

UNIVERSITY OF MANCHESTER  
DEPARTMENT OF COMPUTER SCIENCE

MU5 BASIC PROGRAMMING MANUAL

JULY 1975



## CONTENTS

CHAPTER 1	Introduction
CHAPTER 2	Operand Accessing
CHAPTER 3	The B-Arithmetic
CHAPTER 4	Accumulator Arithmetic
CHAPTER 5	Structure Accessing and Store to Store Orders
CHAPTER 6	Organisational Orders
CHAPTER 7	The Interrupt System
CHAPTER 8	The V Store
CHAPTER 9	The Vx-Store
CHAPTER 10	The Basic Programming Language - KPL





Chapter 11.1 Introduction

The organisation of the machine is reflected in its order code which is essentially of the form:-

F      N

where F defines the function and N the operand. There are four classes of orders:-

computational orders  
B-orders  
Structure accessing and store-to-store orders  
Organisational orders

In the computational orders which are distinguished by a 1 in digit 0, the instruction is divided thus:-

1	cr	f	N
2	4	9	

The cr bits define one of four types of arithmetic:-

signed fixed-point  
unsigned fixed-point  
decimal  
floating-point


In MU5 the signed fixed-point operations use a 32-bit register X, while the unsigned fixed-point, floating-point and decimal operations share a common 64-bit register A. However, the structure of the instruction code allows for four separate registers. The f bits define the operation to be performed and N defines the operand. Computational orders are of the single address type (e.g., A = operand, A + operand).

The B-orders operate on the modifier register B. They have the format:-




The functions provided correspond to those which operate on X but the division orders are not implemented in MU5.

In the structure addressing and store-to-store orders, it is convenient to think of the instruction as being divided in the same way:-



However, the two values of d give a total of 32 possible functions; some of these are used for manipulating registers in the secondary operand unit, which is closely associated with all store-to-store operations.

In the organisational orders, the instruction is divided thus:-



The cr bits are zero, and the 6 f' bits define both the organisational register and the operation to be performed. The organisational orders are mainly concerned with control transfers and the manipulation of organisational registers.

An operand is specified by N (or N'), and is independent of the function. An operand may be a literal, a 'named operand' (more simply, a 'name'), or a secondary operand; the various internal registers (X, B etc.) may also be addressed as operands.

## 1.2 Summary of the Order Code

This section summarizes the overall pattern of the order code. The detail is given in later sections as shown below. Some functions in MU5 differ from general form overleaf, which should be taken only as a statement of the general characteristics.

<u>References</u>	<u>Chapter</u>
Computational Orders	3, 4
B register	3
Accumulators	4
Structure Accessing and Store to Store Orders	5
Organisational Orders	6
Operand Accessing	2
Literals	2.3
Variables	2.4
Internal Register Operands	2.5
Stacked Operands	2.6
Privileged Operands	2.7
Secondary Operands	2.8
Descriptor Types 0 - 3	2.11 - 2.14
Internal Registers	
B, BCD	3
AEX	4
MS	6.2
NB, XNB, SF, (SN)	2.2
CO	6.3
D	2.2
XD, DOD, DT, XDT	5
EN	6.5

Computational and Store-to-Store Orders

F	STS	B	KS	AU	ADC	APL
0	XDC = DO =	=	=	ACD=	DUMMY	= (32)
1	KD = D =	= (-1)	DUMMY	DUMMY	AEK=	= (64)
2	STACK	D *=	*=	X*=	ADD*=	AEK*=
3	KD =>	D =>	=>	X=>	ACD=>	AEK=>
4	KDB = DB =	+	+	A+	DUMMY	+
5	XCHK	MDR	-	-	A-	DUMMY
6	SMOD	KGD	*	*	A*	DUMMY
7	EMOD	EMOD	/	/	DUMMY	DUMMY
8	SLGC	ELGC	≠	≠	A ≠	DUMMY
9	SMVB	BMVB	V	V	A V	DUMMY
10	DUMMY	EMVE	↑	↑ ARITH	A ↑ ECG	A ↑ CIRC
11	SMVF	SMVF	&	&	A &	DUMMY
12	TALU	DUMMY	⊖	⊖	A ⊖	ACDCERP
13	DUMMY	BSCN	COMP	CCMP	A CCMP	CCMP
14	SCIP	BCMP	CINC	AEX=CONVK	DUMMY	UNPACK
15	SUB1	SUB2	∅	∅	DUMMY	DUMMY

3	4	13	16
cr	f	k	n

- k = 0 - LITERAL      n is 6-bit signed integer
- k = 1 - IR            n defines internal register, P.T.O. ->
- k = 2 - V32            Operand is accessed directly at
- k = 3 - V64            (NB) + unsigned n; n is scaled for V32
- k = 4 - S[B]          Operand is accessed via a
- k = 5 - S[B]          descriptor at (NB) + n, using
- k = 6 - S[D]          B or O as an index
- k = 7 - K (Extended Operand) P.T.O. ->

Organisational Orders

2	16	11	16
0	F'	k	n

0	->	EKIT	DUMMY	DUMMY
4	JUMP	RETURN	DUMMY	DUMMY
8	XCO	KC1	KC2	KC3
12	XCA	KC5	KC6	STACK LINK
16	KS =	DL =	SFM	SET LINK
20	XNB =	SN =	XNB +	XNB =>
24	SF =	SF +	SF = NB +	SF =>
28	NB =	NB = SF +	NB +	NB =>
32	= 0	≠ 0	≥ 0	< 0
36	≤ 0	> 0	OVERFLOW	Bn
40	= 0	≠ 0	≥	< 0
44	≤ 0	> 0	OVERFLOW	Bn
48	0	Bn & X	Bn & X	X
52	Bn & X	Bn	Bn ≠ X	Bn V X
56	Bn & X	Bn ≡ X	Bn	Bn V X
60	X	Bn V X	Bn V X	1

-> Denote test result by T (= 0 for NO, = 1 for YES).  
The operand specifies the way in which BN is set.

BN = 0	BN & T	BN / MT	BN = T
BN & / T	DUMMY	BN ≠ T	BN V T
BN / & / T	BN ≡ T	BN /	BN / V T
BN = / T	BN V / T	BN / V / T	BN = 1

Internal Register Operands

The n bits define the internal register to be used.

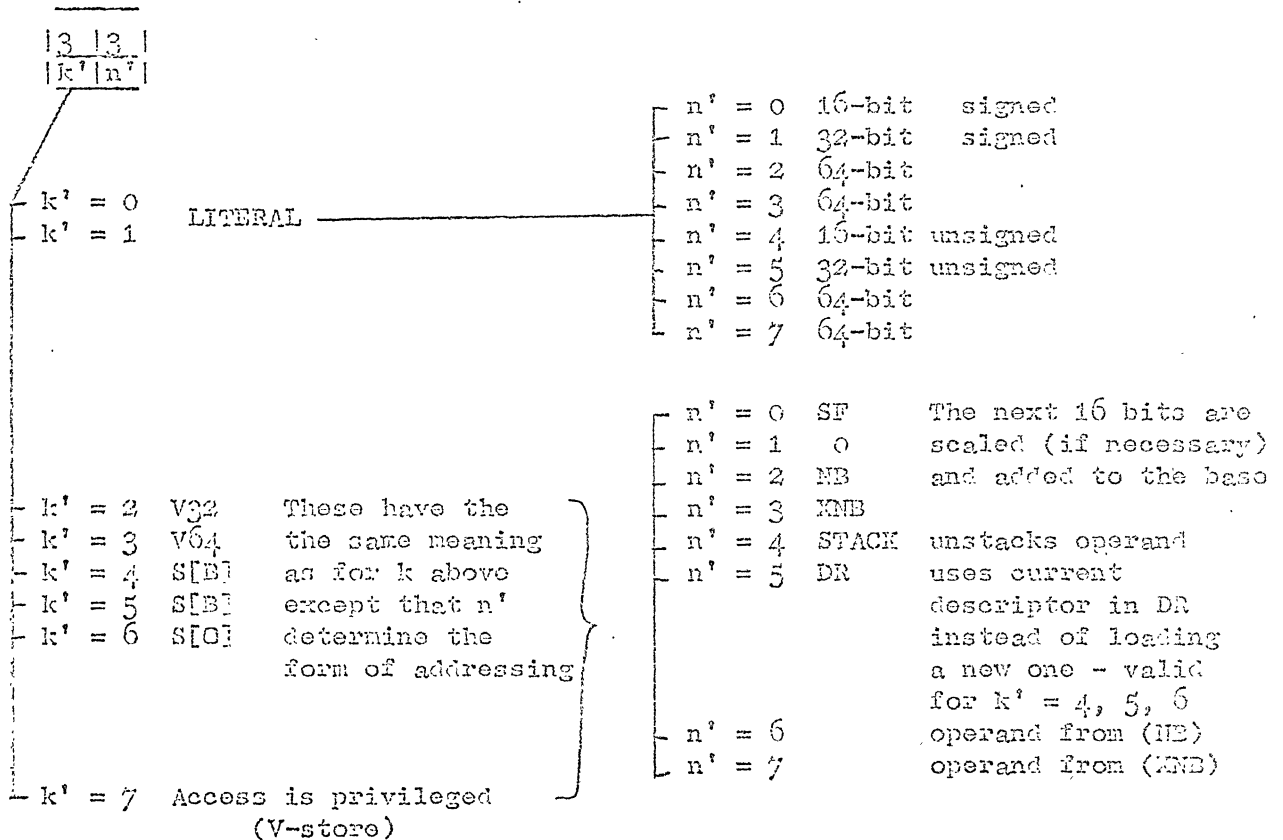
<-16-><-16-><-16-><-16->			
0	FS	NS	CC
1			XNB
2		SN	NB
3		SN	SF
4			BN
5			
6			
7			

16	D
17	XD
18	DT
19	XDT
20	DD
21	
22	
23	


32	BCD	B
33		BCD
34	Z	
35		
36		
37		
38		
39		

<-----64----->	
46	ARX
49	
50	
51	
52	
53	
54	
55	

Extended Operands, K








10	11	UNIT
0	0	PROP
0	1	SEOP
1	0	B
1	1	ACC

0	1	2	central register
0	1		B
1	0		XD
	1		D
0	0		ACC FIXED
	1		ACC LOGICAL
	0		ACC DECIMAL
	1		ACC FLOATING
0	0	0	ORGANISATIONAL

7	8	9	KIND	SOURCE
0	0	0	literal	n (signed)
	0	1	INTERNAL REGISTER	
1	0	0	Var 32	(n/2+NB)
	1	1	Var 64	(n+NB)
1	0	0	S[B]	(D+B) D=(n+NB)
	1	1	S[B]	(D+D) D=(n+NB)
	0	1	S[O]	(D) D=(n+NB)
1	1	1	K	Extended Operand

0	literal	n (signed)
1	K	Extended Operand



10	11	12	KIND	SOURCE
0	0	0	literal	I
	1	0	Var 32	(N/2+Base)
1	0	1	Var 64	(N+Base)
	0	0	S[B]	(D+B) D=(N+Base)
1	1	1	S[B]	(D+B) D=(N+Base)
	0	0	S[O]	(D) D=(N+Base)
1	1	1	V	V Store (N+Base)

13	14	15	BASE	ACTION
0	0	0	SF	16 bit Name
	1	0	O	16 bit Name
	1	1	NB	16 bit Name
1	0	1	XNB	16 bit Name
	0	0	SF	N=0 SF=SF-1
	1	1	O	N=0 D=D
1	0	1	NB	N=0
	1	1	XNB	N=0

13	14	15	LITERAL TYPE
0	0	0	16 bit signed
	1	1	32 bit signed
1	0	0	64 bit signed
	0	0	16 bit unsigned
1	1	1	32 bit unsigned
	1	0	64 bit unsigned

MU5 INSTRUCTION SET - OPERANDS





## Chapter 2 Operand Accessing

### 2.1 Introduction

The operands for all orders are transferred from source to destination via a highway which is 64 bits wide. For a fetch order, the operand part of the order defines how the highway is loaded and the function part defines the destination (and the operation to be performed at the destination). For a store order, the function part defines how the highway is loaded, and the operand part defines the destination.

The function part of an order is described in Chapters 3 - 6; this chapter describes the operand part. With a few exceptions, which will be mentioned when they arise, any function part may be combined with any operand part, so that the two parts may conveniently be described separately.

Operands may be of various sizes up to a maximum of 64 bits. If the operand is less than 64 bits long, then it is loaded on to (or taken from) the least significant end of the highway. On a fetch order the remaining bits of the highway are set to zero (except for literal operands - see Section 1.2.3). On a store order, the remaining bits are truncated; for secondary operands only, the truncated bits are checked for zeros.

In addition to various sizes of operand, there are various kinds of operand:-

- |           |   |
|-----------|---|
| literals  | A literal is specified directly as part of the order, e.g., 'X + 1' would add 1 to the (signed) fixed-point accumulator.      |
| variables | A variable is the value in a store location whose address is specified by a base register and the displacement from the base. |

internal registers      The value in most of the internal registers (B, NB etc.) can be specified as an operand, e.g., 'X = NB' loads the value in NB into the (signed) fixed-point accumulator.

stacked operands      Operands can be sent to and taken from a hardware implemented stack working on a last-in first-out basis, e.g., 'A\*STACK' multiplies the floating-point accumulator by the top operand on the stack, and removes the operand from the stack.

privileged operands      These can only be accessed in executive mode and are described in Chapter 8.


secondary operands      A special mechanism is provided for accessing secondary operands, i.e., operands contained in some data structure. The operand part of the order specifies a data descriptor and a modifier. The data descriptor is a 64 bit animal specified as a variable or stacked operand and is combined with the modifier in D to produce the size and address of the secondary operand. For example, consider the orders 'B = 3, D = FRED, A + D[B] where FRED is descriptor at address (NB + 5). The action would be to load 3 into the modifier register B, then send the descriptor at address NB + 5 to DR, modify by the value in B (i.e., 3) to give the size and address of the secondary operand, and finally add this operand to the floating point accumulator.

## 2.2 Internal Registers Relevant to Operand Accessing

Section 1.2 contains a complete list of the internal registers with references to their descriptions. In this section, only those registers relevant to operand accessing are described.

### The Name Segment Number SN

The name segment number SN is 16 bits long. The two most significant bits are permanently zero, and the remaining 14 bits define the segment currently being used for the names in a program. Any segment (0, 1, 2, . . .  $2^{14} - 1$ ) may be used for this purpose; it is conventional to use segment 0 whenever possible. The value contained in SN may only be altered by calling an executive procedure.




### The Name Base Register NB

The name base register NB is 16 bits long. The most significant 15 bits hold the address of any 64-bit word in the name segment SN, and the least significant bit is permanently zero. When NB is the base register for an operand access, then it is added to the displacement (the name) to give the address of the operand within the segment SN. If the addition overflows out of the segment, there will be an interrupt.

NB is designed to be the base register for local names in a procedure; its value will usually only be changed on entry and exit.

Orders which alter NB are described in Section 6.2 and 6.3.



### The Stack Front Register SF

The format of the stack front register SF is identical with that of NB. SF can be used as a base register in the same way as NB. However, the space in front of SF (i.e., at addresses  $> SF$ ) must not be accessed in this way; interrupt routines use the area in front of SF as working space, and so these locations are liable to change at any time.

The stack is designed both to provide temporary working space within a procedure (e.g., for evaluating arithmetic expressions), and also the space required for procedure calls (see Section 6.4). Certain orders, e.g., STACK B, cause an operand to be stacked - SF is advanced by 2 (32-bit) words and the operand is stored at the 64-bit word specified by the new value of SF. These operands may be unstacked by specifying the STACK as the operand part of an order, e.g., A = STACK - the 64-bit word at SF is loaded on to the highway, and SF is decreased by 2. Note that all unstacked operands are assumed to be 64 bits long.

Orders which alter SF are described in Sections 3.2, 4.4, 6.2 and 6.3.

64-bit word address   0
15
1

#### The Extra Name Base Register XNB

The extra name base register, XNB, is 32 bits long. Bits 2 - 30 hold the address of a 64-bit word anywhere in the virtual store; bits 0, 1 and 31 are permanently zero. XNB is used as a base register in the same way as NB and SF, except that the operand is in the segment defined by the top half of XNB (instead of SN). Note that the addition of the name must not overflow out of this segment, or there will be an interrupt.

XNB is designed to be a base register for non-local names used in a procedure, and will often change its value in a procedure. In many programs, the top half of XNB will be zero (like SN).

Orders which alter XNB are described in Section 6.2.

00	segment	64-bit word address   0
2	14	15
		1

#### The Data Descriptor Register D

The data descriptor register, D, is 64 bits long, and is used to hold the descriptors required for accessing data structures. The operand part of an order which accesses a data structure specifies a descriptor and a modifier. The descriptor is loaded into D and then combined with the modifier to define the size and address of the particular structure element required.

Details of the descriptor types and the mechanism for accessing secondary operands are given in Sections 2.8 and 2.11 - 2.14.

D also plays a major part in the operation of the store-to-store orders. D manipulation orders are described, with the store-to-store orders, in Chapter 5.

### 2.3 Literal Operands

A literal operand appears directly as part of the order; if a literal is specified as the operand part of a store order, there will be an interrupt.

There are several alternatives:-

- (a) 6-bit signed
- (b) 16-bit unsigned
- (c) 16-bit signed
- (d) 32-bit unsigned
- (e) 32-bit signed
- (f) 64-bit

The literals are copied to the least significant end of the highway. The remaining bits of the highway are set to zeros for unsigned literals and to copies of the sign bit for signed literals.

The precise format of orders containing literals is unexpected.

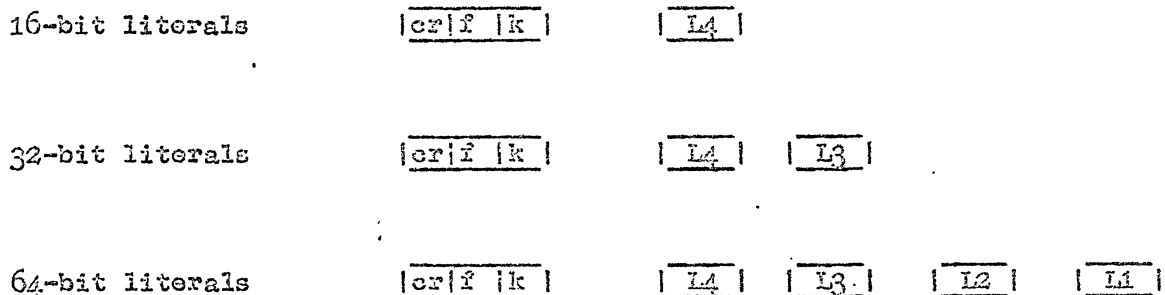
Let L1 denote 16 bits to be loaded to highway bits 0 - 15.

Let L2 denote 16 bits to be loaded to highway bits 16 - 31.

Let L3 denote 16 bits to be loaded to highway bits 32 - 47.


Let L4 denote 16 bits to be loaded to highway bits 48 - 63.

Then the orders appear as follows:-



## 2.4 Variable Operands

There are two kinds of variable, V32 and V64, of sizes 32 bits and 64 bits respectively. The operand part of the order specifies the kind of order and also defines its name and base. NB, XNB, SF or 0 may be used as the base (it is convenient to consider 0 to be a base register which always contains 0). The name is the distance of the variable from the base counting in units equal to the variable size. Some examples are shown below - the diagram represents a section of the virtual store marked out in 32-bit words:-



V64 names are in the range  $-2 \times 15 \leq \text{name} < 2 \times 15$

V32 names are in the range  $0 \leq \text{name} < 2 \times 16$

To calculate the address of the operand, the name is scaled (if necessary) and added to the base. If this addition overflows out of the base segment, there will be an interrupt. If NB, SF or 0 is used as the base, then the variable is taken from the name segment (SN); if XNB is used as the base, then the most significant half of XNB defines the segment. In short instructions, the 6-bit displacement, n, is always unsigned, i.e.,  $0 \leq n < 64$ .

