

HONEYWELL

MULTICS
THE MULTICS
VIRTUAL
MEMORY

GENERAL

MULTICS

SUBJECT:

Technical Papers on the Multics Virtual Memory as Designed for Model 645 Implementation.

SPECIAL INSTRUCTIONS:

These technical papers are reprinted here for theoretical and historical purposes only.

DATE:

June 1972

ORDER NUMBER:

AG95, Rev. 0

PREFACE

This document consists of three technical papers which describe the theory and practice of Multics virtual memory implementation. Multics (Multiplexed Information and Computing Service) is a general purpose computer system which has been designed to be a "computer utility." As such, it is essential that Multics provide its users with sufficient resources to do a wide variety of tasks and that the system be protected from destructive interactions between users. The papers address the theory, the practice, the hardware, and the software used to provide an effectively infinite memory to each user and to protect both the users and the system.

The first paper discusses the concept of a virtual memory and explores several ways in which such a memory could be implemented. The method used to implement the Multics virtual memory on the Series 600 Model 645 processor is developed in detail. This paper is of historical importance and presents valid Multics design theory although the Model 645 is no longer used as the Multics processor.

The second paper extends the discussion further into the subject of protection. The theory of the Multics ring structure is introduced and its implementation on the Model 645 is described. This theory is still valid although the Model 645 is no longer used as the Multics processor.

The third paper shows how the features described in the two earlier papers are handled by hardware under the optional Multics modifications to Series 6000 processors. Several new processor features are introduced and described. Use of these features allows Multics to run on the Series 6000 processors, specifically the Model 6180, with greatly increased efficiency as compared with the earlier implementation of Multics on the Series 600 processors.

The information and specifications in this document are subject to change without notice. This document contains information about Honeywell products or services that may not be available outside the United States. Consult your Honeywell Marketing Representative.

An alphabetized list of abbreviations and acronyms used in all three papers has been included as an aid to the reader.

Papers in this document were written to further the understanding of the Multics design philosophy and practices. They are not intended to be specifications of the Multics system or its components. Authors have made simplifying assumptions at times to make the main point clearer and easier to understand. Persons requiring design specification details are requested to contact the Multics development staff for guidance and assistance.

"The Multics Virtual Memory" was first published as Technical Information Series Report R69LSD3, Copyright 1970 by General Electric Company, U.S. A.

"Access Control to the Multics Virtual Memory" was first published as Technical Information Series Report R69LSD4, Copyright 1970 by General Electric Company, U.S. A.

CONTENTS

	Page
A. The Multics Virtual Memory	v
B. Access Control to the Multics Virtual Memory	119
C. Series 6000 Features for the Multics Virtual Memory	165
D. Abbreviations and Acronyms	191

A. The Multics Virtual Memory

CONTENTS

<u>Chapter</u>	<u>Title</u>
1	General Properties of the Multics Virtual Memory
2	Overview of the Implementation
3	Directory Structure
4	Making a Segment Known to a Process
5	Segment Fault Handling
6	Page Fault Handling
7	Secondary Storage Management
8	Device Interface Modules

PREFACE

In the past few years several well-known systems have implemented large virtual memories which permit the execution of programs exceeding the size of available core memory. These implementations have been achieved by demand paging in the Atlas computer, allowing a program to be divided physically into pages only some of which need reside in core storage at any one time, by segmentation in the B5000 computer allowing a program to be divided logically into segments, only some of which need be in core, and by a combination of both segmentation and paging in the 645 and the IBM 360/67 for which only a few pages of a few segments need be available in core while a program is running.

As experience has been gained with remote-access, multiprogrammed systems, however, it has become apparent that, in addition to being able to take advantage of the direct addressability of large amounts of information made possible by large virtual memories, many applications also require the rapid but controlled sharing of information stored on-line at the central facility. In Multics (Multiplexed Information and Computing Service), segmentation provides a generalized basis for the direct accessing and sharing of on-line information by satisfying two design goals: 1) it must be possible for all on-line information stored in the system to be addressed directly by a processor and hence referenced directly by any computation. 2) it must be possible to control access, at each reference, to all on-line information in the system.

The fundamental advantage of direct addressability is that information copying is no longer mandatory. Since all instructions and data items in the system are processor-addressible, duplication of procedures and data is unnecessary. This means, for example, that core images of programs need not be prepared by loading and binding together copies of procedures before execution; instead, the original procedures may be used directly in a computation. Also, partial copies of data files need not be read, via requests to an I/O system, into core buffers for subsequent use and then returned, by means of another I/O request, to their original locations; instead the central processor executing a computation can directly address just those required data items in the original version of the file. This kind of access to information promises a very attractive reduction in program complexity for the programmer.

If all on-line information in the system may be addressed directly by any computation, it becomes imperative to be able to limit or control access to this information both for the self-protection of a computation from its own mishaps, and for the mutual protection of computations using the same system hardware facilities. Thus it becomes desirable to compartmentalize or package all information in a directly-addressible memory and to attach to these information packages access attributes describing the fashion in which each user may reference the contained data and procedures. Since all such information is processor-addressible, the access attributes of the referencing user must be enforced upon each processor reference to any information package.

Given the ability to directly address all on-line information in the system, thereby eliminating the need for copying data and procedures, and given the ability to control access to this information, then controlled information sharing among several computations follows as a natural consequence.

