Aerospace Division

ON BOARD PROCESSOR SYSTEM STUDY

FINAL REPORT

for

NAS 5-20134

DECEMBER 1970

For

NATIONAL AERONAUTICS AND SPACE ADMINISTRATION
Goddard Space Flight Center
Greenbelt, Maryland 20770

Prepared By

WESTINGHOUSE ELECTRIC CORPORATION
Defense and Space Center
Systems Development Division
Baltimore, Maryland 21203

Aerospace Division

ON BOARD PROCESSOR SYSTEM STUDY

FINAL REPORT

for

NAS 5-20134

DECEMBER 1970

For

NATIONAL AERONAUTICS AND SPACE ADMINISTRATION
Goddard Space Flight Center
Greenbelt, Maryland 20770

Prepared By

WESTINGHOUSE ELECTRIC CORPORATION
Defense and Space Center
Systems Development Division
Baltimore, Maryland 21203

TABLE OF CONTENTS

TABLE OF CONTENTS (Continued)

LIST OF TABLES

LIST OF FIGURES

1. INTRODUCTION

This report summarizes the work performed on NASA Contract NAS-5-20134. The objective of this contract was to perform a study of the present On Board Processor (OBP) System in order to improve overall system performance. The study consists of an evaluation of three areas of endeavor: the selection of a new circuit family, an evaluation of micro-programming, and an evaluation of combining the fixed I/O with the CPU.

The circuit investigation evaluated off-the-shelf circuits, Large Scale Integration (LSI) circuits previously developed by other companies, and special purpose LSI structures for the Advanced On Board Processor (AOP) system. These various approaches were compared for the circuit technologies which initially appeared applicable to the Advanced On Board Processor (AOP) requirements.

Microprogramming was investigated as a means of reducing and systematically organizing the control logic in the CPU. Microprogrammed control was attractive from the standpoint of circuit reduction, improved instruction set flexibility, and improved partitioning feasibility. The application of existing LSI memories as well as memories being developed for the near future was investigated for this approach.

The evaluation of combining the fixed I/O with the CPU was undertaken to improve total system performance. Since many functions performed in the original I/O are fixed and do not change from mission to mission, it appeared desirable to incorporate these functions into the CPU. As a result of combining these units, the I/O unit is greatly simplified and can

be easily changed to meet various mission requirements. In addition to the simplification of the special I/O, the feasibility of combining the CPU and fixed I/O is illustrated by the improved system performance and logic requirements associated with the CPU.

2. CIRCUIT INVESTIGATION

The objective of the circuit investigation is to choose a circuit technology which will improve the packaging, speed, and power characteristics of the OBP without causing an exhorbitant increase in cost. The criteria used to evaluate the various circuit types are cost, size, speed, and power in order of priority. It is assumed that the speed and power characteristics of the Advanced On Board Processor (AOP) will be at least equivalent to those of the present flight Model MOD-I OBP.

The circuits investigated include several types of Metal Oxide Semiconductor (MOS) technologies as well as low power bipolar devices. Because of cost considerations, the availability of off-the-shelf devices is of major importance. The reduction of size on the other hand dictates that LSI devices be used wherever possible. Since LSI circuits are not commonly available on an off-the-shelf basis, an effort was made to discover other ways by which they could be obtained without incurring exhorbitant costs. Table 2-1 lists available circuit configurations.

One hope was that special LSI chips previously developed for other applications would be applicable to the needs of the OBP. Several manufacturers were questioned to determine the availability of such chips.

Also, estimates of the costs associated with developing LSI chips especially for the OBP were solicited from several manufacturers. Even though these costs were expected to be high, it was thought that the packaging advantages of such circuits might justify a limited increase in expense.

The following sections describe the technical applicability and the availability of each circuit type considered. Then, conclusions are drawn as to the advantages and disadvantages of each technology. Finally all the circuits are compared and conclusions are presented.

2.1 P-CHANNEL MOS

The following section describes the characteristics and availability of P-channel MOS (PMOS) circuits.

2.1.1 Applicability

PMOS logic circuits are designed to operate in both a static and dynamic manner. Static logic uses a high impedance PMOS transistor as a load resistor. This is done because a transistor requires much less chip area than a high impedance diffused resistor. The gate of the load transistor is normally tied either to the drain supply voltage or to some other supply at least one threshold voltage more negative than the drain supply. Because of this constantly turned on load device, static logic generally dissipates a large amount of DC power; e.g., as high as 100 mw per gate.

Dynamic logic reduces (depending on the clock scheme used) DC power dissipation by turning on the load device only during clock pulses. Data is stored temporarily with small capacitors on the chip.

It is possible to trade speed of operation for power dissipation by varying the frequency and duty cycle of the clock (or clocks) employed. Since the charge stored on the capacitors will eventually leak off, it is necessary that either some minimum frequency of use be guaranteed for each gate or else provision be made to periodically refresh all outputs through the use of feedback circuits. Use of such a scheme would require a redesign of the AOP logic.

The speed of PMOS circuits is seriously degraded by the large capacitive loading encountered by all signals which must leave an integrated circuit chip. The capacitances which must be charged by interchip signals typically exceed intrachip capacitances by more than an order of magnitude. Small scale integrated (SSI) gates operating at reasonabled DC power loads (less than 10 mw per gate) generally have propagation delays of a few microseconds. However, on LSI chips, gate delays can be reduced to approximately 50 ns per gate.

Because the nature of PMOS integrated circuits obviates the use of area consuming isolation diffusions between transitors, PMOS offers the highest potential chip complexity of any circuit technology in current use. It is at its best where reduction of large complex logic structures to LSI chips is necessary. This technology comes closest to making the "computer on a chip" concept a practical possibility.

Because such great chip complexity can be achieved and because the speed of PMOS devices only becomes acceptable with LSI, most large PMOS systems have been designed using large complex special purpose chips. To save power, dynamic logic has been the most popular mode of operation. It has been found that special purpose chip designs for dynamic circuits rarely function properly initially because of unforeseen parasitic capacitances. It is therefore common to require several costly reworkings of the circuit design and chip layout to obtain a working device.

2.1.2 Availability

A wide selection of small and medium scale integrated (MSI) devices are available from several manufacturers on an off-the-shelf basis. Prices are generally competitive with other circuit types for similar logic func-

tions. Most of the MSI/LSI devices available are large circuits exhibiting a high degree of symmetry, such as memories and shift registers. These are for the most part not applicable to the needs of the AOP system.

2.1.3 Conclusions

It appears that the best speeds achievable with SSI PMOS circuits, are not particularly fast, and can only be obtained at the cost of prohibitive power dissipation. The low power circuits are not fast enough to permit the AOP to operate at even its present speed. The only justification for considering PMOS would be the possibility of building a machine with very complex LSI chips.

Because of the great chip complexity possible with PMOS, it is also usually advertised as inexpensive on a cost per gate basis. However, the low cost per gate for LSI circuits is only realized with quantity production. In the case of the AOP the development and debugging of LSI PMOS design would be prohibitively expensive.

2.2 SILICON GATE MOS

An evaluation of Silicon gate P-channel MOS circuits is discussed in this section.

2.2.1 Applicability

Silicon gate MOS, a variation of the basic MOS transistor, is being investigated by several manufacturers of integrated circuits. The silicon gate technology differs from the conventional MOSFET in that a layer of P-type silicon deposited over the gate dielectric replaces the aluminum normally used as a gate electrode. In addition, the silicon dioxide layer between the gate and N type silicon substrate has in some cases been augmented

by a layer of silicon nitride which acts to increase the dielectric constant of the gate dielectric.

The advantages of the silicon gate technology are as follows. The use of silicon as the gate electrode and increasing the dielectric constant of the underlying oxide layer serve to decrease the threshold voltage of the device. The effect of this decrease is to permit these devices to interface directly with bipolar devices. This may allow circuits to be designed using the most useful combination of MOS and bipolar devices without the problems of voltage level incompatibility presently encountered.

The manufacturing process involves using the gate as the mask for diffusing the P-type source and drain regions. This technique greatly reduces the amount of overlap between the gate and the other two regions with a consequent reduction in input capacitance. It is also possible to reduce the size of the whole device, thus reducing junction capacitances. As a result, speeds are reported to be increased over comparable conventional devices by a factor of approximately three.

2.2.2 Availability

At the present time the only off-the-shelf silicon gate devices are memories. If it were decided to adopt these circuits for use in the Advanced On Board Processor, the only possibility would be to support the customized design of special LSI chips with the attendant high costs.

2.2.3 Conclusions

As with ordinary PMOS silicon gate MOS only shows promise for a completely LSI computer. When compared with other technologies on a small scale integration basis, it is inferior from the standpoints of both power and speed.

2.3 IMOS TECHNOLOGY

The use of ion implantation as a doping technique promises to provide a family of extremely high speed MOS integrated circuits. Ion implantation research has been conducted for several years but has only recently been applied to the manufacture of MOS circuits. Consequently, the devices are in the experimental stages of development.

Ion implantation MOS (IMOS) circuits operate in the same manner as ordinary P-channel MOS circuits. The high speeds predicted for the IMOS circuits are the result of reducing parasitic gate to source and gate to drain capacitances. In normal P-channel devices in which only diffusion is used to establish the source and drain regions, four masking operations are required. The inaccuracies of the mask are such that the gate metal must overlap the source and drain regions to assure that the entire channel region is covered. As a result, parasitic gate to source and gate to drain capacitances, as well as a Miller capacitance proportional to gain, act to degrade the speed of the device.

In the IMOS devices, the source and drain regions are diffused in the normal way, only slightly farther apart than usual. Then the layer of oxide is deposited. The oxide is removed in the places necessary for the source and drain contacts but left in the region where the gate is to be located. The gate is not as wide as the channel between the two diffused P-type regions. The metal is then applied to gate, drain, and source. Finally, the whole structure is bombarded with boron ions which penetrate through to the substrate only in the regions where the oxide layer is not protected by a layer of metal; i.e., between the gate and the other two contacts. The effect of this is to extend both source and drain up to, but

not under, the gate metal. The boron ions do not diffuse laterally, and there is therefore no overlap of the three regions.

Circuit input capacitance is typically reduced by a factor of five through the use of ion implantation techniques. Researchers predict that IMOS MSI circuits will operate three to five times faster than equivalent diffused circuits. It is presumed that the power requirements of these circuits will be of the same order of magnitude as those for ordinary P-channel MOS devices, although no data on this is currently available.

2.3.1 Availability

At present, the only IMOS circuit which is commercially available is a 64-bit shift register which operates at 20 MHz. Hughes Aircraft Company is the only manufacturer known to be actively working on an IMOS development program. The expansion of the types of circuits which become available will depend on the expressed needs of potential customers. Because this technology is presently in the research stage, it will most likely be several years before a variety of circuit types is available on other than a custom built basis. As a result, if the appropriate circuit types are to be fabricated in the near future they will be quite expensive.

2.4 COMPLEMENTARY MOS

The advantages and disadvantages of CMOS circuits are described in this section.

