AN LSI CONTROLLER FOR SATELLITE SWITCHED TDMA

D. P. DOBSON and A. E. RING
Hawker Siddeley Dynamics Limited
Stevenage, U.K.

Abstract. Satellite-switched time division multiple access (SSTDMA) is a high efficiency technique likely to be exploited soon in communications satellites. Besides needing special equipment at the earth stations, two major units are required in the satellite to complete the system, namely, a microwave switching matrix (MSM) and a distribution control unit (DCU). The MSM connects any one uplink to any one downlink in a "telephone exchange" arrangement and the DCU, which can be programmed from the ground, controls the switching of the MSM connection patterns in real time. The DCU also provides the master clock for the complete SSTDMA system. This paper describes an implementation of a DCU using custom large scale integration (LSI), which employs only four chips (two each of two types) for the entire logic to control an 8 x 8 MSM. This offers considerable advantages in terms of mass, power consumption and reliability. Some of the benefits and problems of the LSI implementation are also discussed.

1. Introduction. With the current exponential increase in international communications traffic of all varieties, communications operators are looking for ever more sophisticated, powerful and efficient ways of carrying traffic. In the satellite-link field, techniques have so far been confined to offering what are effectively "transparent" repeaters, each connecting an "uplink" (ground transmitter to satellite receiver) to a "downlink" (satellite transmitter to ground receiver) in a hardwired manner. Several such repeaters may be carried in a single satellite, and techniques such as FDMA, TDMA and antenna cross - polarisation and spot-beaming (Space division multiplex) give each satellite a capability to carry very many channels simultaneously. However, these channels are effectively hard-wired links, with limited capability of re-allocation to allow for varying traffic patterns, such as occur at different times of the day or with the adding of new users to the service, and with only limited reconfiguration to re-optimise in failure situations.

In setting up channel allocations, a compromise has to be reached so that the channel capacities provided for each separate route give the most efficient usage of satellite resources, when averaged over all predicted traffic patterns during the lifetime of the satellite. Thus on a spacecraft serving transatlantic traffic, the numbers of channels allocated to, say, the New York - London route could be set to suit the traffic density in the morning (New York time)/late afternoon (London time), when business hours at the

two ends coincide. This would mean that these channels would be considerably under-used at other times of the day, leading to substantial inefficiency (and hence, indirectly, increased service costs to users).

If these under-used channels could be re-allocated or adjusted at slack times to serve other routes (e.g. intra -European or intra-American during their respective daytimes), this would give a potential for raising channel usage, and hence "plant" efficiency, to nearer the 100% mark. Re-allocation of channels within a satellite also gives considerably increased capability for reconfiguration to bypass faulty units such as a failed repeater, so that the useful channel capacity of the spacecraft can be optimised even when a significant part of the "plant" is lost. This re-allocation calls for new techniques in communication, which are currently being explored by satellite manufacturers, service operators and users; one such technique which shows great promise is SSTDMA (Satellite Switched Time Division Multiple Access).

To implement this technique of SSTDMA, not only did new theoretical studies have to be pursued; new protocols and operating techniques had to be established between users and operators; new equipments had to be developed for the ground communications stations and the control stations; but some totally new and relatively sophisticated hardware had to be developed for incorporation into the satellite, giving the spacecraft itself a measure of "intelligence" and of overall system control which were not required with existing "hardwired repeater link" techniques mentioned above. This satellite hardware became possible with the adoption of technologies which are relatively new in the spacecraft engineering world (with its own particular criteria of very low mass and size, very low power consumption, and extremely high reliability), but which are finding wide application in the commercial/industrial electronics market.

This paper gives a brief description of SSTDMA as a communications technique, and of the major equipment units required by it. It then presents an analysis of the technical problems associated with implementing one of these units, the satellite-borne Distribution Control Unit (which contains all the logic in the satellite needed to control the channel reallocation function and to programme new traffic patterns); and it describes a particular implementation, the way in which the major technical problems are met, and the results achieved.

2. What is SSTDMA? Satellite Switched TDMA is a technique for allocating equipment-receiver, transmitters, repeaters - to establish communications links in a programmable and variable manner, so as to optimise on equipment usage and provide overall channel capabilities for the different communication links which most nearly reflect their traffic densities at any given time. It does this by means of switching the interconnections between receivers (uplinks) and transmitters (downlinks) in real time,

time division multiplexing the various uplink - downlink interconnection patterns in such a way that in any period (say a second), not only can all links be usable, but the ratio of channel carrying capabilities can be varied over a wide range (say up to 100:1) to suit different traffic densities on the different uplink-downlink routes.

To examine the system of operation, consider an SSTDMA satellite with 4 repeaters, that is to say 4 uplinks (comprising receiving antenna and receiver down-converter) and 4 downlinks (comprising amplifier, upconvertor/frequency translator, transmitter and transmitting antenna). Each receiving or transmitting antenna would be arranged to serve one of 4 pre-allocated geographical zones on earth (each of which contains one or more ground-stations). Fig. 1 shows an “elevation view” of such a system, simplified to basic essentials.

