extra 256K of LCM in August.

7614 256K of LCM

7613 512K of LCM

RCC/A/P49
RCC/B/P9
RCC/C/P26
RCC/M/P51
RCC/P/P12
RCC/S/P51

FUTURE PROSPECTS - AS SEEN FROM APRIL 1975

Introduction

1. The provision of computing facilities is a complex, and sometimes a very frustrating task. A large established workload, and the long lead time which must elapse before ordered equipment is delivered, implies inertia. Further, since no manufacturer provides hardware, software and services which are superb in every respect, one must to some extent take the bad with the good. It is simply not realistic to propose to obtain magnetic tapes from one manufacturer, disc drives from another, software from a third party, and so on. In essence we are faced with package deals which, it is hoped, are optimal in some overall analysis - subsequently bad elements must be replaced or overcome. In a short paper it is not possible to provide a comprehensive account of the present position, or of developments which may reasonably be expected in the future. Further elaboration will be given in discussion in committee.

2. The six months since the first ACC meeting was held have seen a steady increase in workload. Contracts have been executed for the installation of improved or novel peripheral systems. The configuration of the central computing elements and of the communications network has evolved. New releases of software have occurred, and ideas have been tested. In isolation none of these factors would have been sufficient to trigger reconsideration of medium term goals based on a revised GEMINI configuration incorporating additional peripheral devices (see Figure 1). The medium term covers the next two to three year period during which saturation of the Cyber 76 processor in prime shift will develop. In combination, however, the factors are sufficient to call into question the adequacy of GEMINI even in the medium term.

3. Production of the Cyber 172 mainframes has been subject to a succession of delays. Further the NOS operating system is delayed, to the extent that there is no alternative to the SCOPE 3.4 operating system for the GEMINI 172s. DCR has formed the view that two Cyber 172s using INTERCOM (with an absolute minimum of modifications) are incapable of outperforming the 3600. Therefore, following acceptance of a mainframe, some months must elapse before TED (file editor) and RIO (remote input/output device driver) are implemented in the SCOPE 3 environment. It is clear that the GEMINI configuration, providing TED and RIO through the network, cannot be available before July 1976. It is obvious that various measures must be adopted in the short term as a matter of expediency. To avoid wasted effort, such developments should be consistent with a long term plan providing full backup of all central services. A configuration termed DOUBLE is provided (see Figure 2) as a basis for discussion.

.../2

4. In the following paragraphs various elements of the computing service are reviewed in terms of status, limitations and proposed developments.

Space

5. The GEMINI configuration, the 3600 and the 7611 station can just be accommodated on the computer floor (after Stage 1 of the building extension programme is completed) and within the capacities of the existing mechanical services (chilled water and air conditioning). Since new equipment must be installed and brought into service before obsolete equipment is de-commissioned, the principal factor to be considered in the provision of a service more powerful than GEMINI, is site preparation. Perhaps the work could be done as an extension to Stage 3 of the DCR building works in progress, and scheduled for completion in late 1976. Even with such special consideration, it is difficult to see completion of mechanical services, needed for the additional computing equipment, before the 1977/78 financial year. It is noted that the 3600 mainframe occupies a dominant proportion of the floor space available for redeployment in the medium term. The 3600 cannot be retained in order to provide some minor facility - in fact the 3600 must be phased out as soon as possible.

3600 load

6. The 3600 is no longer used for job processing. With the introduction of blocked document transmission the interrupt load has been significantly reduced - three principal components remain: (1) CIO (central input/output); (2) disc transfers involved in CIO and RIO spooling, and in C, M and K TED commands; and (3) interactive service requests. At peak periods the interrupt rate approaches 100 (say 60 + 30 + 10) per second. Transfer of the six CIO devices (four printers, card reader and card punch) to the Cyber 76 would remove the dominant term and reduce the second term. Transfer of the RIO function to the Cyber 76, so bypassing the 3600, would further reduce the second term. These transfers would probably extend the interactive service by about 20 users - i.e. 84 instead of 64 users as maximum "workable" load.

CIO station

7. It is proposed to replace the four central printers (501s) by three new printers (580-16s). While it is electronically possible to connect the new printers to the 3600, this configuration is not supported by Control Data Australia (CDA) maintenance services (since the 3600 is obsolete). CIO peripherals (3 580-16 printers, 1 405 card reader and 1 415 card punch) are normally connected to Cyber stations as in the GEMINI configuration. It is believed that CDA would maintain peripherals directly attached to the Cyber 76, as in the DOUBLE configuration, provided that DCR developed the engineering test (SMM) programs. Work is well advanced on the CIO