2.4.1 Applicability

Complementary MOS (CMOS) circuits differ from PMOS devices in that a complementary MOSFET transistor pair is used in place of a P-channel MOSFET switch and its permanently turned on P-channel load transistor. In either the "ONE" or "ZERO" state there is always one transistor turned on and one

turned off. There is therefore no quiescent current flow other than leakage current through the "off" transistor. Since this current is extremely small, quiescent power dissipation as low as a few nanowatts per gate is not unusual. The only time that significant power is dissipated is during switching. The variable component of the power required by CMOS circuits is used to charge the output capacitance of MOSFET transistors. Therefore, it is directly proportional to the frequency of switching and to the square of the supply voltage. This power typically varies from the quiescent value of a few nanowatts per gate at DC to a few milliwatts per gate at 1 MHz. In the OBP the average switching rate for gates is less than 500 KHz implying average power consumption of less than 0.5 mw per gate.

Because there is no high impedance load device required as in PMOS circuits, the output capacitance can be charged and discharged quickly, resulting in faster switching times. Propagation delays vary with both supply voltage and load capacitance. The maximum propagation delay for typical gates with a 14 volt supply and 10 pf load (2 loads) is approximately 50 ns. The maximum propagation delay with a six volt supply and 50 pf load (10 loads) is approximately 400 ns. If a 10 volt supply and a fanout of 8 are assumed, the range of propagation delays (minimum delay to maximum delay) among identical gates (RCA 4000) is from 50 ns to 180 ns. Presumably for a 10 volt supply and a 40 pf load, the average propagation delay of the device lies near the middle of this range, i.e. at about 115 ns. All the above numbers apply to small scale integrated devices. It is understood from discussions with manufacturers that gate propagation times for LSI circuits would be considerably faster.

CMOS circuits require only one voltage supply. The minimum voltage at which the devices will operate satisfactorily is determined by the threshold voltages of the transistors and is approximately six volts. The maximum voltage recommended by the two major manufacturers is 20 volts. The optimum power-speed tradeoff appears to occur in the neighborhood of 10-12 volts.

CMOS devices offer an extremely high noise margin - as much as 45 percent of the supply voltage. Large supply voltage variations can also be tolerated without adversely affecting circuit operation (except to the extent that switching speed is changed).

The chip area required by CMOS gates is larger than that needed for PMOS since there are more transistors, but it still represents considerable reduction over bipolar integrated circuits. Large complex LSI circuits have been built and presumably would perform well for the AOP.

One area of concern for MOS circuits if their radiation susceptibility. This has been a factor in evaluating these circuits. While improvements are being made, there is not enough evidence to indicate that radiation hardened MOS devices are readily available.

Outputs of CMOS gates cannot be tied together to produce the so called "wired or" function since a conflict would cause the supply voltage to be shorted to ground. The loss of this function will cost at least two gate delays everywhere it is presently used if an SSI CMOS approach is adopted. If special purpose LSI chips are designed, gates having any desired number of inputs may be used thereby removing the problem.

2.4.2 Availability

Complete families of small scale devices are manufactured by RCA and Solid State Scientific on an off-the-shelf basis. These include not only discrete gates and flip flops but also many MSI functions such as shift registers, adders, counters, etc. See Table 2-1.

Several manufacturers were contacted to learn if any applicable LSI chips had previously been developed for other applications, but which could be used in the AOP. It was found that RCA and Solid State Scientific had developed four bit arithmetic units capable of performing 16 bit parallel additions in six microseconds and one microsecond respectively. These were judged to be too slow to be satisfactory for the OBP.

Further inquiries were made to determine the expected performance and cost of specially designed LSI circuits. One manufacturer claimed that an 18 bit 100 to 200 nanosecond monolithic adder could be manufactured. The cost of developing a large complex chip (120 mils by 120 mils) was estimated between $25,000 and $50,000 depending on the complexity of the circuit.

2.4.3 Conclusions

CMOS technology has been in use long enough to appear suitable for use in the Advanced On Board Processor. It offers low operating power and good immunity to noise. The speed of CMOS circuits appears too slow in SSI form. While speed can be increased with LSI structures, the cost is very high. The so called "wire or" function cannot be used. However, if a customized LSI chip is designed, gates can be built with the proper number of inputs. Although the power requirements are low, they increase linearly with frequency. The superiority of the CMOS low power figure becomes

questionable as the switching frequency rises to near 1 MHz. Susceptibility to radiation appears to be a disadvantage for the CMOS technology.

2.5 LOWER POWER TTL/DTL

The characteristics of low power bipolar circuits are evaluated in this section.

2.5.1 Applicability

Low power bipolar circuits offer higher speeds than any of the MOS technologies with the possible exception of IMOS. Gate propagation delay times are of the order of 30-50 nanoseconds. Power requirements of approximately 1 mw per gate are typical of the circuits currently available. The speed power product of these devices are almost equivalent to that of CMOS at the switching frequencies encountered in the AOP.

In the event that a medium or high speed TTL adder should be desired for the CPU, there would be no interface problems if the rest of the system were built with low power TTL (LPTTL).

Bipolar circuits are inherently radiation hardened and appear more desirable for space application at this time.

Bipolar circuits require more chip area per logic function than any of the MOS technologies mentioned so far. This potentially limits the possible complexity of an LSI chip. However, all partitioning studies performed for the OBP so far have encountered pin limitations long before all the available gates were assigned.

2.5.2 Availability

The 54L/74L series of low power TTL (LPTTL) manufactured by Texas Instruments is a complete family of small scale integrated circuits. Typical gates are characterized by 30 ns propagation delays and 1 mw power

dissipation. A few MSI devices, including a four bit counter and a four bit register are also available. The only other lower power TTL logic known to be available on an off-the-shelf basis is the Fairchild 93L series. This is a family of MSI functions which is intended to eventually duplicate all the functions available in the higher speed 9300 series. The manufacturer claims that 20 ns, 2 mw gates are typical. The series presently includes 15 devices among which are a 1 of 16 decoder, an 8 input multiplexer, and a 4-bit arithmetic logic unit. These devices as well as others in the group appear to be applicable to the AOP in the event that a hybrid SSI-MSI mixture is used. Fairchild also offers a limited number of LPTTL SSI logic devices.

The possibility of obtaining LPTTL LSI chips which would enable the OBP to operate at high speeds with a considerable decrease in size was carefully investigated. Portions of both the I/O unit and the CPU control logic were partitioned and shown to several TTL manufacturers. It was desired to find out whether logic structures of this type could be built with LPTTL LSI. Cost estimates were also requested. It was hoped that some manufacturer might be interested in absorbing part of the cost of the development of these chips. The results of the investigation were rather discouraging. Although there was no apparent technical reasons why the logic partitions could not be placed on LSI chips, there was no interest on the part of the manufacturers in sharing the cost of chip development. Without such cost sharing, the costs were prohibitively high.

The other alternative which was investigated was the use of LSI gate arrays. These are LSI chips in which the first layer of metallization is used to connect the individual transitors to form gates and is the same

for all chips. The second layer of metal is used to distribute power and to connect the gates into functional logic structures. The second layer mask must be designed from logic drawings and is different for each chip type. Since the original chip development cost has already been paid, this method makes available the advantages of LSI without exhorbitant expense. The gates are claimed to have propagation delays of 60 ns and power dissipation of 0.75 mw.

A preliminary investigation of the feasibility of partitioning the CPU and I/O digit oriented logic was performed assuming a 40 pin package with 120 gates per chip. The results of the investigation indicate that the CPU could be efficiently partitioned forming satisfactory logic groupings.

2.5.3 Conclusions

The relatively high speed and low power of LPTTL devices makes them appear extremely attractive for the Advanced On Board Processor. Apparently, custom LSI chips cannot be obtained without exceeding the cost limitations of the program. However, LSI gate arrays offer the same advantages at a fraction of the price. Since their development costs have already been absorbed and they have a low speed-saver product, these circuits appear very attractive. In the event that the gate arrays prove to be unsuitable or unavailable, there exist adequate families of both SSI and MSI LPTTL devices which may be used to improve the performance of the system.

2.6 FINAL CONCLUSIONS AND RECOMMENDATIONS

To determine the optimum circuit technology for the Advanced On Board Processor, the technical characteristics of each candidate must first be evaluated in the light of the objectives of this study. Then, for the

circuit types judged to be suitable for the system, such factors as cost and availability must be considered.

A few general observations may be in order before listing the characteristics of the circuits which were studied. The most important technical goal of this redesign is to significantly reduce the physical size of the CBP. This is undoubtedly most easily achieved by the use of LSI chips especially designed and manufactured for this system. Unfortunately, the costs associated with the production of a very limited number of special purpose LSI chips is prohibitive. Therefore, the following guideline is imposed upon the circuit evaluation. Any circuit type which is available only in special purpose LSI or which only becomes technically acceptable in the form of LSI cannot be considered for use in the Advanced On Board Processor. Using this and the other criteria of lower power with maximum speed the following conclusions were reached.

PMOS SSI/MSI circuits are characterized by a speed power product inferior to that of the LPDTL circuits used presently. Since the primary advantage of PMOS is its small chip area per logic function, this advantage can only be realized with special purpose LSI. The prohibitive cost of these structures made them unsatisfactory for the AOP.

IMOS is simply not available on other than a custom chip basis. In addition, its power requirements should be similar to those of PMOS.

Silicon gate PMOS is also only available on a custom chip basis with the exception of a few off-the-shelf memory devices. Furthermore, it offers no speed-power improvement over the present system.

CMOS circuits would permit a reduction in the power requirements of the system with some loss in speed, assuming a 12 to 15 volt supply.

They are available off-the-shelf as SSI gates and flip flops and several useful MSI functions. The low speed makes these circuits unsatisfactory. An increase in speed can be achieved with LSI structures, but costs prohibit the development of these chips.

LPTTL will permit the system to operate faster than it presently does and for approximately the same amount of power. It is available in a wide selection of SSI and MSI functions. Since packaging size is a prime consideration, a new packaging arrangement will decrease the size of the OBP. The cost associated with this approach is low. An SSI/MSI hybrid mixture of LPTTL devices is suggested as a backup method for the AOP because of its low cost.

The LPTTL is the only technology which has been used to implement LSI gate arrays. The customized chips permit the advantages of LSI for a fraction of the cost. The development cost has been absorbed by other companies and these structures readily apply to the AOP design. Both reduced power and increased speed will be realized with LSI gate arrays. Thus the speed-power product should be considerably reduced. The inherent radiation hardened feature of these circuits is an added advantage. Since packaging size with minimum cost is the prime consideration, LPTTL LSI gate arrays appear to be the best approach to consider for the AOP.

