EUTECO - European Teleinformatics Conference
T. Kalin (editor)
Elsevier Science Publishers B.V. (North-Holland)

SOME EXPERIENCES WITH PACKET SWITCHED SATELLITE COMPUTER NETWORKS

Peter T. Kirstein

Department of Computer Science
University College London
Gower Street, London. WC1 6BT

In this paper some experiences are discussed with three different broadcast packet satellite computer networks: the DARPA SATNET, the DARPA Wide Band Satellite Network, and the UNIVERSE/STELLA Networks. We concentrate on the unique properties of the satellite portion of these networks. In particular the specific reception problems, channel access experiences, monitoring and control requirements, and protocols experiences are discussed.

1. INTRODUCTION

Satellites have been used for data transmission for many years. For example the Defence Advanced Research Projects Agency (DARPA) was carrying traffic between a seismic array in Norway and a processing centre in Washington as early as 1970. There have been satellite transmission systems in packet switched computer networks for at least a decade; by 1973 the ARPANET [1] had satellite transmission portions then their extensions to Norway and Hawaii. These transmission paths did raise specific problems, because of the long transit times for the individual packets. Comparatively new or local modifications to the transmission algorithms dealt with the immediate problems; more buffers were provided in the sections using the satellite channel. These measures dealt with the establishment of viable networks. However since one did not change the higher level protocols to deal with the small number of satellite sections, the overall performance seen by the end user of such portion of the network could be much worse than those seen by the user of a terrestrial portion. For example, the file transfer speed seen by a single user on ARPANET could reach 35% of the 50 Kbps line speed used in the long distance transmission on the terrestrial portion; the corresponding user of the UK-US portion, while included in 9.6 Kbps satellite hop, never saw more than 10% of the 9.6 Kbps basic speed. In spite of such effects, this paper will not treat such computer networks where satellite portions are merely incidental.

Another class of networks was introduced with the start of the SATNET project in 1975 [2]. Here the attempt was made to allocate a single broadcast satellite voice channel on an INTELSAT satellite dynamically between a number of computers. Moreover, these computers were organised as a packet switched computer network with specific protocols, connection to Host computers, and connection to other networks. The SATNET activity was followed in 1978, by the European STELLA I project; here a 1 Mbps channel on the OTS satellite was shared by computers in a number of European laboratories. These two activities were followed by the DARPA Wideband Satellite Network (WBSN, [3]), European STELLA II project [4], and the UK UNIVERSE project [5]. All three of these started coming live in 1982, and operating experience is only just becoming available.

A brief discussion of all four projects is given in Section 2. Since the satellite portions of the STELLA II project is identical to the UNIVERSE one, and operating experience is even more limited than on UNIVERSE, little is said here about STELLA II

We will concentrate on specific aspects of the four systems: reception problems, processor bandwidth limitations, network monitoring and control, and the impact of protocol choices. These are discussed in Sections 3, 4, 5 and 6. Only SATNET has provided an extensive service for many years, so that it has had to be developed into a real operational tool. For this reason the section on monitoring and control draws particularly heavily on experiences with that network. The author's personal experience has been only with SATNET and UNIVERSE; the discussion of STELLA and the WBSN is based on second hand experience. Moreover, lack of time has precluded the possibility of a comprehensive review of this paper by those with direct experience of these two networks. For this reason, my comments on those networks should be treated with caution.

Our work on the networks mentioned in this paper have been supported by the UK Science and Engineering Research Council, the UK Ministry of Defence, British Telecom, and the US Defence Advanced Research Project Agency; their support is gratefully acknowledged. The work on the development of so many organisations, that individual attribution is invidious. However, one must note the individual efforts of RE Kahn (DARPA) to set up and further SATNET and the WBSN, and of W Hine (CERN), D Thomas (RAL), and J Burren (RAL), who gave STELLA and UNIVERSE a similar impetus.

2. FOUR SATELLITE COMPUTER NETWORKS

In this section we will describe briefly the properties of four Satellite Computer Networks - SATNET [2,6,7], the DARPA Wide Band Satellite Network (WBSN, [4]), STELLA [3] and UNIVERSE [5,9,10]. Each of these are broadcast satellite networks; a number of earth stations share the same satellite channel. Each uses the satellite bandwidth in a Time Division Multiple Access (TDMA) mode; a frame is divided dynamically into slots, which are occupied according to well-defined algorithms by the different earth stations. Figure 1 shows a schematic of the configuration. In each case the Earth Station (ES) has an Earth Station Interface (ESI), attached to a Satellite Interface Message Processor (SIMP). Each SIMP was attached directly or indirectly to another network.

