

Virtual Addressing for the EMMY/360

by Walter A. Wallach

June 1976

Technical Note No. 89

Digital Systems Laboratory
Stanford Electronics Laboratories
Stanford University
Stanford, California

The work herein was supported in part by the Army Research Office-Durham under Grant No. DAAG29-76-G-0001.

Digital Systems Laboratory

Stanford Electronics Laboratories

Technical Note No. 89

June 1976

Virtual Addressing for the EMMY/360

by Walter A. Wallach

Abstract

A virtual addressing technique for expanding the address space of the EMMY/360 to 2 million bytes is presented. The scheme uses a translation table of 512 words of EMMY control store, addressed using bits 20:12 of the logical (virtual) program address. The page table entry reflects status of the referenced page- ie available, or written to. The semantic pointer for each instruction controls whether a logical address is to be translated, or the execution routine is to use the logical address.

The work herein was supported in part by the Army Research Office-Durham under Grant No. DAAG29-76-G-0001.

1.0 Introduction

The EMMY/360 provides a program (logical or virtual) address space of 24 bits, or 16,777,216 bytes (16Mb). The available main store, however, is only 65,536 bytes (64Kb). A method by which the available storage of the 360 may be expanded has been developed. While the entire 16Mb address space will still not be available, a realistic main store address space of 2,097,152 bytes (2Mb) can be supported through the use of a paging disk.

1.1 Background

The IBM technique [1,2] for dynamic address translation, as implemented on the 360 model 67 and various 370 models, use segment and page tables in main store. Page size is 4096 bytes (model 67) and 2048 or 4096 bytes (370 models, chosen under program control).

Bits 23:20 of the virtual address are used to index a segment table (pointed to by a control register). Bits 19:12 of the virtual address are used as an index to a page table, whose address was contained in the segment table entry. The page table entry indicates whether the referenced page is available (in core), or invalid. If the page is available, the page table entry contains the real storage address of the page. This real page address is concatenated on the low side with bits 11:0 of the virtual address, and the resulting address used to access main store. If a page is not available, or the calculated segment or page table entry is invalid (beyond the indicated end of the respective table), an interrupt occurs.

An eight element buffer is maintained with the page table entries of the last eight pages referenced to bypass the multiple main store accesses needed to translate an address. A content addressable memory associates the segment and page addresses of the last eight references with the requested page's segment and page address.

Implementing this scheme would involve the addition of new hardware (a CAM) and a time consuming translation process. It is not clear whether a 16Mb address space is necessary for architectural studies or lab-oriented production work.


The translation process implemented on the RCA3 (and RCA Spectra 70/46) processors [3] appears more attractive for implementation on the EMMY/360 system. This technique uses a 512 halfword (16 bit) translation table memory (TTM) with 300 ns

access to translate storage references into physical addresses. The logical address is partitioned into a 5-bit segment and 6-bit page address, and a 12 bit byte address (line in page). The high order address bit (bit 24) indicates whether this address is to be translated.

Page size is fixed at 4096 bytes. Since only 512 translation table entries exist, the available address space is 2Mb. The high order 2 bits of the virtual address (bits 21:22) are ignored, and storage references beyond 2Mb wrap around location zero.

The 2 Mb address space is partitioned into 8 262,144 byte (64 page) segments. Paging is controlled by the D bit (bit 24 of the virtual address) and the T bit (bit 13 of the Interrupt Status Register).

The TTM consists of 512 words of 16 bits each. Each entry contains control, status, and translation information:



- W - 1=page has been written to
- G - 1=page has been accessed
- U - 1=page is available (in core)
- S - 1=page is nonprivileged
- E - 1=page is read only
- M,H - not used
- xxx - reserved
- Real Page - page of physical main storage which contains the referenced operand

The W and G bits are set by the processor while U, S, and E are set by the software. A reference to a page whose TTM entry contains a reset U bit causes a Paging Queue interrupt. A write reference to a page whose TTM entry contains a set E bit, or a nonprivileged program reference to a page whose S bit is reset cause a Paging Error interrupt.

