

Washington University School of Medicine

Digital Commons@Becker

---

Technical Reports

Institute for Biomedical Computing

---

2-1974

## Macromodular Computer Design, Part 1, Volume 3, Design of Phase 1 Macromodules

Computer Systems Laboratory, Washington University

Follow this and additional works at: [https://digitalcommons.wustl.edu/bcl\\_techreports](https://digitalcommons.wustl.edu/bcl_techreports)

---

### Recommended Citation

"Macromodular Computer Design, Part 1, Volume 3, Design of Phase 1 Macromodules," Computer Systems Laboratory, Washington University (1974). *Technical Reports*. Paper 3.  
[https://digitalcommons.wustl.edu/bcl\\_techreports/3](https://digitalcommons.wustl.edu/bcl_techreports/3)

This Technical Report is brought to you for free and open access by the Institute for Biomedical Computing at Digital Commons@Becker. It has been accepted for inclusion in Technical Reports by an authorized administrator of Digital Commons@Becker. For more information, please contact [vanam@wustl.edu](mailto:vanam@wustl.edu).



MACROMODULAR  
COMPUTER DESIGN

PART 1.

DEVELOPMENT OF MACROMODULES

VOLUME III

DESIGN OF PHASE 1 MACROMODULES

**Technical Report No. 46**

FINAL REPORT - FEBRUARY, 1974

CONTRACT SD-302 (ARPA)

COMPUTER SYSTEMS LABORATORY

WASHINGTON UNIVERSITY

ST. LOUIS, MISSOURI



MACROMODULAR COMPUTER DESIGN

FINAL REPORT - CONTRACT SD-302

FEBRUARY 1974

## Technical Report No. 46

PART 1 - DEVELOPMENT OF MACROMODULES

VOL. III - DESIGN OF PHASE I MACROMODULES

This work has been supported by the Advanced Research Projects Agency of the Department of Defense under Contract SD-302 and by the Division of Research Facilities and Resources of the National Institutes of Health under Grant RR-00396.

The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Advanced Research Projects Agency or the U. S. Government.

Computer Systems Laboratory  
Washington University  
St. Louis, Missouri



ABSTRACT

This volume attempts to reconstruct the rationale for the major design decisions that were made in the development of Phase One macromodules, and presents a review of the internal architecture, intercommunication, electrical and mechanical design. Material describing the interfacing of macromodular systems to other computers is also included. The level of detail presented is intermediate, and assumes a general familiarity with the concept of macromodules and their usage.

Two of the design goals proved particularly difficult to meet. The objective that the system designer should not have to be concerned with any details not directly concerned with the functional definition of his system, and the objective of a system discipline that would place no limits on the maximum system size, both required novel approaches to system control and architecture and required an engineering design that was well outside of the prevailing state of the art.

The unusual or non-routine aspects of macromodule design are emphasized, particularly those points of the design that proved most challenging or intellectually rewarding. A preliminary evaluation of our experience to date in the area of reliability is given.



## FOREWORD AND ACKNOWLEDGMENTS

This volume attempts to summarize the major design and engineering decisions that were made in the course of developing Phase I macromodules. Chapter 2 is based upon a draft of a Computer Systems Laboratory Technical Report by John Newell, supplemented with material on control macromodules written by M. L. Pepper, and material on memory-related macromodules contributed by A. D. Richardson and H. C. Lewis. The mechanical design is described in Chapter 3 by R. J. Arnzen, and the electrical design description in Chapter 4 is by F. U. Rosenberger. G. C. Johns provided the interconnection summary of Chapter 5, and the review in Chapter 6 of power supply and system-wide control signals was written by T. J. Chaney. The discussion of interfacing was written by J. A. Greenfield and T. C. Perry, with the former contributing the detailed description of the PDP-11/40 interface.

Special thanks are due to Christine D. Coaker for assembling, editing, and proofreading this material, to David Shupe for photographic assistance, and to Phillip Lewis for drafting and other documentation support.

Responsibility for accuracy and choice of material, however, rests with the undersigned, who hopes that no serious misunderstandings are created by the inevitable errors and occasional misinterpretations that may have occurred.

Charles E. Molnar, Director  
Computer Systems Laboratory



TABLE OF CONTENTS

1. Introduction . . . . .	1
2. Data and Control Structure . . . . .	3
2.1 Introduction . . . . .	3
2.2 System Considerations . . . . .	16
2.2.1 Function Code . . . . .	22
2.2.2 Word Length Extension . . . . .	24
2.3 Details of Control at the Module Level . . . . .	24
2.3.1 Storage Modules . . . . .	25
2.3.2 Data-Changing Modules . . . . .	41
2.3.3 Decision Modules . . . . .	62
2.3.4 Miscellaneous Modules . . . . .	65
2.3.5 Control Modules . . . . .	68
3. Mechanical Design . . . . .	73
3.1 Introduction . . . . .	73
3.2 Functional Description and Rationale of the Support Set . .	74
3.3 Manufacture of Macromodular Mechanical Hardware . . . . .	82
3.4 Costs of Fabrication . . . . .	84
4. Electrical Considerations . . . . .	86
4.1 Introduction to Electrical Design . . . . .	86
4.2 Classes of Interconnections . . . . .	87
4.3 Transmission Line Terminations . . . . .	87
4.3.1 Sending End or Series Termination . . . . .	87
4.3.2 Shunt or Receiving End Termination . . . . .	91
4.3.3 Double-Ended Termination . . . . .	93



4.4	Bidirectional Bus for Memory . . . . .	97
4.5	Crosstalk . . . . .	99
4.6	Delays . . . . .	101
4.7	Noise Margin vs. Temperature . . . . .	102
4.8	Power Supply Sensitivity . . . . .	103
4.9	An Example . . . . .	103
4.10	Introduction to Electrical Testing and Inspection . . . . .	105
4.11	Evaluation and Qualification of ECL Circuits . . . . .	106
4.11.1	Feedthrough . . . . .	106
4.11.2	Production Testing of Integrated Circuits . . . . .	108
4.12	Printed Circuit Board Testing. . . . .	109
4.13	Other Tests of Components . . . . .	110
4.14	Module Testing . . . . .	110
4.15	Test Results . . . . .	111
4.16	Appendix I - Causes of Feedthrough in ECL Circuits . . . . .	112
4.16.1	Collector to Base Capacitance . . . . .	115
4.16.2	Emitter to Base Capacitance . . . . .	117
4.16.3	Current Source Capacitance . . . . .	118
4.16.4	$C_A$ Discharge . . . . .	118
4.16.5	Wired-OR Feedthrough . . . . .	122
4.16.6	Bias Network and Chip Capacitance Coupling . . . . .	122
4.17	Appendix II - Restrictions on MECL II . . . . .	123
5.	Interconnection Design . . . . .	127
5.1	Scope . . . . .	127
5.2	Background . . . . .	127
5.3	Electronics Package . . . . .	129
5.4	Faceplate Box . . . . .	133
5.5	Cables . . . . .	135



5.5.1	Data Cable . . . . .	135
5.5.2	Control Cable . . . . .	138
5.5.3	Daisy Chain . . . . .	138
5.5.4	Cable-Related Devices. . . . .	139
5.6	Frame Components . . . . .	140
5.7	Performance . . . . .	141
5.7.1	AMPMODU . . . . .	141
5.7.2	AMP 750 . . . . .	141
5.7.3	ASTRO-348 . . . . .	142
5.7.4	Data Cable . . . . .	142
5.7.5	Control Cable . . . . .	143
5.7.6	Power Connectors . . . . .	143
5.7.7	General Performance . . . . .	143
6.	Power and System-Wide Control Signal Structure . . . . .	145
6.1	Introduction . . . . .	145
6.2	Pedestal Control Functional Description . . . . .	145
6.2.1	Design Guidelines . . . . .	145
6.2.2	Interconnection of Pedestals and Controller . . . . .	147
6.2.3	Interconnection of Pedestals and Frame Blocks . . . . .	147
6.2.4	Definition of Daisy Chain Signals . . . . .	148
6.2.5	Definition of Pedestal-to-Frame Signals . . . . .	150
6.2.6	Pedestal Power Malfunctions . . . . .	152
6.2.7	Pedestal Controller Functional Description . . . . .	153
6.2.8	Controller User Description . . . . .	155
6.3	Power and System-Wide Control Signal Structure Design Considerations . . . . .	158



6.3.1	Isolation of Power and Signal Grounds . . . . .	158
6.3.2	Base Pedestal Power System . . . . .	161
6.3.3	Module Supply . . . . .	164
7.	Interfacing . . . . .	168
7.1	Introduction . . . . .	168
7.2	PDP-11 to Macromodule Interface . . . . .	170
7.2.1	Interaction of the PDP-11 with Other Devices . . . . .	170
7.2.2	Summary of Interface Description; Conventions . . . . .	171
7.2.3	Control Ports . . . . .	173
7.2.4	Data Ports . . . . .	175
7.2.5	Synchronization . . . . .	181
7.2.6	System Preset via the Interface . . . . .	183
8.	Conclusion . . . . .	186
9.	References . . . . .	



## 1. INTRODUCTION

The purpose of this volume is to present some of the important details of the implementation of Phase I macromodules. The level of detail of much of the material is intermediate. It is assumed that the reader has substantial familiarity with the general concept and physical and functional description of macromodules, such as might be gained through a reading of Volumes I and II of Part 1 of this report series (also indexed as Technical Reports 44 and 45 of the Computer Systems Laboratory); further details of the implementation are to be found in Part 2 of this report series (Volumes I through XIV). The latter documents are also indexed as Technical Reports 30 through 43.

Most of the material in this volume represents an attempt to "reconstruct the rationale" underlying many of the design choices made during the development of macromodules, and hence represents an after-the-fact assembly of material culled from documents on file, from individual recollections, and occasionally from rethinking and interpolation based on these sources. Hopefully, a major part of the useful thought that has gone into the process of designing and building macromodules has been captured; undoubtedly an attempt such as this to record a collective enterprise will conflict in some details with individual recollections and occasionally with the facts. In particular, an attempt has been made to emphasize understandability, occasionally with the loss of some historical accuracy.

Chapter 2 reviews the essential features of the internal structure of macromodules, particularly some of the important features of the internal signalling and communications structure that is not visible to the user but is essential to making his task safe and easy. Some points of particular interest that are dealt with in this chapter include the use of transition logic for much of the internal control, the handling of control and data timing in structures of unrestricted extent, and the means of defining and sensing a variety of conditions that determine the propagation of data and control signals across the boundaries separating physically adjacent modules.

It is worth noting that the nucleus of Chapter 2 was assembled by Dr. John Newell, who played no part in the design process but set about to describe and document macromodule internal structure from the point of view of a learner and a teacher. His strenuous efforts produced the general approach and style of the Chapter, as well as the description of the basic macromodular data processing structure, the manifold. The remaining material in the chapter was contributed by numerous others.

Chapter 3 reviews important aspects of the mechanical design of macromodules, which in many respects was the most difficult and time-consuming part of the overall task. While absolutely essential to the overall goal of ease of system assembly and use with minimum opportunity for erroneous assembly, an appropriate mechanical design proved to be difficult and expensive. Strict adherence to the system goals proved to be very difficult to maintain in the face of the problems and it is remarkable that so few concessions were made.



Here, as in other parts of the development effort, there were serious conflicts between the intellectual and academic goals of testing a concept in a clean and thorough manner and the more practical and pressing objective of accomplishing an inventory that worked and could be used to solve problems. These conflicts were usually resolved by taking more time so that both kinds of goals could eventually be met.

Chapter 4 deals with electrical considerations in the design. The primary problems in this area were posed by the need to specify interconnection conventions in a manner that would not restrict system size, and to provide a design that was robust and would allow free interchanges of modules without hazard to system integrity. The effort to carry out a true worst-case design with respect to power supply, cooling, ground noise, signal propagation delay, and secondary device properties required extraordinary attention to numerous engineering issues that have not been conventionally addressed; some of the most difficult electrical design problems and some of the most important insights came from this attempt to carry out worst-case design for a system of unlimited extent.

It was recognized almost immediately by everyone concerned with macromodule development that connector technology would present one of the most important limitations to the feasibility and reliability of macromodular systems, and this observation was in no way refuted by our experience. The issues of physical size and convenience of cables and interconnections were in direct conflict with the electrical requirements and those of reliability, and several false starts were made before connectors and cable designs adequate to the task were obtained. Chapter 5 reviews this part of the development, which eventually produced an almost entirely satisfactory solution that in most respects has exceeded the hopes, let alone the expectations, of the designers.

The rather unique power supply system for macromodules is discussed in Chapter 6. This chapter also deals with system-wide control signals that interact very strongly with the signals used to control power supply sequencing and to communicate information about failures of power supply or cooling. The indefinitely extensible nature of macromodular systems posed special problems that were solved through the use of a relatively high voltage (55 volts) direct current primary distribution system with small DC-DC converters in each module producing the particular secondary voltages needed in each module.

One particularly important issue that arose early in the thinking about macromodular design concerns communication between subsystems that do not have a common time reference. While a discussion of this problem (the "glitch") could properly have constituted a chapter of this volume, the amount of material and its widespread interest have led to its publication as a separate volume, Volume IV of Part 1 of this final report. It is also indexed as Technical Report 47 of the Computer Systems Laboratory.



## 2. DATA AND CONTROL STRUCTURE

### 2.1 INTRODUCTION

This report contains a brief functional description of each of the circuit elements employed by Phase I macromodules to generate, manipulate, or propagate control signals within a data-processing manifold, and also within elements of the control modules. A block diagram of the control circuitry of each of the data-processing and storage modules is presented along with a detailed discussion of the control therein. Control signal generation and propagation needed to accommodate the data transfers within and between data-processing manifolds is examined in detail. Word length extension creates a special set of control requirements which are also discussed. In addition, the operation of the control modules is described briefly.

Any macromodular computing system will consist of one or more arrays of modules dedicated to data storage, data changing, decision making, or control. The modules are housed in a cellular frame structure which provides cooling air, operating voltages, and certain implicit pathways to accommodate data and control signals.

The modules each consist of an electronics package, a faceplate box, and (for most) an overlay panel for the faceplate box. The electronics package contains connectors which provide access to signal pathways in the vertical and horizontal channels, and (via data and control cable ports on the faceplate box) to explicit signal pathways and to code switches.

Data are distributed on internal and external pathways as level signals. That is, one of the two permissible logic levels represents a binary zero, and the other indicates a binary one. External control signals are transition signals, by which the following is meant: A control signal is asserted (is made active) by a change in the level of the signal. Thus, the transition from zero to one (or from one to zero) indicates the arrival of a control signal. Internal control signals are sometimes level signals and sometimes transition signals.

The macromodules are designed to manipulate (or store) twelve-bit blocks of data. Horizontal extension capability permits the word length to be any integer multiple of twelve bits. Thus, it is clear that lateral communication (from module to module) is required.

Typically a storage module forms the base of each column of modules within a data-processing manifold. No a priori information exists concerning the height or width of a data-processing manifold. It is therefore necessary to define vertical and horizontal boundaries. The vertical boundaries are referred to as "delivery boundaries" and "transfer boundaries", while the terms "left and right manifold boundaries" apply in the horizontal plane. Incidentally, the right manifold boundary is to one's right when



facing the faceplate boxes of the data-processing manifold. The boundary sensing circuits within the modules behave as though the module were at a boundary position in the absence of a positive signal (from an adjacent module) indicating otherwise. The concept of boundaries will be explored more thoroughly later in this document.

The operation of the macromodules is asynchronous. Those which do not continuously process data respond to a control signal (initiation). At the conclusion of their operation they emit a control signal (completion) which indicates that the data manipulation is complete and that sufficient time has elapsed for the new data to have been transmitted to all destinations.

As indicated previously, data is distributed on implicit and explicit pathways. The implicit pathways are the two vertical busses each of which accommodate thirteen data bits (twelve-bit word plus a flag bit), while the pathways are specially constructed data cables which generally provide twelve data channels and two control signal channels. Control signals are also transmitted by vertical and horizontal implicit pathways.




Figure 1. A data-processing column.



Figure 1 represents one column of a data-processing manifold. The storage module at the base of the column receives data from the data-processing modules above it via the Down Bus and supplies the updated data to the Up Bus. The Up Bus is monitored continuously by those modules which receive data therefrom, but the modules which supply data to the Down Bus must be switched on to the bus. Figure 2 provides a simplified view of the bus circuitry within the modules.

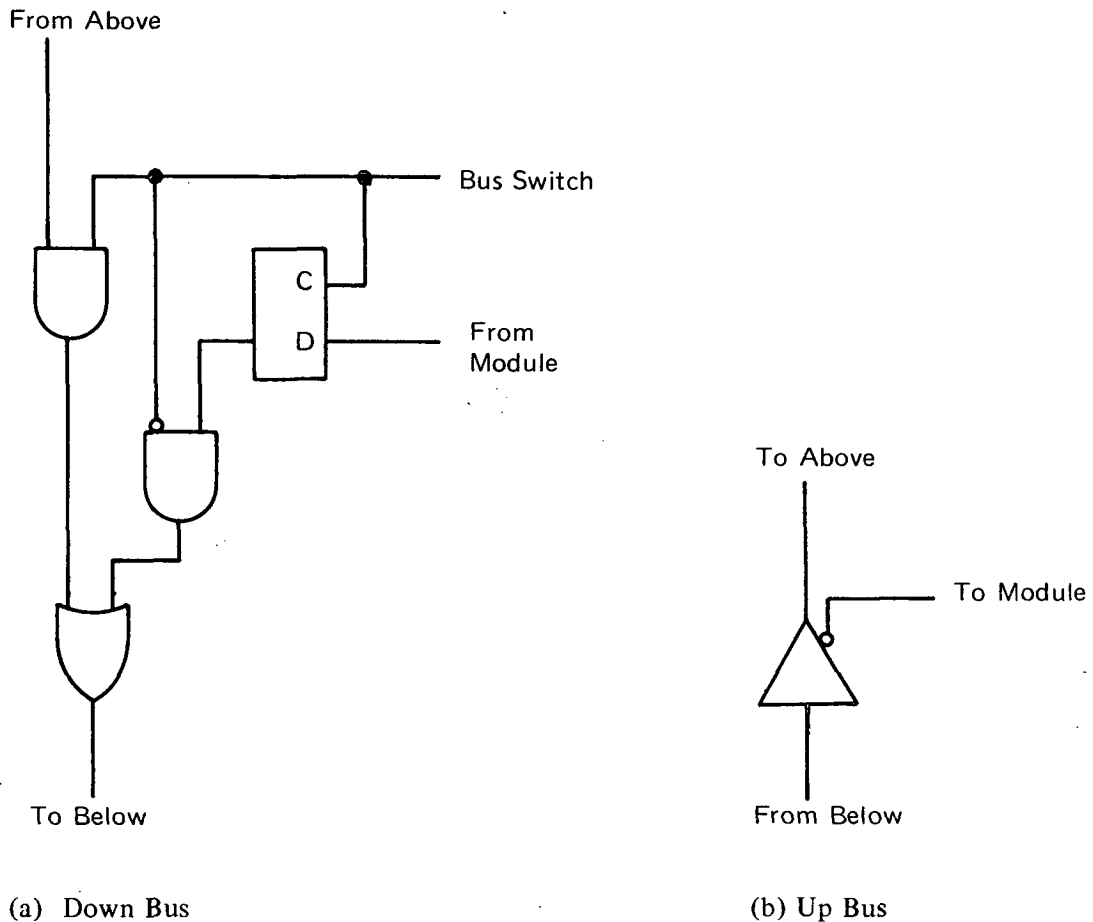


Figure 2. Simplified bus circuitry.



Consider the Down Bus circuit as depicted (per bit) in Figure 2a. The signal labelled "Bus Switch" controls the clock input to the module's output register and allows either the data from above or the data from within to be supplied to the Down Bus. While the clock is enabled, the contents of the module's output register may be changing, but during this condition the data from above is passed through. It is only when the clock is disabled and the contents of the output register are latched that these contents are switched on to the Down Bus. During normal operation the Bus Switch signals within the modules in a column are handled in such a way that not more than one module output register's contents are supplied to the Down Bus at any given time. Note that it is quite possible that no module is supplying data to the Down Bus. As stated previously, and as indicated in Figure 2b, each module continuously monitors and amplifies the contents of the Up Bus. Since the path length between modules is relatively short, and since each data bit is amplified in each module, data are transmitted single rail via the busses.

When a data-changing module transmits data to the storage module, a control signal (called the transfer signal) accompanies the data on a parallel pathway. New data are transmitted via the Up Bus to the data-processing modules in the column as well as to the data cable ports on the faceplate of the storage module. Data delivery signals accompany the data to each destination. It is not until a return signal is received from each of these data delivery destinations that a transfer return signal is sent from the storage module to the data-changing module which initiated the transfer. Figure 3 illustrates the distribution of these transfer and delivery signals.

Each storage module is assumed to occupy the lower boundary of a data-processing column unless a signal to the contrary is transmitted to the storage module from below. Two different modules which would sit below (and communicate with) a register have been conceived (and to some extent designed) but never built.

In data-changing modules, the control signal which called the module to action (the initiation signal) is supplied to the rightmost module of an extended group from an explicit control pathway and from there to extended modules on the bus via an implicit control pathway. Installing a faceplate overlay will expose the initiation and completion terminals, and establish the module in question as a right manifold boundary through the use of a sense switch.

In the paragraphs which follow, the implicit control signal distribution associated with each type of data-processing module will be discussed in detail.

The intent of this report is to present a compact drawing and companion text explaining the control circuitry of each type of module. Of necessity, certain abbreviations will be employed; these are listed in Table 1 at the end of this section. The control drawings will contain functional blocks labelled with a shorthand notation. A brief description of these functional blocks is now presented.



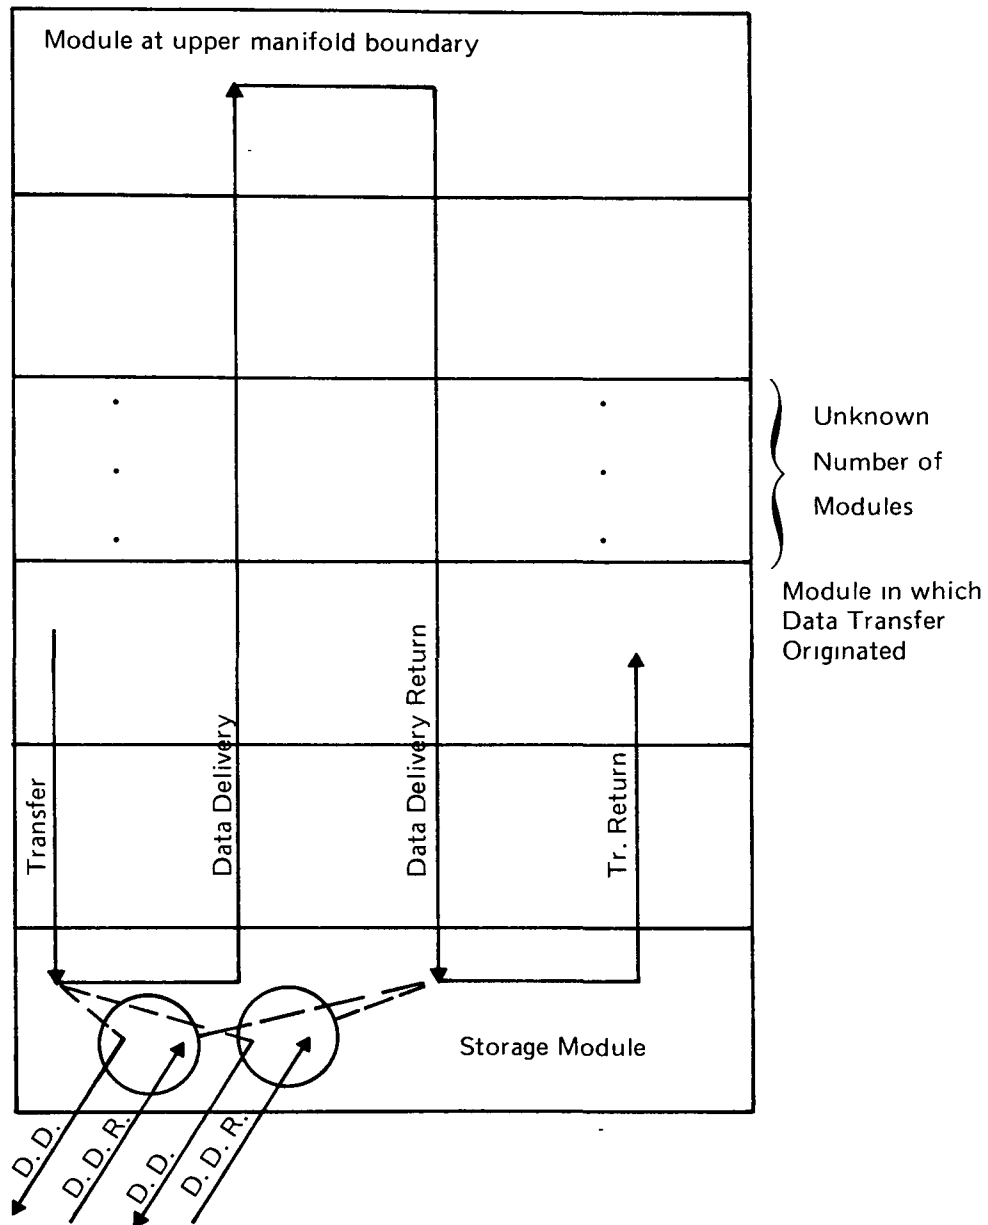


Figure 3. Transfer and data delivery signal distribution.



### Merge Element

As depicted in Figure 4, the merge element may have any number of inputs, and one output. An output signal is produced in response to exactly one input signal. The input and output signals of a merge element are transitions. Although a merge element may be built to accommodate any number of inputs, the typical unit has two inputs.




figure 4. Merge element.

### Rendezvous Element

The rendezvous element produces an output signal after all of its input signals have been activated. Figure 5 represents a rendezvous element with an indeterminate number of inputs. The operation is as follows. All inputs are initially forced to one of the two permissible logic levels. This forces the output to a logic level. Subsequently, as the inputs change levels (undergo transitions) the output remains unchanged until all inputs have changed. There are no constraints on the order of arrival of inputs. Rendezvous elements with any number of inputs may be built, but the typical unit has two inputs.




Figure 5. Rendezvous element.



### Input Selector

The input and output signals of the previous elements have been transitions. This condition is symbolized by the arrowheads on the signal lines. The Input Selector has two transition inputs and one level input (denoted by the diamond-shaped symbol on the line). The value of the level input (input A in Figure 6 ) determines which of the two transition inputs is coupled through to the output. By convention, the input which is physically closest to the level input on the drawing is coupled through to the output when the level input is asserted. Thus, in Figure 6, if input A is asserted a transition on input line X produces an output transition. Otherwise an output transition is produced by a transition on the Y input line.




Figure 6. Input selector.

### Condition Sensor

This element, shown in Figure 7 , is a generalization of two of the control elements to be discussed. In the condition sensor, a transition on input I produces a transition on exactly one of n output lines under the control of n input level signals. In Figure 7 the inputs  $C_1$  through  $C_n$  are levels. When input I becomes active (undergoes a transition) and  $C_i = 1$ , output  $Z_i$  undergoes a transition and the input I becomes inactive. (Note that I does not undergo another transition when it becomes inactive. The circuit is simply made ready to accept another transition on line I). While the order in which I becomes active and  $C_i$  takes on the value 1 is not important, only one of the  $C_i$  may have the value 1 while I is active. If  $C_i = 1$  its value may not change while I is active. (That is, until  $Z_i$  undergoes a transition.)






Figure 7. Generalized condition sensor.

### Output Selector

Figure 8 depicts a special case of the Condition Sensor. In this device, a single level signal determines which of two output signals will undergo a transition in response to activity on a single input line. When  $C$  is asserted, the output line which is physically closest to  $C$  on the drawing responds to signals on  $I$ . Thus, in Figure 8, output  $Z_1$  will respond to transitions on  $I$  when  $C$  is asserted. As indicated above,  $C$  must not change value between a transition on  $I$  and the subsequent output transition.




Figure 8. Output selector.



### Conditioned Pause

Another special case of the Condition Sensor, the single output case, is shown in Figure 9 . The transition input I produces a transition at the output if C = 1 (is asserted). If C is not asserted at the time I undergoes a transition, no output transition is produced until C is asserted.




Figure 9. Conditioned pause.

### Inhibit Gate

The gate shown in Figure 10 passes a transition at the input to the output if the signal C is not asserted. Assertion of C prevents the transition at the input from appearing at the output. The inhibit gate is typically used to block undesired input signals (at boundaries, for example). A transition at I while C is asserted produces no output.




Figure 10. Inhibit gate.



### Conditional Rendezvous

As indicated in Figure 11, two types of conditional rendezvous elements exist. In Figure 11a, when C is asserted a transition on  $I_1$  causes a transition at the output (Input  $I_1$  is physically closest to C in the drawing). If C is not asserted, the device functions as a rendezvous. Figure 11b has two level inputs,  $C_1$  and  $C_2$ . If only  $C_1$  is asserted,  $I_1$  is coupled through to the output; if only  $C_2$  is asserted,  $I_2$  is coupled through to the output. If neither  $C_1$  nor  $C_2$  is asserted, the device functions as a rendezvous.

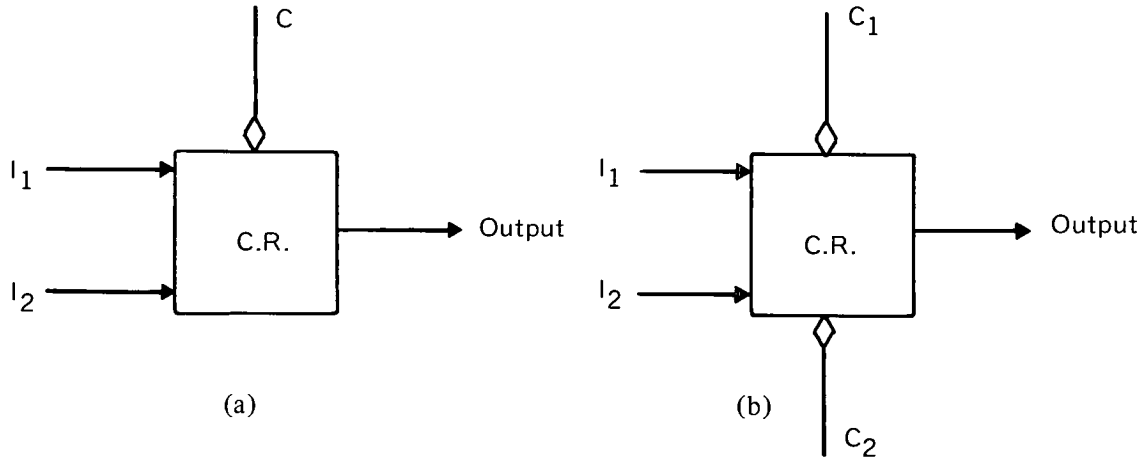


Figure 11. Conditional rendezvous elements.

### F Element

The F element (Figure 12) is a memory element which is set by transition signals on one input line and cleared by transition signals on the other input line. The outputs are levels, and by convention the output nearest a given input line is asserted by a transition on said input line. For example, in Figure 12, a transition on the Set line will assert the F output, while a transition on the Clear line will assert the  $\bar{F}$  output.




Figure 12. F element.



### D Element

Figure 13 represents the general form of a D (Decision) element. As indicated, the element has input signals grouped in bundles of size  $m_1$  through  $m_n$ . The number of outputs equals the product of the numbers of inputs in each input bundle. That is,

$$k = \prod_{i=1}^n m_i.$$

To produce an output, exactly one input per bundle must undergo a transition. When this happens a transition is produced at exactly one output terminal.




Figure 13. Generalized D element.

### Two-by-One D Element

Figure 14 depicts the special case of the D element for which the input bundles contain one and two wires, respectively. In this case Input  $I_3$  must be active and either  $I_1$  or  $I_2$  must be active (the order in which they undergo transitions is not important) to produce a transition on the output line opposite the active member of the two-input bundle. Thus inputs  $I_2$  and  $I_3$  produce an output on  $Z_2$ , while inputs  $I_1$  and  $I_3$  produce an output on  $Z_1$ .






Figure 14. Two-by-one D element.

### Call Element

Consider now an element consisting of a combination of a three-input Merge element and a three-by-two Decision element. Such a device is shown in Figure 15. A transition at any one of the inputs ( $I_1, I_2, I_3$ ) will produce a transition at the output labelled  $Do$ . The input just discussed is also applied to the D element. The signal from the  $Do$  terminal can be used to initiate some sequence of events external to the Call element. If the external device supplies a transition signal to  $R_1$  or  $R_2$  at the conclusion of its operation, a transition will be produced at  $C_{1.1}$  or  $C_{2.1}$  (assuming that the original input was on  $I_1$ ). Thus, for example, an input on  $I_2$  will produce an output on the  $Do$  line, and this input along with a subsequent input on  $R_1$  will produce an output on  $C_{1.2}$ . A return on  $R_2$  along with the original input on  $I_2$  will produce an output on  $C_{2.2}$ . The three-input two-return call element described above is the Call element found in the CALL Module. A simpler two-input, one-return version is used as a control element within some of the data-processing modules.




Figure 15. Call element.



Table 1. Abbreviations Employed

<u>Abbreviation</u>	<u>Meaning</u>
C. INH.	Inhibit (from configuration gate)
C. R. F. B.	Right Field Boundary (from configuration gate)
C. TR.	Configuration Transfer
C. TR.R.	Configuration Transfer Return
D.D.	Data Delivery
D. D. B.	Data Delivery Boundary
D. D. R.	Data Delivery Return
FN. D.	Function Delivery
FN. D. R.	Function Delivery Return
INH.	Inhibit
IN. OP.	Initiate Operation
IN. TR.	Initiate Transfer
L. F. B.	Left Field Boundary
L. M. B.	Left Manifold Boundary
RES. RDY.	Result Ready
R. F. B.	Right Field Boundary
R. M. B.	Right Manifold Boundary
R. LOC.	Local Right Boundary
TR.	Transfer
TR. ANCH.	Transfer Anchor
TR. B.	Transfer Boundary
TR. C.	Transfer Complete
TR. R.	Transfer Return



## 2.2 SYSTEM CONSIDERATIONS

As indicated in the Introduction, the macromodules are housed in a rectangular array of cells with a frame structure. The frame actually consists of a collection of frame blocks which are stacked above supporting units called pedestals. The pedestal units contain energy conversion units, as well as circuitry associated with energy distribution and system initialization signals.

Each frame block contains five columns of four cells each, as shown in Figure 16. The center column contains a Fan module which distributes operating voltages, initialization signals, and cooling air to the other cells in the frame block. Horizontal interconnection is accomplished by fifteen horizontal signal paths which interconnect adjacent cells, skipping past the center column. A special-purpose coupling unit is available for extending the fifteen horizontal signal paths across the frame block boundaries.




Figure 16. A frame block.

Forty-five vertical signal pathways are provided for data-processing modules by means of internal jumpers in the faceplate boxes. When an electronics package is inserted into a frame cell, the connectors on the module mate with connectors on two vertically adjacent faceplate boxes and a connector mounted to the frame block assembly. Figure 17 presents a side view of the interface between a module, the frame-mounted connector and the two faceplate boxes.






Figure 17. Modules, frame, and faceplate connectors.



The names Top, Lateral, Faceplate, and Bottom in Figure 17 refer to the electronics package motherboards to which the respective connectors are attached. The fifteen horizontal signal paths are provided (along with certain energy distribution and system control paths) by the connector labelled L. This connector is mounted to the frame assembly and thus is independent of the faceplate box.

Several different types of faceplate boxes exist. The primary differences between types are the methods of connecting between connector F and the various connectors on the front of the faceplate. If two vertically adjacent modules should communicate with each other via the implicit data pathways, the faceplate box used with the upper module will provide pin-to-pin jumpers between the bottom connector on the upper module and the top connector on the lower module.

Following is a list of the major signal groupings and the symbols used to represent them:

F	A 4-bit operand on the Mode Bus (see Table 2).
Z	A twelve-bit operand on the Down Bus.
Z <sub>f</sub> or Z <sub>12</sub>	Down Bus Flag Bit.
R <sub>f</sub> or R <sub>12</sub>	A twelve-bit operand on the Up Bus (above register).
R <sub>f</sub> or R <sub>12</sub>	Up Bus Flag Bit.
T <sub>f</sub> or T <sub>12</sub>	A twelve-bit operand on the Up Bus (below register).
M	A four-bit function code which may be taken from a vertical parameter bus, from horizontal signal paths (from the right), or from a set of switches on the faceplate box.
D	A twelve-bit operand taken from an input data cable port of the faceplate box.
S	A twelve-bit selection variable supplied via an input data cable port on the faceplate box.

Provision is also made for a two-bit vertical bus (Inhibit and Right Field Boundary) which is referred to as the "Configuration Bus". Tables 2 and 3 present a summary of the vertical and horizontal implicit signal pathways. These tables, however, do not apply to two double-height modules, the General Memory Controller and the Multiply module. Tables 6 and 7 on pp. 31-32 give horizontal and vertical pathways in the GMC module. Table 11 on p. 56 indicates horizontal paths in the Multiplier.

For the purposes of this paper the data-processing modules may be divided into four categories, namely: Storage (Register, General Memory Controller and its associated Memory); Data-Changing (Logic, Load, Arithmetic, Shifter, Multiplier); Decision (Decode, Comparator); and Miscellaneous (Data Branch, Digital to Analog Converter, Data Gate, \* Configuration Gate \*). Table IV is a summary of the data-processing modules with their communication paths. The words "To" and "From" in the Bus columns indicate whether the module accepts data from or supplies data to the appropriate bus. Similarly the entries "Trans." and "Block" mean that the module is transparent to (passes) or blocks signals on the bus involved.

---

\* These two modules have been designed to some extent, but never built.



Table 2. Vertical Signal Paths\*

<u>Pin</u>	<u>Direction</u>	<u>Below Register</u>	<u>Above Register</u>
2, 4, 6, 8	Up	T <sub>3</sub> , T <sub>2</sub> , T <sub>1</sub> , T <sub>0</sub>	R <sub>3</sub> , R <sub>2</sub> , R <sub>1</sub> , R <sub>0</sub>
10, 12, 14, 16	Down	Not Used	Z <sub>0</sub> , Z <sub>1</sub> , Z <sub>2</sub> , Z <sub>3</sub>
18, 20, 22, 24	Up	T <sub>7</sub> , T <sub>6</sub> , T <sub>5</sub> , T <sub>4</sub>	R <sub>7</sub> , R <sub>6</sub> , R <sub>5</sub> , R <sub>4</sub>
26, 28, 30, 32	Down	F <sub>0</sub> , F <sub>1</sub> , F <sub>2</sub> , F <sub>3</sub>	Z <sub>4</sub> , Z <sub>5</sub> , Z <sub>6</sub> , Z <sub>7</sub>
34, 36, 38, 40	Up	T <sub>11</sub> , T <sub>10</sub> , T <sub>9</sub> , T <sub>8</sub>	R <sub>11</sub> , R <sub>10</sub> , R <sub>9</sub> , R <sub>8</sub>
42	Down	Not Used	Z <sub>8</sub>
44	Down	C.TR.R.	Z <sub>9</sub>
46	Down	D.D.	Z <sub>10</sub>
48	Down	FN.D.	Z <sub>11</sub>
50, 52, 54		Not Used	Not Used
56	Up	C.R.F.B.	TR.R.
58	Down	TR.R.	TR.
60	Down	R.F.B.	D.D.R.
62	Both	TR.B.	TR.B.
64	Both	Not Used	D.D.B.
66	Down	R <sub>12</sub>	FN.D.R.
68	Up	Not Used	INH.
70	Up	C.INH.	R. F. B.
72	Down	INH.	Z <sub>12</sub>
74	Up	Not Used	FN.D.
76	Up	TR.	D.D.
78	Up	FN.D.R.	R <sub>12</sub>
80	Up	D.D.R.	F <sub>0</sub>
82	Up	TR.ANCH.	F <sub>1</sub>
84	Up	Not Used	F <sub>2</sub>
86	Up	C.TR.	F <sub>3</sub>
88, 90* *		Not Used	Not Used

NOTE: All odd-numbered pins are grounded.

Pin numbers refer to CSL wiring lists.

\* Vertical signal paths in the General Memory Controller module are given in Table 6.

\*\* Pin 90 is used to distribute -5.2 volts in the Data Branch Module (bottom).



Table 3. Horizontal Signal Paths\*

<u>Pin</u>	<u>Direction</u>	<u>Arithmetic</u>	<u>Compare</u>	<u>Load &amp; Logic</u>	<u>Shift</u>
1, 89	←	F <sub>3</sub>	F <sub>3</sub>	F <sub>3</sub>	F <sub>3</sub>
3, 87	←	F <sub>2</sub>	F <sub>2</sub>	F <sub>2</sub>	F <sub>2</sub>
5, 85	←	F <sub>1</sub>	F <sub>1</sub>	F <sub>1</sub>	F <sub>1</sub>
7, 83	←	F <sub>0</sub>	F <sub>0</sub>	F <sub>0</sub>	F <sub>0</sub>
9, 81	←	Inhibit	Inhibit	Inhibit	Inhibit
11, 79	→	Tr. Comp.	Tr. Comp.	Tr. Comp.	Tr. Comp.
13, 77	→	L. M. B.	L. M. B.	L. M. B.	L. M. B.
15, 75	→	L. F. B.	L. F. B.	L. F. B.	L. F. B.
17, 73	←	In. Op.	In. Op.	In. Op.	In. Op.
19, 71	←	R. D. B.	R. D. B.	R. D. B.	R. D. B.
21, 69		Spare	Spare	Spare	Spare
23, 67	←	C = 0	Not Used	Not Used	Bit from Rt.
25, 65	→	Not Used	Condition 0	Not Used	Bit From Left
27, 63	→	In. Tr.	Condition 1	In. Tr.	Result Ready
29, 61	←	C = 1	Not Used	Result Ready	In. Tr.

The following are common to all module types:

31	Power Down Request
33	Protect (Data Shield)
35	Power Sense (Positive)
37	Power Sense (Negative)
39	Preset
41	Power Down Acknowledge
43	Spare
45	Ground
47, 49, 51	Power (Negative)
53, 55, 57	Power (Positive)
59	Ground

NOTE: Pin Numbers refer to CSL wiring lists.

All even-numbered pins are grounded.

\* Horizontal signal pathways in the Multiply module are given in Table 11; in the General Memory Controller, in Table 7.



Table 4. Module Communication Via Implicit Data Paths

<u>Module</u>	<u>Up Bus</u>	<u>Down Bus</u>	<u>Extendable ?</u>
Register (Above)	To	From	No
(Below)	From	Block**	No
General Memory Controller (and Memory)	To	From	Yes
Logic	Pass & From	To	Yes
Load	Pass *	To	Yes
Arithmetic	From	To	Yes
Shift	From	To	Yes
Multiply	From	To	Yes
Compare	From	Trans.	Yes
Decode	From	Trans.	No
D/A Converter	From	Trans.	No
Data Branch	To	Block**	No
Data Gate***	To	N.A.**	Yes
Configuration Gate***	Pass	N.A.**	Yes

\* The Load module is transparent to the twelve data bits on the Up Bus. However, it takes the Flag from the Up Bus and puts it on the Down Bus.

\*\* The Down Bus doesn't exist below a Register or Data Branch. The Data Gate is designed to sit below a Register.

\*\*\* This portion of the design has been defined for these two modules, although they were never built.



### 2.2.1 Function Code

Some of the data-processing modules are capable of performing more than one function. The four-bit parameter bus carries a signal called the function code, which enables the user to select one of a group of operations to be performed by a multi-purpose module.


As indicated in Tables 2 and 3, the function code is distributed both vertically and horizontally. The vertical parameter bus originates in the Register, where it receives its input from the four least significant bits entered via an input data cable port. This port may be connected to a data cable or to a set of switches (parameter plug) with which the binary values may be established. The function code may also be supplied to a module from vertical or horizontal implicit data pathways or from a set of switches on the faceplate associated with the module receiving the code.

Some modules which are controlled by the function code have a "Mode Control Board" which (among other things) selects the source of function code inputs for the module in which it is located according to the following rules:

- 1) If the module occupies a right manifold boundary position (as indicated by the position of a switch on the faceplate), then:
  - a) If the faceplate parameter switches are set at 17 (octal), the function code is taken from the vertical four-bit parameter bus (mode input).
  - b) If the faceplate switches are set at any octal number other than 17, the function code is taken from the faceplate switches.
- 2) If the module occupies any position other than a right manifold boundary, then:
  - a) If the faceplate switches are set at 17, the function code is taken from the module to the right via the horizontal implicit data paths;
  - b) If the faceplate switches are set at any octal number other than 17, the switch settings provide the function code.

The function code switches and the boundary switch are automatically positioned when the appropriate faceplate box overlay is installed. In the absence of an overlay, the mode control board circuitry behaves as though the switch setting were 0 and the module rightmost.





Symbols:  $\odot$  — local delay;  $\boxed{R}$  — rendezvous.

Figure 18. Propagation of implicit control signals.



### 2.2.2 Word Length Extension

The propagation of implicit control signals in a group of extended modules is depicted (in simplified form) in Figure 18 . A thirty-six-bit manifold is shown in which the currently active group of modules is located at some intermediate vertical location.

An Initiate Operation signal is supplied to the rightmost module of the active group via an explicit control signal pathway and a terminal on the faceplate box. This Initiate Operation signal is passed to the module to the left as quickly as possible. A Result Ready signal is also propagated to the left after a time delay sufficient to guarantee that the current data processing has been completed. The Result Ready signal is transmitted to the left by an extended module when both the Initiate Operation signal (locally delayed) and the Result Ready signal from the module to the right have been received.