N.B. If XNB points to a segment which is not the name segment operands relative to XNB may not be used with the following functions, XNB =>, NB =>, SF =>, SETLINK.

Note that the organisational commands, input/output and CTL use words, in the name segment. Thus when running under the operating system 32-bit words:-

0 - 15 should not be used when writing in KPL

0 - 95 Should not be used when using the Autocode machine.

## 2.5 Internal Register Operands

Any internal register may be specified as the operand for a fetch order; a store order may write to most internal registers, except those within the primary operand unit, i.e., MS, NB, CO, XNB, SN, SF, BN.

A table listing all the registers or combinations of registers that can be accessed in this way is given in section 1.2; note that only complete lines may be accessed, for example, SF cannot be read by itself but only in combination with SN.

Internal register operands may only be used with computational and store to store orders, not with organisational orders.

The Internal Register Z is a dummy operand which is written to as a means of suppressing overlap until the order is complete.

2.6 Stacked Operands

When the operand part of the order specifies STACK, the 64-bit word at SF is loaded on to the highway and then SF is decreased by 2. Note that all operands coming from STACK are 64 bits long; this does not mean that only 64-bit operands may be sent to the stack - shorter operands will be extended by zeros on the way.

[A store order specifying STACK will store the operand at SF and then decrease SF by 2. This is not a sensible order, but is allowed].



## 2.7 Privileged Operands

Privileged operands are used by executive to hold system control information. They can only be accessed in executive mode and are of no interest to the ordinary programmer.

Access can be made in two ways. In the first case, a base register and a name are specified as for a variable operand; the size is always 64 bits, and the address is calculated exactly as for a V64 variable. However, access is made not to the virtual store of the program but to the local V-store (i.e., the address is interpreted as a local V-store address). In the second case, the operand part of the order is STACK; the action is exactly the same as for other stacked operands, but SF is now interpreted as a 64-bit word address in the local V-store. (The local V-store is described in Chapter 8).

## 2.8 Secondary Operands

For a secondary operand, the operand part of the order specifies a 64-bit descriptor and a modifier. Normally, the descriptor specifies the type and origin (i.e., starting address) of the data structure containing the secondary operand, and the modifier defines which particular operand is required. For example, if A is a descriptor specifying a vector of 32-bit elements, then the orders 'B = 25; X = A[B]' would load the 25th element (counting from zero) of A into the fixed-point accumulator.

Descriptors can define vectors or strings of elements of various sizes; miscellaneous special types are also provided. The different types of descriptor are defined in sections 2.11 - 2.14.

A descriptor may be specified in the same way as a variable or stacked operand; it is always 64 bits long, and is loaded into the D register. Alternatively, the operand part of the order may specify that the descriptor is already in D; this avoids unnecessary loading of D if the same descriptor is used for consecutive secondary operands.

The modifier used is normally B or O (i.e., no modifier). However, there are special functions (see Chapter 5) which allow any operand to be used as a modifier and also cause a special type of modification. All modifiers are interpreted as signed 32-bit fixed-point integers.

When access is made via certain types of descriptor (e.g., vectors) it is possible to check automatically that the modifier (if any) lies in the range  $0 \leq \text{modifier} < \text{bound}$ . The bound is held in bits 8-31 of the descriptor.

N.B. A secondary operand may not be used in conjunction with the following functions, D =>, XD =>, XNB =>, NB =>, SF =>, SETLINK.

**2.9**    Length of the Orders

An order may be 16, 32, 48 or 80 bits long. It will be 16 bits long when:-

- (a) Operand is 6-bit literal or internal register
- (b) Operand is variable or secondary; base register is NB and  $0 \leq \text{name} \leq 63$ ; function part is computational or store to store
- (c) Operand is variable, privileged or secondary from STACK.

An order will be 48 bits long if the operand is a 32-bit literal and will be 80 bits long if the operand is a 64-bit literal. In all other cases, an order will be 32 bits long.

2.10 Type 0 - Vector Descriptors

Type 0 descriptors are used for vectors of elements of size 1, 4, 8, 16, 32 or 64 bits. The descriptor defines the origin of the vector, the element size, and an upper bound for the modifier (i.e., the number of elements in the vector). The format is:-

T	SIZE	US	BC	BOUND	ORIGIN IN BYTES
2	3	1	1	1	24
					32

T = 0 Defines type 0.

SIZE Defines the element size as 1, 4, 8, 16, 32 or 64 bits.  
(Coded as follows:- 000 = 1 bit; 010 = 4 bits;  
011 = 8 bits; 100 = 16 bits; 101 = 32 bits;  
110 = 64 bits).

US If US = 0, then the modifier is scaled before being added to the origin - for 1-bit elements, the modifier is shifted down 3 bits, for 4-bit elements down 1 bit, for 8-bit elements none, for 16-bit elements up 1 bit, for 32-bit elements up two bits, for 64-bit elements up 3 bits.

If US = 1, then modifier is not scaled.

BC If BC = 1, then there is no bound check.

BOUND An upper bound for the modifier. If the bound check bit BCHI in DCD (see Chapter 5) is set to 0 and BC = 0, then the modifier (if any) must lie in the range  $0 \leq \text{modifier} < \text{BOUND}$ , otherwise there will be an interrupt. (See Section 5.2).

ORIGIN The origin defines the base address of the vector; it is always a 32-bit byte address. For 16-bit vectors, the least significant bit of the modified address is ignored, so that all elements start at a 16-bit word boundary. For 32 and 64-bit vectors, the two least significant bits are ignored, so elements start on a 32-bit word boundary. Note that vectors of 1-bit and 4-bit elements must start on a byte boundary.

Action: When an access is made, the modifier (if any) is scaled (according to SIZE and US) and added to the origin to give the address of the required element. Provided there is no bound check fail, the element is accessed. On a fetch order, it is loaded on to the highway (operands < 64 bits long are loaded at the least significant end and the remaining bits are zeroed). On a store order, the highway is stored at the element; there will be an interrupt if any non-zero bit is truncated. (See Section 5.2).

2.11 Type 1 - String Descriptors

Type 1 descriptors are used for strings of 8 bit elements. The descriptor defines the origin and length of the string. The format is:-

T	SIZE				LENGTH		ORIGIN IN BYTES	
2	3	1	1	1	24		32	

T = 1 Defines type 1.

SIZE Must define the element size as 8 bits (011), else an interrupt will occur.

LENGTH LENGTH defines the number of elements in the addressed string.

ORIGIN Defines a base address, as in type 0.

Action: The modifier (if any) is added to the origin to give the address of the start of the string. LENGTH defines the length of the string, i.e., the number of elements in the string. There is no bound checking.


The final operand is a string of 8-bit elements. In store to store orders the whole string will be used as the operand (see Chapter 5). In computational orders, the operand can be at most 64 bits long; if the string is less than 64 bits, then it is zero filled for fetch, truncated with zero-checking for store; if the string is longer than 64 bits, just the first 64 bits of the string are loaded on to or stored from the highway.

**2.12 Type 2 - Descriptor Descriptors**

Type 2 descriptors are identical with type 0 descriptors (except that T = 2 instead of 0).

<u>T</u>	<u>SIZE</u>	<u>US</u>	<u>BC</u>	<u>BOUND</u>	<u>ORIGIN IN BYTES</u>
2	3	1	1	1	24
					32

It is conventional to use type 2 descriptors to address vectors containing descriptors; type 0 is used for vectors containing data.


2.13 Type 3 - Miscellaneous DescriptorsType 3.0 Real Address

T, T' = 3, 0 Define type 3.0.

BOUND Upper bound for modifier as in Type 0.

ORIGIN Contains the real store address (the physical address, not the virtual store address) of a 64-bit operand. The three least significant bits are ignored.

Action: The operand is accessed in the same way as a type 0 64-bit element. The modifier is always scaled, and bound checked if bit BCHI in DOD is set to 0. Type 3.0 descriptors may only be used in Executive mode.

Type 3.1 - Read/Store Direct

T, T' = 3, 1 define type 3.1.

BOUND Upper bound for modifier as in type 0.

ORIGIN Defines a 64-bit word address. The three least significant bits are ignored.

Action: Access is made in exactly the same way as for a type 0 64-bit element (assuming US = BC = 0, so that the modifier is scaled, and bound-checked if bit BCHI in DOD is set to 0. Note that the word lies on a 64-bit word boundary.

The accessing mechanism for this descriptor bypasses all operand buffers, and always accesses the real store corresponding to the defined virtual address. This type of access is needed in some executive procedures.

Type 3.2 Read and Mark

T	T'	BOUND	ORIGIN IN BYTES
2	6	24	32

T, T' = 3, 2 define type 3.2

BOUND Upper bound for modifier as in type 0

ORIGIN Defines a 64-bit word address; the three least significant bits are ignored.

Action: Access is made in exactly the same way as for type 3.1 descriptors, bypassing the operand buffers. In addition, for a fetch order, the value of the 64-bit word in the store is finally set to zero.

Type 3.3 Indirect

T	T'	X	ORIGIN IN BYTES
2	6	24	32

T, T' = 3, 3 define type 3.3

X Unused

ORIGIN Defines a 32-bit word address; the two least significant bits are ignored.

Action: The 64-bit element at the origin address is loaded into D and then interpreted according to its type. The new descriptor may be indirect, in which case the whole process is repeated. If a modifier is specified, the modification takes place at the final (not indirect) stage.



Types 3.4 - 3.31 Procedure Call

T	T'	X	ORIGIN IN BYTES
2	6	24	32

T, T' = 3, 4 - 31 Define the procedure call type. (32 - 63 are illegal).

X Not used.

ORIGIN Contains the address of the 'procedure call vector'. The two least significant bits of the origin field are ignored.

Action: When an attempt is made to access an operand, the hardware forces a procedure call to the address held in the first 32-bit word of the vector, with the return link pointing to the instruction attempting to make the access. The origin is not modified (even if a modifier is specified by the operand part of the order). The type bits in D are reset to a type 0 32-bit vector with US = BC = 0. The X and ORIGIN fields are unaltered.

One example of the use of the procedure call descriptor is an implementation of an Algol formal parameter called by name. If the corresponding actual parameter is a simple variable, then the parameter descriptor can be a normal type 0 descriptor. But if the actual parameter is an expression, then the descriptor will be a procedure call to code which evaluates the expression. The value will be stored in some suitable store location and D replaced by a type 0 descriptor pointing to it; finally, the 'D set' bit in the stored link (cf., Section 6.2) is set, and an EXIT obeyed. The order causing the procedure call will be re-obeyed - the 'D set' bit prevents reloading of D and defines that the current value of D describes the required operand. (The 'D set' bit is automatically reset to 0).

Note that a procedure call descriptor may be modified; the modification will take place when the order is re-obeyed after exit from the procedure.



Chapter 3 The B-Arithmetic

## 3.1

There is a separate 32-bit B-arithmetic unit which operates on the modifier register B. Although B is used mainly for modification it is also used for some of the simpler integer arithmetic, for example,  $i = i + 1$ .

The bits of B are numbered from 0 on the left hand (most significant end).

d0	d1	d2																													d30	d31
----	----	----	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	-----	-----

The operand connection to the B-arithmetic unit is from the least significant 32 bits of the highway (bits 32 - 63).

The B-arithmetic unit performs signed 2's complement arithmetic. Thus B may take values in the range  $-2^{31}$  to  $2^{31} - 1$ . If, after any arithmetic operation, the true result is outside this range, the overflow bit is set. The overflow bit and a bit which is used to inhibit the interrupt resulting from overflow are digits 5 and 0 of BOD (see 1.1.2). Thus digit 5 of BOD is set to a one if overflow occurs and the interrupt will be inhibited if digit 0 is also set to a one. All other digits of BOD are not significant.

### 3.2 The B-Instructions

The order code provides for 16 B-functions. Only 14 of these are implemented on MU5 and the rest are dummy instructions. The instructions are:-

#### LOAD (=)

Load B from the least significant 32 bits of the highway.

#### LOAD & DECREMENT (=')

Load B from the least significant 32 bits of the highway then subtract 1. If an overflow occurs digit 5 of BOD is set.

#### STACK & LOAD (\*=)

The stack front register (SF) is first advanced by 2. The contents of B are placed on the highway as for a store order (see below). This is then sent to the 64-bit word whose address specified by the new value of SF. Finally, the operand is loaded into B as in the load order (see above).

#### STORE (=>)

The content of B is placed on the least significant 32 bits (bits 32 - 63) of the highway and zeros are placed on the most significant 32 bits (0 - 31). The operand specifies the destination of this information.

#### ADD (+)

The operand is added to B, leaving the result in B. If an overflow occurs, then digit 5 of BOD is set.

#### SUBTRACT (-)

The operand is subtracted from B leaving the result in B. If an overflow occurs, then digit 5 of BOD is set.

#### MULTIPLY (\*)

B is multiplied by the operand to produce a 32-bit result in B which is the least significant 32 bits of the true 64 bit signed answer. If the true product has more than 32 significant bits, then B contains the least significant 32 bits of the true answer and digit 5 of BOD is set.

## DIVIDE (/)

A dummy instruction.

NON-EQUIVALENCE ( $\neq$ )

B and the operand are non-equivalenced to produce a result in B.

## OR (V)

B and the operand are ored to produce a result in B.

## AND (&amp;)

B and the operand are anded to produce a result in B.

## SHIFT (A)

B will be shifted arithmetically (left) by the number of places specified by the signed integer in digits 57-63 of the operand. If overflow occurs digit 5 of BOD is set.

## COMPARE (CCMP)

The operand is subtracted from B. Bits T1 and T2 of the test register are then set from the result of the subtraction (see section 6.4). A true result is always generated and no overflow may occur. The overflow bit in BOD is copied to bit T0 of the test register. The contents of B are not altered.

## REVERSE SUBTRACT (e)

B is subtracted from the operand leaving the result in B. If an overflow occurs, then digit 5 of BOD is set.

## COMPARE &amp; INCREMENT (CINC)

A compare operation is performed (see above) then B is incremented by 1. If B overflows as a result of being incremented then digit 5 of BOD is set after the compare operation has been completed.

## REVERSE DIVIDE (ø)

A dummy instruction.



Chapter 4 Accumulator Arithmetic4.1 The Accumulator and its Associated Registers

The function code contains a set of 16 functions for each of the following kinds of arithmetic:-

fixed point signed  
fixed point unsigned  
floating point  
decimal

In MU5 there are two associated registers:-

X which is used by the signed fixed point orders

and

A which is used by the unsigned fixed point, floating point and decimal orders.

Each accumulator register is conceptually 64 bits long but digits 0 - 31 of X will not exist on MU5. There are two other visible 64-bit registers in the arithmetic unit, namely AOD and AEX. The bits of AOD are concerned mainly with interrupts whereas AEX (the accumulator extension register) serves to hold the least significant part of double length results. Because the accumulator 'A' is shared, the load and store functions would be the same in the fixed point unsigned, decimal and floating point instruction sets. Therefore the load and store functions in the fixed point unsigned set are made to operate on AOD and those in the decimal set on AEX.

It is convenient to consider the operand for each function to be the 64-bit accumulator input buffer AIB. Thus the operation of the accumulator functions will be described by reference to the registers:-

A, X, AOD, AEX, AIB

4.2 Allocation of Digits in AGDdigit

51	Operand size (O/I meaning 32/64 bits)
52	Inhibit floating point overflow interrupt
53	Inhibit floating point underflow interrupt
54	Inhibit fixed point overflow interrupt
55	Inhibit decimal overflow interrupt
56	Inhibit zero divide interrupt
57	Floating point overflow indicator
58	Floating point underflow indicator
59	Fixed point overflow indicator
60	Decimal overflow indicator
61	Zero divide indicator
62	Inhibit rounding
63	Double length <u>+</u>




4.3 Formats for Arithmetic Data


The formats marked with an asterisk are software concepts only and so have no significance in the hardware.

## (a) Fixed-point signed

Data is signed binary, held in 2's complement form. For multiplication and division, the binary point is at the least significant end, i.e., data is interpreted as an integer.




$$\text{range } -2^{31} \leq x < 2^{31}$$




$$\text{range } -2^{63} \leq x < 2^{63}$$

## (b) Fixed-point unsigned

Data is unsigned binary. For multiplication and division, the binary point is at the least significant end, i.e., the data is interpreted as an integer.




$$\text{range } 0 \leq x < 2^{32}$$




$$\text{range } 0 \leq x < 2^{64}$$

## (c) Decimal

Data is stored in sign-modulus form. The modulus consists of 7 or 15 decimal digits occupying 4 bits each, and the sign occupies 4 bits at the least significant end. Each decimal digit is coded in binary ( $0 \equiv 0000$ ,  $1 \equiv 0001$ , . . . ,  $9 \equiv 1001$ ); the sign code is 1101 means -ve, all other combinations mean +ve (1111 is preferred). For multiplication and division, the decimal point is at the least significant end, i.e., data is interpreted as an integer.



$$\text{range } -10^7 < x < 10^7$$



$$\text{range } -10^{15} < x < 10^{15}$$



(d) Floating-point

Data is stored as 2's complement mantissa *m* with an 11-bit exponent *e* stored at the most significant end. The most significant bit of *m* gives the sign of *m* and the interpretation of *m* assumes a binary point after the sign digit. The exponent has the base 16, and is interpreted as an unsigned 11-bit integer minus 1024, i.e.,

```

00000000000 -> -1024
00000000001 -> -1023
.
.
10000000000 -> 0
.
.
11111111111 -> 1023
    
```

This code has been chosen so that floating-point zero has all bits = 0.



4.4 The Signed Fixed Point Accumulator Orders

The arithmetic functions in this set assume X and the operand to be signed integers.

LOAD (=) Copy digits 32 - 63 from AIB to X.

LOAD DOUBLE (=') Dummy instruction.

STACK AND LOAD (\*=) Stack X in digits 32 - 63 of the next free 64-bit word on the stack making digits 0 - 31 in this word zero. Then operate as for LOAD.

STORE (=>) Copy X to digits 32 - 63 of the Highway and zeros to digits 0 - 31 of the Highway.

ADD (+) Digits 32 - 63 of AIB are added to X and the result is returned to X. If the addition overflows, digit 59 of ADD is set. In this case the result in X will be the least significant 32 bits of a 32-bit answer.

SUBTRACT (-) Digits 32 - 63 of AIB are subtracted from X and the result is returned to X. If the result overflows, digit 59 of ADD is set.

MULT (\*) X is multiplied by digits 32 - 63 of AIB to form a signed single length result in X. If the result overflows, digit 59 of ADD is set, and the result is the least significant 32 bits of the 64-bit answer.

DIVIDE (/) X is divided by digits 32 - 63 of AIB, to form a quotient in X, which will be rounded down. If the divisor is zero, then digit 61 of ADD is set, and X will be unaltered.

NON EQUIVALENCE ( $\neq$ ) The logical non-equivalence of digits 32 - 63 of AIB with X replaces X.

OR (V) The logical 'or' of digits 32 - 63 of AIB with X replace X.

- SHIFT (A) X will be shifted arithmetically (left) by the number of places specified by the signed integer in digits 58 - 63 of AIB. Digit 59 of AOD will be set if the result overflows.
- AND (&) The logical 'and' of digits 32 - 63 of AIB with X replaces X.
- REVERSE SUBTRACT (⊖) X is subtracted from digits 32 - 63 of AIB and the result is stored in X. If overflow occurs digit 59 of AOD is set.
- COMPARE (CMP) The operand in digits 32 - 63 of AIB is subtracted from X. Both are treated as signed integers. Bits  $T_1$  and  $T_2$  of the test register are set from the result of the subtraction. Note that a true result is generated and no overflow may occur. Bit 59 V bit 61 of AOD is copied to bit  $T_0$  in the test register. The content of X is not altered.
- CONVERT (CONV) The only conversion function implemented in the 'X' set is the conversion from integer to floating. The standardised floating result is left in AEX.
- REVERSE DIVIDE (⊘) Except that digits 32 - 63 of AIB are divided by X this function operates as for DIVIDE.