Translation proceeds by interrogating the D and T bits. If the D bit is 0 and the T bit is 1, translation proceeds. Translation is accomplished by using the segment and page address portions of the virtual address to access the TTM. The U, S and E bits of the retrieved entry are interrogated. A reset U bit causes a Paging Queue interrupt. The processor sets the G

(accessed) bit, and if the reference is a write reference, the W (written into) bit. Finally, translation is completed by concatenating the 6 bit page field (7 bits for the RCA3) with the 12 bit line in page field of the virtual address. Bits 31:18 are cleared to zero.

2.0 EMMY/360 Virtual Addressing

The EMMY/360 scheme resembles that of the RCA Spectra 70/46 processor. A 512 word TTM is maintained in control store. Translation control bits appear in the semantic pointer for the various instructions.

Translation is done during the format decode phase of the instruction cycle for data operands, and the operation decode phase for instructions (for a branch instruction, the branch target address is treated as data, and any potential paging queue interrupts deferred until the execute phase. If the branch is not taken, no Paging Queue or Paging Error interrupts can occur).

After the effective address is calculated, the Translate bit (bit 21) of the Semantic pointer is interrogated. If this bit is set, the address is translated. Otherwise, it is given to the execute phase as is.

Translation proceeds as follows:

Logical address bits 20:12 are used to index the TTM in control store. The U bit (bit 31) of the TTM entry indicates, if reset, that the referenced page is unavailable. Unless deferred by bit 19 of the Semantic pointer, a page translation exception interrupt results (program interrupt code 11 hex). Bits 23:21 of the virtual address are ignored. The operand is checked to see if it crosses a page boundary. If it does, a specification error results.

A deferred Paging Queue interrupt is indicated by resetting bit 19 of the Microcontrol register. It indicates to the execution semantic routine that either the translation process was aborted because the referenced page was unavailable, or the operand crossed a page boundary (branch target addresses must not reference a location within 3 halfwords of a page boundary). A Paging Queue interrupt or Specification error interrupt will be taken only if the invalid page is actually used.

The effective virtual address is then translated by replacing bits 17:12 of the virtual address with bits 17:12 of the TTM

entry. Bits 31:18 are set to zero.


2.1 TTM Entry

The page table entry appears in the TTM as follows:

Bit 31 - U bit 1=page available
Bit 30 - W bit 1=page has been written to
Bits 29:18 - unspecified
Bits 17:12 - Real Page (physical storage page)
Bits 11:0 - ignored


The decode routine sets bit 30 of the TTM entry and replaces it in control store if bit 20 (write indicator) of the semantic pointer is reset.

Semantic Pointer



MMC Main Memory Control Byte
 ILC Instruction Length Code
 T Translate Bit 1=translate operand address
 W Write Indicator 1=page will not be written
 x x x - unspecified
 D Deferred Interrupt 0=defer paging queue and specification
 interrupts until execution phase
 Semantic Entry Point - address of execute phase semantic code

TTM Entry



U - Utilization Bit 1=page is available
 W - Written into bit 1=page has been written into
 x x ... x x x - unspecified
 Real Page - physical storage page address

3.0 Programming Considerations

In order to keep overhead to a minimum, certain restrictions are imposed on operand and instruction addressing.

- 1) Address space is restricted to 2M bytes. Addresses above 2Mb will wrap around address zero. No checking is done of bits 23:21 of the virtual address.
- 2) Operands may not cross a page boundary. Otherwise, dynamic translation during execute phase would be necessary. Operands may appear on different pages, however. The microcode will distinguish which address is to be written to in the case of instructions with two storage operands.
- 3) Channel operations must specify real addresses (no translation is performed during channel operations). The real address of a buffer may be obtained by using the Load Real Address instruction.
- 4) The W bit indicates a page has been written, ie the copy in core is no longer the same as the one on backing store. A reset W

bit indicates that, since the copy on backing store is identical to the one in main store, the copy in main store may be overwritten with a page not currently available. The TTM entry for the replaced page must be updated to show it is no longer available.

5) Fixed storage locations are always real addresses-no translation is done during interrupt generation. Real address zero need not correspond to virtual address zero (as seen by the program), however monitor overhead will be reduced if fixed locations always remain available to the program.

