

NOTICE WARNING CONCERNING COPYRIGHT RESTRICTIONS:
The copyright law of the United States (title 17, U.S. Code) governs the making of photocopies or other reproductions of copyrighted material. Any copying of this document without permission of its author may be prohibited by law.

Copyright © 1986 Stanford University

List of Tables

Table 3-1: MIPS-X Pipeline Stages	7
Table 3-2: Delay Slots for MIPS-X Instruction Pairs	9
Table 4-1: Branch Instructions	32
Table IV-1: Number of Cycles Needed to do a Multiplication	78
Table IV-2: Number of Cycles Needed to do a Divide	78

3.2. Delays and Bypassing

A *delay* occurs because the result of a previous instruction is not available to be used by the current instruction. An example is a *compute* instruction that uses the result of a *load* instruction. If in Figure 3-1, instruction 1 is a *load* instruction, then the result of the *load* is not available to be read from the register file until the second half of *WB* in instruction 1. The first instruction that can access the value just loaded in the registers is instruction 4 because the registers are read on phase 2 of the cycle. This means that there is a *delay* of two instructions from a *load* instruction until the result can be used as an operand by the ALU. An instruction delay can also be called a *delay slot* where an instruction that does not depend on the previous instruction can be placed. This should be a *nop* if no useful instruction can be found. Delays between instructions can sometimes be reduced or eliminated by using *bypassing*.

Bypassing allows an instruction to use the result of a previous instruction before it is written back to the register file. This means that some of the delays can be reduced. Table 3-2 shows the number of delay slots that exist for various pairs of instructions in MIPS-X. The table takes into account bypassing on both the results of a *compute* instruction and a *load* instruction. For example, consider the *load-address* pair of instructions. This can occur if the result of the first load is used in the address calculation for the second load instruction. Without bypassing, there would be 2 delay slots. Table 3-2 shows only 1 delay slot because bypassing will take place.

The possible implementations for bypassing are bypassing only to Source 1 or to both Source 1 and Source 2. The implementation of bypassing in MIPS-X uses bypassing to both sources. Bypassing only to Source 1 means that the benefits of bypassing can only be achieved if the second instruction is accessing the value from the previous instruction via the *Source 1* register. If the second instruction can only use the value from the previous instruction as the *Source 2* register, then 2 delay slots are required. Bypassing to both Sources eliminates this asymmetry. The asymmetry is most noticeable in the number of delay slots between compute or load instructions and a following instruction that tries to store the results of the compute or load instruction. Branches are also a problem because the comparison is done with a subtraction of *Source 1* - *Source 2*. Not all branch types have been implemented because it is assumed that the operands can be reversed. This means that it will not always be possible to bypass a result to a branch instruction. This asymmetry could be eliminated by taking one bit from the displacement field and using it to decide whether a subtraction or a reverse subtraction should be used. The tradeoff between the two types of bypassing is the ability to generate more efficient code in some places versus the hardware needed to implement more comparators. Table 3-2 shows the delays incurred for both implementations of bypassing. It is felt that bypassing to both Sources is preferable and the necessary hardware has been implemented.

Instructions in the slot of load instructions should not use the same register as the one that is the destination of the *load* instruction. Bypassing will occur and the instruction in the load slot will get the address being used for the load instead of the value from the desired register.

One other effect of bypassing should be described. Consider Figure 3-1. If instruction 1 is a *load* to *r1* and instruction 2 is a *compute* instruction that puts its result also in *r1*, then there is an apparent conflict in instruction 3 if it wants to use *r1* as its *Source 1* register. Both the results from instructions 1 and 2 will want to bypass to instruction 3. This conflict is resolved by using the result of the *second* instruction. The reasoning is that this is how sequential instructions will behave. Therefore, in this example instruction 3 will use the result of the compute instruction.

4. Instruction Set

There are four different types of instructions. They are memory instructions, branch instructions, compute instructions, and compute immediate instructions. Coprocessor instructions are part of the memory instructions.

4.1. Notation

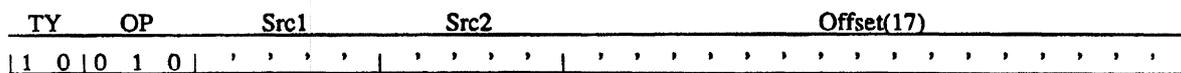
This section explains the notation used in the descriptions of the instructions.