.../3

DOUBLE

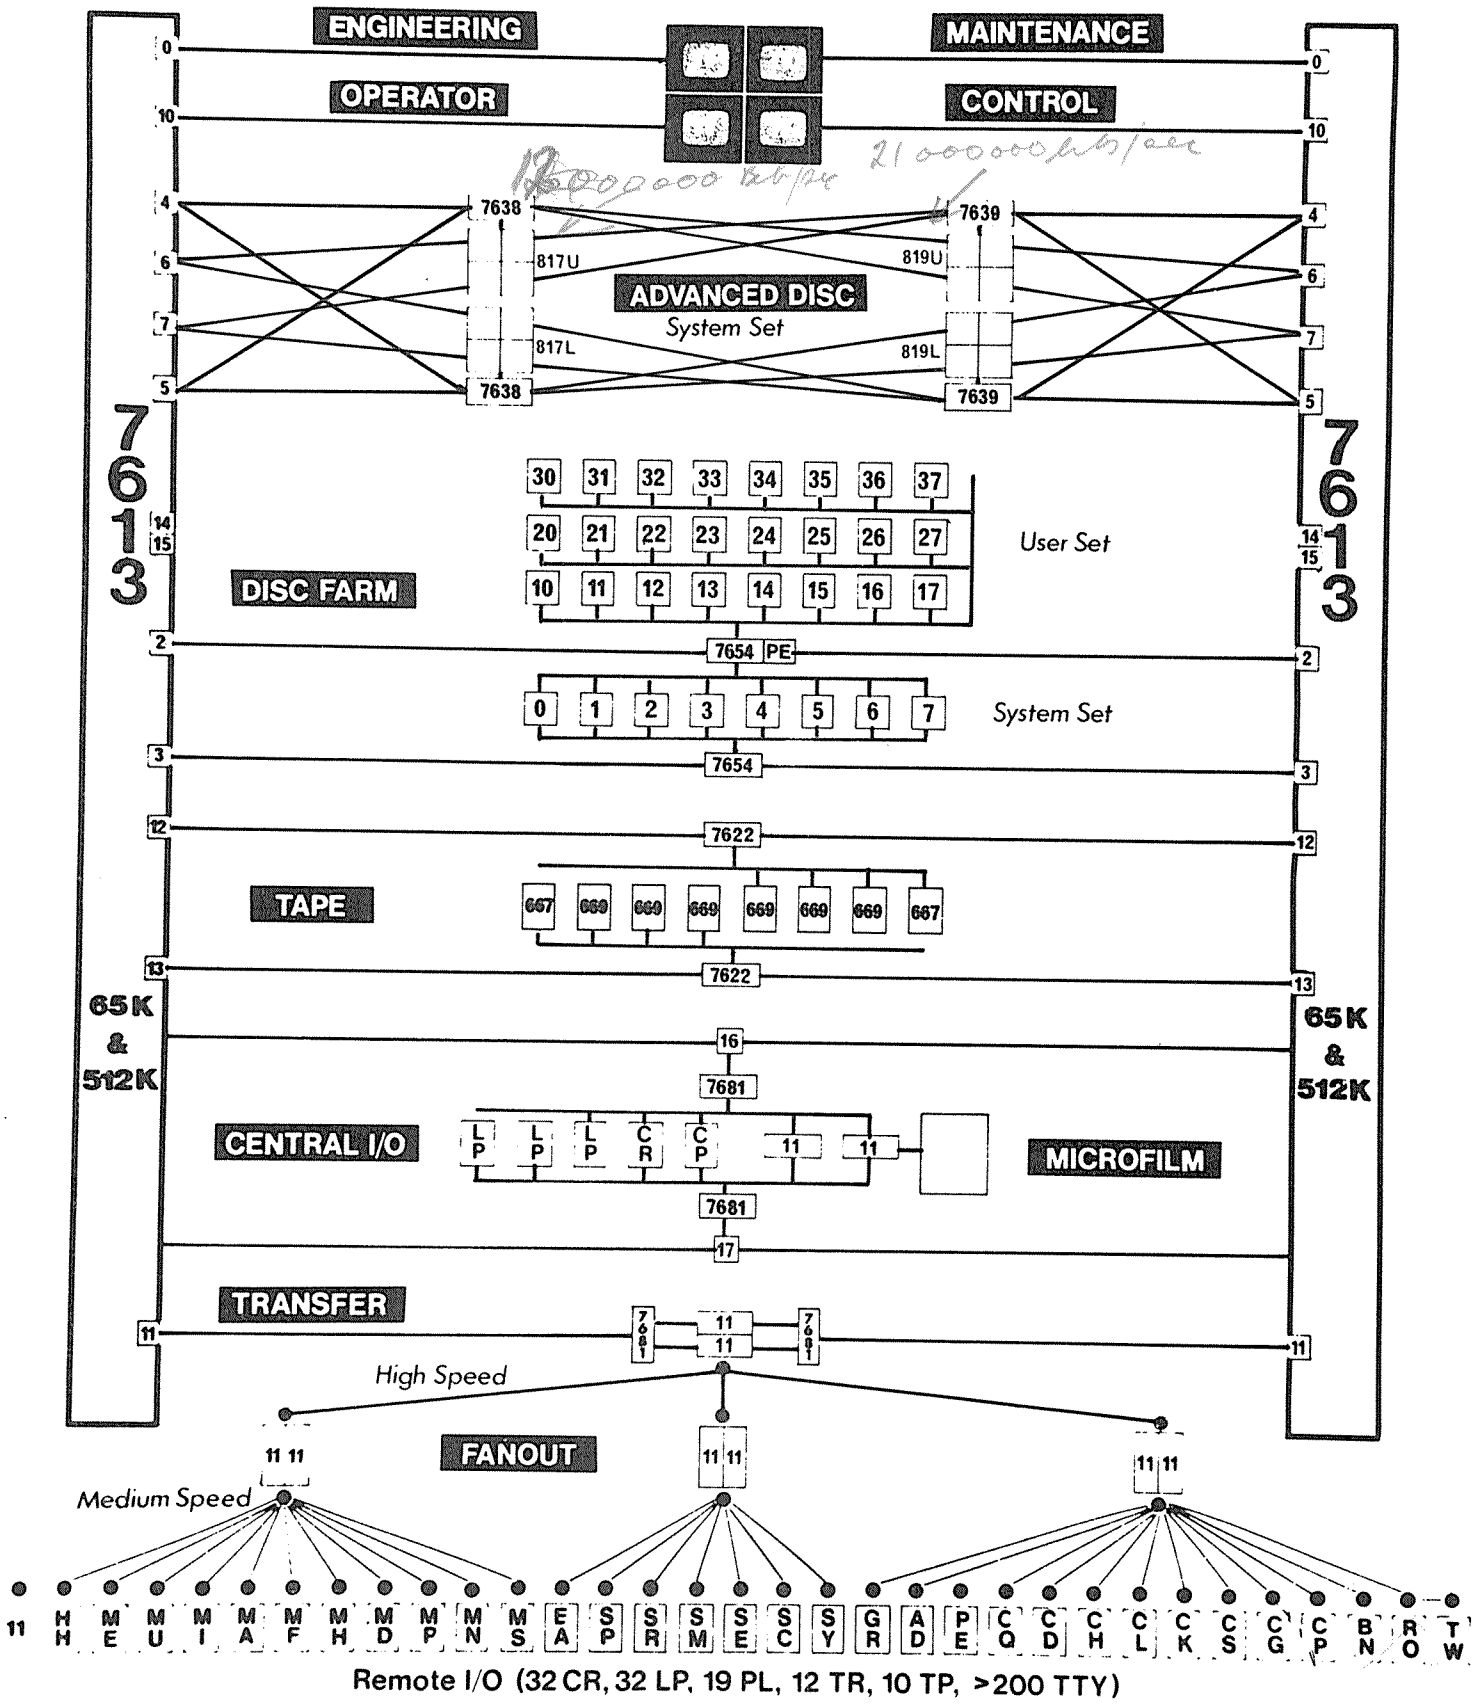
ENGINEERING
MAINTENANCE
OPERATOR
CONTROL
7638
7639
817U
819U
ADVANCED DISC
System Set
817L
819L
7638
7639
7613
7613
30 31 32 33 34 35 36 37
20 21 22 23 24 25 26 27
10 11 12 13 14 15 16 17
User Set
DISC FARM
7654 PE
0 1 2 3 4 5 6 7
System Set
7654
7622
TAPE
667 669 669 669 669 669 669 667
7622
65K & 512K
65K & 512K
7681
CENTRAL I/O
LP LP LP CR CP 11 11
MICROFILM
7681
TRANSFER
7681 11 11 7681
High Speed
FANOUT
11 11
11 11
11 11
Medium Speed
11
HH ME MU MI MA MF MH MD MP MN MS EA SP SR SM SE SC SY GR AD PE CQ CD CH CL CK CS CG CP BN RO TW
Remote I/O (32 CR, 32 LP, 19 PL, 12 TR, 10 TP, >200 TTY)