4.0 Paging Interrupt

When a paging interrupt occurs, the processor returns to the program the operand address, TTM address, and TTM entry for each translation exception recognized during DECODE. The PSW points to the instruction causing the interrupt (operation is suppressed and next instruction address not updated). The ILC reflects the length of the previous instruction.

The operand address, TTM address and entry are returned in a list beginning at main store location 100 hex. The software must perform the necessary paging and modify the TTM entries to reflect any changed status of virtual pages. Since the PSW points to the suppressed instruction, a LPSW will cause that instruction to be retried.

4.1 TTM Support Instructions

Several instructions have been added to support virtual addressing. These are, for the most part, the same (functionally) as those available on the Spectra 70/46. Except where specified, no translation of operand addresses is performed.

Load Translation Memory

9A	C0	b2	d2
----	----	----	----

A block of 512 words beginning at the address specified by b2,d2, is loaded into the TTM.

Condition Codes are not modified. TTM entries are not modified.

Notes:

1) The W bits of the various entries should be zeroed by the program if desired.

Scan and Store Translation Memory

9A	C1	b2	d2
----	----	----	----

TTM is scanned and the address of each TTM entry whose W bit is zero is stored as a 16 bit value. The resulting table of halfword addresses is stored at the location specified by b2,d2. The 9 bit TTM address appears as the low 9 bits of each halfword. The list of halfwords is terminated by a halfword of all ones.

Condition Codes are not modified.

Notes:

1) This instruction is used to determine which pages can be replaced following a Paging Queue interrupt.


Store Translation Memory



The entire Translation Table Memory is stored beginning at the address specified by b2,d2. All 512 entries are stored.

Condition Codes are not modified.

Load Real Address



The virtual address specified by x2,b2,d2 is translated and the real (translated) address replaces bits 23:0 of the register specified by r1. Bits 31:24 are set to zero.

Condition Codes are set as follows:

- 0 - translation successful
- 1 - not used
- 2 - translation referenced an unavailable page
(U bit of TTM entry was zero)
- 3 - not used

References

1. IBM Corp., System/360 Principles of Operation, order no. GA22-6810-11
2. IBM Corp., System 360 Model 67 Functional Characteristics, order no. GA27-2719-2.
3. RCA, Inc., RCA3 Processor Reference Manual, order no. BF-003-2-00.




Figure 5. Data Flow for Dynamic Address Translation (24-Bit)

(360 Model 67, from Reference 2)

Table 1. Format of Registers and Entries

<u>Bits</u>	<u>Meaning</u>	<u>Remarks</u>
1. Segment Table Register Format:		
0-7	Segment Table Length*	Indicates the number of 16-entry groups in the segment table. All zero = One group.
8-31	Segment Table Origin	Since the segment table origin is located on a 64-byte boundary, bits 26-31 must be zero.
2. Segment Table Entry Format:		
0-7	Page Table Length	Indicates the number of entries in the page table. All zero = One entry.
8-30	Page Table Origin	The page table origin is located on a 2-byte boundary.
31	Page Table Availability	1 = Segment translation exception (program interrupt code 16).
3. Page Table Entry (Halfword):		
0-11	Physical Block Address	Starting address of page.
12	Page Availability	1 = Page translation exception (program interrupt code 17).
13-15	Control Bits, Reserved	Must be 000 or specification exception.

*Used only with CPU's that contain the 32-bit addressing feature.

Table 2. Bit Alignment of Address Arithmetic

<u>Bits (24-Bit Mode)</u>	<u>Meaning</u>	<u>Remarks</u>
8-31	Table Origin	Bits 26-31 are considered zero.
8-11	Added to Logical Address	Aligned with bits 26-29 of segment table origin.
8-31	Yields Sum	Segment table entry address (bits 30-31 always zero).
<u>Bits (32-Bit Mode)</u>	<u>Meaning</u>	<u>Remarks</u>
8-31	Table Origin	Bits 26-31 are considered zero.
0-11	Added to Logical Address	Aligned with bits 18-29 of segment table origin.
8-31	Yields Sum	Segment table entry address (bits 30-31 always zero).