MSB(x)	The most significant bit of x .
$x \ll y$	x is shifted left by y bits.
$x \gg y$	x is shifted right by y bits.
$x\#y$	x is a number represented in base y
$x \parallel y$	x is concatenated with y .
PC _{current}	Address of the instruction being fetched during the ALU cycle of an instruction
PC _{next}	Address of the next instruction to be fetched.
Reg(n)	The contents of CPU register n .
FReg(n)	The contents of register n in the floating point unit (FPU).
Reg $\langle n \rangle$, Reg $\langle n..m \rangle$	Bit n or Bits n to m of register Reg .
Memory[<i>addr</i>]	The contents of memory at the location <i>addr</i> . The value accessed is always a word of 32 bits.
SignExtend(n)	The value of n sign extended to 32 bits. The size of n is specified by the field being sign extended.
rSrc1	The register number used as the Source 1 operand.
rSrc2	The register number used as the Source 2 operand.
rDest	The register number used as the Destination location.
fSrc1	The register number used as the Source 1 floating point operand.
fSrc2	The register number used as the Source 2 floating point operand.
fDest	The register number used as the Destination floating point register.
CopI	Coprocessor instruction.
MAR	The memory address register. The contents of this register are placed on the address pins of the processor.
MDR	The memory data register. The address pads of the processor always reflect the contents of this register.

4.2. Memory Instructions

The memory instructions are the ones that do an external memory cycle. The most commonly used memory instructions are *load* and *store*. The other instructions that are part of the memory instructions are the coprocessor instructions. They do not always generate a memory cycle that is recognized by memory. Instead the coprocessor uses the cycle. This is explained in more detail in the individual instruction descriptions.

4.2.2. st - Store



Assembler

```
st Offset[rSrc1],rSrc2
```

Operation

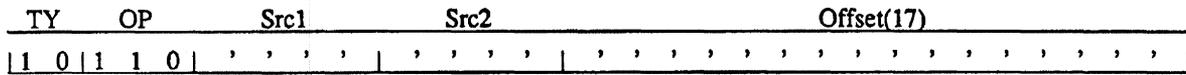
$$\text{Memory}[\text{SignExtend}(\text{Offset}) + \text{Reg}(\text{Src1})] \leftarrow \text{Reg}(\text{Src2})$$

Description

The offset field is sign extended and added to the contents of the register specified by the Src1 field to compute a memory address. The contents of Reg(Src2) are stored at that memory location.

This instruction requires 2 memory cycles, one to read the cache and then one to do the store. To obtain maximum performance, instructions that do not require a memory cycle should be scheduled after a store instruction if possible. Otherwise, the processor may stall for one cycle.

4.2.4. stf - Store Floating Point



Assembler

```
stf Offset[rSrc1],fSrc2
```

Operation

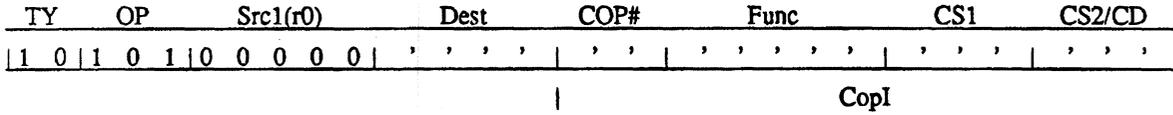
Memory[SignExtend(Offset) + Reg(Src1)] \leftarrow FReg(Src2)

Description

The offset field is sign extended and added to the contents of the register specified by the Src1 field to compute a memory address. The contents of the floating point register specified by Src2 are stored at that memory location. The CPU does not put out any data during this write memory cycle.

Note: If a processor configuration does not have an FPU then different code must be generated to emulate the floating point instructions. Any code that tries to use FPU instructions when there is no FPU will not execute correctly.

4.2.7. movfrc - Move From Coprocessor



Assembler

movfrc CopI,rDest

Operation

MAR \leftarrow SignExtend(CopI) + Reg(Src1)
 Reg(Dest) \leftarrow MDR

Description

This instruction is used to do a Coprocessor register to CPU register move.