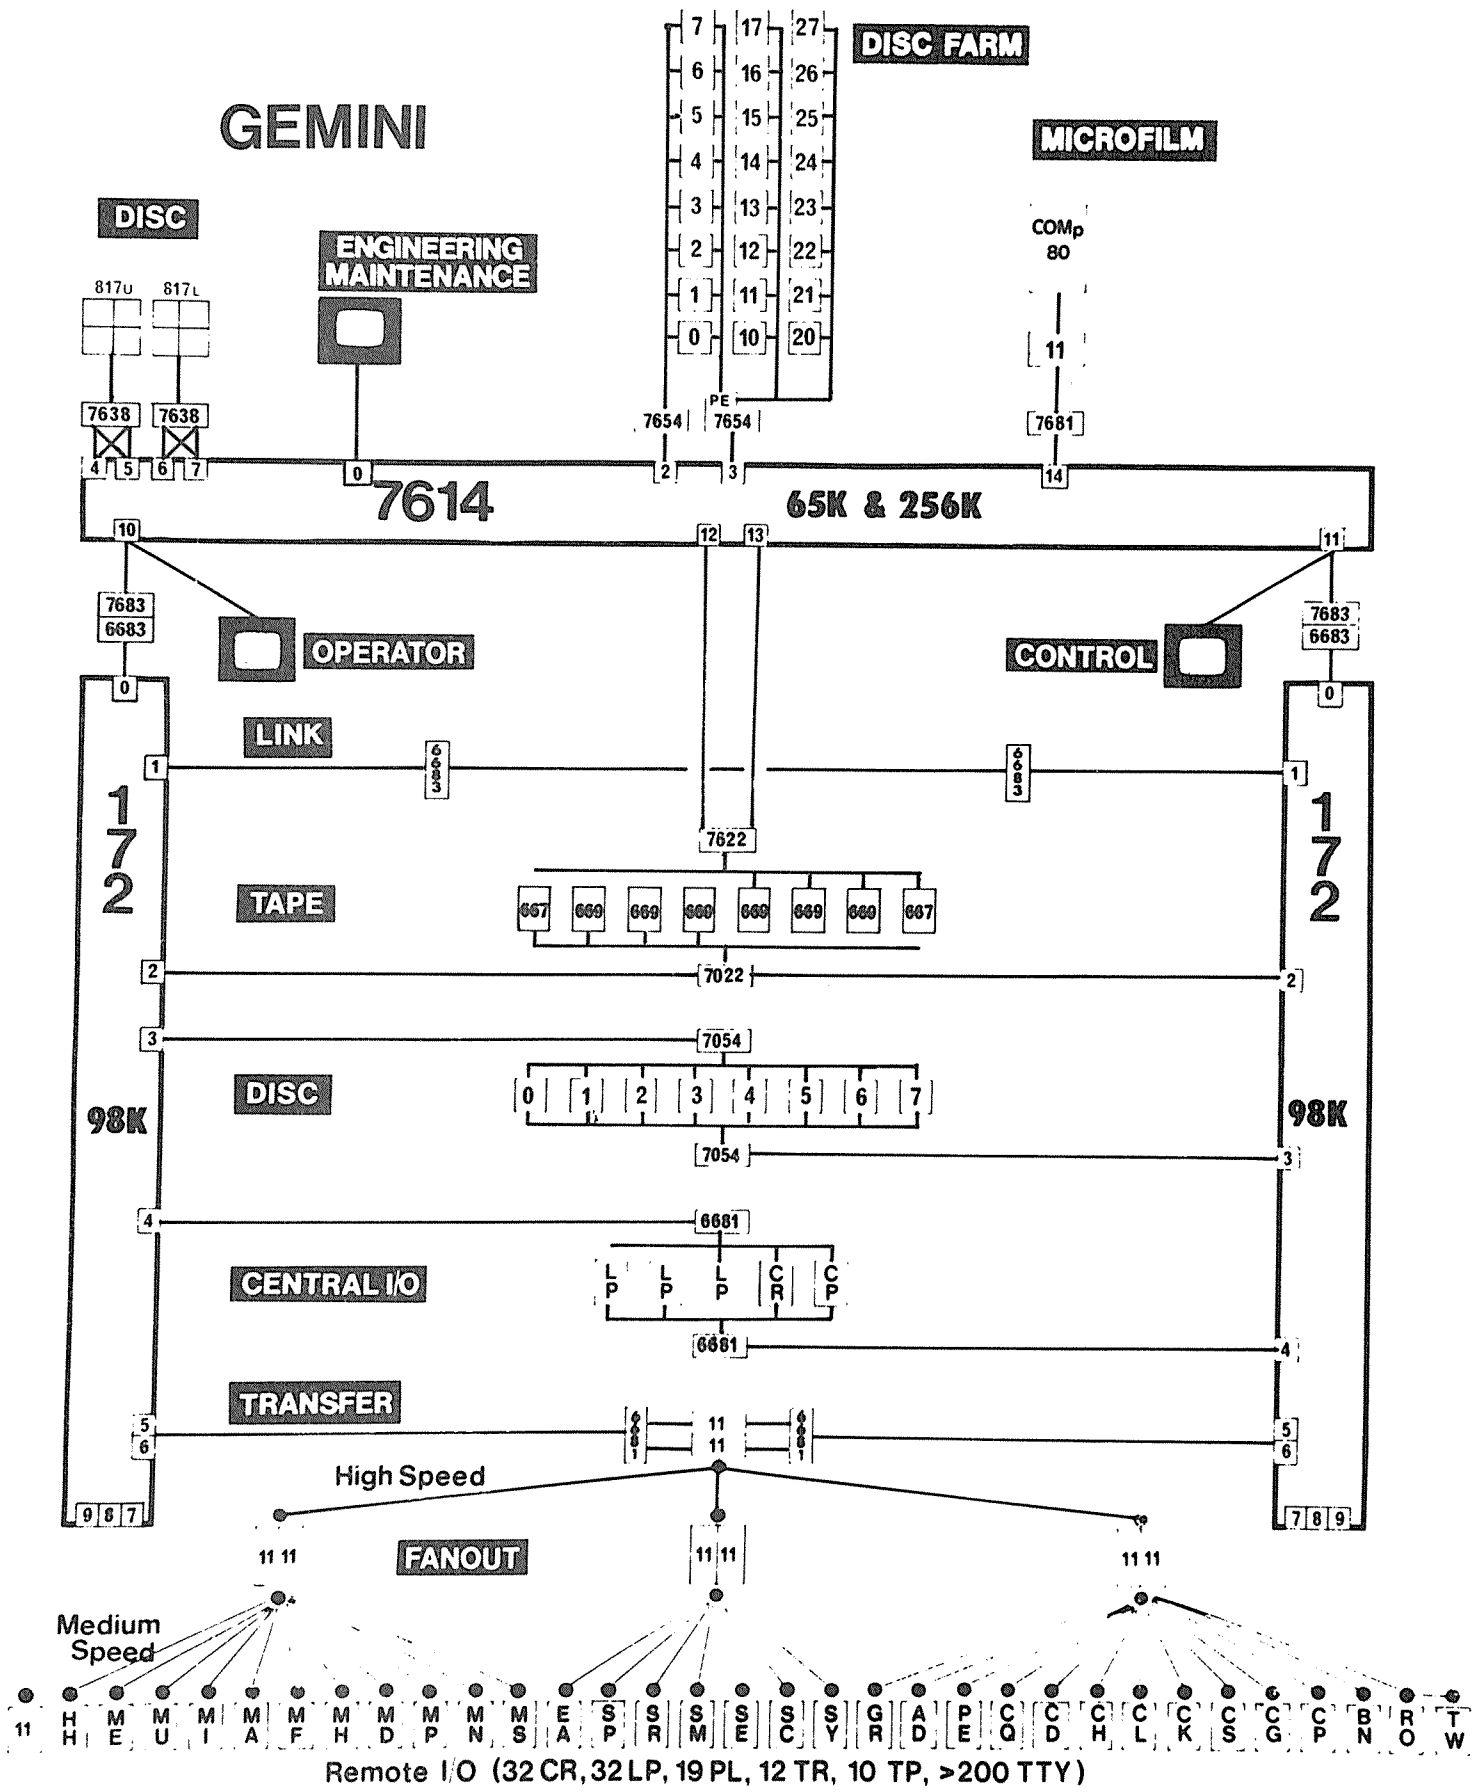
GEMINI
DISC
817U
817L
7638
7638
ENGINEERING MAINTENANCE
DISC FARM
7 17 27
6 16 26
5 15 25
4 14 24
3 13 23
2 12 22
1 11 21
0 10 20
7654
PE 7654
MICROFILM
COMp 80
11
7681
4 5 6 7
0
7614
2
3
14
65K & 256K
10
12 13
11
7683 6683
OPERATOR
CONTROL
7683 6683
0
0
LINK
1
1
6683
6683
172
7622
TAPE
667 669 669 669 669 669 669 667
7022
2
2
3
7054
DISC
0 1 2 3 4 5 6 7
98K
98K
7054
3
4
6681
CENTRAL I/O
LP LP LP CR CP
6681
4
TRANSFER
5
6
6681
11
11
6681
5
6
High Speed
9 8 7
7 8 9
11 11
FANOUT
11 11
11 11
Medium Speed
11
HH ME MU MI MA MF MH MD MP MN MS EA SP SR SM SE SC SY GR AD PE CQ CD CH CL CK CS CG CP BN RO TW
Remote I/O (32 CR, 32 LP, 19 PL, 12 TR, 10 TP, >200 TTY)

