

DESIGN OF THE G-21 MULTI-PROCESSOR SYSTEM

BY

JESSE T. QUATSE

Carnegie Institute of Technology
February 26, 1965

This work was supported by the Advanced Research Projects
Agency of the Office of the Secretary of Defense (SD-146).

Summary

During the summer of 1963, a multi-processor version of the CDC G-20 (the G-21) was developed for the Computation Center of Carnegie Institute of Technology. The purpose was to double the memory capacity and processing speed of the existing G-20 system. Since then, the system has continued to evolve as the needs of the Computation Center continue to grow. The purpose of this paper is to describe the memory control electronics which underlie memory time-sharing. The important related engineering problems are discussed in the terms of the current G-21 system.

ACCESSING SYSTEMS

A collection of computer systems is generally referred to as a multi-processor system if main memory is time-shared by the processors. The G-21 multi-processor system includes three classes of memory users, only one of which is the processor pair. Therefore, the emphasis of this paper is memory sharing rather than multi-processing.

Typical main memories are housed in modules external to the processors. Because of the physical size and high speed of the modules, the intrinsic delay time of the inter-connecting cable is significantly greater than the shortest pulse duration. Line termination is necessary to avoid reception of distinct reflected pulses. The low impedances of terminated lines require expensive driver and receiver circuits. The number required depends upon the type of accessing system as well as the number of devices in the system. Thus, an implicit cost is associated with the choice of accessing systems.

In general, memory sharing control electronics are based upon either a "bus" or a "portal" accessing system [1], [2], [3], [4]. The portal system, often called the "cross-bar" system, is analogous to a set of stepping switches, one per module. As shown in Figure 1(a), each module user is connected to one terminal of the module's stepping switch. When one user is accessing the module through its terminal, access requests from the remaining

FIGURE 1(a)