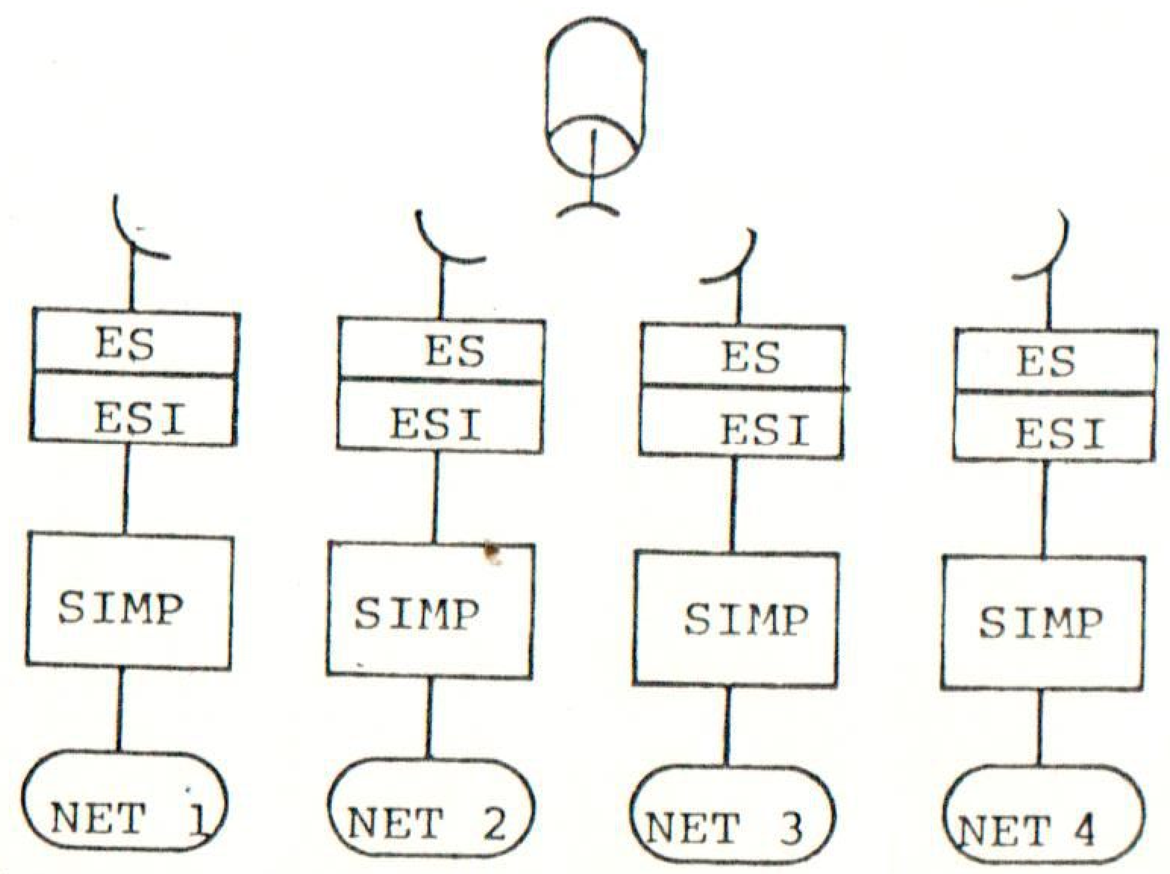

Figure 1.
Schematic of Packet Satellite Configuration

The other parameters of the networks were very different, and will be discussed below:

(a) SATNET
This is the oldest of the three networks. The project started as an experiment in 1975, and was transformed into a routine service about 1979. The satellite channel is a 64 Kbps SPADE channel (at 4-6 GHz) on an INTELSAT V satellite, and the SIMP was a Honeywell 316 (now being replaced by a BBN C30). In this case the ESI was called a PSAT. The SIMPs are attached, via 50 Kbps land-lines, to Gateways (G) which connect into ARPANET [1] in the US, and other networks in Europe; in the early days all the SIMPs were connected via gateways to ARPANET. The situations in 1979 and 1983 are sketched in Figs 2 and 3. (In 1979 there was also an INTELSAT B Earth Station attached at the COMSAT laboratories. Since it does not participate in the current service, it has been omitted from the discussion.)