station which is scheduled for completion shortly after the arrival of the printers in third quarter 1975. It is noted that the central printer load currently fluctuates between 40 and 50 thousand pages daily, i.e. up to a million print lines daily. DCR hopes that central printing will be held at this figure through use of microfilm.

COM station for microfilm

8. The delivery of the Information International Incorporated (III) COMp80 microfilm device together with automatic film processors and ancillary gear will occur as soon as the necessary site preparation is complete, say July 1975. It is proposed to operate the COMp80 on-line to the Cyber 76 in order to simplify interfacing and software problems. The device will be driven by a COM station, i.e. peripheral processor (PP) code closely related to that used in the CIO station. Various aspects of the COMp80 will be discussed in detail at a workshop to be held at DCR Canberra in the near future, (June 3-4).

Permanent files (PFs)

9. The system set continues to be under great pressure in spite of the recent increase in capacity to around 700 Mc (million characters) for PFs plus a dedicated 400 Mc for scratch files. During 1975 it is proposed to progressively increase the set by the addition of disc pack drives (844s) one at a time. Five drives will be added raising capacity from 700 to 1200 Mc. Provided approval is given to an element in DCR 1975/76 Estimates, it is planned to acquire an additional advanced disc file (819) which would make available an additional 800 Mc - i.e. a potential increase to 2000 Mc for PFs in the system set (if some 844s are not redeployed for user sets).

10. A port expander has been incorporated into the Cyber 76 disc farm. This enables 36 (4 x 8 + 4) 844s to be connected via the controller attached to PP3. Unless the 3600 is decommissioned (or floor space provided in another way) it is not possible to install more than about 32 844s! With eight drives assigned to the system set, four assigned to DR and four assigned to demountable user sets, about 16 drives will be available for permanently mounted user sets (PMUS). Under GEMINI the DR drives would move to the 172s, under DOUBLE the DR drives could augment the PMUS. The DCR on-line file base is large by any standard and will be expanded at a rate greater than the rate at which the processing load develops (unless, as some assert, processing load in an unsaturated machine is proportional to on-line file capacity).

Magnetic tape

11. The unsatisfactory nine track tape subsystem, comprising a hardwired 7628 controller and four 659 drives, is being replaced by a programmable 7622 buffer controller and four self-threading 669 drives. At a later date, say August 1975, it is proposed to

.../4

incorporate two 667 seven track drives. Still later it is likely that the number of nine track drives will be increased to six, and that a second controller will be acquired (giving four tape data channels). By provision of additional system set PF capacity, and of PMUS, DCR hopes to limit the number of tape or disc pack mounts to less than 300 daily. In the medium to long term it is planned to install an automated archive system (software and/or hardware) which will make it unnecessary/impossible for users to waste/request off-line media (expect for purposes of information interchange). Note that the 607 tape drives, currently providing an interim tape system using the 7611 station, will be phased out after the arrival of the 667s. The original GEMINI configuration included the obsolete 3623 controller plus four 607s.

Job processing

12. A year ago nearly 2000 jobs were processed daily on the Cyber 76. Six months ago the figure approached 2500. Towards the end of March more than 3000 jobs were processed on one day, and this number had been approached on a number of other days. In the absence of unscheduled interruptions the Cyber 76 has no difficulty in keeping ahead of short jobs requiring no off-line files. Long jobs are scheduled depending on the work load, while very long jobs are run outside prime shift including weekends. Since it is easy for users to generate groups of jobs requiring off-line files, the backlog of such jobs varies markedly from day to day - it is again noted that less than 300 mounts are presently possible daily (the average job must mount less than 0.1 tapes or packs). The backlog may approach two hours at 1330 EST and four hours (say 100 jobs) at 1730 EST. It is normal to complete all jobs by shut-down time each day.