Table 2-1. Characteristics of Available Circuits

RCA COS/MOS		tpd_0/tpd_1 (ns)			Load Cap. (pf)	Quiescent Power
Part No.	Description	V_{DD} = 6V	V_{DD} = 10	V_{DD} = 14V		
CD4000	Dual 3 Input NOR + Inverter		65/100		50	10 nw
CD4001	Quad 2 Input Nor		65/100		50	10 nw
CD4002	Dual 4 Input Nor		65/100		50	10 nw
CD4003	Dual D Flip-Flop	800/800	190/190	130/130	30	50 nw
CD4004	7-Stage Binary Counter	550/600	175/190	110/120	30	50 uw
CD4005	16 bit (16 x 1) R/W Memory		15 ns read time		10	100 nw
CD4006	18 bit static shift register	350/750	150/200	125/150	15	100 nw
CD4007	Dual Comp. Pair + Inverter		65/65		50	10 nw
CD4008	4 bit full adder	(Sum-in to 300/300	sum-out) 400/400	1600/1600	15	5 uw
CD4009	hex inverter/level converter		15/30		50	100 nw
CD4010	hex non-inverting level conv./driver		20/30		50	100 nw
CD4011	Quad 2-input NAND		65/80		50	100 nw
CD4012	Dual 4-input NAND		130/80		50	10 nw
CD4013	Dual D Flip-Flop	800	190	130	30	50 nw
CD4014	8 bit parallel-in/ parallel-out reg.	1000	250	200	15	5 uw
CD4015	Dual 4-bit serial-in/ parallel-out reg.	1000	250	200	15	5 uw
CD4019	Quad AND-OR Select gate					500 nw
TA5577	64 x 1 R/W RAM		Read Time 50			30 uw

All values are typical. Total power = $C_{out} V_{DD}^2 f + P$ quiescent

Supply voltage range: 6-15 volts.

Input Capacitance = 5 pf

Table 2-1 (Continued)

FAIRCHILD 93L SERIES LPTTL MSI

Part No.	Description	Typical Delay (ns)	Typical Power (mw)
93L00	4 Bit Shift Register	60	75
93L28	Dual 8 Bit Shift Register	56	75
93L18	8 Input Priority Encoder	55	70
93L24	5 Bit Comparator	55	55
93L40	4 Bit Arithmetic Logic Unit	85 (135 ns for 16 bits)	400
93L01	One of Ten Decoder	63	35
93L11	One of Sixteen Decoder	70	40
93L21	Dual One of Four Decoders	49	40
93L09	Dual Four Input Multiplexer	48	40
93L12	8 Input Multiplexer	80	34
93L22	Quad 2 Input Multiplexer	44	45
93L08	Dual 4 Bit Latch	53	90
93L14	4 Bit Latch	68	55
93L10	Up Decade Counter	45	75
93L16	Up Binary Counter	45	75

Table 2-1. (Continued)

TEXAS INSTRUMENTS 54L SERIES LPTTL

Part No.	Description	tpd_0(ns)	tpd_1(ns)	Power (mw)
SN54L00	Quad 2 Input NAND	31	35	2.8
SN54L04	Hex Inverter	31	35	2.8
SN54L10	Triple 3 Input NAND	31	35	2.8
SN54L20	Dual 4 Input NAND	31	35	2.8
SN54L30	8 Input NAND	70	35	2.0
SN54L51	Dual 2 Wide AND-OR-INVERT	35	50	4.0
SN54L54	4 Wide 3-2-2-3 Input AND-OR-INVERT	35	50	7.2
SN54L55	2 Wide 4 Input AND-OR-INVERT	35	50	4.0
SN54L71	R-S Master Slave Flip-Flop	35	60	11.2
SN54L72	J-K Master Slave Flip-Flop	35	60	11.2
SN54L73	Dual J-K Master Slave Flip-Flop	35	60	11.2
SN54L78	Dual J-K Master Slave Flip-Flop	35	60	11.2
SN54L86	Quad 2 Input EOR	35	50	21.0
SN54L93	4 Bit Binary Counter	280	280	52.0
SN54L91	8 Bit Shift Register	100	55	52.0
SN54L95	4 Bit Parallel In/Out Shift Left-Shift Right Register	125	115	72.0

C_{load} = 50 pf Fanout Capability = 10

R_{load} = 4Kohms

Power Measured with V_{cc} = 8V

Speed Measured with V_{cc} = 5V

Table 2-1. (Continued)

SOLID STATE SCIENTIFIC CMOS		t_{pd}		load	Quies-cent
Part No.	Description	6V	10V	Cap.(pf)	Power
SCL 5101	Quad 2-Input NOR	200	125	50	50 nw
SCL 5102	Quad 2-Input NAND	200	125	50	50 nw
SCL 5103	Dual 3-Input NOR	200	125	50	50 nw
SCL 5104	Dual 3-Input NAND	200	125	50	50 nw
SCL 5105	Dual 4-Input NOR	200	125	50	50 nw
SCL 5106	Dual 4-Input NAND	200	125	50	
SCL 5201	Quad EOR	200	125	50	
SCL 5402	8-bit Parallel In/Out Register	200	30	30	10 uw
SCL 5401	Presettable 8-bit Counter	175	30	30	10 uw

Supply voltage range = 6-20 volts.

Input Capacitance = 5.0 pf.

Total power = $C_{out}V_{DD}^{2}f + P_{quiescent}$

3. EVALUATION OF MICROPROGRAMMING

This section defines microprogramming and illustrates the advantages gained and limitations encountered when replacing hard wired control with such an approach.

The use of memories to produce the sequence of control signals necessary to execute a machine instruction is referred to as microprogramming. Each machine instruction (macroinstruction) is broken into a sequence of microinstructions, each of which contains all the control signals which are needed during one clock interval. Each output bit of the control memory represents a control signal which may be either high or low depending on the instruction being executed. At the beginning of a microinstruction, the first microinstruction of the microprogram is addressed by the operation code obtained from the main program memory. Successive microinstruction addresses are derived from a combination of a field of control bits of the present microinstruction and control signals dependent on the results of previous operations. The output of the memory, either directly or by conditioning with data dependent signals, generates the signals which perform data transfers and manipulations within the system.

The microprogramming approach was evaluated as a method of improving the overall performance of the Advanced On Board Processor. The basic microprogramming philosophy supports a reduction in hardware, a more feasible approach for logic partitioning and a more flexible design for the instruction set. These factors, along with power, speed and cost were

used to evaluate the implementation of the microprogramming technique. A block diagram of a microprogrammed CPU is presented in Figure 3-1.

3.1 IMPROVEMENTS FROM MICROPROGRAMMED CONTROL

The following sections will discuss the gains that are expected when implementing a microprogrammed approach for CPU control.

3.1.1 Hardware Reduction

The first generation OBP uses approximately 1100 gates for the control logic. Since reduction of the physical size of the machine is one of the prime objectives of the AOP design, it is important to consider any technique that reduces the number of circuits. The use of a memory in the control circuitry of the CPU eliminates the need for an instruction hold register, associated decoding logic, and instruction phase control logic. It has been estimated that roughly 70 percent of the control logic could be replaced by approximately fifteen 256 x 8 LSI memory chips.

3.1.2 Logic Partitioning

Whatever logic family is chosen for the AOP, it is certain that more compact basic packaging units will be used. These will be either MSI/LSI chips or special purpose printed circuit boards containing SSI circuits. In either case, it will be necessary to partition the control logic as efficiently as possible into conveniently sized and configured blocks. This task is greatly simplified when the logic structure to be partitioned is regular and symmetrical. However, computer control logic is inherently almost devoid of symmetry. If a control memory, with its very regular structure, is used to implement a majority of the OBP control functions, the partitioning task is accomplished much more readily.

There is a certain amount of logic (approximately 30 percent of present control logic) which cannot be replaced with memories when implementing the microprogramming approach. This logic generates data dependent control signals and performs such functions as testing the contents of registers or counters, testing sign bits, and testing for the completion of operations before allowing subsequent ones to proceed. This logic presents the same partitioning problems as those encountered with the existing control logic. It is, however, a much smaller group of circuits and fewer partitioned structures and/or structure types will be required for the partitioning task.

A preliminary partitioning design of the required circuitry external to the control memory was delivered to Goddard Space Flight Center personnel. The purpose of this task was to group the logic in such a manner as to reduce interconnects between the partitioned structures. The partitioning design was accomplished with three logic structures of three different types, each type requiring a maximum of 80 pins.

An effort was then initiated by NASA to commit these designs to LSI chips. The response from semiconductor manufacturers was unsatisfactory and no further effort was expended in this area.

3.1.3 Instruction Set Flexibility

One of the more important aspects of implementing a microprogramming approach is the ease gained in changing the instruction set. The instruction set of a microprogrammed computer can be changed in many cases merely by reprogramming the control memory. This is all that is required for the addition of instructions which do not involve data dependent conditions. Reprogramming is particularly easy if a non-volatile read/write memory is used

for the control source. If a read only memory is used, excess word locations can be used to generate control signals for new instructions or for implementing a change to existing instructions. In a new machine requiring a smaller instruction set, the circuitry could be reduced by a reduction in the control memory. If new conditions must be added to control lines for an instruction change, then more substantial hardware changes are required. However, even for conditional instructions, changes appear easier with a microprogrammed machine.

3.1.4 System Testing

The use of microprogrammed control logic will facilitate the task of testing the system simply because it has reduced the number of logic gates by roughly 70 percent. The task of testing the memory for proper operation is much simpler than that of testing the 70 percent of 1100 control gates it replaces.

In addition, the memory generates most of the control signals for the CPU. Many of these control signals are used to execute a variety of instructions. Since multiple use of the control lines is employed in the CPU, once these lines are verified for one instruction, the task of checking the remaining system is simplified.

3.2 CHOICE OF MEMORIES

The choice of memory for the AOP control logic is governed largely by the environment in which the computer must operate. Because of the size, power and weight constraints, it is assumed that plated wire and magnetic core memories are unsuitable. Therefore the following discussion applies only to semiconductor memories.

Since it is desired that the OBP be capable of operating after inactive periods with power removed, it is important that the control memory be non-volatile. If an otherwise acceptable volatile memory required very little standby power, it could possibly be powered while the rest of the system is shut down. However, this is an extra complication and should be avoided if possible.

The memory should also ideally be alterable to allow for design changes, last minute corrections, etc. Alterability is in general a characteristic only of volatile semiconductor memories and so may have to be sacrificed. The only known device which combines the two features is the Metal Nitride Oxide Semiconductor (MNOS) memory. Although Westinghouse is working on the development of this new technology, it would be impossible to obtain an MNOS memory during the time period of interest for the Advanced On Board Processor.

Because bipolar circuits consume so much chip area, it is difficult to produce large bipolar memory chips. Instead manufacturers have concentrated on the speed advantages of TTL circuits to produce several high speed, high power devices of relatively low complexity. It therefore appears that a semiconductor memory suitable for the OBP will be built with MOS circuits. A table shown the performance characteristics of several semiconductor memories is shown in Table 3-1.

It is apparent that the CMOS Read/Write (R/W) and Read-Only (ROM) memories demonstrate definite speed and power advantages over the PMOS devices. The R/W memory should be on the market in the near future. It is expected when the CMOS ROM becomes available, it will equal the chip complexity of the PMOS devices and offer a much better speed power product.

If microprogramming is implemented, it is recommended that the brass board computer be built with a Read/Write control memory. After debugging the brassboard and verifying the microprogram, then a ROM can be substituted.

3.3 ALTERNATIVES TO MICROPROGRAMMING

The adoption of microprogrammed control necessitates a completely new design for the OBP control logic. Therefore, any reasonable alternatives which permit the retention of the present logic must be evaluated.

If the Advanced On Board Processor is built with SSI devices, the control logic could be left in its present form, requiring only those changes dictated by the new instruction set. This approach can be justified by a power savings with CMOS circuits or an increase in speed with LPTTL circuits. In addition, the logic can be partitioned and put on special purpose multilayer printed circuit boards, resulting in a decrease in the physical size of the machine with no drastic design change.