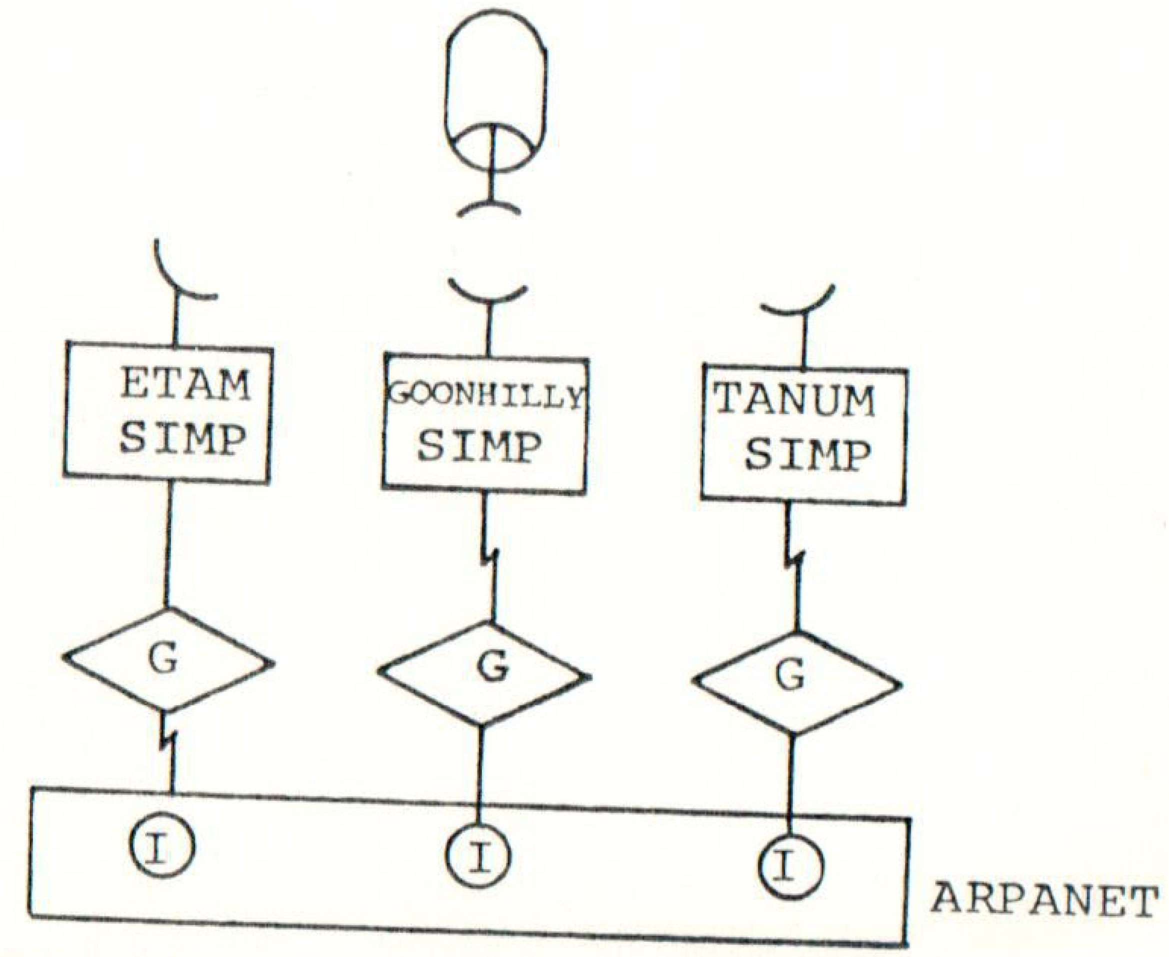

Figure 2. SATNET in 1979

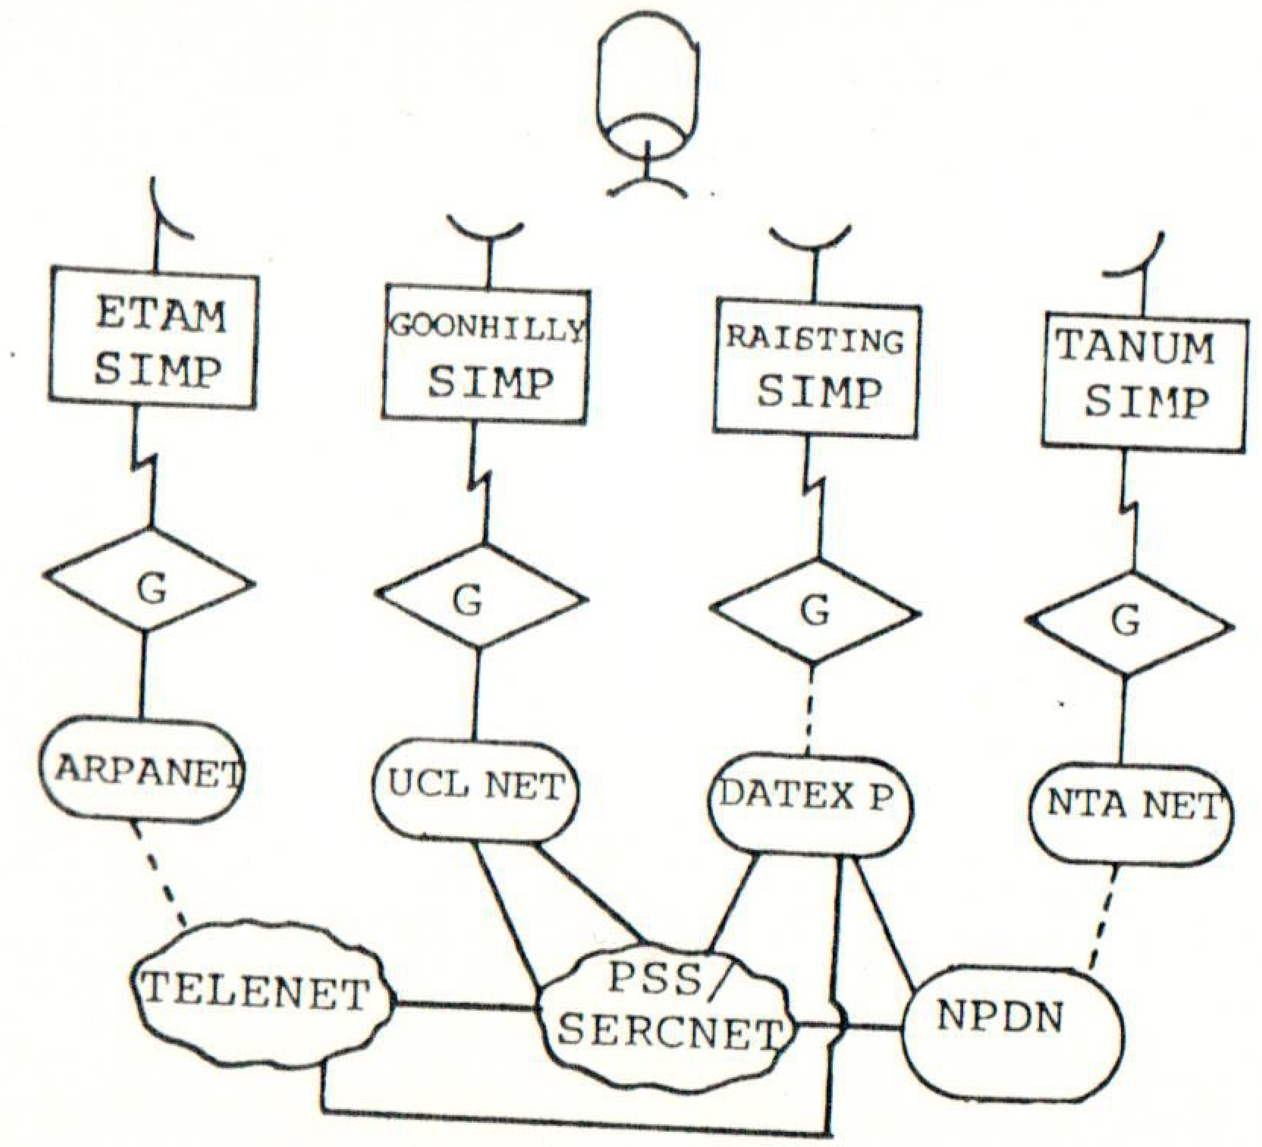

Figure 3. Schematic of SATNET - 1983
----means under development

In 1979 all three remote gateways G were attached to ARPANET (and in fact the ARPANET model computers in the US and Norway, I in Fig. 2, were connected via a satellite circuit of the normal type). For a while, until 1981, direct ARPANET channels were simulated via SATNET. Only by 1982 did we reach the situation of Fig. 3, where there was no direct connection between any of the gateways except via SATNET. Since the networks in all the countries have gateways to their PTT public data networks under development, all the G should be accessible via the public X25 networks by the end of 1983.

In SATNET the bit error ratio is about 10^{-6} under good conditions. We use slots of up to 1 K bit (20 ms) and a guard band of 64 symbols. The CPODA satellite access procedure is used [6]; this provides random contention for short reservation slots, and reservation for larger slots.

(b) The Wideband Broadcast Satellite Network (WBSN)

The WBSN system started as an experiment in 1980. It is still being commissioned. It is a sophisticated system, which uses much the same type of connectivity as SATNET. Now various different local area networks are connected in (Packet Radio, a local broadcast net (LEXNET), a Circuit-switched local net, and Arpanet to all sites.

The SIMPs of Fig. 3 are now PLURIBUS machines (based on 6 Lockheed SUEs and called PSIMPs) [11], and CPODA is again used for satellite channel access. PDP 11/44s are used as gateways, and the same INTERNET Datagrams. Special high performance voice concentrators (2-20 Mbps) are being developed for speech.

A Western Union Weststar III satellite is used with a 3MHz channel at 4-6 G Hz and a 5 m Earth station. Their modems are supposed to use speeds of between 193 K and 1544 K symbols/sec, with BSPK or QSPK coding - and thus should be able to achieve 3.088 Mbps. The ber should be about 10^{-3}, and they use sequential coding with rates of between one and half rate coding. The bursts are short, up to 64 K bits; they use a short guard band of 8 symbols, with a 25 symbol preamble. A schematic of the system is shown in Fig. 4.