The arrival of the Result Ready signal at the Left Manifold Boundary indicates that sufficient time has passed (since the Initiate Operation signal was received) for all of the modules in the currently active row to have completed their data-processing operations. Therefore, the newly computed data value is transferred to the storage module via the Down Bus and a Transfer signal is simultaneously transmitted downward in the leftmost column.

At the Left Manifold Boundary, the Result Ready signal is employed to produce an Initiate Transfer signal. This signal propagates from left to right as quickly as possible, triggering the transfer of data and transmission of the Transfer signal in each data-processing column in turn.

As stated in the Introduction, when the storage module receives the Transfer signal, a Data Delivery signal is transmitted to the upper boundary of the column, from whence a Data Delivery Return signal is sent back to the storage module. Upon receipt of the Data Delivery Return signal (from all data destinations) the storage module produces a Transfer Return signal which is transmitted to the module in which the Transfer signal originated.

When any of the modules in the currently active group receives the Transfer Return signal from below and a Transfer Complete signal from the module to its left, it sends a Transfer Complete signal to the module to its right. At the Left Manifold Boundary the Transfer Return signal from below is sufficient to generate the Transfer Complete signal. At the Right Manifold Boundary the Transfer Complete signal is transmitted to the explicit control signal pathway to indicate completion of the operation.

### 2.3 DETAILS OF CONTROL AT THE MODULE LEVEL

In this section a block diagram of the control circuitry of each of the data-processing modules is presented, along with an explanation of each. The drawings are so arranged that when two (or more) are placed vertically or horizontally adjacent, like-named pathways coincide. The Data Delivery Return (D.D.R.) line and the Initiate Operation (In. Op.) line serve as vertical and horizontal alignment guides respectively.

Functional descriptions of the various circuit elements represented on the block diagrams have been given in the Introduction, along with a list of the abbreviations and their meanings (Table 1 on p. 15).



### 2.3.1 Storage Modules

#### Register Module

The Register is designed to store a twelve-bit data word, an associated flag bit, and a two-bit configuration code word. Its faceplate box provides data cable ports for one input and two output cables.

The twelve-bit data word may be obtained from the Down Bus (from data-changing modules above the Register) or from a module below the Register, and is delivered via the Up Bus and the two output data cable ports on the faceplate. The flag bit input to the Register is from above only, while the flag output is supplied to units both above and below via implicit data pathways.

The two-bit configuration code word is obtained from a module below the Register (a Configuration Gate) and it too is supplied to modules above and below the Register.

A four-bit function code word is obtained from the input data cable port on the faceplate and is distributed via implicit data pathways to units above and below the Register. This code word is not stored in flip-flops (as the data and configuration code words are), but is merely amplified and distributed.

Figure 19 represents the control circuitry in the Register. An explanation of this circuitry follows.

The Register may receive data as a result of three types of transfer, namely:

- 1) A twelve-bit data word from above (plus flag bit),
- 2) A twelve-bit data word from below, or
- 3) A two-bit configuration code word from below.

The latter two transfer types involve hardware which has never been built and are therefore never operationally employed. However, the functions have been exercised in a tester.

Consider first the transfer of a twelve-bit data word from above or from below. Each of these data transfers is accompanied by a Transfer signal from the same direction as the data. The Transfer signal is passed by the appropriate boundary gate (since if a data transfer is requested, that border of the module is not a transfer boundary). The boundary gates are located near the top center and lower left in Figure 19.

The transfer boundary gates supply signals to a cluster of circuit elements, ( $F_1$ , C.P. 1, C.P. 2, and a Call element) near the top center of Figure 19.

The  $F_1$  element asserts the conditioning input to the appropriate C.P. elements and connects the data storage element input circuitry to the



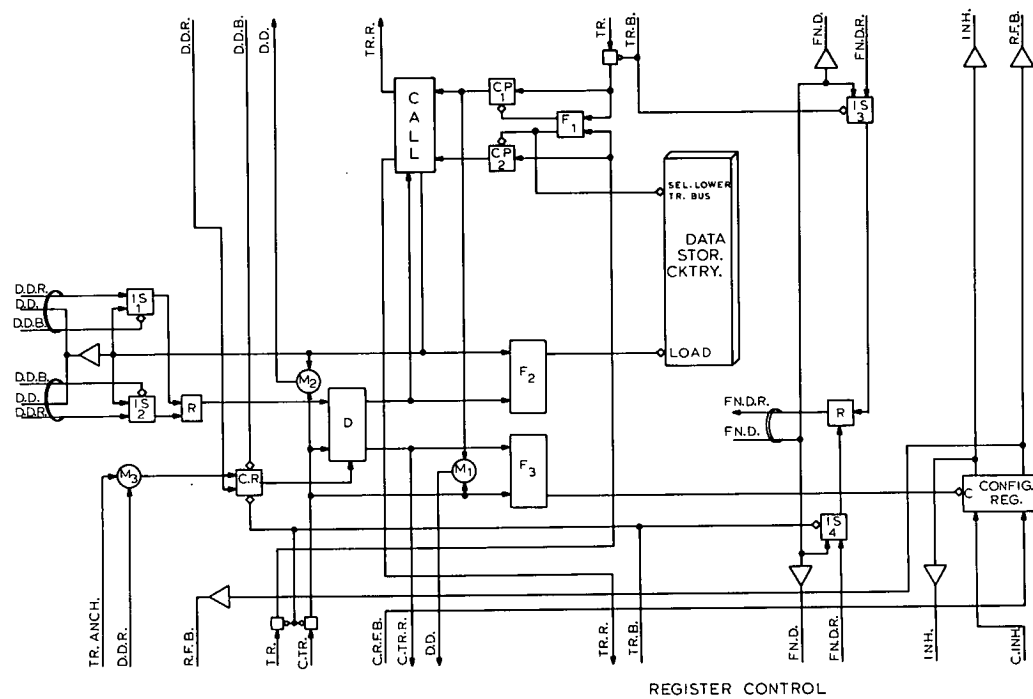


Figure 19. Register module control diagram.



appropriate implicit data pathway. (e.g., in response to a transfer from above, the Down Bus is connected to the flip-flop input circuits.) The C.P. elements delay the Transfer signal until element  $F_1$  has the appropriate output value. C.P. 1 (which responds to a Transfer signal from above) transmits a Data Delivery signal downward via  $M_1$ .

The "Do" output of the Call element provides Data Delivery signals to the output data cable ports (through an amplifier) and to the Up Bus via  $M_2$ . The same output sets  $F_2$ , which turns on the clock inputs to the flip-flops, enabling the data from the selected bus to be stored.

The Data Delivery signals which have been transmitted will (under normal circumstances) be followed by Data Delivery Return signals indicating that the new data values have had sufficient time to arrive at any destination.

The two Input Selectors and Rendezvous element located at the left center of Figure 19 combine the Return signals from the two data cables. The boundary input (D.D.B.) to either of these Input Selectors is asserted if no cable is connected to the data cable port. In this case, the Data Delivery signal is merely turned around and substituted for the Return signal.

The Conditional Rendezvous element (lower left in Figure 19) combines the Return signals from above and below. If the lower edge of the Register is a transfer boundary, the Return signal from above is simply passed by the C.R. element. If the top edge is a transfer boundary, the Return signal from below (or the Transfer Anchor signal, which will be discussed later) is selected. If neither boundary signal is asserted, the output of the Conditional Rendezvous is made active by the Return signal from above and the output of  $M_2$ . The case in which both boundary signals are asserted is a "don't care" condition because it implies an isolated Register.

The composite Return signals created by the Rendezvous and Conditional Rendezvous elements provide inputs to the Decision element such that a signal is generated at its upper output. This signal clears  $F_2$ , thereby turning off the clock signal to the storage elements, and also goes to the "return" input to the Call element. The Call element in turn transmits a Transfer Return signal in a direction consistent with the source of the original Transfer signal.

The response to a Configuration Transfer signal is similar to the above. The Configuration Register clock (lower right) is controlled by  $F_3$ , which is set by the Configuration Transfer signal and cleared by a composite Return signal. Data Delivery signals are transmitted in both vertical directions (via  $M_1$  and  $M_2$ ) but not to the data cables (since configuration data is not transmitted in the data cables). The arrival of all applicable Return signals activates the lower output of the D element, which clears  $F_3$  and provides the Configuration Transfer Return signal.

The preceding paragraphs may seem confusing. Therefore, a brief summary should be worthwhile. A Data Delivery signal is transmitted upward (produced by  $M_2$ ) upon receipt of Transfer signals from either above



or below or in response to a Configuration signal. This is consistent with the distribution of data associated with these signals. Data Delivery signals are provided to the data cable ports in response to Transfers from above or below (but not Configuration Transfer).  $M_1$  transmits a Data Delivery signal downward as a result of Configuration Transfer or Transfer from above.

The fact that the transfer of a twelve-bit data word (no flag bit) from below the Register does not produce a downward-going Data Delivery signal (and the resulting Return signal) has resulted in the creation of the Transfer Anchor signal, which serves as a timing aid.

The data transferred to a Register may have come from a module several positions below. The data will then have been amplified in each intervening module. The problem was to devise a scheme to assure that the current data values are gated onto the Up Bus without wasting time.

The module from whence the data transfer originates must transmit two control signals upward, namely a Transfer signal which is passed through each intervening module as quickly as possible, and a Transfer Anchor signal which is delayed in each module by an amount slightly greater than the worst-case delay for data.

In the Register the following sequence of events takes place.

- 1) The Transfer signal and new data values arrive (the order of arrival is not known),
- 2) The Transfer signal turns on the clock input to the flip-flops in the Register so that the new data may be gated onto the Up Bus, and
- 3) The arrival of the Transfer Anchor signal guarantees that new data have arrived.

Thus, the Transfer Anchor signal may be used in lieu of the Data Delivery Return signals when data words are transferred into the Register from below.

A change in Function Code information appearing at the input data cable port is accompanied by a transition signal on the Function Delivery line. This signal is transmitted (via isolation amplifiers) to units above and below the Register. Input Selectors (No's 3 and 4) and a Rendezvous element produce a composite Function Delivery Return signal which is transmitted to the input data cable port. As in the case of the Data Delivery Return signal, the Function Delivery signal is selected in place of a Return signal at a transfer boundary.

#### General Memory Controller Module

The General Memory Controller (GMC) module, which is 2 cell spaces high, is designed to provide access to a stack of Memory modules, or to any stack of storage access modules that may be of variable sizes. It can perform a number of operations initiated from a cable or a Function Call module, or do a "write" or "insert" operation initiated by a data-changing module in the stack above; it cannot be initiated by means of regular control cables.



The GMC module accesses the Memory modules stacked below it (which must be contiguous) from the top module downwards; an 18-bit address is used. The module contains its own 256-word high-speed integrated-circuit (IC) memory, which may be selected as the first 400<sub>8</sub> addressable locations. If only the IC memory is needed, the GMC module may be used alone, with no Memory modules under it. Also included is a twelve-bit storage buffer, which permits the internal movement of stored data without data delivery. The module has no overlay; four code bits may be supplied through the initiating function-call cable to specify various operations. Four toggle switches on the faceplate box specify additional options, and two lights indicate special error conditions.

The operations of this module may be extended laterally. As in other extendable modules, the control signals and function code are passed down the row from the rightmost module -- here, the eighteen-bit address is also passed. Note that the function-call cable terminals (ADDRESS and CODE) and the address data ports ( $A_L$  and  $A_R$ ) should never be connected on extender modules.

A twelve-bit data word to be stored may be obtained from the Down Bus (from data-changing modules above the GMC) or from the input data cable port on the faceplate. There is no flag bit or configuration code word.

A four-bit function code word is obtained from the input data cable port (Mode) on the faceplate and is distributed via implicit data pathways to units above the GMC. This code word is not stored in memory elements, but is merely amplified and distributed.

The address to the GMC is an eighteen-bit word. It is received by the rightmost module via two cable inputs on the faceplate, and passed to the left via implicit lateral pathways, to be used by an extended GMC as its address. The Data Delivery signal on these cables is also passed to the left. At a left boundary the Data Delivery signal is turned around and returned to the right as the Data Delivery Return signal. At a right boundary this signal is passed to the address cable.

A list of GMC functions and associated codes, with an explanation of each, may be found in Part I, Volume II, Sec. 3.12 of this report. For convenience, the function names are repeated here in Table 5.

A block diagram of the sequence for processing control in the GMC is given in Figure 20. For a description of the symbols employed, refer to Table 1 on p. 15. The block diagram does not indicate the methods of implementation, but rather the sequence in which control is processed. Vertical and horizontal signal pathways for the GMC are given in Tables 6 and 7.

Following is a list of the major signal groupings and the symbols used to represent them:

F	A four-bit operand on the Mode Bus (above GMC)
Z	A twelve-bit operand on the Down Bus (above GMC)
R	A twelve-bit operand on the Up Bus (above GMC)



Table 5. Operations of the General Memory Controller Module

<u>Function Code*</u>	<u>Operation</u>
1	Clear**
2	Read
3	Read and Clear**
4	Write
5	Insert**
6	Exchange
7	Read Modify**
12	Read to Buffer
13	Read and Clear to Buffer**
14	Write to Buffer
15	Write and Clear to Buffer**

---

\* Codes 0, 10, 11, 16, and 17 are invalid; their use will result in an error halt.

\*\* These operations take only one-half of a normal memory cycle.



Table 6. Vertical Signal Pathways in the GMC Module

<u>PIN</u>	<u>DIRECTION BELOW</u>	<u>ABOVE</u>	<u>BELOW GMC</u>	<u>ABOVE GMC</u>
2	Both	Up	Not Used	R <sub>3</sub>
4, 6 8	Down	Up	A <sub>12</sub> , A <sub>13</sub> , A <sub>14</sub>	R <sub>2</sub> , R <sub>1</sub> , R <sub>0</sub>
10, 12, 14	Down	Down	A <sub>15</sub> , A <sub>16</sub> , A <sub>17</sub>	Z <sub>0</sub> , Z <sub>1</sub> , Z <sub>2</sub>
16	Both	Down	Not Used	Z <sub>3</sub>
18, 20	Both	Up	Not Used	R <sub>7</sub> , R <sub>6</sub>
22	Down	Up	Initiate	R <sub>5</sub>
24	Up	Up	No Operation	R <sub>4</sub>
26	Up	Down	OK Operation	Z <sub>4</sub>
28	Both	Down	Not Used	Z <sub>5</sub>
30, 32	Down	Down	C <sub>0</sub> , C <sub>1</sub>	Z <sub>6</sub> , Z <sub>7</sub>
34	Down	Up	C <sub>2</sub>	R <sub>11</sub>
36	Both	Up	Not Used	R <sub>10</sub>
38, 40	Both	Up	D <sub>11</sub> , D <sub>10</sub>	R <sub>9</sub> , R <sub>8</sub>
42, 44, 46, 48	Both	Down	D <sub>9</sub> , D <sub>8</sub> , D <sub>7</sub> , D <sub>6</sub>	Z <sub>8</sub> , Z <sub>9</sub> , Z <sub>10</sub> , Z <sub>11</sub>
50, 52, 54, 56	Both	Both	D <sub>5</sub> , D <sub>4</sub> , D <sub>3</sub> , D <sub>2</sub>	Not Used
58	Both	Down	D <sub>1</sub>	Transfer
60	Both	Down	D <sub>0</sub>	D.D.R.
62	Both	Down	Not Used	TR. B.
64	Up	Down	Bottom Most	D.D.B.
66	Down	Down	A <sub>0</sub>	FN. D.R.
68, 70, 72	Down	Both	A <sub>1</sub> , A <sub>2</sub> , A <sub>3</sub>	Not Used
74	Down	Up	A <sub>14</sub>	FN. D.
76	Down	Up	A <sub>5</sub>	D.D.
78	Down	Both	A <sub>6</sub>	Not Used
80, 82, 84, 86	Down	Up	A <sub>7</sub> , A <sub>8</sub> , A <sub>9</sub> , A <sub>10</sub>	F <sub>0</sub> , F <sub>1</sub> , F <sub>2</sub> , F <sub>3</sub>
88	Down	Both	A <sub>11</sub>	Not Used
90	Both	Both	Not Used	Not Used

Note: All odd-numbered pins are grounded. Pin numbers refer to CSL wiring lists.



Table 7. Horizontal Signal Pathways in the GMC Module

<u>PIN</u>	<u>DIRECTION</u>	<u>TOP GMC</u>	<u>DIRECTION</u>	<u>BOTTOM GMC</u>
1, 89	←	Address Right D.D.	←	A <sub>2</sub>
3, 87	←	Address Left D.D.	←	A <sub>1</sub>
5, 85	→	Address Right D.D.R.	←	A <sub>0</sub>
7, 83	→	Address Left D.D.R.	←	A <sub>8</sub>
9, 81	←	A <sub>5</sub>	←	A <sub>7</sub>
11, 79	←	A <sub>4</sub>	←	A <sub>6</sub>
13, 77	→	L.M.B.	←	A <sub>14</sub>
15, 75	←	A <sub>3</sub>	←	A <sub>15</sub>
17, 73	←	A <sub>11</sub>	←	A <sub>12</sub>
19, 71	←	A <sub>10</sub>	→	Condition
21, 69	←	A <sub>9</sub>	→	Start/Done
23, 67	←	Initiate	←	C <sub>3</sub>
25, 65	←	A <sub>17</sub>	←	C <sub>2</sub>
27, 63	←	A <sub>16</sub>	←	C <sub>1</sub>
29, 61	←	A <sub>15</sub>	←	C <sub>0</sub>

The following are common to all module types:

31	Power Down Request
33	Protect (Data Shield)
35	Power Sense (Positive)
37	Power Sense (Negative)
39	Preset
41	Power Down Acknowledge
43	Spare
45	Ground
47, 49, 51	Power (Negative)
53, 55, 57	Power (Positive)
59	Ground

Note: Pin Numbers refer to CSL wiring lists. All even-numbered pins are grounded.







- A An eighteen-bit address on the Down Bus (below GMC)
- C A three-bit operation code (below GMC)
- D A twelve-bit operand on the Bi-directional Bus (below GMC)

The GMC may be initiated from above, from the faceplate, or (if extended) from the right. An initiate from above implies a write code, since no code lines accompany the transfer signal. There is no lateral communication of control when the GMC is initiated from above.

The writing from above commences with the receipt of a Transfer (TR.) signal. The Transfer signal is passed by the boundary gate (top center of drawing) at which time the code generation begins, and is enabled to the bus below.

Since the data from above may not arrive with the Transfer signal, and writing must not begin until the arrival of the Data Delivery Return (D.D.R.) signal, the first action is to call the Up Bus Data Delivery (D.D.) lines (CALL 1). Upon the return of the D.D.R. signal, data will have been delivered to the Up Bus and be available for storage.

The next sequence is to call circuitry capable of communicating with the storage element. Call element 2 begins this sequence and initiates O.S.6, which will test the function code for validity. Should the code be invalid, the control sequence will die, and a light on the faceplate will be enabled to signal the user. If the code is valid, the address location which has undergone decoding is then tested to determine the location of the storage elements (O.S.5). Should the requested address not be available within the GMC an Initiate Operation will be issued to storage elements below (left bottom of drawing). Accompanying the Initiate to Below will be a function code, the address location, and (if a "write" code) the data to be stored.

If the address requested below is available, a return will be received on the OK terminal. If not, the operation (N.OP.) line will be initiated from below. Notice that these signals can only be received when the boundary "Bottom Most" is not asserted. Also note that in the event an Initiate to Below (IN.T.B.) is generated and the GMC is "Bottom Most" it will generate a N.OP. return via M5.

M7 completes the sequence by directing internal completions or completions from below to Call element 2. A Transfer Return (TR.R.) to above will become eminent from C.P.2 provided that there was not a N.OP. return or that a switch on the faceplate places C.P.2 in a don't care state. Should the control die at C.P.2, a light on the faceplate signals the user.

The functions listed in Table 5 can be implemented when the GMC is initiated via the faceplate using a Function Call module. If the GMC is rightmost, it is initiated via M1 (right top). If it is not rightmost, the inhibit gate passes the initiate operation to M2. In either case, M2 passes the signal to the module to the left as its Initiate Operation signal.

The output of M2 is delayed and directed to call element 2 and C.P.1. If the switch on the faceplate is properly set and a write code is detected,



C.P.1 will gate the data to be stored to the Up Bus and data deliver. A completion from the storage element via call 2 is tested for a write code (O.S.3) and direction to C.R.2 to be rendezvoused with the Data Delivery Return from call 1. Since the code may be "write" and the faceplate switch may order no data delivery, there may be no Data Delivery Return, and the signal from O.S.3 is passed to M4 via C.R.2.

M4 passes the signal to C.R.1 via M3 to be rendezvoused with a completion from the module to the left. Note that the signal from M3 is passed through C.R.1 if this module is a left boundary. The output of C.R.1 generates the completion to the faceplate (O.S.1) and signals the module to the right a completion. Note that there is no danger in giving a completion to the faceplate in modules that are not Right Boundaries, since only the Right Boundary module will have a control cable in the faceplate.

A user may receive one of two completions from the faceplate - Yes, the address location requested was available, or No, it is not. The condition processing circuitry determines which output is to be issued, by testing the type return received from modules below and the type received in the modules to its left. Note that it is not concerned with the condition signal to the left when there is a left manifold boundary. Should any of the modules in the manifold receive a No Operation completion, a No completion is generated to the faceplate (O.S.1), and the condition level will be asserted to the right.

Let us assume the GMC was initiated from the faceplate, the function code is not a "write", and the data received from the storage element is to be data delivered. The output from O.S.3 will be directed to O.S.2 and then to R3.

R3 will rendezvous the request for data delivery with a Start signal. The Start signal indicates that all data parameters needed to complete the operation have been stored, and that data out from this module may safely be changed. Notice that a module generates its own Start signal if it is a left manifold boundary. The output of R3 generates Data Delivery signals for the cable data out and the Up Bus. The bus Data Delivery Return and the cable Data Delivery Return are rendezvoused at R1, and a completion is generated via M3, C.R.1, and O.S.1.

#### Memory Module

The successful utilization of the 4096 X 12 ferrite core, address-extendable Memory module requires the coexistence of a memory controlling module stacked immediately above the array to supply the parameters required by the memory. This controller dictates memory functions, provides addressing data, and is the sole interface between the memory and the other data-processing modules.



The controller must supply the following:

- a) 18 Address bits
- b) 12 Data bits\*
- c) 3 Code bits
- d) 1 Control transition (Initiate Operation)

The controller must accept from the memory:

- a) 12 Data bits\*
- b) 2 Control transitions
  - 1) Initiate Return
  - 2) No Operation Return

In an effort to minimize confusion of terms and to aid the reader in understanding the block diagram (Fig. 21a) with its explanation of control sequences, the following items have been selected for a brief study.

#### A. Bidirectional Data Bus

In order to reduce the interfacing connections between the memory and its controller, a bidirectional data bus is used. These 12 data pathways are switched to gate data out of memory during a read command and to accept incoming data in the event of a write sequence. For details of design parameters refer to [1].

#### B. Selection

Memory selection is achieved in the following manner. Upon receipt of the 18 address bits, each Memory of the manifold in turn decodes this address, and a decision is made as to whether it is the selected module. If not selected it subtracts, from the 18 bit address, the number equal to the number of locations available in its core and presents the remainder to the Memory below. If an address too large for the entire Memory manifold is detected, the bottom Memory in the stack will return any ensuing control signal to the controller above on a "No Operation" line.

#### C. Busy Level

When a Memory is selected, control immediately asserts a "Busy" level in that module. Any ensuing Memory request will be held up at the Memory until Busy is removed by the completion of the operation in progress.

#### D. Deselect

To facilitate the design of other data storage modules which might reside in a Memory manifold and utilize the 18 address bits on the Down Bus, a Deselect Code from the controlling module becomes necessary. Its necessity, designed to facilitate "Source Checking", is not at all obvious; the following explanation is offered.

---

\* The 12 data bits in and out of memory employ a bidirectional bus which allows the sharing of copper paths.



### Source Checking

In order to maintain the philosophy that all data used by a module be guaranteed valid when that module is initiated, a system-wide scheme known as "data delivery" is used. An address destined for use by a data storage module embedded in a stack of Memory modules can therefore be guaranteed valid only as far as the controller at the top of the stack. To ascertain the validity of the address at its inputs, the data-processing module must initiate a transition up the stack to the controlling module as a source check. When this signal arrives at the controller, it is returned down the stack to the requesting module. Its roundtrip propagation time is guaranteed to be equal to or longer than the one-way propagation time for the address; therefore, its return to the data-processing module signifies the validity of the address at its input.

Due to pin limitations of the vertical pathways, the "source check" is designed to use the same copper pathway already used by the Memory as the "Initiate Return" line. The controller senses the return and responds with a transition down the stack on the normal "Initiate" line. If the intervening Memory modules had not been deselected by the 3-bit function code and the address on the Down Bus inadvertently selected such a Memory, then that Memory would respond with an unwanted operation. This would result in a control loop between the unwanted Memory and the controlling module. To compound the problem, the "source check" response would never be returned to the initiating module.

In order to prevent the above problems, it is essential that the controlling module assert the deselect code at least 30 nanoseconds before returning a "source check" response.

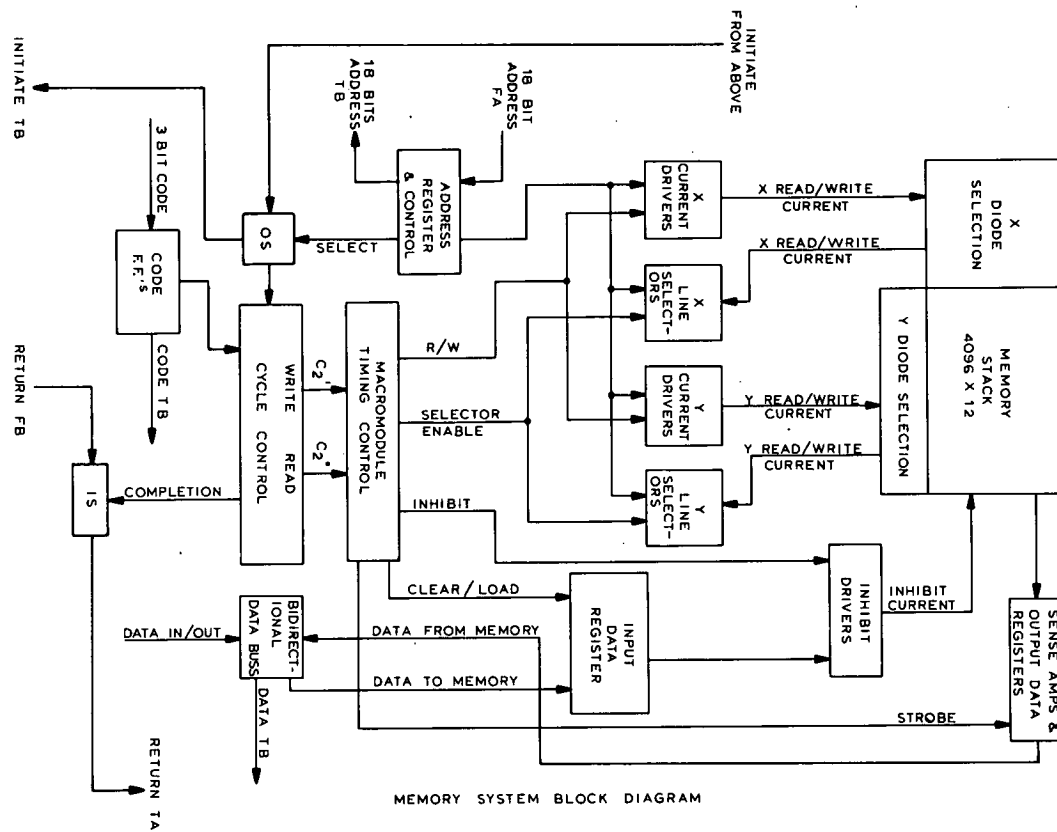
In the following discussion, reference is made to the Memory module control drawing (Fig. 21b).

Upon receipt of an initiate command (INI-FA) at the input of a Memory module, the control sequence branches into two parallel operations. Via M1, the Enable Bus flip flop is cleared or "pointed down" in anticipation of incoming data, and simultaneously, O.S.1 is interrogated. At this time the gating levels which control O.S.1 are asserted, and control either enters the module or is directed to below.

If control is directed to below, the signal is simply buffered and passed on as an Initiate to below (Init TB). Should the module be detected to be "Bottom-Most" by the Inhibit Gate (controlled by Bottom-Most level), then control is directed up toward the controller via M2 as a "No Operation" return. The other input of M2 is for passage of a "No Operation" that might have occurred below a given module. In addition to passing the "No Operation" signal to above, the output of M2 is directed via M5 to the input of M1. The output of M1 (not shown) gates the Bidirectional Bus to its normal state.

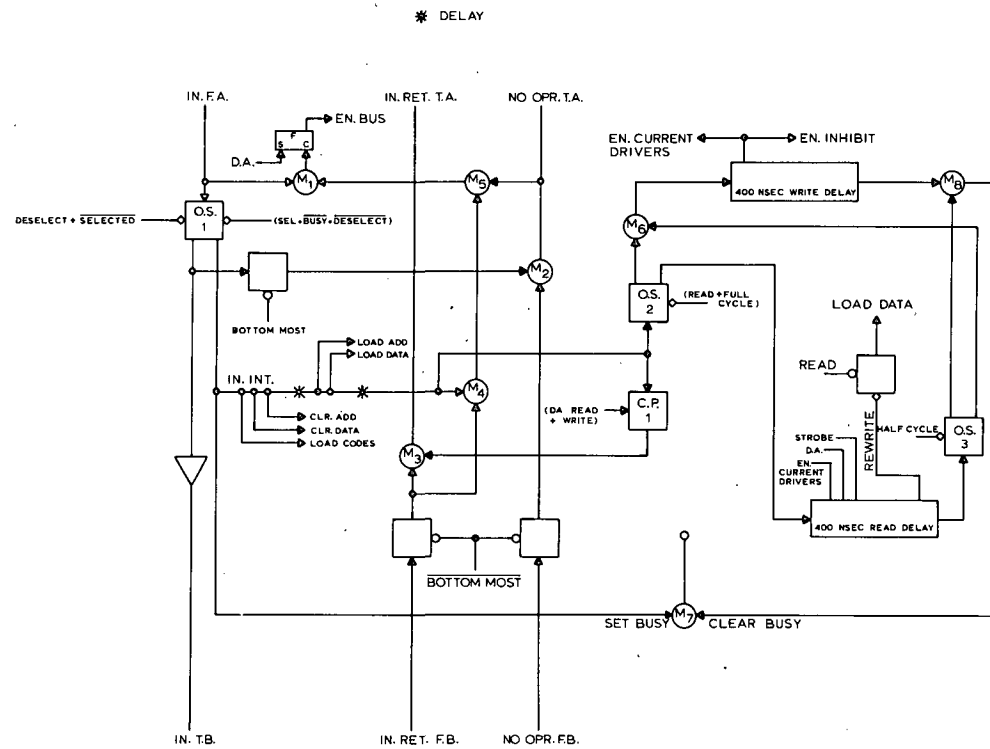
If O.S.1 is gated to direct the control sequence into a module, the memory operation starts. Once again control branches into two paths. The "busy" indicator is set via M7 to lock out any ensuing Initiate (at O.S.1) until the present operation is completed, and a delay is instigated





(a) Block diagram.





(b) Control diagram.

Figure 21. Memory module diagrams.



to allow for the generation and occurrence of the Clear Address, Clear Data, and Load Code pulses. When the delay expires, control then generates the pulses which load the new address (LOAD ADD) and data (LOAD DATA). Control is also delayed at this point to allow time for this event to occur. At the conclusion of this delay, the memory function has been decoded and all necessary parameters have been loaded.

Control branches at this point into three paths as inputs to M4, C.P.1, and O.S.2.

Path 1: Control passes through the input of M4 and ultimately, via M5 and M1, switches the Bidirectional Data Bus to receive data out from memory for transfer up toward the controlling module.

Path 2: The gating of C.P.1 is controlled strictly by the memory function in operation. If the "write" code is asserted, C.P.1 immediately responds with an output which is directed to M3. The output of M3 is the completion signal to above (Init Ret TA). The completion signal in this case does not imply that the write operation has been completed. It does, however, imply that the operation is in progress and will be completed in a known time. Remember that any ensuing Initiate to this module is locked out until the "busy" level is removed at the actual completion of the write operation.

If the function had been a read code, control would have waited at C.P.1 until the assertion of the Data Available pulse (DA) which occurs during the read portion of the cycle. Upon receipt of this pulse, C.P.1 fires the output which (via M3) signals completion to above.

Path 3: Once again the function code directs the two possible paths which are outputs of O.S.2. If the gating level asserted by Read or Full Cycle code is present, control is directed to the input of the 400-nsec lumped constant "read" delay. Otherwise control is passed via M6 to the input of the "write" delay. Consider this possibility first.

Once control enters the write delay, the drive and inhibit currents are switched on which govern the data written or stored in core. The conclusion of this delay turns the driver off (not shown) and then passes through M8, whose output clears the "busy" level. At this time the operation is actually completed.

If control is directed into the read delay from O.S.2, the current drivers are switched on to allow the strobing of data out of core. It is at this time that the DA pulse is generated, which signifies that new data are eminent and must be gated onto the bidirectional bus. Note that DA sets the flip flop F which enables the data bus. Near the end of the read delay, it is necessary to interrogate the Inhibit Gate controlled by the "read" level. If read is asserted, it is necessary to produce a Load Data pulse which transfers the output data register into the input data register via the bidirectional bus (not shown).

The output of the read delay is next delivered to the input of O.S.3. If the operation is determined to be "half cycle", then the control is



simply presented via M8 to the input of M7 to clear the Busy level. Otherwise, control is directed to the write delay via M6, which completes the full cycle operation. The sequence of events that occur then are exactly as before.

The remaining circuit elements for discussion are provided to direct control back through a Memory module when a successful operation has been completed in another Memory below it. The "Initiate Return" from below (Init Ret FB) is gated through the Inhibit Gate governed by Bottom-Most and directed to inputs of M3 and M4. The output of M3 supplies the Initiate Return to above, while the output of M4 prepares the Bidirectional Data Bus to receive data from below (via M5 and M1).

In order to allow sufficient time for Memory selection and function decoding, it is essential that the address and data precede control to the Memory by a fixed time. This time is based upon a worst-case time study which reflects the slowest propagation rate of data versus the fastest propagation rate of control.

The times are:

- a) Address data precede control by 63 nsec
- b) Code data precede control by 30 nsec

To allow for the time required to switch the Bidirectional Bus to gate data into or through the memory, the control input must precede incoming data to be stored by at least 45 nanoseconds.

After receipt of a completion return by the controller from the memory, the controller must delay for 51 nanoseconds before the data on the Bidirectional Bus can be considered valid.

### 2.3.2 Data-Changing Modules

Each of the modules in this group communicates with a storage module via the Down Bus. Each may be extended laterally, and most of them take data from the Up Bus.

A data-changing module behaves as though it adjoins a vertical manifold boundary, in the absence of a signal to the contrary from above. It also sends a voltage level to the module below it indicating the absence of a vertical (Data Delivery and Transfer) boundary between them.

Left and Right manifold boundaries refer to the horizontal extremities of a data-processing manifold. As mentioned previously, the Right Manifold Boundary is established by positioning a switch on the faceplate. Each module is assumed to occupy a Left Manifold Boundary position in the absence of a signal from a module to its left.

Data fields (each having a width equal to an integer multiple of twelve bits) may be established within a manifold. It is therefore necessary to define Left and Right Field Boundaries.



A Right Field Boundary may be established at the right side of a module by asserting the RFB signal on the Configuration Bus, or by setting the faceplate function code switches of a module to some value other than 17 (octal), or by asserting the Right Manifold Boundary condition. When such a Right Field Boundary is established, a signal is transmitted to the module to the right indicating that it occupies a Left Field Boundary position.

Assertion of the Inhibit signal within a module indicates that no change in the data in that column is to appear on the Down Bus. The source of the Inhibit signal (from below or from the right) is controlled by a locally produced signal identified as Right Field Boundary (RFB). The source of the Inhibit signal is the Configuration Bus if RFB is asserted; otherwise the Inhibit signal comes from the module to the right.

Each of the data-changing modules must generate the Transfer signal for its column (when it is active) and respond to the Transfer Return signal. It must also pass these signals when inactive, and pass the Data Delivery and Return signals whether active or passive. The details of these and other aspects of control are presented in subsections relating to specific modules.

#### Load Module

The Load module accepts data from an input cable and supplies this same data to an output data cable port (branch). If the least significant bit of the function code is zero, the Load module supplies twelve zero bits to the Down Bus (when the module receives an Initiate Operation signal); otherwise it transfers the data from the input cable to the Down Bus. The flag value from the Up Bus is transferred to the flag position on the Down Bus. Figure 22, which represents the control circuitry of the Load module, may be a helpful reference when reading the following paragraphs.

If the Load module is rightmost, it is called to action by an Initiate signal applied to the Call element (right center). Otherwise the signal RMB is not asserted, and an Initiate Operation signal is coupled through the boundary gate (also right center). In either case  $M_1$  passes the signal to the module to the left (through an isolation amplifier) as its Initiate Operation signal.

The output signal from  $M_1$  (appropriately delayed) is conditionally rendezvoused with the Result Ready signal from the right. The output of C.R.1 serves as the local Result Ready signal and is also passed to the module to the left. Note that if the Right Manifold Boundary signal is asserted, the delayed Initiate signal becomes the local Result Ready signal; otherwise the Result Ready signal from the module to the right is also required.

The local Result Ready signal sets the F element near the center of Figure 22. This gates the appropriate data onto the Down Bus. The output of C.R.1 is also transmitted to the module to the left, and is also used as one of the input signals to circuit elements I.S.1 and C.R.2.

Each of the latter two circuit elements has a level input signal which is asserted when the module occupies a Left Manifold Boundary position. At



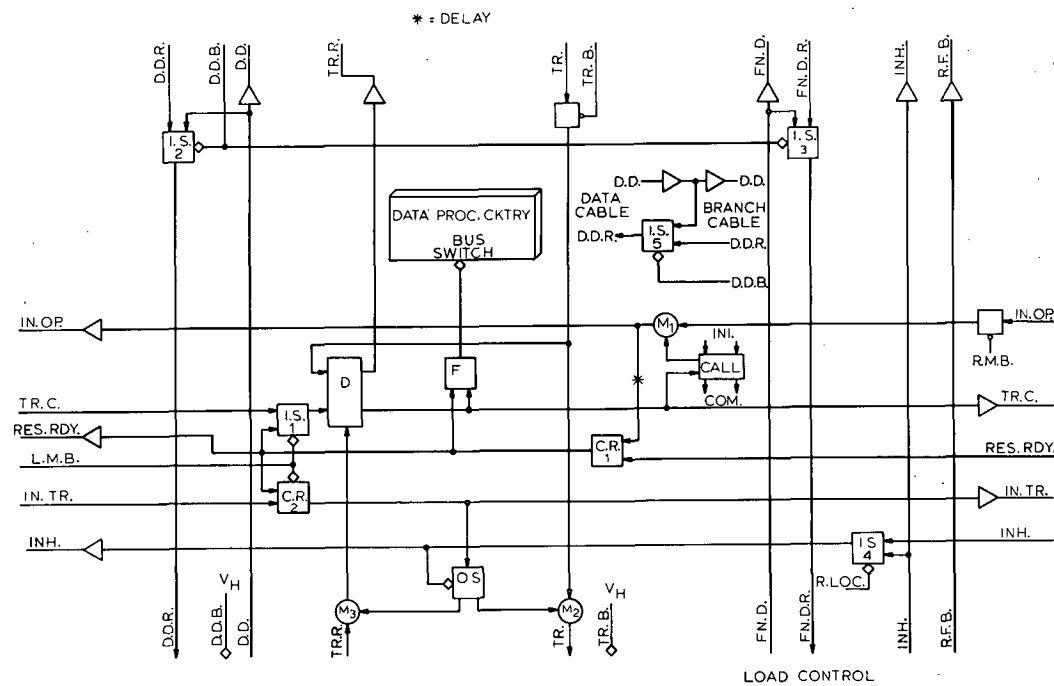


Figure 22. Load module control diagram.



this boundary (L.M.B. signal is asserted) the local Result Ready signal provides one of the inputs to the Decision element (through I.S.1) and creates the Initiate Transfer signal (C.R.2). Note that when the module occupies any position other than leftmost, the Initiate Transfer signal from the left is rendezvoused with the local Result Ready signal to form the local Initiate Transfer signal, and the Transfer Complete signal from the left provides the D element input signal.

The local Initiate Transfer signal is transmitted to the right through an isolation amplifier. It is also directed to Merge element  $M_2$  or  $M_3$  by the Output Selector element (bottom center). The latter element is controlled by a local Inhibit signal. Note that the signals labelled Inhibit are distributed both vertically and horizontally. An Input Selector (I.S.4 in the lower right corner) passes one of these as the local Inhibit signal. Thus if the signal RFB is asserted, the source of the local Inhibit signal is the vertical pathway. As stated previously, RFB is asserted if the module occupies either a Right Manifold or Right Field Boundary position, or if the faceplate function code switches are set at any number other than 17.

If the local inhibit condition is not asserted, the Initiate Transfer signal will be directed (through the Output Selector and  $M_2$ ) to the module below as the Transfer signal. When the Transfer signal arrives at the storage module at the base of the column, it will precipitate a chain of events which will eventually produce a Transfer Return signal at  $M_3$ . Note that had the local Inhibit signal been asserted, the Output Selector would have directed the local Inhibit Transfer signal to  $M_3$  without the Transfer, Data Delivery, Data Delivery Return, Transfer Return signal sequence.

The output signal from  $M_3$  completes the set of input signals required to produce an output signal from the D element. Recall that one input was provided via I.S.1. This output signal clears the F element (turns off the Bus Switch), provides the Return signal to the Call element, and provides a Transfer Complete signal for transmission to the right. At the Right Manifold Boundary, the completion signal produced by the Call element will be available at a control terminal.

The vertical Data Delivery signal is amplified in the Load Module, and either this signal or the Data Delivery Return signal from above is transmitted downward as the Data Delivery Return signal. The Delivery Boundary signal (D.D.B.) assertion causes the turn-around at the boundary. The Function Delivery and Return signals are handled the same way.

Data Delivery signals on the input data cable or the Return signal from the data branch cable (as dictated by the absence or presence, respectively, of the data branch cable) are transmitted as Data Delivery Return signals on the input data cable.

When the Load Module is between the currently active module and the storage module, it will pass a Transfer signal from above through  $M_2$ . Note that this signal also supplies an input to the upper left terminal of the D element. The subsequent arrival of the Transfer Return signal will



produce an output at the corresponding (upper right) terminal of the D element, which is amplified and transmitted upwards.

Since the propagation of control signals in the Load module is representative of that in any of the data-processing modules, a brief summary will be given.

Three cases of interest exist. The module may be in the currently active row, between the currently active modules and the storage module, or above the currently active module. We shall consider these three cases individually.

When the module is in the currently active row, the following sequence of events occurs:

- 1) An Initiate Operation signal is received either from the right or from a terminal on the faceplate.
- 2) This signal is passed to the left, and also delayed to produce the Result Ready signal.
- 3) At the Left Manifold Boundary the local Result Ready signal prepares the D element to eventually issue a Transfer Complete signal.
- 4) The local Initiate Transfer signal sends a Transfer signal downward (or a Transfer Return substitute if the local Inhibit line is asserted) and an Initiate Transfer signal to the right.
- 5) The storage module at the base of the column transmits a Data Delivery signal, which is amplified and sent to modules above. The module at the Data Delivery Boundary will send back a Return signal, which is amplified and sent to the module below.
- 6) The storage module transmits the Transfer Return signal. This produces the Transfer Complete signal (provided the module is leftmost or has received the Transfer Complete signal from the module to its left).
- 7) The Bus Switch (which had been activated by the local Result Ready signal) is deactivated by the local Transfer Complete signal.

When the Load module is between the currently active module and the storage module at the base of the column, the module passes the Transfer signal from above, the Data Delivery signal from below, the Data Delivery Return signal from above, and the Transfer Return signal from below in the order listed.

Finally, when the Load module is above the currently active module it passes the Data Delivery signal, and either creates (at the boundary) or passes the associated Return signal.

In the discussion of the remaining data-processing modules, only those aspects of control which differ from the Load module control will be considered.

#### Logic Module

The Logic Module accepts data from the Up Bus and two input data cables and supplies updated data values to the Down Bus on command. Each bit of



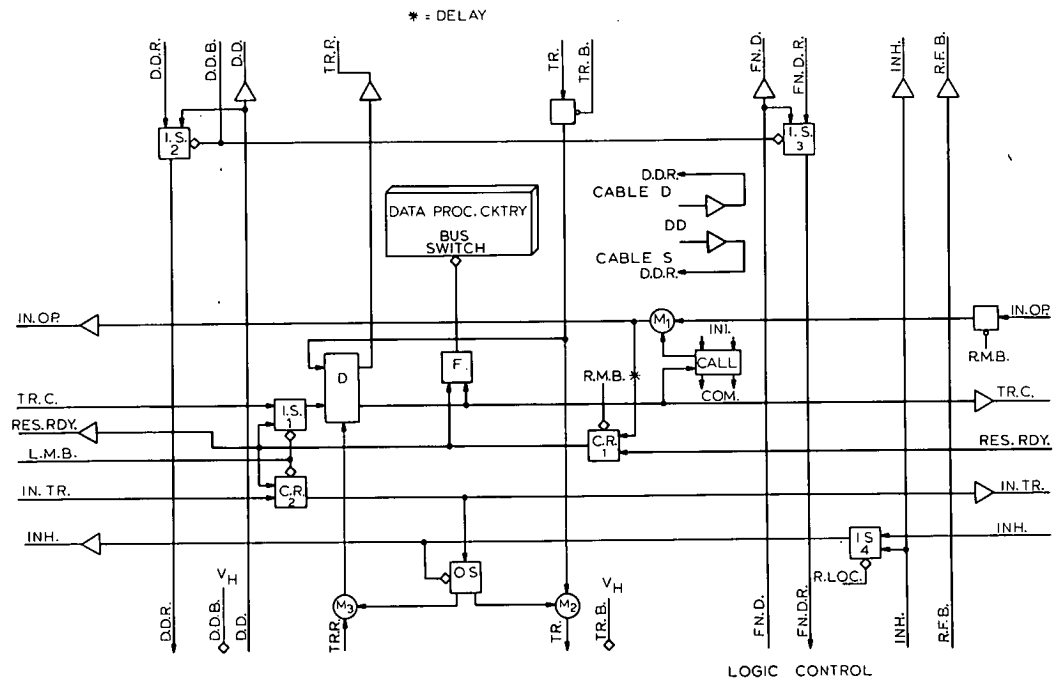


Figure 23. Logic module control diagram.



the output data is determined by the corresponding bits taken from the Up Bus ( $R_i$ ), the input operand ( $D_i$ ), and the selection variable input ( $S_i$ ). For each  $S_i = 0$ ,  $Z_i = R_i$ , where  $Z_i$  is the output bit  $i$  transferred by the Logic module. A summary (in terms of the function code) for the case in which  $S_i = 1$  is given in Table 8.