4.5 The Unsigned Fixed Point Accumulator Orders

The arithmetic functions in the set assume the least significant 32 bits of A and the operand in digits 32 - 63 of AIB to be 32-bit unsigned integers. They return a 64-bit signed result to A. If the most significant 32 bits of A or AIB are initially non-zero they are set to zero prior to arithmetic.

LOAD (=) Copy digits 51 - 63 from AIB to AOD.  
 Note: Floating point LOAD DOUBLE is the correct order to use in conjunction with unsigned arithmetic.

LOAD DOUBLE (=') Dummy instruction.

STACK AND LOAD (\*=) AOD is stacked and loaded.

STORE (=>) Copy digits 51 - 63 of AOD to the Highway setting the other digits of the Highway to zero.

ADD (+) Digits 32 - 63 of AIB are added to digits 32 - 63 of A and the result is stored in digits 0 to 63 of A.

SUBTRACT (-) Digits 32 - 63 of AIB are subtracted from digits 32 - 63 of A and the result is stored in digits 0 - 63 of A.

MULTIPLY (\*) Digits 32 - 63 of A are multiplied by digits 32 - 63 of AIB to form a 64-bit product which is stored in A.

DIVIDE (/) Dummy instruction.

NON-EQUIV (≠) Digits 32 - 63 of A are non-equivalenced with digits 32 - 63 of AIB and the result is stored in digits 32 - 63 of A. Digits 0 - 31 of A are set to zero.

OR (V) Digits 32 - 63 of A are ored with digits 32 - 63 of AIB and the result is stored in digits 32 - 63 of A. Digits 0 - 31 of A are set to zero.

- SHIFT (A) A is shifted logically (left) by the number of places specified by the signed integer in digits 57 - 63 of AIB. This order operates on all 64 bits of A.
- AND (&) Digits 32 - 63 of A are anded with digits 32 - 63 of AIB and the result is stored in digits 32 - 63 of A. Digits 0 - 31 of A are set to zero.
- REVERSE SUBTRACT (e) Digits 32 - 63 of A are subtracted from digits 32 - 63 of AIB and the result is stored in digits 0 - 63 of A.
- COMPARE (CCMP) As for CCMP in the signed fixed point set, except that comparison applies to digits 32 - 63 of A and is on an unsigned basis.  $T_0$  of the test register is set to zero.
- REVERSE DIVIDE (d) Dummy instruction.

4.6 The Decimal Mode Accumulator Orders

LOAD (=)	Dummy instruction.
LOAD DOUBLE (=')	Load AEX from AIB.
STACK AND LOAD (*=)	Stack and load AEX.
STORE (=>)	Store AEX.
ADD (+)	} Dummy instructions.
SUBTRACT (-)	
MULTIPLY (*)	
DIVIDE (/)	
NON-EQUIVALENCE ( <u>≠</u> )	
OR (OR)	
SHIFT (A)	Digits 59 - 63 of AIB are interpreted as a signed binary integer which specifies the number of decimal places by which A is to be shifted (left). The shift is logical over digits 0 to 59. Digits 60 - 63 are unaltered. If a left shift overflows digit 60 of AOD is set.
AND (&)	Dummy instruction.
COMPARE AOD	Digits 51 - 63 of AIB are anded with digits 51 - 63 of AOD. The overflow digit of the test register will be set to 0/1 depending upon the result being non-zero/zero.



COMPARE (CCMP) A is interpreted as a decimal number according to the formats on page 4.3.2. Bit  $T_2$  of the test register is set as the sign (bits 60 - 63) of A. The logical & of A and AIB is formed, and bit  $T_1$  of the test register is set according as the result is = or  $\neq$  to zero. Bit  $T_0$  of the test register is set to bit 60 of AOD (decimal overflow).

UNPACK This instruction sets AEX (32 - 59) = AIB (32 - 59) and AEX (60 - 63) = AIB (60 - 63) V A (0 - 3). AEX (0 - 31) are unaltered. It also shifts digits 0 - 59 of A four places left. Digits 56 - 59 of A are set zero and digits 60 - 63 unaltered.

REVERSE DIVIDE ( $\emptyset$ ) Dummy instruction.

4.7 The ACC Floating-Point Accumulator Orders

For some of the floating-point arithmetic instructions, A and AEX are regarded as a double-length result register, A holding the most significant half, and AEX the least significant half. Both will have the format of page 4.3.3.

In all instructions AIB will form the 64-bit operand if digit 51 of AOD is one. If AOD is zero digits 32 - 63 of AIB will form the most significant part of the 64-bit operand of which the other half is zero.

The operation of the floating point instructions are dependent upon the setting of digits 62 and 63 of AOD. Digit 62 is the inhibit rounding digit. Rounding is performed by forcing 1 into digit 63 of A if the mantissa of AEX is non zero. Digit 63 is set to select the special double length versions of add and subtract and reverse subtract and is ignored by all other operations.

**LOAD SINGLE (=)**

First, digit 51 of AOD is set to zero, then digits 32 - 63 of AIB are copied to digits 0 - 31 of A. Digits 32 - 63 of A are cleared.

**LOAD DOUBLE (=')**

First, digit 51 of AOD is set to a one, then AIB is copied to A.

**STACK AND LOAD (\*=)**

A (or digits 0 - 31 of A) is stacked then A is loaded as in = / =' if digit 51 of AOD is 0/1. Digit 51 of AOD is unaltered.

**STORE (=>)**

A (or digits 0 - 31 of A, as for X) is stored depending on whether digit 51 of AOD is 1 (or 0).

## ADD (+)

The operand from AIB is added to A. First the exponents of A and AIB are compared and the exponent field of A is replaced by the larger. The mantissa associated with the smaller exponent is then shifted right by the number of hexadecimal places given by the exponent difference. Also the digits which are shifted out are placed in the mantissa field of AEX the rest of AEX being cleared. The mantissa field of A is set to the sum of the mantissa fields of A and AIB (one of which may have been shifted). The normalisation shifts which follow apply across the mantissa fields of both A and AEX, with a maximum shift of 13 hexadecimal places. If both mantissa fields are zero a standard floating point zero is generated (4.3.3). The exponents of A and AEX are both set to the exponent of the double length result and rounding is performed as described above.

When digit 63 of AOD is set and A and AEX contain a double length number smaller than AIB a correct unrounded double length result will be formed.

If either of the above cause exponent overflow/underflow digit 57/58 of AOD will be set.

## SUBTRACT (-)

The operand from the highway is subtracted from A. The operation proceeds in the same general way as ADD. However, if the number to be subtracted is the smaller, it is negated and then the add operation is performed.

## MULTIPLY (\*)

A is multiplied by the operand to give a double length result in A and AEX. This result is standardised, and AEX exponent set as above. Rounding will occur if digit 62 of AOD is not set. On exponent overflow (or underflow) digit 57 (/58) of AOD is set.

- DIVIDE (/) A is divided by the operand to give a single length standardised (possibly rounded) result in A. If the divisor is zero digit 61 of AOD is set or if exponent overflow or underflow occurs digit 57 (58) of AOD is set.
- NON EQUIVALENCE ( $\neq$ ) The result of combining A and AIB with logical  $\neq$  is returned to A.
- OR (V) The result of combining A and AIB with the logical V is returned to A.
- SHIFT (A) A is shifted circularly (left) by the number of places specified by digits 58 - 63 of AIB.
- AND (&) The result of combining A and AIB with a logical & is returned to A.
- REVERSE SUBTRACT ( $\ominus$ ) This is the same as SUBTRACT, except that A is subtracted from the operand.
- COMPARE (CCMP) A is compared with the operand in AIB and the test register is set. Both are assumed to be floating-point numbers.  $T_0$  of the test register is set if any of bits 57, 58, 61 are set.
- CONVERT (CCNV) The only conversion function provided in the floating-point set is one which converts the integer part of A to a signed fixed point number leaving the result in AEX. If the result is too big digit 59 of AOD is set.
- REVERSE DIVIDE ( $\oslash$ ) This is the same as DIVIDE except that the operand is divided by A.

Chapter 5 Structure Accessing and Store to Store Orders5.1 Introduction

This chapter defines the registers D, XD and DOD (Section 5.2) and describes the orders associated with the secondary operand unit. The orders fall into three classes:-

- (a) Register manipulation (Section 5.3)
- (b) Structure access (Section 5.4)
- (c) Store to Store (Section 5.5)


The register manipulation orders are concerned with loading and storing the registers D and XD. The structure access orders are concerned with modifying descriptors and accessing elements of data structures. The store to store orders enable operations to be carried out on strings of bytes of any length, e.g., moving one string to another, or comparing two strings. The registers D and XD are used to hold the descriptors of the strings.

The STACK order is described in Section 5.3, chiefly because it isn't described anywhere else.

This chapter assumes that the reader is familiar with the different types of descriptor defined in sections 2.11 - 2.14.

## 5.2 Internal Registers in the Secondary Operand Unit (SEOP)

The two main registers in the SEOP are D and XD. They are both 64 bits long and are used to hold descriptors, so they have type, bound and origin fields as shown below:-



The following notation is used for various parts of D, XD:-

DO	the origin field of D	(d32-63)
XDO	the origin field of XD	(d32-63)
DB	the bound field of D	(d8-31)
XDB	the bound field of XD	(d8-31)
DT	the top half of D	(d0-31)
XDT	the top half of XD	(d0-31)

The only other register in SEOP which can be used by the programmer is DSD; DOD, D, XD, DT and XDT may all be read from or written to as internal register operands. DOD is a 32-bit register which contains the interrupt and interrupt inhibit bits for SEOP as follows:-

d31	XCH	XCHK digit
d30	ITS	Illegal Type/Size
d29	EMS	Executive Mode Subtype used in non-executive mode
d28	SSS	Short Source String in store to store order
d27	NZT	Non-Zero Truncation when storing secondary operand
d26	BCH	Bound Check Fail during secondary operand access
d25	SSSI	SSS Interrupt Inhibit
d24	NZTI	NZT Interrupt Inhibit
d23	BCHI	BCH Interrupt Inhibit

ITS and EMS will always cause an interrupt. SSS, NZT or BCH will cause an interrupt unless SSSI, NZTI or BCHI respectively are set.

5.3 D and XD Manipulation Orders and STACK

STACK	Stack the operand (advance SF by 2, then store operand at new SF).
DO =	Load the origin of D from bits 32-63 of the operand. Bits 0-31 of D are unaltered.
D =	Load D from bits 0-63 of the operand.
D *=	Stack D (advance SF by 2, then store D at new SF). Then load D from bits 0-63 of the operand.
D =>	Store D in bits 0-63 of the operand.
DB =	Load the bound of D (bits 8-31) from bits 40-63 of the operand. The rest of D is unaltered.
XDO =	} As for DO =, D =, D =>, DB = but operate on XD instead of D.
XD =	
XD =>	
XDB =	

Note: Some of these orders may be used with secondary operands. For S[B] and S[O] operands, the effect is as follows:-

- (a) DO =, D=, DB=  
D will first be loaded with the S operand descriptor; then the secondary operand will be accessed and will replace the whole or part of the new value of D.
- (b) D \*=  
The original contents of D will be stacked before the S descriptor is loaded into D.
- (c) XDO =, XD =, XDB =  
Work as expected.

The orders D => and XD => may not be combined with any secondary operand (S[B], S[O], D[B] or D[O]).

5.4 Structure Access Orders

[None of these orders may be used with secondary operands].

- MOD Uses bits 32-63 of the operand as a signed integer modifier for the descriptor in D. The modifier is added to the origin field (after scaling if  $US = 0$  in type 0 or 2 descriptors), and subtracted from the bound field. A bound check interrupt will occur unless  $0 \leq \text{modifier} < \text{bound}$  (assuming bound checking is not inhibited). The bound check applies to descriptors of types 0, 1, 2, 3.0, 3.1 and 3.2. For type 3.3, the indirectly addressed descriptor is loaded into D before the modification takes place. Similarly for types 3.4 - 3.31 the procedure is called first.
- XMOD Exactly the same as MOD except that it works on XD instead of D. (Types 3.3, 3.4 - 3.31 are illegal).
- SMOD As for MOD, but DB is unaltered and there is no bound check.
- MDR Equivalent to MOD followed by  $D = D[0]$ .
- RMOD Bits 0-31 of the operand are loaded into bits 0-31 of D. Bits 32-63 of the operand are added to bits 32-63 of D.
- MOCHI If  $0 \leq \text{operand bits 32-63} < XDB$  then bit 31 of BGD is set to 1, otherwise it is set to 0.



SUB1 A complicated order which works as follows:-

```

XD = operand (bits 0-63)
D = 0      (clear all bits of D)
B = XD[0]  (B = operand addressed by XD)
B * XD[1]
DB = XD[2]
MOD B
XMOD 3

```

XD must be a vector descriptor (type 0 or 2) addressing 32-bit elements.

SUB2 Omits the first two steps of SUB1:-

```

B = XD[0]
B * XD[1]
DB = XD[2]
MOD B
XMOD 3

```

#### Use of structure access orders

MOD (and XMOD) can be used for constructing substrings of longer strings, e.g., 'D = S; MOD I; DB = L' creates a descriptor for the string of length L starting at the Ith byte of the string S. MOD can also be used to step through a vector, since 'D = V; MOD 1' creates a descriptor to a vector consisting of all but the first element of V.

ADR can be used for moving through a list structure or for accessing arrays via Kiffe vectors.

XMOD is used for 'reverse modification'. This is useful when used in combination with the 'dope vector' orders SUB1 and SUB2 as described below. It can also be used to make data structures relocatable.

XCHK is a special order which is used to check for overlapping strings. The only dangerous case is when the start of the destination string (for a move or logical store to store order) lies within the source string. So the idea is to put the source string descriptor in XD and then use XCHK with operand = destination origin - source origin.

Dope Vector Orders

The SUB1 and SUB2 orders are used for accessing arrays via dope vectors. For a general array:-

$$X[l_1 : u_1, l_2 : u_2, \dots, l_n : u_n]$$

we want to access  $X[i_1, i_2, \dots, i_n]$ . The address can be expressed in the form:-

$$XO + (i_1 - l_1)*m_1 + (i_2 - l_2)*m_2 + \dots + (i_n - l_n)*m_n$$

where  $m_1, m_2, \dots, m_n$  are suitable multipliers; and in addition we must have  $l_1 \leq i_1 \leq u_1, l_2 \leq i_2 \leq u_2, \dots, l_n \leq i_n \leq u_n$ . When the array is declared, a dope vector is created which contains a triple of 32-bit elements for each dimension of the array; a triple consists of the lower bound  $l$ , the multiplier  $m$  and a checking value  $c$ . For the array  $X$  above, the dope vector will look like:-

$$\boxed{l_1 | m_1 | c_1 | l_2 | m_2 | c_2 | \dots | l_n | m_n | c_n}$$

A descriptor  $X^*$  is created which points to this vector. To access the element, the appropriate sequence is:-

B = i1	1st subscript
SUB1 X*	Load XD with dope vector descriptor and clear D. Subtract $l_1$ from B, multiply by $m_1$ , check result in range $0 \leq B < c_1$ and add to DO, (Hence $DO = (i_1 - l_1)*m_1$ and $l_1 \leq i_1 \leq u_1$ ).
B = i2	2nd subscript.
SUB2	DO + $(i_2 - l_2)*m_2$ $l_2 \leq i_2 \leq u_2$
.	
.	
B = in	nth subscript.
SUB2	DO + $(i_n - l_n)*m_n$ $l_n \leq i_n \leq u_n$

There are now two ways of accessing the element itself.

RMOD X	or	B = DO
A = D[0]		A = X[B]

[The checking values  $c$  are clearly  $(u_1 - l_1 + 1)*m_1$ ].

### 5.5 Store to Store Orders


The store to store orders fall into three classes, string-string, byte-string and table-string orders, and there is one special table look-up order. The string-string orders operate on a source string and a destination string; operations are provided which move (i.e., copy), compare and logically combine the strings. The byte-string orders use a byte and a destination string; they are the same as string-string orders in which the source string consists of the specified byte repeated as often as necessary. The table-string orders make it possible to translate the characters of a string into a different code specified in a table, or to check a string to see if it contains any of the characters specified in a table. The table look-up order scans the table for a particular element.

#### THE MASK

For all the store to store orders except table look-up (TANJ), bits 48-55 of the operand are an eight-bit mask MASK. In each byte processed by the order, bits corresponding to 1's in the mask are ignored; when any byte is used in an operation the corresponding bits are taken to be zeros, and if a byte is put into store the corresponding bits in the store are unaltered. For example, if MASK = 11000011, then a move order will only change bits 2-5 of the bytes in the destination string.

String-String Orders

For all string-string orders, XD contains the source string descriptor and D the destination string descriptor. The descriptors must have type 0, 1 or 2 and element size 8 bits, otherwise there will be an ITS interrupt. The operand defines the MASK (see above) and a FILLER. The FILLER is not in fact used for all the orders.



SMVB Move one byte from source to destination. If the source string is the null string, move FILLER to destination. If DB = 0 there will be a BCH interrupt. Updates DO, DB, XDO, XDB.

[if DB = 0 then BCH interrupt  
if XDB = 0 then (FILLER => D[0]; MOD 1)  
else (XD[0] => D[0]; MOD 1; XMOD 1)]

\*SMVE Moves source to destination. If source is shorter than destination there will be an SSS interrupt; but note that this will simply terminate the operation if the inhibit bit SSSI is set. Updates DO, DB, XDO, XDB.

[L1: if DB ≠ 0 then (if XDB = 0 then SSS interrupt  
else (XD[0] => D[0]; MOD 1; XMOD 1); -> L1)]

SMVF Moves source to destination. If source runs out, then uses FILLER for remainder of destination. Updates DO, DB, XDO, XDB.

[L1: if DB ≠ 0 then (SMVB OPERAND; -> L1)]

\* See page 5.5.5.

SCMP Compares source and destination strings looking for inequality in two corresponding bytes. If source runs out, FILLER is used. The test register is set = 0 if no inequality is found, > 0 if source byte > destination byte, < 0 if source byte < destination byte; for comparison purposes, the bytes are treated as unsigned integers. DC, DB, XDC, XDB are updated.

```
[L1: if DB = 0 then (T = '='; -> L4);
      if XDB = 0 then (if FILLER ≠ D[0] then -> L2 else
                        (MOD 1; -> L1))
      else if XD[0] ≠ D[0] then -> L3 else (MOD 1;
                                             XMOD 1; -> L1);
L2: if FILLER < D[0] then T = '<' else T = '>' ; -> L4 ;
L3: if XD[0] < D[0] then T = '<' else T = '>' ;
L4:      ]
```

SLGC Source and destination are logically combined and the result stored in destination. The logical operation is the same for each bit of each byte and is defined by bits 44-47 of the operand.

d44	d45	d46	d47
L0	L1	L2	L3

The result of the operation is defined by the table below:-

source bit	0	0	1	1
destination bit	0	1	0	1
result bit	L0	L1	L2	L3

If the source string runs out, there will be an SSS interrupt.

```
[L1: if DB ≠ 0 then if XDB = 0 then SSS interrupt; (XD[0] lge
D[0] => D[0]; MOD 1; XMOD 1; -> L1)]
```

Byte-String Orders

The byto-string orders are the same as the string-string orders except that the source consists of copies of the BYTE specified in operand bits 56-63. The MASK appears as usual in operand bits 48-55. D contains the destination string descriptor which must have type 0, 1 or 2 and element size 8 bits, or there will be an ETS interrupt; XD is not used.

	d48		d63
OPERAND	MASK		BYTE

BMVB Moves one BYTE to destination.