PORTAL ACCESSING SYSTEM

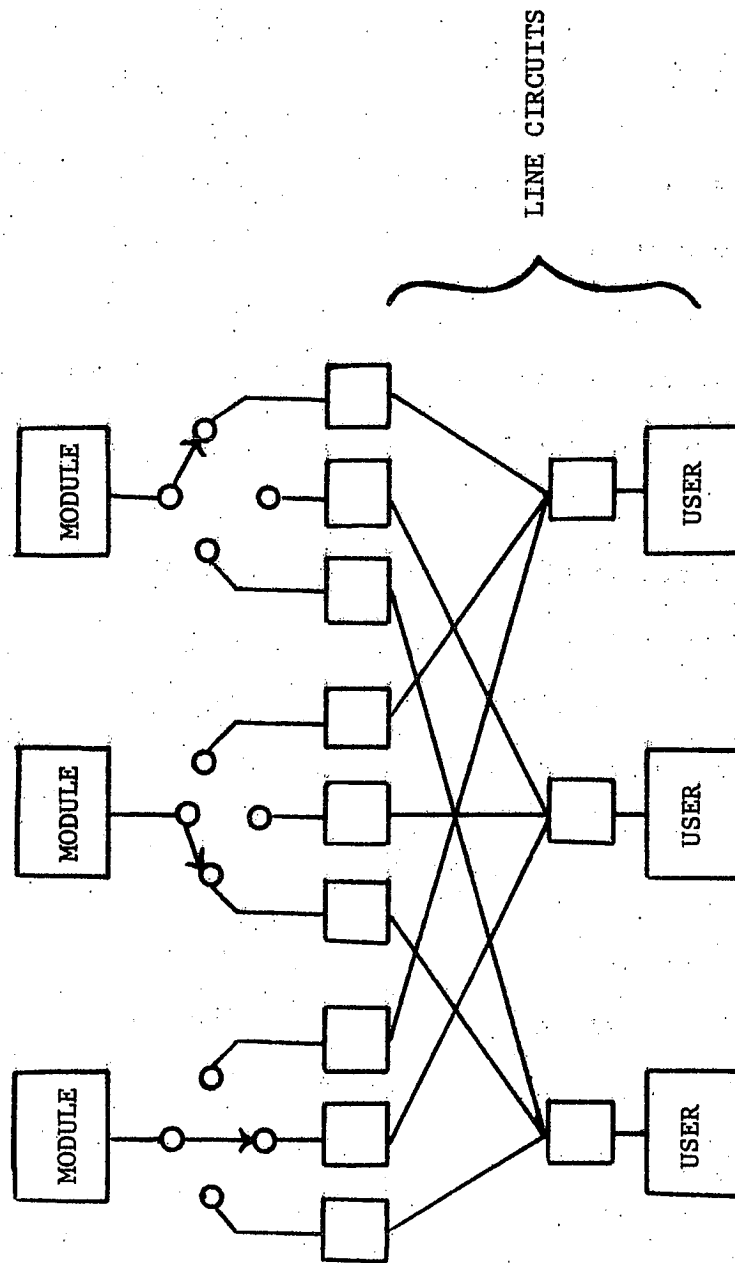
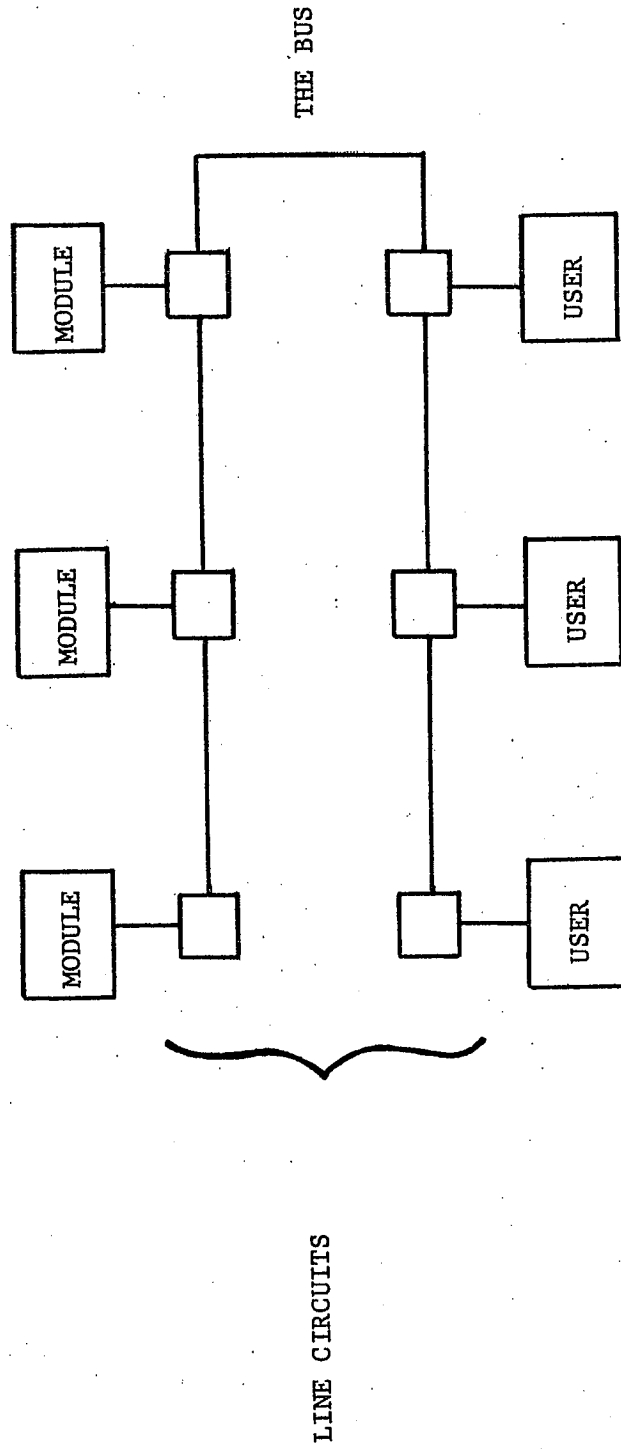


FIGURE 1(b)

BUS ACCESSING SYSTEM



terminals cannot be serviced by the module. Users of the remaining terminals must inactively await promotion by the stepping switch.

The bus system resembles a portal system in which the entire memory bank is serviced by a single stepping switch. The distinction between bus and portal systems entirely disappears where the memory bank consists of a single memory module. Where the bank is an aggregate of modules, the bus system permits only partially overlapped accesses of independent modules. All modules and all users time-share a common set of communication and switching circuits called the "bus" in Figure 1(b). A queue forms when requests for the bus, as well as the module, are simultaneously issued by more than one user. As with the portal system, the bus system depends upon a priority scheme to deplete request queues.

The intrinsic efficiency of a portal system is greater than that of a bus system. Portal systems are characterized by the centralization of control in the modules. Unless users request the same module, users are independent in a portal system. Users of a bus system cannot operate independently. In effect, an access is an uninterrupted conversation between a user and a module. Although the conversation duration may be shorter than the module cycle time, subsequent accesses are delayed by at least that much, and efficiency is degraded. The trade off is efficiency for the cost of driving and receiving circuits. The

number of line circuits for a completely cross-coupled portal system goes as the product of the number of users and the number of modules. For a bus system, the number of line circuits goes as the sum.