Table 8. Summary of Results for Logic Module

$Z_i$  values for  $S_i = 1$

$D_i$	$R_i$	0	1	2	3	4	5	6	7	10	11	12	13	14	15	16	17	Function Code
0	0	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1	0	
0	1	1	0	0	1	1	0	0	1	1	0	0	1	1	0	0	1	
1	0	0	1	0	1	1	0	1	0	0	1	0	1	1	0	1	0	
1	1	1	0	0	1	0	1	1	0	0	1	1	0	1	0	0	1	

Figure 23 depicts the control circuitry of the Logic module. The only difference (from a control standpoint) between it and the Load module is in the absence of an output data cable. The Data Delivery signals from each of the input cables are simply amplified and employed as Return signals in the two cables.

### Shift Module

The Shift module shifts data words to the right or left and provides several options concerning the vacated bit position (Most Significant Bit during right shift and Least Significant Bit during left shift). Table 9 provides the details of these options in terms of the function code. In the table,  $R_m$  and  $Z_m$  refer to the most significant bits on the Up and Down Busses respectively. That is,  $R_m$  refers to bit  $R_{11}$  in the leftmost column of an extended group.

Each of the bits  $Z_{m-1}$  through  $Z_0$  is shifted left one bit position for function codes 10, 12, or 16 (octal), and each of the bits  $Z_1$  through  $Z$  is shifted right one bit position for function codes 0 through 7.  $Z_0$  refers to the least significant bit of the entire word.

Control differs slightly (in the Shifter) from the modules previously discussed. Note that function code 12 describes a rotate left operation, in which the most significant bit in the leftmost column is shifted into the least significant bit in the rightmost column. Since no direct pathway exists, this bit must be shifted through each intermediate module. For this reason the Result Ready signal is originated in the leftmost module (as opposed to the rightmost in the modules discussed previously).



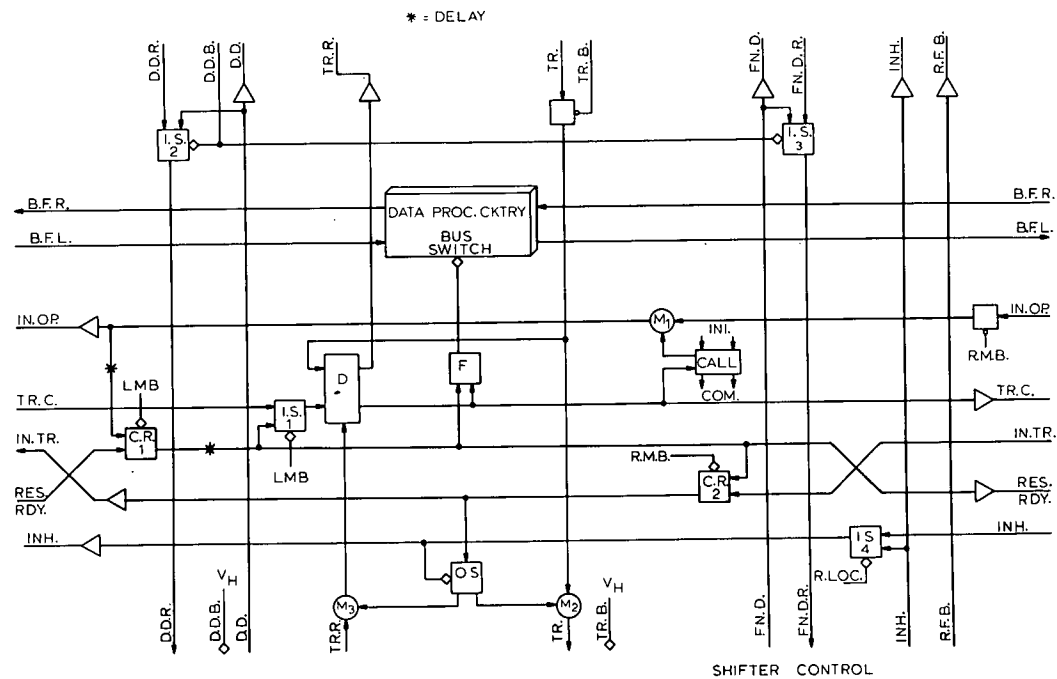


Figure 24. Shift module control diagram.



At the Right Manifold Boundary the local Result Ready signal produces the Initiate Transfer signal (via C.R.2 in Figure 24 ), which propagates from right to left, performing the same function as in the Load and Logic modules. Note that the Initiate Transfer and Result Ready lines are interchanged at each side of the Shifter module.

The two horizontal pathways labelled B.F.R. and B.F.L. carry the bit from the right and the bit from the left when shifting across column boundaries in an extended group of modules.

Aside from the above exceptions and the fact that the Shift module has no data cable connections (and therefore no Data Delivery and Return signals), the control circuitry shown in Figure 24 is identical to that which has been previously described.

Table 9. Values of Flag and Vacated Bit for Shift Module

<u>Funct. Code</u>	<u>Z<sub>f</sub></u>	<u>Z<sub>m</sub></u>
0	R <sub>f</sub>	R <sub>m</sub>
1	R <sub>0</sub> · (R <sub>m</sub> + R <sub>f</sub> )	R <sub>m</sub>
2	R <sub>f</sub>	1
3	R · (R <sub>m</sub> ⊕ R <sub>f</sub> )	R <sub>m</sub> ⊕ R <sub>f</sub>
4	R <sub>f</sub>	0
5	R <sub>f</sub> + R <sub>0</sub> · R <sub>m</sub>	R <sub>m</sub>
6	R <sub>f</sub>	R <sub>0</sub>
7	R · (R <sub>m</sub> ⊕ R <sub>f</sub> )	R <sub>m</sub> ⊕ R <sub>f</sub>

<u>Funct. Code</u>	<u>Z<sub>f</sub></u>	<u>Z<sub>0</sub></u>
10	R	1
12	R	R <sub>m</sub>
16	R	0



# Arithmetic Module

The Arithmetic (Addition) module accepts data from the Up Bus (R) and an input data cable (D), and provides the two's-complement, integer functions of these as listed in Table 10 .

Table 10. Arithmetic Module Output Data

<u>Function Code</u>	<u>Z</u>	<u>Z<sub>f</sub></u>
0	$R + D$	$Z_f = 1$ indicates overflow *
1	$R + 1$	$Z_f = 1$ indicates overflow *
2	$R - D$	$Z_f = 1$ indicates overflow *
3	$R - 1$	$Z_f = 1$ indicates overflow *
4	$D - R$	$Z_f = 1$ indicates overflow *
5	$0 - R$	$Z_f = 1$ indicates overflow *
6	R	1
7	R	$R_m$ (Sign Bit)
10	$R + D$	$R_f$
11	$R + 1$	$R_f$
12	$R - D$	$R_f$
13	$R - 1$	$R_f$
14	$D - R$	$R_f$
15	$0 - R$	$R_f$
16	R	0

---

\* Or underflow.







The adder packages employed in the data-processing circuitry of the Arithmetic module require both carry and no-carry inputs. Naturally these two signals require lateral pathways when Arithmetic modules are extended to increase the word length. Each of these lines is driven low at Preset and at the completion of an addition operation. The presence of a high signal on either of these lines replaces the transitional Result Ready signal. Figure 25 depicts the control circuitry of this module.

Since the presence of a high level on the  $C = 1$  line indicates that the carry from the data column to the right has the value 1, and the presence of a high level on the  $C = 0$  line indicates that the carry from the right has the value 0, the presence of a high level on either of these lines (recall that they were previously driven low) indicates that the module to the right has completed its data-processing operation.

The local Initiate Operation signal (produced by  $M_1$ ) is delayed and applied to a Conditioned Pause element. The level input to the C.P. is the output of a gate which forms the logical OR of the  $C = 0$ ,  $C = 1$ , and Right Manifold Boundary signals. Thus, if the module is rightmost, the delayed local Initiate Operation signal serves as the local Result Ready signal. Otherwise, either the  $C = 0$  or  $C = 1$  input signal must be driven high to satisfy the condition on the C.P. element. Since assertion of the  $C = 1$  line implies a carry input signal to the data-processing circuitry, and since this signal will change the value of the output data, the signal on the  $C = 1$  line is delayed to guarantee accuracy of the output data when the local Result Ready signal is asserted.

The local Result Ready signal enables the bus switch and provides one of the inputs to the special-purpose circuit block labelled "carry circuit" in Figure 25. The other inputs to this block are the carry output from the most significant adder and the Left Field Boundary signal.

Since there is no carry across a Field Boundary, both the  $C = 0$  and  $C = 1$  signals are negated by the assertion of the LFB signal. When the module does not occupy a Left Field Boundary position, the carry circuitry will drive either the  $C = 0$  or the  $C = 1$  signal high following receipt of the local Result Ready Signal.

At the Left Manifold Boundary the local Result Ready signal produces the Initiate Transfer signal, which causes the same sequence of signals as in the modules previously discussed. When the Transfer Return signal activates the D element, which in turn clears the F element, the carry circuit block once again negates the signals on both the  $C = 0$  and  $C = 1$  lines.

The Data Delivery signal in the input data cable is simply turned around and used as the Return signal.

### Multiply Module

The Multiply module forms the product of two twelve-bit two's complement operands. If the least significant bit of the function code is asserted, the operands are treated as binary fractions. Otherwise, the operands are treated as integers.



This module receives twelve bits of the multiplier (MP) on an input data cable, and twelve bits of the multiplicand (MC) from the storage module (at the base of the column) via the Up Bus. A double-length result is generated, of which half is made available on the bus and the other half on the module's data cable outputs.

The product registers within the Multiply module are double ranks of flip-flops (identified as Upper and Lower ranks). Within each module, the twelve most significant bits of the product are referred to as the Left Half Product (LHP), while the least significant twelve bits are called the Right Half Product (RHP).

When fractional multiplication is employed, the LHP is delivered to the Down Bus and the RHP to the output data cable port. For integer multiplication, the destinations of the two product halves are interchanged.

Since the Multiply module is double height, it spans two lateral communications channels, thereby providing thirty horizontal signal pathways as opposed to fifteen in the single-height modules discussed previously. For this reason, Table 3 on p. 20 is not applicable to the Multiply module. Table 11 compares the upper and lower horizontal channels with that of the Load module.

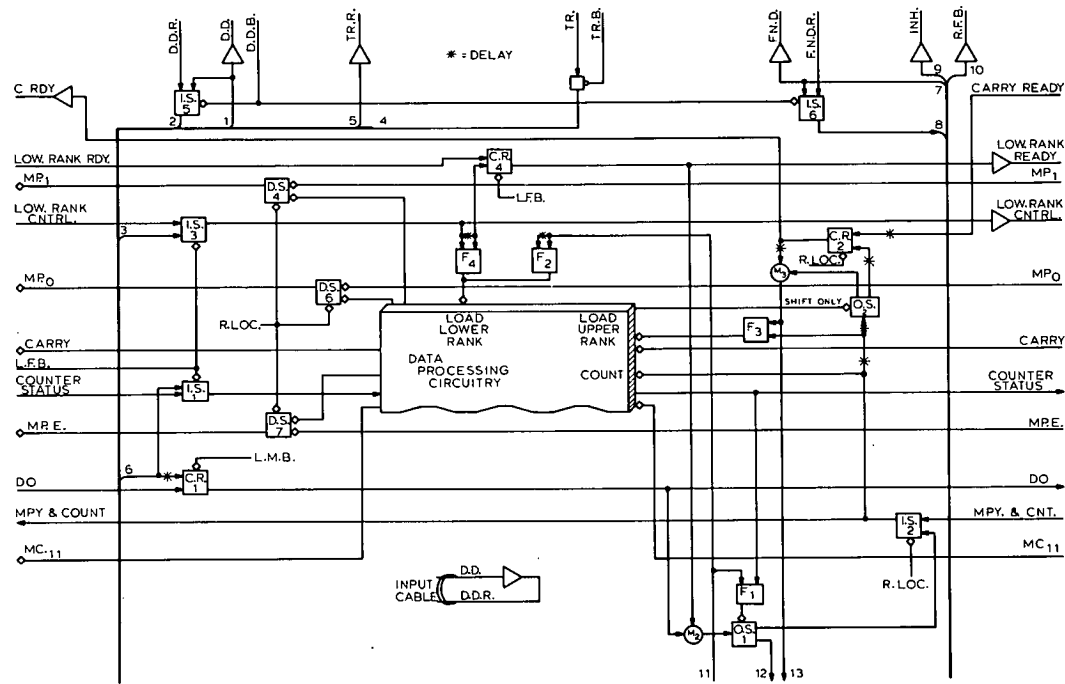
Figure 26 depicts the control circuitry of the Multiply module. The location of circuit elements on the two figures is not intended to infer similar physical locations within the module. However, the signal pathways shown on the upper drawing (Figure 26a) are those from the upper lateral channel. An analogous statement applies to the lower drawing and associated lateral channel.

Vertical module boundaries are represented at the top of Figure 26a and at the bottom of Figure 26b. The intra-module vertical wiring has been condensed into two trunks in the Figures. Signal pathways entering and leaving the trunk are labelled. The three individual lines crossing this boundary are also labelled.

The actual multiplication of the two operands is performed by implementing a modification of Booth's algorithm for the multiplication of signed two's-complement numbers. The multiplier is examined three bits at a time, beginning with the least significant end. Actually, a one-bit overlap exists between adjacent groups of three bits, necessitating the examination of six groups if the multiplier has twelve bits. Table 12 indicates the six groups which are formed from a twelve-bit multiplier. MPE is an extension of the multiplier (least significant end) and is initially zeroed.

Each group of three bits is decoded to determine the current contribution to the partial product. Table 13 lists the results of the decoding operation. The shifting referred to therein is applicable to the multiplier and the partial product. The multiplicand remains in the register throughout the entire multiply operation.







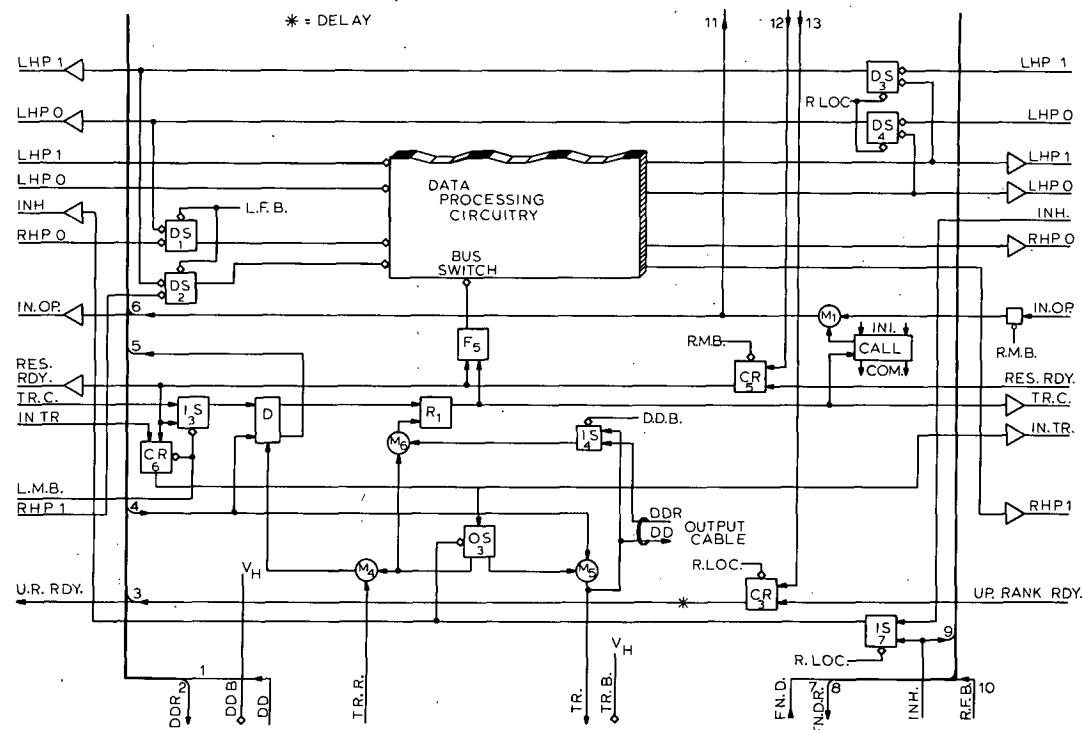


Figure 26. Multiply module control diagram.



Table 11. Horizontal Signal Pathways in the Multiply Module

<u>Pin</u>	<u>Direction</u>	<u>Top Multiply</u>	<u>Bottom Multiply</u>	<u>Load Module</u>
1, 89	←	Carry	RES.RDY.	F <sub>3</sub>
3, 87	←	MC <sub>11</sub>	Up. Rank Ready	F <sub>2</sub>
5, 85	←	Not Used	LHP <sub>1</sub>	F <sub>1</sub>
7, 83	←	Carry Ready	LHP <sub>0</sub>	F <sub>0</sub>
9, 81	←	Mult. and Count	Not Used	INH.
11, 79	→	Counter Status	IN.TR.	TR.C.
13, 77	→	Not Used	L.M.B.	L.M.B.
15, 75	→	L.F.B.	TR.C.	L.F.B.
17, 73	←	MP <sub>0</sub>	IN.OP.	IN.OP.
19, 71	→	Low. Rank Ready	LHP <sub>0</sub>	Not Used
21, 69	→	Not Used	LHP <sub>1</sub>	Not Used
23, 67	←	MP <sub>1</sub>	INH.	Not Used
25, 65	→	Low. Rank Control	RHP <sub>0</sub>	Not Used
27, 63	→	Do	RHP <sub>1</sub>	IN.TR.
29, 61	←	MP <sub>E</sub>	F <sub>0</sub>	RES.RDY.

The following are common to all module types:

31	Power Down Request
33	Protect (Data Shield)
35	Power Sense (Positive)
37	Power Sense (Negative)
39	Preset
41	Power Down Acknowledge
43	Spare
45	Ground
47, 49, 51	Power (Negative)
53, 55, 57	Power (Positive)
59	Ground

Note: Pin numbers refer to CSL wiring lists. All even-numbered pins are grounded.



Table 12. Multiplier Bit Grouping

<u>Group Number</u>	<u>Bits Examined</u>
1	MP <sub>1</sub> , MP <sub>0</sub> , MPE
2	MP <sub>3</sub> , MP <sub>2</sub> , MP <sub>1</sub>
3	MP <sub>5</sub> , MP <sub>4</sub> , MP <sub>3</sub>
4	MP <sub>7</sub> , MP <sub>6</sub> , MP <sub>5</sub>
5	MP <sub>9</sub> , MP <sub>8</sub> , MP <sub>7</sub>
6	MP <sub>11</sub> , MP <sub>10</sub> , MP <sub>9</sub>

Each Multiply module contains a counter to determine when the proper number of iterations (namely six) have been performed. When two or more modules are extended (to accommodate 12n-bit operands), 6n iterations are required. In such cases, the counter in the leftmost module steps through its six states and then enables the counter in the module to the right. At the Right Field Boundary the signal (Counter Status) is instrumental in producing the Result Ready signal. Details of this and other aspects of control in the Multiplier will be discussed subsequently.

Table 13. Results of Multiplier Decoding

<u>Contents of current Multiplier Group</u>	<u>Add to Partial Product Before Shifting Right Two Places</u>
000	Zero
001	1(MC)
010	1(MC)
011	2(MC)
100	-2(MC)
101	-1(MC)
110	-1(MC)
111	Zero

In the Multiplier, as in other modules, the activity begins with an Initiate Operation from the faceplate (for the rightmost module) or the module to the right. As indicated in Figure 26b, Merge unit M<sub>1</sub> produces the local Initiate Operation signal in response to either of these.



In addition to being transmitted to the module to the left, this signal is used to set  $F_1$  and  $F_2$  (Figure 26a) via signal path 11, and is connected to circuit elements  $CR_1$  and  $IS_1$  via path No. 6 in the left-hand trunk.

Note that  $F_2$  will be cleared automatically after a delay (indicated by the \* in the line between its two inputs). The pulse thus produced at the output of  $F_2$  is used to load the multiplier into the right half of the lower rank of flip-flops and zero into the left half of this rank as well as the extension of the multiplier (MPE). The flip-flops of the upper rank are not initialized.

At the Left Manifold Boundary the local Initiate Operation signal is passed through  $CR_1$  to create the Do signal. Since the Left Manifold Boundary is also a Left Field Boundary, the local Initiate Operation signal is passed through  $IS_1$  and is used to enable the modulo six counter in the leftmost module. (Note that the counter will be enabled in any module which occupies a Left Field Boundary position). No significant loss in generality will occur from assuming that only one field is involved in the multiplication operation. This shall be done from this point on.

Consider again the Do signal (created by  $CR_1$ ) in the leftmost module. Figure 26a reveals that this signal is passed to the module on the right and also connected to  $M_2$ , from whence a signal is supplied through  $OS_1$  to one of the inputs to  $IS_2$ . Recall that the conditioning input to  $OS_1$  was asserted by  $F_1$  at Initiate Operation time.

$IS_2$  is controlled by the local Right Field Boundary signal such that at the boundary a Multiply and Count signal is produced. The Multiply and Count signal is transmitted to the module to the left (and from there to the next module, etc.) so that the sequence of events to be discussed takes place in all extended modules at nearly the same time.

The Multiply and Count signal is supplied (see Figure 26a) to the data-processing circuitry as a Count signal, through a delay to the input to  $F_3$ , and through still another delay to the input to  $OS_2$ .

The Count signal will advance the counter which has been enabled (the Initiate Operation signal enabled the counter in the leftmost module). Actually, a Count signal is supplied to the counters in each of the horizontally extended modules. However, initially only those counters in modules occupying Left Field Boundary positions are enabled. After the first six iterations have been accomplished, the modules immediately to the right of the Left Field Boundary modules will have their counters enabled.

$F_3$  simply enables the flip-flop input circuitry such that the Upper Rank flip-flops will accept the output from the data-processing circuitry. In said circuitry, the contents of the left half of the lower rank of flip-flops are added to the proper multiple of the multiplicand (see Table 13). As indicated above, this sum is then placed in the flip-flops in the Upper Rank under the control of  $F_3$ .



The Multiplier employs two twelve-bit ripple carry adders in an arrangement referred to as a conditional sum adder. That is, one adder forms the sum of the correct operands assuming a carry from the module to the right, and the other does the same thing assuming no carry from the right. The carry signal from the module to the right (note that this is a level signal) then indicates which of the sums should be selected.

$OS_2$  is controlled by a signal called Shift Only. Thus, if zero is to be added to the current partial product, there is no need to allow time for the propagation of carries. In this case  $OS_2$  provides an input to  $M_3$ , which in turn supplies a signal to  $CR_3$  (on Figure 26b via path No. 13). At a Right Field Boundary this circuit element creates the Upper Rank Ready signal. Note that in non-Right Field Boundary modules the output of  $M_3$  is rendezvoused with the Upper Rank Ready signal from the module to the right to form the local Upper Rank Ready signal. Thus, when  $CR_3$  provides an output signal, it may be assumed that the data-processing operations have had sufficient time to have been completed in all modules to the right (within the same data field).

Consider now the case when the Shift Only signal is not asserted. This implies that some non-zero multiple of the multiplicand is being added to the current partial product. Under these conditions, the output of  $OS_2$  is rendezvoused (conditioned by the local Right Field Boundary signal) with a Carry Ready signal from the right. The delay in the Carry Ready signal is provided to guarantee that the carry actually is ready. The output of  $CR_2$  is delivered to the module to the left and is also connected to  $M_3$ .

Note that in addition to contributing to the formation of the Upper Rank Ready signal, the output of  $M_3$  resets  $F_3$ , thus latching the contents of the flip-flops in the Upper Rank. The various delays (denoted by asterisks) which have not been discussed in detail are provided to increase the probability of reliable and accurate data manipulation.

It should be apparent that regardless of the status of the Shift Only signal, a Multiply and Count signal applied to  $OS_2$  will eventually provide an input to  $M_3$ , which will in turn clear  $F_3$  and provide an input to  $CR_3$  (Figure 26b).

The local Upper Rank Ready signal is transmitted via path number 3 in the left-hand trunk to  $IS_3$  (Figure 26a). This circuit element thus produces the Lower Rank Control signal at a Left Field Boundary. (Note that  $IS_3$  is controlled by the Left Field Boundary signal.) The output signal from  $IS_3$  is sent to the module to the right and connected to  $F_4$  and (through a delay) to  $CR_4$ .

$F_3$  provides a pulse to the Load Lower Rank circuitry (note that the local Lower Rank Control signal is connected to the set input, and through a delay to the reset input). The contents of the Upper Rank flip-flops are transferred to the flip-flops in the Lower Rank under the control of  $F_4$ . The right shift required by the algorithm is accomplished during the transfer between Upper and Lower Ranks.



The delayed local Lower Rank Control signal is now supplied to  $CR_4$ , which (at a Left Field Boundary) creates the Lower Rank Ready signal. At non-Left Field Boundary locations, the local Lower Rank Control signal is rendezvoused with the Lower Rank Ready signal from the left to form the local Lower Rank Ready signal.

The creation of the local Lower Rank Ready signal indicates that sufficient time has elapsed for data transfers to have taken place between ranks in all modules to the left within the same data field. Ergo, at the Right Field Boundary the arrival of the Lower Rank Ready signal implies completion of the total add-shift iteration. This fact did not escape the attention of the designers of the Multiply mode.

The output signal from  $CR_4$  (the local Lower Rank Ready signal) is connected to  $M_2$ . Thus, a transition at the output of  $CR_4$  will have the same effect on the circuitry as the Do signal had. Each time a transfer between Upper and Lower ranks takes place, another iteration is begun. Let us review the operation to this point by considering a thirty-six-bit data field.

The Initiate Operation signal will set the  $F_1$  element in each of the three modules, serve as one of the inputs to  $CR_1$  in each module, and enable the counter in the module occupying the Left Field Boundary position.

At the Left Field Boundary the Do signal is created and passed to the module at the Right Field Boundary, where the Multiply and Count signal is produced (via  $M_2$ ,  $F_1$ ,  $OS_1$ , and  $IS_2$ ) and delivered to all modules within the field.

In each module the arithmetic operation dictated by the decoded current three bits of the multiplier is performed, and the results of the arithmetic operations are latched in the flip-flops in the Upper Rank. The only counter which advances during this first set of iterations is the one in the Left Field Boundary module.

After six iterations have been performed, the module at the Left Field Boundary will experience a transition on the Counter Status line. This signal clears  $F_1$  in the Left Field Boundary module and enables the counter in the module to the right.

After another six iterations the counter in the rightmost (in the 36-bit field) module will be enabled. Note that when the Counter Status signal in this module clears  $F_1$ , no more Multiply and Count signal transitions can occur.

After  $F_1$  has been reset, the final local Lower Rank Ready signal proceeds via  $M_2$ ,  $OS_1$  and data path No. 12 to one of the inputs of  $CR_5$  as shown in Figure 26b. Thus at the Right Manifold Boundary, the Result Ready signal is created. The Result Ready signal is propagated to the left (and used to gate the new data values on to the Down Bus) exactly as in the data-processing modules which have been discussed previously.

The sequence of Initiate Transfer, Transfer, Data Delivery, Data Delivery Return, Transfer Return, and Transfer Complete signals is identical



to those in the modules previously discussed, with the following exception. When a Transfer signal is transmitted downward via the implicit signal pathway, a Data Delivery signal is simultaneously provided to the output data cable port (see  $M_5$ , Figure 26b). An Input Selector ( $IS_4$ ) directs either this signal or the Data Delivery Return signal from the cable to  $M_6$  (as dictated by the absence or presence, respectively, of a data cable). The Data Delivery Return signal is thus rendezvoused (via  $R_1$ ) with the output of the D element to form the local Transfer Complete signal. This action is required because one-half of the product is delivered to the output cable port as an updated data word.

The circuit elements labelled DS in Figure 26 are Data Selectors. They behave in exactly the same manner as Input Selectors, except that one level input determines which of two other level inputs will be coupled to the output. Thus, for example, when the R.F.B. signal is asserted, Data Selector  $DS_4$  (Figure 26b) will couple the lower input signal to its output terminal.

Note that Figure 26 indicates two sets of signals labelled  $LHP_1$  and  $LHP_0$ . One set is directed from right to left and bypasses the data-processing circuitry, while the other set is directed from left to right and is shown entering and leaving the data-processing circuitry. Another pair of signals labelled  $RHP_1$  and  $RHP_0$  enter and leave the data-processing circuitry.

Since the algorithm employed by the module uses two new bits of the multiplier during each iteration, the entire multiplier is shifted two places to the right per iteration. Thus, the two least significant bits of the multiplier must be shifted across module boundaries. Actually  $RHP_1$  and  $RHP_0$  are shifted into  $RHP_{11}$  and  $RHP_{10}$ , respectively, in the module to the right. RHP stands for the Right Half of the Partial Product. No serious discrepancy in symbolism is involved, because at the completion of the multiply operation the flip-flops labelled RHP actually contain the right (least significant) half of the product. Initially they contain the multiplier, and at intermediate states of the operation they contain a portion of each.

The two least significant bits of the Left Half Product are similarly shifted two places to the right per iteration. For non-boundary modules the shifting is straightforward, with  $LHP_1$  and  $LHP_0$  being shifted into  $LHP_{11}$  and  $LHP_{10}$  in the module to the right.

The data stored in the Left Half Product registers always represents some part of the partial product; hence  $LHP_1$  and  $LHP_0$  are not discarded at the Right Field Boundary. At said boundary, these two bits are turned around via  $DS_3$  and  $DS_4$  and delivered laterally all the way to the Left Field Boundary, where they become (via  $DS_1$  and  $DS_2$ ) the two most significant bits of the Right Half Product. It is thus evident that at the completion of the multiply operation, the most significant half of the product will be in the LHP flip-flops and the least significant half will be in the RHP register.

Data Selectors  $DS_5$ ,  $DS_6$ , and  $DS_7$  (all of which are controlled by the local Right Field Boundary signal) supply the current contents of the two least significant bits of the multiplier as well as the contents of the multiplier extension to the decoder. Thus, while each module has a decoder,



only the rightmost three bits of the multiplier and its one-bit extension are decoded.

### 2.3.3 Decision Modules

Each of the Decision modules monitors the Up Bus continuously, but they do not put any data on the Down Bus. Each bit carried by the Down Bus (twelve data bits plus the flag bit) is buffered by a non-inverting amplifier within the Decision Module.

The Decision modules similarly pass Transfer Boundary signals to the module immediately below them. That is, the Decision Module is transparent to a Transfer Boundary signal from a Data-Changing module above it. Likewise, Transfer and Transfer Return signals are passed by (but never originate in) the Decision modules. The Data Delivery Boundary condition is asserted to the module immediately below a Decision module to indicate the absence of a boundary.

### Compare Module

The Compare module (sometimes called the Comparator) accepts data from the Up Bus and two data cable input ports. The data from the Up Bus(R) are compared to various constants or to the twelve-bit input operand (D) or selected bits thereof, under the control of the four-bit function code and the twelve-bit selection variable (S) input. A summary of the comparisons performed by this module is presented in Table 14.

The control signal ports on the faceplate box employed by the Compare module include one initiation terminal and two completion terminals which are labelled "Y" (yes) and "N" (no). The completion signal is supplied to the completion terminal whose label agrees with the results of the comparison.

Comparators may be extended laterally, and the manifold may be subdivided into two or more data fields. The result of a comparison on a multiple-field manifold is "No" only if the result is "No" for each independent field.

A bit pair is selected by supplying a logical "1" in the corresponding bit position of the selection variable (S).  $R_M$  and  $R_{M-1}$  refer to the two most significant bits of the entire word in an extended case.  $R_f$  refers to the flag bit in the leftmost column of a data field.

Figure 27 depicts the control circuitry in the Compare module.

The Compare module responds to an Initiate signal from the faceplate or from the module to the right and passes an Initiate Operation signal to the module to its left. At the Left Manifold Boundary (after a suitable time delay) a Test Complete signal is created and propagated to the right. At all non-leftmost locations, the Test Complete signal from the left is rendezvoused with the local Initiate Operation signal before being passed on to the right.



Table 14. Operations Performed by the Compare Module

<u>Op. Code</u>	<u>Comparison</u>	<u>Decision</u>
0	$R_i = D_i$	Yes, if no selected bit pairs differ.
1	$R = D$	Yes, if $R = D$ .
2	Flag Set	Yes, if $R_f = 1$ .
3	$R = 0$	Yes, if $R = 0$ .
4	R Not Normalized	Yes, if $R_M = R_{M-1}$ .
5	$R \geq D$	Yes, if $R \geq D$ .
6	$R \geq 0$	Yes, if $R_M = 0$ .
10	$R_i \neq D_i$	Yes, if no bit pairs are selected or if one or more of the selected bit pairs differ.
11	$R \neq D$	Yes, if $R \neq D$ .
12	Flag Cleared	Yes, if $R_f = 0$ .
13	$R \neq 0$	Yes, if $R \neq 0$ .
14	R Normalized	Yes, if $R_M \neq R_{M-1}$ .
15	$R < D$	Yes, if $R \geq D$ .
16	$R < 0$	Yes, if $R_M = 1$ .

Two level signals (as opposed to the more typical transition signals) are transmitted from left to right. These are labelled Condition 0 and Condition 1. These two signals convey the results of the tests which have been performed in the modules to the left. These results are:

- 1) Yes (Condition 1 and Condition 0 both negated)
- 2) Conditional Yes (Condition 1 negated and Condition 0 asserted)
- 3) Conditional No (Both asserted)

A Conditional No becomes a firm No at the Right Manifold Boundary. Since Condition 1 is asserted only for the Conditional No, this signal directs the Test Complete signal (with the aid of an Output Selector) to the appropriate completion terminal. Note that completion signals appear



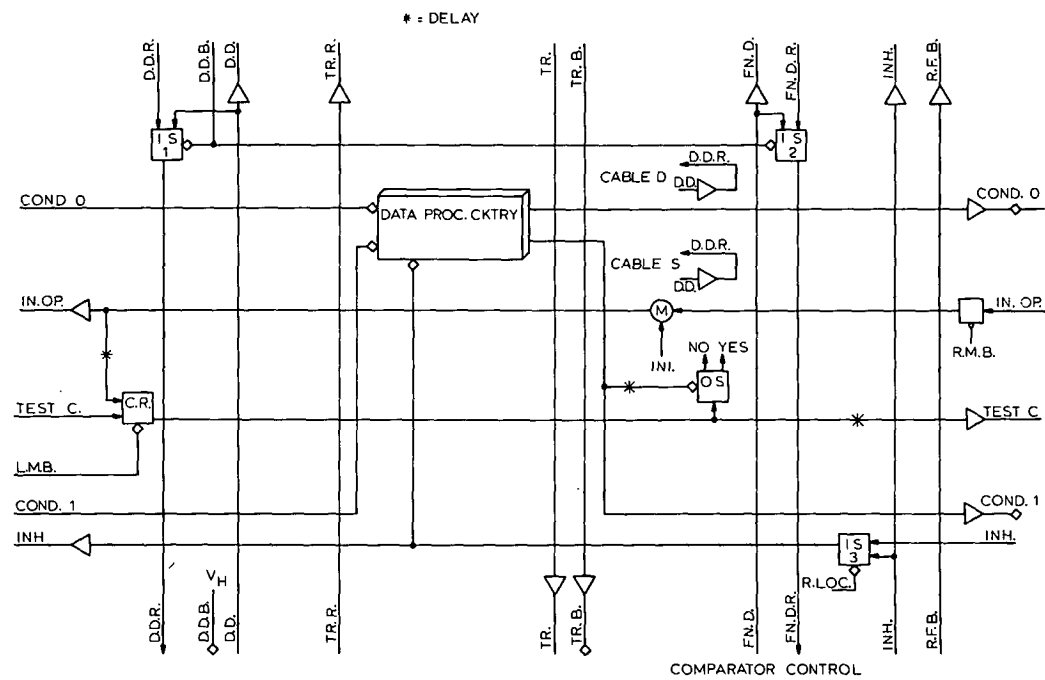


Figure 27. Compare module control diagram.



at control cable ports on the faceplate boxes of non-rightmost modules, but control cables should normally be connected only to the ports on the rightmost module.

Since the Compare module does not put data on the Down Bus, no circuitry is needed to initiate the Transfers.

### Decode Module

The Decode module decodes any selected three bits of data from the Up Bus and provides a signal to exactly one of eight completion terminals on the faceplate.

Bits  $R_{11}, R_{10}, \dots, R_0$  are candidates for selection as dictated by the location of logical ones in  $S_{11}, S_{10}, \dots, S_0$ . The three bits of  $R$  which are in bit positions corresponding to the three lowest-order non-zero bits of  $S$  are selected and labelled  $D_2, D_1, D_0$  (with  $D_2$  the most significant bit). If  $S$  contains  $3+i$  non-zero bits, the  $i$  most significant non-zero bits are assumed to be zero. The selected bits are treated as a three-bit word, which is decoded into the eight possible states.

Since the Decode module does not put any data onto the Down Bus and does not extend laterally via implicit data pathways, the control circuitry is somewhat simpler than that of modules previously discussed. Figure 28 depicts the control.

An Initiate signal, after a delay, produces a completion signal at the appropriate faceplate terminal. The module does not transmit any lateral signals, nor does it initiate any vertical signals.

The Decode module passes Data Delivery and Transfer signals as well as their return signals. It also accepts a Data Delivery signal from the input data cable (selection variable) and provides the return signal.

### 2.3.4 Miscellaneous Modules

#### Data Branch Module

The Data Branch module accepts a 12-bit data word from an input data cable and supplies this same 12 bits to two output data cable ports on the faceplate. The 12-bit data word is also provided to the Up Bus.

The Data Branch unit interrupts the Down Bus, so no data-changing module should be located above it in the data-processing manifold. Decision modules which do not require the flag, or access to the function code via vertical implicit pathways, may be above the Data Branch unit; so may the D/A Converter module (see below).

Figure 29 represents the Data Branch module control circuitry. Since this module operates continuously, no initiation or completion signals are involved.






Figure 28. Decode module control diagram.






Figure 29. Data Branch module control diagram.



The Data Delivery signal from the input data cable is applied to the vertical (upward) implicit data pathway and to the two output data cable ports on the faceplate. The D.D. signal is also applied to the input selector (left center Figure 29 ).

If the module is located at a Data Delivery boundary, the data delivery signal (otherwise the Data Delivery Return from above) is connected to one of the inputs to the Rendezvous element. Data Delivery Return signals from the two output data cables (or the original Data Delivery signal) provide the other two inputs to the Rendezvous element.

Thus, when either a Data Delivery Return signal or an indication that none is required is received independently from each of the three data output paths, the Rendezvous element provides a Data Delivery Return signal to the input and cable.

#### D/A Converter Module

The D/A Converter module accepts a twelve-bit operand from the Up Bus and provides two analog output voltages in the range of -5 V to +5 V. Each of the analog outputs comes from a double-rank register which allows a new digital value to be stored without affecting the current analog output signal. The D/A Converter module also provides (upon request) a 5-volt (1-microsecond) pulse at an output terminal which is labelled "intensity out".

The faceplate box has initiate terminals to load register A, load register B, convert (this updates both analog outputs), and intensify. Output signals appear at corresponding completion terminals after suitable time delays.

The D/A module extends the Up Bus (including the flag) and is transparent to the Down Bus and the Parameter Bus. It handles Transfer, Data Delivery, and Return signals in a manner similar to that of the Compare module. (See Figure 30.)

#### 2.3.5 Control Modules

##### Merge/Rendezvous Module and Call Module

The Merge/Rendezvous and Call modules are trivial with respect to the functional blocks defined in the Introduction (section 2.1). The Merge/Rendezvous module consists of six two-input Merge elements and six two-input Rendezvous elements, with the output of each element available on two lines. The twenty-four control lines leading to faceplate box control connectors are divided into six groups of four control lines each, two inputs and two outputs. Each group of two inputs and two outputs is associated with one Merge element and one Rendezvous element; which element is selected for each group is determined by selection circuitry controlled by the code switches on the faceplate box.

The Call module consists of two three-input, two-return Call elements, exactly as described in the Introduction's definition of the Call element (see Figure 15 ).






Figure 30. D/A module control diagram.



### Function Call Module

Figure 31a shows a faceplate-box view of the control signals in and out of the Function Caller. The module functions as a six-input, two-return call. So that more than one module can be stacked to extend the number of calls, the sixth input to the call is from the module above. The six connections to above and below shown in Figure 31a are carried by the internal up-down bus.

A more detailed view of the control is given in Figure 31b. Before stepping through the control sequence, it would be worthwhile mentioning a few details in Figure 31b. The signal TB is the Transfer Boundary from the module above; if asserted, the Initiate from above ( $I_6$ ) is inhibited. The signal BM denotes "Bottom-Most", and it comes from the code switch on the faceplate box. When BM is not asserted, the  $Y_6$  and  $N_6$  returns from below are allowed to enter; and TB (to below) is not asserted.

When an Initiate signal arrives, the associated S and GS levels are asserted. GS has no effect until later; but S is used to select the appropriate four-bit function code, and then all the S lines are merged (ORed) to form the DO signal. Nothing happens in the module following a DO, until a Y or N is returned. The returning signal is ANDed with the asserted GS, causing the correct completion output and causing a pulse ( $P_Y$  or  $P_N$ ) to be produced. The pulse is used to clear the asserted S and GS lines, and the sequence is finished.

### Interlock Module

The Interlock module consists of three main functional parts, as described below.

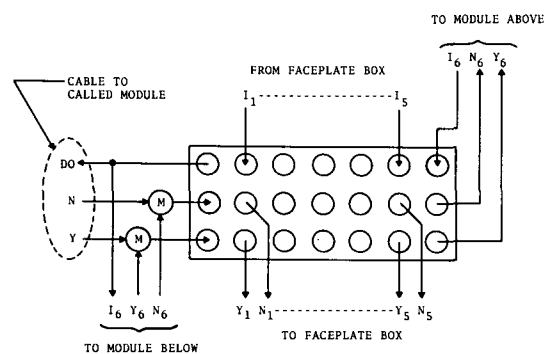
Glitch-Detecting Circuit: On each input flip-flop is a circuit with a GC-1 biasing network and two M35 comparators. The wire-ORed output is low when the flip-flop is in the middle region.

Priority Circuit: Each output flip-flop is controlled by two AND gates (M10 and M05) which toggle the flip-flop when: (a) the corresponding input is active; (b) no inputs of higher priority are active; and (c) the priority circuit is enabled.

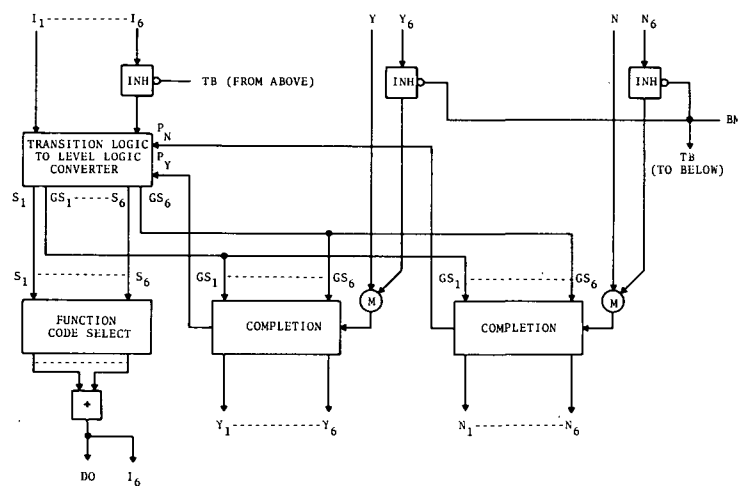
Unlock: The two unlock signals associated with each segment ( $U_Y$  and  $U_N$ ) are merged on the Return Boards and sent back as REL (Release) to the Control Board.

When an Initiate signal arrives, it must pass through a gated flip-flop (1016) in order to be processed. When the module is not busy with any inputs, the gate is left open (on). The first initiate to pass through an open gate causes the gate to close; but there is 15-20 ns in which other Initiate signals may also pass through the open gate. The arrival of an Initiate just as the gate closes may cause a glitch. This condition is detected by Chaney's "Glitch-Detecting Circuit", and a signal is supplied to the priority circuit. As the gate is closed, a signal is sent through a short delay to the priority circuit.





(a)



(b)

Figure 31. Function Call module control diagram:  
(a) simplified; (b) detailed.