BMVE Moves BYTE's to destination until full.

BSCN Scans destination looking for a byte = BYTE. The test register is set = 0 if an equality is found, < 0 otherwise. DB and DE are updated.

```
[L1: if DB = 0 then (T = '<' ; -> L2);
      if BYTE = D[0] then (T = '=' ; -> L2) else (MOD1 ; -> L1);
L2: ]
```

BCMP Scans destination looking for byte ≠ BYTE. The test register is set as for SCMP.

BLCC Combines BYTE with destination string using logical operation defined by operand bits 44-47 as for SLCC. Result is stored in destination.

Table-String Orders

For both the table-string orders, XD contains a string descriptor which must have type 0, 1 or 2 and element size 8 bits. D may contain any descriptor. The operand contains no information other than the MASK, specified as usual in bits 48-55.

\*TRANS Each byte of the XD string is processed in turn. First, it is used as a modifier for the descriptor in D to access a secondary operand. Then the least significant 8 bits of the secondary operand replace the original byte. There may be a BCH interrupt during the D access. D will usually contain a byte vector descriptor.

[L1: if XDB  $\neq$  0 then (D[XD[0]] => XD[0]; XMOD 1; -> L1)]

\*TCHK Each byte of the XD string is accessed as above and used as a modifier for the descriptor in D. If the least significant bit of the secondary operand is a 1, then the operation is terminated with BN = 0. If no 1 is found for the whole of the XD string, then BN = 1. D will usually contain a bit vector descriptor.

[L1: if XDB = 0 then BN = 1 else  
 (<if l.s. bit of D[XD[0]] = 0 then (XMOD 1; -> L1)  
else BN = 0)]

\*TRANS, TCHK, SMVE

These orders are not commissioned. If an attempt is made to execute any of them the effect will be that of a DUMMY order, except that an interrupt may occur if the wrong type, size or length has been specified as described above, and TCHK will set the Test Register in an unspecified manner.

Table Look-Up

TALU This order enables a fast scan to be made for an element equal to the operand. D contains a descriptor which defines the table - origin in DO, length in DB. The length is expressed in byte units. The descriptor must have type 0 or 2 and element size 32 bits. XDR contains a MASK which is used in exactly the same way as the mask in the other store to store orders; bits corresponding to 1's in the MASK are ignored. The least significant 32 bits of the operand are compared with each element of the table in turn for equality (under control of MASK). If no equality is found, then the operation terminates with the test register set > 0, and the descriptor in D updated (DB = 0, DO points after end of table). If equality is found, then the test register is set = 0, and the descriptor in D will point to the element found, with the bound field updated.

[L1: if DB = 0 then T = '>' else  
           (if operand ≠ D[0] then (MCD 1; -> L1) else T = '=')]

N.B. TALU takes operands directly from store.



## Chapter 6 Organisational Orders

### 6.1 Introduction

The format for the organisational orders is shown below.

$$\frac{\text{cr} = 0 \quad | \quad f' \quad | \quad N'}{3 \quad \quad \quad 6 \quad \quad \quad 7}$$

The cr bits are zero and the f' bits define the function to be performed. N' defines the operand for the order as described in Chapter 2; the only kind of operand which cannot be addressed is an internal register.

The organisational orders fall into the following groups:-

- (a) Register operations - orders manipulating NB, XNB, SP, MS.
- (b) Control transfers and procedure call orders.
- (c) Conditional control transfers.
- (d) Boolean orders - operating on the 1 bit Boolean register BI.
- (e) Special orders.

## 6.2 Register Operations

### NB, SF and XNB Orders

The registers NB, SF and XNB are defined in section 2.2, and may be regarded as unsigned integers. In the following orders, it should be remembered that the least significant bit of each register is permanently zero, so that the least significant bit of the operand will have no effect.

NB = Load NB from bits 48 - 63 of the operand.

SF = Load SF from bits 48 - 63 of the operand.

XNB = Load XNB from bits 32 - 63 of the operand.

NB + Add operand bits 48 - 63 to NB; interrupt on segment overflow.

SF + As for NB +, but add to SF.

XNB + Add operand bits 48 - 63 to the least significant 16 bits of XNB; interrupt on segment overflow (carry into top half of XNB is not allowed).

SF = NB + Add NB to operand bits 48 - 63, and store the result in SF; interrupt on segment overflow.

NB = SF + Add SF to operand bits 48 - 63, and store the result in NB; interrupt on segment overflow.

For each of the orders NB+, SF+, XNB+, SF = NB+, and NB = SF+, the base register (or registers) involved are unsigned integers, but the operand is a signed 16-bit integer. The result must be in the range  $0 \leq \text{result} < 2^{\#} 16$ , or there will be a segment overflow interrupt.

NB => The name segment number SN is stored at bits 32 - 47 and NB at bits 48 - 63 of the operand. The operand may not be a secondary operand.

SF => As for NB =>, but store SN and SF.

XNB => As for NB =>, but store (all 32 bits of) XNB.

SN = Load SN from bits 33 - 47 of the operand. This order only alters SN if in Executive Mode.

The Machine Status Register MS

The machine status register contains 16 bits of system information numbered MS0 - MS15. MS8 - MS15 are concerned with the interrupt organization and can only be set in executive mode; they are described in Chapter 7. MS0 is the 'D set' bit whose use is explained under the procedure call descriptor in Section 1.2.14; MS4 - M7 are the three test bits T0, T1, T2 and the Boolean BF described in Sections 6.4 and 6.5.

MS	0	1		4	5	6	7	
	DS	SPARE		T0	T1	T2	EN	EXECUTIVE
	1	1	2	1	1	1	1	8

A  
|  
— Inhibit Program Faults

MS = This order sets various bits of MS to 0 or 1 depending upon bits 32 - 63 of the operand.

Let  $0 \leq i \leq 7$ . Then

- (a) if operand bit  $(56 + i) = 0$ ,  $MS(8 + i)$  is unaltered
- (b) if operand bit  $(56 + i) = 1$ ,  $MS(8 + i)$  is set to operand bit  $(48 + i)$
- (c) if operand bit  $(40 + i) = 0$ ,  $MS(i)$  is unaltered
- (d) if operand bit  $(40 + i) = 1$ ,  $MS(i)$  is set to operand bit  $(32 + i)$ .

If not in executive mode, MS8 - 15 are unchanged, and only (c) and (d) above apply.


Example: The order MS = 00010111 11001111

would set MS8 = MS9 = MS12 = 0, MS13 = MS14 = MS15 = 1, and leave MS10 and MS11 unaltered.

MS is also altered by the EXIT and RETURN functions (Section 6.5). When any order altering MS causes the By-pass CPR's digit to be altered, an Acc => Z instruction must precede it in order that all store operations will be completed before the CPR's are turned on or off. Care must also be exercised in turning the Name Store or Level 0 bit on or off, and when altering bits 12 and 13.

### 6.3 Control Transfers and Procedure Calls

The link is a 64-bit register with format:-



MS and NB are defined in Sections 6.2 and 2.2 respectively. CO is the 16-bit word address of the order currently being obeyed; the most significant bit of CO is always zero.

-> Relative jump; 32 - 63 of the highway are taken to be a signed 2's complement integer and are added to CO. An attempt to transfer control across a segment boundary will cause an interrupt.

JUMP Absolute jump; CO is loaded from bits 33 - 63 of the highway.

STACKLINK CO and bits 32 - 63 of the highway are added as for the relative jump, and the result together with MS and NB is stacked. Symbolically:-

STACK [MS, NB, CO + operand]

The addition CO + operand may give overflow as for ->.

RETURN The operand part of this order must specify the STACK. The order sets SF = NB and then unstacks the link. Symbolically:-

SF = NB

[MS, NB, CO] = [SF]

SF - 2

Bits 8 - 15 of MS are only reset if in executive mode.

If any operand other than STACK is specified, then the order is exactly the same as EXIT. (Note that if the operand specifies that the STACK is to be used as a descriptor, then SF is reset as above).

SEPLINK The link is stored at the address specified by the operand. Symbolically:-

[OPERAND] = [MS, NB, CO]


The operand may not be a secondary operand.

EXIT The link is reset from the operand. Symbolically:-

[MS, NB, CO] = [OPERAND]

Bits 8 - 15 of MS are only reset if in executive mode.

The following example illustrates how STACK LINK and RETURN can be used to call a procedure P with three parameters A1, A2, A3. Note that procedure calls implemented in the compilers are slightly more complicated (see Compiler Writers Manual). Before the call the stack will be:-



The call will look like:-


```

STACK LINK L1
STACK A1
STACK A2
STACK A3
JUMP P

```

L1:

so that after the call the stack looks like:-




The procedure itself will contain orders:-

```

PROCEDURE P
NB = SF - 3      to set NB for use as a base in the procedure
SF + n          for the n local names of the procedure
RETURN

```

The order NB = SF - 3 sets NB -> LINK in the stack; and SF + n advances SF.



The RETURN will reset MS, NB, CO from the LINK in the store and return SF to its original position.

#### 6.4 Conditional Control Transfers

The test bits T0, T1, T2 are bits 4, 5 and 6 of the machine status register MS (see Section 6.2). They are set by the computational orders CCMP and CINC (see Chapters 3 and 4) and by some of the store to store orders. The significance of these bits is generally as follows:-

T0	set to 1 if overflow
T1	set to 0 if result = 0, 1 if result $\neq$ 0
T2	set to 0 if result $\geq$ 0, 1 if result < 0

A set of seven orders is provided which causes a relative jump if MS is suitably set. If the test succeeds then the jump is carried out in exactly the same way as for ->.

IF = 0, ->	jump if T1 = 0
IF $\neq$ 0, ->	jump if T1 = 1
IF $\geq$ 0, ->	jump if T1 = 0 or T2 = 0
IF < 0, ->	jump if T2 = 1
IF $\leq$ 0, ->	jump if T1 = 0 or T2 = 1
IF > 0, ->	jump if T1 = 1 and T2 = 0
IF OVERFLOW, ->	jump if T0 = 1

There is an eighth conditional jump order which may be used to test the Boolean, BN, described in the next section.

IF BN, ->	jump if BN = 1
-----------	----------------

## 6.5 Boolean Orders

The Boolean, BN, is bit 7 of the machine status register MS. There are two kinds of orders which set BN. The first kind combines BN with the result of a test, and uses the operand to define what logical operation to perform; the second kind combines BN directly with the operand.

The first kind of order tests MS in one of 8 ways to produce a result R equal to 0 or 1.

= 0	R = 1 if T1 = 0,	0 otherwise
≠ 0	R = 1 if T1 = 1,	0 otherwise
≥ 0	R = 1 if T1 = 0 or T2 = 0,	0 otherwise
< 0	R = 1 if T2 = 1,	0 otherwise
≤ 0	R = 1 if T1 = 0 or T2 = 1,	0 otherwise
> 0	R = 1 if T1 = 1 and T2 = 0,	0 otherwise
OVFLOW	R = 1 if TO = 0,	0 otherwise
BN	R = 1 if BN = 1,	0 otherwise

Bits 59 - 63 of the operand define the way in which this result R is to be combined with BN as follows:-

0000	BN = 0	set BN = 0
0001	BN &	and BN with R
0010	BN/&	invert BN, then and with R
0011	BN =	load BN with R
0100	BN &/	and with inverse of R
0101	BN = BN dummy order	
0110	BN ≠	not equivalence with R
0111	BN V	or with R
1000	BN/&/	invert BN, then and with inverse of R
1001	BN =	equivalence with R
1010	BN/	invert BN
1011	BN/V (implies)	invert BN, then or with R
1100	BN =/	load BN with inverse of R
1101	BN V/	or with inverse of R
1110	BN/V/	invert BN, then or with inverse of R
1111	BN = 1	set BN = 1

The second kind of order use 4 of the f' bits to specify the function as above, and takes the operand R from the least significant bit of the highway.

## 6.6 Special Orders

- XCO-6 Stack the operand, and jump to segment 8193, 32-bit word locations 0 - 6 respectively.
- DI= The 32 Display Lamps on the Engineers Console are set equal to bits 32 - 63 of the operand. The Display Lamps may also be written to as a V-Line (Chapter 8).
- SPM This function is for use with the System Performance Monitor associated with the MU5 System.



Chapter 7 The Interrupt System7.1 The Interrupt Structure

There are eight types of interrupt divided into two groups of four, the System interrupts and Process based interrupts. The system interrupts are concerned with activities external to the current process (e.g., peripheral control). The process based interrupts occur as a result of specific actions in the current process. The interrupts are shown below, each associated with a three bit interrupt number.

SYSTEM INTERRUPTS	{	000	System Error
		001	CPR Not-Equivalence
		010	Exchange
		011	Peripheral Window
PROCESS BASED INTERRUPTS	{	100	Instruction Count Zero
		101	Illegal Orders
		110	Program Faults
		111	Software Interrupt

When interrupts occur simultaneously, the first to be dealt with is the one with the smallest interrupt number.

When an interrupt occurs, the hardware stops what it is doing and enters an interrupt sequence. During this entry sequence the 'Interrupt Entry Bit' is set. This allows the sequence to run in a special non-interruptible mode of operation described by the table in 7.2.

The first action of this sequence is to retain the state of the machine in a compact form to allow a straightforward return to the current process after dealing with the interrupt. This is achieved by storing a 64-bit link word. The format of this link is the same as the control register consisting of Machine State register (16 bits), Name Base register (16 bits) and the 32-bit control address.

In addition to retaining a link, the interrupt sequence also transfers control to the appropriate interrupt procedure. This control transfer is achieved by resetting the 64 bits (MS, NB, CO) from the second half of the Ith double word entry of a table. (I is the interrupt number). The first word of this entry is used to hold the link. This table is 16 x 64 bits long and starts at word 16 of the first of the common segments (segment 8K).

When a CPR  $\neq$  interrupt occurs the link will point to the next instruction to be obeyed, but the previous instruction may not be complete. Incomplete instructions are held in the OBS buffer and the CPR  $\neq$  interrupt routine must preserve (and restore) this buffer before using any instructions which could alter its content. On other interrupt entries all instructions up to the one addressed by the link will be complete.

Unserviced interrupts may be read in Prop V line 26.

Before discussing the interrupts in detail it is necessary to describe the Machine State Register.

## 7.2 The Machine State Register

The machine state register (MS) is 16 bits long and only the m.s. 8 bits (0,1,4-7) may be directly altered by the user program; bits 4 to 7 are test registers.

Bits 2 & 3 & the l.s. 8 bits are known as the system mode bits and they may only be altered if the executive mode bit is set. This bit is set by interrupt entry or by the functions  $KC_0 \dots KC_6$  (see section 6.6). MS is arranged as follows:-

BIT	
0	Force DR[ ] instead of S[ ]
1	Inhibit program fault interrupts (A,B,D etc.)
2	System Performance Monitor
3	Spare
4	Overflow
5	≠ 0
6	- ve
7	Ecolean
8	Bypass CPR's
9	Bypass Name Store
10	Instruction counter Inhibit
11	B and D faults to System Error in Exec Mode
12	A faults to System Error in Exec Mode
13	Exec Mode flip-flop
14	Level 1 Interrupt flip-flop (LIIF)
15	Level 0 Interrupt flip-flop (LOIF)

For description of bits (0 - 7) see .

Effects of System Mode Bits

Cause	CPR by-pass	Name Store by-pass	Instr. counter off	B or D under Exec control	ACC. under Exec control	Exec Mode flip-flop	L1IF	LOIF	Interrupt entry bit
Effect	8	9	10	11	12	13	14	15	
By-pass CPR's	✓								
by-pass Name Store		✓							✓
Instruction Counter off			✓				✓	✓	✓
Allow V-access						✓	✓	✓	✓
Allow system mode bits of MS to be altered									
Inhibit L1							✓	✓	✓
Inhibit LO								✓	✓
Inhibit System Error									✓
Force relevant program fault to be system error				✓	✓	✓	✓	✓	✓
Use LO name store only								✓	
4x64 bit lines									

B or D interrupts will be forced as system errors if D11 is set and an ACC interrupt will be forced as a system error if D12 is set. All other program faults will be forced as system errors if D13, D14, D15 or the Interrupt bit is set.

(N.B. The Name Store is by-passed when the interrupt entry bit is set since access to common segments cannot be made via the name store).

LO Name Store

In LO mode (bit 15 set) the 8 32-bit lines of name store are used as fast registers. The hardware interprets only the bottom 3 bits of a 'name type' operand address and maps it into this name store. No store accesses occur. These fast registers may be used as 8 x 32-bit names or 4 x 64-bit names.

### 7.3 The System (Level 0) Interrupts

The procedures which service these interrupts must be written so as to not cause any other Level 0 system interrupts, for example, the Peripheral Window Interrupt procedure should not cause a CPR Not Equivalence Interrupt. This requires that a few CPR's are permanently allocated to cover the program and working store used by these Level 0 Interrupt procedures. If a CPR  $\neq$  interrupt occurs while the Level Zero Interrupt flip-flop is set, then a System Error interrupt is caused. This also occurs if an illegal hardware function is executed while the Level Zero Interrupt flip-flop is set. In this sense the System Error Interrupt differs from the other System Interrupts. However, once it does occur it cannot recur until the System Error Status register (see below) is reset, or unless the Engineers 'Interrupt' is pressed.

#### The System Error Interrupt

The System Error Interrupt is caused by hardware or Executive failures. The System Error Status register shown below can pinpoint the exact cause. The software action on this interrupt is to perform error diagnosis, and restart normal system operations if possible.

<u>bit</u>	<u>error indication</u>
48	Engineers Interrupt (Console) - (forces CPR bypass)
49	Early Warning Power Failure
50	SAC Parity
51	Name Store Multiple equivalence
52	OBS Multiple Equivalence
53	CPR Multiple equivalence
54	Spare
55	IBU Multiple Equivalence
56	B or D error & (MS11)
57	Acc error & (MS12)
58	Illegal function & (LOIF + LIIF + EXEC)
59	Name adder overflow & (LOIF + LIIF + EXEC)
60	Control adder overflow & (LOIF + LIIF + EXEC)
61	CPR exec Illegal
62	CPR $\neq$ & (LOIF)
63	Spare

27.11.73/2

### CPR Not-Equivalence Interrupt

This interrupt can be produced by a user program or by an executive Mode procedure. It occurs when an attempt is made to access an address which does not lie within the address field defined by the contents of the CPR's. If the required address lies in local store, the procedure will free a CPR and allocate it to the page containing the address. Control is then returned to the interrupted process.

If the page containing the required address is not in local store, then the procedure will locate the page and organise its transfer to local store. In this case control is not returned to the interrupted process (which is halted awaiting the termination of the transfer) and a process change may occur.

### Exchange Interrupt

This interrupt is set by the Block Transfer Unit on completion (or termination) of a Core to Core Transfer.

### Peripheral Window Interrupt

This Interrupt is caused by an external device (e.g., peripheral processor, drum) writing to the Peripheral Window V-line or by an interrupt from the console (e.g., TTY, CLOCK). The two are distinguished by bits 61 and 60 in V-line  $\%011A$ , the first indicating peripheral window. The information sent to this 32-bit V-line consists of the sending unit number and a message. The Peripheral Window procedure must queue up this message for subsequent processing.

Writing to the Peripheral Window V-line sets it not busy in readiness for another message. The Peripheral Window procedure runs with interrupts inhibited, so a subsequent message won't be acknowledged until the procedure is exited.

This interrupt is also entered for console interrupts (Teletype and Clock). Prop V-line 26 contains the cause of interrupt.