In Multics, segments are packages of information which are directly addressed and which are accessed in a controlled fashion. Associated with each segment is a set of access attributes for each user who may access the segment. These attributes are checked by hardware upon each segment reference by any user. Furthermore all on-line information in a Multics installation can be directly referenced as segments while in other systems most on-line information is referenced as files.

Chapter 1

GENERAL PROPERTIES OF THE MULTICS VIRTUAL MEMORY

1. INTRODUCTION

In recent literature the term "virtual memory" has become quite familiar. The adjective "virtual" suggests that this memory is the image of an ideal memory that one would like to have, since it complies with the actual needs of a multi-programming, multiple-access computer utility. This "ideal memory" is not available as a hardware device and has been simulated by the Multics system using a conventional memory with the assistance of additional hardware and software features.

This chapter describes the properties of the ideal memory, justifies the desire for these properties, and explains the principles of the simulation of this memory.

2. THE IDEAL MEMORY

In order to describe this ideal memory the terms "segment" and "segmented memory" need to be defined first.

2.1. Segments

A segment is an entity defined by:

- 1) A name which uniquely identifies the segment.
- 2) A descriptor which describes the properties or "attributes" associated with the segment.
- 3) A body which is an array of consecutive elements.

The name is a character string of arbitrary length.

The descriptor contains all attributes the system designer needs to attach to the segment: the size and the physical location of the body, access rights for different users with respect to this segment, the date it was created, etc.

The body of the segment is an ordered set of elements, called words, each of which is identified within the segment body by an integer i , its index. The number of elements in the body is called the length of the segment.

2.2. Segmented Memory

A segmented memory will be defined as a memory with the following properties:

- 1) It is capable of containing segments and only segments.
- 2) If it contains a segment named n , then n is the address of the descriptor of this segment and the pair $[n,i]$ is the address of the i th element in the body of this segment.
- 3) It is capable of performing operations on the descriptor and the body of any segment, in accordance with the attributes recorded in the descriptor.

2.3. Ideal Memory

The ideal memory can now be defined as a large, segmented memory directly accessible by the processor, where by "large" it is meant that the maximum number of segments that one can store in it is adequate for the needs of the system.

A simple representation of such a memory is shown in Figure 1; it comprises a memory controller (MC), a large number of descriptors each of which contains the name and the attributes of a segment, and a large number of linear memories each of which is connected to a descriptor and can contain the body of a segment.

The processor can send two types of requests to the MC: requests for operations on descriptors and requests for operations on bodies. In both cases the processor must communicate to the MC the identification of the user on behalf of whom the operation is requested.

2.4. Operations on Descriptors

The general form of a request sent by the processor to the MC for operations on descriptors is

OPCODE n arguments userid

where:

- OPCODE designates operations, such as "create a segment", "change the length of a segment", "change access rights";

- n is the name of the segment. The MC uses it to locate the appropriate segment descriptor;
- arguments are parameters associated with the function defined by OPCODE;
- userid is the identification of the user on behalf of whom the operation is requested. The MC uses this userid in order to determine from the attributes of the segments whether this particular user has the right to perform this particular operation.

2.5. Operations on Segment Bodies

The general form of a request sent by the processor to the MC for operations on segment bodies is

OPCODE [n,i] userid

where:

- OPCODE designates operations, such as "read", "write", "instruction fetch";
- n is the name of the segment; the MC uses it to locate the appropriate segment descriptor. It then uses the segment descriptor to locate the segment body;
- i is the index of the word within the segment to which the operation is to be applied;
- userid is used by the MC as above.




Figure 1. Organization of the Ideal Memory

3. JUSTIFICATION OF THE IDEAL MEMORY PROPERTIES

The ideal memory has been defined as a "large segmented memory directly accessible by the processor". The advantages of such an ideal memory will be explained by successively introducing the advantages of a memory, that is: 1) large (but not segmented); 2) segmented (but not large); and finally, 3) large and segmented.

3.1. Large, Unsegmented Memory

Because the memory is large and directly accessible by the processor, the user is provided with a core memory large enough for any of his computations. Therefore, he can run a program without being concerned with its size. However, no matter how large the core memory is, if it is a linear memory accessible by a single number, no sharing of information in core can be tolerated between programs of different users since no protection mechanism is in effect at the time a word is accessed.

3.2. Small, Segmented Memory

Because the memory is segmented and directly accessible by the processor, the user is provided with several independent linear core memories in each of which he can store one of his segments, deciding who can access it and how. Therefore, the same segment can be shared in core by several user programs without the danger of unauthorized accessed to this segment. However, even though the memory is segmented, if the number of segments that one can store in it is small, the user is faced with the problem of overlays.

3.3. Large, Segmented Memory

By having the two properties "large" and "segmented" a directly accessible memory provides the user with:

- a large machine-independent memory. There is a one-to-one correspondence between the name by which the user references a one-word datum and the physical location in memory where the datum resides. As a consequence, users are provided with a simple means of writing programs such that, when executed, they access common information in core. They merely have to reference this information by its name.

- a protection mechanism. This mechanism is in effect during execution at any memory access and protects segments from unauthorized access.

3.4. Note on Information Sharing

It is worth making some remarks about information sharing. Information to be shared consists of data and procedures.

- Sharing data or procedures in core requires:
 - a) A mechanism by which a reference to a segment by its name X will cause segment X, and not a copy segment X, to be referenced during program execution.
 - b) A mechanism by which the shared information can be protected from unauthorized access while it is in core.
- Sharing procedures in core also requires:
 - c) A mechanism by which one can produce pure procedures that can be executed simultaneously by several programs.