This signal, plus the absence of the glitch signal mentioned before, enables the priority circuit; then the Initiates are processed in a predetermined order (left to right). When an input is processed it produces a D0 output. This eventually comes back as an UNLOCK, which is used as a signal to the priority circuit to process the next Initiate signal. When the last UNLOCK is received, the gate opens again.



### 3. MECHANICAL DESIGN

#### 3.1 INTRODUCTION

The mechanical design process which has ultimately led to one working version of macromodule supporting hardware has, like any endeavor of this kind, been an incremental evolutionary process, each step of which has been gained by coming to grips with and solving the practical problems associated with the conceptual workings of the system. The following material is a reconstruction of the rationale for design and manufacturing goals which have led to the production and utilization of the macromodular system described herein.

The design of mechanical supporting hardware for macromodules has evolved around the notion that, like the modules themselves, the support equipment must be modular in nature to allow practical and economical systems of arbitrary size to be configured. Moreover, support functions must include not only the structural elements that physically house and support the modules, but cooling and integration of power and communications paths as well. Furthermore, the system must be user-oriented to allow assembly of components in a quick and efficient manner, without the need for special assembly tools and requiring a minimum number of personnel.

With these basic goals in mind, the design of mechanical hardware proceeded in close cooperation with the electrical engineering staff engaged in module definition and its hardware implementation. Following a period of preliminary "paper design" studies and several experimental model studies, the module configuration and its operating structure evolved into its present form.

The milestone that had been attained at this point in the design process was probably the most significant with respect to the eventual success of the entire project. Commitment to the Daisy Chain bus structure or manifold, the definition of a module set, and the establishment of the basic printed circuit construction techniques which would be utilized for module assembly allowed mechanical design tasks to proceed in a quantitative way. Design goals were identified with sufficient detail to allow concentrated efforts to be directed at development of specific hardware. At this time the principal items of hardware that were defined and which today are found in the macromodular support set included:

Base Pedestal: the structural foundation upon which a macromodular system is built and in which reside power supplies and system-wide logic hardware.

Electronics Package: the functional module unit.

Faceplate Box: a unit associated with modules that contains switches and connectors to provide codes, data, and control signals to a module. In addition, the faceplate box serves to provide vertical pathways to adjacent modules.

Frame Block: a cellular array used for the construction and support of macromodular systems.



Fan Module: a plug-in unit that provides cooling to macromodules and serves to distribute electrical power and system-wide logic signals.

Channel Coupler:- an electrical coupler used for lateral data extension between horizontally adjacent frame sections.

Figure 32 illustrates the relationship that these mechanical systems share with one another in the assembly of a macromodular array.

### 3.2 FUNCTIONAL DESCRIPTION AND RATIONALE OF THE SUPPORT SET

One of the initial needs with respect to supporting hardware was the design of the electronics case for the modules' printed circuit boards. The electronics case serves three primary functions - first, it is a protective aluminum shell that houses and supports the printed circuit boards and associated components of a module. Second, the geometry of the case serves as ducting to allow air flow over electronic components; and finally, the case provides mechanical alignment for the engagement of electrical connectors through which the module communicates with other parts of a system. Figure 33a is an illustration of the single-cell electronics case. This particular case is the elemental size in which these enclosures are manufactured; however, provision for larger amounts of electronics is provided by expanding the case in the vertical direction. This expansion capability is seen in Figure 33b, where a 4-cell electronics case is illustrated. It will be noted that a number of the subassemblies of the 4-cell case are common to the single-cell case. The reason for this commonality was to economize on the inventory of different components making up the total assembly, thereby reducing the overall cost of the unit.

At this stage in the design process a number of issues related to the size of the cellular array, or frame block, required resolution in order to proceed with the details of mechanical design. The main issue was the size of the frame block that could be easily transported and assembled, yet would provide a sufficient number of cells to avoid hardware redundancy.

The principal criterion for setting the width of the frame block was that it be able to pass through a standard 30-inch door frame. The chosen width, based on the previously established module dimensions, was 5 cells. In the vertical direction the ultimate height was controlled by the minimum typical ceiling height likely to be encountered in use; this was taken to be 8 feet. Subdivision of the column between base-pedestal and ceiling suggested that in order to yield manageable frame block sizes they should probably not exceed 4 cells in height. The result was a frame block of twenty cells which could be stacked eight high, yielding an array of 160 single-high cells.

Of the 20 single cells available in a frame block, only 16 may be randomly used. Four cells, vertically adjacent and centrally located, are occupied by the fan module. As its name would suggest, one of the principal roles of the fan module is to provide a means of cooling the








Figure 32. Overall view of macromodular assembly.





(a) Single-cell case.



(b) 4-cell case.

Figure 33. Electronics case.



modules residing in the respective frame block. Furthermore, the fan module acts as a distributor of electrical power and communications on a local as well as a system-wide basis. As a system-wide distributor, the fan module contains an internal bus structure terminated in electrical connectors at top and bottom. Stacking of frame sections atop one another continues the bus from section to section. On a local basis, the necessary power and communication paths are distributed laterally to each row of cells through the lateral channel assembly. The lateral channel contains not only the printed circuit boards required for electrical pathways to and from modules, but the ducting system to provide cooling as well. The ends of each channel terminate in electrical connectors to continue lateral communication paths between modules. Power and ducting are segregated on a frame block level.

Residing within each fan module is a belt-driven Centraxial blower wheel and motor manufactured by the Rotron Co. of Woodstock, N.Y. The Centraxial blower was chosen for use on the basis of requirements imposed by the frame geometry; i.e., the air handler must overcome the pressure drop resulting from the somewhat tortuous path encountered by the flow, yet supply sufficient air volume to cool the 16 modules it supports. Figure 34 illustrates the flow path at each row of cells.

Based on the absolute minimum air flow requirements to each module to maintain safe operating temperatures, 2.5 cubic feet of air per minute was needed. By belt-driving the fan, which was necessary because of the limited width available, 4.7 cubic feet per minute could be attained while keeping the noise levels generated by the blower wheel within reasonable limits.

It will be noted in Figure 34 that air flow through each module is a result of the decreased pressure in the lateral channel. It had been determined through a series of experimental investigations that pumping air into the channel ducting caused a piling up of the flow at the ends of the channel due to the momentum of the high-velocity stream. The result was a far greater air flow than required to the outside cells and a starving of flow to the inside cells. By "drawing" air through the channel, far greater uniformity of air velocity through the module was achieved. Furthermore, the flow through each cell could be balanced by proper adjustment of duct partitions. The main virtue of this design is the fact that the pressure drop through each module is much less than that through the ducting. As a result, it is of little consequence to the net air flow at each cell location as to whether or not a module is actually present.

It should be pointed out for the historical record that the arrangement just described did not suddenly reveal itself. There were a number of fan designs proposed and some actually built. Most of these early designs centered on a more modular, one cell row / one fan theme. Through lessons learned by trial and error, redesign and experimentation, the consolidation of the cooling unit into a single module per frame block eliminated many redundancies of equipment to allow packaging of not only the fan and motor, but all the bus structures and attendant electrical components, into the space of four single-high modules.






Figure 34. Air flow ducting for a row of cells.



The major strength of the frame block is derived from the six section plates that form the 20 module cells. Each cell is defined by two section plates forming the right-left boundaries and four rails defining the corners of the cell. The rails guide the module into engagement with electrical connectors that provide means for communication and a source of power. One of these connectors resides in the respective lateral channel. The weight of the module is borne by the rails, which transfer their load through the section plate to the front and rear posts. It is the front and rear posts that ultimately transfer the load of the modules and structure to the base-pedestal. It should be noted that the section plates do not contact each other at frame block boundaries, in order that the problems associated with matching post and plate surfaces simultaneously could be avoided. Securement of frame sections to one another is accomplished by inserting a socket head screw at each post location of the frame-to-frame joint. Lateral spacing of frame sections is achieved through the use of notches cut at appropriate locations in the lateral channel wall into which the section plate is fitted. A screw in the front post pushes the notch in the channel against the section plate to secure the assembly.

To begin the construction of any macromodular system requires a base pedestal which contains the necessary power supplies and system-wide logic hardware to support the modules stacked in frame sections above. In addition, this pedestal provides structural support and leveling mechanisms to allow for the alignment of the cell rows or, more importantly, the lateral channels of the frame blocks.

Strength considerations were of prime importance in the design of the pedestal, since it had been estimated that over one-half ton of equipment could be stacked upon this structure. In addition, the rigors that transportation devices impose upon their contents was a loading condition which the pedestal and its electronics in residence would be required to survive.

With these thoughts in mind, the pedestal took on the appearance shown in Figure 35. The principal structural members of the pedestal are its 3/8-inch aluminum plate sides and walls. These members distribute the loads they support to adjustable feet located at the four corners of the structure. Beneath the pedestal all members terminate at the same horizontal elevation, to allow the forks of a lifting device to evenly distribute their effort. On the topside of the pedestal, the cover has been braced for the inevitable load of a person standing on it to reach a module beyond his floor-level grasp. The resident power supply and logic is housed between the front and rear wall, which support the frame blocks and provide protection for the electronics housed within this area. Guide rails serve to align connectors of three plug-in supplies with resident connectors. These supplies are added as needed when expanding a system.

Although provision was designed into the base pedestal for vertical adjustment, no provision was made for side-to-side spacing or coupling of pedestals. The rationale here was that the channel coupler should be designed in such a way as to have sufficient compliance to make up for any dimensional variations resulting from hand positioning of pedestals next to one another. The final version of the channel coupler design to meet these requirements can be seen in Figure 36. The connectors of the device remain








Figure 35. Base pedestal.





(a) Open.



(b) Closed.

Figure 36. Channel coupler.



retracted as the coupler is pushed into the slot formed by the butting together of two channel ends. The handle halves serve as springs to push each side of the coupler into registration with the side of the respective channels, thereby taking up variations in side-to-side placement of frame sections. The coupler is stopped in its fore and aft travel by the back end of the channel, whereupon the handle is pushed down, causing the connectors to engage their mates residing in the ends of the channel. Twisted pairs of wires provide data paths across this boundary.

Associated with every module type is a faceplate box through which data and control signals flow and in which operation codes may originate. Like the modules they serve, faceplate boxes may be constructed in incremental sizes. Figure 37a illustrates the elemental or single-cell faceplate box, while Figure 37b shows a two-cell-high box.

Each faceplate box mates with its respective module through two 90-pin connectors. In addition, it may mate with a module located below itself through another 90-pin connector so as to pass data along in relay-like fashion from module to module. The connectors chosen for this application had very good characteristics with respect to the reasonable forces required for engagement and withdrawal, but they unfortunately had very minimal guidance geometry to assist in mating. As a result, the proper alignment of these connectors relied very heavily upon the dimensional accuracy of frame sections, electronics cases and faceplate boxes.

At each cell location, the frame section rails act as a common reference for registration of modules and faceplate boxes. As a result, careful attention was demanded in the manufacture and assembly of the section plates and rails to insure, under worst-case tolerance conditions, that all connectors involved with engagement at a cell would do so without causing damage to the somewhat vulnerable male pins of the connector.


Each faceplate box is retained in the frame block by its "ears" (see Figure 37). The ears are spring-loaded and snap into place in the front post so that the box is restrained against movement during insertion and withdrawal of cables and modules. In addition, a macro-modular system may be constructed and wired and left undisturbed while the modules themselves are made available to other systems.

### 3.3 MANUFACTURE OF MACROMODULAR MECHANICAL HARDWARE


Various factors dictated very close dimensional control of hardware assemblies, in order that a guaranteed assembly of modular pieces could be expected. One of these factors, mentioned above, concerned the connectors chosen for module construction and their reliance on hardware guidance for proper mating. Another was the additive tolerance problem associated with the stacking of up to eight frame blocks upon one another; this might possibly multiply allowed tolerance variations eightfold, thereby creating lateral mismatch the channel coupler could not tolerate.

To deal with these problems, emphasis was placed not on individual components but upon the overall dimensional accuracy of the final assemblies. With this goal in mind, shops were selected not only on the basis of





(a) Single-height box.



(b) Double-height box.

Figure 37. Faceplate box.



whether or not their facilities could turn out components, but whether or not assembly and inspection procedures could produce workable units. It was much to our advantage in terms of personnel effort and expense to receive finished assemblies ready for use.

To test the feasibility of this notion, an independent machine shop was selected to inspect assemblies which had been produced at another and unrelated shop. This exercise indicated that about 30% of the finished frame section assemblies had some out-of-tolerance condition that fell slightly outside of print specification. These frame sections were made to be usable, but a subsequent purchase contract for mechanical parts was awarded to a new shop which included an extensive inspection facility.

Despite these efforts, a small percentage of equipment received will inevitably require some attention to bring it into specification.

The manufacturing procedure that probably caused the most extensive problems in component production was the finishing process. A great deal of the macromodular system is made from one or more aluminum alloys. A typical finish for these alloys, and the one used quite extensively in macromodular design, is anodized over a surface prepared by a light glass ball peening and etch. If the etching process is not carefully controlled, an excessive amount of material could be removed from the part surface. As a result, screw threads become sloppy, roll pins no longer press fit, parts do not assemble properly, and in general there is an overall degradation of dimensional accuracy. Notes of warning have been liberally included in mechanical documentation to help avoid these kinds of problems.

### 3.4 COSTS OF FABRICATION

Because of the extreme demand for accuracy in this macromodular system, fabrication costs have tended to be somewhat high. The following cost estimate, based upon one fabricator's bids, itemizes the part name and its respective share of the cost of setting up one frame block of equipment to support 16 modules. These costs include manufacture of components, assembly, and (where applicable) circuit boards and wiring.



Table 15. Macromodule Costs

Frame Block (16 cells)

Frame Sections	6 x \$50 = 300
Channel hardware	4 x \$44 = 680
Channel PC boards, connectors, ducting	4 x \$ 170 = 680
Fan Module (shell, PC boards, connectors, and wiring)	1 x \$ 300 = 300
Channel Couplers	<u>4 x \$85 = 340</u>
	\$1796

Base Pedestal (1 per 8 blocks maximum)

Power supplies	4 x \$ 250 = 1000
Pedestal structure	1 x \$1400 = 1400
Pedestal connectors and wiring	<u>1 x \$ 200 = 200</u>
	\$2600/8 = \$325

Faceplate Box (maximum 16 per block)

Shell and wiring	16 x \$150 = 2400
Connectors	<u>16 x \$ 60 = 960</u>
	\$3360

Electronics Package (metal work)

16 x \$ 42 = \$678

GRAND TOTAL

= \$6159 / frame block

= \$ 385 / cell



#### 4. ELECTRICAL CONSIDERATIONS

##### 4.1 INTRODUCTION TO ELECTRICAL DESIGN

Since one of the goals of the macromodule project is to allow systems of arbitrarily large size to be assembled, the electrical design must address itself to arbitrary size. There are clearly practical limits or constraints, based on such factors as primary power available, signal attenuation in cables, reliability of components, cost, etc., that limit the size of an operating macromodular system, but we wished to achieve an electrical design that would, in principle, be adequate for systems of arbitrary size. The necessary electrical constraints are enforced in several ways: sometimes by detecting and indicating unallowed conditions, as in the load sensing for the base pedestal, but usually by making violations of the constraints physically impossible.

Control of parameters such as crosstalk, transmission line reflections, and other parameters that affect signal integrity was the general goal of the electrical design. Because of the frequent reconfiguration that is important to the intended usage of the macromodules, it would not be practical to construct a system and then patch or adjust troublesome signal paths until reliable signal transmission was achieved. Due to the unclocked nature of macromodule logic design, the control of control signal integrity is especially important; whereas in a clocked system only the clock signal must be kept "clean" and well controlled, since the effects of crosstalk or other AC problems on other lines can be reduced by extending the time between clock signals.

Other goals such as speed, power consumption, and size, while important, are secondary to this primary goal of signal integrity. Motorola Emitter Coupled Logic (MECL) was chosen as the logic family because of its ability to satisfy the primary design goal of maintaining signal integrity better than any of the other available logic types, and not for its speed. The discussion in the remainder of this section should explain this choice. The MECL Integrated Circuits Data Book published by Motorola gives a description of the operating principles of MECL logic. Some of the important characteristics of MECL that influenced its selection for the macromodule project were: low input current (100  $\mu$ A MAX) which reduces voltage offsets due to ground return current, capability to drive transmission lines, nearly constant power supply current independent of logic state, availability of differential outputs well matched in level and propagation delay, and a reasonable ratio of noise margin to signal swing. The propagation delay of MECL is smaller than that of most other logic families, which was a desirable but not necessary characteristic. All MECL outputs are NPN emitter followers. They can only source current, so pulldown resistors must be provided, either within the IC package or else external to it, to sink current. In the examples that follow, the value of the pulldown resistor shown will be the parallel combination of external and internal resistors.

The decision to use MECL circuits presented some risks, since it was a relatively uncommon logic family. We still believe it was the best choice, however, despite some of the problems we have encountered.



## 4.2 CLASSES OF INTERCONNECTIONS

There are three distinct classes of interconnections within the macro-modular system: those within a module, those between two adjacent modules via implicit pathways, and those made by cables. The connections within a module are subjected to the fewest sources of degradation, since their drivers and receivers have approximately the same power supply voltage, ground potential and temperature, and the lengths are relatively short. The implicit connections between modules are more susceptible to degradation, since temperature and power supply voltages can differ, and distance, while still relatively short (about 3 feet maximum), is longer. The cables present the most difficult problem, because of their length and the possible physical separation between the connected modules.

## 4.3 TRANSMISSION LINE TERMINATIONS

With the signal rise and fall times of about 5 nsec. characteristic of MECL II logic, many of the connections, even within a module, must be treated as transmission lines and terminated to control reflections if reliable signal transmission with minimum delays or settling times is to be achieved. Three methods of termination are available: series or sending end termination, shunt or receiving end termination, and a combination of both. In addition, each of these methods can be used with a single-ended or differential-signal transmission. The following sections explain the differences between the termination methods, the parameters that affect them, and the reasoning behind the choices made for the macromodules.

### 4.3.1 Sending End or Series Termination

Series termination, shown in Figure 38, propagates a half-amplitude signal voltage to the receiving end, where it is reflected in phase and the reflection returned to the sending end and terminated. For MECL II, the gate output is an NPN emitter follower with nominal high output level of




Figure 38. Series termination.



-.800 volts and nominal low output of -1.600 volts.  $R_p$  represents the total output pull-down resistance and acts as a current sink, since the output emitter follower can only source current.  $R_o$ , the terminating resistance, is equal to the transmission line impedance minus the nominal output resistance of the MECL gate. The nominal line impedance is 130 ohms, and the series terminating resistor is chosen as 121 ohms. Some of the parameters that contribute to signal degradation with this method are tolerances on the transmission line impedance, actual value of driver output impedance vs. current, input current to the receiver, capacitance of the receiver and (to a lesser extent) of the driver, distance of the driver to the series terminating resistance, and impedance of the terminating resistance.

The transmission lines are formed by nominal 12.5 mil lines on 1/16 inch, two-sided FR-4 glass epoxy printed circuit boards. The width is a compromise between narrower lines, which would be more subject to breaks and would have greater percentage variation in width, and wider lines, which would produce more crosstalk and be more susceptible to P.C. board short circuits. If connections are long enough to be considered as transmission lines (about 4"), they are routed on the signal side of the P.C. board, with a ground plane placed opposite them on the component side. The nominal transmission line impedance for this construction is 130 ohms, and manufacturing tolerances give it a range from about 130 ohms to about 108 ohms. The output impedance of the driver is more difficult to characterize, since it is nonlinear and dependent on output current. A compromise value of 8 to 9 ohms is used for this value although it varies, tending to be lower for high outputs since there is more output current and tending to be higher for low level outputs. This plus the tolerance in transmission line impedance means that the reflection from the receiving end may not be terminated completely, but will produce a small re-reflection that propagates back toward the receiver.

The input current of the receiver causes a DC drop across the series terminating resistance. In MECL II circuits the maximum single gate input current for standard gates is 100  $\mu$ A, and this represents a loss of 12 mV for each receiver connected to a series terminated line. Actually the input current is only drawn by the receiver when its input is high, but the effect on high-level signals is approximately the same as if the current was constant. Because there is no input current for low input signals, there is no gain in low-level noise margin due to the series resistor.

The capacitance from the receiver input terminals to ground causes the rise time of the signal at the receiver to be slowed and also causes an initial out-of-phase reflection from the receiving end, which is terminated at the sending end but affects the maximum source and sink current required from the driver. Incident and reflected waveforms are shown in Figure 39 for a negative-going signal. For a low-going signal the pulldown resistor must sink a peak current of  $(V_H - V_L) / R_o$ , or about 6.6 mA, when the first part of the reflection reaches the sending end. This is modified somewhat by the fact that the transmitted waveforms are closer to trapezoidal waveforms than square waves, causing the point of the reflected waveform to be rounded and lowered somewhat. Also, the input capacitance of a MECL II receiver has approximately 60 ohms in series with it. This also limits the peak amplitude






Figure 39. Incident and reflected waveforms for a series-terminated line.

of the initial reflection for a single receiver, but has little effect if a number of receivers are paralleled. As an extreme example, the following waveform shows the effect of a 4 ns ramp and a 20 pf receiver capacitance representing four paralleled receivers.




Figure 40. Effect of a 4-ns ramp and a 20-pf receiver capacitance.

With this condition, the maximum current the pulldown resistor is required to sink is 3.97 mA. The effective resistance of 15 ohms in series with the receiver capacitance and the rounded edges of an actual waveform would further reduce the peak current. The nominal pulldown resistance value used is 750



ohms, which may be a discrete 750 ohm resistor, two paralleled I.C. pull-down resistors of 1500 ohms each, or an internal pulldown and an external 1500 ohm resistor. This will sink a nominal current of 4.93 mA when the driver output is low.

When the driver output is high and the reflection from the receiver is returned, the driver must source 6 mA required by the pulldown plus the current required by the reflected wave. It might appear that the high-level noise margin would be less than the specification, since the output current is higher than the specification point. The output of the gate may be less than the specified high level under this condition; but it is the equivalent voltage in series with the driver output impedance that is significant, and this should remain above the specified high output level. The effective driver output impedance is dependent on the output current, and these variations in output current cause some mismatch and consequent small reflections from the sending end.

Loads other than at the far end of the transmission line cause additional reflections due to their capacitance, and see only a half-amplitude signal until the reflection from the receiver returns to them. In some cases receivers that are very close to the driver are connected between the driver and series terminating resistor and driven as unterminated loads. In this case they add capacitance at the driver, which slows its fall time and may allow the output emitter follower to cut off as the driver output goes low. If the reflected wave from the receiver arrives at the sending end while the output emitter follower is cut off, the termination appears as the terminating resistance in series with a capacitor.

In some cases, receivers may be distributed along a series-terminated line, and their input capacitance taken into account in determining the effective impedance of the line. Some reflections will be caused by the lumped nature of the added capacitance. The longest propagation delay will be for the receiver nearest the driver; and transition or edge-sensitive devices should not be used as receivers, since their input may be in the threshold region for a considerable time.

The series-terminating resistor should be placed close to the driver, since the path connecting them is a transmission line, and its impedance can affect the signal transmission. This consideration limits the usefulness of resistor networks somewhat; and it, plus the fact that discrete resistors could be used as jumpers to cross other printed circuit paths, influenced the decision to use discrete resistors for the series terminations.

The resistors used are 1% metal film, which were chosen for their accuracy, stability and low series inductance. The nonidealness of the resistors, and other effects such as impedance discontinuities at connectors, have not caused noticeable deterioration of the signals.

The general guidelines for use with MECL II in the modules has been to use series termination for lines over 3" long, to allow a maximum fanout of 4 loads lumped at the receiving end for signals within a module, and to allow a fanout of 1 for signals between two adjacent modules.



#### 4.3.2 Shunt or Receiving End Termination

The following figure shows an example of shunt termination used with ECL gates.




Figure 41. Shunt termination.

The terminating resistance  $R_T$  is normally connected to a voltage between -2 and -3 volts, with the exact voltage determined by the gate characteristics and the terminating resistance value. The termination can also be made from a two-resistor thevanin equivalent which eliminates the need for an extra supply, but dissipates considerably more power. Parameters that contribute to signal degradation with shunt termination include tolerance on transmission line impedance, input capacitance of the receiver, and impedance of the terminating resistor.

The effects of receiver capacitance are different than for the series terminated case, since a full amplitude wave is propagated from the driver to the receiving end, and ideally, terminated by the resistor at the receiving end. The receiver input capacitance causes a reflection until the capacitance is charged, and this sends a reflection back to the driver as shown below. Of course, the finite rise time of the signal will round the reflected wave considerably. The pulse reflected by the receiving end is




Figure 42. Incident and reflected waveforms for a shunt-terminated line.



transmitted back to the sending end, and is reflected out of phase by the low impedance of the driver and sent back towards the receiver. When this reflection reaches the receiver, it is in phase with and aids the signal transmitted originally. Thus, the reflections do not directly degrade the signal at the receiver; however, the returned reflection would affect a receiver located other than at the far end of the line. The reflections from receiver capacitance will also degrade the signal if a short pulse is transmitted so that the reflections from the leading edge are present after the trailing edge has reached the receiving end, as shown in Figure 43.




Figure 43. Signal degradation due to reflections from receiver capacitance.

Thus the effect of these reflections must be controlled by limiting the capacitance at the receiving end, controlling the minimum width of signals transmitted, or allowing sufficient propagation delay to attenuate the reflections.

Capacitance at the receiving end causes less rise time degradation for shunt termination than for series termination, since the equivalent source



resistance driving the capacitance is the line impedance in parallel with the terminating resistance, instead of just the line impedance as in the series-terminated case.

Shunt termination requires less power dissipation than series termination since the return voltage is less, but shunt termination requires an additional supply voltage. The termination supply voltage must be well bypassed to ground so that the terminating resistance will properly terminate the line. Loads can be distributed along the line, and their input capacitance used to calculate an effective transmission line impedance. This can be extended further than in the series-terminated case, since the loss in the high DC level due to receiver input current is not as great. The driver must be capable of supplying sufficient current to drive the terminating resistor, and this is greater than the capability of MECL II for reasonable values of terminating resistance. Shunt termination was chosen for use in the restructured macromodules because of the reduced power dissipation, the relatively high input current (265  $\mu$ A maximum) of the ECL 10K circuits, and the high drive current capability of ECL 10K. Another possible approach would have been to use 1000 $\Omega$  pulldowns for short lines (less than 2 or 3 inches), series termination for long lines with fanout limited to 1 receiver, and shunt termination for clocks and gating signals that require significant fanout and moderate length. If the shunt-terminated cases could be kept few enough, they could be made from thevenin equivalents and the need for a termination power supply eliminated. 1000 $\Omega$  pulldown resistors draw less than the rated current from the outputs, and therefore reduce the low-level noise margin slightly. The low-level noise margin for 10K ECL is initially 155 mV, and the high-level noise margin is 125 mV. Use of 1000 $\Omega$  pulldown resistors would give about 125 mV noise margin for both high and low levels. If most of the lines can use 1000 $\Omega$  pulldowns instead of terminations, the power dissipation would be about the same as the situation with all shunt termination and a termination supply voltage. The disadvantage of this method is that three different methods of termination would be in use, and the utilization of resistor networks containing the resistors would be poor because of the variety of types. Another problem would be the possibility of increased crosstalk, as discussed in the section on crosstalk.

#### 4.3.3 Double-Ended Termination

Both sending-end and receiving-end termination can be combined, as shown in Figure 44.




Figure 44. Double-ended termination.



This provides termination at both ends; however, the signal amplitude at the receiver is one-half that which is sent by the driver. Thus, either a stronger driving signal or a more sensitive receiver is required with this termination scheme. The drive current requirements for the sending end are basically the same as for the series termination case discussed before. The driver (or pulldown) resistor must be able to supply enough current for the initial reflection from the capacitance at the receiver, as in the series-terminated case.

The values of  $R_1$  and  $R_2$  are determined by the line impedance and by the driver output impedance, but some freedom exists in choosing  $R_3$  and the terminating voltage  $V_T$ .  $R_3$  may be omitted altogether if  $V_T$  is low enough to bias the driver on under all conditions, but care must be taken to insure that the voltage drop across  $R_1$  does not shift the level out of the specified range. Because of the reduction in signal amplitude this scheme is not practical as it stands, but it can be considered as one-half of a differential terminating scheme as shown in Figure 45.




Figure 45. Differential terminating scheme.

If the two outputs of the driver are exact complements, the circuit can be analyzed by considering each half independently and the midpoint of the receiving end termination to be connected to the average value of the high and low output voltages. For the twisted-pair cable used by the macromodules, the characteristic impedance is nominally 125Ω, and the resistor values shown in the figure are used. The equivalent circuit for one side of the circuit is shown in Figure 46.

The current driving requirements are similar to the series-terminated case discussed earlier, since the signal swing on the transmission line is one-half normal and the impedances are one-half those discussed for the series-terminated case. The capacitance at the receiving end must be limited; otherwise, the reflection of a negative-going signal from the receiver may draw more current than the 750Ω pulldown resistor can sink, and thereby allow the signal to be reflected from the driver back to the receiver.






Figure 46. Equivalent circuit for one-half of differential terminated line.

Driving both inputs of the receiver with a differential signal compensates for the one-half amplitude in each signal, and in the typical case provides improved noise margin since there is no  $V_{BB}$  tolerance as in the single-ended reception. The reduced amplitude of the signals improves the positive common mode range of the receivers, since a greater common mode signal is required before the receiver input transistor saturates. Note that if the two outputs of a differential driver are not exactly complements, a common mode signal is generated which drives the two sides of the twisted pair in phase. This generates a common mode signal with the equivalent circuit shown in Figure 47. The two series-terminating resistors in parallel provide a somewhat mismatched termination for the common mode signal, and several round trips on the line are necessary before the signal dies out. Some of the early experiments used a bifilar wound choke at the output, which provided a low impedance for the differential signal and a high impedance for the common mode signal. Fortunately, the common mode signal generated at the driver is usually much less than 100 mV, and for this reason, the difficulty of fabricating and mounting the chokes, and the possibility that the chokes would introduce resonant circuits, they were not included in later work.




Figure 47. Equivalent circuit for common mode signal.



Differential transmission schemes that use only series or only shunt termination are also possible, as shown in Figure 48.




Figure 48. Differential transmission scheme using (a) only shunt termination and (b) only series termination.

Both of these methods transmit about full signal swing, which does not allow quite as much common mode range. In addition, the shunt-terminated scheme has greater problems with the reflection from the receiver capacitance, since there is no termination at the transmitting end. The current from the drivers exceeds the specifications of MECL II, since they are driving the equivalent of a  $65\Omega$  line.

The series-terminated design can be modified, as shown in Figure 49, to drive a half-amplitude signal within the capabilities of MECL II. This scheme was used in some early experiments. It requires the same number of resistors for terminating as the combined series shunt method. However, if the inputs to an unused receiver are to be held to a known state, two resistors are required (one for each input); while the series shunt scheme requires only one resistor for this purpose, since the receiver inputs are effectively tied together by the shunt-terminating resistor. The series-terminating scheme has the advantage that parameter plugs or other non-ECL or differential signals can drive the high input impedance of the receiver much more easily than the  $130\Omega$  terminating resistance of the series shunt method.

Series-shunt terminated twisted pairs as shown in Figure 45 are used






Figure 49. Modified series termination.

for all cable connections in the macromodules. The differential noise margin is about 200 MV, which is more than adequate.

The worst-case crosstalk from all other pairs to a single pair in the data cable is less than 60 MV for a 20-foot cable, and the maximum attenuation (for minimum signals) is less than 15 MV. A maximum length of 20 feet has been established, and no cables are allowed beyond this length. The connector configurations prevent cables from being cascaded without amplification. An early version of the control cable had a male connector on one end and a female on the other end. This prevented the logical error of connecting two control outputs or two control inputs together, since one end of the cable would plug only into inputs and the other end would plug only into outputs. Unfortunately this allowed control cables to be cascaded without intervening amplification; so the cable design was changed to have the same connector at both ends, since the elimination of electrical or loading errors was a requirement, but the elimination of possible logical errors is only desirable.

#### 4.4 BIDIRECTIONAL BUS FOR MEMORY

A bidirectional bus was developed for use in transmitting data below the General Memory Controller module, since the addition of the 12 address bits to the normal up-down signals would have exceeded the pin capacity of the connectors. The bidirectional bus is used to transmit data up or down, depending on whether the operation is a read or write. It operates with series termination when driving down and shunt termination when driving up. Figure 50 shows the connections used for one section.

When the bus is used to transmit data down, the output of gate B is held low, and gates A1 and A2 drive the series-terminated transmission line. When data is to be transmitted up, the outputs of gates A1 and A2 are held low, and gate B drives the line as a shunt-terminated line with the terminating resistor return voltage determined by the low-level output voltage of gates A1 and A2 with the 375Ω pulldown resistor. Two drivers are necessary,



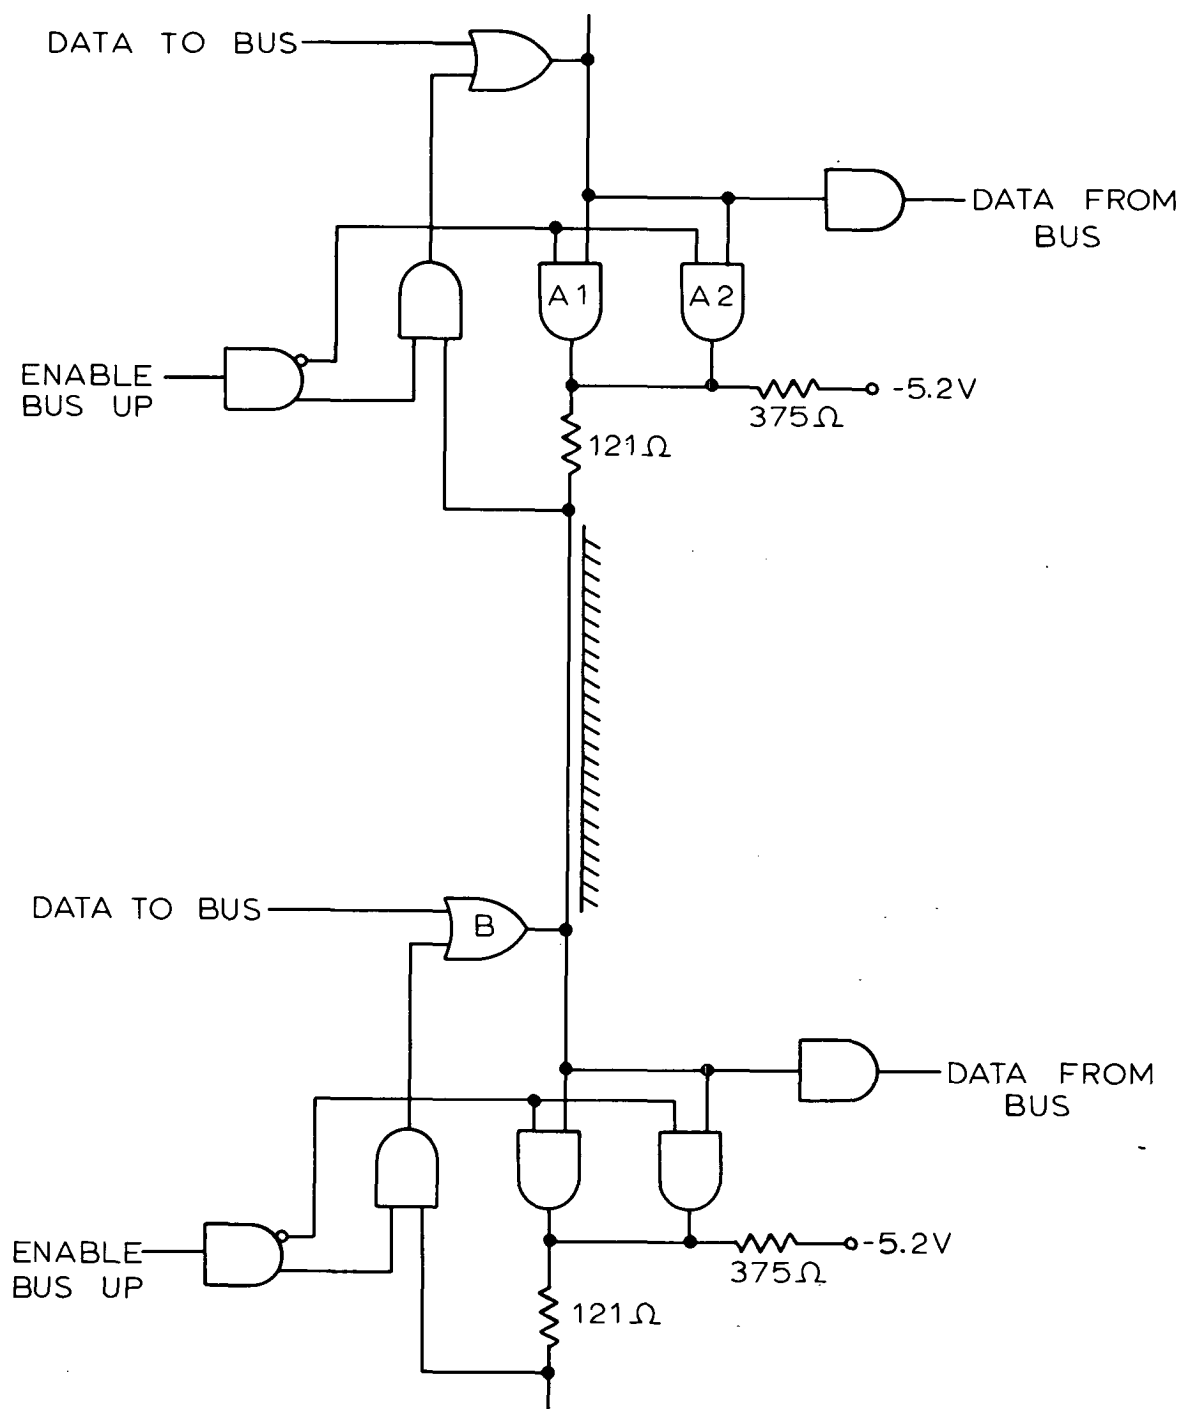


Figure 50. Connections used for one line of the bidirectional memory bus.



since one driver could not drive the  $375\Omega$  pulldown resistor required to sink all of the current that might be supplied through the terminating resistor. Several nonstandard specifications must be applied to gates used in this configuration to assure valid operation.

The minimum low level for A1 and A2 and the maximum high level for gate B are constrained to  $-1.700\text{ V}$  and  $-750\text{ MV}$  respectively, so that when gate B is driving the bus high, the pulldown resistors will conduct enough current to keep the output impedance of A1 and A2 low. Package B is tested for low output voltage with  $.5\text{ mA}$  load current, and for high output voltage with  $7\text{ mA}$  load current. In addition, because of the large number of loads on the receiving side of the series terminating resistor when A1 and A2 are driving high, the input current to the gates used as receivers is tested so that their total input current will not exceed  $150\text{ }\mu\text{A}$  with inputs high.

With these additional tests, which are an annoyance, and are required because of the low drive capability of MECL II, the bus is properly terminated for data transfers in either direction, and requires only one gate delay per module for signals propagating down and two gate delays per module for signals propagating up.

#### 4.5 CROSSTALK

Adequate separation must be provided between P.C. board lines to keep crosstalk to acceptable levels. Figure 51 shows forward and backward crosstalk levels versus line separation and length for the board thickness and line width used for microstrip lines in the macromodules. The maximum length of connections between modules is about 3 feet, and based on the desire to keep crosstalk below  $80\text{ MV}$  the separation was restricted to  $.100''$  minimum. The worst case is the backward crosstalk from the wave reflected from the receiver, since the backward crosstalk is larger, wider, and occurs later. It may appear that series termination offers an advantage over shunt termination, since the transmitted signal is only one-half amplitude. Conversely, it may appear that shunt termination offers an advantage, since coupled signals are terminated at the receiving end while for a series terminated line they are reflected and doubled in amplitude. Actually the two cases are equivalent. In the series-terminated case (assuming all lines propagate in the same direction) the backward crosstalk propagates back to the driver of the coupled line, is reflected toward the receiver by the low driver impedance, and is terminated there. In the series-terminated case the backward crosstalk from the forward wave is terminated at the driver, but the backward crosstalk from the reflected wave propagates to the receiver of the coupled line where it is doubled by the reflection. It is finally terminated at the driving end. Since the coupled signal is one-half as large as the shunt-terminated case and is doubled at the receiver by reflection, the effective crosstalk amplitude is the same.

The data cables contain twelve twisted-pair paths for data. Because of the terminations at both ends, the backward crosstalk between data lines is not a problem; it is absorbed by the sending end terminations. The data delivery return signal, however, propagates in the opposite direction from the data and is given a unique pitch, or number of twists per inch, in order






Figure 51. Forward and backward crosstalk levels vs. line separation and length, for the board thickness and line width used for microstrip lines in the macromodules.



to minimize the crosstalk coupled to it. More details of the cable construction are given in Section 5.5 of this volume.

#### 4.6 DELAYS

Because of the unclocked nature of the logic design, delays are necessary to allow sufficient settling time for logic in many sections of the macromodules. Although cascaded logic gates are used for delays in some cases, they are not attractive for delays of more than a few nanoseconds because of the ratio of maximum to minimum propagation delays. For longer delays, series RC elements with gates as buffers have generally been used as shown below. Delay lines were generally avoided because of their cost, size, and the large number of delay values that would be required.




Figure 52. Series RC delay circuit.

The output of the RC element has a very slow rise and fall time, and the second buffer is used to reshape the signal somewhat. Due to the relatively low voltage gain of ECL (about 6 to 8), the output of the buffer may not have normal rise and fall times. Another disadvantage of this delay element is that unless a significant recovery time (several times the delay value) is allowed before the next initiation, the delay value will be reduced. Also, the delay value depends on the driver levels and receiver thresholds, and can vary considerably depending on the particular gate characteristics. Furthermore, the delay value may be different for positive and negative-going signals, if the threshold voltage of the receiver is not midway between the high and low output voltages of the driver. Because of the possible delay variations with a given RC value, all delays are measured and verified before circuit boards are installed in a module.

Where unsymmetrical delay values are desired, the series resistor is omitted and the delay is generated by the time required for the driver pulldown resistor to discharge the capacitor. In the positive-going direction, the output emitter follower supplies enough current to charge the capacitor relatively quickly.



An alternative design that overcomes some of the deficiencies of the series RC delay is shown in Figure 53.




Figure 53. Feedback capacitor delay circuit.

With this arrangement, the delay is longer than normal if full recovery has not been reached, and the positive feedback causes sharp rise and fall times on the output. The disadvantages are that the input to the second gate is driven into saturation on a positive-going signal, and there is some change in output value before the gain becomes high enough to cause regenerative action. Some newer receiver types are designed with increased common mode range and could be used without danger of saturating the input transistor. The lack of recovery time requirement would reduce module operating time in many cases, since at present the delay value must be set longer than the minimum allowed in case the module is initiated twice in succession.

#### 4.7 NOISE MARGIN VS. TEMPERATURE

MECL circuits are specified at three distinct temperatures, and the specified input and output voltages vary with temperature. For MECL II the DC noise margin between two circuits at the same temperature and a -5.200 volt power supply is 175 MV. There is a loss of 1.5 MV in  $V_{NH}$  (high-level noise margin) and 1.3 MV in  $V_{NL}$  (low-level noise margin) for each degree centigrade difference between connected circuits. In a macro-modular system, the worst case occurs when a cold module is placed in an operating system and must communicate with a warm module. The maximum outlet air temperature from a module is about 50°C., and the junctions in a package dissipating 125 MW with a thermal resistance of .1°C/Watt would be 12.5°C. above the outlet temperature, or at 62.5°C. If modules are stored at 20°C., the loss in high-level noise margin is 64 MV, and the loss in low-level noise margin is 55 MV. This loss in noise margin occurs only for single-ended signals transmitted between adjacent modules. Only the common mode range of differential signals transmitted on cables is affected, since both halves of the signal are affected equally.



#### 4.8 POWER SUPPLY SENSITIVITY

The MECL input and output voltages are a function of power supply voltage, and there is a loss of about 7 MV in  $V_{NH}$  and 18 MV in  $V_{NL}$  for each percent of power supply error allowed.

The module power supplies are allowed a  $\pm 5\%$  deviation from their nominal value of -5.200 volts and in the case of two adjacent modules this represents a possible loss of 35 MV for  $V_{NH}$  and 90 MV for  $V_{NL}$  for signals between the two modules. As in the case of temperature difference, the only effect on cable connections is on the common mode range.

#### 4.9 AN EXAMPLE

As an example of a limiting constraint on system size, consider a structure composed of  $n+1$  fully loaded pedestals as shown in Figure 54. Each pedestal has 128 data cables (one per cell), each 20 feet long, connected to the next higher numbered pedestal, and a single 20-foot control cable from pedestal  $n+1$  to pedestal 1. The maximum DC current in each cable is 1.3 mA, and the DC resistance of the ground return (shield) is  $.1\Omega$  per cable.




Figure 54. Macromodular configuration for example.

For the  $128$  cables between two pedestals, this gives a total of 166 mA and  $7.8 \cdot 10^{-4}$  ohms. The resistance of a single control cable shield is 0.125 ohms, and the voltage between the  $n+1$ st pedestal and the 1st pedestal is given by

$$166 \text{ mA} \cdot \frac{(.125)(n \cdot 7.8 \cdot 10^{-4})}{.125 + n \cdot 7.8 \cdot 10^{-4}}$$

The limit as  $n$  approaches infinity is 166 mA times  $.125 \Omega$  or about 21 mV, and for  $n=10$  is only 1.2 mV. Of course each pedestal in the above configuration could be replaced by a group of pedestals to allow more data cables



to be connected, thereby developing higher potentials across the control cable. It seems, however, that such calculations are moderately ridiculous, and that the probability of failure due to the DC signal currents can be well approximated by zero. The same configuration can be used for the A.C. case. If we assume a differential unbalance of 50 mV in the output of each driver and all outputs switching in phase, the equivalent circuit for a data cable is a 50 mV signal driving a  $3.5\Omega$  line (considering the twisted pairs lumped together to form the inner conductor, and the cable shield the outer conductor, of a coaxial cable). The equivalent series-terminating resistance is about  $2\Omega$ , formed by the parallel combination of the individual twisted-pair terminating networks. The 50 mV signal will propagate to the receiving end and be reflected to the sending end, where it will be partially terminated by the  $2\Omega$  series-terminating resistance. The reflection factor will be

$$\frac{2-3.5}{2+3.5} = .272$$

If the data outputs complement each time the reflection reaches the sending end, successive signals will add in phase and give the following maximum transmitted voltage:

$$50 \text{ mV}(1 + .272 + (.272)^2 + \dots) = 69 \text{ mV}.$$

The 69 mV induces a current of  $69 \text{ mV} / 3.5\Omega = 19 \text{ mA}$  in a cable. The total AC current in 128 cables is 2.45 amps.