THE G-21 ACCESSING SYSTEM

The particular variation of accessing systems devised for the G-21 is a combination of bus and portal.* As shown in Figure 2, eight memory modules (MM) are coupled to two central processors (CP) by the processor bus. (The bulk memory is still in the planning stage. It appears only as an illustration of the bus system's generality.) Emanating from each MM is a three-terminal portal system. One terminal couples a display controller (DC) which acts as a bus system for a maximum of four CRT display consoles. The other terminal couples a general exchange controller (GXC) which acts as a bus system for a maximum of four telephone channels (TC). Each TC interfaces the general exchange through a 202A data-phone subset.

The three portals of a module are not identical. Each is specialized to the needs of the particular category of facility it services. Names are given, in Figure 2, to suggest the nature of each category. The processor category is the most expensive to delay and receives the highest grade of service: "high capacity". The display system category,

* A brief specification of the G-21 system is given in the Appendix.




FIGURE 2

G-2I MEMORY CONFIGURATION

"continuous demand", receives the lowest grade of service. The module is used as a display regeneration memory by the DC. Module access requests are continuously generated by the display system [5]. Failure to obtain accesses for extended periods of time causes no loss of information. At worst, flicker becomes visible. In contrast, information can be lost by sufficiently poor service of the "low demand" terminal. The general exchange facility is unbuffered beyond the single character input register. In order to avoid the necessity of further buffering, the facility must be guaranteed service within a fixed time interval following an access request. The guarantee is possible only if neither of the other terminals can preempt the module indefinitely. (Because of the processor speed, the processor bus is capable of continuously requesting memory accesses at full memory speed.)

Permanent module preempting (by the processor bus) is prevented by allowing one of the other two terminals to gain access to the module between contiguous high capacity terminal requests. The high capacity terminal is demoted for a single access every time it wins a tie. The other terminals have fixed priority levels. Access requests by the high capacity terminal are always granted in preference to requests by the continuous demand terminal. The high capacity terminal is

of higher priority than both when promoted and of lower priority than either when demoted. Thus, a continuing condition of three-way ties results in alternate accesses by the high capacity terminal and the low demand terminal. Were persistent ties possible, accesses by the continuous demand terminal would be entirely excluded. All display operations would cease. The situation is avoided by limiting the channel capacity of the low demand terminal. Four TC channels degrade the flicker rate by less than 2 percent.

Permanent preempting must be prohibited in the bus system also. Fixed processor priorities would permit one processor to exclude the other from accessing the same module. Consequently, the priority alternation of the bus system is similar to that of the portal system. A manual switch designates one processor as the "master" and the other as the "slave". Priority initially resides in the master. Each tie lost by the slave causes the slave to promote itself and demote the master. Thus, the master always wins the first of a series of contiguous ties, but wins every second tie thereafter.

The G-21 bus and portal priority schemes can be characterized as simple queue control mechanisms which obey the following rules:

- (1) The queue consists only of access requests.

Addresses and data are not issued by the user until the request is granted. (This leads to more elaborate synchronizing circuitry, but fewer register flip-flops, than would otherwise be needed.)

- (2) Requests are queued, instead of serviced, in two situations: whenever the module is busy or whenever a request is issued simultaneously by a higher priority user.
- (3) A request is serviced when received if no higher priority request is queued, the module is not busy, and no higher priority request is received simultaneously. The order in which queued requests are promoted and serviced depends upon the facility category, rather than the order of request reception.

Other priority schemes in current use range from elaborate queueing circuitry to fixed cyclical service. The assumption here is that elaborate queue control could not pay its way in the G-21. An access request is never denied unless another is simultaneously granted. Thus, the average access rate, and therefore the efficiency, cannot be increased by any other scheme. The distribution of accesses can be effected by an elaboration upon the queue control, but to no worthwhile advantage. For example, the high capacity terminal could be

controlled so as to win four out of five ties instead of one out of two.

COMMUNICATION AND SYNCHRONIZATION OF REQUESTS

The communication and synchronization tasks inherent to programmed time-sharing are performed at the hardware level by the queue control electronics. Since programmed queue control resides in the users, and not in the shared facility, programmed time-sharing is analogous to a bus rather than a portal system. In a bus system, whatever the control level, independent user requests must be communicated to the users as well as the time-shared device. Synchronization is necessary to resolve ties.

The time-shared use of G-21 disc memory is typical. Program records stored on the disc must be available to both processors. Each processor effects record changes by reading from disc, computing, and replacing on disc. A period exists in which the record in one processor core region differs from the record on disc. A disc reading by the other processor, at this time, will not be accurate. The other processor must wait until the modified record is returned to disc.