The memory described here provides a) and b), but not c). In fact the memory itself cannot provide c); writing a pure procedure implies the ability of communicating as parameters to this procedure the names of any information private to the program on behalf of which the procedure is executing. These names cannot be stored in the memory itself; they have to be stored in processor registers whose names are invariant. During execution of a pure procedure by a processor on behalf of a program, the names of data segments private to the program are stored in processor registers whose names are stored in the pure procedure. The processor requests the data from the memory controller using the name found in the appropriate processor register.

4. PRINCIPLE OF THE SIMULATION

The memory presented here is simulated in the Multics system, this simulation being achieved by a combination of hardware and software features. Hardware segmentation has been implemented in the 645 and constitutes the most important of the hardware features mentioned above. Paging has also been implemented in the 645; although of immense help to the implementation, we do not regard paging as a concept fundamental to a description of the principles of the ideal memory simulation and shall postpone the discussion of paging until the end of this section.

Let us first examine how much of the ideal memory capability has been integrated into the hardware. Then a discussion of the software functions needed to compensate for those capabilities which are not provided by the hardware will follow.

4.1. Hardware Segmentation in the 645

Concepts of segment name, segment descriptor, and segment body have been integrated into the hardware as follows.

4.1.1. Segment Names. A segment name for the hardware is an integer s , called segment number, such that $0 < s < 2^{18}$.

4.1.2. Segment Descriptors. The segment descriptor of segment " s " is the s th entry of a table called a Descriptor Segment. The descriptor segment is in core memory and its absolute address is kept in a processor register. A descriptor segment entry is called a Segment Descriptor Word (SDW). SDW number s will be designated by the notation SDW(s).

Attributes that can be recorded in an SDW area are:

- The absolute core address of the head of the segment body.
- The length of the segment body.
- Access rights for only one user with respect to the segment body.
- An invalid attribute flag. F , which, when ON, signals the absence of the above attributes in the SDW and causes the processor to fault.

Since an SDW can contain access rights for one user only, each user program must be provided with a private descriptor segment. (See Figure 2.)




Figure 2. Hardware Segment Descriptors

4.1.3. Segment Bodies. The segment body is an array of contiguous words in core memory and its maximum length is 2^{18} words.

4.1.4. Address Transformations. Word number i of the body of segment s is addressed by the pair $[s,i]$ and is accessed through $SDW(s)$ by the processor.

Provided that the absolute core address m_0 of word 0 of the segment is stored in $SDW(s)$, the processor transforms -

- the processor segment name s into the core memory address m_0 using the descriptor segment which provides the mapping $m_0 = Z(s)$.
- the processor address $[s,i]$ into the core memory address m_i by the translation $m_i = m_0 + i$, that is $m_i = Z(s) + i$.

4.1.5. Access Rights Checking. Before accessing word m_i the processor performs a check on -

- the length of the segment by comparing i to the length recorded in $SDW(s)$.
- The access rights for the user with respect to segment s by using the access rights recorded in $SDW(s)$.

This hardware organization presents the following advantages over more conventional hardware.

- The set of processor addresses $[s,i]$ is sufficiently large that all words referenced by a program can be assigned unique processor addresses. The user does not have to organize a large program into overlays provided that he uses no more than 2^{18} segments.
- Processor addresses are independent of physical memory addresses. Addresses which appear in the instructions of a program are invariant when segments are moved from one location to another in core memory.
- Each access to core memory is subject to access rights checking.

However, the hardware has only a restricted understanding of the concept of segments and needs to be complemented by appropriate software features.

4.2. Software Segmentation

Given the foregoing hardware segmentation capabilities, the corresponding software segmentation capabilities required to implement the Multics virtual memory can be described.

4.2.1. Segment Names. A segment name is a character string called a symbolic segment name. The set of symbolic segment names is larger than 2^{18} . Therefore, the supervisor must map a large set of symbolic segment names into a smaller set of segment numbers.

4.2.2. Segment Descriptors. The hardware does not permit one to -

- retrieve attributes of a segment given the symbolic name of the segment. The software provides this capability.
- store all attributes of a segment in a hardware segment descriptor or SDW. The software provides complete segment descriptors for each segment and stores them in a catalog. See Figure 3.

Segment name	Segment Attributes			
a				
b				
c	core/secondary address	length	Access for user 1 Access for user 2 Access for user 3	other attributes
d				

Figure 3. Representation of a Catalog

4.2.3. Segment Bodies. The body of a segment is an array of contiguous words in core or in secondary memory. Since the processor can fetch data and instructions only from core-resident segments, the software must intercede when a segment is found to be missing from core.

4.2.4. Address Transformations. Assuming for the moment that all segments are in core memory, the supervisor performs the following three transformations to make segments accessible by the processor.

First, for any segment in the system, the supervisor must provide a one-to-one mapping from its symbolic name n into its memory address m_0 , where m_0 is the address of the beginning of the segment. This mapping $m_0 = X(n)^0$ is recorded in the catalog. See Figure 4.

Next, for each segment referenced by a user program, the supervisor must provide a one-to-one mapping from its symbolic name n into the segment number s assigned to it in this user program. This mapping $s = Y(n)$ is recorded in a table associated with the user program and called the Known Segment Table (KST).