For convenience, call the 4 uplinks A, B, C, D, and the four downlinks W, X, Y, Z. Then in a normal “hardwired repeater” (non SSTDMA) satellite, the links A - D would be routed to W - Z in a fixed manner, for instance:

A → W
B → X
C → Y
D → Z

A - D and W - Z could be wide-coverage or spot-beams, and could have overlapping service zones, so that a reasonable variety of links can be offered by this system, but once the satellite is launched, only very minor variations can be made to this allocation.

In the SSTDMA satellite, instead of the uplinks and downlinks being hardwired together, they are connected by a “matrix” of switches which allows any one of A - D to be routed to any one of W - Z at a given time (as with a crossbar-switch). In fact four such 1 -to - 1 links can be supported simultaneously, so that at one time the routes may be A→W, B→X, C→Y, D→Z (as before), and at another time they may be A→Y, B→W, C→Z, D→X, or any one of 22 other possible patterns.

The strength of the SSTDMA system lies in the fact that this switching is controlled in real-time, time-division-multiplexing the various connection patterns. This is achieved by “framing”, whereby time is divided into main repetition periods or “frames”, each of which is subdivided into a large number of time slots or “frame intervals”: for example, a frame of 750μS, comprising 125 intervals each of 6μS.

This frame interval of 6 μs is the basic unit for establishing links; at the beginning of a 6μs interval (say interval “m”) a pattern of switches is set up which provides paths for 4 links

(say A→W, B→X, C→Y, D→Z). These connections are maintained through the 6µs frame interval to the next switching time, when a new set of paths may be established for interval "m +1I" (say A→Y, B→W, C→Z, D→X). Thus within a 750/µs frame, up to 125 such sets of paths will have been connected. (The number may be less than 125 if desired, since individual paths or sets of paths can be left established for multiple frame intervals; this is completely programmable to meet the needs of the user). Fig. 2 shows part of an example connection pattern.

By repeating this frame of 125 connection time slots every 750µS, it is possible to establish a useful protocol whereby the ground stations can match their traffic transmission schedules and reception schedules with the programmed availability of connected paths. Thus with the above example, a ground station in zone A will know to transmit traffic for zone W in every time slot "m", for zone Y in every slot "m + 1", and so on. Conversely, receiver X will expect traffic from transmitters B and D respectively at these times.

For the ground transmitting and receiving stations and the satellite on-board hardware to be able to cooperate in this way, so that the correct traffic is always sent through the correct connection paths, two additional factors are needed. The first is a timing-reference, so that all elements can synchronise frame and frame-interval timing exactly (including compensating for transmission path delays). This is achieved by the satellite acting as the master timing reference, and by allocating one fixed frame interval in every frame to synchronisation activity (a discussion of this is outside the scope of this paper, but see refs. 4, 6, 8).

The second requirement is that all parties concerned (operators and users) should at all times have full knowledge of the interconnection pattern currently employed. At present it is envisaged that the patterns be reprogrammed only as required, so as to follow major changes in traffic pattern, such as described above in the Introduction. This is the approach adopted for the work described in this paper, where it is applied to an 8 x 8 switch matrix situation (8 uplinks, 8 downlinks). There is no theoretical reason why in the future this cannot be improved to something approaching on-demand assignment - however the problems here will not be solved until at least one generation of satellites later.

A further discussion of the theory of SSTDMA is outside the scope of this paper. The interested reader is referred to refs 1 - 10.

3. Equipment Required. The major equipment units required to operate SSTDMA are:

a) Data storage/multiplexer/transmitter)
b) Data receiver/storage/demultiplexer) at ground stations
c) Ground station synchroniser
d) Microwave switching matrix (MSM) in satellite
e) Distribution Control Unit (DCU) in satellite

(This is in addition to the usual RF uplink and downlink units).

This paper is concerned with the satellite-borne equipment and in particular with the DCU.

4. DCU Functions. The DCU is shown in schematic form in Fig. (3) and has the following major elements:

i) a memory to store the current up/down connection pattern for each frame interval;

ii) a second identical memory for use in reprogramming operations to avoid disturbance of normal operation and to use in a degraded back-up mode in the event of a failure in the first memory;

iii) reprogramming logic interfacing with the satellite telemetry and telecommand (TT&C) system to receive new patterns from ground control and to echo these for verification;

iv) a set of output decoders and drivers to convert the pattern data into a form suitable for driving the individual switches of the MSM;

v) a system of "static junction control bits" (SJCBs); these override dynamic MSM control signals for static control over MSM junctions thus allowing frequency division multiple access (FDMA) for parts of the system;

vi) the system master clock and timing logic;

vii) a means of continuously checking for correct operation of the system in order to warn the satellite operator of any equipment failures;

viii) a means of switching equipment states, both to bring into service a newly programmed MSM switching pattern (in the correct synchronisation) and to recover from equipment failures.

5. DCU Major Requirements. The DCU under development has been designed to drive an 8 x 8 matrix, with a 750µsec frame of 125 x 6µsec frame intervals (as described earlier). Switching times at the outputs of the unit are <75nsec including all jitter, risetime and skew between outputs.