Table 2. Bit Alignment of Address Arithmetic (Cont)

<u>Bits</u>	<u>Meaning</u>	<u>Remarks</u>
2. Page Table Entry Address - Either 24- or 32-Bit Mode		
8-31	Page Table Origin	Bit 31 is considered zero.
12-19	Added to Logical Address	Aligned with bits 23-30 of page table origin.
8-31	Yields Sum	Page table entry address (bit 31 is always zero).
3. Physical Address Result		
0-11	Page Table Entry	The high-order portion.
20-31	Logical Address	The low-order portion.
8-31	Physical Address	Both portions taken together.

Table 3. Associative Register Format

24-Bit Addressing:

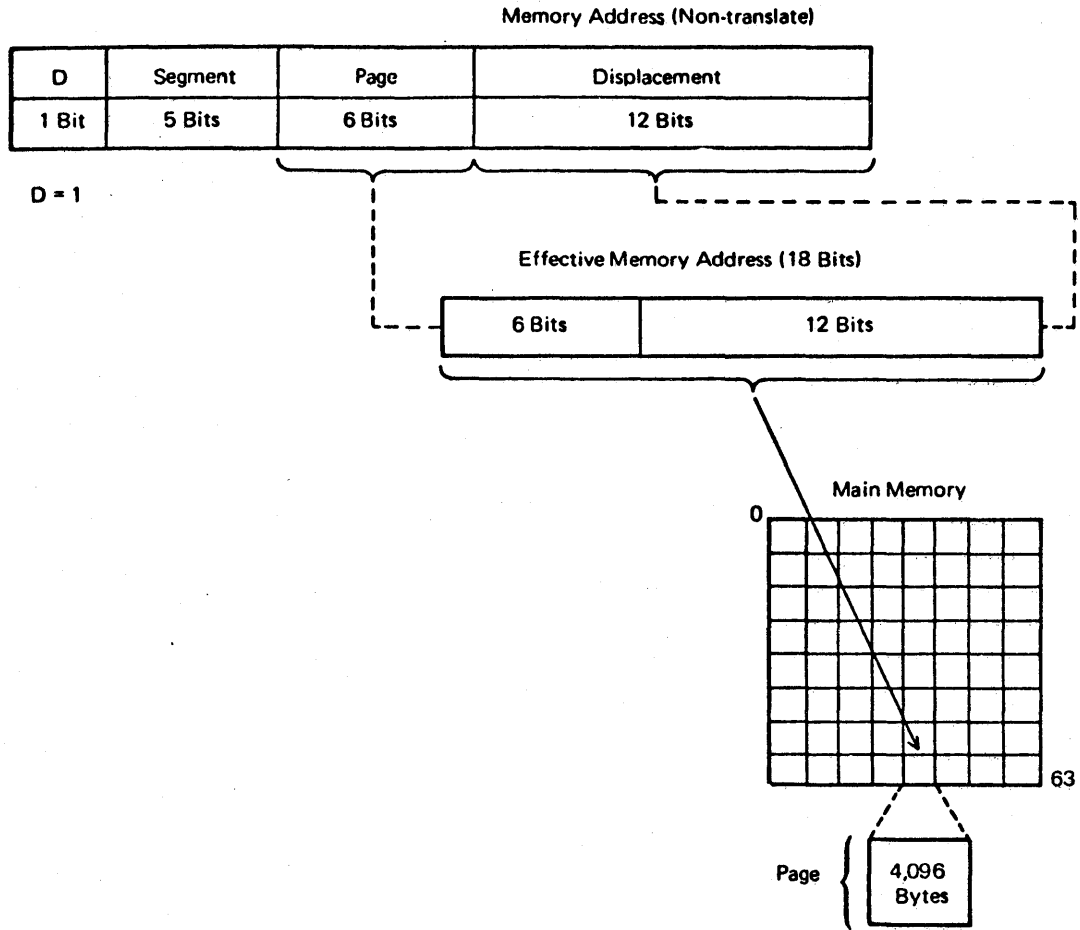
<u>Bits</u>	<u>Content</u>	<u>Remarks</u>
8-19	Virtual Address	The page address from a previous translation which corresponds to the virtual address in bits 8-19.
20-31	Physical Address	
32-35	Unassigned* Register Valid	Set to 1 upon loading the register.
37	Recent Usage - "Load"	Set to 1 upon loading the register and upon any use thereafter.
38	Disable	Set with special diagnose codes 8-15.