Finally, for each segment that has been assigned a segment number s in a user program, the supervisor must provide a one-to-one mapping between the segment number s and its memory address m_0 . This mapping $m_0 = Z(s)$ is recorded in the descriptor segment associated with the user program.

The transformation X is independent of the user program; transformations Y and Z for user program u are user-dependent and will be denoted as Y_u and Z_u .

In order to permit several user programs to share the same segment by merely referencing it by the same name, these transformations must be such that, for any user program, $X(n) = Z_u(Y_u(n))$.




Figure 4. Address Mapping Tables

To this point we have assumed that all segments are in core. In fact, core memory being limited, the supervisor has to move segments between core and secondary memory.

The transportation of segment n from core memory address m_0 to secondary memory address M_0 must be associated with the following address mapping modifications:

- m_0 must be replaced by M_0 in the catalog entry for n .
- m_0 must be replaced by an undefined value in any SDW in which it appears. This is done by setting the invalid attribute flag ON in the SDW.

Note that the mapping between n and s remains unchanged in any user program.

A subsequent reference to segment n by segment number s in a user program will cause the processor to fault since the invalid attribute flag is ON in SDW(s). This fault will be referred to as a missing segment fault. Using the KST associated with this user program, it is possible to determine the name n of the segment s . Knowing n , the catalog entry for n can be found. The segment must be moved from secondary memory address M_0 to some (generally different) core memory address m'_0 . This move must be associated with the following address mapping modifications:

- M_0 must be replaced by m'_0 in the catalog entry for n .
- The undefined value (Flag) in SDW(s) must be replaced by m'_0 .

Note again that the mapping between n and s remains unchanged by the move.

4.2.5. Access Rights Checking. We have seen how the supervisor responds to a missing segment fault occurring in a user program but the description was not complete. A missing segment fault is a signal to evaluate the segment attributes in a specific SDW. Only the evaluation of the core address attribute has been described. Moreover, when the supervisor extracts core address information from the catalog, it also extracts the length and access rights attributes and stores them in the SDW. Each subsequent hardware reference to the segment by this user program is made through the SDW with the hardware performing access checking.

However, when performing operations on segment attributes the supervisor itself must do the necessary validation for any operation requested by a particular user since the hardware does not provide for access checking on attributes.

4.3. Paging

In a system in which the maximum size of any segment were very small compared to the size of the entire core memory, the "swapping" of complete segments into and out of core would be feasible. Even in such a system, if all segments did not have the same maximum size, or had the same maximum size but were allowed to grow from initially smaller sizes, there remains the difficult core management problem of providing space for segments of different sizes.

Multics, however, provides for segments of sufficient maximum size that only a few can be entirely core-resident at any one time. Also, these segments can grow from any initial size smaller than the maximum permissible size.

By breaking segments into equal-sized parts called pages and providing for the transportation of individual pages to and from core as demand dictates, several practical problems encountered in the implementation of a segmented virtual memory are solved.

First, since only the referenced page of a segment need be in core at one instant, segments need not be small compared to core memory.

Second, "demand paging" permits advantage to be taken of any locality of references peculiar to a program by transporting to core only those pages of segments which are currently needed. Any additional overhead associated with demand paging should of course be weighed against the alternative inefficiencies associated with dedicating core to entire segments which have been swapped into core but which may be only partly referenced.

Finally, since pages are all of equal size, space allocation is immensely simplified. The "compaction" of information in core and on secondary storage characteristic of systems dealing with variable-sized segments or pages is thereby eliminated.

The basic principles of paging in the Multics virtual memory may be briefly summarized as follows.

When a segment is not paged, the memory location of its element i is defined by relation (1), where m_0 is the memory location of element 0.

$$(1) \quad m_i = m_0 + i$$

When a segment is paged into pages of K elements, the memory location of its element i is defined by relation (2), where m_{pK} is the memory location of element pK ; that is, the memory location of the page number p of the segment.

$$(2) \quad \begin{cases} m_i = m_{pK} + j \\ j = i \bmod K \\ p = (i-j)/K \end{cases}$$

If N is the number of pages in a segment, paging this segment requires -

- a segment map with N entries, one for each page.
- a relocation capability in the hardware.

In the 645 the N entries of the segment map are provided by a "page table" and the relocation is performed by the processor itself. Furthermore, a page table entry contains a missing-page flag such that, if found ON by the processor while attempting to perform relocation, causes the processor to trap to the supervisor.

The missing-page flag is ON when the corresponding page is not in core. When, upon attempting to access a missing page, the processor traps to the supervisor, the supervisor must move the requested page into core. In order to do so the supervisor must maintain a segment map of N entries in the software descriptor, i.e., in the directory entry. Each time page p is moved from one location to another, this move must be associated with the following address mapping modifications.

- Update the mapping in entry p of the segment map located in the directory entry.
- Update the mapping in entry p of the page table.

Although paging need not be considered essential to a description of the simulation principles of an ideal memory, it is a basic feature for the implementation of such a memory.

The next chapter describes in some detail how the ideal memory has been simulated in the Multics system, using hardware segmentation and hardware paging as implemented on the GE-645.