The CopI field is sign extended and added to the contents of the register specified by the Src1 field. The Src1 field should be Register 0 if the CopI field is to be unmodified (hackers take note). The CopI field will appear on the address lines of the processor where it can be read by the coprocessor. The coprocessor will place a value on the data bus that will be stored in Reg(Dest) of the CPU. The memory system will ignore this memory cycle.

The CopI field is decoded by the coprocessors to find the coprocessor being addressed (COP#) and the function to be performed. A possible format is shown above. The fields CS1 and CS2/CD show possible coprocessor register fields. The format is flexible except that all coprocessors should find the COP# in the same place.

Note: An instruction in the slot of a movfrc instruction that uses the same register that the movfrc instruction is loading is not guaranteed to get the correct result. Do not try to use the slots in this manner.

4.3. Branch Instructions

As described previously in Section 3.4, all branch instructions have two delay slots. The instructions placed in the slots can be either ones that must always execute or ones that should be executed if the branch is *taken*. There are two flavours of branch instructions that must be used depending on the type of instructions placed in the slots. They are:

No squash:	The instructions in the slots are always executed. They are never squashed (turned into <i>nops</i>).
Squash if don't go:	All branches are statically predicted to <i>go</i> (be taken). This means that the instructions in the branch slots should be instructions from the <i>target</i> instruction stream. If the branch is not taken, then the instructions in the slots are squashed.

The instructions in the slots must be both of the same type. That is, they should both always execute or both be from the target instruction stream. If squashing takes place, both instructions in the slots are treated equally.

Note that for best performance, it is best to try to find instructions that can always execute and use the *no squash* branch types.

Branch instructions can be put in the slot of branches that can be squashed.

The branch conditions are established by testing the result of

$$\text{Reg}(\text{Src1}) - \text{Reg}(\text{Src2})$$

where *Src1* and *Src2* are specified in the branch instruction. The condition to be tested is specified in the *COND* field of the branch instruction. The expressions used to derive the conditions use the following notation:

N	Bit 0 of the result is a 1. The result is negative.
Z	The result is 0.
V	32-bit 2's-complement overflow has occurred in the result.
C	A carry bit was generated from bit 0 of the result in the ALU.
⊕	Exclusive-Or

Some branch conditions that are *usually* found on other machines do not exist on MIPS-X. They can be synthesized by reversing the order of the operands or comparing with *Reg(0)* in Source 2 (*Src2=0*). These branches are shown in Table 4-1 along with the existing branches.

Branch	Description	Expression	Branch To Use If Synthesized
beq	Branch if equal	Z	
bge	Branch if greater than or equal	$\overline{N \oplus V}$	
bgt	Branch if greater than	$\overline{(N \oplus V) + Z}$	blt (rev ops)
bhi	Branch if higher	$\overline{C} + Z$	blo (rev ops)
bhs	Branch if higher or same	C	
ble	Branch if less than or equal	$(N \oplus V) + Z$	bge (rev ops)
blo	Branch if lower than	\overline{C}	
blos	Branch if lower or same	$\overline{C} + Z$	bhs (rev ops)
blt	Branch if less than	$N \oplus V$	
bne	Branch if not equal	\overline{Z}	
bpl	Branch if plus	\overline{N}	bge (cmp to Src2=0)
bmi	Branch if minus	N	blt (cmp to Src2=0)
bra	Branch always		beq r0,r0

Table 4-1: Branch Instructions

4.3.3. bhs - Branch If Higher Or Same

TY		Cond		Src1				Src2				SQ	Disp(16)																									
0	0	0	1	0									s																									

s = 1 \Rightarrow Squash if don't go

s = 0 \Rightarrow No squashing

Assembler

```
bhs   rSrc1,rSrc2,Label      ; No squashing
bhssq rSrc1,rSrc2,Label     ; Squash if don't go
```

Operation

```
If [Reg(Src1) – Reg(Src2)]  $\Rightarrow$  C
then
    PCnext  $\Leftarrow$  PCcurrent + SignExtend(Disp)
```

Description

This is an unsigned compare.

If $\text{Reg}(\text{Src1})$ is higher than or equal to $\text{Reg}(\text{Src2})$ then execution continues at *Label* and the two delay slot instructions are executed. The value of *Label* is computed by adding PCcurrent + the signed displacement.

If $\text{Reg}(\text{Src1})$ is lower than $\text{Reg}(\text{Src2})$, then the delay slot instructions are executed for *bhs* and squashed for *bhssq*.