In order to achieve acceptable budgets for the satellite (which ultimately determine whether it is economically viable to launch an SSTDMA system rather than one of the simpler fixed systems) the complete DCU must be achievable in a unit which has the following target parameters:

mass	< 1.5kg
size	$<(200mm)^3$
power consumption	< 1.5 watts
reliability	>.88 for 7 years

and the unit must be able to withstand satellite environments, which include wide temperature range (-15°C to +55°C), launch vibration and shock and high incident radiation.

6. Selection of means of implementation. For the 8 x 8 MSM, 64 control signals are required at the DCU/MSM interface and 125 patterns of these signals are required sequencing at a rate of 160kHz. In order to achieve the reconfigurability discussed in sections 1, 2 it must be possible to modify the sequence of patterns. The data storage requirements of the DCU are then that the system shall be capable of reading, storing and outputting 125 words of 64 bits. Because of the fact that only one microwave switch per column of the MSM can be on at any one time, it is possible to eliminate redundant information and store only 125 words of 24 bits; the 24 bits are organised as 8 x 3 bit groups and each 3 bit group uniquely selects one switch from a particular column of the MSM.

A storage requirement of 125 words x 24 bits is large in conventional medium scale integration (MSI) and printed circuit board terms, when considering spacecraft budgets. Alternative means of data storage must be considered, possibilities being: core memory, plated wire memory, hybrid film packaging of “off the shelf” MSI and LSI elements, and custom LSI design. The support circuitry overhead for core and plated wire systems are too heavy for this size of memory; also the truly random access nature of core and plated wire storage is not necessary or ideal for this application.

To build the system in hybrid film form using standard storage and logic elements is feasible. However the memory support circuitry is fairly random in nature and requires a large number of standard SSI, MSI and LSI chips and the number of pins required to

interconnect individual hybrid packages is high. An estimate based on the use of conservative 40 pin packages is that five hybrid circuit designs would be required with seven packages being used for the system. Additionally the output components would have to be 54 series TTL consuming considerable power and having a poor overall MTBF. The system would still require a highly reliable high stability oscillator.

To build the system with custom monolithic LSI is feasible with the present state of advancement of this technique. To make full use of the ability to reduce package count it is however advantageous to use LSI chips whose size is greater than that normally considered acceptable for commercial viability based on a high volume market. It is then possible to partition the system such that only two chip designs are required, again in addition to a highly reliable high stability oscillator.

Within the field of custom monolithic LSI there are many different types of process each of which has its own attributes. The possibilities considered for use in the DCU were CCD, PMOS, NMOS, CMOs* and bipolar. CCD is a relatively new technique and is ideal for serial type memories but is not capable of operating at the speeds required nor is it suitable for the support logic. PMOS is the most established process; and the one furthest advanced towards space qualification; it is capable of packing the largest number of circuit functions on to a single chip but is slow in comparison to NMOS, CMOs and bipolar. Nevertheless it is fast enough for operation at the speed required for the DCU. By use of the design technique of dynamic logic, the power dissipation can be minimised and problems of asynchronism and skew within the logic can be avoided; this requires a two phase-clocking circuit as an overhead. NMOS has all the performance capabilities of PMOS but it is not possible as yet to integrate as much circuit function on to one chip; it is capable of faster operation than PMOS and is more easily interfaced to other types of component. CMOs has the capability of very low power operation with no support circuit overhead; it is inferior to NMOS on grounds of packing density of circuit function per chip. Bipolar processes have the capability of extremely fast operation but this cannot readily be married with high packing density per chip and low power consumption except by use of the relatively unproven design technique of Injection Logic.

* CCD = Charge Coupled Device
PMOS = P-channel MOS
NMOS = N-channel MOS
CMOS = Complementary MOS

The process selected was PMOS since those areas in which it is inferior to the other candidates are not important for this application. In the interest of power economy and packing density the technique of four phase dynamic logic was adopted; the impact of

requiring extra clocking circuitry is minimal in this application as the main part of this is already required for the overall reference clock.

7. Problems Associated with the use of Custom LSI. There are various problems associated with the combination of custom LSI design and space usage and these fall into the two broad categories of purely technical problems and project management problems. These are in addition to the more widely-recognised design considerations (such as pin limitations, thermal design, chip area, new design methodology) associated with all LSI.

These are discussed in several standard works (ref. 11).

7.1 Technical Problems. The problem of degradation of MOS devices due to radiation damage has received considerable attention (refs. 12, 13, 14). However this tends to be oriented towards short term nuclear environments and in particular electron bombardment whereas the radiation environment of a geostationary orbit where a communications satellite is likely to reside is relatively benign but of much longer term (7 years). The philosophy of screening out incident radiation using the metal of the unit case has been adopted, on the basis that a level of about $1gm/cm^2$ reduces the radiation experienced to an acceptable level.