A cell in common memory acts as a busy switch. Notices posted in the cell, by either processor, communicate the "record busy" status to the other processor. All disc-read

operations are preceded by this switch interrogation procedure in which the requesting processor both reads and notifies the state of the busy switch. A timing conflict arises whenever both processors simultaneously require program records. Between the time that one processor reads the switch and sets it, the memory module is free for at least one access. Unless a synchronizing procedure is followed, the intervening access might be used by the other processor to read the switch. In that case, both processors will interpret the record status as "not busy", and both will attempt record reading. The current G-21 executive program resolves the ambiguity and synchronizes the requests by arbitrarily establishing one processor as the "master" and the other as the "slave". Each processor requests use of the disc from the other processor by posting a notice in a common memory cell and generating an interrupt. The slave processor denies the request only if it is currently using the program records. The master processor also denies the request when it is about to use the program records. As with the queue control circuitry, the synchronizing program has established one of the two processors as the priority user in order to resolve the tie.

One major difference between the circuitual and programmed versions of queue control affects the distribution of facility accesses. As previously described, the bus user priority

alternates between processors. Permanent module preemption is impossible. In contrast, the fixed priority established by the executive program theoretically permits the master to exclude the slave from ever reading the disc. Precautions are unnecessary with disc switching because of the large fraction of running time in which the disc is available.

THE BUS QUEUE CONTROL ELECTRONICS

The queue control electronics are physically located in both of the processors. Module requests are interpreted by the queue control electronics as requests for the processor bus. If the bus is unavailable, the request is queued until the bus is available. Otherwise, the requested module is accessed.

The modules are accessed as asynchronous 8,192 word core banks with self-contained read/restore electronics. When given a "start" signal, a 13-bit address, and a read/write control signal, a module will commence the indicated read or write access-cycle at the indicated address (See Figure 3). If a write operation was requested, a 32-bit data word follows the start signal within 2 μ s. If not, a 32-bit data word is read by the module and issued between 2 and 4 μ s. after reception of the start signal. In either case, a "data available" signal is generated by the module 3 1/2 μ s. after




FIGURE 3

MODULE ACCESS TIMING

reception of the start signal. The read portion of the read/restore cycle requires 3 μ s. But the internal start signal is delayed 1 μ s. by the module to allow address transmission time. The 3 μ s. restore cycle is internal to the module and does not require processor bus information. The final synchronizing signal issued by the module is a "finish cycle" signal. It precedes the cycle completion by 1 1/2 μ s.

The address, read or write control, data write, and data read lines are common to all modules and all users. They constitute 80 of the bus lines. The start, data available, and finish cycle lines are common to both users but private to the modules. Thus, one trio is provided for each module. Together with the 80 common lines, the eight trios bring the total number of module control lines to 104.

The module control signals are issued or received by the set of queue control electronics located in each processor. The set consists of a control sequencer, a bank of nine status flip-flops, and the bus driver and receiver circuits. Eight of the nine status flip-flops record the individual module states by monitoring the eight start lines and the eight finish cycle lines. The simplicity of the queue control eliminates the necessity to record which processors issued the module start signal. The remaining status flip-flop

determines the processor priority. It is accompanied by one promotion line and one demotion line. Along with two "bus request" lines and one "bus busy" line, they bring the total number of bus lines to approximately 108. (A few more lines are necessary for unimportant circuitual reasons.) One request line is driven by the master and interrogated by the slave. The other is driven by the slave and interrogated by the master. The bus busy line is driven by either and interrogated by both.

A module access-request by a processor is interpreted by the associated queue-control electronics as a request for use of the processor bus. Depending upon the status of the processors, modules, and bus, one of five timing sequences occur. The five sequences appear in Figure 4, Figure 5, and Figure 6 as five distinct cases. Figure 4 shows the basic timing sequence resulting from a request issued by one processor when the other processor has not also issued a request and both the module and bus are free to grant the request. By comparison with Figure 3, it is evident that the delay in converting the start signal to an internal start signal is absorbed by the early finish cycle signal during contiguous accesses of the same module. The processor "sees" a string of contiguous accesses as requiring 7 μ s. for the first access cycle time and 6 μ s. for




FIGURE 4

BASIC BUS ACCESS TIMING

all subsequent cycle times.

The basic cycle is modified, in the second situation, by the presence of another request. Assuming again that the requested module is free for access, the other processor may issue a request at one of three times. A request issued at time 2 in Figure 4 produces the timing sequence shown in Figure 5(a). The bus is not busy, and the other request has disappeared. The situation is identical to that of Figure 4, and bus access is granted. The access of the second processor follows the access of the first processor by 2 μ s. The greatest variance of any signals shown in Figure 3 is 2 μ s. for the read lines. Therefore, a bus unavailability period of 2 μ s. is certain to delay accesses long enough to stagger successive outputs to the read lines. In fact, any request time in the "safe" 4 μ s. interval, shown in Figure 4 as $t = 2$ to $t = 4$, is equivalent to the situation represented by 5(a).

