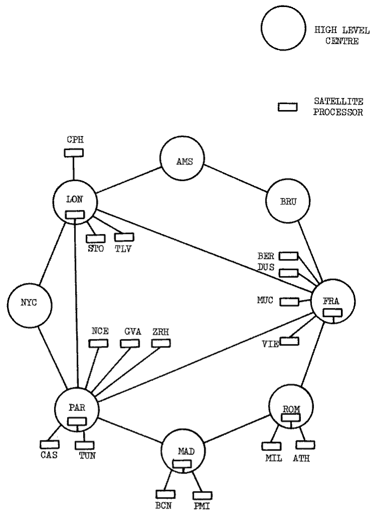

FIGURE I The SITA Medium Speed Network

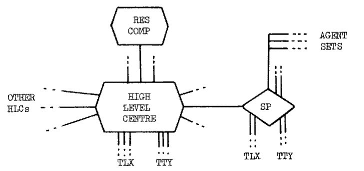

FIGURE II High Level Centre and Satellite Processor Interfaces

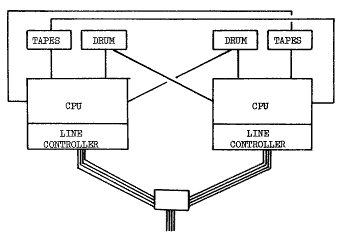

FIGURE III A Typical High Level Centre Configuration

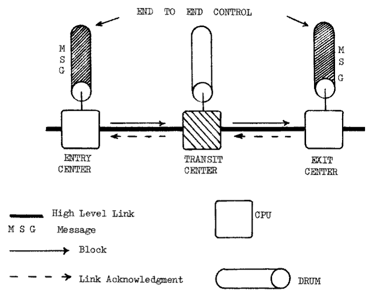

FIGURE IV End to End Acknowledgment

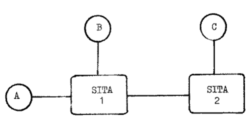

FIGURE V Message Responsibility

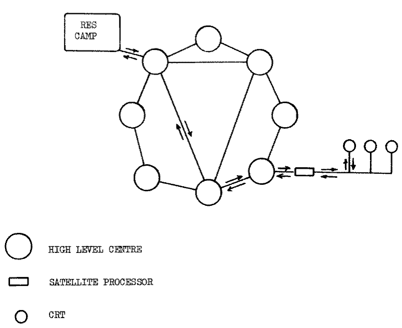

FIGURE VI Type a Message Path