It is customary to assess the reliability of an electronic unit for use in spaceborne equipment in numerical terms starting from the individual component level. The boundary at which a specific circuit function has a numerical reliability attributed to it becomes much more vague for LSI devices since the close proximity of individual devices on a chip destroys the basic assumption of independence between one device and another. Early attempts at determining the reliability of a chip were somewhat crude being based either on chip size or number of “Gates” but the procedure now adopted (ref. 15, 16) attempts to take account of environment, circuit complexity, chip temperature and the difference in confidence levels between market standard devices and custom designed devices.

With the use of LSI the problem of rigorous commissioning and testing becomes far more complex. This is due in part to the pure increase in amount of logic associated with a unit of sensible physical size and in part to the fact that access to the internal paths of the LSI device is extremely difficult and limited. To combat the problem of increased complexity it is virtually essential to use a computer aided test system but it is still impracticable to test the operation of the devices for all theoretical initial states, input sequences, and real-time situations. The philosophy which has been followed is to test only from those conditions which can be expected as initial states in operational use. As an aid to this, test reset inputs have been included as extra circuitry on the chip. Also wherever possible the chip design has been carried out to allow areas with identifiable forbidden states or lockout conditions to count out of these conditions or be accessible to telecommand thus giving only a

temporary malfunction. This technique is however not always compatible with optimum circuit design. An example of this is the use of a pseudo-random code generator as a counter instead of the normal binary counter. This pseudo random generator consumes considerably less silicon area but has a lock-out condition which does not count out; to overcome this a reset input is included.

The sensitivity of the memory to various sequences of patterns on a memory of 125 words by 24 bits cannot be checked out rigorously (it would take 10^{30} years to do this once). There is no real way round this except to accept that any pattern sensitive failure will be repeatable and can be combatted by rescheduling the pattern sequence in the memory. This rescheduling would have no detrimental effect on the end user.

Malfunctions experienced during the commissioning phase of an LSI device can be investigated to a limited extent by microprobing the chip but this action may itself so modify chip operation due to capacitance loading as to render the results meaningless. In this event a solution is to move away from the traditional method of bench testing and to analyse a malfunction on a computer transient simulation or logic simulation by varying circuit parameters to simulate the fault.

7.2 Project Management Problems. The manufacturers of LSI are in general not interested in using their design engineers to carry out low volume custom design work since their priority inevitably lies with high volume long market lifetime products. One alternative solution to this problem of device design is for the customer to buy in expertise or receive training to a stage where he can carry out the specialised design work himself (which may be a very expensive exercise, especially when the capital costs of special equipment are added in). Another possibility is to employ one of the specialist custom LSI design companies, although this can possibly present problems in that these companies may not be familiar with spacecraft design disciplines, so that a close liaison between system and subcontractor engineers is required. This last approach was employed in this case. In all events it is necessary for the system design engineers to acquire for themselves a knowledge of the new design methods of LSL

As discussed previously, complex test equipment is required to commission a system using LSI adequately. If the proprietary systems are considered it must be realised that these are very costly (£200K to £ 500K) and if special computerised test equipment is designed for a particular application (as in this development) it must be remembered that the cost of this may be high, possibly comparable to the cost of the design and manufacture of the LSI devices.

Quality assurance of LSI devices for spacecraft use is a difficult problem. To qualify each individual chip design would be prohibitively expensive. The alternative approach is to qualify a particular process and a particular manufacturer of chips using that process and then to establish a set of standard design layout rules and process parameters. Each individual chip design is then carried out using these rules and parameters and is treated much as the design of a unit from discrete components has been previously. Present philosophy in the U. K. is to base work on the British Post Office D4000 or BS9400 series specifications; these are blanket specifications with appendices added to cover particular processes. A significant amount of work is also involved in the qualification of suitable packages for LSI devices.

8. Details of DCU System Design. As discussed above, a system using custom-designed four-phase PMOS chips was adopted. The detailed system design in this technology will now be discussed. A total of 4 chips, two each of two types, is needed to fulfil the requirements outlined earlier. These contain the whole of the logic function; in addition two high-stability temperature-compensated crystal oscillators (TCXO's) are provided, one main and one standby in case of failure of the main; and a number of discrete and analogue components are needed for power regulation and smoothing, level-shifting etc. The whole unit (apart from two chips) is contained in a spacecraft type housing measuring approximately 140mm x 240mm x 40mm. (The two chips not included are the two custom decoder/driver chips, which are physically mounted within the Microwave Switching Matrix unit, although logically they may be regarded as an extension of the DCU.)

Fig. 4 shows the partition of the main function of the DCU among the four chips. There are two Memory and Control chips, called A1 and A2, and two Decoder/Driver chips B1, B2. Chips A1 and A2 form an online/standby pair: each is fully equipped to drive the chips B, and hence the MSM unit, on its own. When chip A1 is "online", its memory contains the data pattern currently being used to control the matrix, while A2 is available offline for reprogramming (to implement the next traffic change); or, if A2 is programmed with an identical copy of A1's memory, it is able to act as a constant self-check of the correct functioning of A1, and to flag the ground should any inconsistency occur. Similarly chip A2 may be online and A1 offline - the chips are identical.