The third situation corresponds to the request issued to $t = 1$ in Figure 5(b). The queue-control electronics of the second processor interprets the presence of the bus busy signal as an "idle" state. All states remain unchanged until the bus busy signal disappears. Thus, the situation degenerates to that of Figure 5(a).

The fourth situation, shown in Figure 5(c), occurs when

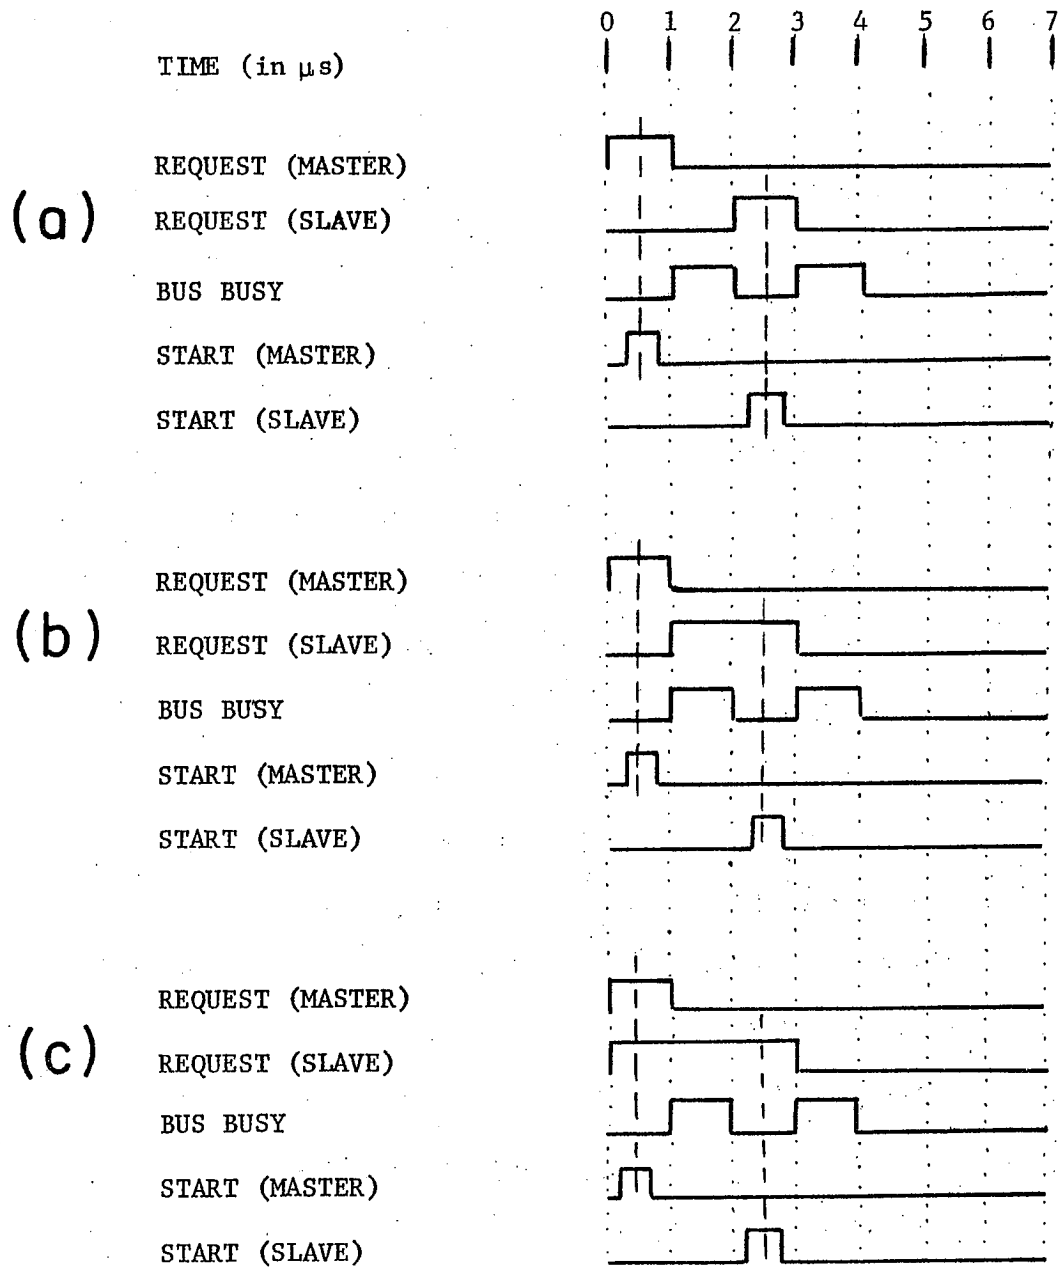


FIGURE 5

TWO PROCESSOR BUS ACCESSING

both requests are issued simultaneously. The low priority processor, in this case the slave, interprets the presence of the high priority request as equivalent to the presence of the bus busy signal. The situation degenerates to that of 5(b) which in turn degenerates to that of 5(a).