## 7.4 Process Based Interrupts

### The Instruction Counter Zero Interrupt

This interrupt is set when the instruction counter becomes zero. It may be inhibited by the 'instruction counter inhibit' bit being set in the Machine State Register.

### The Illegal Order Interrupt

This interrupt is caused by program fault conditions detected by the hardware; these conditions set bits 48 - 53 in the program fault status V-line. (See next section).

### The Program Fault Interrupt

This interrupt is caused also by program fault conditions detected by the hardware which sets bits 56 - 58 in the program fault status V-line.

The following table relates the assignments of bits in the Program Fault Interrupt Status V-line to specific fault conditions:-

<u>V-line bit</u>	<u>condition</u>
48	Illegal function & $(\overline{LOIF} + \overline{LIIF} + \overline{EM})$
49	Name Adder overflow & $(\overline{LOIF} + \overline{LIIF} + \overline{EM})$
50	Control Adder overflow & $(\overline{LOIF} + \overline{LIIF} + \overline{EM})$
51	Illegal V store access
52	CPR $\overline{EXEC}$ illegal (illegal access via the CPR's when not in the exec mode)
53	Parity
54	System Performance Monitor
55	Spare
56	B fault & $\overline{MS1}$
57	D fault & $\overline{MS1}$
58	Acc fault & $\overline{MS1}$

When a program fault occurs, the hardware may not wait until the arithmetic and control units have finished their current operation(s), so that multiple fault conditions may not be completely recorded in the Program Fault Status register.

### The Software Interrupt

This interrupt occurs in user mode only when the software interrupt bit is set. (See 8.3.3).





Chapter 8 The V-Store8.1 Introduction

The V-store contains hardware registers used to control and/or diagnose parts of the MU5 processor. The V-store does not include the Internal Registers, which are addressed by a different mechanism. The V-store is not generally accessible to a user program but is accessible from within MU5 to Executive, interrupt routines (7.2.1) and certain control or hardware diagnostic programs.

The figure in 1.2.3 shows the instruction format required to obtain a V-store operand.

The V-store is divided into 128 blocks of 256 lines each. Each line is normally a 64 bit quantity but many of the lines will contain less than 64 bits, in which case the bits will appear right justified in a 64 bit word. The bits in a V-line are numbered as they would appear on a 64-bit wide highway. The following table gives the allocation of block numbers to sections of V-store in the central part of the machine.

<u>BLOCK NO.</u>	<u>V-STORE TYPE</u>
0	SYSTEM V-STORE (S8192)
1	PROP V-STORE
2	CBS V-STORE
3	CONSOLE V-STORE
4	SAC V-STORE
5	YBU V-STORE
6	PERIPHERAL WINDOW
7	PARITY V-STORE

## 8.2 System V-store (S8192)

These 256 V-line addresses are mapped into the first 512 x 32 bit words in segment 8192, the first of the common segments, by the PRDP before using them to access store. This gives the executive a simple means of communicating between processes and approximates the Atlas working store.


Starting at the 32nd 32-bit word of segment 8192 are 8 pairs of new and old links used by the interrupt entry sequence.

<u>V-line</u> (Decimal)	<u>Virtual Address (S8192)</u> Hex specifying 32-bit boundaries	<u>Name</u>	<u>size</u>
16	20	System Error Old Link	64
	22	System Error Entry Link	64
18	24	CPR $\neq$ Old Link	64
	26	CPR $\neq$ Entry Link	64
20	28	Exchange Old Link	64
	2A	Exchange Entry Link	64
22	2C	Peripheral Window Old Link	64
	2E	Peripheral Window Entry Link	64
24	30	Instruction Count Old Link	64
	32	Instruction Count Entry Link	64
26	34	Illegal Order Old Link	64
	36	Illegal Order Entry Link	64
28	38	Program Fault Old Link	64
	3A	Program Fault Entry Link	64
30	3C	Software Old Link	64
	3E	Software Entry Link	64

### 8.3 Primary Operand Unit V-Store (Block 1)

The following is a list of the registers in the V-store of the Primary Operand Unit, giving size, type of access and references to more detailed descriptions of usage and construction:-

<u>Address (Decimal)</u>	<u>Name</u>	<u>size</u>	<u>Access</u>
specifying 64-bit boundaries			
0	PROGRAM FAULT STATUS	16	R/W=Reset
	This line records fault reasons for the Program Fault Interrupt (bits 56 -> 58), the Illegal Order Interrupt (bits 48 -> 53) and the System Performance Monitor (bit 54). (See 7.4.1).		
1	SYSTEM ERROR STATUS	16	R/W=Reset
	The line contains the flip-flop used to record system error conditions for the System Error Interrupt. (See 7.3.1).		
2	PROCESS NUMBER	4	R/W
	This line contains the number of the current process and is used to generate the 'P' part of the virtual address. Writing to this line clears the JUMP TRACE in the Instruction Buffer Unit by resetting the Valid bits for each line (8.7).		
3	INSTRUCTION COUNTER	16	R/W
	This contains the number of instructions remaining to be executed before the Instruction Counter Interrupt occurs, (i.e., 64K machine instructions/Instruction Count Interrupt). The counter may be stopped by setting MS10.		
8	SEARCH ADDRESS	12	W




Bits 48 -> 59 specify the block address of a 16 x 32-bit word block in the name segment of the process specified in PROCESS NUMBER (LINE 2). The line number is ignored. Writing to this line causes an associative search of the name store using the SEARCH MASK (line 9) which must have been set up previously. If a line of the specified block exists in name store the test register is set non zero.

9

SEARCH MASK

12

W



The mask operates on the search (block) address. A '1' specifies that bit to be ignored.


16/17

(%10/%11)

NS LINE COPY

16+16 R

These lines contain a bit significant indicator showing which name store entry received the last valid name store access.




These lines have only READ access. However, writing to these addresses results in hardware action in the PROP V store without disturbing NS LINE COPY.

Writing to the address line 16 causes LINE POINTER to be reset. This is a hardware pointer to an entry in the name store. Resetting sets the pointer to the first entry in the name store. The pointer can only be altered by reading line 24 which causes it to be incremented by 1 (modulo 28). Writing to address line 17 causes all entries in the name store to be marked unaged and unaltered. The core copies of any existing entries are not updated.

24/25  
(%18/%19)

NS NEXT LINE VIRTUAL ADDR 4+15 R

These lines contain the virtual address in the associative name store entry pointed to by the LINE POINTER.



The normal virtual address format contains a segment number. Here the segment number is implied. It is the name segment of the process given by line 24. The line address in line 25 references a 64 bit word boundary in the name segment.

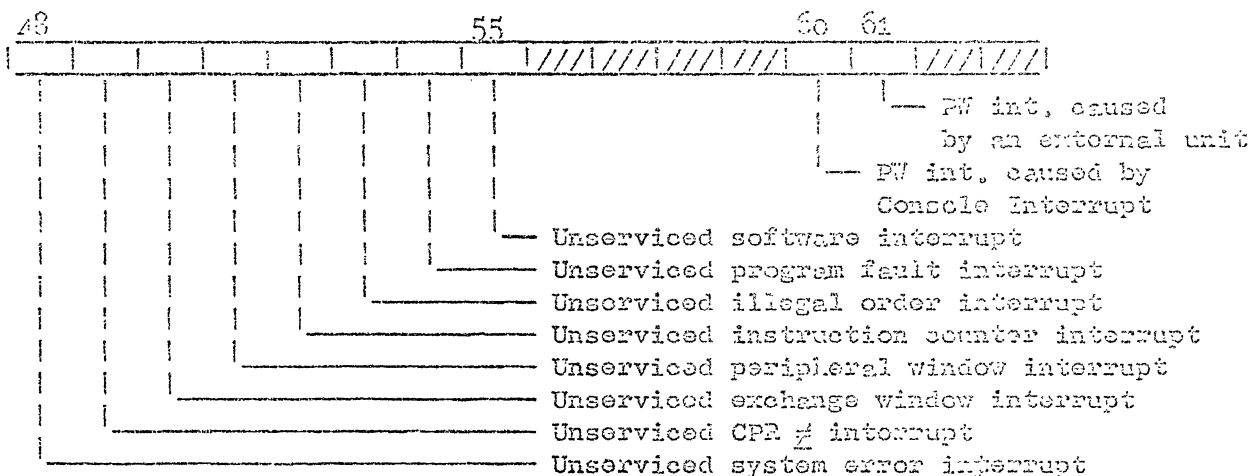
Reading line 24 causes the name store LINE POINTER to be incremented (cyclic modulo 28).

Although access to these lines is limited to READ only, writing to the address line 24 causes special action without the contents of line 24 being disturbed. Writing to line 24 causes the name store to be purged throughout.

26  
(%1A)

DISPLAY LAMPS 32 R/W

Writing sets the engineers display lamps. When READ, the following bits have the meaning:-



27

SOFTWARE INTERRUPT

1

R/W

(9.1B)

Bit 63 of this line is set by the software trapping mechanism and by process-based interrupts (see 7.1). It only causes an interrupt when in user mode.

#### 8.4 The OBS V-Store (Block 2)

The following short description of the main features of the OBS system will clarify points in OBS V-store control.

The OBS contains an operand store similar to the PROP name store but which is not restricted to dealing with only name segment operands (names). All operands which are not names and all operands in accumulator orders (names included) are buffered in the OBS operand store. (This means that the OBS has a name store of its own and the PROP name store only buffers names associated with 'non-accumulator' orders).

An entry in the operand store consists of virtual address and contents; this means an operand may be altered in the operand store without the main store version having been updated. The update only takes place when the operand is displaced by a new operand request. This replacement is normally cyclic depending on other activities in the OBS.

A request to the OBS consists of a function and its operand. All accumulator orders are passed to OBS which queues the function part and if necessary buffers the operand. The Acc queue has a maximum of six entries each of which references an entry in the operand store. Operands referenced by the Acc queue are avoided by the operand replacement mechanism described above.

A 'non-accumulator' order sent to the OBS bypasses the queuing mechanism and may be dealt with out of program sequence provided there is no possibility of a clash between its operand and those referenced by the Acc queue.

When an Acc order causes a CPR  $\underline{p}$ , processing of the Acc queue is halted. The leading entry is responsible for the  $\underline{p}$  and the remaining functions can only be processed sensibly when the  $\underline{p}$  has been serviced.

The following is a description of the OBS V-lines.

<u>Address</u> (Decimal)	<u>Name</u>	<u>Size</u>	<u>Access</u>
-----------------------------	-------------	-------------	---------------

0	OBS.CLEAR OBS.PURGE	1	W
---	------------------------	---	---

Writing a '0' to bit 63 causes the OBS to be CLEARED. This means all 'altered' operands in the OBS operand store are written back to main store. They are also retained in the operand store and reset to 'unaltered'.


Writing a '1' to bit 63 causes the OBS to be PURGED. This means all 'altered' operands are written back as in the CLEAR but in addition all lines in the operand store are set empty. i.e., the OBS operand store is left in a RESET state.

1	(Unassigned)		
2	OBS.MASK	12	W




This is used to mask the X field of the OBS.FIND line (see below). A '1' in the mask causes the corresponding bit in the Find operation to be ignored.

3	OBS.FIND	30	R/W
---	----------	----	-----




This address defines a 8 x 64 bit block boundary. Writing to this V-line initiates a masked associative search of the OBS operand store (see also line 2). If association equivalence occurs bit 63 of this V-line is set to '1'; otherwise it is set '0'. If association equivalence occurs the test register is set non-zero otherwise the test register is set equal to zero.





Writing to this line causes the functions on the Acc queue and their operands to be written to the specified 16 x 64 bit block. All useful information in the operand store is retained (i.e., the lines are not all RESET) but the Acc queue is returned to the 'empty' state.


Each dumped entry consists of 2 x 64 bit words.



This contains the Acc function (FUNC) and some hardware information (DOP). The operand virtual address is P, S, X, L and the mode at the time of access is in E (executive). The T bits specify the type of operand:-

Type no	OPERAND
0	Invalid (Empty entry)
1	Literal
2	Vector
3	Name

If the operand is a literal the virtual address field is irrelevant and the literal is held in the second word of the entry.




Note that when the operand is not a literal type, word 1 is irrelevant.

5

OBS.UNDUMP

29

W



The above address specifies the 16 x 64 bit block from which the OBS is to be reloaded. Writing to this V-line firstly causes a CLEAR operation to be performed (see line 0). Then the Acc queue is retrieved. Operands (other than literals) associated with the new Acc queue consist only of a virtual address part (see line 4 - OBS.DUMP). Any such operand which does not exist in the operand store at this time is now inserted in an 'unfilled' state, i.e., its 'contents' are not accessed from main store at this stage (see line 6 - OBS.RESTART).

No OBS orders may be executed between the Undump and Restart/Exit orders.

6

OBS.RESTART

-

W

An UNDUMP operation can leave the operand store with 'unfilled' operands. Such operands have to be filled by causing accesses to main store before normal operations may continue. Writing to this V-line causes the above action to be initiated when the next EXIT or RETURN order is obeyed. No order which makes use of the OBS must appear during this period.

The restart sequence forces all its operand accesses in executive mode. This means that if a parity occurs at this time a System Error will be generated.

7


(Unassigned)

8


OBS.INSPECT

28

R/W



Writing to this line causes all virtual addresses in the operand store to be written to the specified 32 x 64 bit block of store. Each virtual address will have the following format.




P, S, X and L are as described earlier (see line 4). The C bits have the following meaning:-

bit 31 = 1 means 'in use'

bit 32 = 1 means 'referenced' from the Acc queue.

Reading this V-line yields the following information.



— = 1 when Acc queue is full = not empty  
 — = 1 when Acc queue is partly full = not full

#### NOTES

OBS V-store addresses should not be used with organisational functions.

In REMOTE mode, if PRCP is sending orders to OBS, any attempt to access OBS V-store from a PPU through exchange may produce spurious results.

15

OBS.RESET

(3/5 F)

Writing to this V-line causes a reset of the OBS Buffer Store. Lines will be set 'not in use' and unaltered. No updating of the main store will take place.


8.5 Control Console V-Store (Block 3)

The console V-lines are as follows:-

<u>Address</u> (Decimal)	<u>Name</u>	<u>Size</u>	<u>Access</u>
-----------------------------	-------------	-------------	---------------

Specifying  
64-bit boundaries

0	CONSOLE INTERRUPT	4	R/W
---	-------------------	---	-----



Each of these is '1' when set.

F.C.I. is the fast clock interrupt. (Currently 1/100 sec).

T.C.I. is the teletype character interrupt.

T.E.I.I. is the teletype external incident interrupt.


S.C.I. is the slow clock interrupt. (Currently 1 sec).

External incidents are 'accept', 'cancel' and 'input request'. (See line 7).

Writing to this line cannot cause the T.E.I.I. bit to be reset. This can only be achieved by resetting line 7.

Writing a '1' to bit 62 resets bit 62. Writing a '1' to bit 63 resets bits 60 and 63.


2	TIME UPPER	13	R
---	------------	----	---



Hours and minutes appear in binary coded decimal in tens and units as shown.

3

TIME LOWER 15 R



Seconds and fractions of seconds, in b.c.d. as shown. This line is staticised by reading TIME.UPPER.

4

DATE LOWER 11 R



Months and days appear in binary coded decimal in tens and units as shown.

5

DATE UPPER/HCOOPER 8 R




The year in b.c.d.

Although this line has read only access, writing to bit 63 at this address will operate the hcooper.

6


TELETYPE DATA 8 R/W



To output a character, TELETYPE CONTROL must be in output mode, then writing to this line starts the transfer. On reading BIT 56 is the character parity.

7

TELETYPE CONTROL 8 R/W



Each bit has individual significance and is '1' when set active.

Digit

56	PRINT ON/OFF (OFF = '1')
57	CANCEL INSTRUCTION
58	Input/Output Teletype ('1' for input)
59	'Input Request'
60	'Accept'
61	'Cancel Message'
62	Teletype Start
63	Teletype online

When the TTY is online, bits 59, 60, 61 will cause an interrupt.

ON LINE (Peripheral)

PRINT	___	LIT	CANCEL	ACCEPT	INPUT
___	LIT	MESS	MESS	REQUEST	

OFF LINE (Instruction Source)

PRINT	LIT	___	CANCEL		MUST BE
___	LIT	INST			LIT

Depress 'OFF LINE'

10  
(%A)

MODE SWITCHES

16

R

The least significant 8 bits and the most significant 4 bits of this line specify various modes of operation when set to '1'

Digit

48 -> 55	These are used for switching the stacks of the Local and Mass stores off-line.
56	Level 0
57	Inhibit Clock Interrupt 1
58	Inhibit Interrupts
59	No overlapping of instructions.
60	Bypass Name Store.
61	Inhibit Clock Interrupt 0.
62	Allow Exchange resets.
63	Reset parity.

11 ENGINEERS HANDSWITCHES 16 R  
 (% B) All 16 bits (48-63) are used to control hardware diagnostic programs and error recovery procedures.

12 ENGINEERS CONTROL SWITCHES 10 R  
 (% C) Those appear in the 10 l.s. bits of the line and have the following significance. In 'Auto' all bits are zero.

Digit

54	Remote OFF/ON (See 1.9)
55	Reset
56	Interrupt
57	Single Shot
58	MCs. (i.e., not TEST)
59	STEP (i.e., not AUTO)
60	Increment CFF
61	PREPULSE ON
62	HANDKEYS for instruction source. (i.e., not TELETYPE).
63	Instruction buffer/Manual instruction


N.B. The above description of the Console V-store is only true as long as the REMOTE switch is OFF (Digit 54 line 12). If REMOTE is 'ON', lines 10, 11 and 12 have their access permission increased from READ only to READ/WRITE with the exception of bits 54 and 62 of line 12 which remain READ only.

8.6 SAC V-Store (Block 4)

The following is a list of registers in the V-store of the SAC unit. The list gives the address size, access and references to more detailed descriptions and constructions:-

Address                      Name                                      Size                                      Access  
 (Decimal)  
 Specifying  
 64-bit boundaries


0                      CPR SEARCH                                      30                                      W




The use of this line is described later under line 5 the CPR FIRD vector.

1                      CPR NUMBER                                      5                                      W  
 This contains the 5 bit CPR number (0 - 31).

2                      CPR VA                                      30                                      R/W



3                      CPR RA                                      31                                      R/W



The 4 LZ bits represent the page sizes 64K - 16 words as the values 12 -> 0 respectively. The AC bits are the access control on the page. (Shown below).

33	24		
0	0	Obey Only	Program
0	1	Obey and Read	
1	0	Read Only	Data
1	1	Read and Write	

bit 35 = 0 - Executive Mode only

1 - Any



Lines 1 -> 3 are used for reading from and writing to one of the 32 CPR registers. Each of these is conceptually divided into Virtual Address part and Real Address part of the format illustrated in CPR VA and CPR RA. Writing to either CPR VA or CPR RA initiates a write operation to the VA half or RA half of the CPR register specified by the contents of CPR NUMBER. Similarly, reading from CPR VA or CPR RA initiates a read from the VA or RA half of the CPR specified. A restriction involved in loading a CPR is that the Real Address part must be written to immediately before the Virtual Address part. This restriction does not apply when reading from a CPR.