The two chips B1 and B2 provide line decoders and output drivers to convert the stored output from the online chip A into the right levels, and on the right number of lines, to control directly the individual switches of the MSM. Here the two chips B are not online/standby, but are both online in direct parallel operation, each driving 32 of the 64 matrix points. (The chips could not readily be combined for practical reasons: their combined pin count would exceed that on an acceptable space-approved package, and their power consumption and size would present unacceptable problems.) Within chip A,

the memory section is required to output data patterns in a fixed sequence manner rather than a random access manner. This allows the use of shift register type memory which is ideally suited to MOS dynamic logic, and the memory has been designed as 24 shift registers of 125 bit length; the principle of operation is shown in Fig. (5).

Circuit operation of the four phase dynamic logic allows data to be output from the memory either at the beginning of each frame interval or half way through each frame interval, and by alternating the output time for successive shift registers it is possible to multiplex the 24 output signals on to 12 lines hence reducing the package pin count and interconnecting cable mass. The clocked synchronous nature of the four-phase circuits also overcomes all problems of asynchronism and logic skew.

9. Results. At the time of writing (April ’76), the DCU is in the process of manufacture; the development version is scheduled to be commissioned by the autumn of 1976, and it is likely that a development MSM unit (from another source) will be available by summer/autumn 1977.

The main parameters arising from the DCU design are:

Mass	700 gms)	For the development version
Size	140mm x 240mm x 40mm)	which is approximately flight-
Power consumption	0. 95 Watts)	representative.

Reliability 0. 899 for 7 years (this is limited to a large extent by the failure rates of the TCXO’s, and could be improved if necessary; it also includes the type B chips).

The Memory and Control chip A is approximately 7.1mm x 7.5mm (0. 28" x 0. 295") and contains some 20,000 transistors, chip B being somewhat smaller (4mm x 4mm). The unit design meets its functional specification completely, and offers small enhancements in one or two areas to simplify in-operation use. A composite drawing of chip A is shown in Fig. 6.

10. Conclusions. This paper has described the technique of SSTDMA and its advantages for future generations of communications satellites (it is being considered for versions both of the Intelsat range and of the ECS range of spacecraft).

To implement this technique, a number of technological advances in spacecraft engineering are called for, and one of them, involving the use of custom LSI, has been described. This may be seen as a forerunner of a series of applications for custom LSI in space, where its advantages of low mass, size and power consumption and high unit reliability make it very attractive.

Acknowledgements. The work described was performed by Hawker Siddeley Dynamics Limited, under contract to the International Telecommunications Satellite Organization, Intelsat (contract IS-563).

The authors wish to thank Hawker Siddeley Dynamics and Intelsat for permission to publish this paper, and Messrs. R. S. Cooperman and I. H. Williams for their assistance in its preparation and in the work described.

Views expressed herein are those of the authors, and not necessarily those of Intelsat.

References.

1. R. S. CoQperman and T. Dobyns, “A Distribution Control Unit for Satellite - Switched Communications”, International Telemetry Conference, Los Angeles, Oct. 72.

2. W. G. Schmidt, “An on-board switched multiple-access system for millimeter wave satellites.” Proc of the Intelsat/IEE International Conference on Digital Satellite Communications, Nov. 69, London, pp399-407.

3. R. Rapuano and N. Shimasaki, “Synchronisation of earth stations to satellite-switched sequences,” Proc of the AIAA Fourth Communications Satellite System Conference, April, 72, Washington, AIAA Paper No. 72-545.

4. W. G. Schmidt, “Satellite -Switched TDMA : Transponder -Switched or Beam-Switched?”, Proc of the AIAA Fifth Communications Satellite Systems Conference, April 74, Los Angeles, AIAA Paper No. 74-460.

5. W. L. Pritchard and P. L. Bargelhini, “Trends in Technology for Communication Satellites”, Astronautics and Aeronautics, pp 36-42, April 72.

6. W. G. Schmidt and R. S. Cooperman, “ Satellite -Switched SDMA/ TDMA System for a Wideband Multibeam.Satellite”, 1973 International Conference on Communications, Seattle, pp 12-7 to 12-12.

7. K. Nosaka and T. Muratani, “New Satellite Transponder with Regenerative Repeater and Higher Order DPSK for TDMA/SDMA System”, Proc of the Second International Conference on Digital Satellite Communications, Paris, 1972.

8. M. Asahara et al., “Synchronisation and acquisition problems for SDMA/SS-TDMA: results of hardware studies”, Proc of International Conference on Satellite Communications Systems Technology, London, April 75, IEE Conf. Pub. No. 126 pp 220-226.

9. Y. Ito and M. Kyogoku, “SDMA on-board switching system using rearrangeable multi-stage network”, ibid pp 227-232.

10. “Study of On-board switching for European TDMA Telecommunications Satellites”, publication by Laboratoire Central de Telecommunications, Velizy-Villacoublay, France, under ESRO Contract 1768/72 SW.