Because of cost considerations it is unlikely that completely custom designed LSI devices of any type will be used for this computer. However, a further reduction in size and in the speed-power product can be obtained by using LPTTL LSI gate arrays. This approach is being investigated at the present time and it appears that the technology will be available for the Advanced On Board Processor.

Preliminary estimates indicate that approximately 25 of these chips will be required to implement the OBP control logic in its present form. The cost of these chips is of the same order of magnitude as the cost of ROMs in small quantities. This approach would permit nearly the same size

reduction possible with microprogramming and would obviate designing and debugging a completely new control section.

3.4 CONCLUSIONS ON MICROPROGRAMMING

If a microprogramming approach is adopted, the best choice of memories seems to be the CMOS ROM being developed by RADIATION. This device outperforms all the PMOS ROM's in terms of both speed and power. In the event that it does not become available, the second best choice is probably the CMOS R/W memory. Even allowing for its smaller size and volatility, its high speed and extremely low standby power justify its selection. Other companies are presently developing CMOS R/W and ROM memories and these units may be available in the near future.

The choice between changing to microprogramming and keeping the present control design depends largely on the type of circuits used for the rest of the computer. If small scale integrated circuits are used for the register logic, the advantages of reduced size, increased regularity of structure, ease of partitioning, and instruction set flexibility can be claimed for microprogramming. It should be noted that there will be a loss in operating speed with this approach due to the access time required for the ROM output control lines. Use of a control memory dictates that time be allowed to set a synchronous address register and then access the memory. The time necessary for these operations is considerably longer than that required to setup control signals with the present system.

If LSI-Gate arrays are used to implement the register logic, it seems reasonable to also use them for the control logic. They will apparently realize the same savings in physical size as an LSI control memory and at no

greater cost. The LSI-Gate arrays offer the advantage of providing a proven design which has been thoroughly tested while improving the speed-power product for the system.

Table 3-1. Semiconductor Memories

Manufacturer	Part No.	Dimensions	Technology	R/W or ROM	Power	Access Time
INTEL	3101	16 x 4	TTL	R/W	550 mw	40 ns
INTEL	3301	256 x 4	TTL	ROM	650 mw	60 ns
FAIRCHILD	Mu9035	16 x 4	TTL	R/W	649 mw	36 ns
INTEL	1101	256 x 1	Si Gate	R/W	225 mw	1 us
NATIONAL	MM421	256 x 4	PMOS	ROM	240 mw	600 ns
NATIONAL	MM423	256 x 8	PMOS	ROM	290 mw	850 ns
PHILCO-FORD	pM S1024C	128 x 8	PMOS	ROM	130 mw	2 us
UNION-CARBIDE	ROM1K	128 x 8	PMOS	ROM	---	1 us
FAIRCHILD	3501	128 x 8	PMOS	ROM	120 mw	2.5 us
SOLID STATE SCIENTIFIC	SCL5553	256 x 1	CMOS	R/W	600 u watt/ MHz	250 ns
HARRIS	RAM0256	32 x 8	CMOS	R/W	5.0 mw/ MHz	200 ns
HARRIS		256 x	CMOS	ROM	---	200 ns

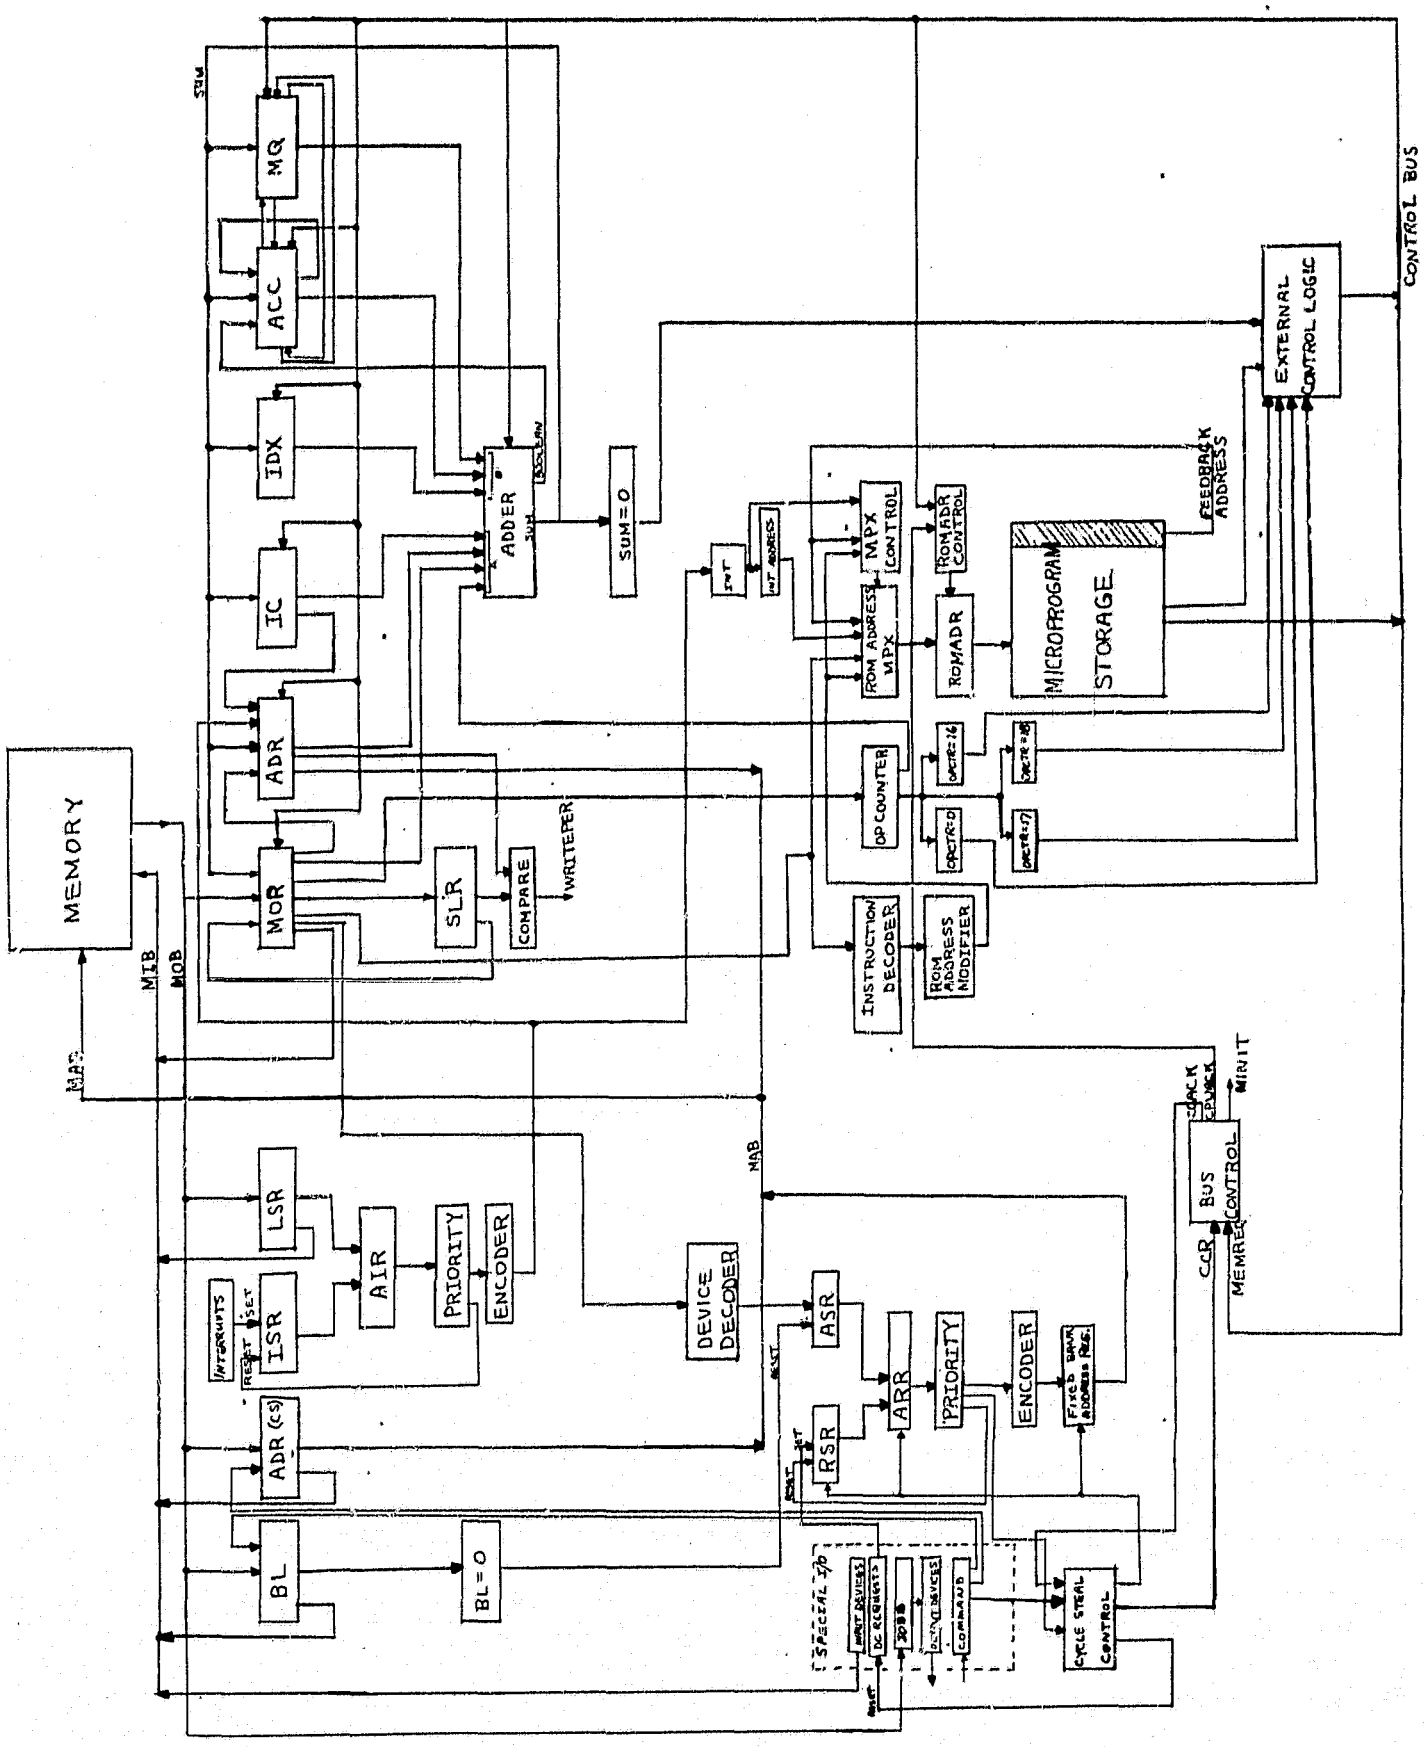

Figure 3-1. Microprogram Control

4. EVALUATION OF INTEGRATING THE FIXED I/O WITH THE CPU

The integration of a section of the I/O with the CPU was proposed for two reasons. First, the present I/O has certain circuitry which will remain fixed from mission to mission while other areas of the I/O must be changed to meet the various mission requirements. By combining the fixed portion of the I/O with the CPU, the fixed section would not be redesigned for each mission and in addition a much less complex special I/O will result. Second, the combination of these units should reduce the total number of circuits used when compared to the present CPU and I/O. The reduction in circuitry is accomplished by packaging the communicating sections of the CPU and I/O together. This logic reduction is expected to improve packaging feasibility, systems performance and reliability.