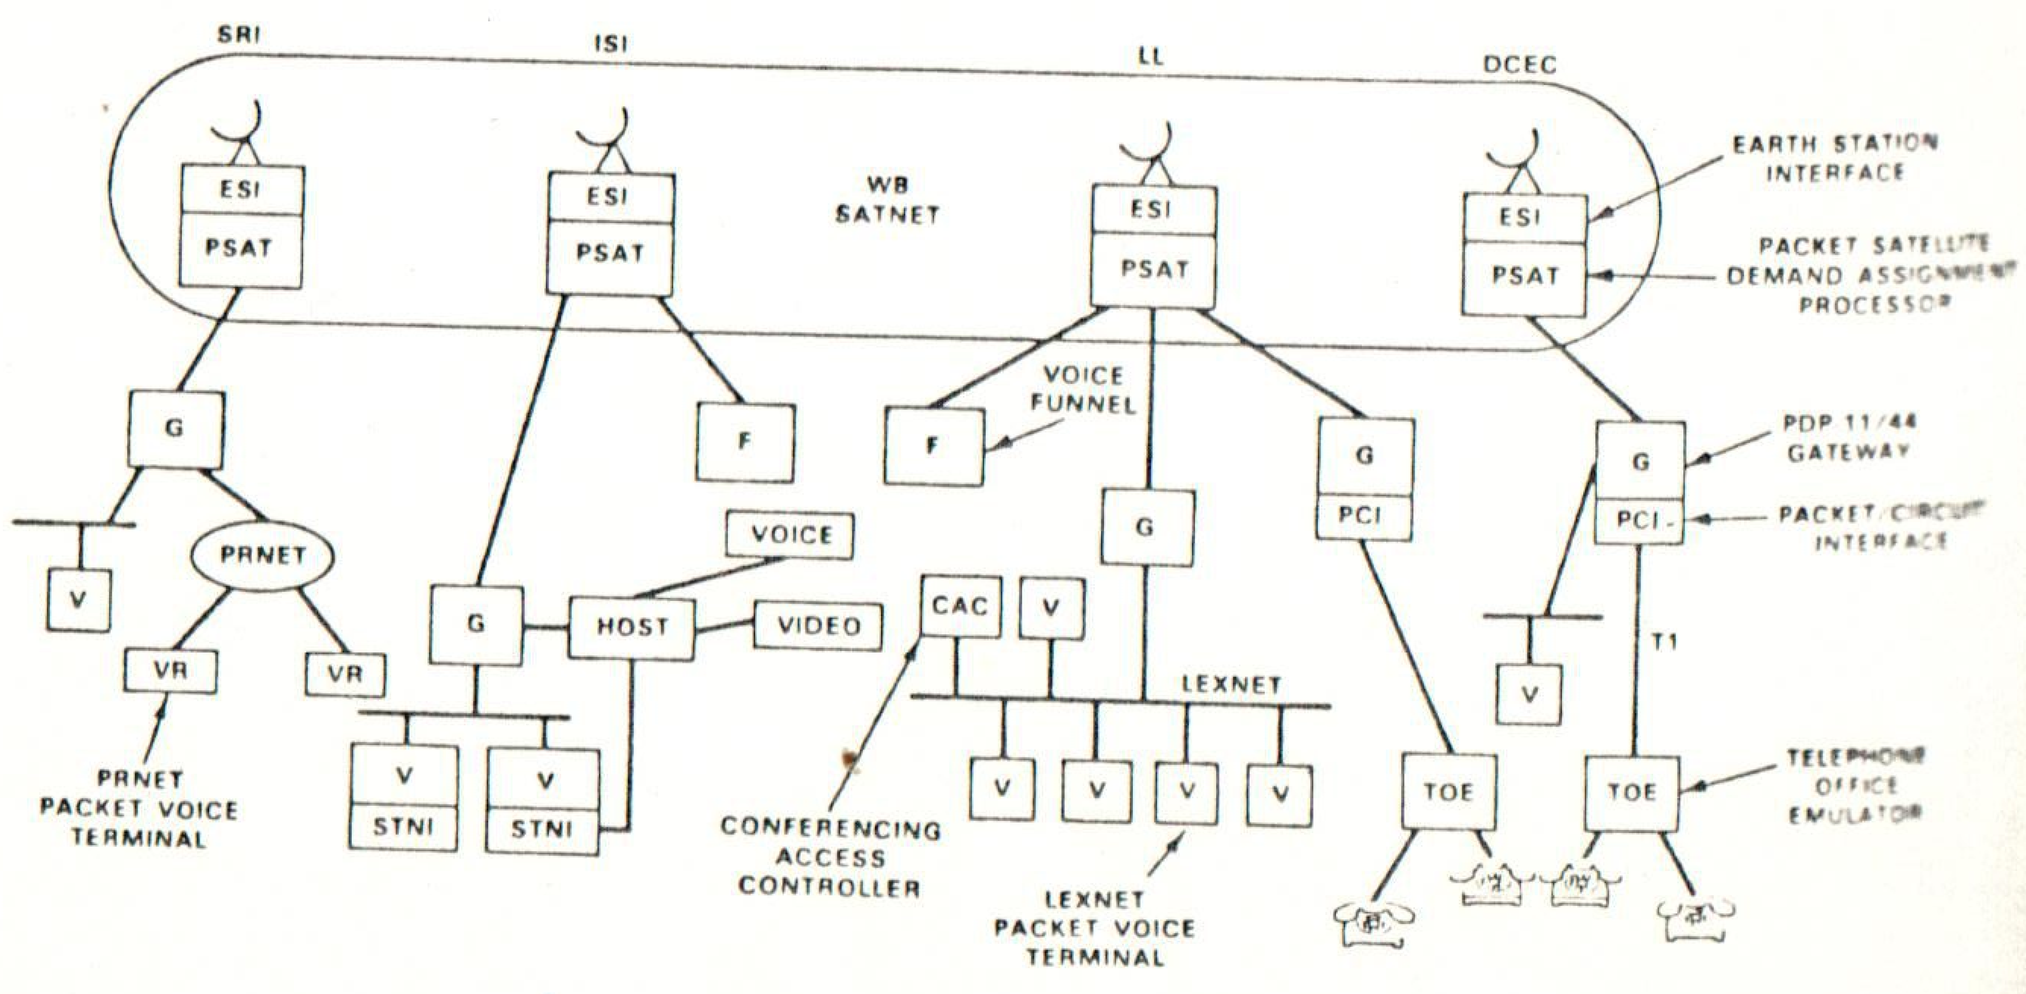

Figure 4. Schematic of WBSN at end 1982

(c) STELLA II/UNIVERSE

The STELLA I was a European experimental network started in 1978, and was replaced by STELLA II in 1981/82. UNIVERSE, a British experimental system, uses the same satellite access system as STELLA II, but has somewhat different terrestrial equipment. The STELLA II system has earth stations at CERN (Geneva), CNUCE (Pisa) and the Rutherford and Appleton Laboratory (RAL, near Oxford). UNIVERSE has earth stations at RAL, Cambridge U, University College London, Loughborough U, British Telecom Research Laboratory (Martlesham), and GEC - Marconi Laboratories (Chelmsford); it started operating in 1982. Both projects use 2 MHz channels on OTS at 11/14 GHz, 3m earth stations, and 1 in 2 Viterbi encoders giving an error rate of 10^{-9} (under good conditions). One important difference between STELLA and UNIVERSE, is that in the former different hardware is used at each earth station, in UNIVERSE only a common hardware and software are used. This makes software maintenance and systems commissioning much more difficult in STELLA than UNIVERSE. Both systems have local area networks attached to the earth stations. The systems are sufficiently similar (particularly in their satellite portions being developed collaboratively) that we will discuss only UNIVERSE. A schematic of that network is shown in Fig 5. The satellite access is by centralised control. Each earth station has a fixed small window for reservation purposes; this takes 10 ms out of a 140 ms frame. The remaining 130 ms can be allocated in variable sized data windows - but only one window is available to each transmitting earth station each frame. Transmissions for several destinations can be carried in the window used by one transmitter.

All the SIMPs (called LDCs in UNIVERSE) are connected to an X25 network; there are also gateways to that network from the UNIVERSE local area networks - mainly Cambridge Rings. The UNIVERSE system contains some 15 local area networks, and 150 computers (many small micro computers). Full descriptions are given in [5] [9] [10].

