```
Inter-Office Memorandum
```

To	File	Date	October 6, 1978
From	Roy Levin	Location	Palo Alto
Subject	On the Harmonious Cooperation of XMesa and User Microcode (version 2)	Organization	CSL

XEROX

Filed on: [Ivy]<XMesa>Doc>Microcode.bravo (and .press)

This document describes an interface that permits users to load their own microcode into the RAM of an Alto II XM and execute it without interfering with the Mesa emulator. These facilities are available only under XMesa. Higher-level facilities of XMesa are described in a separate document filed on [Ivy]<XMesa>Doc>XMesa.press (and .bravo).

Loading User Microcode

The XMesa emulator occupies all of ROM1 and about 25 words of RAM. The remainder of the RAM is available for use by client programs. To simplify loading of user microcode, a file containing the overflow portion of the XMesa emulator is included in the user's assembly, producing a single output file that can be loaded in its entirety. Addressing requirements are specified in the file; no locations are reserved "by convention" alone.

Assembly Procedure

Users should carefully follow the steps below in constructing RAM images to be loaded under XMesa.

1) Obtain `MesaXRAM.mu` and `MesabROM.mu` from the dump file `[Ivy]<XMesa>System>XMesaMu.dm`. The former is the actual microcode, the latter contains the definitions for all registers and other symbols used by the Mesa emulator.

2) Prepare a source file, `Driver.mu`, of the following form:

```
#AltoConsts23.mu;                 standard Alto definitions

; Sufficient pre-definitions should be inserted here to
; reserve locations 0-17 inclusive.

#MesaXRAM.mu;                     includes MesabROM internally
#UserMicrocode.mu;                user microcode source
```

3) Assemble `Driver.mu` with MU, obtaining `Driver.mb`.

4) Use the PackMU subsystem to convert `Driver.mb` to `Driver.br`.

The code in `MesaXRAM.mu` must be assembled to begin at RAM location 20. The user must assure this condition by including a predefinition (with no omitted labels) at the indicated place in `Driver.mu`; e.g.,

```
%17,1777,0,L0,L1,L2,L3,L4,L5,L6,L7,L10,L11,L12,L13,L14,L15,L16,L17;
```

If the user microcode includes the code to control a non-standard I/O device (e.g., a Trident disk), at least one of the labels in this predefinition will refer to an instruction used as part of the silent boot sequence (see Alto Hardware Manual, sections 2.4, 8.4, and 9.2.2). Prudence suggests that the locations corresponding to tasks that are not intended to execute in the RAM be filled in with dummy instructions of the form

```
Li:          TASK, :Li;
```

Thus, if `Driver.mu` contains no device control microcode, L0-L17 should all have this form. If there is a space crunch, however, locations 0-17 may be used for ordinary microcode, though this is usually unnecessary.

Loading Procedure

`Driver.br` can now be loaded using the (Mesa) RamLoad package. The facilities of this package are defined in `RamDefs.bcd`, and are exported by the module `RamLoad.bcd`. These files may all be found on [Ivy]<XMesa>Utilities. `RamDefs.mesa` contains adequate documentation for the use of this package; the following Mesa code illustrates a typical use.

```
BEGIN OPEN RamDefs;
driver: MuImage _ ReadPackedMuFile["Driver.br"];
IF LoadRamAndBoot[driver,FALSE] ~= 0 THEN SIGNAL BogusMicrocode;
ReleaseMuImage[driver];
END;
```

The second parameter to **LoadRamAndBoot** controls whether a silent boot is performed and should be **FALSE** unless it is necessary to alter the set of running tasks (as is required when additional device control microcode is initialized). **LoadRamAndBoot** returns the number of mismatches between the constants specified in the microcode file and those present in the Alto's constants ROM. If this number is non-zero, the microcode cannot execute properly on this machine.

Note to the nervous: Since RamLoad is a Mesa program, it must exercise caution to avoid smashing the emulator on which it is running. As long as the procedure above is meticulously followed in constructing Driver.br, RamLoad will be able to load the RAM without committing suicide.

There and Back Again

This section describes mechanisms for transferring control between the Mesa emulator and user-supplied microcode. It assumes that this microcode is, in effect, implementing "extended instructions", and the linkage mechanisms described are suitable for this purpose. Device driver microcode (*e.g.,* for a Trident disk) executes in a different hardware task and therefore does not require these linkages.

The hardware-defined linkage mechanism between control memory banks is somewhat precarious. Opportunities for errors arise because of way the "next address" field of the microinstruction is interpreted when SWMODE has been invoked (see section 8.4 of the Alto Hardware Manual). The thing to remember is that whether you are in ROM1 or RAM, to get to the other memory you must specify an address with the 400 bit on (*i.e.,* BITAND[address,400B] = 400B). If you preserve

this invariant, your Alto will probably avoid most black holes.

The Mesa JRAM Instruction

Mesa has a bytecode, JRAM, that is a straightforward mapping of the Nova JMPRAM instruction. JRAM takes the top element off the stack and dispatches on it, doing a SWMODE in the same instruction. The microcode therefore looks approximately like this:

```
JRAM:      IR_sr17, :Pop;           pops stack into L, T; returns to JRAMr
JRAMr:     SINK_M, BUS, SWMODE;
           L_0, :zero;
```

Thus, at entry to whatever microcode JRAM jumps to, L=0 and T has the target address of the jump.