4.4.8. bic - Bit Clear

TY	OP	Src1	Src2	Dest	Comp Func(12)
0 1 1 0 0	' ' ' '	' ' ' ' ' '	' ' ' ' ' '	' ' ' ' ' '	0 0 0 0 0 0 0 0 0 0 1 0 1 1 1

Assembler

bic rSrc1,rSrc2,rDest

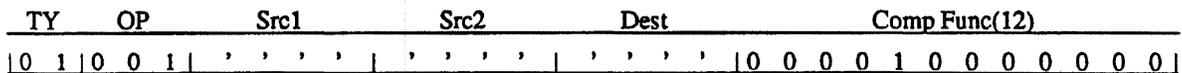
Operation

$\text{Reg}(\text{Dest}) \leftarrow \overline{\text{Reg}(\text{Src1})}$ bitwise and $\text{Reg}(\text{Src2})$

Description

Each bit that is set in Source 1 is cleared in Source 2. The result is placed in Destination.

4.4.15. rotlcb - Rotate Left Complemented by Bytes



Assembler

```
rotlcb rSrc1,rSrc2,rDest
```

Operation

$\text{Reg(Dest)} \Leftarrow \text{Reg(Src1)} \text{ rotated left by } \text{BitComplement}[\text{Reg(Src2)} < 30..31 >] \text{ bytes}$

Description

This instruction rotates left the contents of Source 1 by the number of bytes specified by using the bit complement of bits 30 and 31 in Source 2. For example,

```
Reg(Src1) = AB01CD23#16
```

```
Reg(Src2) = 51#16
```

```
rotlcb rSrc1,rSrc2,rDest
```

Rotate amount is $\text{BitComplement of } 01\#2 = 10\#2 = 2.$

```
Reg(Dest) = CD23AB01#16
```


4.5.8. hsc - Halt and Spontaneously Combust

TY	OP
1 1	0 0 1 1 1 1 1 1 0

Assembler

hsc

Operation

Reg(31) ← PC

The processor stops fetching instructions and self destructs.

Note that the contents of Reg(31) are actually lost.

Description

This is executed by the processor when a protection violation is detected. It is a privileged instruction available only on the -NSA versions of the processor.

I

This routine will jump through low core to flush the cache by setting all the tags to be in system space. Note that this routine will also blow away the entry in the cache that called this routine but to make it be general it will have to since you don't want to have to figure out where you came from. This is called from a trap so it knows where to return to.

The sequence of jump locations is designed to account for the behavior of the ring counter that is used to determine the next instruction cache block to be replaced. It is not sufficient to access the locations in sequence.

The ".makenop n" means that "n" nop instructions should be inserted.

This module should be loaded starting at address 0x1800.

```

text
moreorg
clicaps:
0x1800:
    jpc1 r0,#0x1810,r0
    jpc1 r0,#0x1820,r0
    jpc1 r0,#0x1830,r0
    .makenop 13
0x1810:
    jpc1 r0,#0x1840,r0
    .makenop 15
0x1820:
    jpc1 r0,#0x1850,r0
    .makenop 15
0x1830:
    jpc1 r0,#0x1860,r0
    .makenop 15
0x1840:
    jpc1 r0,#0x1870,r0
    .makenop 15
0x1850:
    jpc1 r0,#0x1890,r0
    .makenop 15
0x1860:
    jpc1 r0,#0x18a0,r0
    .makenop 15
0x1870:
    jpc1 r0,#0x18b0,r0
    .makenop 15
0x1880:
    jpc1 r0,#0x1940,r0
    .makenop 15
0x1890:
    jpc1 r0,#0x18c0,r0
    .makenop 15
0x18a0:
    jpc1 r0,#0x18d0,r0
    .makenop 15
0x18b0:
    jpc1 r0,#0x18e0,r0
    .makenop 15
0x18c0:
    jpc1 r0,#0x18f0,r0
    .makenop 15
0x18d0:
    jpc1 r0,#0x1880,r0
    
```

```

    .makenop 15
0x18e0:
    jpc1 r0,#0x1920,r0
    .makenop 15
0x18f0:
    jpc1 r0,#0x1930,r0
    .makenop 15
0x1900:
    jpc1 r0,#0x19c0,r0
    .makenop 15
0x1910:
    jpc1 r0,#0x19d0,r0
    .makenop 15
0x1920:
    jpc1 r0,#0x1950,r0
    .makenop 15
0x1930:
    jpc1 r0,#0x1960,r0
    .makenop 15
0x1940:
    jpc1 r0,#0x1970,r0
    .makenop 15
0x1950:
    jpc1 r0,#0x1900,r0
    .makenop 15
0x1960:
    jpc1 r0,#0x1910,r0
    .makenop 15
0x1970:
    jpc1 r0,#0x19b0,r0
    .makenop 15
0x1980:
    ; start return
    movfcs pcm4,r31 ; save this for restart
    .makenop 15
0x1990:
    movtos r0,pcm1 ; prepare for return
    .makenop 15
0x19a0:
    movtos r0,pcm1 ; restart instruction after trap
    movtos r31,pcm1
    pcra
    jpc
    .makenop 11
0x19b0:
    jpc1 r0,#0x19e0,r0
    .makenop 15
0x19c0:
    jpc1 r0,#0x19f0,r0
    .makenop 15
0x19d0:
    jpc1 r0,#0x1980,r0
    .makenop 15
0x19e0:
    jpc1 r0,#0x1990,r0
    .makenop 15
0x19f0:
    jpc1 r0,#0x19e0,r0
end