Figure 5

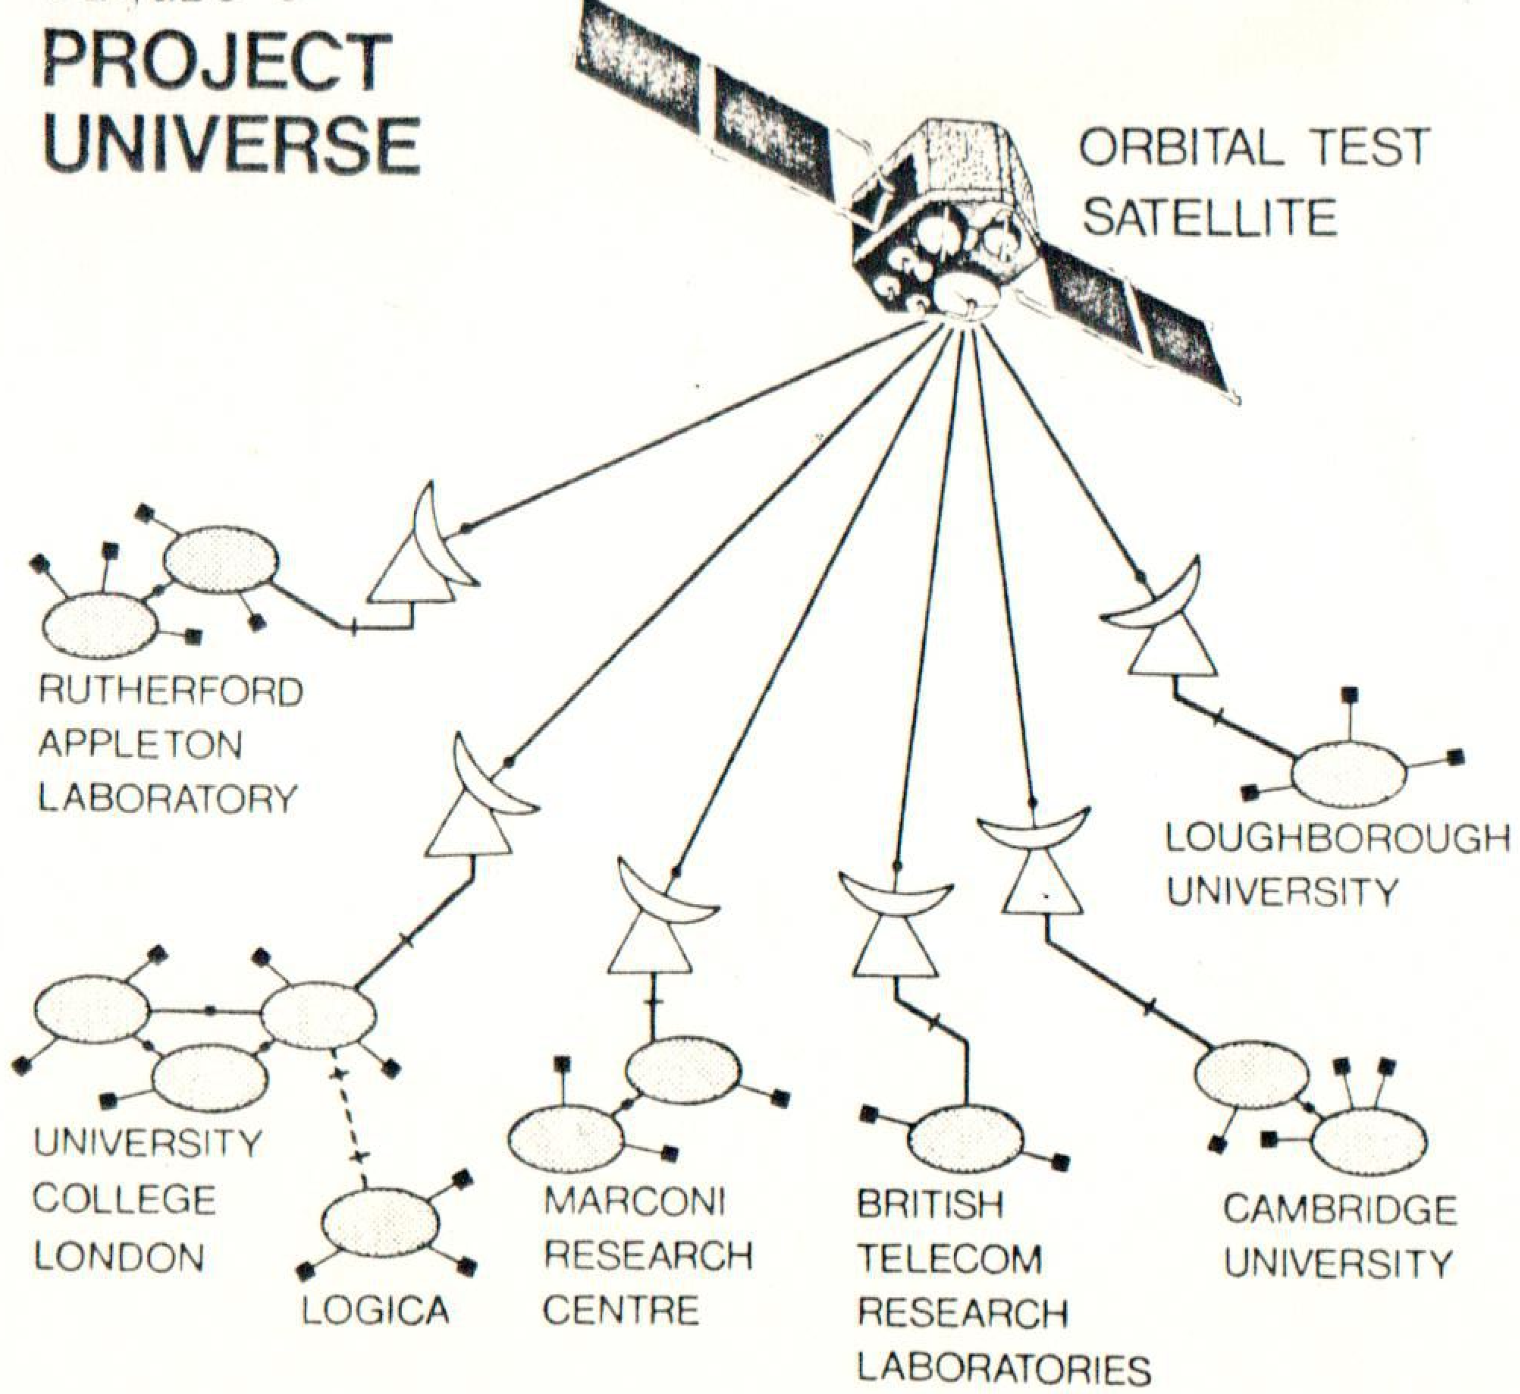

96

3. RECEPTION PROBLEMS

In all the systems, reliable operation of the satellite portion has been a problem - though the reasons and symptoms differ. In the SATNET case there have been four principal reasons:

(i) Drift of the whole power level of the SPADE transmitting outside specifications, has caused one earth station to shut out others.

(ii) Transmission on another SPADE channel has caused interference due to faulty adjustment of the SPADE equipment at another earth station.

(iii) Frequency drift of oscillators.

(iv) Sundry equipment failure.

The first of these is much more serious with the SATNET type of traffic than voice for which the SPADE system is normally used. For the normal voice, only one SPADE terminal is assigned a particular frequency band; for SATNET all participating SIMPs are assigned the same frequency. This variation in transmission power can actually cause erroneous operation when two SIMPs are transmitting at the same time. The second is an important operational and managerial problem; it is due to poor operating procedures, and has occurred also on the WBSN. The third can be difficult to distinguish from the others' causes. There are clearly many causes for the fourth.

In the WBSN case there have been similar problems in particular due to:

(i) Drastic change in power levels due to variations in the power levels allocated to the WBSN.

(ii) Interference from terrestrial microwave sources.

(iii) Electromagnetic interference on various of the interfaces.

With the high bit error rate, 10^{-3}, it does not take much degradation for the headers to become corrupt. The CPODA system then becomes unworkable.

There has been pattern sensitivity in the STELLA/UNIVERSE modems, problems in maintaining power levels and frequency drift of oscillators. Otherwise the reception problems have been less in this case - partly because the signal to noise ratio was so much higher, there was almost no sharing of channels with other applications, and there were very large guard bands.

4. SIMP PROCESSOR BANDWIDTH

In each of the cases, the SIMP ran out of processor bandwidth earlier than expected. It was always necessary to put in extra facilities to deal with unexpected events, and extra diagnostics were required. In the case of SATNET, this limited the satellite current speed to 64 Kbps with the Honeywell 316: one had to change to a C30 before 128 Kbps could be attained with CPODA. In the case of the WBSN, the shortage of SIMP bandwidth meant that all six processors of the PLURIBUS [11] were required to attain 1.544 Mbps on the satellite circuit with CPODA; this reduced the reliability of that SIMP, because it has been expected that the redundant processor architecture would be more resistant to failures.