The effect of a high-frequency voltage between two ends of a control cable is greatly attenuated by the coupling between the shield and the center conductors of the cable. At the driving end, the shield and the center conductor are coupled together by the driver. A high-frequency voltage between the two ends of the shield is coupled from the shield to the center conductors by the capacitance of the cable and the voltage difference between the center conductors, and the shield at the receiving end is reduced by a factor of about 100 from the voltage difference between the cable ends. Thus a 50-volt signal would be required between two base pedestals to cause a 0.5-volt common mode signal. The impedance between pedestals is not known, but two limiting cases can be considered. If the impedance is zero and 69 mV is developed between each pair of pedestals, the 0.5 V common mode signal in the single control cable would require about 730 pedestals each with 128 data cables to the next one, each signal of each cable with 50 mV common mode difference, and all switching in the proper phase so the voltages would add to produce the maximum voltage across the single control cable between the end pedestals. If the impedance from pedestal to pedestal is finite, then an infinite chain of pedestals would develop a 2.45-amp signal. The peak common mode error voltage induced in a control cable by a 2.45-amp peak-to-peak 20 MHz signal has been determined experimentally as about 2.5 volts. It is less than the voltage induced from one end of the cable to the other, since the voltage on the cable shield couples to the signal conductors, thus reducing the difference between shield and signal conductor at the receiving end. Of course, with a finite impedance, more than 730 pedestals would be required to generate a 0.5 V common mode error.



#### 4.10 INTRODUCTION TO ELECTRICAL TESTING AND INSPECTION

The goal of testing and inspection is to reduce the total cost of a project by reducing costs in the use and application phases by a greater amount than is expended in testing and inspection. It is (fortunately) not necessary (and probably not possible) to find all faults and problems during or before assembly and testing; but only a few should remain after that point, since they can become very difficult to isolate and identify when embedded in a system of many modules. Faults that remain after assembly and testing are due to intermittent faults that do not occur during testing, marginal or other conditions that are not tested, wearout or failures occurring after assembly and testing, and testing errors. Intermittent faults are a problem to locate in any system and are frequently due to mechanical problems with connectors, or are due to marginal operation of electronics which may then be susceptible to temperature, power supply variation, or induced noise. The cure is to use, as extensively as possible, connectors that are reliable and have sufficient tolerance and assembly constraints to minimize marginal conditions. In electronics the cure is to test components to try to insure enough margin to operate under all temperature, power supply, interconnection, and timing conditions.

It is difficult to determine what expenditure of money and effort for testing and inspection will be the best choice. Some intangible factors are involved, such as user frustration and confidence, in addition to problems of predicting the number of undetected faults versus expenditure.

The philosophy adopted by CSL was to test incoming IC's for margins and to test assemblies for function. There are several reasons for this. If qualified and used correctly, most marginal problems (not intermittent) are due to electronic components out of tolerance, or in some cases to improper assembly (components interchanged, etc.). Except for those that degrade with time, these out-of-tolerance elements can be identified and rejected at the component level. To identify them at the circuit level, after assembly, is very difficult, since many of these are not readily accessible. For this reason only functional tests are performed on the assembled modules, and component tests are relied upon to identify marginal units. Even functional tests require many combinations and significant time to perform. If additional parameters were required to be varied, the complexity of the testing apparatus would be drastically increased and the test time would also be drastically increased.

The three steps taken to try to achieve reliability are: evaluate components based on their design and construction to insure that they are really suitable for their intended use; restrict the design so that components are used within their intended ratings and specifications; test and inspect components and assemblies to eliminate or repair faulty components and assembly. Most of the attention given to testing and inspection was devoted to the ECL integrated circuits and the P.C. boards, since they are the two vulnerable components that are in widespread use and whose careful testing and inspection offers a saving in total effort.



#### 4.11 EVALUATION AND QUALIFICATION OF ECL CIRCUITS

Emitter Coupled Logic circuits possess a number of peculiarities, some of which are described by the manufacturer and a number of which are not. In addition, some circuits are not specified by the manufacturers under realistic worst-case conditions, due either to the difficulty of performing the tests, or more likely, due to the lowered yields under those conditions. We have also found a number of cases where IC's were shipped to us that were faulty in some manner that might not be identified in testing, and about which Motorola made no attempt to notify us when they became aware of the problem. There also have been cases where the design of a particular IC was changed without any notification or information released by the manufacturer. The change may have been intended to be an improvement, but customers are unaware that there is more than one circuit with the same part number. In addition, even if the customer manages to learn of the different designs and their effective date, Motorola has taken old stock returned from a distributor and remarked the date codes, thereby giving a new date code to an old device, possibly unsuitable due to design errors. Even a contract clause requiring notification of mask changes might not, in practice, protect against altered date codes.

Early plastic packages used by Motorola had very troublesome characteristics. Five or ten thermal cycles would cause a few percent of the packages to fail, and further thermal cycling would cause additional failures. It did not appear possible to perform enough thermal cycles to eliminate potential failures, so our decision was to use only ceramic-packaged circuits. One of the early ceramic packages (white ceramic) also had the problem that when the leads were stressed to insert the package in a socket or circuit board, the top might pop off. We have had no significant problems with the gray ceramic packages which we commonly use except that they are sometimes difficult to obtain in the commercial temperature range. The plastic and epoxy packages have undoubtedly improved since our early experience, but we have not felt that the risk of significant numbers of package failures was worth the small price difference at this point in the project.

##### 4.11.1 Feedthrough

One of the ECL peculiarities never mentioned by the manufacturer is feedthrough, which is undesired coupling from inputs to outputs that can cause short pulses up to half the logic swing in amplitude to occur at the circuit outputs. As an example where feedthrough can be a problem, consider the circuit shown in Figure 55. Data is available from two sources, and the select line determines which of them is selected for gating into the flip-flop. If the AND gates suffer from feedthrough, changes in the data applied to the unselected AND gate may set the flip-flop incorrectly. Note that only feedthrough from the AND gate inputs labelled D is of concern, since the flip-flop clock signal is off when the C input to the AND gate changes. Another case where feedthrough must be prevented is shown in Figure 56, where the flip-flop output is to remain undisturbed when the clock changes and the input level is the same as the flip-flop state. One application of this circuit is a conditional pause for transition logic signals.






Figure 55. An example where feedthrough may incorrectly set a flip-flop.




Figure 56. An example of feedthrough in a flip-flop.



In this application, feedthrough from the clock input might appear as two transition signals to circuits following the flip-flop, although there was no change in the data input level.

There are a number of causes of feedthrough, most of them peculiar to ECL, and in some cases feedthrough has seriously restricted our use of many ECL circuits, thereby requiring increased package counts and/or propagation delays in circuits. In one particular IC which contained four flip-flops, feedthrough on internal stages was severe enough that some flip-flops occasionally dropped "1's" when the clock was turned off. The "1's" dropping was influenced by power supply voltage, clock pulse width, clock rate, clock rise and fall times, and temperature. Of course many of the circuits would pass acceptance tests and then fail, perhaps intermittently, in use. This problem was predicted, based on the circuit diagram and our studies of feedthrough, before any of the packages were placed in service, but unfortunately after some boards were manufactured that included those packages. The boards were redesigned to use a different flip-flop package, since the manufacturer's lead time to redesign and produce suitable circuits of the original pinout was much longer than could be tolerated. Section 4.16 gives an explanation of the causes of feedthrough and an estimate of the effects of each cause. Section 4.17 gives the usage restrictions that we have applied to MECL II circuits in order to limit the effects of feedthrough.

A different sort of problem occurred with another package type in which a significant number (several percent) were rejected at incoming testing because one or more inputs were shorted to ground. Investigation showed that the bond wires to the inputs were lying almost parallel to the chip surface and contacting the ground metallization in some cases. Several thousand circuits were returned to the manufacturer in exchange for ones with a better bonding technique or ones with a redesigned mask that eliminated the close proximity of the bond wires and ground metallization.

In evaluating 10K ECL for use in restructured macromodules, a number of plots of output voltage versus input voltage and of output current versus output voltage were made. Some outputs exhibited a peculiar characteristic at low currents that could be modeled as a normal output in parallel with a resistor to ground. According to the manufacturer, this is due to small channels or pipes that connect the collector and emitter regions of the output emitter followers together. This has implications in cases where several outputs are ORed together, as in the bus, since the outputs cannot be depended on to follow normal transistor characteristics at low currents. In order to help insure proper bus operation, a test of output voltage at low current (100  $\mu$ A) has been added for 10K ECL gates.

#### 4.11.2 Production Testing of Integrated Circuits

Our production testing consists of measuring the DC parameters of input current and output voltage with worst-case load current and input levels. The tests are performed using an ALMA model 480B bench top tester. This has generally proven to be flexible enough for the tests we perform, since our volume does not justify a more elaborate or flexible computer-controlled



tester. With MECL II circuits this guarantees a noise margin of 175 mV with 25°C temperature, -5.200 volt power supply, and worst case DC loading. In order to increase yield, manufacturers have changed specifications on some circuits to allow less noise margin (0 volts in some cases), but our policy has been to test to the standard levels regardless of the individual specifications for a circuit. In a few cases our yield has been so poor (or non-existent) that we have had to examine the particular circumstances of use for circuits and to relax the tests somewhat. This was never done as a blanket policy, but only based on the requirements of a particular application.

Another source of problems was the fact that Motorola's test specifications call for worst-case levels to be applied to only one pin at a time. In high fan-in gates, the difference between only one pin and all pins at worst-case low levels can be significant. Since in most applications it is the large number of outputs, not inputs, that are used, the circuits were tested in two categories and so marked and used. One type could be used in all applications and the other could be used only where a single input was required.

Differential input receivers are used for receiving cable inputs because of their common mode noise rejection. The manufacturer specifications for some receivers call for measuring bias voltage provided by the package and then testing the receivers with one input pin tied to the bias voltage. Because of the wide tolerance allowed on the bias voltage, the differential sensitivity and common mode range of the receivers is not very well tested by this method. Therefore, our tests include a specified differential voltage at both high and low levels to provide a test of differential sensitivity over a reasonable common mode range.

No AC tests have been performed on the circuits as part of the production testing. This is based on our experience in evaluation of ECL circuits at CSL, where practically no circuits have been found with propagation delays that are outside the specified limits. In contrast, our DC tests typically find about  $\frac{1}{2}$  of 1% bad packages (open pin, output stuck high, etc.) and 5-10% marginal failures (insufficient noise margin, excessive input current, etc.). Some of these rejects are caused by difficulties in correlating our tests with the manufacturer's tests.

The other components in the macromodules, except for the power supply, are resistors and capacitors, and these are not tested based on the extremely small number of defective ones that we have found.

#### 4.12 PRINTED CIRCUIT BOARD TESTING

P.C. boards are inspected under low-power magnification for shorts, breaks, and raised lines. This is a very important inspection, since a large fraction of our defects with assembled modules is due to shorts from poor etching and solder slivers on the P.C. boards.

The only operating test typically performed on boards before assembly into a module is to measure the value of delays established by RC networks.



The tolerances are such as to make worst-case design unreasonable for these delays, so they are checked by measurement.

#### 4.13 OTHER TESTS OF COMPONENTS

The critical components, or those that are suspect or have given problems in the base pedestal and module power supply cards, are tested before assembly. After assembly the power supply cards are tested at worst-case load and line conditions, and output ripple is checked to verify power supply operating margins.

Because of the high output current capability of the ECL packages, they can be damaged by short circuits to the power supply. Thus in some cases, individual boards have been powered for tests from a current limited supply and the power supply voltage measured. This allows detection of package outputs shorted to the power supply voltage, but does not allow enough power dissipation to damage the package.

Parts such as data cables, control cables, channel boards, and couplers are checked for continuity, short circuits, and transposed connections by special-purpose test equipment that we have designed.

#### 4.14 MODULE TESTING

The first test performed on each module is a test for power supply current from the 54-volt input. As the following graph for the Load module (Figure 57) shows, the current variation from module to module is extremely small, and a deviation of more than a few percent from the average value is regarded with suspicion. Possible faults detected by this current measurement are incorrect assembly, package output short circuits, and power supply failure.




Figure 57. Module load current vs. frequency of occurrence.



Some of the early modules were tested for function by operating them in conjunction with other modules and a LINC [2], but some combinations and/or sequences of operations cannot be generated and/or tested in this manner. For this reason a test module was developed for testing assembled modules that allows more complete tests and that greatly aids diagnosis of failures. It has a set of flip-flops that are controlled by a LINC computer and used to drive module inputs and a group of multiplexed inputs used to sense module outputs. Initially 4 high faceplate boxes were used to connect between the tester and the module under test, but a somewhat different mechanical framework was subsequently developed to allow easier access to the module under test and to the interconnecting wiring harnesses. Adapters with wiring harnesses allow the desired connections to be made between any module and the tester.

A table-driven program has been written for the LINC which applies test patterns to the module under test and compares the module outputs to the expected outputs. The test tables used to date have been manually generated, since the primary use to date of the test set has been on modules after assembly rather than after failure in use, and test generation procedures that are based on single failures and on stuck at high or low assumptions are not valid. A large percentage of the failures encountered after initial assembly are shorts between supposedly independent circuits of the module. The assumptions of simple failure and stuck at high or low would be much more reasonable for modules that fail in service after having operated correctly. Some work has been done on developing such test patterns for fault isolation in the COMPARE module, a particularly difficult one to analyze since it has only a single output reporting either success or failure of the comparison.

A program called PROBE-1 is available which allows an operator to control the inputs to a module via the LINC keyboard and switches and to observe the module inputs and outputs on the LINC display. This program has been extremely useful in isolating faults in modules and in initial debugging and verification of new module types.

#### 4.15 TEST RESULTS

Table 16 shows the results of the test and inspection program in terms of the number and types of faults that have been found after module assembly.

The modules constructed to date contain over 50,000 ECL integrated circuits and of these only 26 IC's have failed, 25 of these between incoming component test and module test, or about one bad IC in two thousand. Only one IC has failed thus far in service, although others may be out of tolerance but not bad enough to cause failure in their particular operating environment. Over 5-10 hours of IC operation have been accumulated in the three-year period since January 1971.



Table 16. Results of Test and Inspection Program

TOTAL NUMBER OF PRODUCTION UNITS CONSTRUCTED	667
NUMBER WHICH FAILED INITIAL TEST	250
NUMBER WHICH FAILED IN SERVICE	26

REASONS FOR FAILURE

SHORTS

ETCHING	28
SOLDER	55
METAL SLIVERS	22
OTHER	23

OPENS

BROKEN LINES	16
NO SOLDER/COLD SOLDER JOINT	30
OTHER	12

MISCELLANEOUS

IMPROPER COMPONENT	16
MISSING COMPONENT	16
INCORRECT INSTALLATION	16

IC FAILURE	26
------------	----

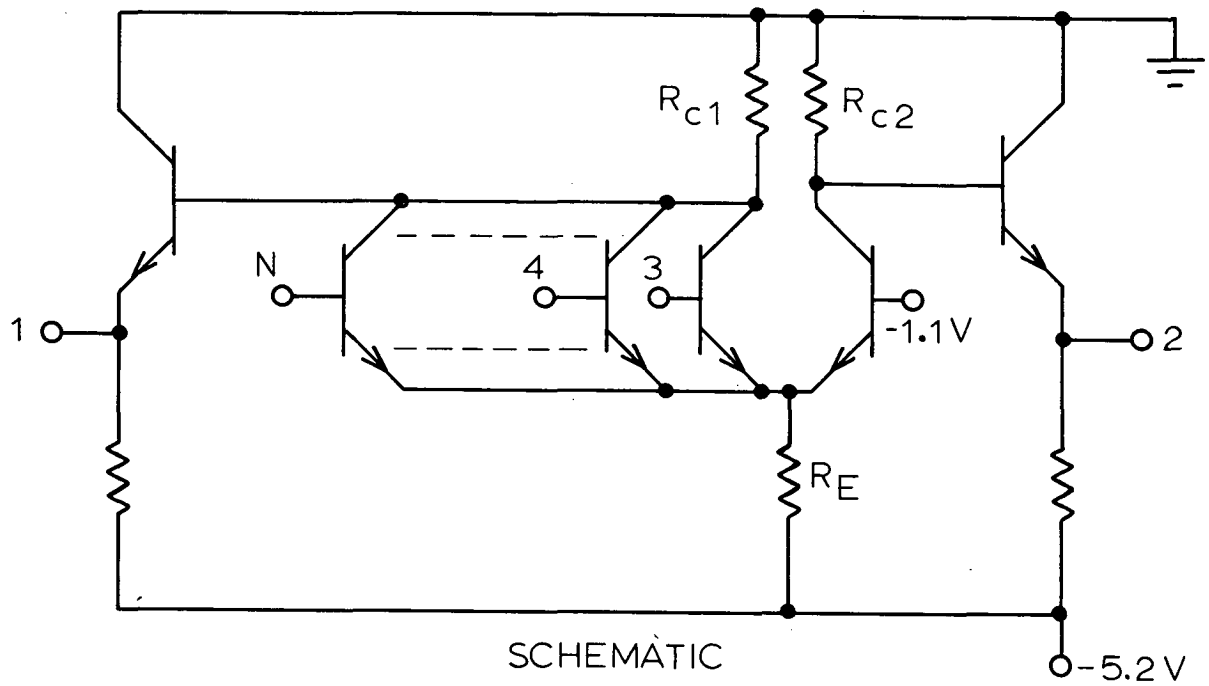
PUSHED IN OR MANGLED PINS	85
(Includes 43 damaged during testing by one bad faceplate box)	

The reliability of the macromodule electronics has been quite satisfactory. Most problems have involved connectors - particularly in faceplate boxes, since a number of these were put in service without testing due to immediate need and the fact that an automatic tester was not available.


4.16 APPENDIX I - CAUSES OF FEEDTHROUGH IN ECL CIRCUITS

Figure 58 shows a schematic for a standard ECL OR/NOR gate and Figure 59 shows one for an AND gate using two levels of current steering.  $R_E$  in Figure 58 acts as a current source of about 3 mA, and this current flows through either  $R_{c1}$  or  $R_{c2}$  depending on the levels of the two inputs. If any input is high the current flows through  $R_{c1}$ , causing the "NOR" output to be low, while if all inputs are low the current flows through  $R_{c2}$  and the output level is high. Thus the gate operates by "steering" the current through one





SCHEMATIC



LOGIC SYMBOL

Figure 58. MECL OR/NOR gate.



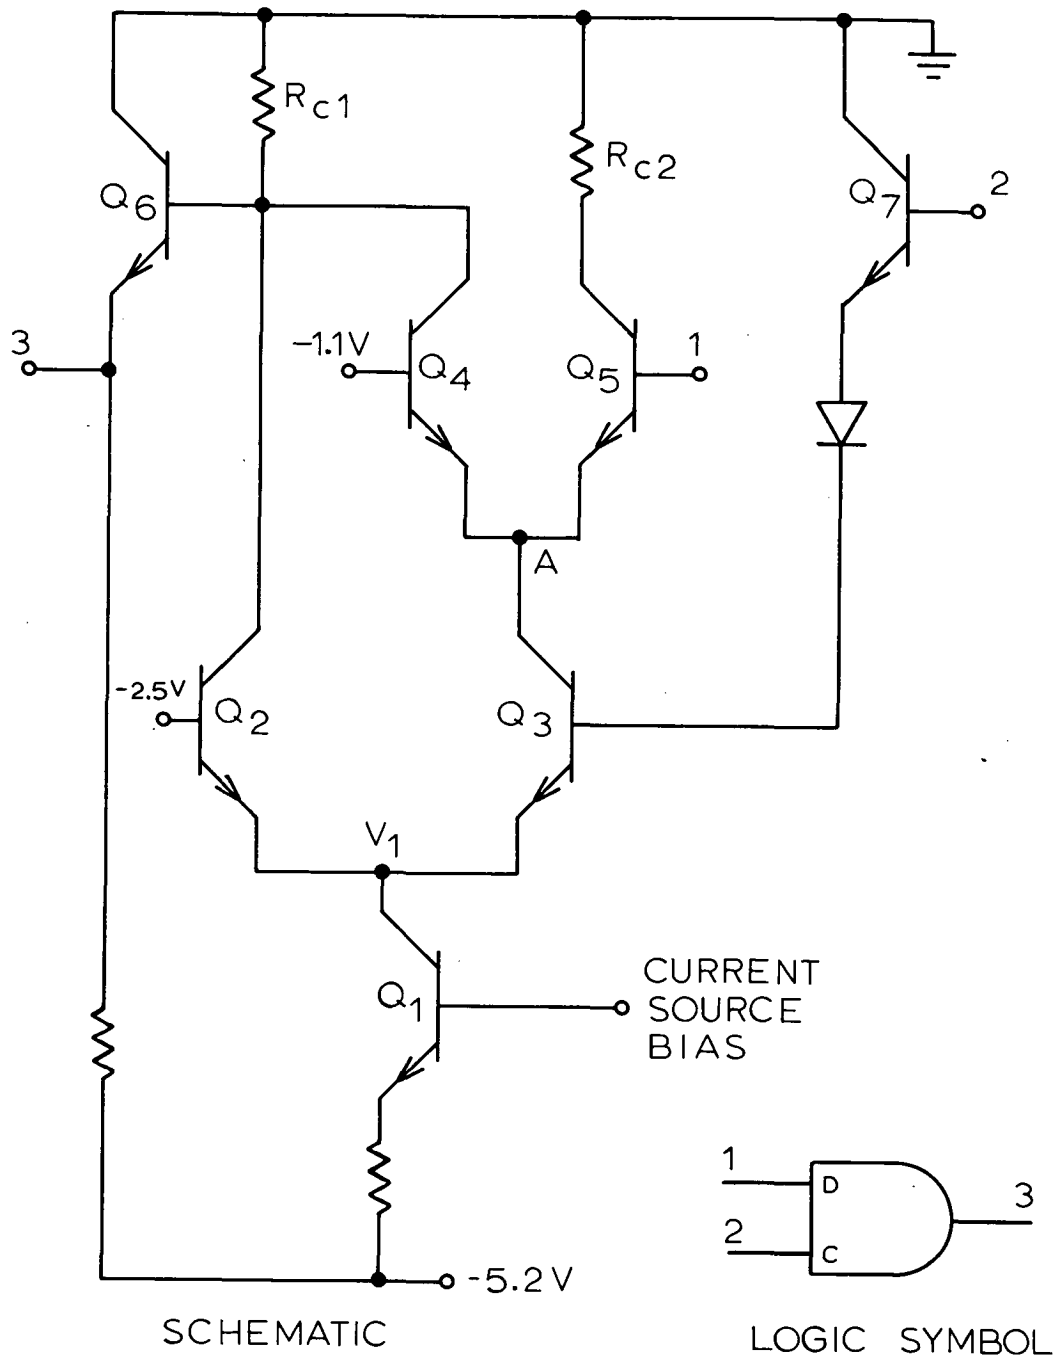


Figure 59. MECL AND gate.



collector resistor or the other.

In the "AND" gate of Figure 59 there are two levels of current steering and the output is high only when both inputs are high, thereby steering the current through  $R_{C2}$ . In this circuit a transistor is used to develop the constant current, since the percentage change in voltage across the current source is greater than for the OR/NOR gate.

In the following discussion we will ignore the switching time of the transistors and assume that they have a constant .7 volt  $V_{be}$  drop when conducting. The following junction capacitance values will be used and will be assumed to be independent of reverse bias:

$C_{cb}$	1.5 pf	(Collector to base capacitance)
$C_{be}$	1.5 pf	(Base to emitter capacitance)
$C_{cs}$	2 pf	(Collector to subtract capacitance)

#### Causes of Feedthrough

There are several different effects which cause feedthrough, and all of them are not present in every circuit. The following sections give explanations of the types of feedthrough and estimates of their effects.

##### 4.16.1 Collector to Base Capacitance

The most obvious cause of feedthrough is due to  $C_{cb}$  of the input transistors shown in Figure 60. If one input is a high level holding the NOR output low, the equivalent circuit for transitions applied to the other inputs is shown in Figure 61. The output voltage for an input ramp of  $\alpha$  volts/sec and a duration of  $\beta$  seconds is given by the following equation, where  $u(t)$  represents the unit step function.

$$V_o = C_1 \alpha R_{c1} \left\{ \left[ 1 - \exp \left( - \frac{t}{R_{c1}(C_1 + C_2)} \right) \right] u(t) - \left[ 1 - \exp \left( - \frac{t - \beta}{R_{c1}(C_1 + C_2)} \right) \right] u(t - \beta) \right\}$$

where:  $N$  = number of inputs

$M$  = number of inputs switched  
from low to high

$$C_1 = M C_{cb}$$

$$C_2 = N \cdot C_{cs} + (N - M + 1) C_{cb}$$

$\alpha$  = rate or rise of input  
in volts per second

$\beta$  = duration of input ramp  
in seconds



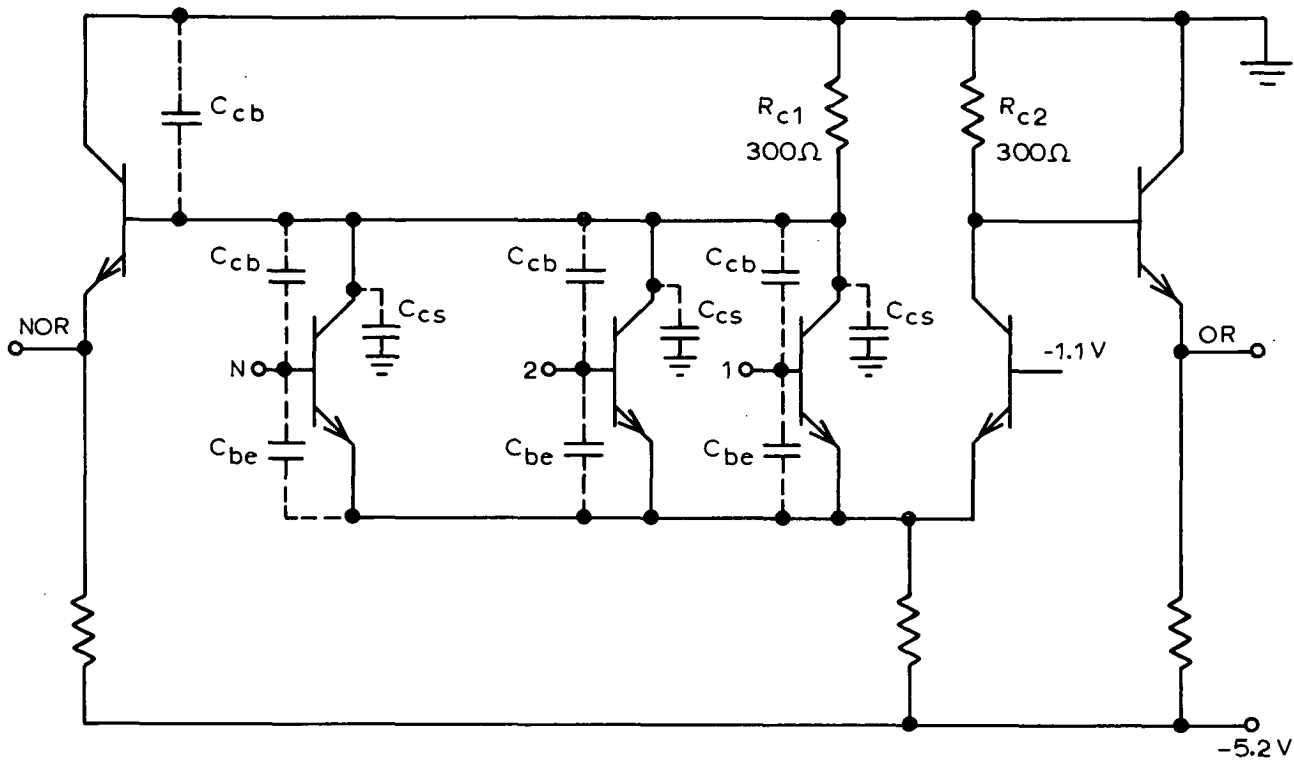


Figure 60. MECL OR/NOR schematic, showing junction capacitance.




Figure 61. Equivalent circuit for M input switching.



Substituting the appropriate values for a 4-input gate, with three of the inputs switched, into the above equations gives the following output voltage when

$$\alpha = 2 \cdot 10^8 \text{ and } \beta = 5 \cdot 10^{-9} \text{ (a one-volt transition in 5 ns):}$$

$$V_0 = 270 \text{ mV} \left[ \left[ 1 - \exp\left(-\frac{t}{6 \cdot 10^{-9}}\right) \right] u(t) - \left[ 1 - \exp\left(-\frac{t-5 \cdot 10^{-9}}{6 \cdot 10^{-9}}\right) \right] u(t-5 \cdot 10^{-9}) \right].$$

The peak amplitude of the output occurs at  $t = 5 \text{ ns}$  and is 154 mV. In a gate with larger fan-in, even larger signals could be coupled to the output of the gate.

#### 4.16.2 Emitter to Base Capacitance

Another mode of coupling in the gate shown in Figure 60 is due to emitter-base capacitance. The transistor whose base is held high acts as a common base amplifier to couple the signal from the  $C_{be}$  capacitance of the other input transistors to the collector resistor  $R_{cl}$ . If we assume that the common base amplifier has zero input impedance and a current gain of one, then the output current from the common base collector is simply:

$$I_{e_b} = \alpha C_3 u(t) - u(t-\beta),$$

where  $C_3$  is the total base-to-emitter capacitance. This current, applied to  $R_{cl}$  and the associated capacitance  $C_1 + C_2$ , gives the following output voltage:

$$V_0 = \alpha C_3 R_{cl} \left[ \left[ 1 - \exp\left(-\frac{t}{R_{cl}(C_1+C_2)}\right) \right] u(t) - \left[ 1 - \exp\left(-\frac{t-\beta}{R_{cl}(C_1+C_2)}\right) \right] u(t-\beta) \right].$$

Again, using a gate with a fan-in of four as an example, if three of the inputs are switched from low to high the maximum amplitude of the output pulse due to  $C_{be}$  coupling will be 154 mV.

By adding the output signal generated by  $C_{cb}$  coupling to that for  $C_{be}$  coupling, the total peak voltage at the output due to both  $C_{be}$  and  $C_{cb}$  is found to be 308 mV.



In this example the peak amplitude of the output pulse is considerably larger than the worst-case noise margin (175 mV) of MECL. Figure 62 shows a photograph of the output signal from a 4-input ECL gate operated under the conditions described. The output resembles the predicted waveform, except that it is somewhat rounded due to the bandwidth limitations of the output emitter follower and of the Tektronix 454 oscilloscope used. The amplitude of the pulse is, however, close to the value predicted by the simple model.

#### 4.16.3 Current Source Capacitance

Capacitance from the current source to ground can cause feedthrough in two-level circuits like the AND gate shown in Figure 59. If the D input is held low, the output level should remain low in spite of any changes in the level of C. However, as input C changes from low to high,  $V_1$  will change by about half of the logic swing. If the current source was ideal there would be no change at the output due to the change in  $V_1$ . Actually the current source is very good at DC, but is shunted by some capacitance which affects its high-frequency performance. The capacitance functions as an emitter peaking capacitor for  $Q_2$  and causes the transition on input C to be coupled to the output. A simplified equivalent circuit is shown in Figure 63; the calculated output waveform is:


$$V_o = \alpha C_e R_{c1} \left[ [1 - \exp\left(-\frac{t - \beta/2}{R_{c1} C_1}\right)] u(t - \beta/2) - [1 - \exp\left(-\frac{t - \beta}{R_{c1} C_1}\right)] u(t - \beta) \right].$$

For the capacitance values used previously, the peak output pulse amplitude would be 243 mV, a large output transient with only one input switched. Note that in this case there is no direct capacitance coupling to the output, and also that the output pulse is in the opposite direction from the input transition. Another interesting point is that the propagation delay and rise and fall time measurements commonly made for logic circuits may not be sensitive to capacitance at this point.

#### 4.16.4 $C_A$ Discharge

Another effect which can produce feedthrough is the discharge of the stray capacitance,  $C_A$ , from the node labelled A in Figure 59. If both inputs are low, then transistors  $Q_3$ ,  $Q_4$ , and  $Q_5$  are all off and conduct only leakage currents. This makes the voltage at node A indeterminate. It cannot be less than -1.8 volts or transistor  $Q_4$  would conduct, but it may be almost zero volts if the leakage current of  $Q_5$  is high enough. The capacitance from node A to ground will be charged to the voltage determined by the leakage currents. If input C is changed from low to high, transistor  $Q_3$  will begin conducting instead of  $Q_2$ , but transistor  $Q_4$  cannot conduct until  $C_A$  is discharged to -1.8 volts. The current for the current source will be supplied by  $C_A$  instead of by  $R_{c1}$  during this time, and a positive-going pulse will be produced at the output. The amplitude of the pulse will depend on the relative values of the leakage currents, and can vary widely from circuit to circuit. Ratios of two to one in output pulse amplitude have been observed in circuits within the same package.





Time Scale: 5 ns/div.  
 Upper Trace: Input, 500 mV/div.  
 Lower Trace: Output, 500 mV/div.

Figure 62. Output of a 4-input OR/NOR gate, showing feedthrough to the NOR output from three inputs.

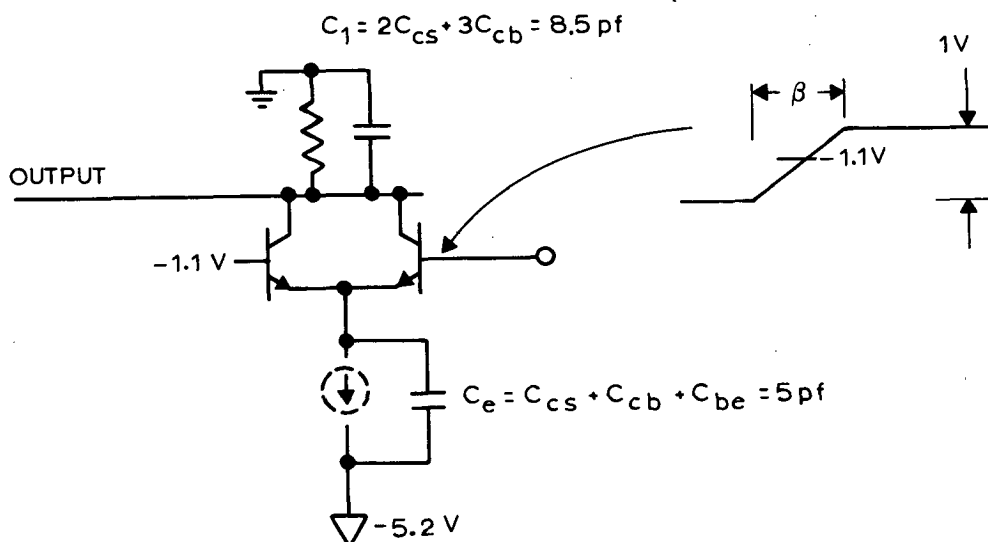


Figure 63. Equivalent circuit for AND gate with one switched input.



Assume that both gate inputs are initially low and that the C input is switched from low to high.  $Q_4$  will not conduct until  $C_A$  is discharged. The time required for this will be

$$t_o = \frac{C_A \cdot \Delta V}{I_E}$$

where  $\Delta V$  is the voltage to which  $C_A$  is charged and  $I_E$  is the current source value. Substituting 5 pf for  $C_A$ , 400 mV for  $\Delta V$  and 3.33 mA for  $I_E$  gives a value of  $t_o$  equal to 0.6 ns. The output pulse is then given by the following formula:


$$V_o = I_E R_{c1} \left[ \left[ 1 - \exp\left(-\frac{t}{R_{c1} C_1}\right) \right] u(t) - \left[ 1 - \exp\left(-\frac{t-t_o}{R_{c1} C_1}\right) \right] u(t-t_o) \right].$$

This gives a peak amplitude of -330 mV. A more complete calculation, taking into account the switching time of the transistors, would predict a pulse with the same total area but having a smaller peak amplitude. However, it is possible that in some cases  $C_A$  would be charged by more than 400 mV and thereby produce a larger pulse on discharge. It is interesting to note that in contrast to the capacitive coupling and emitter peaking feedthrough, the output pulse amplitude and shape are not a function of the rise time of the input signal. In a more complete analysis the output waveform would only be affected slightly by the input rise time. This occurs because there is one stage of gain between the input signal and the capacitance causing the feedthrough. Figure 64 is a photograph showing the feedthrough pulse from an AND gate. This large an output pulse appears to occur in about 3% of the small number of gates which have been investigated. Almost all other gates exhibit an output pulse amplitude about half this large.

There are two distinct differences between  $C_A$  discharge feedthrough and the ones discussed previously. First, it only occurs when the input transition is positive, while the other forms of feedthrough function with either positive or negative transitions. Second, the output pulse amplitude is not a strong function of input transition time or amplitude, since there is effectively a stage of common emitter gain between the input and node A. Also the peak output pulse amplitude is much larger for this form of feedthrough than any other with a single switched input.


To verify the effect of  $C_A$  discharge, a 100K resistor was added from node A of an AND gate to a variable power supply,  $V_x$ , and measurements were taken of peak pulse amplitude versus  $V_x$ . The results are shown in Figure 65.






Time Scale: 5 ns/div.  
Upper Trace: Input, 500 mV/div.  
Lower Trace: Output, 100 mV/div.


Figure 64. Feedthrough caused by an extremely bad case of  $C_A$  discharge.




$V_x = 0$



$V_x = -.50 \text{ V}$



$V_x = -1.5 \text{ V}$



$V_x = -2.5 \text{ V}$

All photographs taken with a Tektronix 454 oscilloscope, 5 ns/div.  
Upper trace, input signal, 500 mV/div.; lower trace, output signal, 200 mV/div.

Figure 65. Feedthrough with different values of bias applied to the second-level capacitance.



The area under each of the curves is almost proportional to the difference between  $V_x$  and the -1.8 volt level at which we assume  $Q_4$  starts to conduct. There is some feedthrough coupling due to the collector-to-base capacitance of  $Q_3$  and due to emitter peaking from the current source shunt capacitance. This is very apparent in the fourth photograph, where  $V_x$  keeps  $Q_4$  slightly forward biased at all times.

The voltage across the 100K resistor was less than a few millivolts, even with  $V_x = 0$ , indicating that very little current was required to supply the leakage currents of the transistors. The probe used to connect the 100K resistor to the logic circuit adds some capacitance, but this is small compared to the junction capacitances.

This experiment confirms the cause of the large feedthrough signals occurring for positive transitions of the C input. Some circuits in newer ECL logic families incorporate resistors of about 50K $\Omega$  from nodes such as A of Figure 59 to a negative voltage, to reduce the effects of this form of feedthrough.

If the C input is held low and the D input changed there is no feedthrough to the output due to  $C_{cb}$  coupling,  $C_{be}$  coupling, emitter peaking, or  $C_A$  discharge.


#### 4.16.5 Wired-OR Feedthrough

Another form of feedthrough is present in wired-OR circuits that depend on one of several output emitter followers to hold a line high despite switching of other outputs wire-ORed to the same line. If the emitter follower that remains high was originally conducting the major share of the current, there is little effect by the other emitter followers going low. However, if the emitter follower that is conducting most of the current goes low, some other emitter follower must assume the load, and the finite response time of the transistors causes a short low-going pulse at the output. This is the only form of feedthrough that affects a high output, and the maximum pulse output depends on the DC load current, the number of outputs, the response time of the output emitter followers, and the distribution of load current between the outputs. Figure 66 shows the effects of wired-OR feedthrough for two MC1010 outputs wire-ORed together, a 750 $\Omega$  pulldown resistor, and the high level of the switched output adjusted to be slightly higher than the other output to give maximum current transfer between outputs. The peak amplitude of the feedthrough is about 150 mV for this case and is about 80 mV for two outputs with equal high output levels (a likely case for two outputs in the same package) with a 1500 $\Omega$  pulldown resistor.

#### 4.16.6 Bias Network and Chip Capacitance Coupling

There is some coupling between circuits in the same chip due to the use of a common bias voltage source or sources, the common ground and power leads, and capacitance between elements on the chip. These effects appear to be significantly smaller than those discussed above. The worst case found was simultaneously switching three sections of a quad EXCLUSIVE OR and monitoring the fourth output. The signal induced in the output was less than 75 mV peak amplitude.





Top Trace: Input 500 mV/div.

Bottom Trace: Output 100 mV/div.  
Time Scale: 5 ns/div.

750 pull-down resistor  
two outputs of a 1010 package,  
wire-ORed

Figure 66. An example of wired-OR feedthrough.

Gates which are propagation delay and rise and fall time limited in performance by RC values and not intrinsic transistors characteristics should have the same peak value of feedthrough as a percent of logic swing regardless of the R and C values if  $C_1$  and  $C_2$  maintain the same ratio. The rate of change of the input waveform<sup>1</sup> (the output from a previous gate) will be proportional to  $1/RC$ , and this compensates for the change in sensitivity with R and C. Also, using a larger logic swing will not reduce feedthrough as a percentage of logic swing, since if the same propagation delay is maintained the transition rate will increase proportionately. The larger the contribution of intrinsic transistor parameters to propagation delay, however, the larger feedthrough should be encountered due to  $C_{cb}$  and  $C_{be}$ .

#### 4.17 APPENDIX II - RESTRICTIONS ON MECL II

##### MC1001/MC1002/MC1003/MC1201/MC1202/MC1203 Six-Input Gate

When tested with all 6 inputs held to  $V_{ILmax}$ , only about half of the packages pass the noise margin tests, and the percentage passing varies widely from lot to lot. The 1200 series packages appear to be better than the 1000 series packages in this respect. Feedthrough is a problem on the inverting outputs if more than two inputs can change simultaneously.

##### MC1001B/MC1002B/MC1003B/MC1201B/MC1202B/MC1203B\* Six-Input Gate

These are MC1001 to MC1203 gates which fail the CSL noise margin test when all inputs are connected to  $V_{ILmax}$  simultaneously. They are retested with  $V_{ILmax}$  applied to one input pin at a time and the other inputs held at -5.2 volts. These units may only be used as one-input gates, with the other 5 inputs tied to -5.2 volts. Almost all units failing the test for MC1001 pass as MC1001B units. No problems with feedthrough occur, since only one input is used.

---

\* CSL designation.



MC1004/MC1005/MC1006/MC1204/MC1205/MC1206 Dual Four-Input Gate

Feedthrough is a problem on the inverting output if more than one input can change simultaneously. Some recent tests have failed a significant percentage of the packages with all inputs held to  $V_{ILmax}$ .

MC1004B/MC1005B/MC1006B/MC1204B/MC1205B/MC1206B\* Dual Four-Input Gate

These are MC1004 to MC1207 gates which fail the CSL noise margin test when all 4 inputs are used, but pass with any one input used and the other three inputs tied to -5.2 volts.

MC1007/MC1008/MC1009/MC1207/MC1208/MC1209 Triple Three-Input Gate

Feedthrough is a problem if more than one input can change simultaneously.

MC1010/MC1011/MC1012/MC1210/MC1211/MC1212 Quad Two-Input Gate

No known problems.

MC1013/MC1213 85 MHz AC-Coupled J-K Flip-Flop

Requires care in application due to A.C. coupling.

MC1014/MC1214 Dual R-S Flip-Flop (Positive Clock)

MC1015/MC1215 Dual R-S Flip-Flop (Negative Clock)

No test experience. These types could probably be tested easily, but would add additional package types to the inventory. Feedthrough is a problem.

MC1016/MC1216 Dual Latch (Positive Clock)

Feedthrough is a problem; otherwise very satisfactory performance.

MC1017/MC1217 Dual Saturated Logic (DTL-TTL-RTL) to MECL Converter

No known problems.

MC1018 MECL to Saturated Logic (DTL-TTL-RTL) Converter

No known problems. MC1039 is recommended for new designs, since it contains 4 translators per package.

MC1019/MC1219 Full Adder

The new design (starting with about January, 1970) has apparently solved the noise margin problems experienced with previous packages. Feedthrough is a problem.

MC1020/MC1220 Quad Line Receiver

No known problems.

MC1021/MC1221 Full Subtractor

Same as MC1019/MC1219.

MC1022/MC1222 Type D Flip-Flop

No test experience. Feedthrough is a problem.

---

\* CSL designation.



MC1023/MC1223 Dual Four-Input Clock Driver

No test experience. 240 MW power dissipation. Problems with crosstalk may be accentuated because of the short rise and fall times. Feedthrough is a problem on the inverting output.

MC1024/MC1224 Expandable Dual Two-Input OR/NOR Gate

No test experience. Feedthrough is a problem on the inverting output if more than one input can change simultaneously.

MC1025/MC1225 Dual 4-5 Input Expander

No test experience.

MC1026/MC1226 Dual 3-4 Input Transmission Line and Clock Driver

No test experience. Problems with crosstalk may be accentuated because of the short rise and fall times. Feedthrough is a problem on the inverting output.

MC1027/MC1227 120 MHz AC-Coupled J-K Flip Flop

No test experience. Difficult to test because of AC coupling. 250 MW power dissipation.