```


Bit 25 For byte rotate instructions, set to 1 if rotlb, 0 if rotcb
Bits 25-31 Shift amount for funnel and arithmetic shift operations (sh and asr instructions). The range is 0 to 31 bits. Although this can be encoded in five bits, the two low-order bits are fully decoded; therefore, the field is seven bits. The two low-order bits are decoded as follows: 0 = bit 31, 1 = bit 30, 2 = bit 29, 3 = bit 28. For example, a shift amount of 30 would become 1110100 in this seven-bit encoding scheme.


```

//////////////////////////////////////////////////////////////////
;
; MUL
; fast, unchecked, signed multiply
;           rLink = link
;           rMand = src2
;           rDest = rMer = src1/dest
;           rTemp1 = temp
;           rTemp2 = temp
;
; Note: This code has been reorganized
;
//////////////////////////////////////////////////////////////////
MUL:
    asr     rMer,rTemp2,#7           ; Test for positive 8-bit number
    bne    rTemp2,r0,lnot8
    sh     r0,rMer,rTemp1,#24        ; assume 8 bit
    movtos rTemp1,md
    mstart rMand,rDest              ; may need nop before this
    mstep  rDest,rMand,rDest
lmul8bit:
    mstep  rDest,rMand,rDest
    mstep  rDest,rMand,rDest
    mstep  rDest,rMand,rDest
    mstep  rDest,rMand,rDest
    jspci  rLink,#0,r0
    mstep  rDest,rMand,rDest
    mstep  rDest,rMand,rDest
lnot8:
    addi   rTemp2,#1,rTemp2
    beqsq  rTemp2,r0,lmul8bit        ; 8 bit negative
    mstart rMand,rDest
    mstep  rDest,rMand,rDest
    movtos rDest,md                  ; do full 32 bits
    mstart rMand,rDest              ; may need nop before this
    mstep  rDest,rMand,rDest
    mstep  rDest,rMand,rDest
    mstep  rDest,rMand,rDest
    mstep  rDest,rMand,rDest
    .
    .
    .
    24 msteps
    .
    .
    .
    mstep  rDest,rMand,rDest
    jspci  rLink,#0,r0
    mstep  rDest,rMand,rDest
    mstep  rDest,rMand,rDest

```

Figure IV-1: Signed Integer Multiplication

Appendix V Multiprecision Arithmetic

Multiprecision arithmetic is not a high priority but it is desirable to make it possible to do. The minimal support necessary will be provided. The most straightforward way to do this would seem to be the addition of a carry bit to the PSW. However, this turns out to be extremely difficult.

The following program segments are examples of doing double precision addition and subtraction. The only addition required to the instruction set is the *Subtract with No Carry (subnc)* instruction. This is only an addition to the assembly language and not to the hardware.