Both the SATNET and WBSN had put very considerble diagnostics into their SIMPs and ESIs. In STELLA/UNIVERSE, the shortage of processor bandwidth led to an inability to put much status monitoring or ability to handle remote access into its statistics while it was running. Clearly the fact that even SATNET runs out of processor bandwidth with only a 64 Kbps satellite channel, while the WBSN and STELLA/UNIVERSE aim at a further factor of 20-40 in satellite channel bandwidth, has implications on the SIMP bandwidth needed.

In all cases the new designs are putting much greater processor power into the SIMP.

5. NETWORK MONITORING AND CONTROL

Multiple access chanels are much more sensitive to variations in parameters than ones in which only one device uses the channel. While the satellite channel and even earth stations may seem to be operating satisfactorily macroscopically, the single multi-access channel may be unusable. For this reason it is important to put in adequate monitoring facilities - and preferably control to allow remote testing.

In the case of SATNET, there were two iterations. The second produced the Packet Satellite Program (PSP) Terminal [6] in the place of the ESI of Fig. 1. This is specifically designed to provide on-line performance monitoring of the different parameters under the control of the SIMP, and off-line testing under control of either the SIMP or a Data Test Set simulating the SIMP. The on-line monitoring capability provides information on progressive performance degradation before catastrophic failure. The off-line testing capability, coupled with the ability to set signal loopbacks at various points within the PSP terminal, is a powerful fault diagnostic tool. Commands to the PSP terminal from the SIMP, and certain non-dynamic status information, can be passed between SIMP and PSP under program control. Examples of the types of parameters that can be controlled are [7]:

Change preamble length of packets
Change SOM patterns
Place PSP in various modes of reception
Place PSP in various loopback or bypass modes
Allow the PSP to operate with Data Test Set and remotely
    control the operation of the Data Test Set
Dynamically monitor the automatic gain control of the modems,
    and the correlation erros detected in the PSP
Control the interfaces, and receive various types of error
    detection information
Reset the PSP terminal and set it in a predefined mode.

The data associated with the real-time on-line monitoring and the data obtained during off-line testing, are transmitted to the SIMP over the normal data paths. This data includes information on channel conditions, modem performance, and receive interface operation, as well as specific parameters such as frequency off-sets, acquisition and demodulation AGC values, signal power and noise power. The receive interface, as part of its processing, performs certain consistency checks on the packet, including verifying the cyclic redundancy check sum (CRC), and appends the pertinent information, and its own status word onto each packet. Further details of fault operation and of the burst mode Data Test Set are given in [7].