The basic evaluation of the CPU-I/O combination began with a study of the I/O drawings. After a brief study period, meetings were held with NASA personnel to evaluate combination possibilities. The meetings were organized to discuss timing problems in the present I/O. Some of the solutions to these problems were not obvious from the logic drawings.

The original consideration, in combining these units, was to utilize a serial IN-OUT data transfer between the special I/O and the CPU. This method of transfer would require a high speed shift register for the input and output channels. A cycle steal operation or I/O instruction could perform the data transfer within a typical CPU clock cycle by utilizing a 20 MHz clock for the shift register. The availability of low power, high

speed shift registers and the prospect of reducing the number of interface lines demonstrated the practicality and attractiveness of this approach.

It was subsequently found that a parallel operation utilizing the bus for transfers was the best approach to take for the CPU-I/O design. At this point in time the design study evolved into a redesign of the CPU and the I/O, incorporating the fixed section of the I/O into the CPU. During this phase it became evident that the parallel transfer technique would eliminate logic and improve reliability.

4.1 CHANGES ASSOCIATED WITH THE CPU-I/O COMBINATION

The combination of the CPU and fixed portion of the I/O introduced changes to the OBP system. These changes affected three areas: Interface circuitry and signals; I/O circuitry and redesign; and System philosophy. The objective of the combination of these changes was to improve the overall system performance. This was achieved by hardware changes and procedural improvements which are discussed in detail below.

4.1.1 Circuit Reduction

Circuit changes were incorporated in two areas: interface circuitry and reductions through a CPU-I/O redesign. One restriction placed on the combination was the limitation of interface signals. Thus, any design change must taken this area into consideration.

The reduction in circuitry was limited and pertained only to the I/O design. While the I/O redesign reduced existing I/O logic, new logic was added in the CPU for sequencing and control functions. The sum total of this effort was a large circuit reduction in the I/O with an increase of circuitry in the CPU-I/O unit. A small total hardware decrease resulted from the redesign.

The interface circuits were a relatively small part of the circuit changes incorporated in the combination of the units. Any savings from the change in the interface structure was consumed by the new system requirements. As a result, an increase in interface signals and circuitry was required in the final system configuration. Approximately 54 lines are required to communicate between the CPU and I/O. These lines include the 32 request-acknowledge lines, 16 interrupt lines, 4 command lines, clock lines and possibly a few control lines not completely defined at this time.

One area of circuit reduction through design change was in the cycle steal logic. The initial design utilized four registers and associated control logic for the cycle steal routine. These four registers were replaced by one register used as a block length/address register. This register is an up/down counter that is shared for those operations. During a command request the addresses are generated by the 12 least significant bits of this register in conjunction with four gates tied on the address lines.

Since only one block length register is used now, only one set of block length test gates is required to test the block length for zero. With this change came the addition of a separate address register for use in accessing the fixed memory locations reserved for the cycle steal address and block length values. This is an 11 bit register with seven fixed bank address bits and 4 bits for selecting the desired 4K block of memory. Certain bits may be hard wired onto the bus. A six bit shift register was added to the design to sequence the cycle steal operation. Also added to the design was an activation status register which allows requests from selected channels to be honored by the CPU. This register acts on requests before the priority logic selects the request to be processed.

A further reduction in circuitry was achieved by the elimination of the I/O buffer register. The original I/O unit utilized a buffer register (IOBB) to hold data that was being transferred either between the CPU and I/O or between the memory and I/O. The combination of the CPU and I/O eliminated the need for such a buffer in the CPU since data could be clocked directly into the I/O registers located in the CPU. With the elimination of the IOBB, it was felt that data transfer between the units could be accomplished via the bus. Since the bus is necessary for data linkage between the CPU, memory and I/O, it is reasonable to use it for all data transfers. In this manner, a reduction in interface connections and hardware could be achieved. Thus, the idea of utilizing shift register and serial transfers was discarded for the direct transfers between the memory and CPU, and the memory and special I/O.

The bus control logic has also been reduced. This results from the sharing of common memory request logic by all I/O cycle steal channels. This reduction also includes the elimination of those circuits used to expand or stretch the acknowledge signals.

4.1.2 Timing

The timing associated with the new system design was investigated for two separate reasons. First, the memory cycle time has been set up as the limiting factor in the timing chain. The use of the adder in the fetch-add operation is the only function which may require more than a memory cycle time for completion. The worst case situation occurs with a one bit carry ripple adder and a microprogrammed control unit. The present investigation indicates that a three bit adder will be used and that the microprogramming technique will not be used. If this approach is taken, the memory will be

the critical timing chain unless the LSI gate arrays discussed in the circuit area are much slower than indicated by the manufacturers.

The second area of timing consideration has been with the memory initiate and bus control logic. These signals may increase the clock period if delays are incorporated into the system to eliminate the possibility of transient address data on the busses.

4.1.3 Fixed Memory Location

The changes made to the CPU and I/O also required changes in the fixed memory locations. The original system specified fixed memory locations for the interrupts, the EXIT instruction and the I/O instruction. These locations were changed to allow the interrupt storage to start in location zero. The cycle steal fixed bank will start in location 200 and EXIT will start in location 240. No fixed locations are required for the new I/O instruction. Additional gating will be required since the cycle steal locations do not begin in location zero.

4.1.4 Procedural Changes

The major system improvement comes from the improved system performance obtained by procedural changes in servicing interrupts, cycle steal routines, and command servicing routines. A brief description of each function is given below.

The basic interrupt routine is quite similar to that used in flight model OBP system. The one major change was the addition of control logic for two instructions necessary to override interrupts.

The cycle steal routine now utilizes a fixed memory bank for storing address and block length values for cycle steal routines. Thus

one cycle steal register allows access to 16 different channels. A total of 16 devices can be decoded for I/O instruction and cycle steal routines.

The LOAD and DUMP command requests now have their own associated control lines and are completely independent of all other CPU-I/O operations. These commands have top priority over any CPU operation.

4.2 INSTRUCTION SET CHANGES

The combination of the fixed I/O and CPU brought about the following changes to the instruction set of the flight model OBP. First, two new minor op-code instructions were added to the instruction repertoire to set or reset the INT OVERRIDE flip-flop. The INT OVERRIDE flip-flop was added to the design for use in overriding all interrupts (except INITIATE) through program control. Since one instruction is permitted between interrupts, the SET OVERRIDE instruction will override the interrupts and give control of the processor to the programmer so that he may perform necessary housekeeping operations. The final instruction used by the programmer will reset the INT OVERRIDE flip-flop prior to turning processor control over to the next interrupt.

In addition to the override instructions, the complete I/O instruction organization has been changed. Initially, the same I/O instruction organization as that used in the flight system was used in the Advanced On Board Processor except for the "connect to" operation. The "connect to" operation was used to set or reset the appropriate activation status flip-flop prior to I/O operations. Bit 16 of the "connect to" word determined the set/reset condition, and the remaining 15 bit code was used to select the desired flip-flop. It was decided that two major op-code instructions

would be used to execute all I/O operations including the "connect to" function. The "connect to" can be replaced by treating the activation status register as a separate output device and decoding it as such.

Two major op-code instructions, LET OUTPUT TO and LET INPUT FROM, have been added to the CPU instruction repertoire. These two instructions will perform all of the original I/O insturctions. The LET OUTPUT TO and LET INPUT FROM instructions use location (n) and (n+1) for execution. The device code is stored in (n) and is used to select a particular input or output device. The next memory cycle either outputs the contents of location (n+1) to the selected output device or stores data from an input device in (n+1). The implementation of these two instructions required the deletion of two existing major op-code instructions due to the lack of available major op-codes. These two instructions replaced the original I/O instruction and the "EXECUTE" instruction. The I/O instruction was removed since it was no longer used. The "EXECUTE" instruction was deleted since it received limited use in the existing system programs. The elimination of the "EXECUTE" instruction will also eliminate the special hardware required to generate this instruction.

The "connect to" instruction initially set or reset the selected activation status flip flop. While this operation is still required, a new procedure is used. A LET OUTPUT TO instruction is decoded, and then the activation status register is selected by the device code in location (n) as one of the I/O devices. The following fetch of location (n+1) is decoded through the same decoder as that used for device decoding and is used to select the desired activation status flip-flop. One bit of the (n+1) word will be used for set/reset control.

The LET FUNCTION TO operation was initially used to set up discretes and relays. This operation is still required, and when the LET FUNCTION TO operation was deleted, an alternative was necessary. The discretes and relays are now treated as a separate device and are decoded as such. They are set or reset according to the field present in location (n+1) of a LET OUTPUT TO instruction.

In addition to these changes in the instruction set, the accumulator and location seven are no longer used for the I/O instruction. A complete list of instructions for the Advanced On Board Processor are given in Table 3-1.

4.3 INTERRUPT STRUCTURE

The basic AOP interrupt philosophy will not change appreciably from that implemented in the flight model OBP system. In the original system, the interrupt logic was contained in the I/O unit. The interrupt code was generated on a priority basis and sent to the CPU as a four bit address code for servicing the selected interrupt.

The new interrupt structure permits the use of up to 16 interrupts. These 16 interrupts are honored according to a set of priorities determined by mission requirements. The interrupts are stored asynchronously in the interrupt storage register. Following each interrupt routine will be one CPU instruction execution. In the flight model design, this instruction was available to give the programmer some control over the system and to perform the necessary housekeeping tasks. As a result of previous program experience, it was decided that the programmer could not exercise sufficient control over the system. Therefore, two new instructions were added to the CPU instruction set. These instructions set and reset a flip-flop used to

override all interrupts except "INITIATE." The interrupt structure now permits the programmer to have complete override capabilities on all interrupts so that he may perform his necessary housekeeping functions. While the INT OVERRIDE flip-flop is set, any interrupts which occur will be stored in the interrupt storage register for future servicing.

In addition to the override capability given to the programmer, a lockout capability still exists in the interrupt structure. When an interrupt or EXIT instruction is executed, the lockout status register is updated by the value contained in a given memory location. The value of the word is controlled so that the programmer may lockout specified interrupts. The status of this register is stored during each interrupt routine.

A total of 16 lines will be necessary to interface with the special I/O for interrupt specification.

4.4 CYCLE STEAL ROUTINE

The original system utilized two cycle steal channels for faster I/O operation and reliability through redundancy.

A block length register and address register were associated with each channel. When a cycle steal operation was initiated (either by an I/O request or by command), one of the cycle steal devices was given control of the channel. The block length register was decremented by one and the address register was incremented by one during each cycle steal. Requests from the device were honored until the block length was decremented to zero. When zero was reached, an interrupt was generated.