MC1028/MC1228 Dual Four-Channel Data Selector

Test results on the latest shipments have been satisfactory. Control inputs cannot be tied off to -5.2 V. Feedthrough from the control inputs is a problem.

MC1029/MC1229 Data Distributor

Not used. Tests on a limited number of units indicates that almost all units pass.

MC1030/MC1230 Quad EXCLUSIVE OR Gate

No known problems.

MC1031/MC1231 Quad EXCLUSIVE NOR Gate

No known problems.

MC1032/MC1232 100 MHz AC-Coupled Dual K-K Flip-Flop

Test results on 20 packages with date code 6946 (46th week of year 1969) showed that the packages were extremely sensitive and that they would toggle on signals as small as 150 mV. Some of them would toggle twice on a single 800-mV 10-ns transition. A few packages with a newer date code were obtained and tested. These appeared to be generally satisfactory, and this circuit can be used in some applications if care is taken in layout and clock signal generation.

MC1033/MC1233 Dual Latch (Negative Clock)

No test experience. This type could probably be tested easily, but would add another package type to the inventory. Feedthrough is a problem.

MC1034/MC1234 Type D Flip-Flop

No test experience. Feedthrough is a problem.



MC1035/MC1235 Triple Differential Amplifier

Both inputs of unused sections cannot be tied off to -5.2 volts, or left open, because other sections will be affected. Recommended procedure is to tie one input of an unused section to  $V_{bb}$  (pin 9) and one to -5.2 volts. Motorola specifies the noise margin tests for this package with one of the inputs of each section tied to the internal  $V_{bb}$  supply. Since the internal  $V_{bb}$  supply can vary over the range -1.10 to -1.26 volts at 25°C, the amplifier noise margin for differential signals and common mode signals is not adequately tested by this method. The test used by CSL subjects each amplifier to a 90 mV differential signal and independently tests the  $V_{bb}$  output.

MC1036/MC1037/MC1236/MC1237 16-Bit Coincident Memory

250 MW power dissipation. Most of these units do not even pass the Motorola tests which allow a -15 mV (i.e., less than zero) noise margin. A new design for this circuit has been produced by Motorola and 5 samples were tested. These appear to have approximately a 100 mV noise margin and could be used if reasonable precautions were observed.

MC1038/MC1238 Eight-Input Data Selector

Not used. Tests on a limited number of units are satisfactory. Feedthrough from the data inputs is a problem. Control inputs cannot be tied off to -5.2 V.

MC1039/MC1239 Quad MECL to Saturated Logic Translator

No known problems. Positive output rise time may be improved for low fan-out by connecting a resistor from the output to the positive power supply.

MC1040/MC1240/MC1070/MC1270 Quad Latches with Output Gates

The packages produced in 1971 and later are satisfactory. Feedthrough still exists but should no longer cause the flip-flops to drop a "1". Caution must be exercised with the package since Motorola has shipped some packages with a 1971 date code that were made with pre-1971 masks. They were apparently distributor returns that were scrubbed and re-marked. Circuits with the old design, identifiable by the chip size, may drop "1"s when the clock goes off.

MC1042 Dual Binary to 1 of 4 Decoder

Not used. Unused inputs cannot be tied off to -5.2 V.

MC1043/MC1243 3-Bit Binary to 1 of 8 Decoder

Not used. Unused inputs cannot be tied off to -5.2 V. Preliminary test results are favorable.

MC1044/MC1244 Decoder - Nixie Driver

No test experience.

MC1046/MC1246 Eight-Bit Parity Checker and Generator

No test experience.

MC1047/MC1247 Quad Two-Input AND Gate

Feedthrough from pins 3,4,10, 11 only, to outputs is a problem.

MC1048/MC1248

No known problems.



## 5. INTERCONNECTION DESIGN

### 5.1 SCOPE

This section gives a short history and description of the devices and techniques used for electrical interconnection in the first generation of macromodules. The functional descriptions of the various components are covered elsewhere in "A Macromodule User's Manual", which is included as Part 1, Volume II of this report; and the manufacturing information is contained in a series of documents whose numbers will be referenced as the individual components are discussed. The manufacturing documents are included as the 14 volumes of Part 2 of this Final Report, and reference will be made to the volume and page number where applicable. The reader is expected to be familiar with the above User's Manual in order to understand the functions, usage, placement, and identification of the various macromodular components.

An "inside out", "outside in" organization will be followed for the remainder of this chapter. After a general background discussion, there will be sections on the Electronics Package and Faceplate Box. These will trace the path of electrical connections from inside the Electronics Package to the front of the Faceplate Box, from which the electrical signals are sent to other modules in the system. The various Cables will then be discussed, followed by a description of the "implicit" pathways which route power and signals through the Frame structure. The last subsection will summarize the performance and reliability of the completed system.

### 5.2 BACKGROUND

The design objectives of macromodules manifested the explicit threat of a high density of electrical interconnection. The level of modularity and the requirement to suppress all extraneous electrical engineering details led to somewhat unique requirements for these many interconnections.

The very fact that a user could (and does) have the capability to build, tear down, and rebuild a large digital system several times a day leads to a very stressful environment for the primary connectors. While enduring the stress of constant handling, the primary connectors and cables must be small, light, and flexible enough to remain in the background of the user's thoughts during the design, assembly, debugging, and use of a macromodular system. The primary connectors include the cable connectors (Control and Data Cables) and the 90-pin block connectors used as the standard interface between Electronics Packages, Frame members, and Faceplate Boxes. This primary class receives the greatest handling from the users.

There is a secondary class of connectors which are only partly visible to the user and which are almost invisible in the functional description of the system. These carry implicit vertical and lateral data and control signals and certain system-wide control signals such as Preset and Shield. These secondary connectors also carry the system power and power-sequencing circuits. Secondary connectors include the 90-pin blocks on the Fan Module



and the heavy power connectors on the Fan Module and Pedestal Slug Power Supplies. The Channel Coupler and Daisy Chain Cable may also be considered in this class.

Finally, there is a tertiary class which exists internally in the modules, frame, and pedestal. These connectors couple both signals and power among boards, and between boards and the functional boundary of the unit where the secondary connectors take over.

Internal interconnections have only to withstand the process of manufacture and assembly and the rare disassemblies for repair. The primary connections on the Frames, Faceplate Boxes and Electronics Package, as well as the cable connectors, must endure constant handling, with the possibility of multiple insertion/withdrawal cycles in each day of use. Coupled with this high usage is the requirement that each pin be able to pass a logic signal of less than two volts.

Interconnections involving such low voltages usually require gold plating on the mating surfaces [3]. Concern was therefore directed to the selection of gold-plated connectors which could retain the integrity of the gold surface after at least 500 cycles of insertion/withdrawal. Most connector specifications limited testing to 25 cycles, so experiments were performed on a wide variety of connectors to determine extended service life.

Comments on individual connectors will be made in later sections, but it may be simply stated that all macromodular interconnections are able to maintain a gold-to-gold contact after 500 or more cycles of usage.

Connector contact testing was performed mainly by mating and unmating whole connectors or single contacts either by hand or in simple alignment fixtures. The mating surfaces were then examined for deformation and condition of the gold plating. Freshly scored base metal is very difficult to distinguish from gold with an ordinary "dissection" type microscope, so most contacts were immersed in a corrosive solution to color the base material. For beryllium-copper contacts, a solution of equal parts  $\text{CuSO}_4$  and  $\text{NH}_4\text{Cl}$  dissolved in several volumes of  $\text{H}_2\text{O}$  at about  $160^\circ\text{F}$  gives a light gray deposit which quickly shows any abraded area.

Most connections involving wire-to-contact terminations in a macromodular system are crimped, and several tests were performed on each type of crimped connection to ascertain or verify crimp-tool settings or less obvious variables (see section 5.6) such as the number of conductors loaded into a crimped splice.

Crimped contacts were sectioned, polished with a fine abrasive stone, and etched lightly to show wire boundaries. Usually a dilute  $\text{FeCl}_3$  solution was used, as this leaves a clean, bright surface. The selection criteria included: 1. Lack of voids; 2. Deformation of every wire in a multi-wire bundle, but no excessive crushing; 3. Integrity of the retaining wall of the connector, with no cracks or excessive thinning of the wall section.



The sectioning results are somewhat subjective unless rigorous area measurements are attempted. They are rendered more useful when combined with pull tests to determine the actual strength of the crimped connection. A proper crimp will exhibit about 70% of the tensile strength of the wire, which in practice means that a #30 copper wire will elongate about 150% before breaking, either at the crimp or near it.

The early selection of a family of Emitter-Coupled-Logic Motorola MECL II to perform all internal logic functions led to the design rule that all signal pathways have a constant impedance. The rise times of the ECL were about 5-10 nanoseconds for lightly loaded lines, and the asynchronous "transition" logic control architecture was very intolerant to overshoot and ringing. Therefore, all internal signal pathways over about 3 inches long were source-terminated and routed over a ground plane (stripline) for an impedance of about 125 ohms.

These stripline etched lines were determined to be  $0.012 \pm 0.003$  inches, for an impedance of 120 to 130 ohms on an 0.062 inch thick epoxy glass board. No attempt was made to control the impedance of the connectors, since all were electrically "small" at a rise time of 5 nanoseconds, and the presence of a connector could just barely be detected on a sampling oscilloscope display. All signals conveyed by cable are sent differentially on twisted pairs with a balanced impedance of about 125 ohms.

The cable designs are covered in detail in section 5.5, in addition to a few devices which share the Data Cable connector and which provide some of the user-machine interface necessary for setting up and debugging a new macromodular configuration.

Two deviations from this control of impedance will be mentioned later as specific devices are discussed (see Sections 5.4 and 5.5). A more complete discussion of the design rules for using ECL has been given in Section 4.

In addition to the low-voltage signal connections, there are a class of heavy, large connectors which are used to distribute power and certain system-wide sequencing signals. (These make up part of the secondary class mentioned above). These connectors have the dual role of carrying high-voltage, high-amperage AC and DC circuits, in addition to a few low-voltage control signals. These power connectors must withstand considerable physical abuse as Frames and Power Slugs are assembled. More about this will be found in Section 5.6.

The modular aspect of macromodular interconnection creates special conditions of size, signal environment, power density, and durability which will be taken up in the succeeding subsections.

### 5.3 ELECTRONICS PACKAGE

The electronics package is the computational heart of the system, and the various functions available are well covered in the "User's Manual" (Part 1, Volume II) while the manufacturing information is covered in the



200 series documents in volumes II, III, and IV of Part 2.

The typical electronics package internal structure of 7 vertical boards and 4 motherboards (Figure 67) needs some explanation. At the time of the original design work (1965-66), large multilayer boards were being introduced in a few systems, and a study was done to determine their suitability to modules with about 100 integrated circuits and 300 to 400 external connections. Several suppliers were contacted, visited, and studied. Attempts were also made to produce small multilayer prototypes in our nascent printed circuit facility, with very poor results.

The commercial suppliers suffered from low yields and very high prices. In addition, a functioning in-house prototype facility was considered crucial because of the complexity of the individual module logic structures, and the many unknowns connected with layout rules for controlled-impedance signal paths.

Small boards of the type shown in Figure 67 could be produced in a modest prototype shop, and after verification, these board designs could be economically produced in a plated through-hole version by a commercial supplier.


Once the decision was made to adopt this "Chinese Puzzle" configuration of circuit boards, the board-to-board interconnections became very important components, and the first selection effort was directed toward finding an individual path connector.

The first prototype modules of the "Puzzle" configuration were built with ELCO subminiature "VARICON" fork contacts for board-to-board connection and for the block type connectors on the motherboards. (See Figure 68.) These connectors were chosen for the high density of contacts (0.1-inch centers) and for the wide variety of contact tail styles. We felt that the use of a card-edge type of connector would force an unnaturally large size on the Electronics Package.


The ELCO connectors were replaced by two families of AMP connectors due to certain failures during evaluation. The fork contacts used a very high mating force to ensure a clean surface, with the result that the thin gold plate was removed after as few as five cycles of use. This would be of little consequence if the high pressure were really able to scrape the oxides from the mating surfaces. Unfortunately, the forks had a tendency to relax, and this lowered pressure coupled with a few weeks in a corrosive environment led to high contact resistances and erratic performance at low voltages. In addition to the electrical problems detailed above, normal handling of the boards during manufacture and assembly caused the individual contacts to crack. The block connectors on the ends of the motherboards were more durable mechanically, but the insertion force for a single-high electronics package was approximately 300 pounds, and no mounting method could be devised which would keep the connectors mounted firmly on the motherboards. (The mounting "ears" provided by ELCO broke instantly).

The individual path connections were finally made with the AMP AMPMODU






(a) Partially assembled.




(b) Fully assembled.

Figure 67. "Chinese Puzzle" circuit board configuration.





(a) Detail.



(b) Assembled in place.

Figure 68. ELCO "Varicon" fork contact connectors.



system. The male is a 0.025-inch-square stamped, coined, and plated post. The female is a stamped, folded, and plated cage with two cantilevered leaf contacts which apply very light pressure to the mating post surfaces. (see Figure 69 ). The AMPMODU pair is much more tolerant to misalignment than the ELCO mentioned above, and has the property of nearly undetectable wear after hundreds of usage cycles. Several other connector products were evaluated, but were clearly inferior for this application. (Among these were pin socket sets from CAMBION and a "tuning fork" design from AMPHENOL.)

The block type connectors on the ends of the four motherboards provide a standardized mechanical and electrical interface between the Electronics Package and the remaining macromodular components. (See Figure 67 for identification of the four motherboards.) The Top Motherboard connector communicates with the Bottom Motherboard of the module above through the V-Bus connections on the rear of the Faceplate Box. The Faceplate Motherboard connector mates with the third connector on the rear of the FP Box and carries data and control pathways for explicit connection by the user. The Lateral Motherboard transfers power and power sequencing signals from the service column and Channel, and carries the implicit data and control signals for lateral communication between modules. The system chosen for these connectors is the AMP 750 Series Box Contact connector. This system was developed for use on helicopters, and has very good tolerance to shock and vibration. The 750 Series shares the low insertion force, longevity of gold plate, and tolerance of small misalignments of the AMPMODU series, but achieves this through a different mechanical structure.

The male is a stamped and formed U channel, while the female square cage contains four hyperbolic leaf springs which give several possible contact surfaces with a low-pressure wiping action during insertion. The 360 contacts provided by the four motherboards' connectors require only about 20 pounds initial insertion force, as opposed to the 300 pounds of the ELCO system.

These are high-density connectors with contact spacings of 0.075 inches, and as a direct result, they are fragile. A summary of the performance to date will appear in the last section.

#### 5.4 FACEPLATE BOX

As mentioned before, the 750 Series connectors on the Electronics Package motherboards provide a standard interface to the remainder of a macromodular assemblage. Typically, three of these mate with Faceplate Boxes, and one mates with the Channel member of the frame. The Channel provides power, global signals such as Preset, and the implicit lateral pathways between adjacent macromodules.

The three 750 Series connectors which mate with the Faceplate Box function as follows: Two connectors are wired together on a pin-for-pin basis to form the Vertical Bus (see 300.6 in Volume VII, Part 2) which provides the implicit vertical communication in a manifold. The third connector is wired to a variety of connectors on the front panel, or Faceplate, of the box.






Figure 69. Detail of AMP AMPMODU connector pairs.



The Faceplate Box is simply an adapter between the standard system connectors and a wide variety of Faceplate connectors which allow explicit communication between modules and between a module and the external world.

The major Faceplate connectors allow the use of the Data Cable and Control Cable, and as these connectors are highly interdependent with their respective cables, they will be discussed in conjunction with the cables (section 5.5). The Faceplate provides a ready carrier for specialized connectors such as the dual BNC outputs of the D/A module and the board edge scheme of the Function Call module, in addition to providing space for the function selector switches on the same module.

The internal wiring of the Faceplate Box consists mainly of individual twisted pairs routed from back to front. This is one instance where the impedance of the path is not strictly controlled. The pairs are about 6 inches long, and the impedance is about 140 to 160 ohms, depending upon their position in the bundle and their nearness to the cover plates.

This length and the magnitude of the impedance change are quite small for the rise times encountered in MECL II. The waveforms in the system exhibit a kind of textbook smoothness, with no overshoot and uncomplicated transition regions. This state of affairs starts to degrade with MECL 2.5 circuits (2-3 ns), and becomes marginal with ECL 10K circuits with 1 to 2 ns rise times. The Restructured Macromodules (see Part 3 of this report) have eliminated the Faceplate Box and have retained tighter control of impedance.

Manufacturing details for the series of Faceplate Boxes will be found in the 300 series documents in volumes VII, VIII, and IX of Part 2. Descriptions of the Overlays will be found in Part 1, Volume II; and little need be said of the Code Switches (document 300.7, Volume VII), since these are standard gold-contact switches whose main characteristics are mechanical rather than electrical.

## 5.5 CABLES

The two major system cables (Data and Control) must carry controlled-impedance differential signals over small distances, and must also be easy to handle. These cables and other data and control devices will be discussed in turn.

### 5.5.1 Data Cable

The Data Cable has 12 data pathways, 2 control paths, and 3 spare pairs. The design is somewhat unusual for a computer cable, and the various parts will be discussed from the standpoint of macromodular design objectives. Refer to document 010 (Part 2, Volume I) for the cable specification and to document 370 (Part 2, Volume XI) for the assembly procedure. The cable and the associated connector mutually influence the selection of each other; so a few words are necessary about the Amphenol ASTRO-348 connector which is used for the Data Cable.



Many different shapes and styles of data connector were studied, but the final selection was determined largely by the natural sizes that had been assumed by the Electronics Package and Faceplate Box. The implicit density and the necessity of easy handling characteristics eliminated rectangular rack-and-panel types with their attendant jackscrews in favor of subminiature circular designs.

An inexpensive subminiature circular connector (Amphenol 222 series) was selected for evaluation. These literally crumbled while under study. Among other problems, the plastic shells were unshielded, the contacts did not retain their gold plate and had a tendency to migrate rearward during repeated mating cycles, in addition to breaking the wires just behind the crimp, and the retention ears would break after a few gentle operations from the user's hands.

The military subminiature circulars had nearly as many problems. Some had contacts which had a 0.004 inch internal wall. These could be broken between the thumb and opposed fingers of one hand. Others had contact retention schemes based on tiny spring slips which were equally fragile.

The final decision was to use a modified version of the AMPHENOL ASTRO-348. The locking ring was removed as too large and too hard to operate at close connector spacings. The bulkhead connector was left virtually unchanged, but the cable end was additionally fitted out with a housing and cable-capturing scheme.

Here starts the interaction between connector and cable, since the ASTRO-348 was designed to terminate a bundle of individual wires, not a jacketed cable. The problem was to get a sufficient length of wire to allow insertion of the contacts without having a long unwieldy, housing or a big wad of crushed wires after assembly.

The ASTRO-348 is normally swaged into a one-piece assembly for very good reasons in the military environment. The modified version allows complete removal of the contact retention disc, which then allows contact insertion with a very short length of wire free from the cable. (See the 370 document in Volume XI, Part 2 for the detailed pictures of the assembly procedure). The heavy braided shield of the cable provides electrical and mechanical termination of the cable to the rear of the housing.

The Data Cable had to be easy to handle, and have a small diameter and good crimping characteristics for ease of assembly. The only commercial cable available with nearly the correct impedance was the National Wire Dynatronic D200 series, which had a large-diameter, stranded wire, and was extremely stiff. These characteristics are not important for long runs of hidden cable, but were decisive in the projected macromodular user environment.

The individual conductors for the Data Cable pairs are made of #30 AWG solid wire, since solid wire is easier to crimp and a bundle of 17 pairs of #30 solid was much more flexible than the bundle of 15 pairs of #24 stranded wire available from National.



Controlled impedance is not a normal manufacturing specification for a multi-conductor cable. The manufacturer will build nearly anything specified by the customer, but the final characteristics are the responsibility solely of the customer. The wire makers are, however, quite used to one close-tolerance wire, wire-wrap wire, for the reason that the insulation diameter is very important to automated wire-wrap machines. Wire-wrap wire uses a close tolerance Oxygen Free High Conductivity (OFHC) copper wire with a thin silver coating.

Choice of insulation is another matter, since one of the major parameters of wire-wrap wire is resistance to cut-through, and the dielectric constant (K) and dissipation factor are secondary to mechanical needs. The only available wire-wrap insulation with low K and loss was polysulfone. A trial cable made of polysulfone had good electrical characteristics, but terrible mechanical problems when twisted into pairs and jacketed into a small cable. (When the jacket was stripped back, all pairs instantly unravelled.) Negotiation with a friendly manufacturer (BRAND-REX) led to the selection of polyethylene as the primary insulation. This was becoming popular for telephone cables, but was not suitable for wire-wrap applications. BRAND-REX was willing to apply wire-wrap tolerances to polyethylene, and a successful cable was produced.

The wall thickness of the primary insulation was partly determined from empirical data obtained on small hand-made cables, and partly from the dielectric constant and a packing factor known only to cable manufacturers.

Reference [4] contains a formula for the impedance of parallel lines in an infinite uniform dielectric medium:

$$Z_0 = 120 \cosh^{-1} \frac{D}{d},$$

where d is the diameter of the lines, and D is the center-to-center distance between the lines. This formula may be restated for a medium of dielectric constant K as:

$$\frac{D}{d} = \cosh \frac{Z_0 K}{120}$$

Since the wire diameter d and the impedance  $Z_0$  are fixed and K becomes fixed when an insulation is chosen, the relation yields the interwire spacing D. The wall thickness of the insulation is then  $1/2 (D-d)$ .

Experience has shown that the packing factor of a cabled bundle causes the  $Z_0$  to rise a few ohms, so a value of 120 ohms was used to calculate the insulation wall. The finished cable measures  $125 \pm 5$  ohms by Time Domain Reflectometry.

The individual pairs are twisted with four pitches in an emulation of large telephone cables, which have larger numbers of pitches to suppress crosstalk. One unique feature of this cable is an inner jacket of insulation under the shield to protect the outer pairs from the impedance variation



which would result from close proximity to a metallic ground. (In retrospect, we discovered that National Wire achieves this by increasing the primary insulation diameter on the outer pairs in the D200 cable series. This variation is too expensive for custom color-coded cables of the type discussed here).

In the first polyethylene cable (Data Cable Model C) this inner jacket was made from polyethylene for low loss. This proved a little stiff, so the Restructured Data Cable (see Part 3) has an inner jacket of PVC which has negligible electrical effect, and the cable is more flexible in spite of 6 more pairs.

The shield is braided copper for good mechanical strength, since the connector housing retains the cable and maintains ground continuity by effectively crimping the shield. A high braid angle of 50 degrees was specified for increased flexibility at the expense of easy push-back. The assembly procedure thus requires that the shield be combed out and trimmed before the crimp operations. The outer jacket is thick enough to provide insulation and abrasion protection for the braided shield and a smooth comfortable surface for the user.

The present Data Cable is about 0.312 inches in diameter, and is quite flexible, with a bending radius of about 2 inches.

#### 5.5.2 Control Cable

The Control Cable is a simple twisted pair, with shield and outer jacket of pressure-molded PVC to give good mechanical termination to the connector. The balanced impedance was controlled at 120 ohms by the manufacturer (BRAND-REX). (See Volume I, 010 section.) The internal conductors are Copperweld steel, to give sufficient strength to withstand the withdrawal force of the connector.

The control connector is an individual contact from a large block structure designed for simultaneous termination of large numbers of coaxial cables. The AMP Twin Standard Coaxicon is crimped to the Control Cable in a single operation (two wires and the shield), and the other jacket is then crimped under the ferrule. (See 360, Volume XI.)

Some difficulty has been encountered in the manufacture of Faceplates, because AMP does not maintain close control over the outer diameter of the coaxicon (see 300.5, Volume VII). Each lot of Coaxicons must be measured, and the ream dimension adjusted for a good press fit.

#### 5.5.3 Daisy Chain

There is another system cable (Daisy Chain) which is handled by users, but only during the initial phase of configuration. This Daisy Chain cable provides a pathway between a Pedestal Controller and each Pedestal used in a system.



Only a few signal pathways were required (10), and the logic family chosen for the Pedestal had slow but large signal swings, so individually shielded conductors were selected for low crosstalk. The individual #24 AWG stranded conductors are insulated with PVC and shielded with an aluminum/mylar film with the metal on the outer surface. Two drain wires and an outer shield with the metal facing inward provide the necessarily excellent ground return required by single-rail logic signals. The cable was made by National Wire to the specification in document 361, Volume XI.

These Daisy Chain cables are terminated by small ASTRO-348 connectors from the same family used in the Data Cable. The same rear housing and compression nut are used, but an additional adapter is required to mate to small connector shell to the outer housing. (See 361.8 Volume XI).

The tined ferrule (see 361.9 Volume XI) differs from the Data Cable in that the oval rear opening is changed to a 3/8 - 16 internal thread to capture the outer jacket of the Daisy Chain Cable.

#### 5.5.4 Cable-Related Devices

Three useful programming and debugging aids have been implemented as special-purpose extensions of the Data Cable connector.

The Parameter Plug and the Parameter Switch allow the user to input fixed and variable data values respectively.

##### Parameter Plug

The Parameter Plug inputs a fixed value. Two popular values have been documented, but the device is so simple that several special plugs have been made in other useful patterns. Documents 352 and 353 in Volume X give manufacturing information for the 7777(octal) pattern and the 0000 pattern respectively. Other patterns which have proved useful in small quantities are 0001, 0007, 4000 and 6000.

##### Parameter Switch

The Parameter Switch has four thumbwheel switches and a somewhat elaborate internal circuit which allows the user to set any 12-bit data value or 4-bit mode code as a quasi-static value. (See the 351 document in Volume X for construction information.) The Interswitch thumbwheel switches are not the best possible in terms of contact design and durability, but they were the only switch thin enough to allow two Parameter Switches or a Data Cable plus Parameter Switch to share the limited space available on the Faceplate. The Parameter Switch design came long after basic dimensions had been fixed, so the basic fragility of the contacts had to be overcome by special fixturing and handling during assembly. (See 351 document, Volume X.)

The circuit of resistor networks and two selected diodes provides sufficient differential voltage to the data receivers of the host macro-module over a range of variations, which include power supply regulation, resistor tolerance, and 4 to 12 data loads. The values chosen are the result



of an iterative computer analysis which showed the sensitivity of the data voltage to variations in each of the components. The plastic case and connector mounting were designed for the Parameter Switch, but were then quickly adopted for the LED Data Indicator box, Miniconsole and other associated equipment.

#### LED Data Indicator Box

Data values can be read out quickly in octal by using a LED Data Indicator box. This box uses the Parameter Switch housing hardware, with the addition of a panel to support the light-emitting diodes (LED's). The internal circuitry converts the MECL II logic levels to voltages sufficient to drive the LED indicators.

#### Mini-Console

Control sequences may be initiated or examined with the aid of a Mini-Console. Again using the basic housing scheme of the Parameter Switch, the Mini-Console has circuitry to initiate control signals, and two indicator drivers which allow the user to monitor the state of the control pathway.

### 5.6 FRAME COMPONENTS

The Frame components do not use any of the primary connectors which define the explicit data and control pathways in the implementation of an algorithm.

Two types of secondary-level connectors appear in the form of the AMP 750 series box contact connector and a heavy power connector (AMP type W). The AMP box connectors have been described in section 5.3. Four 90-pin versions pass lateral data information between adjacent modules and transfer power and system-wide control signals between the Lateral Channel and each module. Three pins are used for power and ground respectively; no arcing damage has occurred, since most users remember to power down the system before inserting or removing modules.

Two 30-pin 750 series connectors pass data through the Channel Coupler (see 500 in Volume XIV) to an adjacent channel. The signal lines in the Channel are controlled impedance, but the path through the Coupler consists of short twisted pairs with one member of each pair grounded. No attempt has been made to measure the impedance of these short lines, since no degradation has been noticed in the lateral pathway waveforms.

The Channel contains two circuit boards sandwiched together with soldered jumpers through the boards. A good pictorial sequence in the manufacturing document (see document 403 of Volume XII) should fully explain the structure. The two boards provide the necessary constant impedance for the data pathways and wide, low-resistance buses for the primary 55-volt power, zero-volt return, and signal ground.

Heavy power connectors are used by the vertical Service Column which passes through the Fan Module, and by the modular power slugs which insert into the Pedestal. These connectors (AMP type W) carry large currents on some pins (50 amps), large voltages on others (120 V AC), and a low-level control signals on a few contacts. The type W has a wide choice of contact sizes (and wire crimp sizes) contained in a physically rugged housing. In



both applications, Fan Module and Power Slug, the cast metal housing of the connector serves directly as a mechanical guide and protects the insulator block from the forces caused by misalignment of heavy frame parts during system configuration.

Within the Fan Module and Pedestal there are numerous tertiary-level connectors whose form should indicate their respective function. AMPMODU pairs are used for board-to-board interconnection and FASTON terminals make the motor-capacitor and circuit-breaker line connections, and various small crimp splices connect individual wires. The large crimp splices on the power and ground buses are used to gain a common point in each frame block for all wires of a bus. Differences in contact resistance could lead to voltage differences in the individual wires, so the power for the frame block is tapped at the common point defined by the large splice (441 document in Volume XIV). Numerous sectioning and pull-out tests were performed to determine the proper number of wires to load into these splices, since the crimp tool was not adjustable for fine variations. The final loading of the splice gave uniform deformation of all wires and no detectable voids when examined by the sectioning and etching technique mentioned earlier.

## 5.7 PERFORMANCE

Since late 1971, when significant quantities of macromodules were first placed in service, the various interconnection devices have seen the rigorous service predicted by the designers.

Two classes of failure have become evident during a period of use exceeding two years: early failure due to improper manufacture, and random failure due to usage. Each type of connector will be treated separately below.

### 5.7.1 AMPMODU

The AMPMODU system used in the Electronics Package has not suffered a single usage failure. There have been manufacturing defects such as females inserted backwards, females filled with solder, or improperly staked male contacts. These imperfections were generally caught during visual inspection of the boards prior to final assembly. There have been occasional instances of a mis-assembly of the board stack allowing a male pin to brush against the outside of a female instead of being inserted within the square cage of the female contact. This type of fault is insidious, since the good initial contact will allow the module to pass all electrical tests, only to fail later after handling or vibration have rendered the connection intermittent or open. The only protection against these faults is visual inspection of the completed board stack, pin by pin, before the stack is inserted into the metal wrap-around, using a fibre-optic inspection probe (CYSTOSCOPE).

### 5.7.2 AMP 750

The AMP 750 series box contact connectors were recognized as fragile from the earliest evaluation experiments. A little over 100 individual



male contact pins have been forcibly pushed into the plastic housing or have been mangled. Nearly half of these incidents were traced to a single Faceplate Box being used in the automated module tester.

This Faceplate Box had a foreign material (probably grease pencil wax) filling one of the female contact positions. As each module was inserted into the tester the filled hole would cause the associated male pin to be bent or crushed. These pins are replaceable by a tedious desoldering operation.

The remaining failures have not been traced to any single cause, and they may be due to simple carelessness when handling the modules. There are nearly 400,000 AMP 750 contact pins in our current inventory of macro-modules, so the repairable failure of 100 pins has not been a major problem.

#### 5.7.3 ASTRO-348

The ASTRO-348 connectors used for the Data Cables and associated data ports on the Faceplate Boxes have suffered very few failures. There are over 2000 bulkhead ASTRO-348 connectors in our current Faceplate inventory, and only 2 have failed due to mechanical abuse. These had severely bent male pins, and the damage could only have resulted from insertion or a foreign object.

There is one manufacturing failure mode which is common to all the ASTRO-348 contacts. The length of the bare wire remaining after stripping is very critical for #30 AWG solid wire. If the bare end is too long, the wire may break as the assembly is handled during manufacturing. A small length of insulation must be inside the contact barrel at the time of crimping to provide a cushion for the wire. However, if the bare wire is too short and the insulation is crimped with the wire, intermittent and position-sensitive connection failures will result.

#### 5.7.4 Data Cable

About 9 Data Cables have failed due to broken or intermittent connections, from a population of about 550. These cables are marginally repairable, but complete rebuilding is recommended since it is hard to discover which end of a cable is at fault without disassembling both ends. Many Faceplate Boxes were rejected for the above problem when first received from the manufacturer. These were repaired and placed in service.

The first manufacturing run of Data Cables had a design flaw. They used a round ferrule (see 361-9, Volume XI) to retain the outer cable jacket. This ferrule allows the cable to rotate within the housing, and eventually the outer jacket becomes free, exposing the shield to damage. About 8 short cables of this design have failed, and a slow attrition is expected. The connector parts may be salvaged for use in special interfaces and adapter cables between the original and restructured macromodules.



#### 5.7.5 Control Cable

The Control Cables and their associated AMP Coaxicon connectors have exhibited very few electrical failures, but there have been 40 or so cables which have allowed the outer jacket to slip out of the crimped ferrule. A better design would use a longer ferrule or a stiff rubber strain relief. Some retrofitting may become necessary if large numbers of cables develop mechanical difficulties. There are approximately 1900 control cables presently in service.

#### 5.7.6 Power Connectors

The heavy Power Connectors (AMP type W) have suffered from 4 instances of broken retention springs on male contacts, probably due to mishandling during the crimping operation. There has been no evidence of other electrical or mechanical damage.

The Channel circuit board assembly has shown some 19 cases of high-resistance (40-ohm) short circuits. These tend to heal when the affected pad is resoldered, but a more permanent repair requires enlargement of certain clearances on the ground plane side of the board. This design error, and the failing Channels, will be repaired as necessary.

#### 5.7.7 General Performance

The performance of the interconnections in macromodules has been very good. Most failures described above can be attributed to manufacturing errors which remain repaired once discovered. There have not been any random or intermittent "spooks" which could not be finally identified with a particular device.

The Molecular Modeling System described elsewhere is a good example of connector reliability. This system flushed out several initially defective devices, but then continued to operate on a heavy schedule with no connector failures, in spite of much reconfiguration and one 3000-mile trip.

The system contains 141 modules and associated parts as detailed in the table below. This large assemblage has been torn down and rebuilt 7 times with no apparent interconnection failures. The long trip mentioned above involved a teardown in St. Louis, Mo., a trip by truck to New Hampton, N.H., and a successful rebuild and operation for a week. The system was then torn down, trucked to Yale University in New Haven, Conn., and again rebuilt and operated, this time in an extremely humid location. Finally the system was returned to St. Louis, rebuilt, and is still operating successfully at the time of this report.

Numerous smaller systems have been used both here at CSL and at remote locations with similar reliability.



Table 17. Molecular Modeling System Components

Electronics Packages	141
Faceplate Boxes	141
Data Cables	52
Control Cables	285
Frames	12
Pedestals	2
Pedestal Slugs	5
Mini-Consoles	6
L.E.D. Data Indicators	6
Parameter Switches	12



## 6. POWER AND SYSTEM-WIDE CONTROL SIGNAL STRUCTURE

### 6.1 INTRODUCTION

The sketch in Figure 70 illustrates a typical macromodular setup with the power and system-wide control signal structure emphasized. The controller normally originates control signals, which are passed to each base pedestal via the daisy chain cables. Each pedestal can draw power from one or two 115V AC, 20-ampere power outlets. Each pedestal has a plug-in supply handle cover door with a position sensor (open or closed) so that plug-in supplies cannot inadvertently be plugged in or removed when the power is on. The location of the set of lights that depict user-correctable undesirable conditions (such as a supply handle cover door open) is also shown.

### 6.2 PEDESTAL CONTROL FUNCTIONAL DESCRIPTION

With this brief overview, let us consider the functional description of the pedestal control structure, which has three primary functions. First is the distribution of the system-wide control signals, Preset and Shield, from the console to all modules. Second is the control and sequencing of power turn-on and turn-off in response to signals from the source controller. Third is the detection of undesirable conditions that are hazardous to equipment, users, or proper system operation. If such conditions exist, turn-on of power is inhibited. Under some fault conditions, the pedestal will generate shield signals and distribute them system-wide and turn off its own primary power, and in one case, will turn off power in other pedestals.

#### 6.2.1 Design Guidelines

The pedestal control circuitry design was based on the following guidelines:

- A. The pedestal units shall be expandable to any number, all controlled from one source.
- B. Turn-on of power will be inhibited if there are macromodular system configurations above and in each pedestal unit which will, when power is turned on, result in circuit damage, have insufficient power to supply the load, or produce personal hazards. These conditions, which may be corrected by the system user, will be indicated.
- C. Equipment failure detection shall be kept to a minimum. In general, the complexity of the circuitry protected shall be greater than the complexity of the detector, and only those failures which will damage additional equipment shall be detected. A trained observer shall be able to find the defective pedestal by visual inspection.
- D. The MECL control signals between the pedestal unit and the frames shall withstand shorts to any voltage from zeroV to -5.2V DC, and the sense lines shall withstand shorts to each other and to any voltage from zeroV to -5.2V DC without damaging any pedestal circuitry.






Figure 70. Power and system-wide control devices and connections in a macromodular system.



- E. An attempt will be made to preserve stored information in the event of an equipment failure or a user indiscretion (restructuring the system, other than cables, with main power on). However, the first priority shall be given to personal and equipment safety.

#### 6.2.2 Interconnection of Pedestals and Controller

System-wide control signals are assumed to be generated by a controller unit, which is normally part of a console. The controller is connected by a single cable (daisy chain cable) to a pedestal unit, which may be connected in a daisy-chain fashion to a second pedestal unit, which in turn may be daisy-chained to additional pedestal units. The final pedestal unit in the chain is referred to as the end unit.

There is a bi-directional communication between the controller and a chain of pedestal units. The controller may originate the following signals to the chain of pedestals:

Shield  
Preset  
Auxiliary Power  
Main Power

Each pedestal unit receives these input signals and, when sufficient conditions are met, relays them to the next pedestal unit in the chain until the signals reach the end unit.

Various pedestal units may originate control signals (returns) that pass through the daisy chain toward the controller. However, under normal operating conditions, the following four signals directed toward the controller are generated only by the end pedestal:

Shield Return  
Preset Return  
Auxiliary Power Return  
Main Power Return

Two additional signals are used in the chain of pedestals.

System OK  
System OK Return

Any pedestal may originate signals on these pathways.

#### 6.2.3 Interconnection of Pedestals and Frame Blocks

The pedestal unit communicates with the column of frame blocks above it. In addition to supplying 54.62 volts DC for main macromodule power, and 110 volts AC to drive fan motors, the pedestal unit sends the following signals to the frame blocks above:



Column Preset  
Column Shield  
Power Down Request

The pedestal unit obtains information from the frame blocks above it by sensing the following pathways:

Power Down Acknowledge  
Sense  
Frame Sense  
Cap Sense  
Cooling Sense

The Sense, Frame Sense, and Cap Sense signals are used to detect conditions under which main power should not be turned on. The Power Down Acknowledge signal is used to delay the main power turn-off until all modules have done all necessary internal power down sequencing. The Cooling Sense signal is used to detect a condition (cooling failure) under which main power should be turned off.

#### 6.2.4 Definition of Daisy Chain Signals

The nominal daisy chain cable signal levels are zeroV DC (high) and -15V DC (low). When the daisy chain signals are high, the auxiliary and main power are off and the Preset and Shield signals are asserted. Normal control of the daisy chain signals involved changing a control signal only when its return is in the same state. The pedestal logic was designed to produce the expected logic levels only under normal operation. While random switching of the control signals can produce no physical or electrical damage, the logic levels may produce undesirable results (for example, changing stored information). Each signal and its return performs a different function.

#### A. Auxiliary Power and Auxiliary Power Return

The Auxiliary Power signal, when switched low (-15 volts) by the controller, turns on a small auxiliary pedestal circuitry power supply in the first pedestal. This supply then turns on the auxiliary supply in the second pedestal in the daisy chain, etc. The end pedestal unit initiates an Auxiliary Power Return signal (signal to the low state), which is relayed back along the daisy chain to the controller. Receipt of the Auxiliary Power Return signal by the controller ensures that the pedestal system is ready to respond to other control signals.

The Auxiliary Power signal, when switched high, is relayed along the daisy chain to the end pedestal unit. The end pedestal unit then turns off its auxiliary power supply only if the main power supply is off. When the auxiliary power supply has switched off, the Auxiliary Power Return signal (signal to the high state) is passed on to the next pedestal. Thus the receipt of the Auxiliary Power Return (to the high state) at the controller indicates that all power is off in all pedestals.



## B. Main Power and Main Power Return

The Main Power signal is normally asserted (signal to the low state) only if the System OK Return signal is high and the Auxiliary Power Return has been low for 1.5 seconds minimum. When the Main Power signal is asserted by the controller, the first pedestal unit initiates a one-second-long attempt to turn on the main power supply. If there are no undesirable system conditions in the pedestal unit, the AC input power to the main power supplies is switched on, and the main power DC bus voltage then has one second to reach 54.62V DC. When the bus voltage has reached 54.62V DC, the Main Power signal is passed to the next pedestal unit. If there are undesirable system conditions in the pedestal, the AC input power is not turned on. (The undesirable system conditions must be corrected and then the Main Power signal switched high, and after 1.5 seconds minimum, low in order to turn on the main power supplies.) The end pedestal unit initiates the Main Power Return signal (signal to the low state) which is passed along to the controller, latching the Auxiliary Power Return signal on. Thus the main power supplies must be off before the small auxiliary power supply can be turned off. Also, as the Main Power signal is received at each pedestal, the pedestal-to-frame block signal Power Down Request is turned off.

The Main Power signal, when switched high, is passed along to the end pedestal unit, asserting the pedestal-to-frame block Power Down Request signal at each pedestal. The end pedestal unit turns off its main power supplies when the pedestal-to-frame block signal Power Down Acknowledge is asserted. When the main DC power supply has been switched off, the pedestal passes the Main Power Return signal (signal to the high state) to the next pedestal unit. Thus, the receipt of a high state of the Main Power Return signal at the controller indicates that all the main supplies in all the pedestal units are off.

## C. Preset and Preset Return

The Preset signal is normally turned off (signal to the low state) after the main power supplies have been turned on. The Preset signal is relayed along the daisy chain of pedestal units until it reaches the end unit, which removes the Preset Return signal. The Preset Return signal is then passed back to the controller. The pedestal-to-frame Column Preset signal is derived from the Preset signal.

## D. Shield and Shield Return

The Shield signal can only be turned off (signal to the low state) if the main power supplies in all the pedestals are on. The Shield signal is relayed along the daisy chain to the end unit, and then returned to the controller as the Shield Return signal.

If a main power supply is turned off by a fault sensor in any pedestal, that pedestal asserts both the Shield and the Shield Return signals (signals to the high state). The pedestal-to-frame Column Shield signal is derived from the logical OR of Shield and Shield Return. Thus all frames will have Shield asserted, and the Shield Return signal at the controller will be in the high state, if any main power supply fails.



### E. System OK and System OK Return

The System OK and System OK Return signal lines indicate the condition of the pedestals. Before the Main Power signal is turned on, the System OK Return signal at the controller will not be asserted if there are any user-correctable undesirable system conditions. These conditions are:

1. Not enough power supplies are plugged into the pedestal.
2. Too many frame blocks are stacked on the pedestal.
3. Not enough AC power is being supplied to the pedestal (power cord not plugged in).
4. The cover cap for the service column connector is missing on the top frame block.
5. The plug-in supply handle cover door is not closed.

Each pedestal has a series of five lights which indicate if that pedestal has one or more of the above undesirable conditions. In addition, each pedestal has a sixth light which indicates that the System OK signal is not asserted.

When the System OK Return signal to the controller is asserted (high), all pedestals are ready for the Main Power signal.

After the main power supplies have been turned on, the System OK Return signal at the controller may turn off because of reason 3, 4, or 5 listed above or because a cooling fan is not working. If the System OK Return signal is turned off for reason 3, 4, or 5 listed, then the Shield Return signal was also asserted (high) because the main power supply in a pedestal unit was turned off. If a cooling problem exists, the pedestal power supply is not turned off, and thus the Shield Return signal is not asserted. The controller may then power down the system in an orderly manner when a cooling problem occurs.

### 6.2.5 Definition of Pedestal-to-Frame Signals

The pedestal-to-frame signals pass up through the service column bus located in the fan modules. The signals use different logic levels; some are MECL compatible, some are MHTL compatible, and some are analog signals.

#### A. Column Preset, Column Shield, Power Request

The Column Preset, Column Shield and Power Down Request signals are MECL compatible and are asserted high (-0.67V to -0.80V DC). The signals have slow rise and fall times (0.5  $\mu$ sec. to 1.0  $\mu$ sec.) and are terminated (AC termination) in the characteristic impedance of the signal bus to attenuate reflections and coupled external signals. In addition, the low state of the signals is -2.9V to -4.9V DC, to enhance the noise margin. Each signal is capable of a fanout of 256 MECL II gate inputs (100  $\mu$ A maximum input current per gate). The signal lines may be shorted to any voltage from zeroV to -5.2V DC without damage to the pedestal circuitry.

The Column Preset signal is asserted when the daisy chain Preset signal is high. The Column Shield signal is asserted when the daisy chain Shield or Shield Return signals are high or if the pedestal main power supply is off. The Power Down Request signal is asserted when the daisy chain Main Power signal into the pedestal is high.



### B. Power Down Acknowledge

The Power Down Acknowledge signal is a bus which is at -15V DC if no modules are connected. If a module will require power down sequencing before the 54.62V DC bus is turned off, then that module must clamp the Power Down Acknowledge but to zeroV DC when the 54.62V DC bus is turned on. When the Power Down Request signal is asserted, the module may sequence down its power internally before releasing the Power Down Acknowledge bus.

The switch that clamps the Power Down Acknowledge bus must be capable of supplying 4 mA when the switch is closed, and must supply no more than 4  $\mu$ A when the switch is open. A 2N5086 PNP transistor in a circuit which reverse biases the base-emitter junction by more than 0.10V DC when the transistor is cut off is an acceptable switch.