The WBSN has incorporated the same techniques into its ESI, and similar techniques are used. The UNIVERSE/STELLA ESIs (called CIMs in the references) have no similar remote control facilities or on-line monitoring. In some cases (e.g. AGC, signal level and noise level) the parameters are measured on the earth station control rack; however, no mechanisms are provided for passing them to the SIMP. Both Graz U under STELLA I and UCL under UNIVERSE have built computer interfaces to collect data on signal/noise ratio; at UCL this data can be passed over the local area network under remote program control.

In each of the networks, each earth station periodically sends HELLO packets. Since these should be recovered by each earth station and hence SIMP, it is possible for any one SIMP to gather information not only on its own ESI and earth

... In SATNET, the SIMPs collect all this relevant information, and send it to a specific computer, called the Internet Operations Centre (INOC) as normal data. The data passes either through SATNET or the terrestrial lines, depending on the connectivity. Under normal conditions the SIMP sends this information periodically and autonomously. When the SMCC does not hear from a SIMP for a predetermined interval, it polls the SIMP over all available communication paths at frequent intervals. Whenever the SIMP receives a poll over a communication path, it responds over the same path. The INOC will then continue to use this path until normal communication is restored. This strategy allows the INOC to maintain end-to-end connectivity in times of problems, while maintaining efficient reporting elsewhere. An elaborate data base of the status messages, suitably time stamped, is kept on the INOC. A number of auxiliary programs then allow on-line access to the history of various parameters, and to allow routine or specific correlations to be carried out. Other programs process parts of the data base autonomously, and send individually tailored regular digits as ordinary electronic mail to specific individuals concerned with SATNET operations.

In SATNET we have already said that ESI and SIMP parameters can be set up remotely. A set of programs have been devised in the INOC which allow the following operations to be performed on-line:

- change any part of a SIMP program
- setting of ESI parameters
- setting of SIMP or ESI interfaces
- snapshots of SIMP internal tables
- polling of ESI status parameters

Particular data files have been set up on the INOC to allow specific experiments, tests or diagnostics to be run routinely, without the need to enter all the necessary parameters. The data files usually contain default parameters, which can be modified before a particular test is run.

The WBSN is building on the same monitoring and control infrastructure as SATNET, and uses the same INOC. Neither UNIVERSE nor STELLA have a similar infrastructure. UCL will put in some similar facilities [12], but the SIMP, ESI, and earth station have not had this concept of remote diagnostics built into their basic implementations. Thus, for example, the signal to noise ratio of the satellite channel is now being monitored - but only at one site by building a separate microprocesssor input from an existing monitor port on the earth station control to a separate interface to the UCL local area network. All three networks can remotely load the SIMP software from a central point; only SATNET/WBSN can do much on-line control or modification of the SIMP.

Over a few years, the INOC has become an important operational tool in the maintenance of the SATNET. It is particularly valuable in its role of allowing several different people to collaborate in trouble shooting - e.g. SIMP programmer, modem designer, ESI builder, satellite control, and earth station control. The traffic statistics are read daily by the operation staff, and preventative action can be taken on oscillator drifts or AGC changes. Occasional serious outages still occur. However, usually MTTRs are now hours rather than the days we saw in earlier times. The WBSN is not really in this state yet, because some of its equipment is still being commissioned. It is much more ambitious in its design aims; it is working at a very low signal to noise ratio; it uses very sophisticated channel access algorithms. This makes it much more sensitive to any malfunction. Since the alternate path, by-passing SATNET, was removed fault diagnosis becomes much more difficult in the event of equipment failure. We anticipate using X25 networks to allow this fault route to be used when the dashed lines of Fig. 3 become a reality.

The UNIVERSE satellite portion usually works fairly well; it has a much higher signal to noise ratio, and a comparatively conservative design. However, if there are failures anywhere, faults are hard to detect, let alone correct. The diagnosis can be very difficult, and the meantime to repair long. One important reason for this is the lack of diagnostic tools available. However, there is one important aid to diagnosis. Each SIMP is attached to an X25 network in such a way that a remote terminal can act as its console. This allows both monitoring and program modification of all SIMPs from a central site. This greatly aids fault diagnosis and repair.

6. EXPERIENCE WITH NETWORK PROTOCOLS

All four networks adopted protocols with some similarities and some differences. All use a datagram as the fundamental transport unit. In SATNET and WBSN, this datagram has a specific Internet Protocol (IP) format described in [13]. This format allows priority, type of service, fragmentation and reassembly. While gateways do not maintain state information and can discard datagrams, they will also send such quenching messages to the source if packets are discarded. The SIMPs do allow for streams to be set up, thus allowing buffers and bandwidth to be reserved for data requiring speedy delivery - like real-time voice. End-to-end reliability is achieved by a standard US Defence transport protocol called the Transmission Control Procedure (TCP, [14]), which has elaborate facilities for error control and flow control. These streams have also proved useful if there is heavy traffic between two earth stations; with short packets of data, since it allows data aggregation and the transfer of a smaller number of larger packets without the need for reservation. This is particularly important, since we find the number of packets/sec is usually the limiting bottleneck in SIMP and Gateway throughput. We have also found that adaptive protocol implementation can greatly increase end-to-end performance as seen by the user; an example is when the retransmission timeout is made to depend on the round-trip delay.

In STELLA and UNIVERSE a somewhat similar datagram mode [15] is defined - with significant divergences of the Transport level. For STELLA the ISO level 4 [16] and CERNET [17] procedures are being implemented. We have not studied their properties in the context of Satellite networks. For reliable transmission in UNIVERSE a light-weight control call is established - without flow control or error detection at the gateways. In its early versions an end-to-end protocol with a window of one is used; at a later stage a window of eight will be adopted.

The datagram modes of all four networks can probably be made reasonably compatible - though the lack of fragmentation, time stamping, and type of service would be serious disadvantages if the STELLA/UNIVERSE datagram was used in large concatenated networks. (STELLA does allow time stamping) The lightweight virtual call used in UNIVERSE will be very inefficient if a substantial number of packets are discarded. Moreover, the protocol structure used for virtual calls will reduce reliable traffic to a maximum station-to-station bandwidth of 64 Kbps - but will probably be much less.