13. Overheads inevitably incurred in the provision of the computing service make serious inroads on the availability of processing time during scheduled operating hours - currently 0830 to 2130 EST (approximate). It is remarked that most SCOPE systems testing and development is done outside scheduled operating hours and therefore need not be considered as an overhead until scheduled operating hours are widened. Users make little use of the configuration before 0930 EST. Two, say hour, periods beginning 1230 and 1630 EST are used for systems checkout or to catch up on backlog. Thus six prime hours are potentially available to the user (0930-1230, 1330-1630). About 50% of these prime hours are used by the 3000 odd user jobs, i.e. a thousand an hour, while 25% are used for system work (e.g. generation of SCOPE, DAD or NODECODE systems, accounting, performance evaluation, etc.). Given the present growth rate 4500 jobs per day will be reached in 18 months time. Assuming all systems work could be done outside prime shift (which is unreasonable since systems programmers require turnaround like others) 6000 jobs per day would be reached in April 1978 - i.e. prime shift will be saturated with an increasing proportion of work suffering an over-night turnaround - by the 1977/78 financial year.

.../5

General purpose interaction (GPI)

14. Everyone has views on interactive computing! For present discussion GPI is defined to occur when a user job interacts with terminal user(s). Clearly it is pointless for GPI jobs to request access to off-line files - but no other restriction is imposed. It was never proposed to offer GPI on a machine other than the Cyber 76. An experimental system has been developed, and is being used as a vehicle to transfer FOCAL out of the 3600 environment. Selected users have employed the method to provide interactive access to data bases, and to provide graphics systems. A bare minimum has been provided under SCOPE 2 - i.e. no work has yet been done on the scheduler to optimise response time, or to build in facilities which enable multiuser jobs. It is now proposed to undertake a limited enhancement of SCOPE to provide improved GPI facilities.

RIO processor

15. A large number (currently about 110) of relatively slow file input/output devices are attached to the communications network. RIO devices include printers, card readers, plotters and paper tape readers and punches. The number of devices will exceed 150 in the short term since about ten more nodes will be added in the period. The RIO activity is dominated by printing - remote printing now is consistently ahead of central printing for the first time, i.e. 50,000 pages or a million lines daily. The output channel of the Canberra-Melbourne group band line (48kb/s channel) exceeds 30% average loading over prime shift. The Canberra fanout complex bears a load approaching that of Melbourne, while the Canberra-Sydney group band line is lightly loaded. Remote output is averaging about 30 kb/s over prime shift. Coding techniques (blank compression and line truncation) are employed to minimise bandwidth requirements. The January release of NODECODE has reduced document retransmission to acceptable levels, and further improvements will be seen when repositioning is introduced in the next NODECODE release (say June).

16. Performance of the RIO function depends on the existence of a block of "core" memory. Otherwise the discs containing the input/output files are thrashed - i.e. the same sectors are repeatedly accessed for successive records. Currently one bank of 3600 memory (8 * 32768 $\doteq$ 256k characters) is used for the special RIO buffer pool, i.e. about 2k characters per remote device. In the GEMINI configuration about 300k characters could be assigned to the function either in one mainframe or split between the two in a load sharing mode (when software was developed and assuming enough PPs are available). The function could also be performed on special purpose stations employing partial file transmission to avoid the necessity for a full file system - say a minicomputer with access to much core (e.g. PDP11/70 with 256k 16 bit words).

.../6

17. RIO could be provided using a block of 300K characters of Cyber 76 LCM (memory). This approach avoids wasting resources by respooling files to station discs - overheads which it is alleged could swamp the two disc channels of the GEMINI configuration. Certain Cyber 73 configurations are said to suffer this defect, and these machines have only a fraction of the throughput, or of the peripheral devices of the CSIRO configuration. It is my opinion that the Cyber 73 configurations (98K but doing user jobs) showing this symptom have insufficient buffer space, and more core dedicated to the RIO function should be considered. It is noted that the pro-rated cost of 30K LCM words is less than that of the disc controllers and disc drive capacity used for spooling on the GEMINI configuration. Given the necessary software the question reduces to consideration of reliability - i.e. in the environment of DOUBLE a stable performance could be achieved.

TED processor

18. TED is implemented on the 3600 as a resident multiuser job which operates on dynamically relocatable workspaces which are swapped to drums. Four workspaces are currently used - five will be available when FOCAL is transferred to the Cyber 76. Workspaces are held in core to enable: (1) edit operations on data in the workspace, or (2) input/output of records to/from the workspace. The latter activity holds the workspace in core for "long" periods, and it is proposed to further improve performance by permitting input/output with a swapped out area (i.e. increase the utilisation of the available workspaces).

19. Cyber 73 98K configurations have difficulty in supporting more than 30 INTERCOM editors at a two second response time for edit commands. The 3600 does much better than this! Drums are a key - Cyber 172 PPs drive drums (which in any case are obsolete) at half speed, and for similar reasons drive 844s at half speed. Thus Cyber 172s swap workspaces at 3Mb/s against 12Mb/s for the 3600. One answer to this problem is to have more workspaces in core. It is noted that the Cyber 76 drives 844s at 6Mb/s. Further, the 819 advanced disc file, only available on a Cyber 76, has a transfer rate of 21Mb/s. It is the only device which can exceed the performance of the obsolete drums.

20. Performance of the TED function again depends on the existence of a block of "core" memory - about 20K characters per 16 users on 3600 experience, i.e. if it is desired to support 128 users at least 160K characters storage are required, plus adequate processing power and swapping bandwidth. Increasing the size of workspaces increases the bandwidth required. Clearly a block of LCM, equal to that proposed for RIO, would support 128 users with double sized workspaces given an 819 swapping medium, and use less than 5% of the Cyber 76 processing power. Again subject to the availability of software the question reduces to considerations of reliability.

.../7

In the environment for DOUBLE the approach is sound. In the short-term there is no alternative since the 3600 is not capable of extension to 128 users - the processor power is just not there!

Queue interrogation

21. Various interactive programs enable the interrogation of system queues, and the like. These queues are stored on the drums as long linked lists, e.g. the system may handle 20,000 documents in a 24 hour period. It is noted that the SCOPE 3 operating system holds this type of information in core, and that less than 2000 system documents can be accounted for. In order to provide facilities on the Cyber 172s which match those presently available under DAD an extensive software effort will be required. The SCOPE 3 operating system is very limited since configurations are supported which include no ECS (the analog of LCM). In essence the system comprises disc overlays of PP code with minimum resident central code. SCOPE 2 is based on the concept of overlays from LCM to central - a more open arrangement. Frequently used disc resident overlays lead to inefficiency and hangups. It is clear that it will be more easy to modify SCOPE 2 than SCOPE 3 - in fact it is probably unwise for DCR to attempt to maintain two major operating systems since each requires a minimum of four people.

LCM

22. The performance of the Cyber 76 will be greatly increased by the acquisition of an additional bank (256K words). Frequently used SCOPE 2 overlays currently residing on disc will be moved to LCM - thus removing bottlenecks and avoiding hangups. Multi-programming efficiency will be increased leading to increased throughput. The maximum number of jobs ever done in an hour is about 750 - in that hour 15 minutes were unused due to conflicts. The LCM will enable full utilisation of the processor for the prime time job mix. In initially released systems following the upgrading it is proposed that the maximum permitted user LCM field length to be at least doubled to around 256K words (1000_8 blocks). The remainder will be assigned to TED, RIO and SCOPE. While the new TED and RIO processors may be developed as user jobs in advance of the delivery of LCM, it is not proposed to release these systems until the additional memory is installed - i.e. the user job field length will not be reduced in the interim period. Further it is not proposed to switch over to the Cyber 76 based systems until the 3600 based systems are outperformed.