Chapter 2

IMPLEMENTATION OF THE MULTICS VIRTUAL MEMORY: OVERVIEW

1. INTRODUCTION

As we have seen in Chapter 1, the Multics virtual memory is a large, segmented memory. Each segment can be referenced by its name in a user program; a reference by name will cause the segment to be accessed by the processor according to the access rights of the user with respect to that segment. The memory is called "virtual" because it is not available as a hardware device. Instead, it is simulated using a conventional non-segmented memory, a set of processor registers which provide the second dimension of a segmented memory and a supervisor which compensates for the difference in capabilities between the 645 hardware and the ideal memory described in Chapter 1.

Although the hardware checks each user's access rights to a segment whenever it accesses that segment, a certain number of additional functions must be provided by the supervisor in order to give the illusion that all segments are directly accessible by name by the processor.

- The hardware cannot retrieve the attributes of a segment using its symbolic name; the supervisor organizes segment attributes into "directories" where it can retrieve them.
- The hardware cannot interpret access rights for segment attributes; all operations on segment attributes are done by the supervisor.
- The hardware cannot reference a segment by a symbolic name; it does it by a segment number. The supervisor translates all symbolic segment names into segment numbers.
- The hardware cannot access a segment if it is not in core memory; each reference to a segment which is not in core will cause the supervisor to move the segment from secondary memory to core memory. In order to help the supervisor in core memory allocation, the hardware provides a paging capability.

This chapter builds upon the ideas developed in Chapter 1 to show in some detail how the ideal memory is simulated. The major topics covered are:

- Segmentation and paging on the 645 processor.
- The organization of segment attributes into hierarchically ordered directories and the manipulation of these attributes by the supervisor.
- Segment accessing and all the supervisory functions needed to make a segment directly accessible by the processor.
- The structure of the supervisor itself, showing how parts of the supervisor are able to utilize the virtual memory provided for user programs.

2. THE 645 PROCESSOR

This paper discusses only those features of the 645 processor which are of interest for the implementation of a virtual memory. They can be grouped into two different classes -- segmentation and paging -- and are treated separately below.

2.1. Segmentation

Any address in the 645 consists of a pair of integers $[s,i]$. The range of s and i is 0 to $2^{18}-1$. s is called the segment number, i the index within the segment. Word $[s,i]$ is accessed through a hardware register which is the s^{th} word in a table called a descriptor segment (DS). This descriptor segment is in core memory and its absolute address is recorded in a hardware register called a descriptor base register (DBR). Each word of the DS is called a segment descriptor word (SDW); the s^{th} SDW will be referred to as SDW(s). See Figure 1.

The DBR contains the following values:

- DBR.core which is the absolute core address of the DS.
- DBR.L which is the length of the DS.

Segment descriptor word number s contains the following values:

- $SDW(s).core$ which is the absolute address of the segment s .
- $SDW(s).L$ which is the length of the segment s .
- $SDW(s).acc$ which describes the access rights for the segment.
- $SDW(s).F$ which is a flag that can be ON or OFF. This is the invalid attribute flag mentioned in Chapter 1.

The algorithm used by the hardware for executing an instruction of the type $OPCODE [s,i]$ is as follows:

- If $DBR.L < s$, generate a fault.
- Access $SDW(s)$ at absolute location $DBR.core + s$.
- If $SDW(s).F = ON$, generate a missing segment fault.
- If $SDW(s).L < i$, generate a fault.
- If $SDW(s).acc$ is incompatible with $OPCODE$, generate a fault.
- Apply $OPCODE$ to the word whose absolute address is $SDW(s).core + i$.




Figure 1. Hardware Segmentation in the 645

The above description assumes that segments are not paged; in fact, paging is implemented in the 645 hardware.

2.2 Paging

A bit in an SDW indicates whether the corresponding segment is paged or not. Another bit in the SDW indicates whether the page size is 64 or 1024 words. Analogous bits in the DBR serve the same purpose for the descriptor segment.

However, in the Multics implementation, all segments are paged and the page size is always 1024 words. Therefore, this description makes the following two assumptions:

- All segments are paged.
- The page size is a constant, K , equal to 1024 words.

No further reference will be made to these two bits in the SDW and DBR.

Element i of a segment is the y^{th} word of the x^{th} page of the segment, x and y being defined by:

$$\begin{cases} y = i \bmod K \\ x = (i-y)/K \end{cases}$$

where K is the page size.

Since $K = 1024 = 2^{10}$, the processor can compute x and y from the 18 bit-binary representation of i by merely dividing i into two parts. The right part, which consists of the 10 least significant bits of i , represents the binary value of y ; the left part, which consists of the 8 most significant bits of i , represents the binary value of x . See Figure 2.




Figure 2. Hardware Interpretation of the Word Number

The page table (PT) of a segment is an array of physically contiguous words in core memory. Each element of this array is called a page table word (PTW).

Page table word number x contains the following items.

- $PTW(x).core$ which is the absolute core address of page $\#x$.
- $PTW(x).F$ which is a flag that can be ON or OFF. This is the missing page flag mentioned in Chapter 1.

The meaning of $DBR.core$ and $SDW(s).core$ is now as follows:

- $DBR.core$ = Absolute address of the PT of the DS.
- $SDW(s).core$ = Absolute address of the PT of segment $\#s$.