The basic change to the cycle steal routine consisted of utilizing only one cycle steal register arrangement, i.e., one block length register and one cycle steal address register. To keep track of information required

for cycle stealing, certain fixed memory locations are used to store the address and block lengths associated with each of 16 selectable channels. A variety of methods for cycle stealing were considered, with each method utilizing a separate address register for accessing the fixed location of the block length and address.

First, five memory cycles were used for each cycle steal operation. The first cycle fetched the address from (n) and stored it in the address register. The second cycle fetched the block length from (n+1) and stored it in the block length register. The third cycle either fetched or stored the desired word and decremented the block length. The fourth cycle stored the block length in (n+1) and updated the address. Finally the address was stored in (n) to complete the cycle steal routine. If a block length equal to zero was detected during the routine, an interrupt was generated.

The second approach to the cycle steal routine dealt with a change to the interrupt generation. It was suggested that one interrupt be used to detect a block length of zero for all channels. This approach was not satisfactory since there was not enough information available for the programmer when trying to detect which block length went to zero. As a suppliement to this approach, the block length was stored in its fixed memory location, with a bit being set to indicate that the block length had reached zero and had not been examined by the programmer. Since only 12 bits are required for the block length, the most significant bit was used for this purpose. This approach was too involved and did not allow the programmer fast access to the correct interrupt. It was necessary to scan the entire set of fixed cycle steal locations everytime these interrupts were examined. The final

decision on the block length interrupts was to generate a separate interrupt for each channel when the associated block length goes to zero.

A third approach evaluated for the cycle steal routine was arranged to save time when honoring a cycle steal request whose block length had not gone to zero. If the same request was being honored consecutively, it was suggested that only the data word fetch or store cycle be executed. This operation was performed when a request from the same device and a block length not equal to zero were detected. If the block length goes to zero after the fetch or store, the address and block length are stored in the fixed cycle steal locations and an interrupt is generated. The control associated with the detection and execution of this cycle steal organization is undesirable, because it increases the required amount of non-symmetrical hardware, and is therefore not amenable to efficient partitioning.

In the final approach considered, the fetch and zero detection of the block length was done first. If the block length was zero, then no further execution time was taken for the cycle steal. The associated channel activation status flip-flop was cleared, thereby locking out further requests from that device. As each block length reached zero, a unique interrupt was generated. When utilizing this approach, it was decided that one register could be used as both a block length and an address register. This is possible since the new approach does not require that both registers be in use at the same time. The five memory cycle routine is maintained through the following sequence. The block length is fetched in the first clock cycle, the block length is tested for zero and an interrupt is generated, or decremented during the second clock cycle, the block length is stored during the third clock cycle, the address is fetched during the fourth clock cycle,

the data word is fetched or stored and the address is updated during the fifth cycle, and the address is stored during the sixth clock cycle. Following each cycle steal routine, one CPU clock cycle will be honored before servicing another cycle steal request. The CPU will function normally during a cycle steal routine if executing a non-memory request instruction; e.g., completing a divide, multiply or normalize routine.

4.4.1 Activation Status Register

The change of the cycle steal routine led to the decision to have some type of program control over the requests generated in the special I/O. The first consideration was the implementation of a register that could be set under program control to allow or inhibit requests from being honored. This register was first set with a "connect to" instruction. This allowed the programmer to enable all channels or disable all channels at one time.

Bits 17 and 18 of the "connect to" word were used to decode a "connect to" and bit 16 was used as a set/reset indicator for the selected flip-flops. The remaining 15 bits were used to select the flip flops that were to be set or reset. To protect the activation status register, it was decided that only one bit of the 15 bit field could be used for flip-flop selection. This still did not protect against an error which could change all the flip-flops of the register. As a final solution, a decoding scheme was chosen to insure that only one flip-flop could be set or reset at a time during a "connect to" instruction. The use of an activation status register adds circuitry to the CPU and for this reason was temporarily discarded during one point in the design.

As an alternative to this approach, the possibility of using only a priority scheme to resolve conflicts between requests was considered. This scheme would operate in the same manner as the interrupt logic without a lockout status register. In addition, a flip-flop, under program control, would be available to lockout all cycle steal requests except command requests. This flip-flop was to be set with a IET FUNCTION TO operation. It was later decided that this approach did not provide enough flexibility.

As the investigation of the I/O organization progressed, it was decided that the activation status register would be incorporated in the CPU design. This register will be set or reset under program control (decoded as a device in a IET OUTPUT TO instruction) with one bit at a time being acted upon through a decoding network. An activation status flip-flop will be reset whenever a block length equal to zero is detected. Once a flip-flop is set, it will remain set until changed by program controlor by a block length going to zero for that channel.

4.5 COMMAND REQUEST EXECUTION

The command execution is changed considerably from the flight model OBP system. Two approaches to honoring command requests were evaluated during the study and are explained below.

The first method considered, for executing command operations, was integrated into the initial cycle steal routine. This routine required five clock cycles. Three normal clock cycles were inhibited during the execution of a command request. The execution of the command inhibited the address and block length fetch, and clocked these registers with zeros. The useable address was stored in location zero and the block length register was decremented so that all 1's were stored in the (0+1) block length

location. The next command cycle fetched the contents of location zero and the command load begins at the address specified during the first command cycle. If a command dump was to be performed, a command load was executed to program the memory for the command dump routine. The command was initially taken as the highest priority request which could not be locked out or inhibited.

The above mentioned approach to command operations did not appear to be the most satisfactory one for executing command routines. Further study led to the adoption of the following command load and dump routines.

Both the load and dump commands are controlled in the system with separate request lines. Each command request line has an associated acknowledge line. The command load operates as follows. First, a master clear is performed on the CPU. Secondly, the command load request line goes high. After this line goes high, the CPU acknowledges and writes one word in location zero. The command request line must drop and go high again prior to the CPU honoring a second command request. The command load overrides all CPU operations and loading continues until no more requests are received. The block length/address register associated with the cycle steal routine is used as an address register for the command load and dump routines.

The command dump operates in a manner similar to the command load. A master clear is performed on the CPU, the command dump request goes high, and the word from location zero is dumped. The dump operation continues as long as the command dump request line falls and rises until the same 4K of memory is dumped twice. The double 4K dump is controlled by bit 14 of the address register. When bit 14 of the register is set in the dump mode, dumping is inhibited until a master clear is performed on the CPU and the

command dump request line goes high again. The use of bit 14 is possible since address bits 13-16 are generated by the fixed bank address lines which select the desired 4K memory block.

The command lines have lockout control latches on each line. These latches prevent the commands from taking control of the system if they fail. The line must go low and high for the execution of each word load or dump.

Table 4-1. Advanced On Board Processor Instruction Set

MAJOR OP-CODES	MINOR OP-CODES
1. ANDED WITH	1. CLOSE EXTENSION WITH DECISION
2. CYCLED BY	2. COMPLEMENTED
3. DIVIDED BY	3. EXIT
4. DOUBLE CYCLED BY	4. HALT
5. DOUBLE SHIFTED BY	5. IS ZERO
6. EORED WITH	6. IF OVERFLOW
7. GO TO	7. IF PARITY ODD
8. IND LET	8. INC SUBSCRIPT NE
9. IND YIELD	9. INC EA NE
10. IS EQUAL TO	10. IS FALSE
11. IS GREATER THAN	11. IS POSITIVE
12. IS LESS THAN	12. NEGATED
13. LET	13. NORMALIZED
14. LET INPUT FROM	14. PASS
15. LET LOCATION OF	15. PLUS CARRY
16. LET OUTPUT TO	16. RESET D
17. MINUS	17. RESET INT OVERRIDE
18. ORED WITH	18. RESET OVERFLOW
19. PERFORM	19. REVERSED
20. PLUS	20. SET INT OVERRIDE
21. RESUME FROM	21. SET PAGE
22. SAVE EXTENSION IN	22. X A SUBSCRIPT
23. SAVE SUBSCRIPT IN	23. X A EA
24. SET EXTENSION WITH	24. X EA SUBSCRIPT
25. SHIFTED BY	
26. STEP SUBSCRIPT BY	
27. SUBSCRIPT NOT GREATER THAN	
28. THEN GO TO	
29. TIMES	
30. USE SUBSCRIPT	
31. YIELD	

Table 4-1. Advanced On Board Processor Instruction Set

MAJOR OP-CODES	MINOR OP-CODES
1. ANDED WITH	1. CLOSE EXTENSION WITH DECISION
2. CYCLED BY	2. COMPLEMENTED
3. DIVIDED BY	3. EXIT
4. DOUBLE CYCLED BY	4. HALT
5. DOUBLE SHIFTED BY	5. IS ZERO
6. EORED WITH	6. IF OVERFLOW
7. GO TO	7. IF PARITY ODD
8. IND LET	8. INC SUBSCRIPT NE
9. IND YIELD	9. INC EA NE
10. IS EQUAL TO	10. IS FALSE
11. IS GREATER THAN	11. IS POSITIVE
12. IS LESS THAN	12. NEGATED
13. LET	13. NORMALIZED
14. LET INPUT FROM	14. PASS
15. LET LOCATION OF	15. PLUS CARRY
16. LET OUTPUT TO	16. RESET D
17. MINUS	17. RESET INT OVERRIDE
18. ORED WITH	18. RESET OVERFLOW
19. PERFORM	19. REVERSED
20. PLUS	20. SET INT OVERRIDE
21. RESUME FROM	21. SET PAGE
22. SAVE EXTENSION IN	22. X A SUBSCRIPT
23. SAVE SUBSCRIPT IN	23. X A EA
24. SET EXTENSION WITH	24. X EA SUBSCRIPT
25. SHIFTED BY	
26. STEP SUBSCRIPT BY	
27. SUBSCRIPT NOT GREATER THAN	
28. THEN GO TO	
29. TIMES	
30. USE SUBSCRIPT	
31. YIELD	