We have little experience yet with the performance throughput of STELLA/UNIVERSE. We have measured a little over 40 packets/sec can be transmitted to a non-active site [10] (giving up to 650 Kbps for 16 K bit blocks), and also a 60 Kbps throughput on an end-to-end basis - but have not investigated whether there is any particular limit of our system, or the cause of any such limit. With SATNET we have identified throughput limits at the SIMP, Gateway, and Gateway Arpanet interface. However, the main performance limits are that the end Hosts often optimise their buffer strategies for high speed terrestrial links - and are therefore much too low for satellite situations. In addition the timeouts are usually set for the higher performance, lower delay, terrestrial routes. While the satellite channel operates reliably, the system can usually cope. If the channel becomes more noisy, many timeouts can be triggered at different protocol levels, which can render the system unusable. Great care must be taken in specifying timeout, retry and reporting features to avoid such an occurance and pinpoint it if it occurs.

7. CONCLUSIONS

From our direct experience with two, and our indirect knowledge of two other broadcast packet satellite networks, we can draw some general conclusions.

They can be made to work, and work reliably. It is vital to put in the maximum of remote diagnostic and control facilities. The efforts on collecting the diagnostics and routinely processing it is well spent. The design performance of these networks can be more difficult to achieve than expected. It is again very important to have detailed information on transit times, packets dropped, maximum processing power of SIMPs, terrestrial paths and gateways. The basic raw processing or transmission power may not be adequate. Comparatively small unanticipated poor performance in such segments (e.g. higher noise) can cause queue build ups, time outs, and dramatic deterioration in systems or user throughput. Only a clear awareness of the causes of the problem, and its consequences, provides the information to change algorithms and make the whole system less sensitive.

The broadcast nature of the communication, and the simultaneous use of the same channel by several sites, makes this type of channel use more sensitive than the normal TDMA or FDMA use of communication channels - even satellite ones. This sensitivity makes it very difficult to analyse faults from one, or even just two, locations. A systematic inbuilt error reporting, performance monitoring and system control facility is extremely important as an operational diagnostic and maintenance tool. Moreover, many administrative procedures may need to be put in place to ensure that the people with the information can talk directly to those who must effect the repairs.

Care must be taken in the architecture and design of the SIMPs and ESIs. They are already usually short of processor bandwidth - yet the monitoring and control features are vital to their reliable operation.

The protocol structure options over the satellite portion can often not be divorced from those in other portions of the network. It is vital to have a clear idea of design goals in error performance, traffic characteristics and application requirements in the choice of the protocol architecture. In view of the critical dependence of performance on protocol algorithms, adaptive modification of key parameters, depending for instance on round-trip delay, can be very important.

REFERENCES

1. LG Roberts, et al: "The ARPA network", Computer Communications Networks, Prentice Hall, pp 485-499, 1973

2. RE Kahn: "The introduction of Packet Satellite communications", Proc Nat Telecom Conf IEEE, pp 45.1 1-8, 1973

3. N Celandroni et al: "STELLA - satellite interconnection of local area networks: present state and future trends", Proc ICCC, ONLINE, pp 425-430, 1982

4. HM Heggestad et al: "Voice and data communication experiments on a wideband satellite/terrestrial internetwork system", to be presented at the IEEE International Conf on Communications, Boston, June 1983

5. PT Kirstein et al: "The UNIVERSE Project", ibid, pp 442-447

6. IM Jacobs et al: "Packet Satellite network design issues", Proc Nat Telecom Conf, IEEE, pp 45.2.1-12, 1979

7. DA McNeill et al: "SATNET monitoring and control", ibid, pp 45.3.1-4

8. WW Chu et al: "Experimental results on the Packet Satellite network", ibid, pp 45.4.1-12

9. CJ Adams et al: "Protocol architectures of the UNIVERSE project", Proc ICCC, ONLINE, pp 379-383, 1982

10. AG Waters: "Project UNIVERSE - Implementation and experience of the satellite bridge", this conference proceedings

11. SM Ornstein et al: "Pluribus, a reliable multi-processor", Proc Nat Computer Conf, AFIPS, pp 555-560, 1975

12. SR Wilbur et al: "An outline of network measurement in project UNIVERSE", ibid, pp 390-394

13. JB Postel: "The DoD standard internet protocol", RFC 791, Information Services Inst, Marina del Rey, 1981

14. JB Postel: "The DoD standard transmission control protocol, ibid, RFC 792

15. R Cole et al: "An interworking architecture for STELLA and UNIVERSE", Proc Conf Satellite and Computer Communications, Versailles, North Holland, pp 369 - 378, 1983

16. N Celandroni et al: "Satellite interconnection of local area networks: architecture and protocols", ibid, pp 25-40

17. JM Gerard: "The development and use of a high speed satellite packet switching network in a high energy physics laboratory", Proc EURO IFIP 1979, North Holland, 1981

96

EUTECO - European Teleinformatics Conference
T. Kalin (editor)
Elsevier Science Publishers B.V. (North-Holland)

96

ISO REFERENCE MODEL AND SATELLITE COMMUNICATION

Christian HUITEMA

Nadir Project
INRIA
P.O. BOX 105
Rocquencourt
78153 LE CHESNAY
FRANCE

The ISO model aims at protecting software investment against changes of technology, like the introduction of satellite networks. It is shown that, if these networks could offer an X25 like service, the existing application would only need simple tunings to use them. The specification of a new service for new applications like packet voice or images is discussed. It should be connection oriented, with emphasis put on fixed throughputs and delays. The connection will be either point to point or point to multipoint ; this new service will have to be introduced in the OSI Model.

ISO model and Satellite communication

Rapid changes in network or computer technologies are hazardous for the user who may have to rewrite, periodically, considerable amounts of software. Thus, the seven layers OSI model has been standardized. Its goal is to protect the users ; its means are layering and standardization.

The concept of layering, and independance of the different layers, has become well known. Each of the seven layers offers a service, which is performed by enhancement of the lower layer service, through the use of a protocol. Not surprisingly, the standardization process began in a point-to-point, connection oriented context, meeting the needs of most existing applications. However, some provisions have been made for the future introduction of point-to-multipoint or connection less services.

Satellite communication is a new technology. It offers new possibilites : high bandwidth and broadcasting, to the cost of a long transmission delay (around 300 ms). So, two questions arise : to what extent shall this technology lead us to rewrite existing software, and what new features shall be introduced in the OSI model, in order to facilitate a "safe" development of new applications ?

The "Standard" point-to-point applications

Most existing applications are point-to-point, connection oriented. Even on a connectionless datagram network, the user is given a connection oriented service by means of a "transport" protocol. The satellite networks may allow the users to perform those applications at lower cost and higher speed. It may also extend their range : instead of transfering a "small" file of a few 100 kilo octets on an X25 network, the user will consider the transmission of a whole database of some 100 Mega octets.