### C. Sense, Frame Sense, and Cap Sense

The Sense signal, the Frame Sense signal, the Cap Sense signal, the plug-in supply handle cover door signal, the plug-in supply sense signal, and the AC power sense signal form the inputs for the circuit which inhibits the turning on of the main power supply.

The Sense signal is a resistance between two signal lines, +SENSE and -SENSE. Each module has a resistor which is connected across those two signal lines when the module is plugged in. The resistor values are selected such that the resistance of the SENSE lines is a function of the power required by the modules that have been plugged in. The resistance value selected for each electronics package module is based on the maximum 54.62V DC input power that a module may require and a conversion factor of 1.85  $\mu$ mhos per watt. The Sense resistance is compared with the plug-in supply and resident supply power sense resistance. If there are sufficient supplies plugged in to handle the module load, this portion of the inhibit circuit will not be asserted.

The Sense signal network greatly improves the efficiency of the power system. The power requirements of each module and the number of modules plugged into any given column of frame blocks will vary a great deal. If the power system in each pedestal had to be capable of supplying the maximum power required by a column of frame blocks filled with the most power-hungry modules, the most modular systems would utilize only part, probably less than half, the available pedestal power. The power sensing network eliminates this problem by comparing the maximum load in the column of frame blocks with the available power in the pedestal. If a pedestal is overloaded, part of the modules can be moved to another pedestal and frame block column. This feature allows a pedestal power design which can only supply a typical mix of modules. If a lot of power-hungry modules are used in a computer system design, then a given pedestal frame-block column can only be partly filled.

The Frame Sense signal is a resistance between two signal lines, +FRAME SENSE and -FRAME SENSE. Each frame block has a register across these signal lines. The resistance of this bus is compared with a fixed resistor in the pedestal. If the signal line resistance is too low, this position of the inhibit circuit will be asserted.



The Cap Sense signal insures that the exposed vertical power bus connector 110V AC fan power contacts are covered before the main power is turned on. The connector cover (the "cap") contains a jumper which connects the Cap Sense bus to signal ground.

#### D. Cooling Sense

The Cooling Sense signal is asserted by grounding the cooling sense bus if one of the fans in the frame blocks fails. This signal will only inhibit the System OK Return signal (signal to the low state) if a fan fails.

#### 6.2.6 Pedestal Power Malfunctions

In most cases, if there is a power system malfunction in a pedestal, the main power supply in that pedestal will turn off, and that pedestal will assert the Column Shield, Shield and Shield Return signals. If the malfunction is operator correctable, the System OK signal will also be turned off. No other action will be taken by the pedestal. The Main Power and Main Power Return signals will remain on, and the Preset and Preset Return signals will remain off.

The controller is notified of a malfunction by a spontaneous Shield Return signal. (If a malfunction occurs between the time the controller turns off the Shield signal and gets a normal Shield Return off signal in reply, then the controller will never receive an acknowledgement for the removal of Shield. In this case, the controller could signal an error condition after a few seconds.)

If there is a cooling malfunction, the System OK signal will be turned off, but the Shield Return signal will remain low.

Some power system malfunctions trip the pedestal circuit breaker. This removes the power from the pedestal logic, which results in all of that unit's control signals reverting to their power off states. All pedestal units down the daisy chain from the faulty unit will also turn off. Units between the controller and the faulty unit are still under controller control except for the Shield, Shield Return, and Preset Return signals, which are asserted.

Most power system malfunctions in a pedestal are signalled to the console in the form of abnormal inputs from the pedestal daisy chain. A summary of possible patterns and their interpretation follows:

##### A. The System OK signal off:

The System OK Not signal indicates that one or more operator-correctable abnormal conditions exist or that there is a cooling malfunction. (See section 6.2.4-E.)

##### B. The Shield Return signal asserted with the Shield signal not asserted:

This condition means that a pedestal has autonomously turned off its main power supply. If the System OK signal is still asserted, the condition is



not operator-correctable. Before the Main Power signal is turned off, the faulty pedestal may be found by noting the unit with its fan module lights off.

C. The Shield Return and Preset Return signals asserted with Shield and Preset signals not asserted:

This condition means that a pedestal auxiliary supply has lost power. This can be caused by an AC power failure, a tripped circuit breaker, or a blown internal fuse. All pedestals down the daisy chain from the faulty unit will be turned off. The faulty pedestal will be the nearest pedestal along the daisy chain with its fan module power indicator lights off. The AC power line should be checked before replacing the pedestal. The remaining pedestals may be turned off by the console in the normal manner.

D. The Main Power Return signal does not turn off:

When turning off the main power supplies, the Main Power Return signal will not turn off at the controller until all main power supplies are off. The faulty pedestal column will be the farthest pedestal along the daisy chain with its fan module power indicator lights on. The problem may be the result of failure of a module to provide Power Down Acknowledge or a failure of the Power Down Acknowledge pathway in a frame block or in the pedestal. Removal of controller AC power will cause the Auxiliary Power signal to turn off, which will override the Power Down Acknowledge signal. If the main power supply in the faulty pedestal column now turns off, all the fan module lights will turn off. The fault is then in the frame blocks or the modules. If all the fan module lights do not turn off, then the main power supply in the pedestal is still on, and it will be necessary to pull the circuit breaker on the faulty pedestal. The remainder of the pedestals will then turn off.

6.2.7 Pedestal Controller Functional Description

The pedestal controller is that part of a macromodular console which deals with the pedestal control signals, both the power control signals and the system-wide Preset and Shield signals. However, with logic level interfacing, the controller portion of a console may also function as a stand-alone control system for the pedestal daisy chain signals.

The pedestal controller generates and responds to the pedestal control signals in a manner that assures no change in macromodular data (if no macromodular control is active when the pedestal control is active). The controller functional description is constrained both by the sequencing requirements of the pedestal and by the user.

Under normal operating conditions, the console should originate signals in a prescribed sequence, and wait for proper acknowledgement of each signal from the daisy chain of pedestal units before initiating the next signal in sequence. The normal power on sequence is:



1. Turn controller power on;
2. Turn Auxiliary Power on (low);
3. Wait for Auxiliary Power Return on (low);
4. If System OK Return is asserted (high), wait 1.5 seconds, then turn Main Power on (low);
5. Wait for Main Power Return on (low);
6. If System OK Return is still asserted after three seconds, turn Preset and Shield off (low);
7. Wait for Preset Return and Shield Return off (low);
8. After a few microseconds wait, assert the controller System Ready signal.

Power is now on and the macromodular system may be used. The normal power off sequence is:

1. Stop all macromodular logic signals.
2. Turn Shield on (high). (Removes controller System Ready signal.)
3. Wait for Shield Return (high).
4. After a few milliseconds, turn Preset on (high).
5. Wait for Preset Return (high).
6. Turn Main Power off (high).
7. Wait for Main Power Return (high).
8. Set single shot so Main Power cannot be turned back on in less than two seconds.
9. Turn Auxiliary Power off (high).
10. Wait for Auxiliary Power Return (high).
11. Turn Controller Power off.

The controller must also ensure some interacting of daisy chain signals, and must react to some abnormal conditions.

#### A. Console Power

The controller must ensure that its Auxiliary Power output signal remains in the off (high) state when it is not energized, and during the turn-on and turn-off of console power.

#### B. Control Sequence and Timing

The complete power and control sequence has been given. The sequence may be reversed at any point when the signals and their returns are equal.

The controller must ensure that the main power control logic in the pedestal has more than 1.5 seconds to stabilize after the small auxiliary supplies in each pedestal have turned on, and has more than two seconds to recover after the main power supplies have been turned off.

The Preset Return and Shield Return signals must be delayed by the controller before they are allowed to produce other control commands, by the amount of time necessary for each pedestal to react to the signals. This time may be a few milliseconds after assertion of the signals, and tens of microseconds after removal of the signals.



### C. Power Malfunctions

In addition to receiving return signals during power and control sequencing, there are two other daisy chain control signals which inform the controller of pedestal conditions. These signals are System OK Return and autonomous assertion of Shield Return.

To assure no additional physical damage to pedestals or modules, the presence of a low System OK Return signal must result in the Main Power signal turning off (high) in less than one minute. A low System OK Return signal may mean that modules and/or pedestal supplies are not being cooled.

#### 6.2.8 Controller User Description

To put section 6.2 into proper perspective, a user's guide for controlling the macromodular power system-wide control signals is given below.

The controller provides an interface between the macromodular control cables and the daisy chain cable, a set of lights to indicate the state of the daisy chain control signals, and four toggle-switch governed control signal ports. (See Figure 71.)

The four toggle-switch governed control signal ports may be used to control the four daisy chain initiate signals (they are labelled for that use). The four switch ports, however, are each independent and have switch bounce elimination circuits, and thus may be used as macromodular control signals. The toggle-switch governed ports are in the preset state in the "stop" position.

There are nine lights, which indicate the four controller-initiated signals and their four returns and the System Ready For Main Power signal line. The System Ready For Main Power light indicates operator-correctable problems. If, after the Auxiliary Power On Initiate and Return lights are on, the System Ready For Main Power light is not on, then some of the correction indication lights on the back of the pedestals are on. (If the Auxiliary Power Initiate light is on and the Auxiliary Power Return light is off, then either one or more pedestals have no AC power cable plugged into the line 1 receptacle, or a circuit breaker is off.) The pedestal correction indication lights will describe the action required to turn the lights off. This action may be carried out without turning Auxiliary Power back off. When the System Ready For Main Power light on the controller is on, Main Power may be turned on. If Main Power is turned on when the System Ready for Main Power light is off, System Condition Light #1 will turn on. Turn the Main Power switch back to the stop position, and proceed to find out why the System Ready for Main Power light is off.

The four daisy chain control signals that may be initiated by the controller are interlocked to force a proper sequence of control to the pedestals. The sequence from "stop" to "run" is: Auxiliary Power on, then Main Power on, then Preset off, and then Data Shield off. The sequence from "run" to "stop" is: Data Shield on, then Preset on, then Main Power off, and then Auxiliary Power off. The System Ready output signal is asserted (to the non-preset state) only when all eight Initiate and Return control signals are



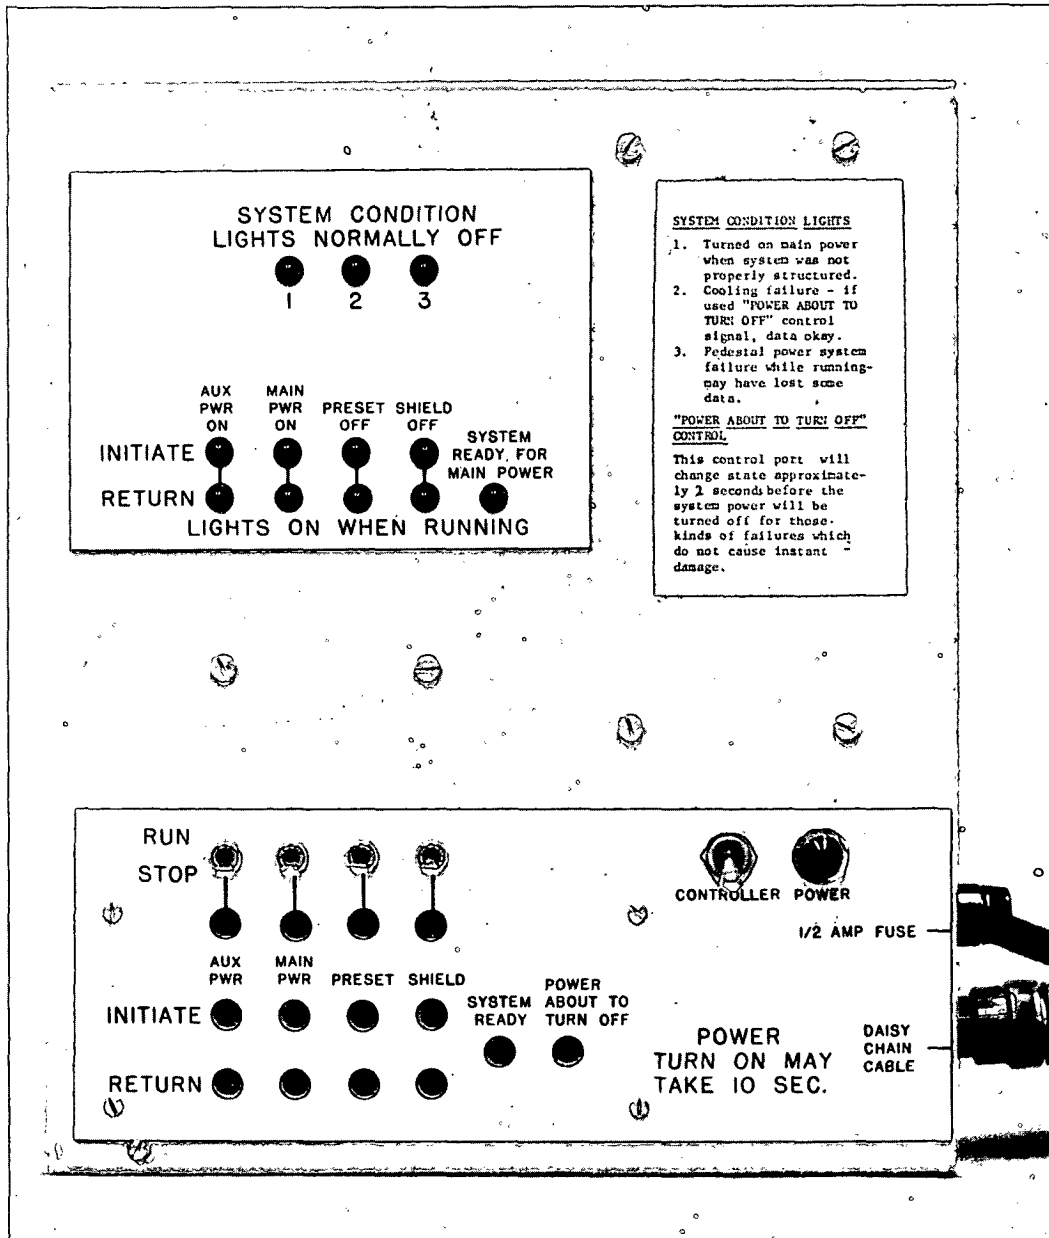


Figure 71. Macromodular system controller.



in the "run" state. (The "System Ready" signal and the "System Ready for Main Power" light are two independent signals.)

There are two typical ways the controller has been used. The first way is to use the controller to turn on power, preset the macromodules, and then remove Preset and Data Shield. This is accomplished by jumpering all four toggle switch ports to the four initiate ports. Before turning on the controller power, set the Main Power, Preset and Data Shield switches to the "run" position and the Auxiliary switch to "stop". Turn on the controller power and then throw the Auxiliary toggle switch to "run". The controller will turn on Auxiliary Power, wait until it receives an Auxiliary Power return, then turn on Main Power, wait for a return, turn off Preset, wait for a Preset off return, then turn Data Shield off, wait for a Data Shield off return, and then assert the System Ready signal.

The macromodular system may now be preset by simply moving the Preset switch to the "stop" position. (If the macromodular machine control is running, this action may destroy data). The controller will turn Data Shield on, then wait for a Data Shield Return, wait an additional 6 msec., then turn Preset on. The System Ready signal goes back to the preset state when Data Shield is turned on.

A second way the controller has been used employs the LINC interface to operate Preset and Data Shield. The Auxiliary Power, Main Power, and Data Shield initiate ports are jumpered to toggle switch ports. The Preset initiate and return and System Ready return ports are connected to the LINC interface. The Main Power and Data Shield toggle switches are set to the "run" position. The Preset initiate line from the LINC interface is set to a zero. The Auxiliary Power toggle switch is then set to the "run" position. The system then has power.

The Preset control signal may now be turned off by the LINC interface by setting the interface Preset bit to a one. A System Ready signal must be received at the interface (the appropriate bit read as a one) before the macromodular system is ready to run. When Preset is turned back on by the interface, a Preset Return signal must be received before the rest of the interface-initiated control signals are preset. (The Data Shield Return signal will turn on before the Preset Return signal.) Preset can then be turned off. The Preset signal will remain on in the pedestals for approximately another 6 msec. before it is removed.

The controller has an additional output signal control port and 2 additional system lights. The control port, Power About to Turn Off, is useful if the user wishes to save his data in case there is a cooling failure while his program is running. When there is a cooling failure, the Power About to Turn Off signal will go to the non-preset state 2 seconds before the daisy chain Main Power signal is turned off. The user can wire the Power About to Turn Off signal into his system as an interrupt (using an Interlock module). The processor then has 2 seconds to store in non-volatile memory whatever information is needed to restart. System Condition Lights 2 and 3 then inform the user whether or not the power shut-down was the result of a cooling failure or some other type of power failure. A note on each controller describes the functions of the System Condition Lights and the Power About to Turn Off control port.



### 6.3 POWER AND SYSTEM-WIDE CONTROL SIGNAL STRUCTURE DESIGN CONSIDERATIONS

This section deals with the design of the power system used in macro-modules. The electrical separation of the signal ground and the power distribution network ground is discussed, along with some details of parts of the power system which may be of use in other systems.

#### 6.3.1 Isolation of Power and Signal Grounds

First we consider the basic structure of the power distribution and signal ground networks. The signal and power grounds are interconnected at a single common point in the base pedestal of each frame block column. The signal ground net is connected to this common point via: 1) 3 signal ground lines that pass up through each fan module via the service column bus; 2) 2 connections to the base pedestal frame; 3) a connection to the base pedestal daisy chain cable shield; 4) 4 connections to the base pedestal control logic. The more negative output terminal of each of the 54.62V DC base pedestal power supply units is connected to this common point, as is the negative sense lead of the power supply. From this common point, the power ground, or the zeroV line, as it is called, is passed up through the fan modules via 10 number 16 wires in the service column bus. The zeroV line is treated as a power line beyond this common point, and is never again connected to the signal ground net.

The two AC power input line safety grounds are also connected to this common point through a rectifier bridge with a shorted DC output, as shown in Figure 72, to allow the various AC power line safety grounds to "float" with respect to each other up to  $\pm 1V$  without inducing significant currents in the macromodular signal ground net. These rectifier bridges are capable of passing AC power line fault currents, up to a maximum expected current of 200 amps, long enough to open a 20-amp circuit breaker without damaging the bridges. It should be noted that this bridge arrangement means that an AC power line fault in one base pedestal may pass through the pedestal daisy chain cable shield, then through a rectifier bridge in another base pedestal with a lower forward voltage drop to ground. Therefore the daisy chain cable shield and the daisy chain connectors must be capable of passing 200 amps for a few milliseconds without damage.

The details of the common point connections are shown in Figure 73. The large crimp containing 13 number 16 wires in one end and 15 number 16 plus 4 number 24 wires in the other end with a "4W" in the middle of the crimp is the common point.

Beyond this common ground point in each base pedestal, the zeroV and the +54.62V DC power lines pass through each fan module via the service column bus with taps into each frame block channel, and then into each module through the lateral connector to a DC to DC converter, which transforms the 54.62V DC to the voltages needed by that particular module and isolates the 54.62V DC bus from the signal and local power ground.

Thus the power and signal ground currents share a common path only in the confines of a module electronics package, except for some high-frequency







Figure 72. Power rectifier bridge arrangement.







currents created by the DC to DC converter. There is some current coupled through the converter power transformer by the capacitance between the primary and secondary windings. This current must travel through the signal ground net to the base pedestal common ground point, then through the zeroV line to the DC to DC converter primary. Later DC to DC converter designs, including the D/A module, the Memory module, the Memory Controller module, and the Multiply module converters, incorporated a small-value bypass capacitor connected between the center taps of the primary and secondary windings of the converter power transformer. The high-frequency 54.62V DC bus to signal ground isolation was improved in these modules by placing inductors in series with both the +54.62V DC line and the zeroV line in each module. The improvement may be seen by comparing Figures 74 and 75. In Figure 75 the added bypass capacitor is C19 and the input inductors are L1 and L2.

Isolation of the AC power input lines from the common point ground is provided by a small high-frequency input filter for the AC power input and by the low primary to secondary coupling capacitance of the ferroresonant transformers in the base pedestal supplies. The coupling capacitance of a fully loaded pedestal capable of supplying 1850 watts is about 400 pfd. The ferroresonant transformers also greatly attenuate difference mode AC power line noise.

### 6.3.2 Base Pedestal Power System

The power system consists of the base pedestal, the power distribution network, and the module DC to DC converter. The power distribution has been discussed. The base pedestal contains a set of power supplies capable of supplying 1850 watts to up to 8 frame blocks, a load sensing network, and a digital control section.


#### The Power Supply

The base pedestal power supply is made up of 4 power units, which can each supply about 460 watts. Each power unit contains a ferroresonant transformer, a rectifier bridge, a capacitor filter, and a series regulator power stage. The resident supply also contains the circuitry necessary to control the regulator power stages in all the power units in a manner that assures not only the proper bus voltage but also current sharing between the power units. Current limiting and over/under voltage detection are also provided by the resident supply circuitry.

One of the units (resident) is built into the pedestal, and the other 3 (slugs) plug in. The power supply is modularized for two reasons. The first is so that the base pedestal, with the 3 plug-in supplies removed, can be lifted by two men. The second reason is that a single 120V, 20-amp AC power line can only power 2 power units. The base pedestal logic is arranged so that a base pedestal can be operated using only the resident, or the resident and 1 plug-in supply, with only one AC power line.

A ferroresonant transformer series regulator scheme was chosen over an isolation transformer L-C filter switching regulator scheme [5] for several reasons. In addition to the lower primary to secondary coupling capacitance obtainable because ferroresonant transformer windings are physically separated, the ferroresonant transformer provides some voltage





5	2-15-74	ECO 0305	<i>Age</i>	
CHANGE NO.	DATE	DESCRIPTION		
<p align="center"><b>COMPUTER SYSTEMS LABORATORY</b>          WASHINGTON UNIVERSITY          ST. LOUIS, MISSOURI</p>				
<p align="center"><b>MACROMODULAR PROJECT</b></p>				
TITLE <b>15 WATT DC/DC 5 VOLT CONVERTER SCHEMATIC</b>				
APPROVED BY <i>TJC</i> FOR <i>TJC</i> DATE			END. <i>TJC</i> DRAWN BY <i>PLL</i> CHECKED <i>Age</i>	
			DRAWING NO. <b>200.5 D3</b> DATE <b>10-31-68</b>	

-162-



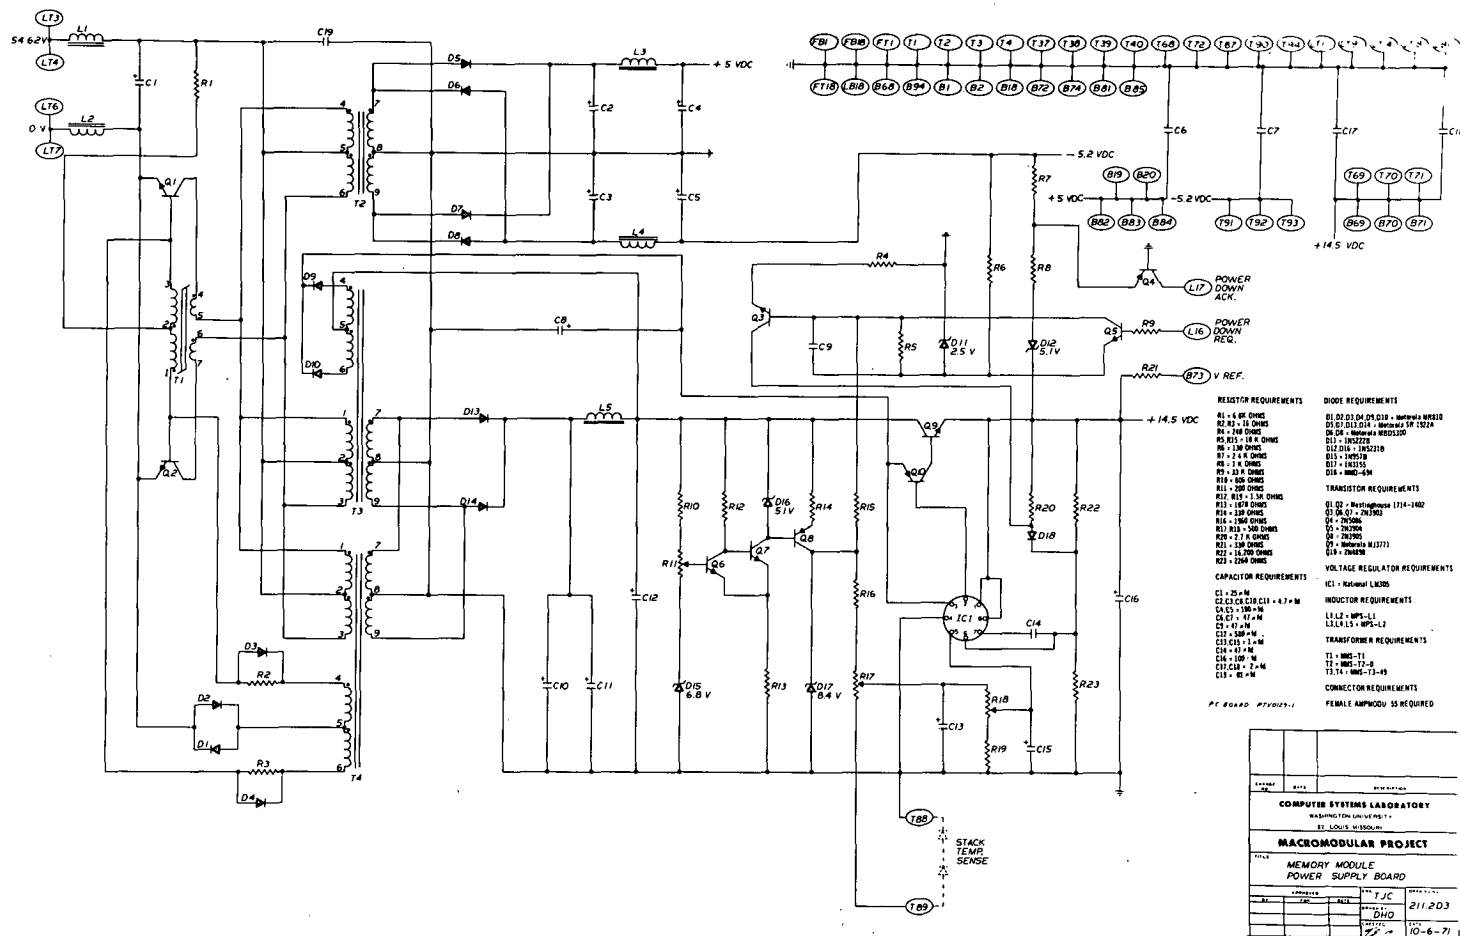


Figure 75. Memory module power supply board.



regulation and a good AC input line power factor in a unit smaller than an isolation transformer and filter inductor assembly of equal power rating. The coarse voltage regulation of the transformer allowed the design of a series regulator with a full load efficiency close to that of a switching regulator. A series regulator can respond to load changes faster than can a switching regulator; in addition, there are no high-frequency voltage spikes to contend with in a series regulator design.

Two negative comments must be made about the ferroresonant transformer. The first is acoustic noise. Ferroresonant transformers have a hum and accompanying vibration that can be quite loud, especially if the transformer is mounted such that the vibration is coupled to the structure. The base pedestal transformers are mounted on vibration isolators and cannot be heard over the cooling fan noise. The second comment has to do with the transient load response of the ferroresonant transformer. The series regulator design used in the base pedestals is capable of functioning properly with as little as 3V DC across the series regulator power stage at full load. The transient response of the transformer to a load change of 25% is such that an additional 5 volts must be allowed across the regulator power stage so the regulator will continue to function properly when the transformer output voltage dips some 50 to 200 milliseconds after the load change.

Even allowing for these problems, the base pedestal power system has an input power factor of 0.9 and an efficiency of 72% at full load. The 0.9 power factor is achieved in part by the incorporation of power factor correction capacitors in each base pedestal.

#### A.C. Power Switching

The turning on and off of high-power equipment in a manner that does not cause neighboring equipment to malfunction was considered an important problem. To achieve reasonably limited AC power line inrush currents during turn-on, solid-state AC switches were used so that the timing of switch closure with respect to the 60-Hz line waveform could be controlled. The load was divided into capacitive and inductive loads. The capacitive load was switched on during zero-volt voltage crossover of the AC line and the inductive load was switched on at a voltage peak of the AC line, for reasons presented in the Fairchild handbook [6]. If an inductive load is switched on at a zero-volt voltage crossover, the magnetic core of the inductor (in this case, the ferroresonant transformer) may saturate during the first half cycle. If the core does saturate, then the line current can rise very rapidly. A base pedestal power unit was tested with a zero volt crossover voltage turn-on applied; the first half cycle current exceeded 100 amperes. The same unit with a peak voltage turn-on had a maximum input current of under 20 amperes. The circuit that controls the base pedestal AC power line switches may be found in Figure 76.

#### 6.3.3 Module Supply

The module supply consists of a DC to DC converter that transforms the 54.62V DC power supplied by the pedestal into the voltage or voltages







required by the module, and provides DC isolation between the 54.62V bus and the signal ground. The first and most widely used converter is described in detail in [5]. A more detailed discussion of this type of converter may also be found in the Restructured Macromodules section (Part 3) of this report. This converter is used in modules that require from 7.5 to 15 watts of -5.2V DC power. This same converter design, using power transformers with slightly different turn ratios to compensate for the different load requirements of the modules, resulted in a set of 3 DC to DC converters that are used to provide power for all the single-height modules except the D/A module.

The D/A module DC to DC converter is the standard design described above with additional secondary windings on the power transformer to provide all the voltages needed in the module. Except for the problem of switching noise coupled through the power transformer, this converter was a straightforward extension of the standard design.

There are three modules that require more than 15 watts of power. These are the Memory, the Multiply, and the Memory Controller modules, which are all double-height modules and require from 40 to 80 watts of power. The DC to DC converter designed for these modules is a scaled-up version of the converter design, with some improvements. The core used for the standard timing transformer is a tape-wound core, which has a much tighter tolerance on the saturation flux. The switching transistors are much faster than the ones used in the 15-watt converter. The power transformer could not be made larger because of physical constraints in the existing module outline, so up to 3 transformers connected in parallel are used to provide the needed power. In addition, each transformer can pass more power because a higher converter frequency is used, from 20 to 30 KHz. An additional power saving was made in some modules by using Schottky type rectifiers, which have a lower forward voltage. The higher-power DC to DC converters include inductor isolation from the 54.62V DC power bus and a power transformer bypass capacitor.

The Multiply and Memory Controller modules use ECL logic exclusively, and thus the converter output is the standard rectifier-filter type circuit. The Memory module, however, requires +14.5V DC regulated power for the core memory stack, -5.2V DC power for the ECL circuits, and +5V DC power for the TTL core stack driver circuits. The memory module DC to DC converter must also assure that the +14.5V DC stack driving supply turns on last and off first to assure that ECL and TTL circuits don't cause data changes in the core stack during turn-on and turn-off of module power. The -5.2V DC and +5V DC power is derived from one center-tapped transformer by using semiconductor junction fast recovery diodes to produce the +5V DC voltage and Schottky type power rectifiers (which have a lower forward voltage drop) to produce the -5.2V DC voltage. Two other transformers are connected in parallel to provide the power necessary for the +14.5V supply. The +14.5V supply is regulated using a series regulator, and the power up and down sequences are controlled by clamping the reference voltage for the +14.5V DC regulator to a low value. The power down sequence utilizes the Power Down Request and Power Down Acknowledge signals (see section 6.2.3) to assure that the +14.5V DC supply has reduced to a safe value (less than 8V) before the 54.62V DC bus



is turned off in the base pedestal. The +14.5V DC regulator includes a memory core stack temperature feed-back via 2 diodes mounted in the stack cooling air stream, so that the +14.5V DC supply can vary its output voltage to compensate for the stack temperature.



## 7. INTERFACING

### 7.1 INTRODUCTION

The ease with which modules may be interfaced with other computer hardware is a very important consideration. Modules can perform few meaningful tasks standing alone because of their lack of I/O capability. The following paragraphs discuss techniques used to interface various hardware to the modules. Equipment which has been interfaced to the modules includes: 1) LINC computer; 2) PC computer; 3) IBM System 7 computer; 4) DEC PDP-11/40 computer; 5) EVANS & SUTHERLAND LDS-1 Matrix Multiplier and Line Drawing Scope; 6) DATACRAFT Scientific Arithmetic Unit; 7) FABRI-TEK memory systems; 8) DEC Register Transfer Modules; 9) movie camera controller for making movies of molecular graphics; 10) A/D converters.

Before a discussion of the details of macromodule interfacing, one byproduct of interfacing the abovementioned equipment to the modules should be noted. Interfacing to the modules has provided a standardized interconnection structure which allows various devices to be connected, not only to the modules, but directly among themselves. This has proved useful on a number of occasions. A specific example is the interconnection of the LINC and the EVANS & SUTHERLAND display system. Another potentially desirable combination would be either the LINC or the PDP-11 to the Scientific Arithmetic Unit.

The majority of the following discussion is devoted to the PDP-11 interface. It was chosen for several reasons: 1) it is the most recent design, and incorporates features conceived through experience with previous interfaces; 2) it is one of the most flexible interfaces - having the ability to initiate macromodule control sequences and thus being a controlling device, and the ability to be initiated by the modules, thus being a controlled device; 3) the PDP-11 is the most widely known device among those interfaced to the modules. This interface actually serves both the Phase I modules and the restructured macromodules (see Part 3 of this report). Most aspects of this discussion, however, refer to either set of modules. Where reference is occasionally made to features not relevant to Phase I macromodules, the text will so indicate.

A question often asked by someone considering how to build a macromodule interface is how to handle the conversion between transition logic and the more conventional level or pulse logic. The answer is provided by an EXCLUSIVE OR of the module initiation and completion signals. The macromodular Preset signal initially insures that all control lines (initiation and completion) are at the same logic level. A transition on the input control line to a module signals it to perform its operation. When this input control transition occurs, the EXCLUSIVE OR of the input and output control lines will yield a high level. Upon completion of the operation, the module's output control line will change state, so once again the input and output lines are the same and their EXCLUSIVE OR is low. Thus it is not the absolute level of either the input or the output line, but



their relative state, which conveys the information. Figure 77 illustrates how this information can be used to convert between transition logic and level logic.




Figure 77. Conversion between transition and level logic.

First consider the use of this circuit when the modules initiate an operation of the interfaced device. Initially, the L.L. OUT (LEVEL LOGIC OUT) line is low. When the T.L. IN line switches, the L.L. OUT line will go high. The positive transition of this line may then be used to signal the interfaced device that it has been asked to perform its operation. Upon completion of the operation, a pulse on the L.L. IN line will toggle the flip-flop, cause the T.L. OUT line to go back low, and indicate a completion on the T.L. OUT lines. If the interfaced device is to initiate a module operation, a pulse on L.L. IN will initiate the sequence, and the negative transition of L.L. OUT will indicate completion of the sequence.

Preset is a system-wide signal which must be handled properly by any interfaced device. If Preset is not handled properly, data may be lost, or some parts of the macromodular system may not be preset, or both may happen. A detailed discussion of the macromodular Preset signal has been covered in the section on the base pedestal controller (Section 6.2). Here we will simply present the minimal information required to be able to handle and react to Preset properly.

If the interfaced device is merely to receive Preset from the modules or some other device, all that is required is for the device to receive the signal on its Preset In line, react to it (force all control out lines to the Preset state regardless of the state of the inputs and preset internal states), and then pass the signal on to its Preset Out line. If the interfaced device is to initiate a macromodular Preset sequence, the process is slightly more involved. The sequence is:



- 1) The interfaced device asserts Preset on its Preset Out line.  
(The interface does not yet preset its other control out lines.)
- 2) After receipt of the Preset Return signals, the device presets its control out lines and internal states and releases its Preset Out line.
- 3) When the assertion of System Ready by the pedestal controller is detected by the interfaced device, the entire system has been preset, and the device may then initiate the first control sequence.

If the interfaced device has data which must be preserved during a Preset sequence, the Shield signal from the pedestal controller may be used to indicate that a Preset assertion is forthcoming. Upon receipt of Shield, the interface must inhibit data from changing, to preserve it against spurious control signals generated during Preset. Once the device has properly reacted to Shield, it may pass Shield out; and upon receipt of Shield Return by the controller, the controller will then assert Preset.

The word "data" at the beginning of the preceding paragraph can mean more than just digital data stored in a register. It could just as well be the physical state of a mechanical device interfaced to the modules. Consider a movie camera taking pictures of a display on a CRT. A reasonable control sequence would be for the modules to advance a frame of the movie camera, open the camera shutter, display a frame on the CRT, and then close the shutter again. If the interface for the movie camera did not respond to Shield, a transition due to Preset on one of its control in lines could be interpreted as a request to advance a frame, instead of ignored because it was merely that line returning to the Preset state. Shield (also referred to as Data Protect) is included in macromodular systems so that such a problem may be avoided.

We will now proceed to the discussion of the PDP-11/40 interface. A more detailed explanation may be found in [7].

## 7.2 PDP-11/40 TO MACROMODULE INTERFACE

### 7.2.1 Interaction of the PDP-11/40 with Other Devices

The PDP-11/40 [8] communicates with peripheral devices through the UNIBUS. The UNIBUS consists of various open-collector signal lines which carry a logic "1" when pulled to a low voltage (approximately .5V), and a logic "0" when held at a high voltage (approximately 2.4V) by a resistive divider at each end of the signal line. The signal lines are divided into several groups according to their function. These include address lines (A<17:00>), data lines (D<15:00>), control lines (MSYN, SSYN, C0-C1), and various other signal lines which are not presently relevant to the discussion.

Communications on the UNIBUS are accomplished through a master-slave relationship among devices. The master controls the UNIBUS when communicating with another device (the slave). Generally, the processor is the master device. However, another device may gain control of the UNIBUS to access other devices or to interrupt the processor. The interface



between the PDP-11/40 and the macromodules usually functions as a slave device, but it is capable of interrupting the processor.

When the PDP-11 is the master device, it is capable of executing four basic types of data transfers. These are determined by the levels on the C0 and C1 control lines. They are described as follows:

<u>NAME</u>	<u>MNEMONIC</u>	<u>FUNCTION</u>
Data in	DATI	Data from slave to master.
Data in, pause	DATIP	Inhibits restore cycle in destructive read-out devices. Not used in connection with the interface.
Data out	DATO	Data from master to slave.
Data out, byte	DATOB	Data from master to a single byte in slave. Data transferred on D<15:08> for A00=1 and D<07:00> for A00=0.

Each device connected to the UNIBUS is assigned addresses corresponding to one or more bytes in the PDP-11 External Page Addresses (EPA). In order to address a device, the processor establishes the device address on A<17:00>. Each device monitors A<17:00> to determine if it is being selected. If so, the assertion of MSYN signals the device to perform the operation indicated by C0-C1. The slave device asserts SSYN to indicate that it has recognized its selection and is performing the desired operation. The processor cannot resume operation until the slave device removes its assertion of SSYN. Thus, the device may delay the resumption of processor operation as long as necessary for the completion of its task.

### 7.2.2 Summary of Interface Description; Conventions

The interface has four functional sections, which are outlined below. More detailed descriptions of each section follow.

#### Data Output Ports(PDP-11 → MM)

There are eight data output ports which receive words or bytes from the UNIBUS and transfer 12 or 16-bit words via a data cable to macromodules. Each port has an associated pair of control connectors. Each time a data transfer is directed to a port, a control signal may be sent out, and the UNIBUS will be released only when a control return is received.

#### Data Input Ports (MM →PDP-11)

There are eight data input ports which allow the UNIBUS to access the 12 or 16 data bits from a data cable. The associated pair of control connectors allows the PDP-11 to sense the presence of a macromodular control signal to initiate data transfers. If the proper ENABLE byte is set, the presence of a macromodular control signal will cause a vectored interrupt to the PDP-11.



### Concurrent Control Ports (MM ↔ PDP-11)

These eight pairs of control connectors allow the PDP-11 to sense the presence of macromodular control signals, or the PDP-11 may initiate macromodular control signals. The UNIBUS is released immediately on either a "read" or a "write" operation to these ports.

### Pausing Control Ports (MM ↔ PDP-11)

Six pairs of control connectors are wired to initiate and receive macromodular control signals, with one difference from the concurrent control. When the PDP-11 initiates a control signal, the UNIBUS is frozen until the corresponding control completion is received.

Two other functions are of interest to this outline.

### Preset

The macromodular Shield-Preset sequence may be initiated in two ways; any PDP-11 console START or RESET instruction will preset the macromodules, or the PDP-11 program may write into a byte address and cause a macromodular Preset without initializing the UNIBUS.

### Function Calling

For convenience when restructured macromodules are used, the interface provides a feature to allow function calling. When the restructured data cable is used in the Function Caller mode, the NO completion sets a particular byte to all ones (377<sub>8</sub>).

The PDP-11 interface is assigned 100<sub>8</sub> contiguous bytes in the EPA. All of these bytes have the same value of A<17:06>. The interface monitors only A<17:06> to determine if it is being selected. The value of A<05:00> is used to determine which byte is being addressed within the interface address space. The PDP-11 may execute a DATI to any byte in the interface address space. The data received may be data sent by the macromodules, the status of a control port, the status of an interrupt ENABLE byte, or the status of a previously initiated Preset sequence (these will be described later).

A DATO or DATOB may also be executed to any byte in the address space. However, in some cases such an action is meaningless (as in the case of sending data to an input data port); in these cases the interface merely returns control to the processor without performing any functions. The meaningful cases include sending data to the macromodules, sending out a control signal, starting a Preset sequence, and setting or clearing an interrupt ENABLE byte.

Before describing the interaction between the interface and the modules, the establishment of several definitions is necessary. A Control Line refers to a control cable joining the output (sending) connector of one module to the input (receiving) connector of another module. The action of sending a transition signal corresponds to switching the voltage level on the cable conductor. A Control Port, as considered here,



consists of two connectors which each accept a control cable. The Control Port may receive a transition signal over a cable inserted into the upper connector (input) and may send a transition signal over a cable inserted into the lower connector (output). Thus, each connector corresponds to a unidirectional control path. The input and output control lines are said to "match" if the levels on the cables are the same.

The interface is designed so that each Data Port will accept a data cable from either the original modules or the restructured modules. Except when explicitly noted, the following discussion applies to both types of modules.

### 7.2.3 Control Ports

There are two types of Control Ports used for communicating between the PDP-11 and the modules. They are referred to as Concurrent Control Ports and Pausing Control Ports. The Concurrent Control Ports allow the most flexible control linkage between the PDP-11 and the modules. The status of a Control Port refers to whether the Control Lines are matched, and is indicated in two ways. First, a light on the front panel of the interface is illuminated if the Control Lines do not match (this aids the tracing of control sequences by the user). Secondly, if the PDP-11 performs a DATI from the byte address assigned to the Control Port, the data received will consist of all zeroes if the Control Lines are matched, and all ones if they are not matched.

Each Concurrent Control Port is bidirectional in the following sense. In the first case, with the Control Lines initially matched, the PDP-11 performs a DATO or DATOB to the CONTROL PORT byte, and the interface responds by sending a transition signal on the output Control Line. The Control Lines are now mismatched. The PDP-11 immediately resumes operation, and may test the Control Port status at any time to determine whether a module has sent a completion signal on the input Control Line. The Control Lines will remain mismatched until such a completion signal arrives. In this case the PDP-11 is acting as the controlling device. In the second case, with the Control Lines initially matched, a module sends a transition signal to the interface on the input Control Line. The Control Lines are now mismatched, and the PDP-11 may recognize this condition by testing the Control Port status. After the PDP-11 has responded to the signal from a module, it may perform a DATO or DATOB to the CONTROL PORT byte, causing a completion transition signal to be sent to the modules on the output Control Line. In this second case, a module is acting as the controlling device.

The circuit used to accomplish the interaction described above is shown in Figure 78. The logic elements employed include an EXCLUSIVE OR gate, two translators for conversion between TTL levels and MECL levels, and a D type flip-flop. The flip-flop may change state only on the positive-going edge of the C input. At this time, the Q output is set to the value of the D input. A low level on the S input sets Q to a high level. In the configuration shown, the flip-flop switches state on every positive-going transition of SEND H. Thus, it is equivalent to the toggle flip-flop used earlier to describe the conversion between transition logic and level logic.



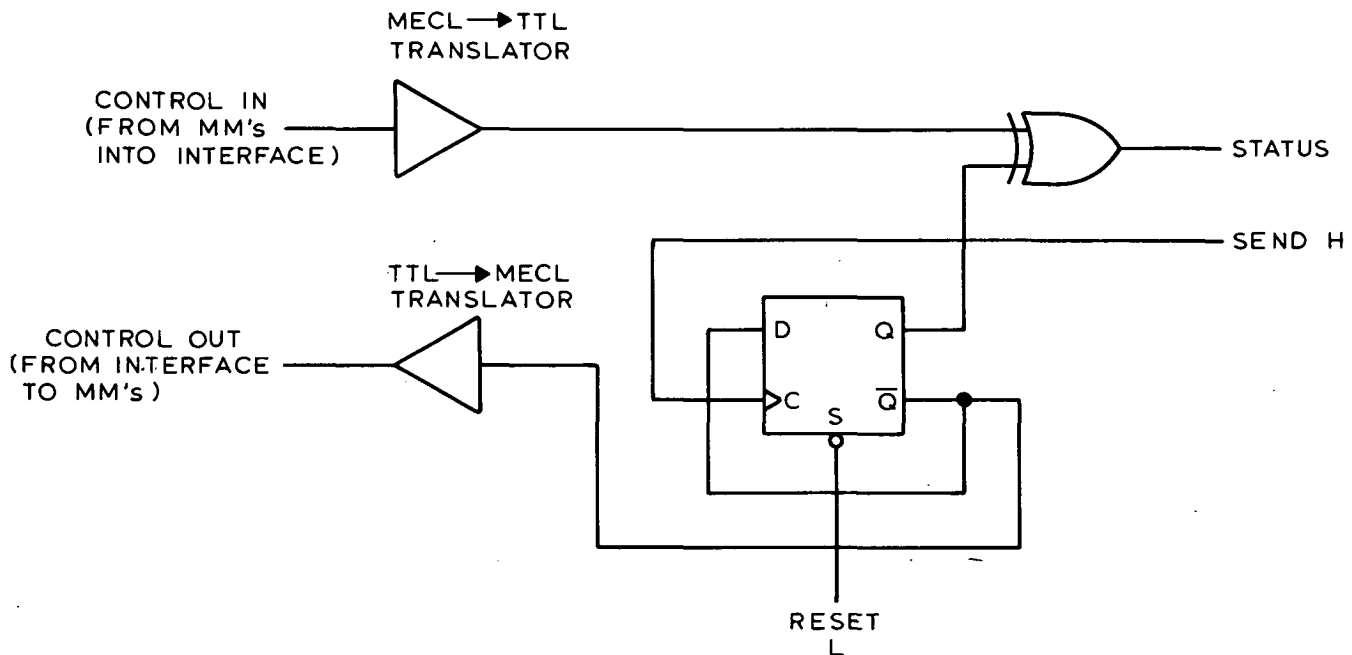


Figure 78. Control circuitry for Concurrent Control Port.