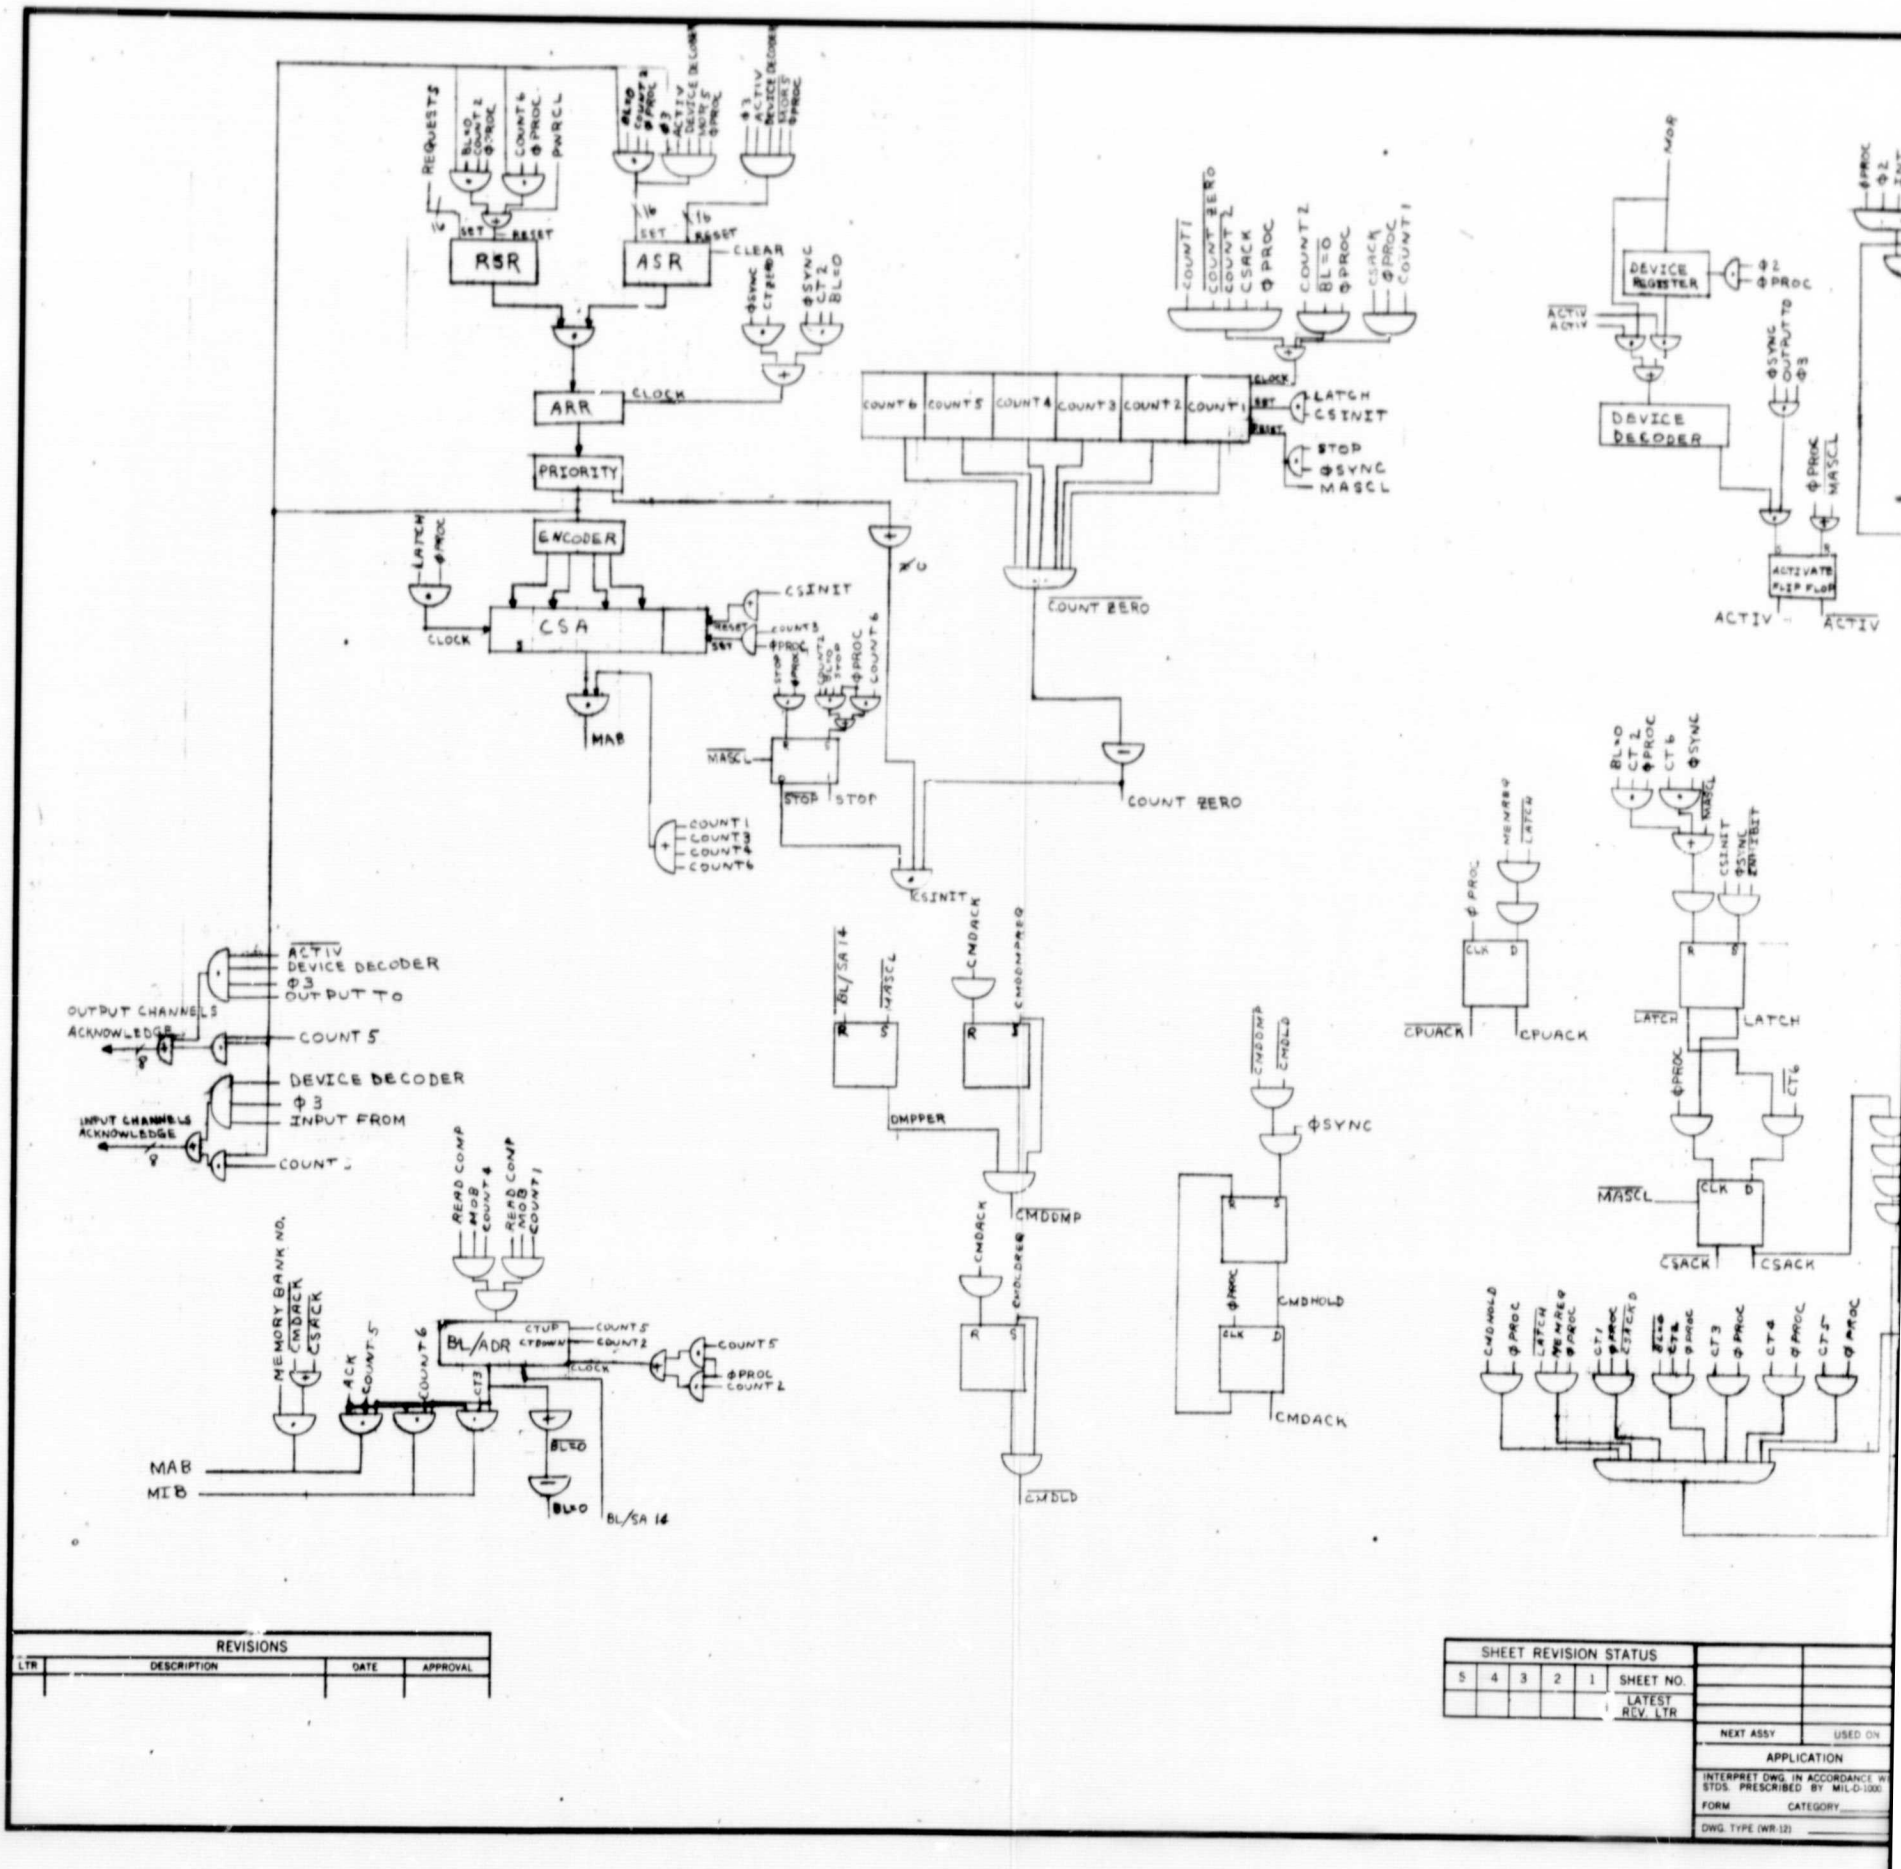

REVISIONS			
LTR	DESCRIPTION	DATE	APPROVAL

SHEET REVISION STATUS					
5	4	3	2	1	SHEET NO.
					LATEST REV. LTR

NEXT ASSY | USED ON

APPLICATION

INTERPRET DWG. IN ACCORDANCE WI STDS. PRESCRIBED BY MIL-D-1000

FORM | CATEGORY

DWG. TYPE (WR-12)