- |   |  |    |     |
|---|--|----|-----|
| 4 | CPR IGNORE   | 32 | R/W |
|   | This consists of 32 bits each of which corresponds to a CPR and when set to '1' means that CPR is empty. This vector of bits is ordered so that the most significant refers to CPR line 0.   |    |     |
| 5 | CPR FIND   | 32 | R/W |
|   | This has the same format as CPR IGNORE but is used in conjunction with CPR SEARCH (line 0) and CPR FIND MASK (line 9) in an equivalence search through all the CPR registers. The CPR FIND MASK line specifies which bits in the PSX part of the virtual half of the CPR's are not used in the equivalence check and CPR SEARCH specifies the required bit pattern. Writing to CPR SEARCH initiates the operation. Each CPR which causes equivalence has a '1' 'OR'ed into its corresponding bit in the CPR FIND line. |    |     |
| 6 | CPR ALTERED  | 32 | R/W |
| 7 | CPR REFERENCED   | 32 | R/W |
|   | Lines 6 and 7 are vectors of the Altered and Referenced bits. These lines have the same formats as lines 4 and 5. An attempted access via a CPR causes a bit to be set in line 7 (& 6 if write access) even on access violation. Writing to a CPR (Vline 2) resets its bit in lines 4 -> 7.  |    |     |

9

CPR FIND MASK

16

W



The use of this line is described under line 5 CPR FIND. The Find mechanism operates over each bit of the segment field whereas the P and X fields both merely have a 'do' or 'do not' single bit specification. A '1' means do not search on this bit.


11

CPR X FIELD

24

R

(% B)




Reading this register initiates a read from the Virtual Address half of the CPR specified by CPR number. This line is required for engineering purposes and is fully described in the Engineering MU5 Manual, Chapter 6.

16

CPR NOT EQUIVALENCE PSX 30

R/W=Reset

(% 10)




17

CPR NOT EQUIVALENCE S 14

R

(% 11)



Line 16 holds the virtual address of the 16 word block which contains a line address that is to be presented to the CPR's for equivalence. Line 17 holds the segment field of this address. When a CPR Not Equivalence occurs, these lines remain set although further addressing through the CPR's can take place.

The action of writing to the PSX line returns both to their normal state. Any CPR  $\neq$  interrupts which occur during the CPR  $\neq$  interrupt procedure will be monitored as system errors. If these occur before the PSX line has been reset, there will be no information in PSX about the address causing the system error CPR  $\neq$ .


20

(%14)

SAC PARITY

8

R/W=Reset Bit 57



Bits 56 and 57 ARE RESET BY:- GENERAL RESET  
 WRITING TO THIS LINE  
 WRITING TO LINE 23  
 WRITING TO BLOCK 7 LINE 1

A data parity error locks out bits 57 - 63 until reset.  
 Bit 56 records the occurrence of all exec parity fails.

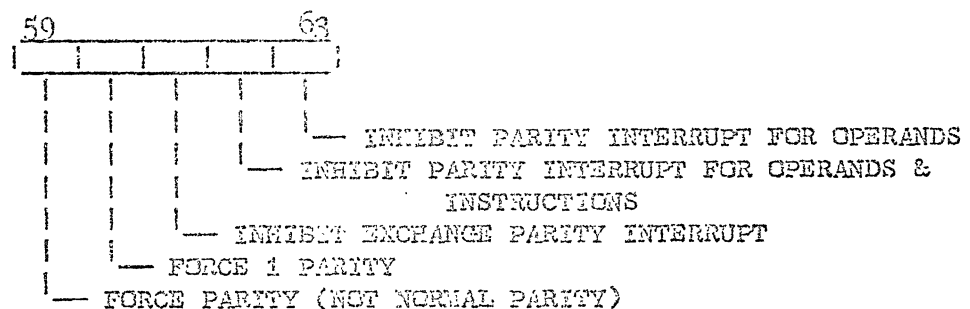
21

(%15)

SAC MODE

5

R/W



Interrupt inhibits do not inhibit the setting of the parity fail bits in lines 20 and 23.

All bits in this line are cleared by general reset.


22

(% 16)

ACCESS VIOLATION

3

R/W=Reset Bits 61 &amp; 63



Bit 61 is only valid when bit 62 is set, and refers to program faults only.

Bits 60, 62 and 63 are reset by:- General reset

Writing to this line.

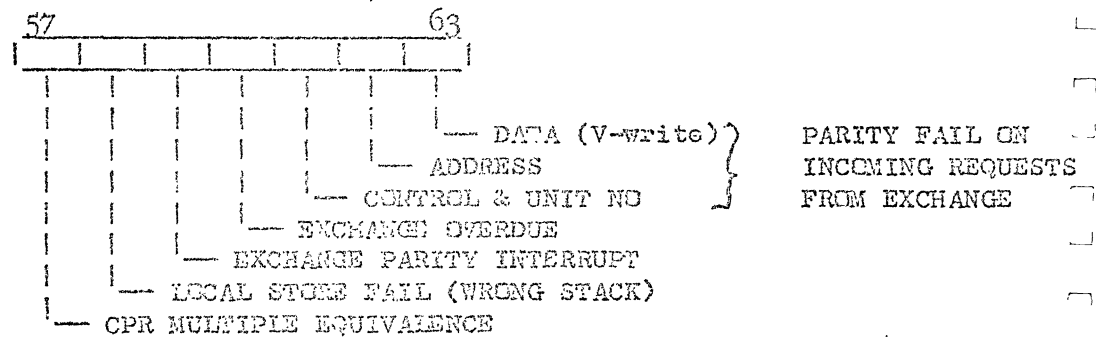
23

(% 17)

SYSTEM ERROR INTERRUPTS

7

R/W=Reset (see below)



These bits are all set independently of one another, as their respective faults are detected.

Writing to this line resets all bits except bit 58, and also resets line 20, and block 7 line 1. Bit 58 detects a hardware fault that cannot be cleared by software, and can only be reset by general reset.

General reset resets all bits.

Writing to block 7 line 1 clears the bottom 4 bits.

24

(% 18)

UNIT STATUS

R

1 - 1905E OPERATIONAL (Bit 63)

2 - EXCHANGE OPERATIONAL (Bit 62)

25

(% 19)

1905E INTERRUPT

W


Writing to this line causes an interrupt signal to be sent to the 1905E.

### 8.7 The IBU V-Store (Block 5)

The instruction buffer unit maintains a record of the eight most recent control transfers in the form of a 'jump from' address with a 'jump to' address. This table is known as the JUMP TRACE and software communicates with it by means of the IBU V-lines. The JUMP TRACE is not maintained in any interrupt mode. Since V store access can only be obtained in these modes this ensures that IBU V store is used sensibly when JUMP TRACE is static.


<u>Address</u>	<u>Name</u>	<u>Size</u>	<u>Access</u>
0	FILL-POINTER	4	R/W

64-bit  
boundaries



The line points to the entry in JUMP TRACE which is the next to be filled. When the hardware fills an entry the FILL-POINTER is incremented by 1 (modulo 8). Reading this line gives the format shown above. The Valid bit (bit 32) indicates whether that entry has been filled by the hardware since the last process change (see PROP V store 8.3).

When writing to this line the format is as follows:-



Bit 59 causes the FILL-POINTER to be written to when bit 60 (Trace On) is zero, i.e., the Trace is switched off. To switch the Trace on the line must be written to again with bit 60 set to a one, or a general reset given.

1

JUMP FROM

32

R

V1	'JUMP FROM' ADDRESS	
32		63

This line contains the 'JUMP FROM' address (the address of the last 16 bit section of the JUMP instruction) in the entry in JUMP TRACE pointed at by the FILL-POINTER.

N.B. The 'JUMP TO' addresses in JUMP TRACE cannot be read as V store.

It cannot be read unless bit 60 in line 0 has been set to zero.


8.8 Peripheral Window V-Store (Local Block 6)

<u>Address</u>	<u>Name</u>	<u>Size</u>	<u>Access</u>
0	MESSAGE WINDOW	32	R/W=Reset

This register belongs to MU5 V-store but in addition may have information written to it from other units in the system. The writing of this information causes an interrupt in MU5. The information may be read but any attempt to write to the line from within MU5 will merely cause the line to be set not busy.

8.9 Parity V-store (Block 7)

<u>Address</u> (Decimal)	<u>Name</u>	<u>Size</u>	<u>Access</u>
0	MU5 RIPP	1	R/W
	This provides a means of inhibiting further requests from MU5 to exchange (see also 9.2.1). General reset resets this to zero.		
1	EXCHANGE REQUEST PARITY	4	R/W=Reset (see below)



Writing to this line resets all bits and resets the l.s.

4 bits of block 4 line 20.

This line is reset by:- General Reset

Writing to this line

Writing to Block 4 line 23.

This line is duplicated in block 4 line 20.



Chapter 9 The Vx-Store9.1 Introduction

Vx-store consists of registers which control and/or diagnose system hardware and which are communicated with from MU5 and other units in the system. It defines the means of communication between MU5 and the other units of the system. MU5 may only access Vx-store when in executive mode or any interrupt mode. Access is achieved by a real address in a real address mode descriptor or by CPR bypass. Below is a list of the system Vx-stores:-

9.2	MU5 Vx
9.3	Disc (Drum) Vx
9.4	Block Transfer Unit Vx
9.5	1905E Vx
9.6	Local Store Vx
9.7	Mass Store Vx
9.8	System Performance Monitor Vx

N.B. A problem exists when writing to the Vx-store and then changing the status of the machine. As far as MU5 is concerned, a write order is complete when it is accepted by the store access control and it is possible for a large number of instructions to be obeyed before the order is actually executed. A software interlock must be applied in cases where this could cause trouble, e.g., writing to reset a BTU interrupt and then releasing the interrupt flip flops in Machine Status may result in a second BTU interrupt, which will apparently vanish and may look like a message interrupt in the worst case. A similar problem could arise with parity interrupts, and local store fail soft.

There are many ways of providing such an interlock and two such ways are illustrated:-

a)	=> Vx-Store	b)	=> Vx-Store
	Bn = same Vx-store		B $\neq$ same Vx-store
	(destroys Bn)		B $\neq$ 0
			(innocuous)

A and X orders do not provide a satisfactory interlock.

## 9.2 The MU5 Vx-Store

The normal MU5 Vx-store consists of the following:-

### The Peripheral Window (Block 6, Line 0)

This line falls into the Vx category because it is the means by which other units in the system communicate with MU5. These have write only access and use this to put information into the 32 bit line. This causes an interrupt in MU5 which allows the information to be read.

### The Parity V-Store (Block 7)

MU5 RPPF is bit 63 of line 0 of MU5 V-store block 7. It has read/write access from all units in the system including MU5. When any bit in the exchange Vx line UPF becomes set (see 9.4), exchange sends a signal to every unit in the system. This is ignored if the unit has the manual 'ignore parity' switch set. If not, and the RPPF bit is set to '1' then no further accesses are permitted from MU5 to exchange until appropriate action is taken.


In normal circumstances this is a complete description of MU5 Vx-store. However, if the REMOTE switch is set in MU5 Console V store (Line 12 bit 54 see 8.5) then blocks (2 -> 5) of MU5 V store becomes available as Vx-store in addition to the two lines just described. The MU5 V-store is completely described in chapter 8. All V lines in blocks 2, 3 and 5 become treated as Vx lines with the same access as before. However, block 3 the Console V-store (8.5) does have some changes on access. Lines 10, 11 and 12 as V store have only READ access. The permissible access as Vx lines is READ/WRITE with the exception of bits 54 and 55 of line 12 which stay READ only.

Access to the Peripheral Window and MU5 RPPF is unaltered on REMOTE.

9.3 The Disc (Drum) Vx-Store

The Disc Vx lines exist in one block (Block 0) which consists of 32 lines of 64 bits.

<u>Address</u> (64-bit word)	<u>Name</u>	<u>Size</u>	<u>Access</u>
0	DISC ADDRESS	22	R/W



P is the internal read request bit and if set (i.e., = '1') overrides bit 33.

(Therefore normally a zero).

R/W specifies whether reading from or writing to the disc. ('1' = READ).


D specifies the disc number 0 -> 3. Each has 64 bands containing 37 blocks of 256 words (32 bits + 4 parity bits).

SIZE specifies the number of blocks requested for transfer.

Writing to this line initiates a disc transfer.

Note The block and size digits are updated during a transfer.

1	STORE ADDRESS	29	R/W
---	---------------	----	-----



The 28 bits specify the real address in what is the receiving or sending unit. Hardware ignores the l.s. eight bits of the address; it is assumed to point to at least a 256 word block boundary (minimum transfer size). This line is altered during a transfer.

DISC STATUS 32 (see below)

The description below is of the status bits when set = '1'.

Digit

32	-	Decode the rest of the status line. This is examined by software on completion of each transfer. If set, set, some further action is required. This digit is reset by writing a '1' to it.
33	-	Decode Vx line 7. This indicates an operators request to go onto SELF TEST. Further information about the request is held in line 7.
34 - 41	-	Eight bits reserved for discs 2 & 3 having the same significance as bits 42 -> 49.
42	-	Disc 1 absent. Set manually to indicate the disc is off line i.e., cannot be read from or written to by a CPU.
43 - 44	-	Spare
45	-	Disc 1 on Self Test.
46	-	Disc 0 absent.
47 - 48	-	Spare
49	-	Disc 0 on Self Test.
50	-	Illegal request to the disc.
51 - 52	-	When input parity error occurs these bits define whether it occurred in data, address or control information.
53	-	(PFO). Input parity error. This causes a bit to be set in the exchange Vx line UPF. Resetting of both bits is achieved by writing a '1' to this bit.
54	-	Bound locked out (see line 5).
55	-	Data late (Hardware error).


- 56 - Column parity error (internal to disc).
- 57 - Row parity error (internal to disc).
- 58 - Ignore parity fault - applies to 'input parity error' (bit 53) only.
- 59 - End transfer. '1' = ENDED.
- 60 - 63 - Disc unit number (= 0).

Access

All bits of STATUS can be read. Only bits 32, 50, 53, 58 and 59 can be written to. Each is reset by writing a '1' to its bit position.

3


CURRENT POSITIONS 24 R



This line gives the current positions of each of the discs, and also records which packing density is current ('0' - HALF P.D., '1' - FULL P.D.).

4

COMPLETE ADDRESS 28 R/W



This line holds the address to be written to on disc transfer complete. The information written is the STATUS line 2. This address must not be a disc address.

5

LOCKOUT 01 32 R



This contains 16 lockout switches for each of discs 0 and 1. These are set manually. Each bit locks out 4 bands.

6 LOCKOUT 23 32 R  
As for line 5 applied to discs 2 and 3.

7 REQUEST SELF TEST 5




Digits

- 59 -> 60 Reserved for Request Self Test on discs 2 and 3.
- 61 'Request self test' on disc 1.
- 62 'Request self test' on disc 0.
- 'Request self test' is set manually.
- 63 'CPU permission to self test'.

Access


All bits can be READ. Only bit 63 can be written to.

8 SELF TEST COMMAND 7 R/W



Digit

- 57 - 59 Margins
- 60 A/D Required
- 61 Self Test
- 62 - 63 Disc number



A/D holds the A/D conversion value. The current BAND, BLOCK and TRACK is maintained.

Digit

59	-	Max/Min Signal
60	-	Print A/D
61	-	Surface Error
62	-	Address Error
63	-	Self Phasing Error

9.4 The BTU Vx-Store

The Block Transfer Unit is designed to perform autonomous block transfers between six possible stores in the system (2 mass stores and 4 local stores). Up to 4 block transfers may be specified concurrently and these are carried out on an equal priority, time-shared basis.

A block transfer from mass to local for instance is carried out one word at a time through the exchange. The word to be transferred is buffered in the BTU before being sent on to the local store. Thus the transfer mass -> local actually consists of two transfers:-


Mass -> BTU followed by  
BTU -> Local

A block transfer has 4 controlling V-lines associated with it. There are 4 sets of these V lines allowing 4 concurrent transfers. Each set is situated in one of the 4 block addresses 0 -> 3 of this units address field. A block of BTU Vx-store consists of 32 lines of 64 bits. Within blocks 0 -> 3 the Vx lines have the following significance:-

BLOCKS 0 -> 3 (Transfer control V lines)

<u>Address</u> (Decimal)	<u>Name</u>	<u>Size</u>	<u>Access</u>
0	SOURCE ADDR	32	R/W

Specifying  
64-bit boundaries



The hardware interprets this address as referring to a boundary that is a multiple of the block size obtained by rounding the transfer size up to the nearest power of 2. In addition the 4 least significant address bits are always interpreted by the hardware as zero (16 word minimum boundary). A transfer of all zeroes (null transfer) is achieved when bit 41 of the source R.A. is set to 1 and the unit no (bits 36 - 39) are 9 (i.e., local).




1

DESTINATION ADDR

32

R/W




This address is interpreted by hardware in the same way as the source address.

2

SIZE

20

R/W



At the start of a block transfer N specifies the transfer size as 2 less than the number of 32 bit words due for transfer. (N/2 will always be odd). The maximum transfer size is 64K and the minimum theoretical size is 2 words.

The transfer is carried out from the final word of the block backwards to the first. Each time a 64-bit word is transferred N is decremented by 2. On completion of the transfer N will be = -2.


Ui is the number of the unit which is to be interrupted on completion of the transfer.

3

TRANSFER STATUS

4

R/W



Bit 60 is the Transfer in Progress bit. Setting this bit initiates a block Transfer. It may be reset by software to terminate the transfer midway. Hardware resets this bit on transfer complete (successful or not).

Bit 61 is the Transfer Complete bit. Hardware sets this bit on completion (successful or not) of the block transfer. This is what causes the Block Transfer Complete interrupt in Ui. The interrupt may be turned off by resetting this bit. Bit 62 is the Parity during Transfer bit. When set this bit indicates that the current transfer has been terminated by hardware because of a parity fault (see B.T.U. Block 4 - Parity V-lines). Bit 63 is a spare fault bit.

BLOCK 4 (PARITY MANAGEMENT)


0	PFO	1	R=Reset
			Bit 63 of this line is set if parity is detected on address or control bits sent to the B.T.U. from exchange. This signal is passed back to exchange and sets a bit in the UPF line in the Exchange Vx-store (see BTU Vx-store Block 5 line 2). Both of these bits are reset by reading the PFI bit only.
1	B.T.U. RIPP	1	R/W
			Bit 63 of this line is a means of inhibiting further requests from the B.T.U. to exchange. When any bit in the UPF Vx line (see block 5 line 2) becomes set as a result of some parity fault exchange sends a signal to each unit in the system. When this signal appears in the B.T.U., provided parity interrupts are uninhibited, the RIPP line if set will stop further requests to exchange.
2	TRANSFER COMPLETE	4	R
			Bits 56 - 59 contain the transfer complete bits for channels 0 -> 3 respectively. N.B. Early morning reset sets these bits to zero.

BLOCK 5 (Exchange Vx Lines)

The exchange does not have the status of a unit and its V-lines are addressed via the B.T.U. The exchange Vx lines constitute block 5 of the B.T.U Vx-store.


0 SUPF 13 R/W=Reset

This line indicates sending unit parity fail. For each transfer through exchange, the data, address and control information is checked for correct parity. If the check fails a bit is set in SUPF corresponding to the unit which sent the information.



Bits 52 -> 63 correspond to wrong parity from units 0 -> 11 respectively. Thus, exchange does not stop the transfer if a parity failure is detected, but merely notes who was responsible for it. Reading this line causes it to be cleared. Bit 51 is known as CAP. If set, it indicates a control or address parity.

1 UPF 13 R



Bit 51 of this line is set if any of the bits of SUPF (line 0) are set. It represents any parity fail on information entering exchange. Bits 52 -> 63 represent any parity signal sent to Exchange from units 0 -> 11 respectively. All units parity check information coming from exchange. Any failure causes a bit to be set in the unit and a signal returned to exchange which sets the appropriate bit in UPF.

Bits 52 -> 63 of UPF are reset by resetting the parity fail bit in the appropriate unit. Bit 51 is reset by reading SUPF.

9.5 The 1905E Vx-Store

This Vx-store consists of four lines in block 0.