The promotion and demotion circuits function when the addressed module is busy. When the addressed module's status flip-flop is in the busy state, the queue control enters an idle state in which bus and module requests are not issued. While in the idle state, the slave processor promotes itself and issues a demotion signal to the master processor. The assumption is that the master processor is competing with the slave processor for use of the same module. When the module finishes its current cycle, the slave, and not the master, will be the priority user. The module status flip-flop is interrogated at 1 μ s. intervals until the not-busy state is detected. One of the previously defined situations then exists. As the status flip-flop changes state, the queue control leaves the idle state and initiates the appropriate access sequence. After obtaining use of the bus, the slave processor returns priority to the master by demoting itself and issuing a signal on the promotion line.

The final situation, as shown in Figure 6, illustrates alternating priority. Both processors are requesting the




FIGURE 6

PRIORITY ALTERATION

use of the same module at the maximum module rate. Initially, the slave processor, lacking priority, is unable to obtain use of the bus. It retracts its request, at $t = 2$, and enters the idle state described above. The situation is then indistinguishable from the one just described. When the slave obtains use of the bus, at $t = 6$, the master is returned to priority status. Thus, accesses proceed at full memory rate, but alternate between the two processors.

THE PORTAL QUEUE CONTROL ELECTRONICS

The portal queue control electronics consists mostly of driver and receiver circuits. Its simplicity is attributable to its lack of efficiency. The slow access and synchronous access terminals were designed for specific facility categories, neither of which imposed critical timing restrictions upon the portal. Therefore, the basic cycle time for either portal was extended to 7 μ s. The additional 1 μ s. is used by the queue control electronics to detect and act upon access requests by the rapid access portal. The inefficiency alleviates the need to communicate portal request information to the bus queue control circuitry. Communications with the high capacity terminal are conducted entirely by the standard module control lines.

The queue itself consists of two flip-flops, one for each

of the inefficient terminals. An access request by the high capacity terminal is always serviced, regardless of the queue condition. The information that the high capacity terminal is engaged in an access is recorded by the priority flip-flop. A non-empty queue, during this high capacity demotion interval, inhibits the transmission of the finish cycle signal. The portal queue control detects the cycle completion and services all current requests. Only queued requests, or newly received requests from the inefficient terminals, are possible at this time. The processors are unable to issue further requests, to the module, until the finish cycle signal is received.

If no requests are available for service, the portal queue control issues the missing finish cycle signal. If a request is granted, the signal will appear automatically when the granted access cycle attains completion. It should be evident that only one "sneak" is possible. Requests from the low demand terminal always mask requests from the synchronous access terminal. Thus, no ties can occur, and the low demand terminal is able to permanently exclude the synchronous access terminal.

SYSTEM IMPROVEMENTS

The purpose of this paper has been to describe the G-21

as it was developed. Use of the system has suggested areas for improvement. The four improvements described below vary in cost and importance, but none are being considered as future G-21 developmental projects. They would be practical to consider only if much of the equipment involved were not already designed and operating. At the current status of G-21 development, the cost of each would greatly exceed its value to the CIT Computation Center.

First, the efficiency of time-shared memories can always be increased by decreasing the number of words per module. Cost increases can be minimized by resorting to a bus system and by housing more than one module per cabinet [2]. The G-21 houses 65k of memory in eight 8k modules. An immediate improvement can be had by treating each module as the two independent 4k stacks that it is.

Second, the programmed time-sharing of equipment can be greatly expedited by a few hardware communication aids. The G-21 is equipped with a processor-to-processor interrupt. Without it, the previously cited disc switching example would be unmanageable. The synchronization requirements can be satisfied by a simple modification to the memory modules. A new mode of accessing can be installed which does not restore after reading. If it were, a new "clear add" opcode, using the non-restore mode, would then leave the contents of the

accessed location cleared to zero. One cell in common memory would correspond to each facility. A zero value, for the cell, would indicate the facility busy state. Any other value would indicate the not-busy state. In order to request use of a facility, a processor would read the corresponding cell with the new opcode. It would then test the accumulator for zero. If the accumulator were zero, the facility would be recognizable as being used by the other processor. If not, the facility would be both free for use and declared busy to the other processor. The resulting priority scheme would be "first come first served". Since a single access is required to read the busy switch and to change its state, a tie could not occur. One processor must obtain the single access before the other. The switch would be set to not-busy by a standard store command. Again, no ties could occur.