32-Bit Addressing: Same as above except for virtual address, which is in bits 0-19.

*For expansion; not physically implemented.


Definition of Control Fields
(from Reference 2)

Example A

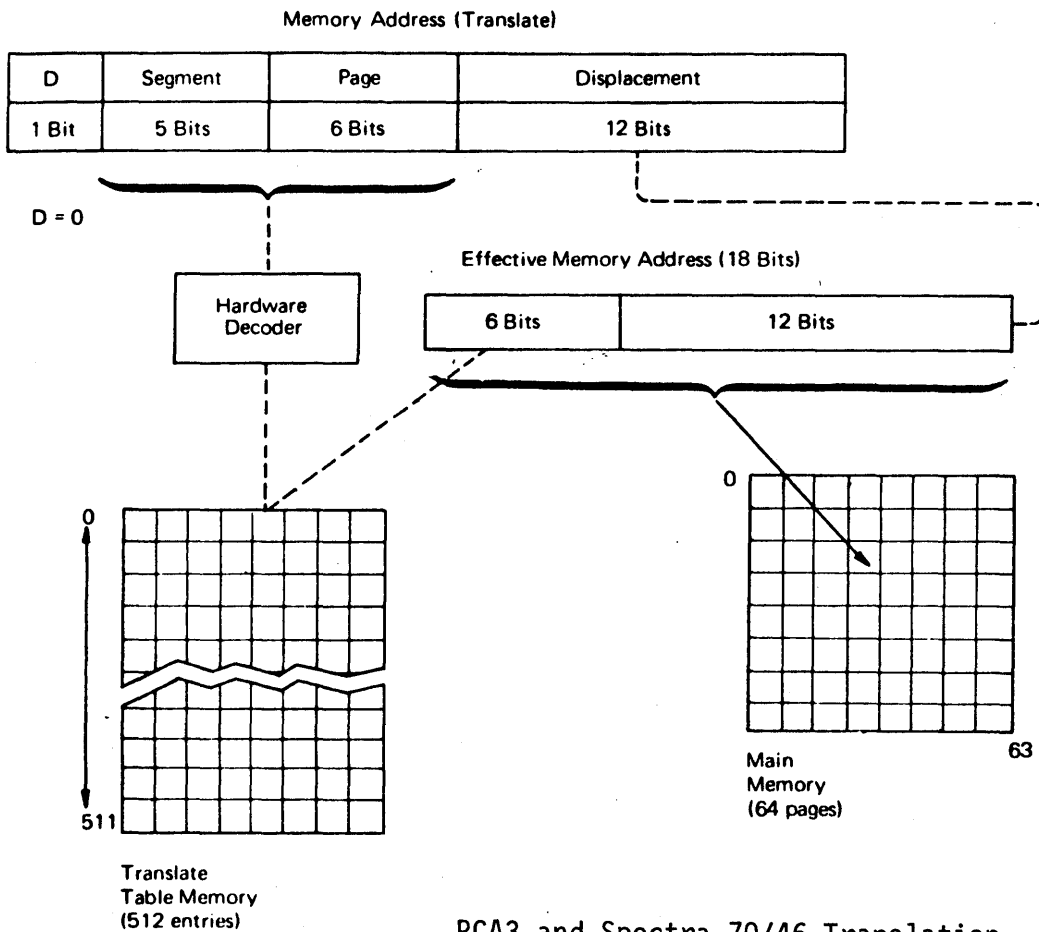


RCA3 and Spectra 70/46 Virtual Address (from Reference 3)

Example B



Example C



RCA3 and Spectra 70/46 Translation
(from Reference 3)