<u>Address</u> (Decimal)	<u>Name</u>	<u>Size</u>	<u>Access</u>
0	VXINT	1	W
			Writing to this line interrupts the 1905E at the next-occurring Normal-Mode instruction-fetch time. (The value of bit 63 is irrelevant). Writing to this line also sets bit $2^{20}$ in the 1905E's internal V-line 129, (known as SR129).
1	5ERIPF	1	W
			Writing a 0 to bit 63 of this line <u>resets</u> the 5ERIPF flip-flop, (thus allowing requests through Exchange in the event of a parity fail). Writing a 1 to bit 63 <u>sets</u> 5ERIPF, (thus inhibiting requests in the event of a parity fail). 5ERIPF appears as bit $2^{21}$ in the 1905E's internal V-line 129.
2	RPS	1	W
			Writing to this V-line <u>resets</u> the 1905E's PFO flip-flop. (The value of bit 63 during writing is irrelevant). PFO, which indicates a parity fail detected by the 1905E on information received via Exchange, also appears as bit $2^{18}$ in the 1905E's internal V-line 129.
3	(Spare)	1	W
			(Writing to this line causes no action).

Notes

- a) Writing to higher Vx addresses causes the address to be decoded modulo 4.
- b) Incoming Read, and Read-andMark requests are ignored by the 1905E, except that 'Buffer Free' signals are returned to Exchange. Note that the 'Store Free' flip-flop does not exist in Exchange for the 1905E.
- c) In order to aid system development, there exist further signals between the 1905E and MU5. The following is a list of relevant bits in the 1905E's internal V-Line 129.

<u>digit in V-line 129</u>	<u>meaning</u>
$2^1$	Advance Warning of Power Failure - (similar to bit 49 in MU5 System Error V-line).
$2^2$	MU5's Remote switch on/off.
$2^3$	Exchange operable/inoperable - (similar to MU5 V-line 24 in block 4).
$2^{15}$	Allow/Inhibit MU5 communication - (a manual switch on the 1905E).
$2^{16}$	Unit fail - (similar to the Exchange Overdue signal at bit 60, MU5 V-line 23 in block 4).
$2^{17}$	DINTO = (Diagnostic Interrupt Outwards), this is set by writing to MU5 V-line 25 in block 4.
$2^{22}$	E.T.U. End-of-Transfer interrupt.

The 1905E also sends a signal which appears in bit 63 of MU5 V-line 24, block 4, to indicate that the 1905E is operational. Finally, the 1905E software may produce a signal DINTI - (Diagnostic Interrupt Inwards) - which interrupts MU5 at Level 0, and appears as bit 54 in the System Error V-line. (As yet not implemented).


Further description of the 1905E/MU5 interface may be found in the relevant 1905E documentation.

9.6 The Local Store Vx-Store

This Vx-store consists of 5 lines in Block 0.

<u>Address</u> (Decimal)	<u>Name</u>	<u>Size</u>	<u>Access</u>
0	PFO	1	R/W=Reset
<p>Bit 63 is set when a parity failure occurs on incoming addresses or control bits from exchange. This results in a bit being set in exchange V line UPF. Both bits are reset by writing to PFO.</p>			

8	FAILSOFT	8	R/W
---	----------	---	-----



Digit

56 -> 59

These bits can be set to specify a failsoft mode. (See below).

60 -> 63

These bits are set manually to indicate a stack (0 -> 3 respectively) is OFF LINE. Writing to these bits has no effect.


Bits 56 - 59 of Line 8

These may be set to 11 meaningful numbers. Each number corresponds to a mode of operation of the Local Store, as below:-


<u>NUMBER</u>	<u>MODE</u>	0	1	2	3
0	Normal - all 4 stacks interleaved in order	0 1 4K	2 3 4K	4 5 4K	6 7 4K
4	Non-interleaved - Each stack contains 4K of sequential addresses, in order	0- 4K	4- 8K	8- 12K	12- 16K
5	Non-interleaved - Each stack contains 4K of sequential addresses, in order	8- 12K	12- 16K	0- 4K	4- 8K
6	Non-interleaved - Each stack contains 4K of sequential addresses, in order	4- 8K	0- 4K	12- 16K	8- 12K
7	Non-interleaved - Each stack contains 4K of sequential addresses, in order	12- 16K	8- 12K	4- 8K	0- 4K
8	Interleaved in pairs, stacks 0 & 3 and 1 & 2, addressed in order	0 1 8K	0 1 16K	2 3 16K	2 3 8K
9	Interleaved in pairs, stacks 1 & 2 and 0 & 3, addressed in order	0 1 16K	0 1 8K	2 3 8K	2 3 16K
10	Interleaved in pairs, stacks 0 & 2 and 1 & 3, addressed in order	0 1 8K	0 1 16K	2 3 8K	2 3 16K

9.6.3


11 Interleaved in pairs, stacks 1 & 3  
and 0 & 2, addressed on order



12 Interleaved in pairs, stacks 0 & 1  
and 3 & 2, addressed in order



13 Interleaved in pairs, stacks 3 & 2  
and 0 & 1, addressed in order



Modes 1, 2 and 3 also give normal interleaving.

Suitable precautions must be taken to ensure that the setting of this line is not changed unless the Local Store is completely inactive.




16  
( $\frac{9}{10}$ )

SELF TEST CONTROL

11

R/W

Digit


53 - 54	Margins - these bits specify three states as follows:- 00 - Normal. 01 - Normal 10 - Inverse of console switches 11 - Obey console switches
55	End Action - '0' - Continuous self test '1' - One cycle of the stack
56 - 57	Fixed Address Bit These bits specify the range of 64-bit addresses to be tested on this stack. 00 - All 64 bit words 01 - Only the odd 64-bit words 10 - Only the even 64-bit words
58 - 60	Pattern for testing
61	Function '0' - Clear Write '1' - Read Restore
62 - 63	Stack on test

24  
( $\frac{9}{13}$ )

READ ADDRESS

12

R



This allows the current self test address to be read (e.g., after stop on error). The address is in the form of 12 bits referencing a 64-bit word in the stack on test. The stack in question must already be on self-test.

32  
( $\frac{7}{8}$  20)

PST

5

See Below



Digits

59 - 62 READ ONLY Fault bits - one per stack

63

R/W - this bit initiates and terminates  
self test when set = 0.

9.7 Mass Store Vx-Store

<u>Address</u>	<u>Name</u>	<u>Size</u>	<u>Access</u>
0	POWER STACK 0	8	R/W
1	POWER STACK 1	8	R/W
2	POWER STACK 2	8	R/W
3	POWER STACK 3	8	R/W

These first 4 Vx lines define the operation of power supplies in the stacks 0 -> 3 with the following format:-




<u>digits</u>	
56 -> 57	Power Supply 0
58 -> 59	Power Supply 1
60 -> 61	Power Supply 2
62 -> 63	Power Supply 3

Each stack has 4 power supplies and the digits above define whether these have to work at nominal values or on increased or reduced margins.

<u>Value</u>		
00	-	Nominal
01	-	Reduced margin
11	-	Increased margin

4 OFF LINE STACKS 4 R



These bits when set to '1' define a stack to be OFF LINE. This is achieved manually.

5 WORKING STACKS 4 R/W

This line has the same format as line 4. Each bit when set specifies a stack to be working normally. Reset to put a stack on test.


6 STACK ON TEST 2 R/W

Bits 62 and 63 of this line define which of stacks 0 to 3 is on self test.

7 ON TEST INDICATOR 1 R/W

Setting bit 63 initiates self test on the stack specified in line 6. To reset testing after a stoppage, stop on error, for example, OT1 must first be reset to zero, then set back to '1'.

8 SELF TEST CONTROL 6 R/W



Before putting any stack on self test, patterns (58 -> 60) and operating modes (61 -> 63) can be written to this line.

9

END ACTION

3

R/W

Bit 61 is SEQ and specifies only one cycle of all addresses of the stack on test ('1') or continuous cycling ('0').

Bit 62 is ST and for the stack on test specifies 'stop at end of current store cycle' ('1').

Bit 63 is CON and specifies (for the stack on test) 'stop on error' ('0'). To continue after stop on error this bit can be reset to '1' then set back to '0' again.

10

TRANSFER ADDRESS

17

R

(%A)

This line holds the address at which an error occurred when on self test.

11

TRANSFER DATA OUT

36

R

(%B)

This line may be read to obtain 9 bits, the data involved in an error on self test. Data is regarded as 36 bits, (32 data bits + 4 parity bits). This is conceptually divided into 4 sets of 9 bits. When line 11 is read, line 17 defines which set of 9 bits is to be read.


12

PARITY FAIL

6

R=Reset

(%C)



<u>Digit</u>	<u>Parities</u>	
58 (PFV)	-	V store I/P data
59 (PFO)	-	Stack 0 I/P or O/P data
60 (PF1)	-	Stack 1 I/P or O/P data
61 (PF2)	-	Stack 2 I/P or O/P data
62 (PF3)	-	Stack 3 I/P or O/P data
63 (PFAC)	-	Address or Control bits from Exchange
57 (INVAL)		INVALID ADDRESS sent to interface i.e., to high for number of stacks working

If any of these bits is set a signal is sent to exchange and to the Engineers Door and Console.

13                    STACK 0 STATUS                    3                    R  
(%D)

14                    STACK 1 STATUS                    3                    R  
(%E)

15                    STACK 2 STATUS                    3                    R  
(%F)

16                    STACK 3 STATUS                    3                    R  
(%10)

The formats of these 4 lines are the same.

Digit

61            -            Power Supply on margin  
62            -            Stack on self test  
63            -            Error during self test

17                    TRANSFER PART                    2                    R/W  
(%11)

This specifies which section of the data or address word is to be read when reading TDO or TA lines 24 and 26.

'0' refers to the most significant section.

'3' refers to the least significant section.

9.8 The System Performance Monitor Vx-store

The System Performance Monitor Vx store consists of the following:-

System Performance Monitor V-line

This line is accessed by using a real mode descriptor, with V-store specified in the origin. The real address part specified is irrelevant. The top 32 bits of the V-line are write only (if read, they return 0). The bottom 32 bits are read only.

BIT	0	Enable bits 1, 2, 3	
	1	RC1	Resets and initiates histogram logic
	2	RCVAL	Overrides validation lines for histogram input pulses
	3	INTOFF	Turns off interrupt
BIT	4	Enable bits 5, 6, 7	
	5	RCNT1	Initiates fast counters
	6	RCNTVAL	Overrides validation lines for fast counters
	7	RCNTR	Resets fast counters
BIT	8	Enable bits 9, 11	
	9	RAIO	Identifies the activity required for software monitoring
	10	Not used	
	11	RAI1	Used as RAIO
BIT	12	Enable bits 13, 14, 15	
	13	RIMA	} Control the input mode on the two histogram input channels
	14	RIMB	
	15	RDYW	Allows Dwell Histogram logic to continue upon overflow of Y-Counter
BIT	16	Enable bits 17, 18, 19	
	17	RSP2**4	} Set the scale factor on the Prescaler
	18	RSP2**2	
	19	RSP2**1	

BIT 20		Enable bits 23, 24, 25, 26, 27, 28, 29, 30
21		Not used
22		Not used
23	RSAC	Allows CPU to write to or read the real store via Exchange
24	RSAD	Initiates the display of the store contents on the VDU
25	RSAIN	Interval Histogram Mode
26	RSAIDH	Increment/Decrement Histogram Mode
27	RSARCH	Clears Title store
28	RSADH	Dwell Histogram mode
29	RSADAD	Direct Addressing mode
30	RSARD	Clears Data store
27 & 30		Clears title store and data store
28 & 29		Snap-shot Mode

NOTE: Except for those pairs noted, only one bit in the above group should be set at a time.

BIT 31	PFLR	Turn off parity fail indicator
BITS 32 -> 59		Not used
BIT 60		Not used
BIT 61	RM*C	Monitor is in Manual Mode
BIT 62	RIGNEX	Monitor has ignored an Exchange request
BIT 63	GIGNEX	Monitor is ignoring Exchange except for V-reads i.e., Monitor is in one of the following modes: MANUAL, DISPLAY, CLEAR-STORE



System Performance Monitor Real Store

This consists of 3 blocks of store addressed as follows:-

<u>Address</u> (64-bit word)	<u>Name</u>	<u>Size</u>	<u>Access</u>
0	SPM.DATA.STORE	256 x 16 bits	R/W
	To read/write from this store bits 48 - 63 of the highway are used.		
512	SPM.FAST.COUNTERS	16 x 32 bits	R
	To read these counters bits 48 - 63 of the highway are used. The more significant half of a fast counter is at the address 512 + 7 + 16*N, the less significant half is at the address 512 + 15 + 16*N where N is the fast counter required.		
768	SPM.TITLE.STORE	256 x 7 bits	R/W
	To read/write from the store bits 49 -> 55 of the highway are used.		



## Chapter 10 The Basic Programming Language - XPL

XPL compilers exist for MU5 and the 1900. Compatibility between the compilers has been aimed for, however, there are unavoidable differences.

Since the MU5 compiler will be used more extensively than the 1900 version, this description applies to the MU5 version, any differences between the compilers will be mentioned.

### Table of Contents

The Metalanguage . . . . .	10.1
Program & Statements . . . . .	10.2
Names, Literals & Labels . . . . .	10.3
Tables & Texts . . . . .	10.4
Blocks . . . . .	10.5
Declaratives . . . . .	10.6
(1) Variable Declarations	
(2) Literal Declarations	
Instructions . . . . .	10.7
(1) Computational	
(2) Store to Store	
(3) Organisational	
(4) Conditional	
Procedure Call Facilities . . . . .	10.8
Special Directive Statements . . . . .	10.9
Alternative Punching Conventions for the VDU's . . . . .	10.10
The Test Bits in MS . . . . .	10.11

### 10.1.1 The Metalanguage

Modified BNF (Backus Naur Form) is used to define any XPL syntactic element.

The modifications to the BNF are as follows.

(1) In the order of alternatives and of elements within alternatives:-

- a. Any alternative which is a stem of another comes after it.
- b. If one alternative is a special case of another, it must come first.
- c. In recursive definitions, there must be at least one left-most element not recursive.

(2) Metalinguistic Bracketing:-

Several alternatives may be specified as an element of another by enclosing them in square brackets.

10.2.1 Program & Statement

<XPL.PROGRAM> ::= <PROGRAM.OF.A.SEGMENT> [<XPL.PROGRAM> | <NIL>]

<PROGRAM.OF.A.SEGMENT> ::= \*SEGMENT<SP><SEGMENT.NO><NL>  
 BEGIN <NL>  
 <program> <NL>  
 END <NL>  
 \*END OF SEGMENT<NL>

The <program> consists of a number of statements, and hence,

<program> ::= <STATEMENT> [<program> | <NIL>]

The statements are:-

<STATEMENT> ::= <LABEL> |  
 <LABEL><SEP> |  
 <TABLE><SEP> |  
 <TEXT><SEP> |  
 <BLOCK><SEP> |  
 <DECLARATIVE><SEP> |  
 <INSTRUCTION><SEP> |  
 <SPECIAL.DIRECTIVE.STATEMENT><SEP> |  
 <SEP>

where <sep>, a separator is:-

<SEP> ::= <NL> | <COMMENT>

and where comment commences with two colons and terminates with a newline. For line continuation purposes  $\pi$ <NL> is ignored.

The seven types of XPL statements are explained in the following sections.

10.3.1 Names, Literals & Labels

Since the basic operands in the language are names and literals, it is convenient to define them first.

(1) Names & Literals

```

<LITERAL> ::= <DECIMAL> |
              % <HEX.DIGITS> |
              ½ <CHARACTER.STRING> ½ |
              <NAME> |
              <DR.LIT>

<DECIMAL> ::= [+|-|<NIL>] <INTEGER> [ . <INTEGER> | <NIL> ]
<INTEGER> ::= <DECIMAL.DIGIT> [ <INTEGER> | <NIL> ]
<HEX.DIGITS> ::= [ <HEX> | <HEX> ( <INTEGER> ) ] [ <HEX.DIGITS> | <NIL> ]
<HEX> ::= ø | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | A | B | C | D | E | F
<CHARACTER.STRING> ::= [ <CHARACTER> | <HEX.PAIR> ] [ <CHARACTER.STRING> |
                                <NIL> ]

<HEX.PAIR> ::= <VB> <HEX> <HEX> <VB>
<VB> ::= a vertical bar
<NAME> ::= <LETTER> | <LETTER> <NAME.SYMBOLS>
<NAME.SYMBOLS> ::= [ <LETTER> | <DECIMAL.DIGIT> | . ] [ <NAME.SYMBOLS> |
                                <NIL> ]
<DR.LIT> ::= D <QUA> <OCT> <OCT> / [ <INTEGER> | <NAME> ]
              / [ <INTEGER> . <INTEGER> [ . <INTEGER> | <NIL> ] |
              <NAME> ]

<QUA> ::= ø | 1 | 2 | 3
<OCT> ::= ø | 1 | 2 | 3 | 4 | 5 | 6 | 7

```

(2) Labels

A label takes the form of a name followed by a colon, i.e.,

```
<LABEL> ::= <NAME>:
```

Any number of labels may precede an instruction, and references to labels are optimised by XPL.

(3) Further Notes on Literals

There are five different types of literals, and each of them is explained below. The XPL compiler will attempt to optimise on their lengths.

- a. DECIMAL NUMBER (not currently implemented on MU5) in which a decimal point implies that the number is to be coded as a floating-point literal. All decimal numbers are signed.

Examples:-

## Signed fixed-point:-

2	::	6-bit
51	::	16-bit
+51	::	16-bit (the positive sign is optional)
327680	::	32-bit

## Floating-point (all 64-bit):-

1.0	::	1. will be faulted
-0.23	::	-.23 will be faulted

- b. BINARY LITERAL is a % (percent) sign followed by a string of hexadecimal digits. Right justification is adopted.

Examples:-

%1F	::	6-bit signed
%2F	::	16-bit unsigned (6-bit is impossible, since the sign bit will be propagated)
%F903	::	16-bit unsigned
%AF903	::	32-bit unsigned
%8000000000004F903	::	64-bit
%80(10)4F903	::	gives the same literal as in the above example

- c. CHARACTER LITERAL is a string of characters enclosed by single characters which represent double quotes. Hexadecimal pairs can be used to represent characters not available on the input device. A compiled literal is packed eight bits per character (in ISO code) and right justified.

Examples:-

$\frac{1}{2}$ ABCDE $\frac{1}{2}$	:: Double quotes is represented by $\frac{1}{2}$
	:: on most Flexowriters
$\frac{1}{2}$ ABCDE OA XYZ $\frac{1}{2}$	:: The hexadecimal pair  OA  is
	:: compiled as a newline character

- d. NAME LITERAL can either be a name declared as a literal, or the name of a label, in which case the absolute address is used.

Examples:-

MAX.VALUE	= 100
MASK	= %FCF
ENTRY.POINT	= $\frac{1}{2}$ START $\frac{1}{2}$
NINE	= 9

BASE:

- e. DR LITERAL is a descriptor literal. In XPL, the bound field can either be a previously defined name or an integer, whereas the origin field can be either a name or 'SEGMENT.WORD.BYTE', which specifies the address of the start of the string (in the absence of the '.BYTE', the byte position is taken to be zero).



Examples:-

D033/19/8196.74.3

:: This is a vector descriptor (type 0), defining  
 :: a string of 8-bit elements. Modifier is not  
 :: scaled and there is no bound check. There are  
 :: 19 elements, starting at byte position 3 of line  
 :: 74 in segment 8196.

D260/NINE/BASE

:: This is a descriptor (type 2), defining a  
 :: string of 64-bit elements. Modifier is scaled  
 :: and there is bound check. There are 'NINE'  
 :: elements, starting at the label 'BASE'.  
 :: (The name literal 'NINE' must be declared  
 :: prior to this DR literal, whereas, the label  
 :: 'BASE' can be a forward reference).