Getting back to ROM1 is equally straightforward. User microcode should terminate with:

```
           SWMODE;                  sigh...SWMODE and TASK are both F1s
           :romnext;                (... and we can't TASK here)
```

`romnext` is defined in `MesabROM.mu`, as are all the registers and other symbols used by the Mesa emulator.

The safest way to force a microinstruction to a specific address in the RAM is to use MU's '%' pre-definition facility. For example,

```
%1,1777,440,MyCode;

MyCode:    . . . ;                  start of user-written microcode
```

would enable the Mesa procedure **MyCode** (defined just below) to transfer control to the microcode at `MyCode`.

Setting Up a JRAM in Mesa

The easiest way to access microcode in the RAM is by declaring a procedure to invoke it as follows:

locationInRAM: CARDINAL = 440B;

MyCode: PROCEDURE[arg1: AType, arg2: AnotherType] RETURNS [result: AThirdType] =
      MACHINE CODE BEGIN
      Mopcodes.zLIW, locationInRAM/256, locationInRAM MOD 256;
      Mopcodes.zJRAM
      END;

When **MyCode** is invoked, the arguments are pushed on the evaluation stack, then a transfer to location 440 in the RAM occurs. The microcode is responsible for computing **result** and placing it on the stack (details appear below). The return sequence described above will then cause execution to resume immediately after the invocation of **MyCode**.

The Mesa Stack

The Mesa stack is implemented by 8 S-registers named `stk0`, `stk1`, ..., `stk7`, with `stk0` being the base of the stack. An R-register, `stkp`, indicates the number of words on the stack and is thus in the range [0..8]. To obtain values from the stack in the general case, therefore, requires a dispatch on `stkp`, but there is an important special case. If **MyCode** is invoked in a statement context (*i.e.*, one in which it stands alone and is not a term of an expression), the stack will be empty except for **arg1** and **arg2**, which will therefore appear in `stk0` and `stk1`, respectively. In this case `stkp` will be 2 at entry, and the microcode should set it to 1 before returning to ROM1. **result** should also be stored in `stk0`. It is frequently useful to restrict the use of **MyCode** to statement contexts so that the microcode can take advantage of the known stack depth. Note: this assumes that values of type **AType**, **AnotherType**, and **AThirdType** are each represented in a single word. Multi-word quantities are stored in adjacent stack words, with consecutive memory cells being pushed in increasing address order.

An argument or return record exceeding 5 words in length requires special handling that is beyond the scope of this document. See Roy Levin for details.

Other Emulator State

The Mesa emulator has a number of temporary registers that may be freely used by code entered via JRAM. The following list identifies all registers used by the Mesa emulator, their intended function, and whether they may be used as destroyable temporaries by user-supplied RAM microcode:

Register	Type	Destroyable by RAM code	Function
mpc	R	No	program counter
stkp	R	No	stack pointer
XTSreg	R	No	Xfer trap status
ib	R	No	instruction byte
brkbyte	R	No	Xfer state variable; overlaps AC3
clockreg	R	No	high resolution clock bits
mx	R	Yes	temporary during Xfer; overlaps AC2
saveret	R	Yes	holds return dispatch values; overlaps AC1
newfield	R	Yes	used by field instructions; overlaps AC0
count	R	Yes	temporary for counting
taskhole	R	Yes	temporary for holding things across TASKs
temp	R	Yes	general temporary
temp2	R	Yes	general temporary
lp	S	No	local frame pointer (+6)
gp	S	No	global frame pointer (+4)
cp	S	No	code segment pointer
stk0-7	S	No	8 evaluation stack registers
wdc	S	No	wakeup disable counter (interrupt control)
ATPreg	S	Yes	alloc trap parameter
XTPreg	S	Yes	xfer trap parameter
OTPreg	S	Yes	other trap parameter
mask	S	Yes	used by field instructions; smashed by BITBLT
index	S	Yes	used by field instructions; smashed by BITBLT
alpha	S	Yes	temporary for alpha byte; smashed by BITBLT
entry	S	Yes	Xfer and field temporary; smashed by BITBLT
frame	S	Yes	ALLOC/FREE temporary; smashed by BITBLT

my	S	Yes	Xfer temporary; smashed by BITBLT
unused1	S	Yes	smashed by BITBLT
unused2	S	Yes	smashed by BITBLT
unused3	S	Yes	smashed by BITBLT

User microcode should endeavor to TASK within 5 microcycles of entry. It should also TASK as close to the end as possible. Ideally, the penultimate instruction before returning to the emulator would TASK, but unfortunately SWMODE must appear in the same instruction and both are F1s. The Mesa emulator tries to TASK at least every 12 microcycles; user microcode should observe the same guideline.

The Mesa emulator does not check for pending interrupts on every instruction. It does so only when it must fetch a new instruction word from memory. Therefore, user microcode that sets a pending interrupt condition must not expect that the interrupt will be noticed by the emulator immediately upon return to ROM1.

Distribution:
    XMesa Users
    Mesa Group