The full algorithm used by the hardware to access word $[s,i]$ is (see Figure 3):

- If $DBR.L < s$, generate a fault.
- Split s into s_x and s_y such that $s_y = s \bmod K$ and $s_x = (s - s_y)/K$.
- Access $PTW(s_x)$ at absolute location $DBR.core + s_x$.
- If $PTW(s_x).F = ON$, generate a missing page fault.
- Access $SDW(s)$ at absolute location $PTW(s_x).core + s_y$.
- If $SDW(s).F = ON$, generate a missing segment fault.
- If $SDW(s).L < i$, generate a fault.
- If $SDW(s).acc$ is incompatible with $OPCODE$, generate a fault.
- Split i into i_x and i_y such that $i_y = i \bmod K$ and $i_x = (i - i_y)/K$.
- Access $PTW(i_x)$ at absolute location $SDW(s).core + i_x$.
- If $PTW(i_x).F = ON$, generate a missing page fault.
- Apply the $OPCODE$ to the word whose absolute location is $PTW(i_x).core + i_y$.




Figure 3. Hardware Segmentation and Paging in the 645

3. SEGMENT ATTRIBUTES

3.1. Directory Hierarchy

The association between the name of a segment and its attributes is recorded in a catalogue. This catalogue consists of a table with one entry for each segment in the system. An entry contains the name of the segment and all its attributes (length, memory address, list of users allowed to use that segment with their respective access rights, date the segment was created, etc.).

In Multics this catalogue is divided into several segments called directories, which are organized into a tree structure. A naming convention permits one to search the tree structure for a given name without having to search all directories.

A segment name is a list of subnames reflecting the position of the entry in the tree structure, with respect to the beginning of the tree or root directory. By convention, subnames are separated by the character ">". Each subname is called an entryname and the list of entrynames is called a pathname.

There are two types of directory entries called branches and links. A branch is a directory entry which contains all attributes of a segment while a link is a directory entry which contains the pathname of another directory entry. This chapter will deal only with entries of the branch type.

The pathname is the only name by which a segment can be searched for in the directory hierarchy.


The attributes associated with a segment whose pathname is ROOT > A > B > C are found as follows (see Figure 4):

- Search the root directory for an entry whose entry name is A. This entry contains attributes for the directory segment whose pathname is ROOT > A. These attributes permit one to locate the directory ROOT > A in memory.

- Search directory ROOT > A for an entry whose entry name is B. This entry contains attributes for the directory segment whose pathname is ROOT > A > B. These attributes permit one to locate the directory ROOT > A > B in memory.

- Search directory ROOT > A > B for an entry whose entry name is C. This entry contains attributes for the segment ROOT > A > B > C.

ROOT > A > B > C



Squares are directory segments.
Circles are non-directory segments.

Figure 4. Directory Hierarchy

3.2. Operations on Segment Attributes

All operations on segment attributes are done by supervisor primitives. There is a set of primitives available to the user which allow him, for example, to:

- Create a segment.
- Delete a segment.
- Change the entryname of a directory entry.
- Change the access rights of a segment.
- List a directory.

Any of these operations is performed on behalf of a user by the supervisor only if the user has the right to perform them.

Some further details about one of these operations, segment creation, are important to an understanding of the topic of segment accessing developed in the next section.

Creating a segment whose pathname is ROOT > A > B > C consists basically of taking the following actions:

- Check, by searching the directory hierarchy, that this segment does not exist already in the system.
- Allocate space for a branch in directory
ROOT > A > B.
- Store in the branch the following items:
 - . The entry name C.
 - . The access list, given by the creator.
 - . The segment map which consists of a secondary storage address for each page of the segment. This segment map is manufactured by the supervisor.
 - . The segment status "inactive", meaning that there is no page table for this segment.

Once the segment has been created, the user can reference it. Note that no segment number has been assigned to the segment at creation time, and the only way to refer to it is by the pathname.

4. SEGMENT ACCESSING

We are now in a position to understand a description of the functions that are provided by the supervisor in order to make accessible by the processor segments which are referenced by name in a user program. Figure 5 is key to an appreciation of the Multics virtual memory implementation. Although frequent references to Figure 5 follow, the full implications of its contents will not be apparent until the entire section has been read.

4.1. Concept of Process and Address Space

A process is generally understood as being a program in execution. A process is characterized by its state-word defining, at any given instant, the history resulting from the execution of the program.

A process is also characterized by its address space. The address space is the set of processor addresses that this process can use to reference the memory. In Multics the address space of a process is defined as the set of segment numbers that the process can use to reference segments in the virtual memory. As explained in Chapter 1, a segment number can be used to reference the virtual memory only if it has been associated with a segment name, i.e., a pathname. This association [pathname, segment number] is recorded in a table called the Known Segment Table (KST) which defines the address space of the process. There is a one-to-one correspondence between Multics processes and address spaces. The action of adding a new pair [pathname, segment number] in a KST is referred to as making the segment with that pathname known to the process.

4.2. Making a Segment Known to a Process

Each time a segment is referenced in a process by its pathname, the pathname must be translated into a segment number in order to permit the processor to address the segment. This translation is done by the supervisor using the KST associated with the process. The KST is an array organized such that the entry number "s", KSTE(s), contains the pathname associated with segment number s. See Figure 5.

If the association [pathname, segment number] is found in the KST for this process, then the segment is known to the process and the segment number can be used to reference the segment.

If the association [pathname, segment number] is not found it means that this is the first time the segment is referenced in the process and the segment has to be made known. This is done by assigning an unused segment number "s" in the process and by establishing the pair [pathname, segment number] in the KST by recording the pathname in KSTE(s). Furthermore, the directory hierarchy is searched for this pathname and a pointer to the corresponding branch is entered in KSTE(s) for later use (see Section 4.3.).