I 1. The Engineering Staff of AMI Inc., “MOS Integrated Circuits”, Van Nostrand Reenhold 1972.

12. B. L. Gregory and C. W. Gwyn, “Radiation effects on semi conductor devices,” Proc. IEEE, Vol. 62 No. 9, pp 1264-1273; Sept. 1974.

13. A. G. Holmes Siedle, “Radiation -induced ionisation effects in solids: a review of research problems,” Proc IEEE, Vol. 62 No. 9, pp1196-1207; Sept. 1974.

14. Burrell, Wright and Watts, “An analysis of energetic space radiation and dose rates,” NASA TN D4404; Feb. 1968.

15. J. L. Flood, “Quality Assurance aspects of Custom LSI, Proc.” AGARD Lecture series 75; April 1975.

16. E. D. Colbourne, G. P. Coverley, S. K. Behara, “Reliability of MOS LSI Circuits”, Proc IEEE, Vol. 62 No. 2, pp 244-259.

This paper is based upon work performed under the sponsorship of the International Telecommunications Satellite Organization (INTELSAT). Any views expressed are not necessarily those of INTELSAT.

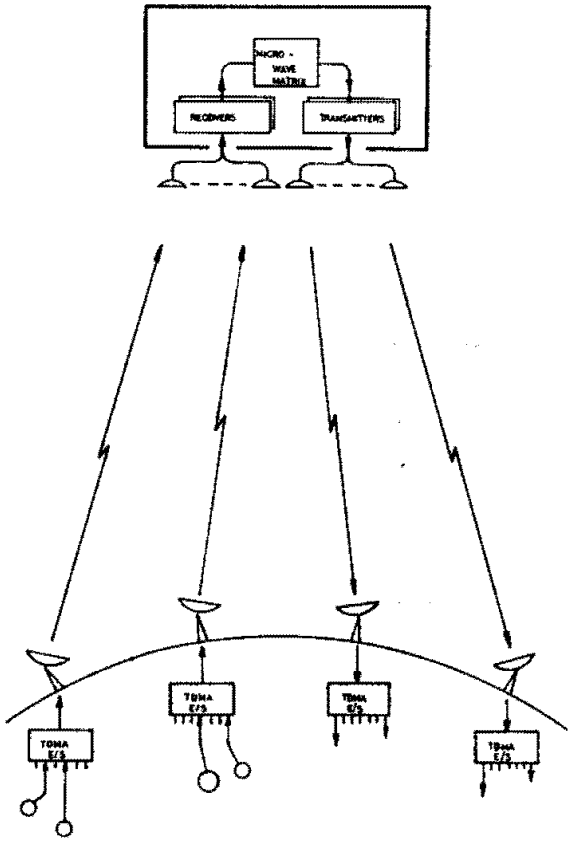

FIG. 1 TOTAL SSTDMA SYSTEM

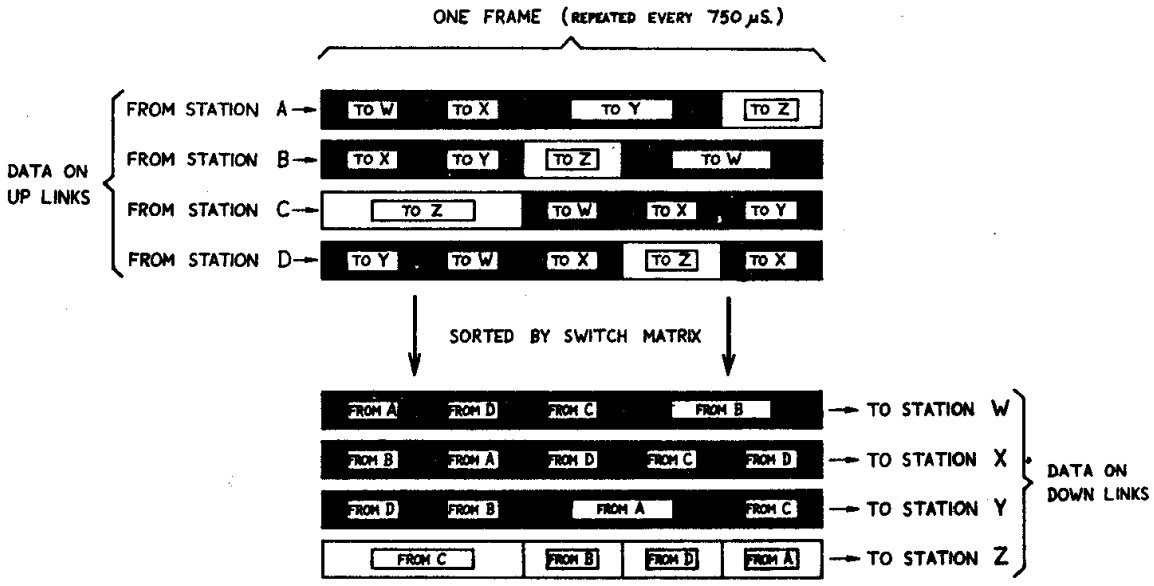

FIG. 2 PART OF EXAMPLE CONNECTION PATTERN

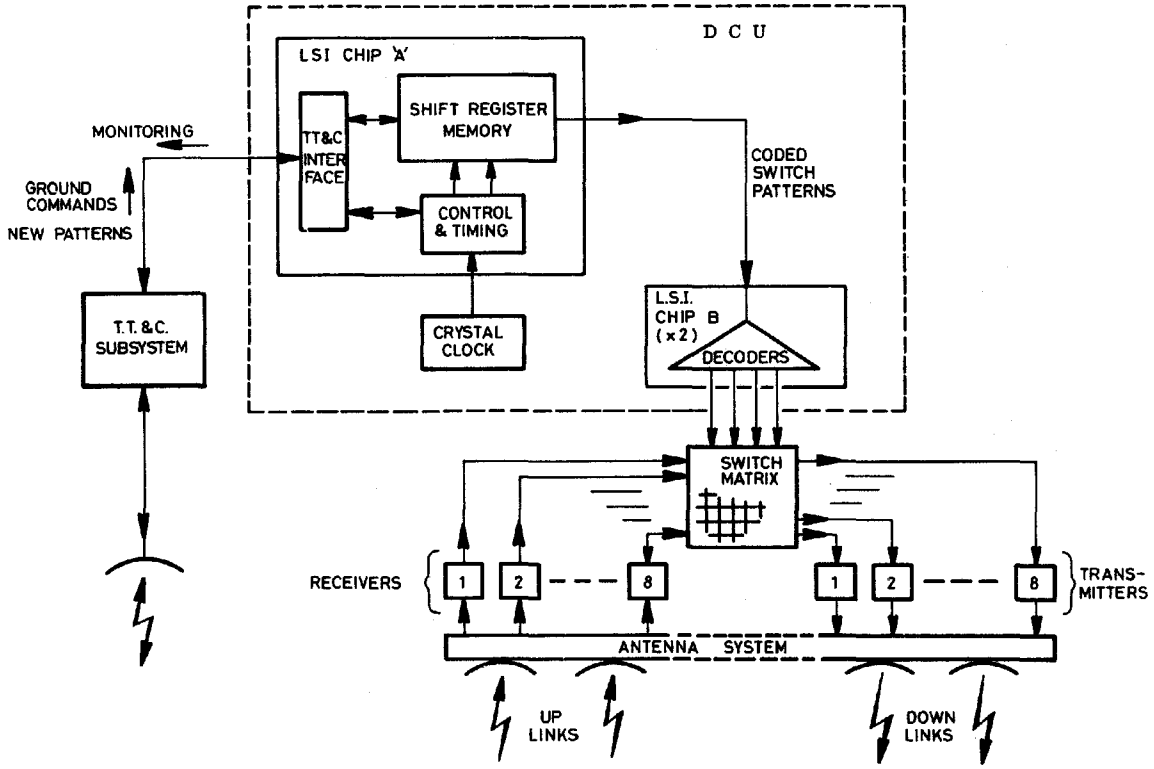

FIG. 3 MAJOR DCU FUNCTIONS WITHIN SATELLITE SYSTEM

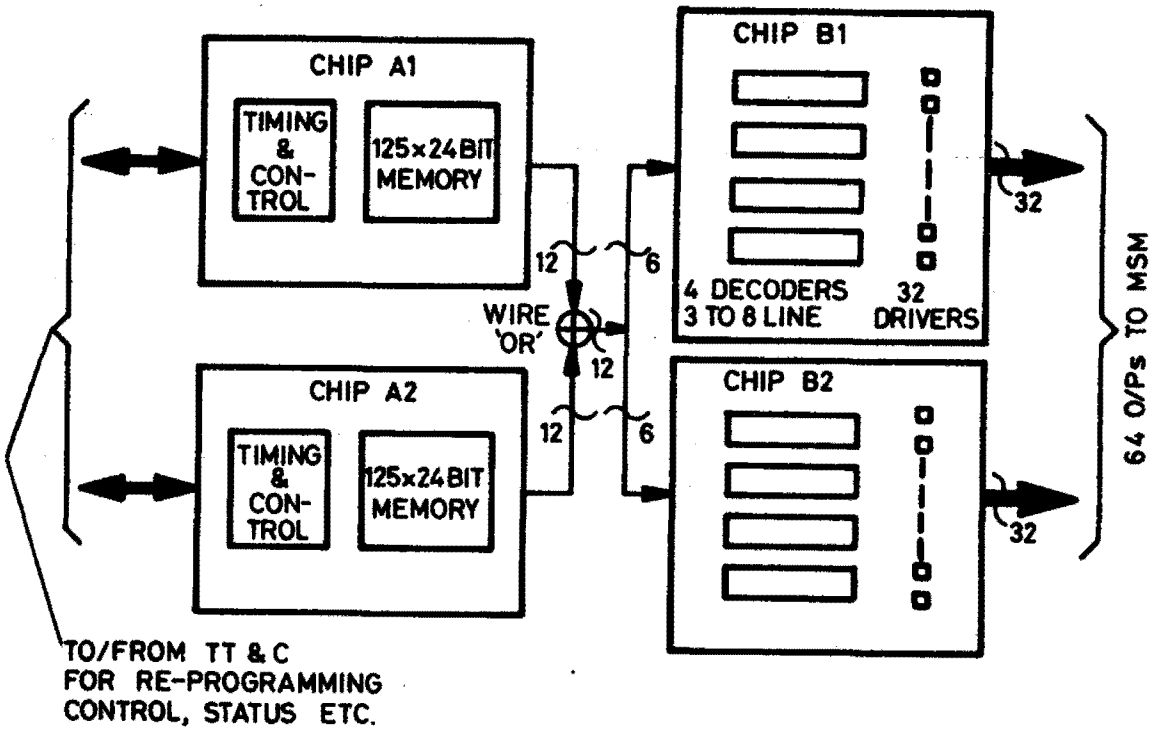

FIG. 4 PARTITION INTO FOUR CHIPS

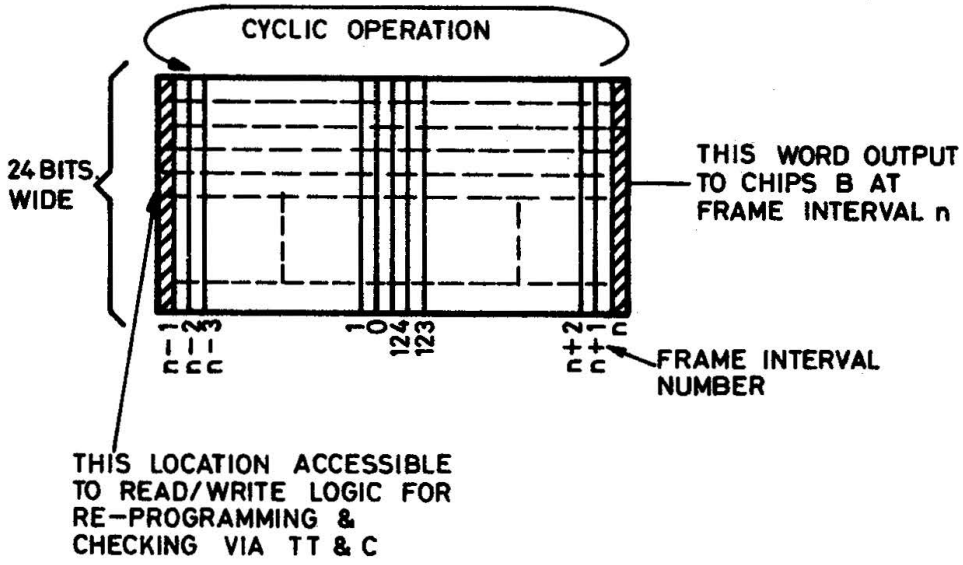

FIG. 5 OPERATION OF CHIP A MEMORY

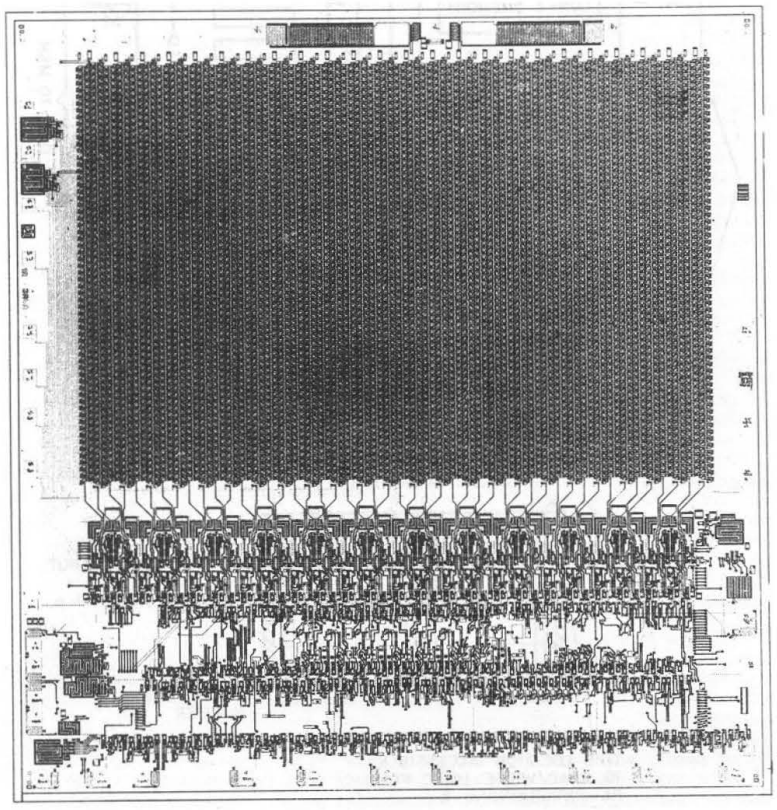

FIG. 6 CHIP A LAYOUT DRAWING