Figure 4-1. Fixed I/O

FOLDOUT FRAME 1

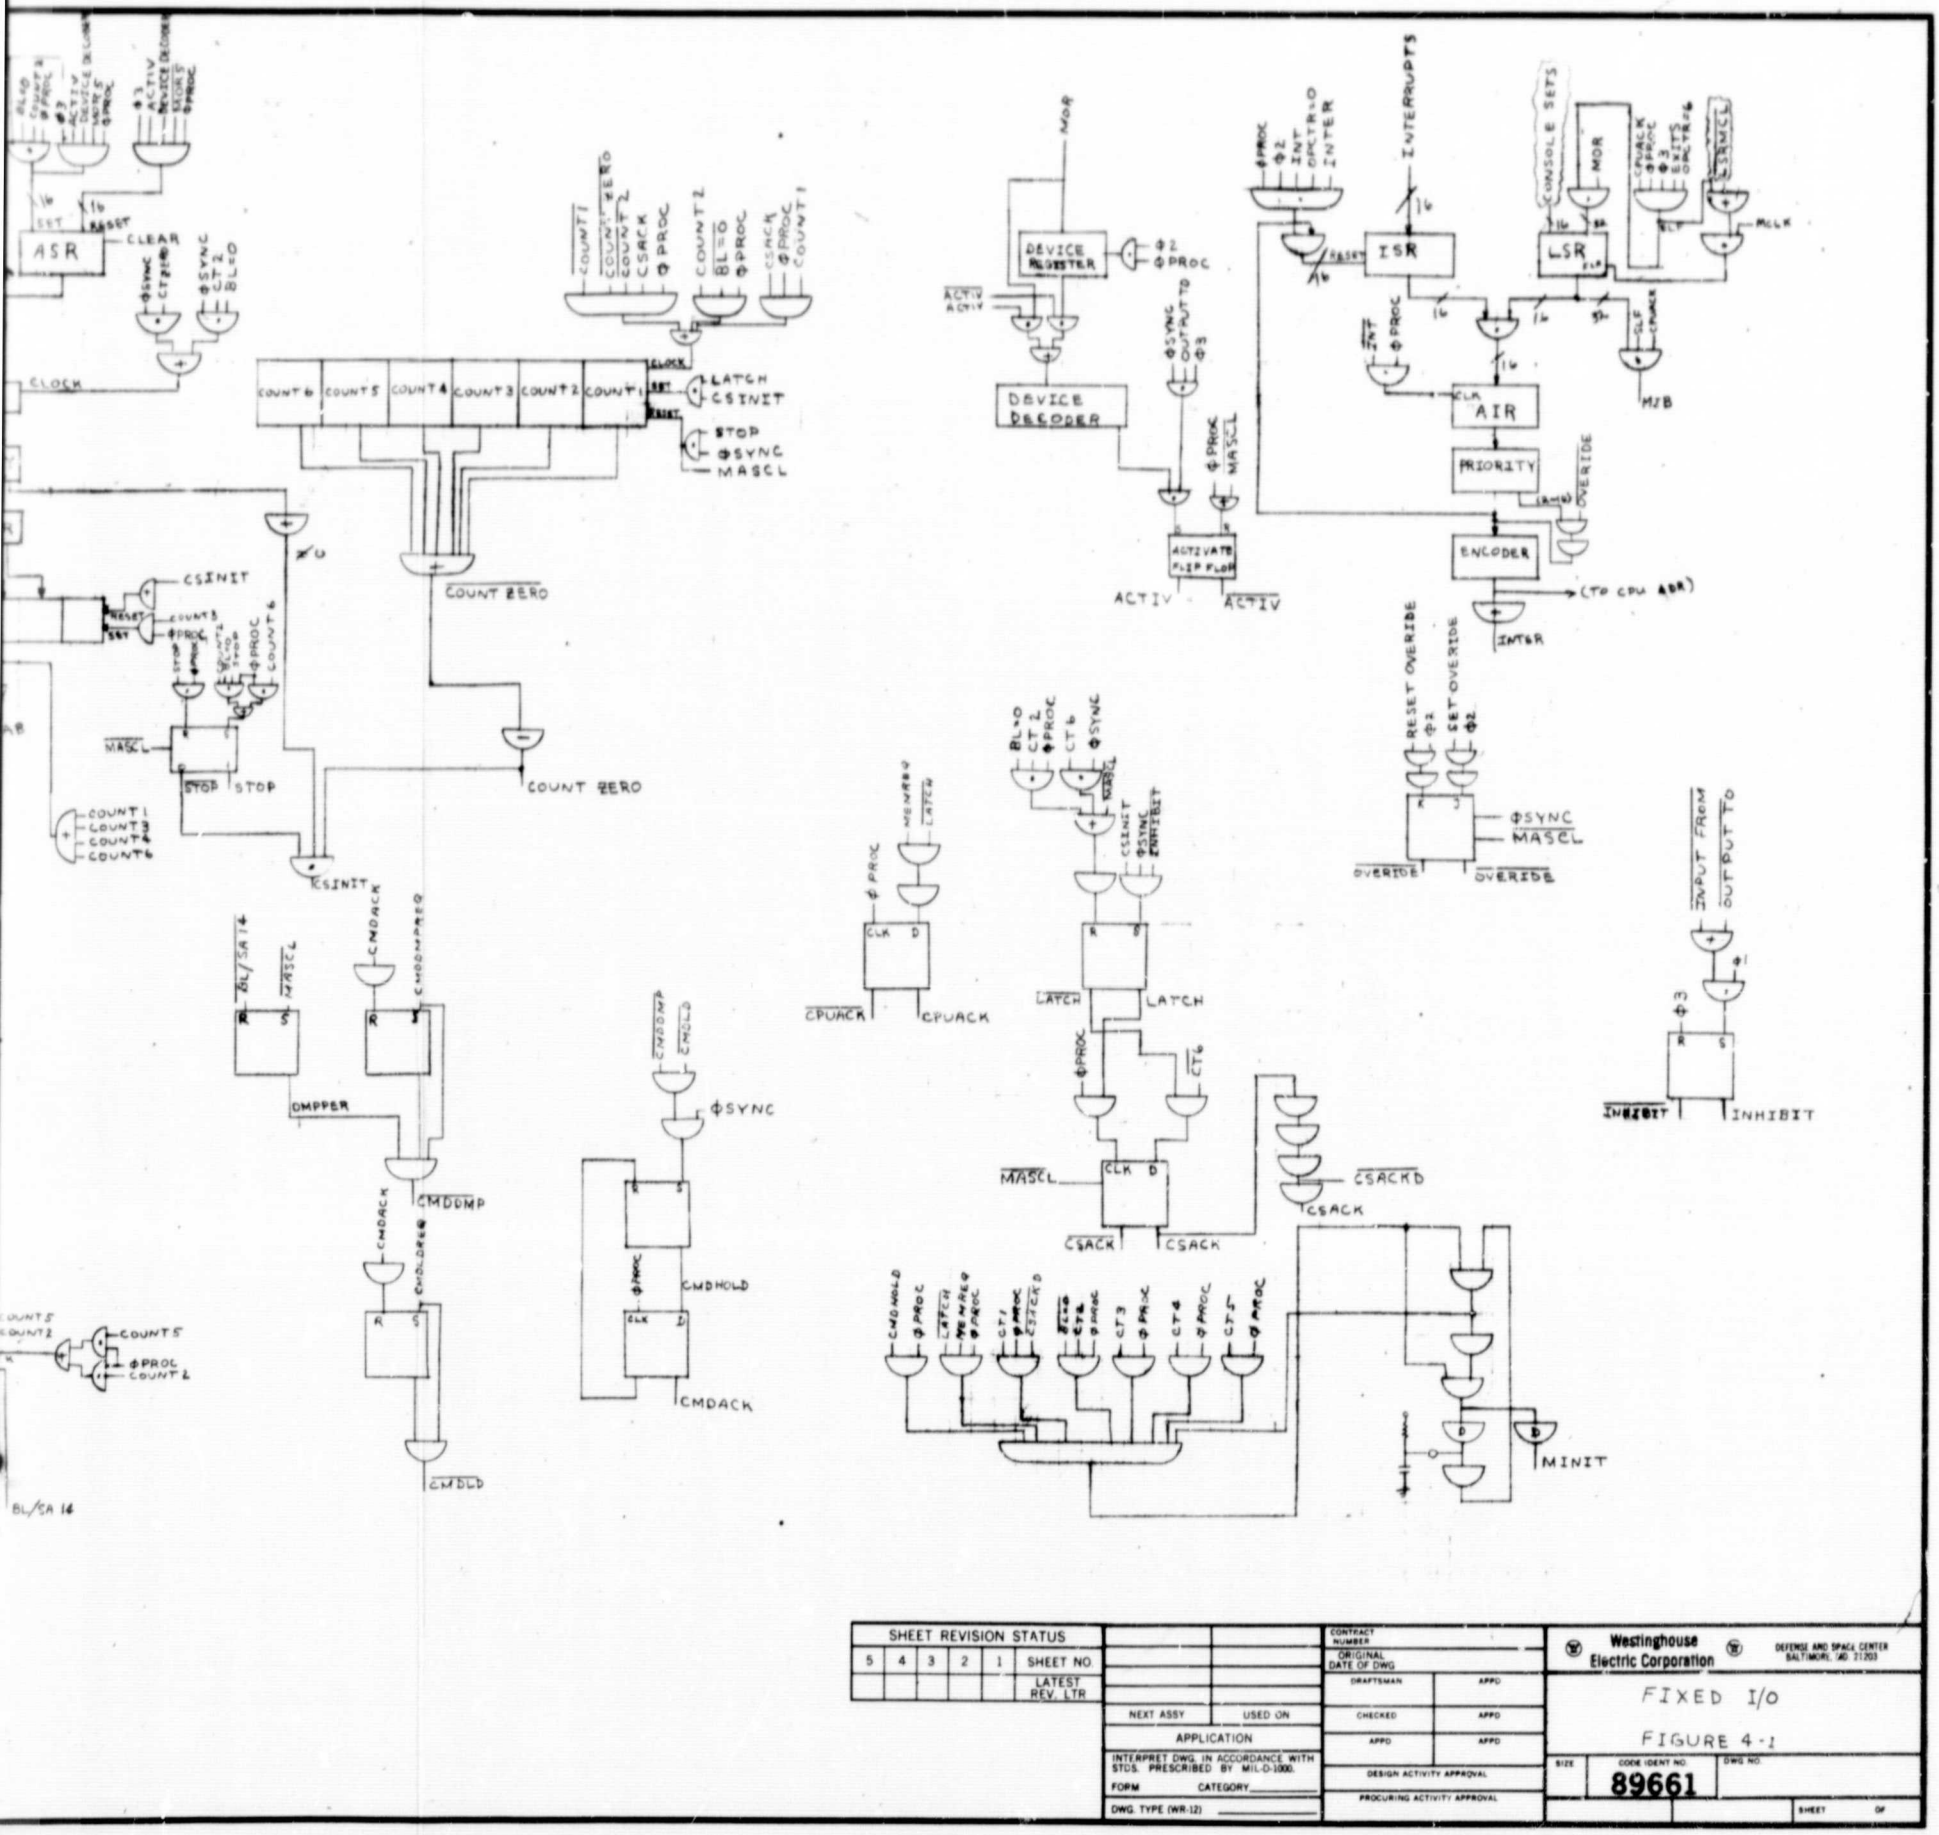

SHEET REVISION STATUS					
5	4	3	2	1	SHEET NO.
					LATEST REV. LTR

NEXT ASSY | USED ON

APPLICATION

INTERPRET DWG. IN ACCORDANCE WITH STDS. PRESCRIBED BY MIL-D-1000.

FORM | CATEGORY

DWG. TYPE (WR-12)

CONTRACT NUMBER

ORIGINAL DATE OF DWG

DRAFTSMAN | APPD | CHECKED | APPD | APPD | APPD

DESIGN ACTIVITY APPROVAL

PROCURING ACTIVITY APPROVAL

Westinghouse Electric Corporation — DEFENSE AND SPACE CENTER BALTIMORE, MD. 21203

FIXED I/O

FIGURE 4-1

SIZE | CODE IDENT NO. 89661 | DWG NO.

SHEET OF

Figure 4-1. Fixed I/O

FOLDOUT FRAME 2