This stage is fundamental because, in the Multics system, it is impossible to assign a unique segment number to each segment. The reason is that the number of segments in the system may be larger than the number of segment numbers available in the processor.

When a segment is made known to a process by segment number "s" its attributes are not placed in SDW(s) of the descriptor segment of that process. SDW(s) has been initialized with an invalid attribute flag. Therefore, the first reference in this process to that segment by segment number "s" will cause the processor to generate a fault. In Multics this fault is called a "missing segment fault" and transfers control to a supervisor module called the segment fault handler.

4.3. The Segment Fault Handler

Upon the occurrence of a missing segment fault, control is passed to the segment fault handler whose function is to store the proper segment attributes in the appropriate SDW and to set the invalid attribute flag OFF in the SDW.

This information, we recall, consists of:


- The page table address.
- The length of the segment.
- The access rights of the user with respect to the segment.

The information initially available to the supervisor upon occurrence of a missing segment fault is:

- The segment number s .
- The process identification.

The only place where the needed information can be found is in the branch of the segment. Using the process identification, the supervisor can find the KST for this process. It can then search this KST for the segment number s . Having found the KST entry for s , it can find the required branch since a pointer to the branch has been stored in the KST entry when the segment was made known to that process. See Section 4.2.

Using the active switch (see Figure 5) in the branch, the supervisor can determine whether or not there is a page table for this segment. Recall that this switch was initialized in the branch at segment creation time. If there is no page table, one must be constructed. A portion of core memory is permanently reserved for page tables. All page tables are of the same length and the number of them is determined at system initialization.



The supervisor divides these page tables into two lists: the "used list" and the "free list". Manufacturing a page table (PT) for a segment could consist only of selecting a PT from the free list, putting its absolute address in the branch and moving it from the free to the used list. If this were actually done, however, then the servicing of each missing page fault would require access to a branch since the segment map is kept there.

Since all directories cannot be core-resident, page fault handling could thereby require a secondary storage access in addition to the read required to transport the page itself into core. Although this mechanism works, efficiency considerations have led to the "activation" convention between the segment fault handler and the page fault handler.

4.3.1. Activation. A portion of core memory is permanently reserved for recording attributes needed by the page fault handler, i.e., the segment map and the segment length. This portion of core is referred to as the active segment table (AST). The AST contains one entry (ASTE) for any segment that has a PT. A PT is always associated with an ASTE, the address of one implying the address of the other. They may be regarded as a single entity and will be referred to as the [PT,ASTE] of a segment.

A segment which has a [PT,ASTE] is said to be active. The property of being active or not active is an attribute of the segment and, therefore, has to be recorded in the branch. When this active switch is set ON it means that both the segment map and the segment length are no longer in the branch but are to be found in the segment's [PT,ASTE] whose address has been recorded in the branch during "activation" of the segment.

To activate a segment the supervisor must:

- Find a free [PT,ASTE]. Assume temporarily that at least one is available.
- Move the segment map and the segment length from the branch into the ASTE.

- Set the active switch in the branch.
- Record the pointer to [PT,ASTE] in the branch.

Having defined activation, the actions taken up to now by the segment fault handler can be summarized as:

- Use the segment number *s* to access the KST entry.
- Use the KST entry to find the branch.
- If the active switch is OFF, activate the segment. If it is ON, then activation is unnecessary at this time as the segment was already activated for another process.

By pairing an ASTE with a PT in core, the segment fault handler has guaranteed that the segment attributes needed by the page fault handler are core-resident, thus permitting efficient page fault servicing.

4.3.2. Connection. Now that the segment is active, the corresponding SDW must be "connected" to the segment.

To connect the SDW to the segment the supervisor must:

- Get the absolute address of the PT, using the [PT,ASTE] pointer kept in the branch, and store it in the SDW.
- Get the segment length from the ASTE and store it in the SDW.
- Get the access rights for the user from the branch and store them in the SDW.
- Turn off the flag which caused the fault from the SDW.

Having defined activation and connection, segment fault handling can finally be summarized as:

- Use the segment number *s* to access the KST entry.
- Use the KST entry to find the branch.
- If the active switch in the branch is OFF, activate the segment.
- Connect the SDW.

Note that segment sharing in core is "automatically" guaranteed by the use of the active switch and [PT,ASTE] pointer kept in the segment branch since all SDW's describing this segment will point to the same PT.

Now that the segment has an SDW pointing to the PT, the hardware can access the appropriate page table word. If the page is not in core, a missing page fault occurs, transferring control to the supervisor module called the page fault handler.

4.4. The Page Fault Handler

When a page fault occurs the page fault handler is given control with the following information:

- The PT address.
- The page number.

The information needed to bring the page into memory is:

- The address of a free block of core memory into which the page can be moved.
- The address of the page in secondary memory.

A free block of core must be found. This is done by using a data base called the core map. The core map is an array of elements called core map entries (CME). The *n*th entry contains information about the *n*th block of core (the size of all blocks is 1024 words). The supervisor divides this core map in two lists; the used list and the free list.

The job of the page fault handler is to:

- Find a free block of core. (Assume temporarily that there is at least one free block in the free list.)
- Access the ASTE associated with the PT and find the address in secondary memory of the missing page.
- Issue an I/O request to move the page from secondary memory into the free block of core.
- Upon completion of the I/O request, store the core address in the PTW and remove the fault from the PTW.