Assume that there are 2 double precision operands (*A* and *B*) and a double precision result to be computed (*C*). Assume that the necessary registers have been loaded.

```

;Double precision addition
      add    rAhi,rBhi,rChi      ;add high words
      sub    r0,rBlo,rClo       ;get -rBlo; branch does subtract
      bhssq  rAlo,rClo,l1       ;check to see if carry generated
                                      ;branch if carry set
      addi   rChi,#1,rChi       ;add 1 to high word if carry
      nop
l1:   add    rAlo,rBlo,rClo      ;add low words

;Double precision subtraction
      subnc  rAhi,rBhi,rChi     ;subtract high words
      bhssq  rAlo,rBlo,l1       ;check if subtract of low
                                      ;words generates a carry
                                      ;branch if carry set
      addi   rChi,#1,rChi       ;add 1 to high word if carry
      nop
l1:   sub    rAlo,rBlo,Clo      ;subtract low words

```



```

| macroState
label      : ID :      { ID must be in column 1 }
binALUState : binALUOp reg,reg,reg
binALUOp   : ADD
           | SUB
           | AND
           | OR
           | XOR
           | ROTLB
           | ROTLCB
           | MSTEP
           | DSTEP
           | SUBNC
           | BIC
monALUState : monOp reg,reg
           | MSTART reg,reg
monOp       : NOT
           | MOV
specState   : MOVTOS reg,specialReg
           | MOVFRS specialReg,reg
specialReg  : MD
           | PSW
           | PCM4
           | PCM1
nopState    : NOP
addiState   : ADDI reg,#exp,reg
jspciState  : JSPCI reg,#exp,reg
shiftState  : ASR reg,reg,#exp
           | SH reg,reg,reg,#exp
           | LSR reg,reg,#exp
           | LSL reg,reg,#exp
loadState   : LD exp[reg],reg
           | LD #exp,reg
           | { adds constant to literal pool and loads it }
           | LDT exp[reg],reg
           | LDF exp[reg],freg
storeState  : ST exp[reg],reg
           | STT exp[reg],reg
           | STF exp[reg],freg
branchState : branchOp reg,reg,ID
           | branchSqOp reg,reg,ID
           | BRA ID
branchOp    : BEQ
           | BNE
           | BGE
           | BGT
           | BHI
           | BHS
           | BLE
           | BLO
           | BLS
           | BLT
branchSqOp  : BEQSQ
           | BNESQ
           | BGESQ
           | BGTSQ
           | BHISQ
           | BHSSQ
           | BLESQ
           | BLOSQ
           | BLSSQ
           | BLTSQ
copState    : MOVTOC exp,reg

```

```

| MOVFRC exp,reg
| ALUC exp
| floatBinOp freg,freg
| floatMonOp freg,freg
| MOVIF reg,freg
| MOVFI freg,reg
floatBinOp : FADD
| FSUB
| FMUL
| FDIV
| IMUL
| IDIV
| MOD
floatMonOp : CVTIF
| CVTFI
miscState  : TRAP exp
| JPC
| JPCRS
directState : TEXT
| DATA
| END
| EOP
| ASCII STRING { string: ".*" }
| WORD exp
| FLOAT FLOATCONSTANT
| ID = exp
| DEF ID = exp
| REORGON
| NOREORG
| COMM ID,INT
| GLOBL ID
| LIT liveList
| LIF liveList
liveList   : reg
| liveList,reg
macroState : PJSR ID,#exp,reg
| IPJSR reg,#exp,reg
| IPJSR exp,reg,#exp,reg
| RET
exp        : exp addOp term
| - factor
| term
addOp      : +
| -
term       : term multOp factor
| factor
multOp     : *
factor     : ( exp )
| ID
| INT
| HEXINT   { like C: 0x12fc }
reg        : REG      { r0..r31 }
freg       : FREG     { f0..f15 }

```

notes:

- 1) only labels and directives may start in column 1
- 2) Keywords are shown in upper case just to make them stand out. In reality, they MUST be lower case.
- 3) directives begin with a '.'

References

- [1] Cohen, Danny.
On Holy Wars and a Plea for Peace.
IEEE Computer 14(10):48-54, October, 1981.
- [2] Gill, J., Gross, T., Hennessy, J., Jouppi, N., Przybylski, S. and Rowen, C.
Summary of MIPS Instructions.
Technical Note 83-237, Stanford University, November, 1983.
- [3] Lamport, Leslie.
A Fast Mutual Exclusion Algorithm.
Technical Report 7, DEC Systems Research Center, November, 1985.