Conclusion

23. The DAD system based on the 3000 configuration is now highly specialised and is performing very efficiently. Nevertheless the 3600 is obsolete! The cost of maintenance services approaches that of the Cyber 76. Clearly there is a real possibility that critical elements of the 3600 may become unmaintainable even in the

.../8

short term. The Cyber 170s provide no solution to the problem in the short term: Even if a Cyber 172 were commissioned by fourth quarter 1975 this would not displace the 3600 until mid 1976 at the earliest time, when suitable high performance software could be available.

24. Work to implement the key elements of the DAD system in the Cyber 76 has commenced. It can be confidently predicted that the systems can be developed to outperform the 3600 well before December 1975. In this way we become independent of the 3600. Backup of all services is a design objective - funds must, of course, be generated to provide the necessary equipment redundancy. The trade-offs between performance and reliability are sensitive issues since human behaviour is called into question - e.g. interactive vs. batch style of working. Unless something can be made of the Cyber 170s, the principal problem seen in the medium term is the MTBF of the Cyber 76. Perhaps we must live with this until duplication occurs.

Peter J. Claringbold
750407

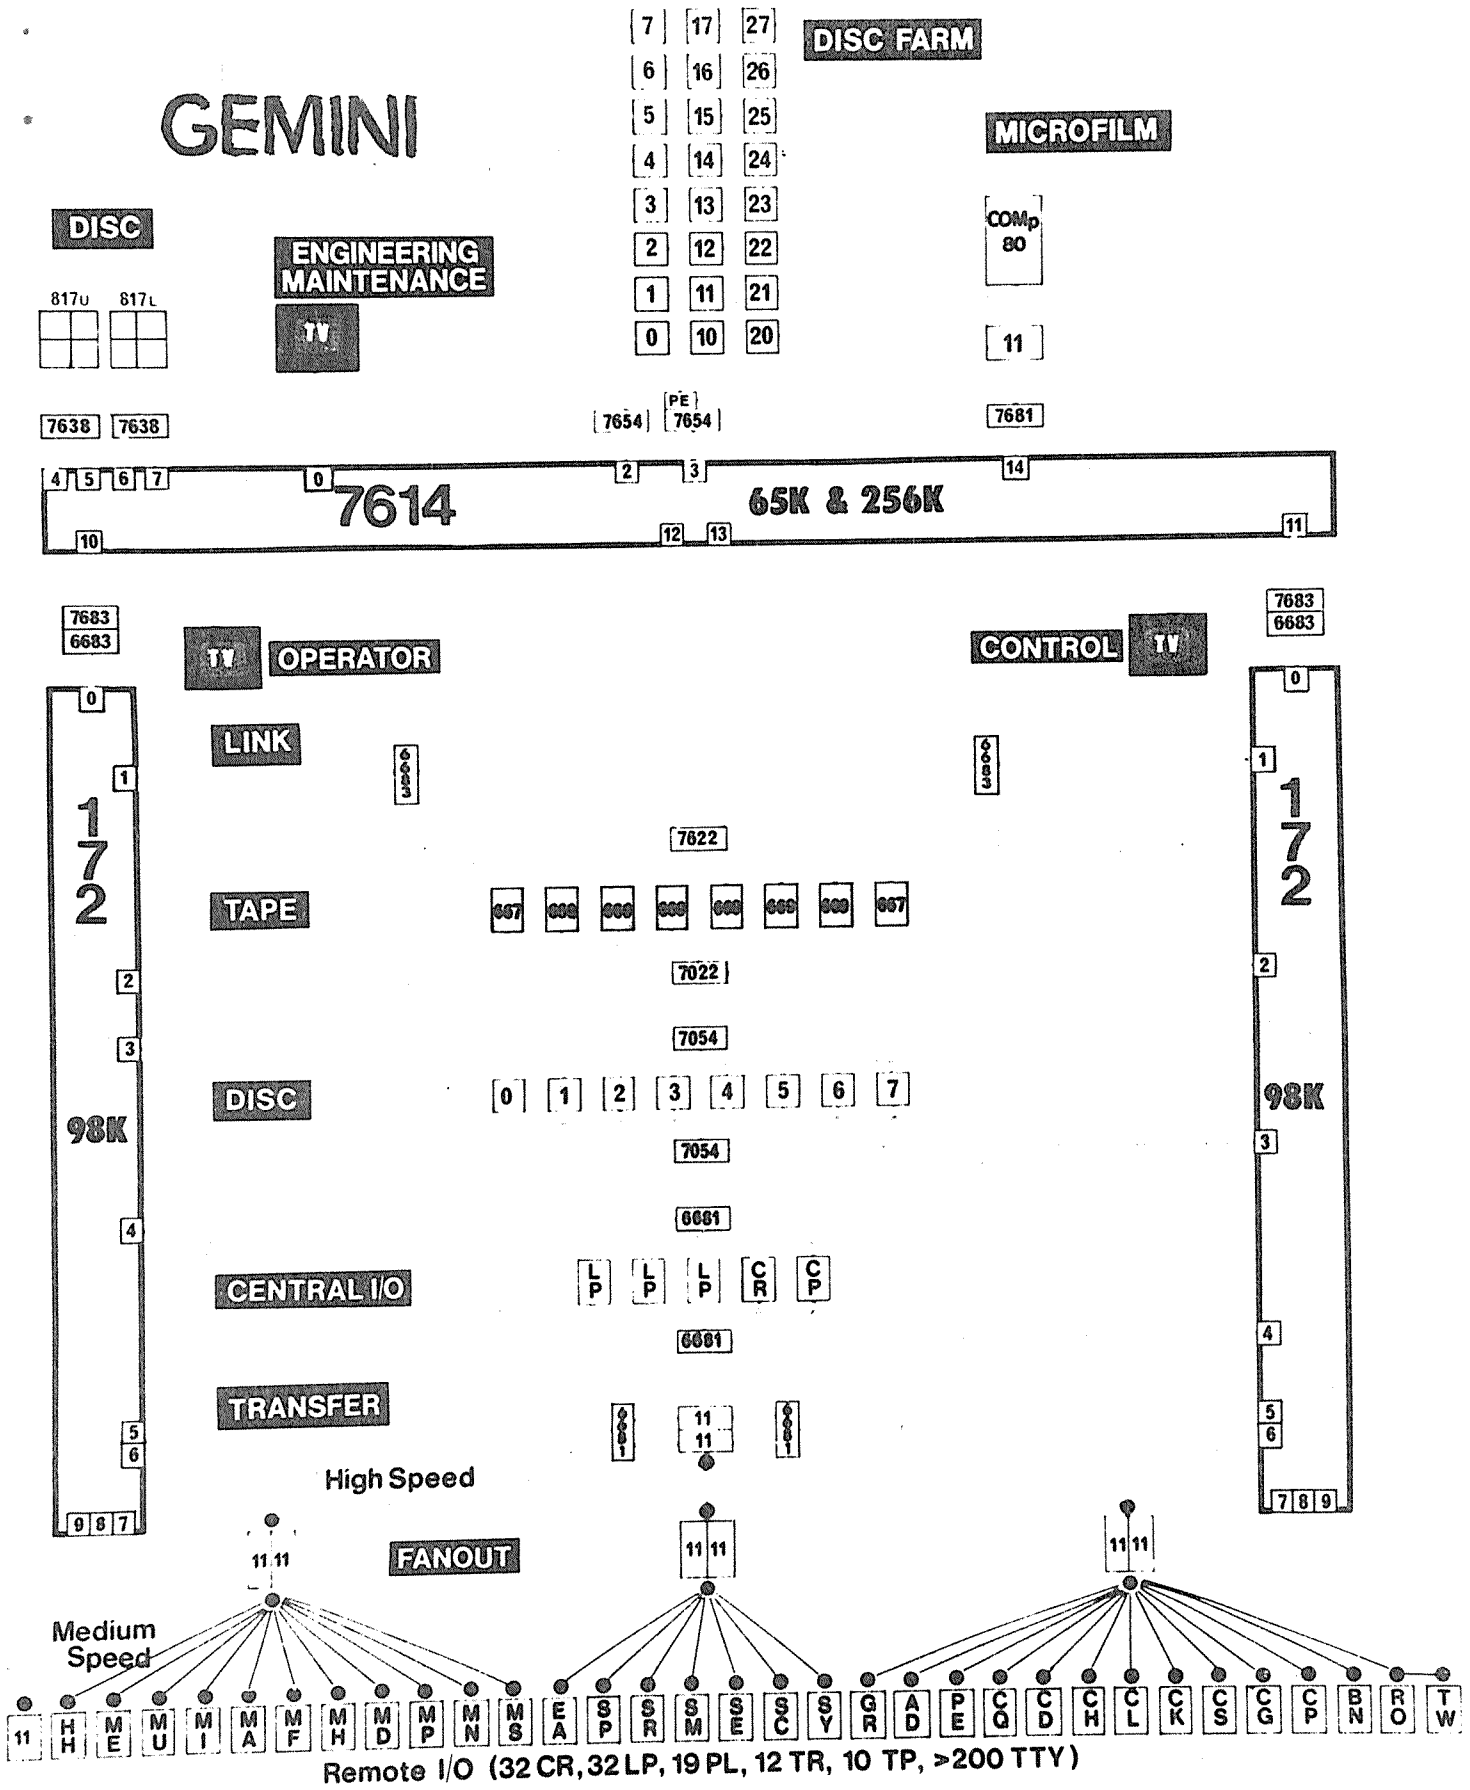
GEMINI
DISC
817U
817L
ENGINEERING MAINTENANCE
TV
7638
7638
DISC FARM
7 17 27
6 16 26
5 15 25
4 14 24
3 13 23
2 12 22
1 11 21
0 10 20
7654
PE 7654
MICROFILM
COMp 80
11
7681
4 5 6 7 0 2 3 14
7614
65K & 256K
10 12 13 11
7683 6683
TV
OPERATOR
CONTROL
TV
7683 6683
0 1 2 3 4 5 6 7 8 9
172
98K
LINK
6683
7622
TAPE
667 666 666 666 666 666 666 667
7022
7054
DISC
0 1 2 3 4 5 6 7
7054
6681
CENTRAL I/O
LP LP LP CR CP
6681
TRANSFER
6681
11 11
6681
High Speed
FANOUT
11 11
11 11
11 11
Medium Speed
11 HH ME MU MI MA MF MH MD MP MN MS
EA SP SR SM SE SC SY
GR AD PE CQ CD CH CL CK CS CG CP BN RO TW
Remote I/O (32 CR, 32 LP, 19 PL, 12 TR, 10 TP, >200 TTY)