More elaborate mechanisms can be devised which conserve memory locations. One example is a pair of new commands "unite and jump" and "selectively reset". The effect of the first command is to unite the accumulator with the specified cell, in one operation, and jump out of sequence if the value changes. The second command simply clears every bit position in the specified cell that corresponds to a zero positional value in the accumulator. The use of the command pair is analogous to the new clear add. The difference is that each

bit of the common cell can correspond to a different facility.

Third, a system efficiency gain can be had by more thorough exploitation of the module portals. Modules are ideal rate matching devices. They communicate at any rate up to the system maximum. If necessary, they can be independently accessed without degradation of system performance. They are the axis of the system in the sense that they are common to all operations. An important use to which G-21 portals could be put is the coupling of low speed telephone channels. Sixteen TTY channels enter the system through expensive delay line circuitry. Buffering beyond one character is unnecessary to portal coupling.

Finally, hardware relocation and memory protection are of obvious utility to a multi-processor system. The G-21 is equipped with a mild form of memory protection and no form of hardware relocation. Briefly stated, the protection circuits transform an illegal access into an interrupt. An access is illegal if it consists of any write operation into a protected region. Regions are assigned on an 8k basis so that an entire module is either fully accessible or fully protected. Any program may declare any module as unprotected by loading a commonly accessible flip-flop register. Thus, undebugged programs can write into any region by a suitable

sequence of errors. Since no read protection whatever is available, "malicious" users are free to inspect any information in core.

An adequate addressing scheme for the G-21 should incorporate both protection and relocation. The efficiency of time-shared use is greatly increased by the hardware provision for locating any program segment at any absolute address without object code modification. Segments may then be loaded into any sufficiently large block of available memory and more than one user may be simultaneously resident in memory. Systems which provide paging as well as segmenting hardware obviate the necessity to compact available memory into large contiguous blocks [3], [6], [7], [8]. Memory is partitioned, by the hardware, into small "pages" so that all segment replacements proceed a single page at a time. Only the page called need be loaded, and it may overlay any page not in use. The assignment of protection status, to each page, is restricted by the hardware to the debugged and benign executive routine.

A further efficiency is gained by the provision for independent relocation of data and procedure. System programs such as the executive, the languages, and the library, may be written as common procedures and thereby need be overlaid only when memory space is needed. Segment swaps,

and the associated costs, are then reduced in magnitude by the number of memory cells occupied by the required common procedures.

In essence, the purpose of elaborate relocation and protection hardware is to delay until the latest possible time (i.e., call time) the decision of which information must be resident in memory, of what size it must be, and of where it must be loaded. The net effect is to reduce the costs of unexpected or unpredictable events which precipitate large information swaps. The effect is achieved in two ways: by reducing the quantity of information involved in each swap and by efficiently allowing for common procedures and for more than one user's program to simultaneously reside in main memory (thereby reducing the expected number of memory swaps). Without such hardware, multi-processor time-shared systems such as the G-21 are susceptible to high memory swap costs whenever, shortly after loading, large program segments either abort, call upon input/output, call upon other large segments to be overlayed, or attempt to interact with the user in a conversational mode.

CONCLUSIONS

A multi-processor system can be classified as either a "bus" system, a "portal" system, or a combination of the two.

Each independently accessible memory module of a portal system is equipped with a set of independently accessible line and queue circuits. Accesses of one module in no way interfere with accesses of another. A bus system makes less efficient use of memory because the bus line and queue circuits are time-shared by the memory users. The economy of circuital time-sharing is exploited by bus systems in which the memory is partitioned into many independently accessible modules so as to reduce the probability of module access conflicts.

The G-21 is a combination of the two systems. The cost of the portals is reduced by the non-uniformity of portal terminals. Three classifications of transfer rates and accessing modes are available at each module portal. Expensive circuitry is thereby made available only to facilities requiring expensive performance. In addition, not all modules require all three classifications. The portals themselves may differ from module to module.