[N.B. When a vector is accessed by using a name corresponding to a descriptor literal, two orders will be compiled, e.g.,

A = DESCRIPTOR.1[B]

will be compiled into:-

D = DESCRIPTOR.1

A = D[B]

10.4.1 Tables & Texts

Table is used to plant literals within the compiled code, whereas text is used to plant a string of symbols.

(1) Tables

<TABLE> ::= DATAVEC <SP> <NAME> (<LENGTH>) <NL> <LIT.LIST> <NL> END

where <NAME> ::= a literal descriptor set up to access the content of the TABLE.

<LENGTH> ::= an <INTEGER> specifying the length of the literals in the TABLE. (The length can be 1, 4, 8, 16, 32 or 64-bit).

<LIT.LIST> ::= <LIT.LINE> <NL> <LIT.LIST> | <LIT.LINE>

<LIT.LINE> ::= <LIT.ITEMS> [<LB> <INT> <RB> | <NULL>]

<LIT.ITEMS> ::= <LITERAL>, <LIT.ITEMS> | <LITERAL>

Where <integer> indicates the number of times that the preceding items on this line are to be repeated. Nested repetitions are not allowed.

Example:-

```
DATAVEC TABLE.1 (64)
    99
    %14A76F02F
    ½AB½
    -1,[5]                :: -1 to be planted 5 times
    4,FRED,[2]           :: 4, FRED to be planted twice
    0
END
```

(2) Texts

<TEXT> ::= DATASTR <SP> <NAME> ½ <CHARACTER.STRING> ½

where <NAME> ::= a string descriptor set up to access the <CHARACTER.STRING>

Example:-

```
DATASTR CAPTION.1 ½ **FIX LINEPRINTER** ½
```

10.5.1 Blocks

```

<BLOCK> ::= [BEGIN|PROC<SP><PROC.NAME>][(<LABEL.LIST>)|<NIL>]<NIL>
          <program>
          END

```

where <PROC.NAME> ::= the <NAME> of the procedure.  
 <LABEL.LIST> ::= <NAME>[,<LABEL.LIST>|<NIL>]

Basically, a block consists of two declaratives BEGIN and END, and it serves to define the scope of labels. Declarations other than labels have a global scope equivalent to that of a forward reference. Redefinitions of names within this global scope are not allowed. Blocks can be nested to any depth.

If PROC is used instead of BEGIN, a jump instruction is planted by XPL to jump round the procedure.

If the BEGIN is followed by some names of labels, e.g., BEGIN (L31, L32), then entries to that block can be made to these labels from the enclosing block.

Example:-

```

BEGIN
  L1: - - -
      -> L31
      - - -
      -> L32
      - - -
      BEGIN (L31, L32)  :: implying L31 and L32 are in the
                       :: same level as L1 and L2.
                        - - -
                        L31: - - -
                        - - -
                        L32: - - -
                        - - -
      END
  L2: - - -
END

```

### 10.6.1 Declaratives

There are two types of declaratives both of which give the name a global scope, namely:-

- (a) variable declaration <VAR.DEC>, and
- (b) literal declaration <LIT.DEC>.

<DECLARATIVE> ::= <VAR.DEC> | <LIT.DEC>

#### (1) Variable Declarations

The variable declarations assign a name to a displacement relative to a base, which can either be NB, XNB, SF, 0 or STK (for accessing the stack).

<VAR.DECL> ::= V[32|64|V]/[NB|XNB|SF|0|STK]<VAR.SPEC>  
 <VAR.SPEC> ::= <NAME>:<DISPLACEMENT>[,<VAR.SPEC>|<NIL>]  
 <DISPLACEMENT> ::= [-|<NIL>]<INTEGER>| $\%$ <HEX.DIGITS>

A name declaration must start with a V, followed by the size of the variable. The size is either 32- or 64-bit, or a V indicating that the variable will be used in privileged mode to access a V-store location.

#### Example

V32/NB FRED:3           :: FRED can be found three  
                           :: 32-bit words away from  
                           :: NB and is of size 32-bit.

V32/STK TOP.32.bits:0   :: The displacement must be  
                           :: zero with STK.

Notes

1. The hardware of the machine treats all V store as 64-bit quantities. This means that VV quantities are equivalent to V64 quantities.
2. The position of variables in MU5 are governed by a base address contained in one of the registers NB, XNB, SF and a displacement. It should be noted that the value of the Base address and the value of the displacement cannot be freely interchanged as may have been expected.

Thus:-

```
V64/NB FRED.1:6      :: sets up a Base of 0
NB = 0                :: and a displacement of 6
```

and:-

```
V64/NB FRED.2:0      :: sets up a Base of 6
NB = 6                :: and a displacement of 0
```

will not access the same location in the store. This is because the base registers always count in units of 32 bits whereas the displacement counts in units of the size of the variable, in this case 64 bits. The example therefore says that zero units of 32 bits plus 6 units of 64 bits is not the same as zero units of 64 bits and 6 units of 32 bits.

(2) Literal Declarations

A literal declaration assigns a value to a name. When the name appears as an operand, its value will be coded as a literal in the instruction. (As already shown in (3) of section III, XPL will optimise on the lengths of the literals).

<LIT.DECL> ::= L/<LIT.SPEC>

<LIT.SPEC> ::= <NAME>=<LITERAL>[,<LIT.SPEC>|<NIL>]

Examples:-

```
L/          MAX.VALUE = 100
L/          MASK = % 44420F
L/          NAME.STRING = ½ACCUMULATOR½
```

[N.B. a. A number of declarations can be placed on a line, e.g.,

```
V32/SF      VAR.0:0, VAR.1:1, VAR.2:2, VAR.3:3
L/          MAX.VALUE = 100, MIN.VALUE = -100
```

b. Newline or a comment would terminate the sequence, hence,

```
V32/SF VAR.0:0, VAR.1:1, VAR.2:2, VAR.3:3
V32/SF VAR.4:4, VAR.5:5
```

would have to be written. ]

10.7.1 Instructions

There are four types of instructions, namely,

- (1) computational <COMPUT>,
- (2) store to store <STS>,
- (3) organisational <ORG> and
- (4) conditional <CONDIT>.

<INSTRUCTION> ::= <COMPUT> | <STS> | <ORG> | <CONDIT>

Each type of instruction is dealt with separately, however, it is more convenient to define the syntax of an operand first.

<OPERAND> ::= <SIMPLE.OPERAND> | <NAME><LB>[B|O]<RB>

where <SIMPLE.OPERAND> ::= <NAME> | <LITERAL>

<LB> ::= left square bracket

<RB> ::= right square bracket

(1) Computational Instructions

<COMPUT> ::= [<B.ORD> | <X.ORD> | <A.ORD> | <ADD.ORD> | <AEX.ORD>] <OPERAND>

<B.ORD> ::= B[|=|'|\*|=|>|+|-|\*|/|≠|V|<=|&|-:|COMP|CINC]

<X.ORD> ::= [X|S|X][|=|'|\*|=|>|+|-|\*|/|≠|V|<=|&|-:|COMP|CONV|/:]

<A.ORD> ::= [AFL|A][|=|'|\*|=|>|+|-|\*|/|≠|V|<=|&|-:|COMP|CONV|/:] |  
 [AU|AX][+|-|\*|/|≠|V|<=|&|-:|COMP] |  
 [ADC|AD][<=|COMP|CONV]

<ADD.ORD> ::= AGD[|=|'|\*|=|>|COMP]

<AEX.ORD> ::= AEX[|=|'|\*|=|>]

IMPORTANT:-

WHERE THE OPERATOR -: is reverse subtract  
 and /: is reverse divide

17-4-73/1

(2) Store to Store Instructions

This group of orders uses the secondary operand unit, and they operate on strings.

```

<STS> ::= <FN.1><OPERAND> |
          <FN.2><SIMPLE.OPERAND> |
          SUB1 <NAME> |
          SUB2

<FN.1> ::= D=|D*=|DC=|XD=|XDC=|STACK
<FN.2> ::= D=>|XD=>|DB=|XDB=|
          MOD|RMOD|SMOD|XMOD|MDR|XCHK|
          BMVE|BMVB|BCMP|BLGC|BSCN|
          SMVE|SMVB|SCMP|SLGC|SMVF|TALU|TCHK|TRNS

```



(3) Organisational Instructions

This group of orders defines internal register operations.

```

<ORG> ::= RETURN |
        [EXIT|JUMP|XJUMP|STKLINK]<OPERAND> |
        SETLINK<SIMPLE.OPERAND> |
        <MS.ORD> |
        <XC.ORD> |
        <SF.ORD> |
        <NB.ORD> |
        <XNB.ORD> |
        <MESC.ORD> |
        <DUMMY.ORD>

<MS.ORD> ::= MS = <OPERAND>
<XC.ORD> ::= [XC0|XC1|XC2|XC3|XC4|XC5|XC6]<OPERAND>
<SF.ORD> ::= SF[=|+=|NB+]<OPERAND> |
            SF => <SIMPLE.OPERAND>
<NB.ORD> ::= NB[=|+=|SF+] <OPERAND> |
            NB => <SIMPLE.OPERAND>
<XNB.ORD> ::= XNB[=|+] <OPERAND> |
            XNB => <SIMPLE.OPERAND>
<MESC.ORD> ::= [SN =|DL =|SPM =]<OPERAND>
<DUMMY.ORD> ::= D[1|2|3|4]
                (for coding up dummy organisational orders)

```

The XJUMP order (MU5 only) will search the Common Procedure Name List for the name and plant an absolute jump to it.

On the 1900 XJUMP will cause LB to be abutted to the front of the operand. The user is then to use the \*NON LOCAL NAMES 'N8221' option to present the entry address for the library calls.

(4) Conditional Instructions

This group of orders deals with control transfers and the setting of the BBOOLEAN, BN.

```

<CONDIT>      ::= <JUMP.SPEC> <NAME> |
                IF <COND>, <JUMP.SPEC> <NAME> |
                BN <B.FN> IF <COND> |
                BN <B.FN><OPERAND>
<COND>        ::= =0|≠0|<0|<0|>0|>0|
                OV|BN
<B.FN>        ::= /|
                ≡|≠|
                =|≠|
                &|&|/|&|&|/|
                V|V|/|V|/V|
<JUMP.SPEC>   ::= [<LONG>|<SHORT>|<NULL>] ->
<LONG>        ::= +
<SHORT>       ::= -

```

The jump instructions (->) are relative; and XPL will compile the optimum code.

As indicated by the + or - preceeding the jump either a long (32-bit) operand or a short (6-bit) operand is assumed. The default option of NULL will result in a 16-bit operand.

10.8.1 Procedure Call Facilities

There are 6 types of procedure calls available:-

CALL            ::: Plants a relative jump  
 ACALL           ::: Plants an absolute jump  
 XCALL           ::: The MU5 Compiler will look the procedure  
                   ::: name up in the common procedure name list  
                   ::: and plant a jump to it  
                   ::: The 1900 compiler will abutt LB to the  
                   ::: operand. See XJUMP

CALL <PROC.NAME>(<LINK>, <PARAMETERS>)

Where <PARAMETERS> is any operand permitted after the instruction STACK.

The above will be compiled into:-

```

      STKLINK <LINK>
      STACK PARAMETER.1
      STACK PARAMETER.2

      STACK PARAMETER.N
      -> <PROC.NAME> (MU5)  -> LB<PROC.NAME>
  
```

10.8.2

ENTER            :: Plants a relative jump  
AENTER (or SCALL) :: Plants an absolute jump  
XENTER            :: The MU5 Compiler will lock the procedure  
                  :: name up in the common procedure name list  
                  :: and plant a jump to it  
                  :: The 1900 compiler will abutt LB to the operand

However there must be at least 1 parameter

ENTER <PROC.NAME>(<PARAMETERS>)

The above will be compiled into:-

STKLINK L1  
STACK PARAMETER.1  
STACK PARAMETER.2  
.  
.  
STACK PARAMETER.N  
JUMP    <PROC.NAME> (MU5) JUMB LB<PROC.NAME>(1900)  
L1:

10.9.1 Special Directive Statements

\*SEGMENT<SP><EXECUTE.SEG.NO>[,<COMPILE.SEG.NO><NL>|<NL>]

1905 ONLY

where <EXECUTE.SEG.NO> has a value of -1 or 1 to  $2^{14}-1$

and <COMPILE.SEG.NO> has a value of 1 to  $2^{13}-1$

:: <EXECUTE.SEG.NO> specifies the segment  
 :: in which the code is to be executed.  
 :: <COMPILE.SEG.NO> specifies the segment  
 :: in which the code is to be compiled,  
 :: if this is unspecified the compiler  
 :: will select the next available segment.  
 :: If <EXECUTE.SEG.NO> is equal to -1,  
 :: the compiler will select the execution  
 :: and compilation segment numbers.

\*SEGMENT<SP><EXECUTE.SEG.NO><NL>

1900 ONLY

where <EXECUTE.SEG.NO> is an integer in the range 0 -  $2^{14}-1$

:: This specifies the segment in which the  
 :: the code is to be executed.

\*LINE<SP><LINE.NO><NL>

:: Specifies the line (in 16 bit quantities)  
 :: within the segment, at which the next  
 :: instruction is to be planted.

\*END

:: Prints a list of unmatched references  
 :: on the currently selected output stream  
 :: sets the CTL information and returns.

\*PRINT ON

:: Control the listing of the program text  
 :: and compiled code.

\*PRINT OFF

\*NL<SP><NLADDR>

:: <NLADDR> is the address in 32 bit words of  
 :: the LIBRARY NAME LIST.

\*POPI<NL>

1900 ONLY

:: Print out planted instructions  
 :: (giving a hexadecimal dump of the  
 :: code compiled so far, or since the  
 :: last \*POPI).

\*NAME LIST<NL> 1900 ONLY       :: Prints the NAME list.  
  
 \*MAP ON<NL>       1900 ONLY       :: Controls the printing  
 \*MAP OFF<NL>      1900 ONLY       :: of the compile map  
  
 \*INCLUDE FILE '<FILE,NAME>'       :: XPL will now compile from the  
                   1900 ONLY        :: file specified.  
  
 \*RETURN           1900 ONLY       :: Causes XPL to return to the previous  
                                       :: input stream after the \*INCLUDE FILE CMD.  
  
 \*FILE CODE IN '<FILE,NAME>'       :: Creates the files containing the code  
                   1900 ONLY        :: and name list respectively providing  
 \*FILE NAME LIST IN '<FILE,NAME>'   :: the compilation is successful,  
                   1900 ONLY  
  
 \*NON LOCALS '<FILE,NAME>  
 <NAME>,<NAME>,<NAME>, etc.'  
                   1900 ONLY        :: Causes the names specified to be  
                                       :: added to the compilers namelist.  
  
 \*FILE ANYWAY      1900 ONLY       :: Causes the name list and code to be  
                                       :: filed even if the program is faulty.  
  
 \*REF ON 1900 ONLY        :: Causes all the unmatched references  
 \*REF OFF           1900 ONLY       :: to be printed at the end of each block

10.10.1 Alternative Punching Conventions for the VDU's

		<u>Paper Tape</u>	<u>VDU</u>
(1)	Not Equal	$\neq$	$\neq$
(2)	Equivalent	$\equiv$	$\equiv$
(3)	Not Equivalent	$\not\equiv$	$\neq$
(4)	Greater than or Equal to	$\geq$	$\geq$
(5)	Less than or Equal to	$\leq$	$\leq$

PUBLICATIONS RELATING TO M.U.5.

1. "Integrated Circuits for Large-scale Computers": Aspinall & Edwards.  
I.E.E. Colloquium, November, 1967.
2. "A System Design Proposal": Kilburn, Morris, Rohl & Sumner
3. "Associative Memories in Large Computer Systems":  
Aspinall, Kinniment & Edwards
- and
4. "An Integrated Associative Memory Matrix":  
IFIP Congress August 1968, Booklet D, Hardware 1.
5. "An Integrated Associative Storage System": Kinniment,  
Knowles & Edwards.  
I.E.E. Conference Publication No. 54, Microelectronics,  
June 1969.
6. "Sequential-state Binary Parallel Adder": Kinniment & Steven  
Proc.I.E.E. Vol. 117, No. 7, July 1970.
7. "Review of High Speed Addition Techniques": Gosling  
Proc.I.E.E. Vol. 118, No. 1, January 1971.
8. "An Implementation of a Segmented Virtual Store":  
Morris & Detlefsen
9. "Information Coding on Magnetic Drums and Discs":  
Whitehouse & Warburton.  
I.E.E. Conference Publication No. 55, Computer  
Science & Technology, July 1969.
- and
10. "Influence of High Level Languages on Computer Design":  
R. A. Brooker
11. "A Virtual Processor for Real Time Operation": Morris &  
Detlefsen  
Software Engineering Vol. 1, 1969 - Academic Press.
12. "The Nature and Benefits of Modular Operating Systems":  
Morris  
INFOTECH Conference - The Fourth Generation, 1970.
13. "Instruction Fetching in High Speed Computers": Sumner  
INFOTECH Conference 1970, Giant Computers,  
published 1971.
14. "Computer Development in Great Britain": Edwards  
Computers and their Future, Llandudno Computer  
Pioneer Conference, July 1970.



15. "Making Hardware Match the Language": Lindsay  
North Holland Publishing Company in the Proceedings  
of the Algol 1968 Implementation Conference -  
J.E.L. Peck (Ed.)
16. "A System Program Generator": Morris, Wilson & Capon  
Computer Journal, August, 1970.
17. "An Experimental Paging Unit": Lavington, Kinniment &  
Knowles  
Computer Journal, Vol. 14, No. 1, February 1971.
18. "Design of Large High-Speed Floating Point Arithmetic Units":  
Gosling.  
Proc.I.E.E. 1971, Vol. 118, pp 493-498.
19. "Design of Large High-speed Binary Multiplier Units":  
Gosling  
Proc.I.E.E. 1971, Vol. 118, pp 499-506
20. "The MU5 Instruction Pipeline": Ibbett  
Vol. 15, No. 1, Computer Journal, February 1972.
21. "Demand Paging in an On-Line Environment": Morris  
Software 71 Conference, Transcript Books 1971.
22. "Operand Accessing in the M.U.5. Computer": Sumner  
Seminaires IRIA, Structure et programmation des  
calculateurs, 1971, I.E.E.E. Conference Boston,  
March, 1971.
23. "System Ideas in M.U.5. : Computer Structures, Past,  
Present, Future": Edwards  
Fall Joint Computer Conference, Las Vegas, November  
1971.
24. "Computer Design and Performance Monitoring": Edwards  
Infotech State of the Art Lecture on Computer Design,  
March 1972.
25. "Simulation and Production Automation for Logic": May &  
Kahn  
Pittsburgh Conference on Modeling and Simulation,  
April, 1972.
26. "The Structure of the MU5 Operating System": Morris,  
Detlefsen, Frank & Sweeney,  
and
27. "The MU5 Compiler Target Language and Autocode":  
Capon, Morris, Rohl & Wilson  
Computer Journal, May 1972.

28. "The Use of Logic Simulation in the Design of a Large Computer System": Kahn & May
  29. "The MU5 Secondary Operand Unit": Standeven, Lanyado & Edwards
  30. "Circuit Technology in a Large Computer System": Kinniment & Edwards
  31. "Communications in a Multi-Computer System": Morris, Frank & Sweeney
  32. "The Compiler Writer's MU5": Capon & Wilson  
and
  33. "Control of the MU5 Instruction Pipeline": Ibbett, Phillips & Edwards
- ALL Joint Conference on Computers - Systems & Technology,  
I.E.R.E. October, 1972.
34. "The Application of Paging, Segmentation and Virtual Memory":  
Sumner  
INFOTECH Conference 1970, Giant Computers, published 1971.