FIGURE 1.

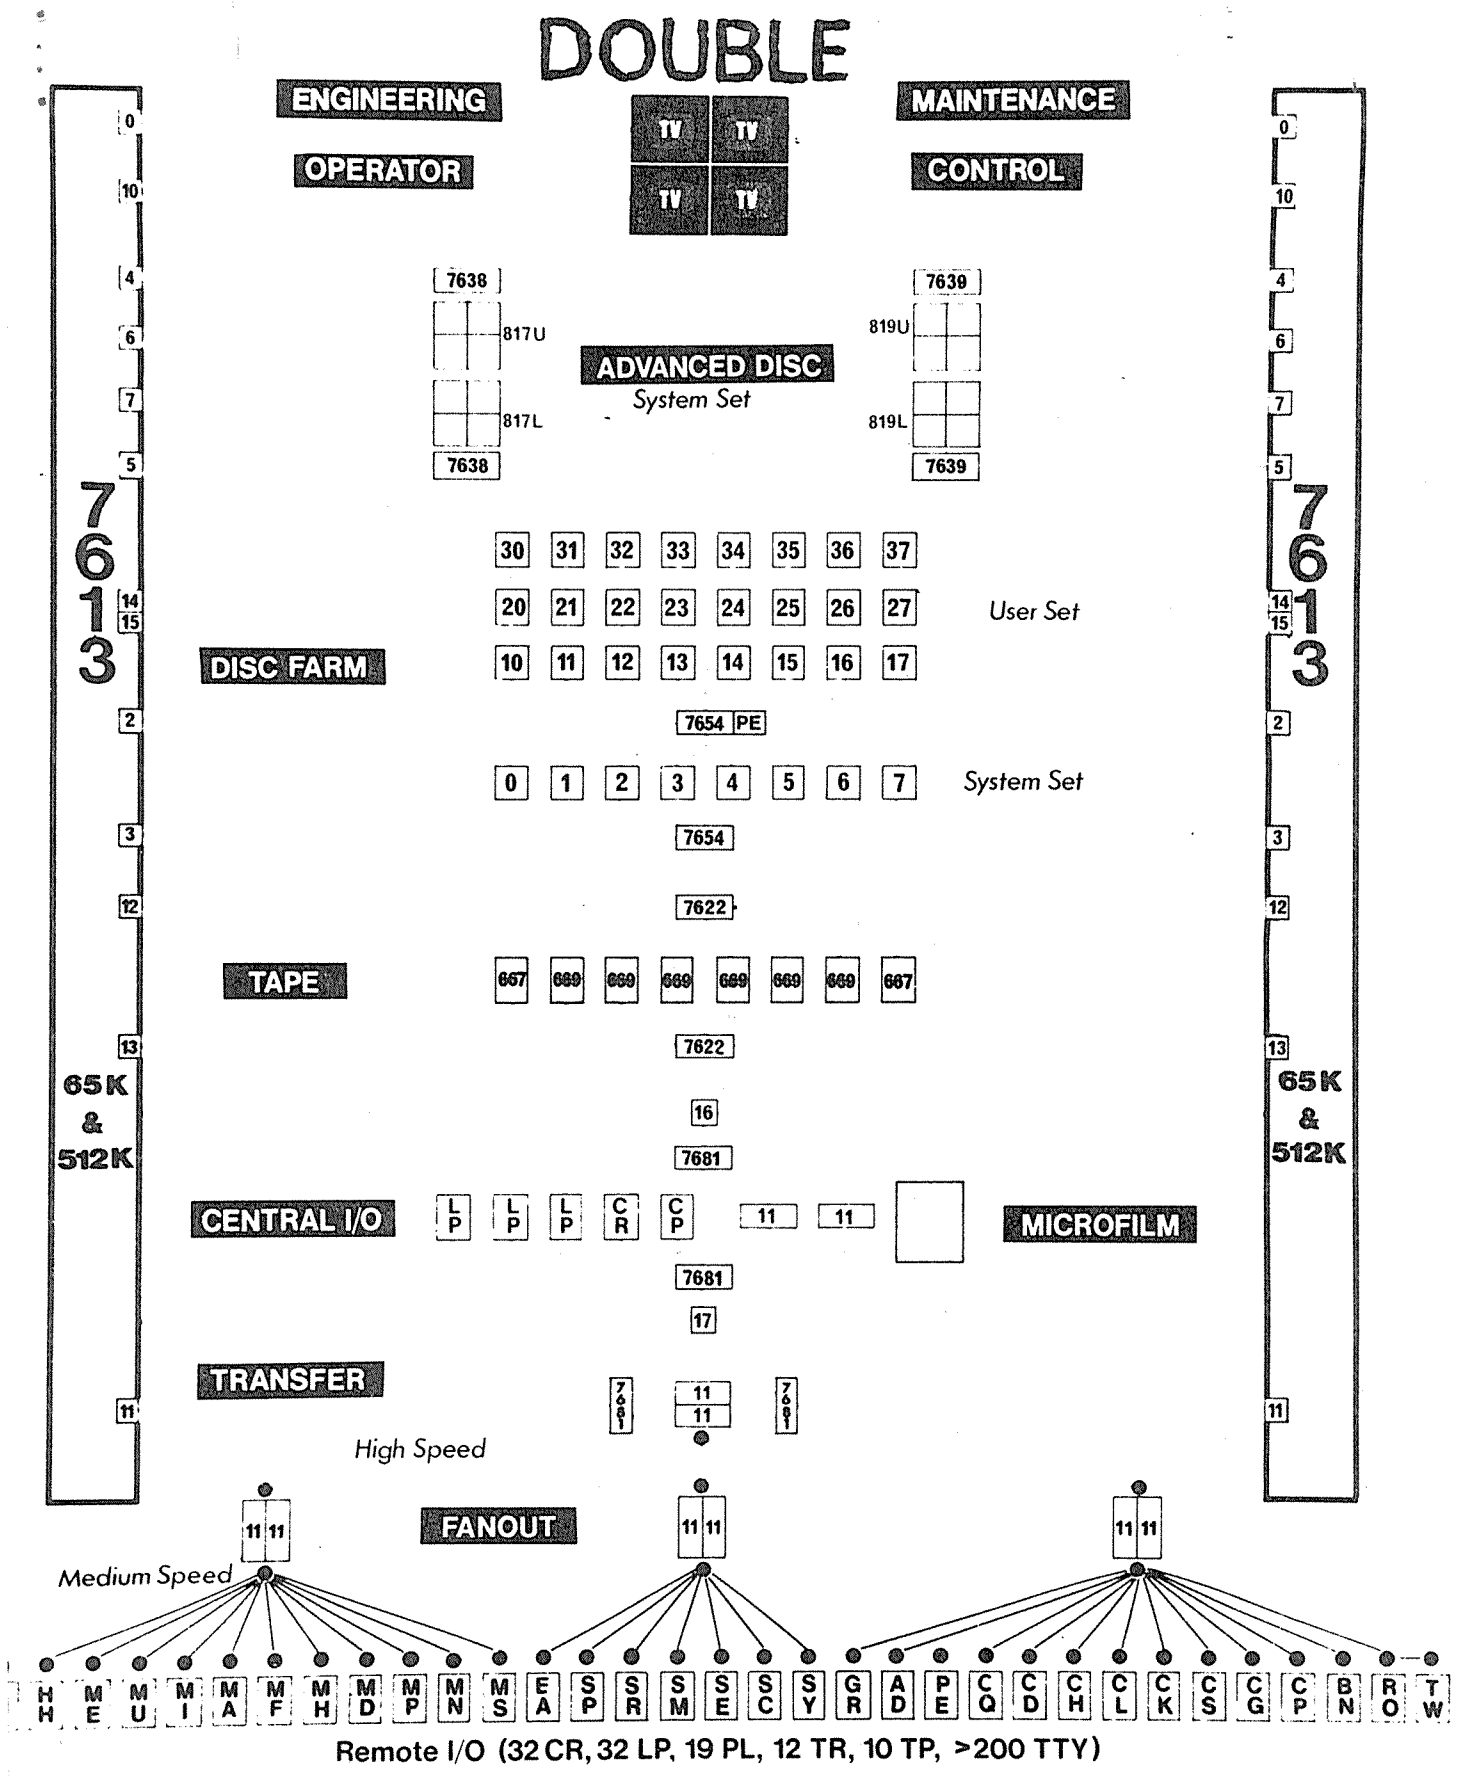
DOUBLE
ENGINEERING
OPERATOR
TV
TV
TV
TV
MAINTENANCE
CONTROL
7638
817U
817L
7638
ADVANCED DISC
System Set
7639
819U
819L
7639
30 31 32 33 34 35 36 37
20 21 22 23 24 25 26 27
User Set
10 11 12 13 14 15 16 17
DISC FARM
7654 PE
0 1 2 3 4 5 6 7
System Set
7654
7622
TAPE
667 669 669 669 669 669 669 667
7622
16
7681
CENTRAL I/O
LP LP LP CR CP 11 11
MICROFILM
7681
17
TRANSFER
7681
11
11
7681
High Speed
FANOUT
11 11
11 11
11 11
Medium Speed
HH ME MU MI MA MF MH MD MP MN MS EA SP SR SM SE SC SY GR AD PE CQ CD CH CL CK CS CG CP BN RO TW
Remote I/O (32 CR, 32 LP, 19 PL, 12 TR, 10 TP, >200 TTY)
7613
65K & 512K
0
10
4
6
7
5
14
15
2
3
12
13
11

FIGURE 2.