The G-21 bus circuitry is the high performance access mechanism of the system. By requiring only one-third of the total memory cycle time, the asynchronous bus circuitry can be efficiently shared by at least three users (although only two exist). The bulk of the queue control circuitry is centralized in the two users. Only the communication circuitry is reproduced in each of the right modules. Elaborate queue control circuits are not considered necessary; but certain over-simplifications

can be disastrous. In particular, any fixed priority accessing scheme can lead to facility preempting. With the G-21, fixed priority accessing is avoided by the alternation of priority whenever the possibility (and not the certainty) of conflicting accesses is detected. Thus, one of the two processors is a preferred user. The first of a sequence of conflicting module access requests is always won by the preferred user; but preferential treatment ends at the first conflict. Subsequent contiguous conflicting requests are granted in alternation so that both processors run at half-speed but neither is totally excluded from the module.

In general, time-sharing needs encountered at the program level also appear at the hardware level. Some, such as synchronization, relocation, and protection, are greatly reduced at the program level by provisions at the hardware level. Some hardware features increase efficiency, but not enough to warrant expensive G-21 redesign. Improvements in neither area are being seriously considered for future G-21 development, but should be considered during initial time-shared system design.

APPENDIX

The first reference of this paper does not fully describe the G-21 system. It assumes prior knowledge of the G-20 system. The reader who is interested in G-21 details not covered by this paper may consider the first reference, and the related G-20 material, worthwhile. This appendix is provided for the reader whose only interest in either system is confined to the scope of this paper.

The CDC G-21 is the 64K dual-processor computer system of the CIT Computation Center. It presently supplies the bulk of the general computational service offered by the Center. The configuration includes a 16-channel general telephone exchange interface which handles 40 TTY consoles in the vicinity of the Center. Buffered card equipment, magnetic tape transports, and disc memory are each programmably switchable from one processor to the other. A 450-lpm printer, a paper tape station, and the general exchange interface are manually switchable from one processor to the other. Each processor is equipped with an unswitchable 900-lpm line printer. The visual display system, a 640 million character RACE file system, the higher capacity telephone interface, and a few other extensions are scheduled for the fall of 1965. The million word core memory is expected a year from then.

RCA -
Random
Access
CARD
Equipment

The seven main G-21 memory modules, of 8K words each, house the 56K common memory. In addition, one 8K module houses 4K of private memory for each processor. Another 4K of private memory is housed within each processor main-frame. Thus, each processor can access 56K of common memory and 8K of private memory. The internal 4K of memory operates on 3 μ s read/3 μ s restore cycle. The external modules require 4 μ s for the first read interval of an access sequence. Subsequent contiguous accesses proceed at the 6 μ s internal cycle rate. No direct entries to memory are available other than those to be described in this paper. The general purpose input/output route to memory is through either of the processors. A processor is completely absorbed in the input/output task so that no other commands can be obeyed, by the processor, during the data transfer.

The processors use transistorized, synchronous, d-c coupled logic. The clock rate is 1 mc., two-phase. All arithmetic is performed floating point on 42-bit coefficients with 6-bit octal exponents. The 32-bit memory word is formatted as either a number, a command, a 32-bit boolean operand, or an input/output word. Full 42-bit precision requires two memory accesses. An operand, operand address, or address of an operand address occupies 16 bits of the command word. Six additional command word bits specify one out of 63 core index registers whose

contents may be added to the operand or operand address. The input/output operand is formatted either four 8-bit characters or one 8-bit and four 6-bit characters to the memory word. Command sequences commonly require two to three memory cycles. Thus, the average operation rate is near $20 \mu\text{s}/\text{command}$.

REFERENCES

- [1] Quatse, Jesse T., "The G-21 Multi-Processor System", Engineering Specification 40-19, Control Data Corporation, July 25, 1963.
- [2] Control Data 6600 Computer System Reference Manual, Control Data Corporation, August 1963.
- [3] Bright, Herbert S., "A Philco Multiprocessing System", Proceedings Fall Joint Computer Conference 1964, AFIPS, Volume 26, Part II, Spartan Books, Washington, D.C., 1965.
- [4] GE-635 System Manual, General Electric, July 1964.
- [5] Quatse, Jesse T., "An Integrated Memory Display System", Computation Center Working Paper C.W.O.-15, Carnegie Institute of Technology, Pittsburgh, Pennsylvania, October 5, 1964.
- [6] Kilburn, T., Edwards, D.B.G., Lanigan, M. J., and Sumner, F. H., "One-Level Storage System", IRE Transactions on Electronic Computers, April 1962, pp. 223-35.
- [7] Corbato, F. J., "System Requirement for Multiple Access, Time-Shared Computers", MAC-TR-3, Massachusetts Institute of Technology, Boston, Massachusetts.
- [8] Arden, B. W., Galler, B. A., O'Brien, T. C., Westervelt, F. H., "A Segment Relocation Hardware Proposal and Associated Software", Computing Center, University of Michigan, January 20, 1965.