When the modules are being preset, PRESET L is low (the PDP-11 initiates all Preset sequences in the modules) and CONTROL IN and CONTROL OUT are both low. This makes STATUS high, indicating that the CONTROL lines are matched. The STATUS line controls the light on the front panel, and is used to determine the data sent to the PDP-11 when a DATI is performed from the CONTROL PORT byte. Note that STATUS is low when the CONTROL lines do not match. A pulse is sent on the SEND H line whenever the PDP-11 performs a DATO or DATOB to the CONTROL PORT byte. This toggles the flip-flop and changes the level of CONTROL OUT and STATUS. A corresponding change in the level of CONTROL IN will restore STATUS to its original level.

Several points should be made about the above circuit. The MECL-TTL translator circuit is designed to yield a low output level whenever a control cable is not inserted in the input control connector. This assures that the CONTROL IN signal will appear to be at the Preset level when an input cable is missing. There is no actual requirement that each transition signal on the output Control Line be accompanied by a transition signal on the input Control Line. Consequently, the PDP-11 may send repeated transition signals on the output Control Line without waiting for a signal on the input Control Line. Similarly, the modules may send repeated transition signals on the input Control Lines. Of course, most macro-modular systems operate with the restriction that each signal into a module is accompanied by a signal sent out of the module. Another point to be noted is that the interface is unable to distinguish the presence or absence of a control cable in either of the Concurrent Control connectors. In the case of other types of control, the absence of a cable must be recognized to help prevent system failures.

#### Pausing Control Ports

The Pausing Control Ports are very similar in basic operation to the Concurrent Control Ports. Each Control Port has been assigned to it a byte address in the interface address space. The status of the Control



Port is indicated by a light on the front panel, and may be obtained by the PDP-11 when a DATI is performed from the CONTROL PORT byte. However, the Pausing Control Ports are different in one major respect. Suppose cables are inserted in both the input and output connectors, and the Control Lines are matched. The PDP-11 may perform a DATO or DATOB to the CONTROL PORT byte, causing a control signal to be sent on the output Control Line. The interface now maintains the assertion of SSYN, halting any further operation of the processor. The SSYN signal is not negated until a completion signal arrives from the modules on the input Control Line. This mode of operation obviously corresponds to having the PDP-11 act as the controlling device.

As in the case of the Concurrent Control Ports, the Pausing Control Ports may be operated with a module as the controlling device. With the Control Lines initially matched, a module sends a transition signal on the input Control Line. This causes the Control Lines to become mismatched, and the status of the Control Port is changed. The PDP-11 may perform a DATI on the CONTROL PORT byte at any time to determine the Control Port status. After responding to the arrival of the signal from a module, the PDP-11 must perform a DATO or DATOB to the CONTROL PORT byte, causing a completion signal to be sent to the modules on the output Control Line. In this case, the interface allows the processor to continue operation immediately. Note that the Pausing Control Ports require that each transition signal on the output Control Line be accompanied by a transition signal on the input Control Line.

The circuit shown in Figure 79 accomplishes the function of the Pausing Control Port. This circuit includes an additional feature which was not described above. The presence of the input control cable is sensed to generate the CABLE PRESENT H signal. If the input control cable is not present, it would never be possible to receive a signal from the modules in response to a signal sent out on the output Control Line. Thus, the absence of an input control cable forces the CONTROL OUT signal to remain at the Preset level. In addition, whenever the flip-flop is toggled (when SEND H makes a positive-going transition) a negative pulse is generated on the STATUS line. A positive-going transition of STATUS is used to negate the SSYN signal, so a short negative pulse on the STATUS line causes the interface to immediately allow the processor to resume operation.

#### 7.2.4 Data Ports

The two types of Data Ports used for data transfers between the PDP-11 and the modules are referred to as Input Data Ports and Output Data Ports. Each Data Port has associated with it a Control Port.

##### Input Data Ports

The Input Data Ports allow the modules to transfer 12 or 16 bits of data to the PDP-11 (the number of bits depends on the type of modules being used). Each Input Data Port is assigned four bytes in the interface address space. Two of these are used for the input data being transferred. The third byte is used for the interrupt enable. The fourth byte is used to indicate the status of the Input Data Control Port.



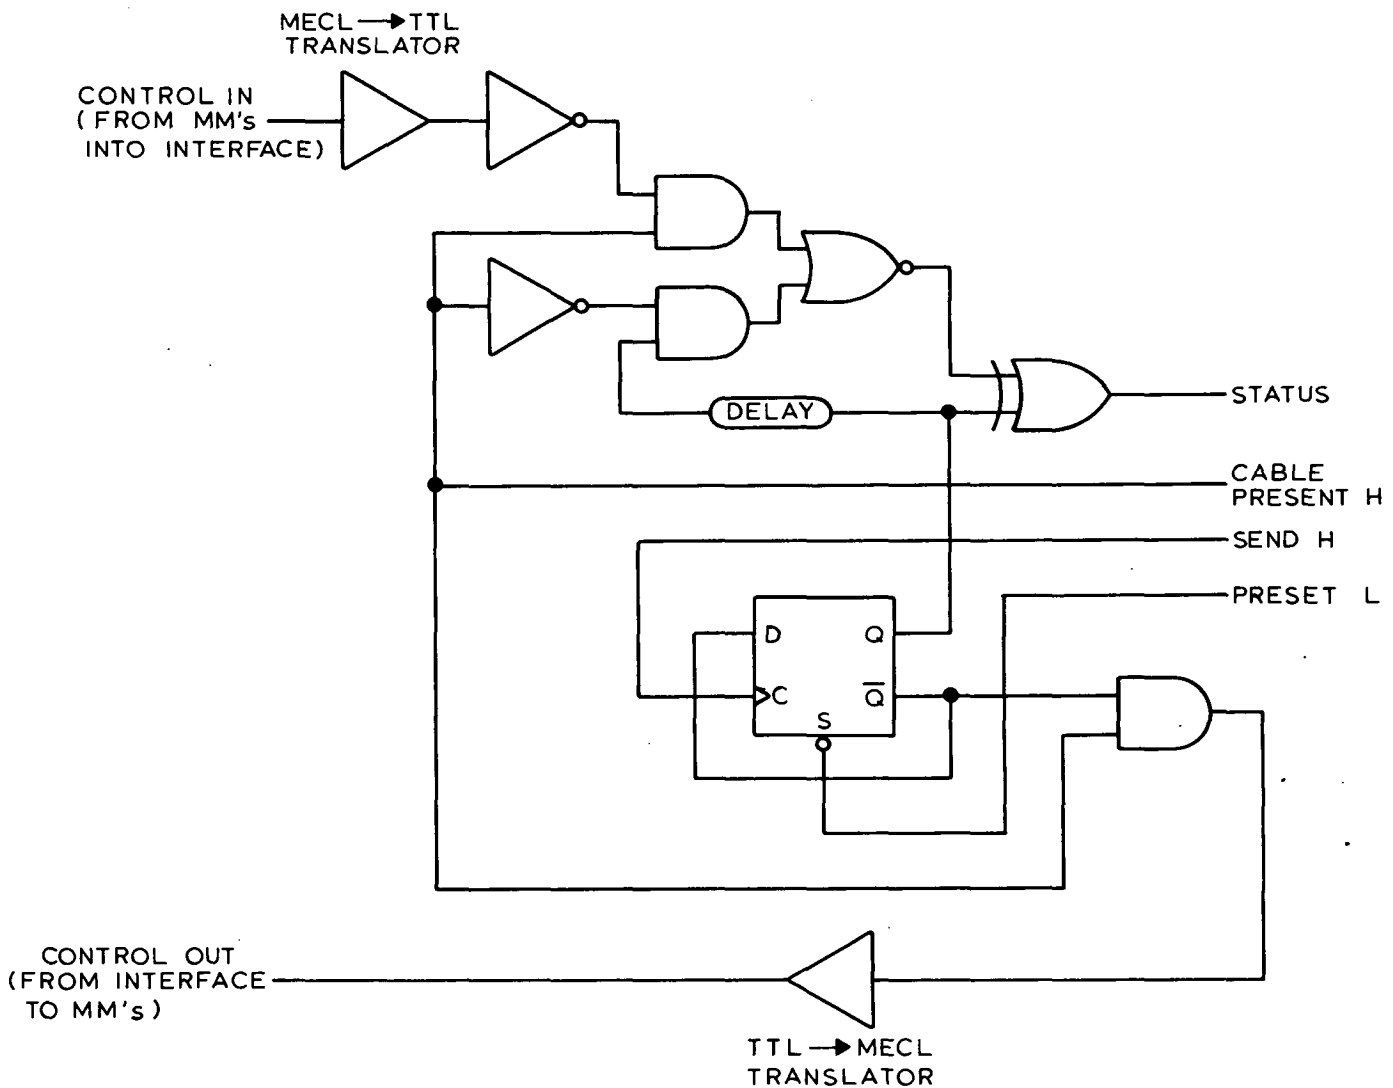


Figure 79. Control circuitry for Pausing Control Port.

The arrival of data sent by a module may be sensed by the PDP-11 in two ways. A module must follow the sending of a data word by a control signal sent on the input Control Line. This causes a mismatch of the Control Lines, which may be tested when the PDP-11 performs a DATI from the STATUS byte. If the ENABLE byte is set to all ones, then the arrival of the control signal will cause the interface to perform a vectored interrupt of the processor. Generally, the interrupt service routine in the PDP-11 would include an instruction which performs a DATI from the INPUT DATA bytes. When such an instruction is executed for the first time after the arrival of the data, the interface sends out a signal on the output Control Line, indicating that the data has been received by the PDP-11. The Control Lines are then matched so that testing the STATUS byte will yield all zeroes. Any later DATI performed from the INPUT DATA bytes will not affect the output control line, but the input data may not be assumed to still be valid.

The ENABLE byte may be set or cleared by performing a DATO or DATOB to the ENABLE byte. The contents of this byte may be examined at any time



by performing a DATI from the byte. The STATUS byte may also be examined at any time by performing a DATI from the byte. In addition, the contents of the STATUS and ENABLE bytes are indicated by lights on the front panel of the interface.

The circuit in Figure 80 performs the control interaction necessary for the Input Data Port. When the PDP-11 performs a DATI from the INPUT DATA

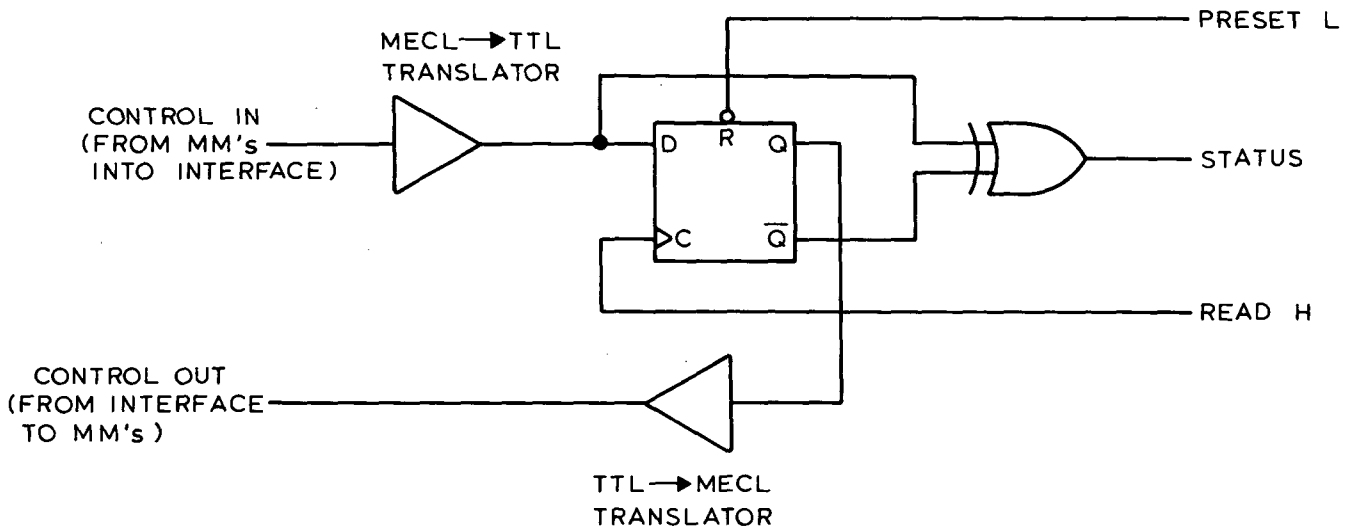


Figure 80. Control circuitry for Input Data Port.

bytes, a positive pulse appears on the READ H line. If the Control Lines are matched, the CONTROL OUT signal does not change. However, if an input control signal has arrived, a pulse on READ H will change CONTROL OUT, sending a transition signal on the output Control Line. If there is no input control cable present, then the MECL-TTL translator output will remain at the Preset level (low), and a pulse on READ H will have no effect on CONTROL OUT. Note that the conversion between level logic and transition logic is accomplished somewhat differently than in the Concurrent and Pausing Control Ports.

It is necessary to recognize that the handling of data and control are electrically independent. When data is sent to an Input Data Port on a data cable, it is accompanied by a Data Delivery signal. This signal is buffered and returned as the Data Delivery Return signal. The PDP-11 must determine that a control signal has arrived, indicating that data is available, before the data may be accessed. However, this control signal does not have to be sent to the Input Data Control Port associated with the data which has arrived. It may arrive at any Control Port in the interface (except the Output Data Control Ports). The only requirement is that the control signal arrive after the data has been data delivered. For example, it may be desirable to send several words of data to several Input Data Ports, and access all of them from the PDP-11 after a single final control signal has arrived.



## Output Data Ports

The Output Data Ports allow the PDP-11 to transfer 12 or 16 bits of data to the modules. Each Output Data Port is assigned two bytes in the interface address space. These are used for the output data being transferred. There is no status byte associated with an Output Data Port.

In normal operation, both the data cable and a pair of control cables are present in the connectors associated with the Output Data Port. These conditions will be assumed throughout the following discussion. When the PDP-11 performs a DATO or DATOB to the OUTPUT DATA bytes, the interface first sends the data out on a data cable, and accompanies it by a DATA DELIVERY signal. When the DATA DELIVERY RETURN signal arrives from the cable, the interface sends a transition signal on the output Control Line, to indicate that data is ready. The interface maintains the assertion of SSYN, preventing the resumption of processor operation, until a signal arrives on the input Control Line, causing the Control Lines to again be matched. Since the processor is never operating when the Control Lines are mismatched, there is no need to provide for the PDP-11 to test the status of an Output Data Control Port. However, the status is indicated on the front panel of the interface as an aid in the tracing of system errors.

There are three other modes of operation corresponding to the absence of a data cable and/or an input control cable. These modes, along with the mode described above, are summarized in the table below.

Table 18. Interface Modes of Operation

DATA CABLE	INPUT CONTROL CABLE	ACTION
Present	Present	Entire sequence; data delivery; control initiated
Present	Absent	Data delivery; no control initiated
Absent	Present	No data delivery; data stored in buffer register; control initiated
Absent	Absent	Data stored in buffer register; processor continues operation immediately

In the second case, with no input control cable, the interface data-delivers the data, but does not have to wait for a control sequence to be completed. An example of the use of this mode is when several words of data are sent to the modules, followed by a single control signal. All data words except the last sent would be delivered under the above mode of operation. The last data word would have a control signal associated with it, requiring delivery under the first mode of operation.



The third mode of operation, with only the control cable present, is very similar to the operation of the Pausing Control Ports. The only difference is that the Output Data Control Port may only operate in a single direction. The interface must send a control signal out to the modules before a corresponding control signal may arrive on the input Control Line. This is a result of the somewhat different method used to convert between level logic and transition logic.

The fourth mode of operation requires no explanation, since no interactions take place between the interface and the modules when there are no cables present. However, there is an additional feature of the third and fourth modes of operation which requires some clarification. In order to maintain the data on the data cable in the first two modes of operation, it is necessary to store the data in a buffer register within the interface. This action takes place whenever the PDP-11 performs a DATO or DATOB to the OUTPUT DATA bytes. It is entirely independent of whether or not a data cable is connected. The PDP-11 may perform a DATI from the OUTPUT DATA bytes at any time, causing the contents of the buffer register to be sent to the processor. Thus, the Output Data Ports may be used as general registers by the PDP-11, in addition to their function as a means for transmitting data to the modules.

When the restructured modules are being used, the data cable may be used in the Function Caller mode. In this situation, the Data Delivery signal is used to initiate an operation which yields a YES/NO decision. The outcome of the decision is returned on one of two DATA DELIVERY RETURN lines. A YES decision is returned on the DDR YES line while a NO decision is returned on the DDR NO line. The receipt of either of these signals indicates that the data has been received. However, it is desirable to allow the PDP-11 to determine which decision was reached. Consequently, one byte of the interface address space is assigned to the status of the DATA DELIVERY RETURN. After each DATO or DATOB to an Output Data Port, this status byte is set (NO decision) or cleared (YES decision). The PDP-11 may access the status of the last output sequence by performing a DATI from the DDR status byte. Note that the status is automatically cleared at the beginning of each output sequence and is set only at the end of an output sequence which results in a signal arriving on the DDR NO line.

The circuit used to accomplish the Output Data Port interactions is shown in Figure 81. Several preliminary observations will ease the analysis of this circuit. When the PDP-11 performs a DATO or DATOB to the OUTPUT DATA bytes, WRITE H makes a positive-going transition. The completion of the output data transfer sequence is indicated by a positive-going transition of the COMPLETION signal. Whenever a transition signal arrives on the DDR NO or DDR YES line, DDR makes a transition. When DDR NO makes a transition, a pulse is generated on SET DDR STATUS, setting the DDR status byte.

Now suppose that the data cable and the control cable are present, and that all macromodular signal lines are in the Preset state (the MECL-TTL translator outputs are low). The presence of the cables opens two signal paths and closes two signal paths through DATA CABLE PRESENT H and CONTROL



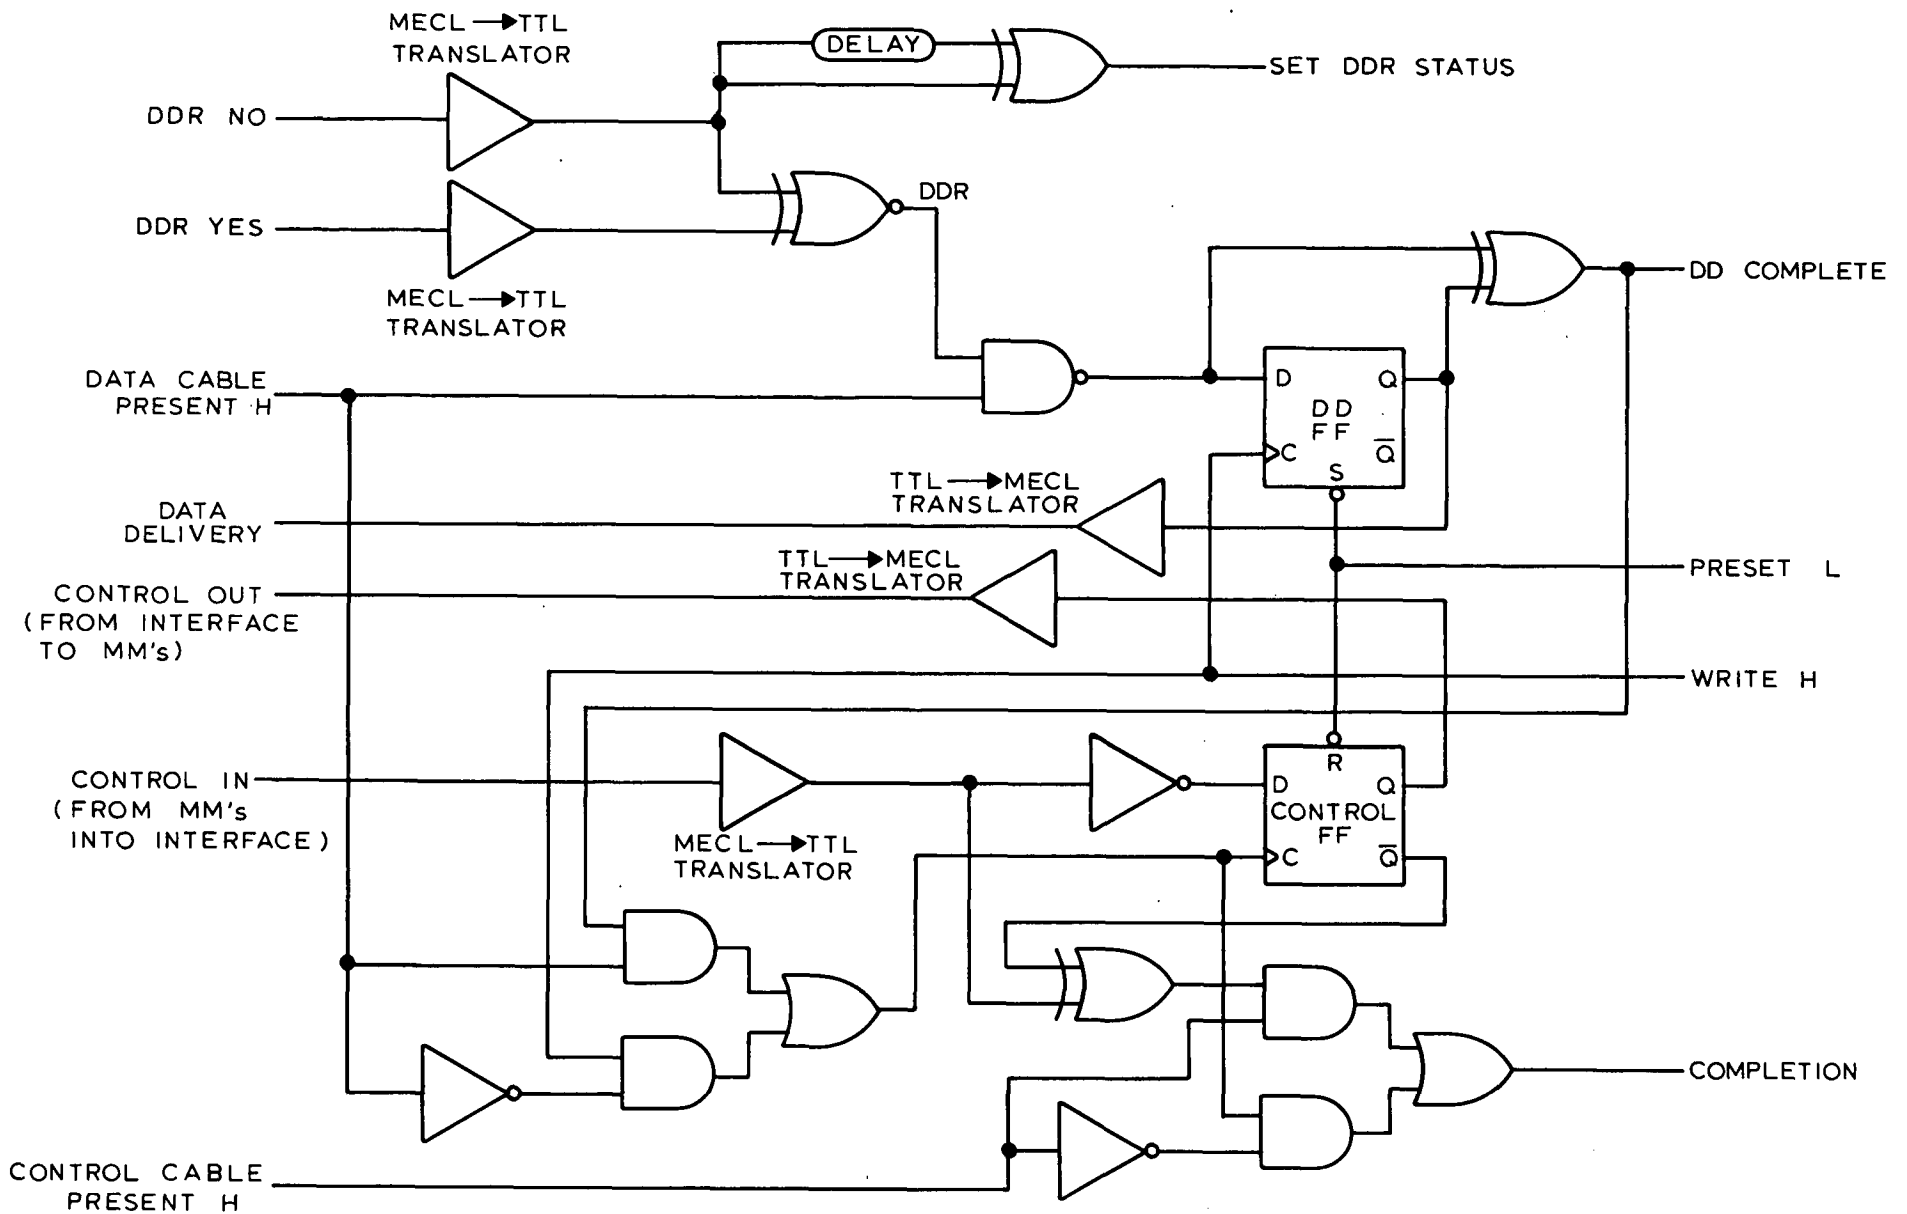


Figure 81. Control circuitry for Output Data Port.



CABLE PRESENT H. The D input to the DD flip-flop is low and the Q output is high. This makes DD COMPLETE high. The output of the CONTROL IN translator is low and the Q output of the CONTROL flip-flop is high. This makes COMPLETION high. When WRITE H makes a positive-going transition, the DD flip-flop is set low, sending out DATA DELIVERY and making DD COMPLETE go low. The low transition of DD COMPLETE is passed through to the C input of the CONTROL flip-flop. When a DATA DELIVERY RETURN signal arrives, DDR makes a transition and DDR goes high. This makes DD COMPLETE go high, and the positive-going transition is passed to the C input of the CONTROL flip-flop. Since the D input of the CONTROL flip-flop is high, the flip-flop is set and a transition signal is sent out on the CONTROL OUT line. In addition, the negative-going transition of the Q output of the CONTROL flip-flop causes COMPLETION to go low. When a transition signal arrives on the CONTROL IN line, COMPLETION goes high again, causing the interface to release SSYN and allow the processor to continue operation.

If the control cable is not present, then any transition of DD COMPLETE is passed directly to COMPLETION, bypassing the CONTROL sequence entirely. When DD COMPLETE goes high, indicating the receipt of the DATA DELIVERY RETURN, COMPLETION also goes high, and the interface allows the processor to continue operation.

If the data cable is not present, then a positive transition of WRITE H does not change the DD flip-flop, and the transition is passed to the input of the CONTROL flip-flop. If a control cable is present, then the CONTROL sequence is executed as described above. If the control cable is not present, the transition of DD COMPLETE is passed on to generate the COMPLETION signal. Thus, if both the data cable and the control cable are absent, the COMPLETION signal is a delayed replica of the WRITE H signal.

#### 7.2.5 Synchronization

A major difficulty in the design of any interface is the proper handling of synchronization problems. A lack of synchronization arises whenever two communicating devices do not share a common time reference, such as a periodic clock. In the PDP-11 to module interface, synchronization problems arise whenever the PDP-11 performs a DATI to obtain the contents of a status byte. Note that in all other data transactions, the PDP-11 is either controlling the transaction or is required to test a status byte before proceeding with the transaction.

Since the problem is restricted to a group of very similar operations, it is relatively easy to provide a simple and reliable solution. There are four types of status which may change independently of the PDP-11 processor operation. These are the status of the Concurrent Control Ports, the status of the Pausing Control Ports, the status of the Input Data Control Ports, and the status of a Preset sequence. All of these are handled identically, so that it is only necessary to describe a single circuit which is applicable to all of the cases.

Figure 82 shows a circuit for transferring the level of a status signal (STATUS) to the PDP-11 UNIBUS. In this simplified implementation,



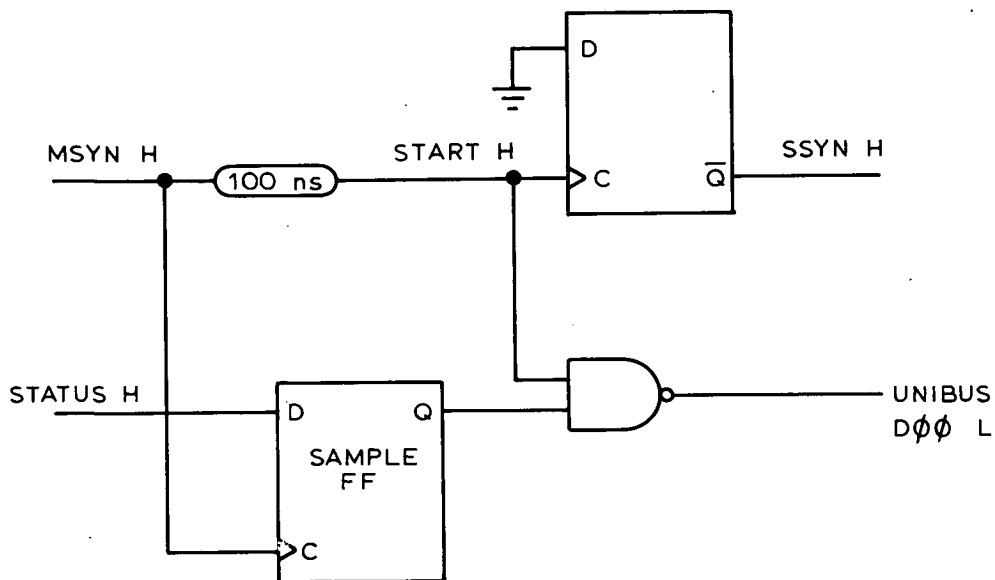


Figure 82. Circuit for transferring status level to PDP-11 UNIBUS.

the arrival of a positive-going transition on MSYN H signals the interface to respond to a DATI operation. The SAMPLE flip-flop is immediately clocked, and the level of STATUS H is presumably stored by the flip-flop. However, if STATUS H were to make a transition at approximately the same time as the positive transition of MSYN H, then it is possible for the output of the SAMPLE flip-flop to reach and remain at a metastable level somewhere between the well-defined high and low values. After a delay of 100 ns, the probability that the output is still at a metastable level is approximately  $10^{-8}$ , assuming that the flip-flop originally entered a metastable state. (See Part 1 Vol. IV of this report.)

On the circuit of Figure 82, START H makes a positive transition at least 100 ns after the positive transition of MSYN H. Since START H is used to pass the output of the SAMPLE flip-flop to the UNIBUS, there is a very small probability that a metastable level will be gated onto the UNIBUS. Since the processor uses the assertion of SSYN as a signal to accept the UNIBUS data, the assertion of SSYN must occur after START H gates the data onto the UNIBUS. Thus, the positive transition of START H is used to assert SSYN. Under worst-case conditions (MSYN H and STATUS H varying at 100 KHz clock rates), and assuming that UNIBUS L is used by the processor immediately after SSYN is asserted, the mean time between the transmission of metastable levels is approximately  $10^9$  seconds. However, the processor waits an appreciable



time after SSYN is asserted before using the UNIBUS data. This greatly increases the mean time between the receipt of metastable levels by the processor, since the probability of observing a metastable level decreases exponentially as the time delay before observing the output increases linearly.

It is important to note that the final level to which the metastable state settles is somewhat arbitrary, but this presents no serious difficulties. An example will illustrate this point. Suppose that STATUS H was originally low and makes a positive transition as SAMPLE H goes high. Further, suppose that the SAMPLE flip-flop enters a metastable state. If the output settles to a high value, then the PDP-11 receives the correct value of STATUS H. If the output settles to a low value, then the PDP-11 will determine that STATUS H has not yet changed. This is partially correct, since the PDP-11 examined STATUS H just as it was changing. If the attempt had occurred 20 ns earlier, an entirely correct result would have been obtained. In a normal situation, the PDP-11 would be repeatedly testing the value of STATUS H, so that the next attempt would surely yield a correct value of STATUS H. The only unpleasant effect of this situation is a small delay in determining the correct value of STATUS H.

#### 7.2.6 System Preset Via the Interface

The interface has the ability to start a Preset sequence in a macro-modular system. This may be started in several ways. Whenever the PDP-11 performs a BUS INIT, the entire interface is initialized, and a Preset sequence is begun. As long as the Preset sequence is being carried out, the modules may not be accessed from the PDP-11. Consequently, it is desirable for the PDP-11 to be able to determine when the Preset sequence has been completed. The PRESET status provides this information. One byte of the interface address space is assigned to the PRESET status. The PDP-11 may determine whether a Preset sequence has been completed by performing a DATI from the PRESET status byte. The PDP-11 may also initialize the interface and start a Preset sequence without executing a BUS INIT by performing a DATO or DATOB to the PRESET status byte. The third method of starting a Preset sequence is by pressing a switch on the front panel of the interface. This is useful when manual control by the user is desired.

The circuit of Figure 83 accomplishes the generation of a Preset sequence through the use of a controller. The PRESET RETURN H and SYSTEM READY H signals come from the controller, while the PRESET H signal goes to the controller. The INITIALIZE H signal is used to initialize the interface, including all macromodular output control cables connected to the interface. The level of the PRESET STATUS H line is sent to the PDP-11 when a DATI is performed from the PRESET status byte. When PRESET STATUS H is high, the modules are in the process of being preset. The START PRESET H signal goes high when a BUS INIT occurs, or when the PDP-11 performs a DATO or DATOB to the PRESET status byte.

The signals between the interface and the controller are related in the following manner. When START PRESET H goes high, the PRESET H signal is sent to the controller, and PRESET STATUS H is set to high to indicate that the Preset sequence has begun. Some time later, SYSTEM READY H goes low. When the controller sets PRESET RETURN H to high, the interface



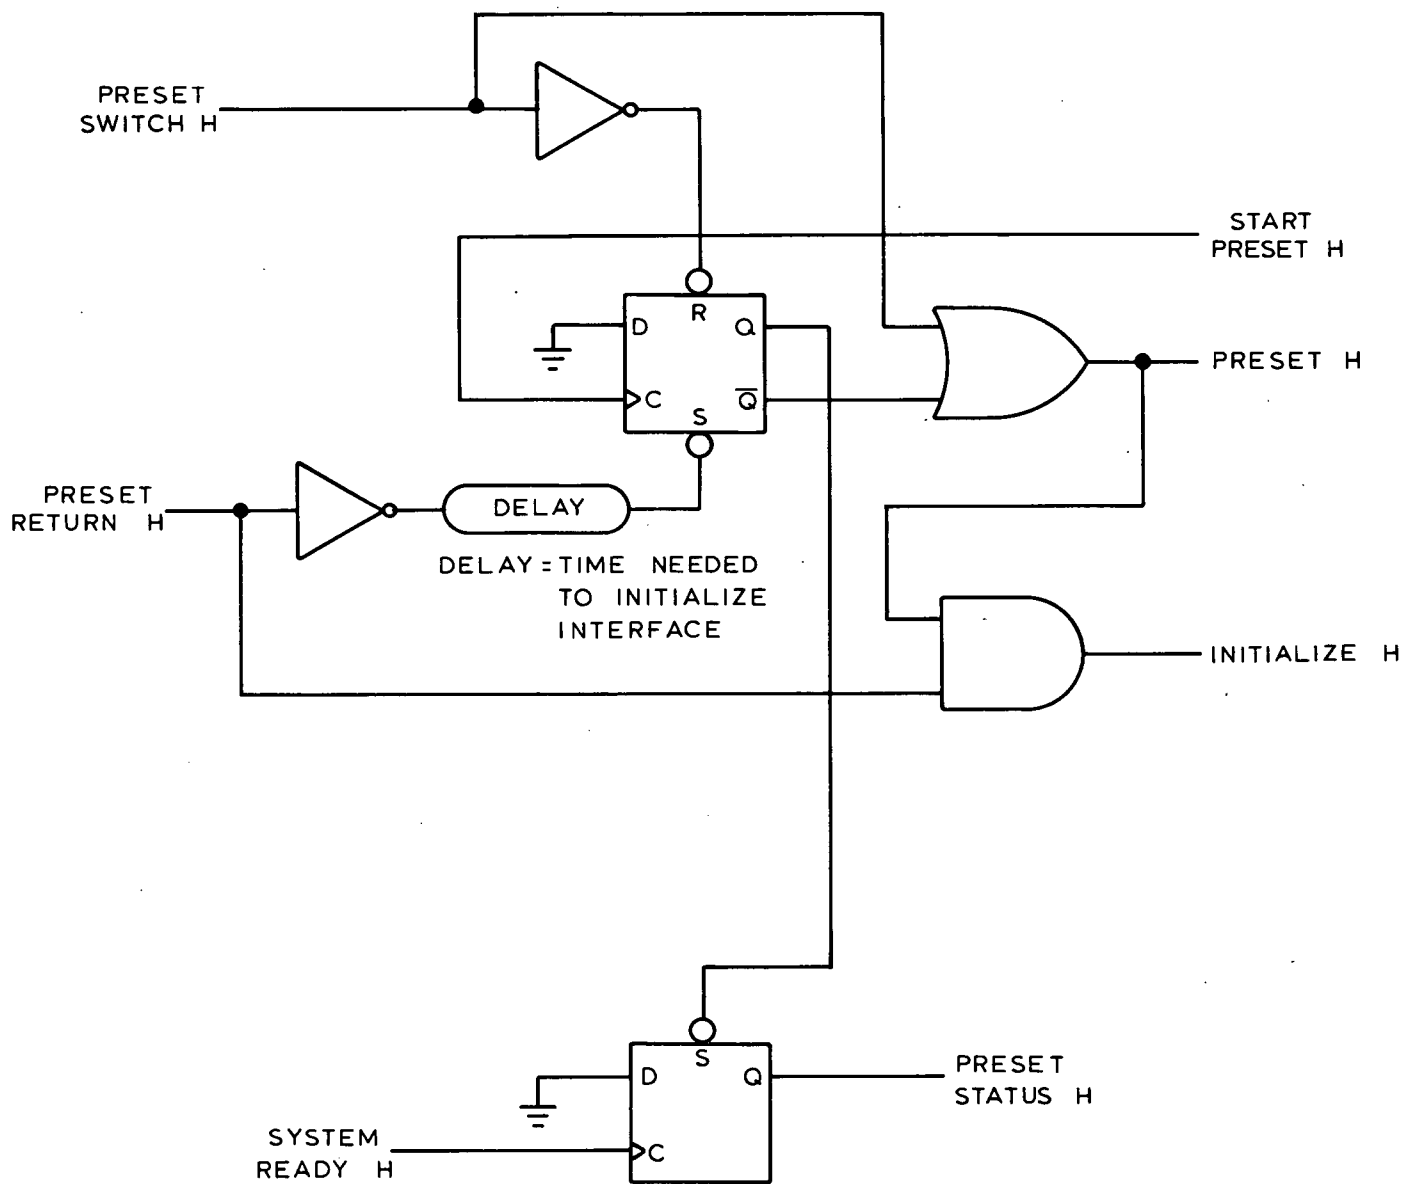


Figure 83. Circuit to generate a preset sequence via controller.



begins its internal initialization. After this is complete, the interface may let the PRESET H line go low. At this time INITIALIZE H goes low. If the Preset sequence was started by pressing the front panel switch, PRESET H will probably remain high long after the PRESET RETURN H goes high. Some time after PRESET H goes low, the controller responds by making PRESET RETURN H go low. After an indeterminate delay, SYSTEM READY H goes high. This transition sets PRESET STATUS H to low, indicating that the Preset sequence has been completed.



## 8. CONCLUSION

We believe that the macromodule design that has just been described meets successfully the technical objectives that were established at the beginning of the program. Our experience with the working inventory has demonstrated that the power of macromodules makes them a unique tool, and that their ease of use and reliability are adequate and do not seriously limit their applicability.

The primary limitations to the full use of macromodules are in their high cost and limited availability, and in the need to develop stored program macromodular components that can handle the non-critical parts of a system task in a standardized and efficient manner. Thus, on the one hand a prospective user is faced with the fact that his access to macromodules is limited; on the other hand, the investment in learning to use them in a serious way that is integrated with the use of other kinds of computer tools is substantial. As a result, practically all of the serious users of macromodules to date have been members of our staff or colleagues at Washington University.

Both of these limitations are presently being attacked under the continuing support of the National Institutes of Health. Part 3 of this Final Report gives a status report on the effort to develop a restructured form of macromodules that is less expensive and manufacturable by more conventional processes. Other work, not reported on here, is now under way to develop means for providing programming and debugging support for individual small stored-program computers that can be embedded in large numbers in systems designed in the macromodular style.



9. REFERENCES

1. H. Lewis and D. Richardson, "Macromodule Core Memory Timing Considerations," Computer Systems Laboratory Technical Memorandum No. 152, Washington University, St. Louis, Mo., April 1972.
2. W. A. Clark and C. E. Molnar, "A Description of the LINC," in Computers in Biomedical Research, Vol. II, R. W. Stacy and B. D. Waxman, eds., Academic Press, New York, pp. 35-66, 1965.
3. Ragnar Holm, Electric Contacts, Springer-Verlag, Berlin/New York, 1967.
4. Reference Data for Radio Engineers, Fifth Edition, International Telephone and Telegraph Co., Indianapolis, Ind., 1968.
5. T. J. Chaney, A DC to DC Converter for Macromodules, Computer Systems Laboratory Technical Report No. 13, Washington University, St. Louis, Mo., Aug. 1969.
6. "Applications of the  $\mu$ A742 Trigac, a Zero Crossing AC Trigger," Fairchild Semiconductor, Mountain View, Calif., pp. 15-16, 1969.
7. Gerald C. Johns, "Interfacing the PDP-11/40 with Restructured Macromodules," Computer Systems Laboratory Technical Memorandum No. 191, Washington University, St. Louis, Mo., June 1973.
8. Peripherals Handbook for the PDP-11, Publication 112-01071-1854, Digital Equipment Corp., Maynard, Mass., 1973.



Unclassified

Security Classification

## DOCUMENT CONTROL DATA - R &amp; D

(Security classification of title, body of abstract and indexing annotation must be entered when the overall report is classified)

1. ORIGINATING ACTIVITY (Corporate author) Computer Systems Laboratory Washington University St. Louis, Missouri		2a. REPORT SECURITY CLASSIFICATION Unclassified	
3. REPORT TITLE  Design of Phase I Macromodules		2b. GROUP	
4. DESCRIPTIVE NOTES (Type of report and inclusive dates) Final Report 4/1/65 through 12/1/73			
5. AUTHOR(S) (First name, middle initial, last name)  Christine D. Coaker, Editor			
6. REPORT DATE February, 1974		7a. TOTAL NO. OF PAGES 187	7b. NO. OF REFS 8
8a. CONTRACT OR GRANT NO. DOD (ARPA) Contract SD-302		9a. ORIGINATOR'S REPORT NUMBER(S)  Volume III of Part 1	
b. PROJECT NO. ARPA Project Code No. 655		9b. OTHER REPORT NO(S) (Any other numbers that may be assigned this report)  Technical Report No. 46	
10. DISTRIBUTION STATEMENT  Distribution of this document is unlimited.			
11. SUPPLEMENTARY NOTES		12. SPONSORING MILITARY ACTIVITY ARPA - Information Processing Techniques, Washington, D.C.	
13. ABSTRACT <p>This volume attempts to reconstruct the rationale for the major design decisions that were made in the development of Phase One macromodules, and presents a review of the internal architecture, intercommunication, electrical and mechanical design. Material describing the interfacing of macromodular systems to other computers is also included. The level of detail presented is intermediate, and assumes a general familiarity with the concept of macromodules and their usage.</p> <p>Two of the design goals proved particularly difficult to meet. The objective that the system designer should not have to be concerned with <u>any</u> details not directly concerned with the functional definition of his system, and the objective of a system discipline that would place <u>no</u> limits on the maximum system size both required novel approaches to system control and architecture and required an engineering design that was well outside of the prevailing state of the art.</p> <p>The unusual or non-routine aspects of macromodule design are emphasized, particularly those points of the design that proved most challenging or intellectually rewarding. A preliminary evaluation of our experience to date in the area of reliability is given.</p>			

DD FORM 1 NOV 65 1473

(PAGE 1)

S/N 0101-807-6801

Unclassified

Security Classification



Unclassified

Security Classification

14. KEY WORDS	LINK A		LINK B		LINK C	
	ROLE	WT	ROLE	WT	ROLE	WT
Macromodular computer systems Modular computer systems Engineering design of macromodules Mechanical design of macromodules Logical design of macromodules Interfacing to macromodules Interconnection design of macromodules Power supply for macromodules						

Unclassified

Security Classification