4.5. Page Multiplexing

It was assumed that a free block of core was available in the core map free list; however, this is not always the case since there are many more pages in the virtual memory than there are blocks of core. Therefore, in order to get a free block of core, the page fault handler may have to move a page from core to secondary memory. This requires:

- An algorithm to select a page to be removed.
- Knowing the address of the PTW which holds the address of the selected page in order to set a fault in it.
- Knowing where to put the page in secondary memory.

The selection algorithm is based upon page usage. The hardware provides valuable assistance by the fact that, each time a page is accessed, a bit is set ON in the corresponding PTW. This bit is called the used bit.

The selection algorithm will not be described here; however, it should be noted that candidates for removal are those pages described in the core map used list. Therefore, each core map entry which appears in the used list must contain a pointer to the associated PTW in order to permit one to examine the used bit. The action of storing the PTW pointer in the core map entry must be added to the list of actions taken by the page fault handler when a page is moved into core (see Section 4.4.).

- A fault is stored in the PTW.
- The secondary storage address for the page is found in the ASTE whose address can be computed from the PTW address.
- An I/O request is issued to remove the page to secondary storage.
- Upon completion of the I/O request, the core map entry is removed from the used list and put in the free list.

By this mechanism, blocks of core are multiplexed among all pages of all active segments in the system.

It is important to realize that a page is either in core or in secondary storage. There is no such thing as a "copy" of a page. When a page is moved from secondary storage to core, its secondary storage address, located in the ASTE, could be freed; it is no longer needed since the address of the page is now in the PTW. When the page is to be removed, a free block of secondary storage could be assigned to it. It is only for practical reasons that the block of secondary storage is not freed each time a page is moved into core.

Page multiplexing maintains a "perpetual motion" between core and secondary storage of pages of active segments. If the set of active segments in the system were invariant, then pages of other segments would never have a chance to be in core.

4.6. [PT,ASTE] Multiplexing

In the description of segment fault handling, when a segment had to be activated, a pair [PT,ASTE] was assumed available for assignment to that segment. In fact, the number of [PT,ASTE] pairs is limited in the system and is, by far, smaller than the number of segments in the virtual memory. Therefore, these [PT,ASTE] pairs must be multiplexed among all segments in the virtual memory.

This means that making a segment active may imply making another segment inactive thereby disassociating this other segment from its [PT,ASTE]. Since each process sharing the same segment will have the address of the PT in an SDW it is essential to invalidate this address in all SDW's before removing the page table. It is also essential to move to secondary memory all pages of that segment which are in core before removing the [PT,ASTE], since the ASTE is needed to remove a page. Then, and only then, can [PT,ASTE] be disassociated from the segment.

This operation requires:

- An algorithm to select a segment to be deactivated.
- Knowing all SDW's that contain the address of the page table of the selected segment in order to invalidate this address.
- The removal of all pages of the selected segment that are still in core.
- Moving the attributes contained in the ASTE back to the branch and changing the status of the segment from active to inactive in the branch.

The selection algorithm is here again based on segment usage. The only thing of interest at this point is that selection is done by scanning the ASTE used list. Therefore, the ASTE must provide all the information needed for removing the [PT,ASTE]. This means that during activation and connection this information must be made available as explained below.

During activation, a pointer to the branch must be placed in the ASTE; during connection, a pointer to the SDW must be placed in the ASTE. Since more than one SDW is connected to the same PT when the segment is shared by several processes the supervisor must maintain a list of pointers to connected SDW's. This list is called a connection list. See Figure 5.

Now we are in a position to understand how a [PT,ASTE] can be disassociated from a segment. After the selection algorithm decides on an ASTE to be freed, actions to be taken consist of two steps called "disconnection" and "deactivation".

Disconnection consists of storing a segment fault in each SDW whose address appears in the connection list in the ASTE.

Deactivation consists of removing all pages of this segment that may be in core, moving the segment map from the ASTE back to the branch, resetting the active switch in the branch and putting the [PT,ASTE] in the free list.

4.7. Segment Number Multiplexing in a Process

The number of segments that a process can describe in its descriptor segment is limited to 2^{18} . It is unlikely that a process needs to access more than 2^{18} segments from the time it is created to the time it is destroyed. However, if this should happen, a facility is provided to a process to remove an association [pathname, segment-number] by an explicit call to the supervisor. This action is referred to as making a segment unknown. When segment A which is known to a process by the segment number "s" is made unknown to that process, no attempt is made by the supervisor to remove residual [s,i] pairs that may have been generated and stored during the time that s was assigned to A. Making segment A unknown to the process implies freeing KSTE(s). If subsequently another segment, say B, is made known to the process, the supervisor may assign this unused segment number s to segment B, entering the pathname B in KSTE(s). From this point on, any reference by segment number s in this process will cause segment B to be accessed. Therefore, it is entirely the responsibility of the programmer, after segment A is made unknown, not to reuse any residual pair [s,i] that was generated for accessing segment A.

4.8. Directory Entry Multiplexing

When a segment is deleted, the branch of that segment is deleted. No attempt is made by the supervisor to remove residual KST entries that contain a pointer to this branch. However, the supervisor can detect references by residual segment numbers, to segments which have been destroyed as follows:

- When segment A is created, the supervisor assigns a unique number N_A to segment A and stores it in the branch.
- When segment A is made known to a process P by the segment number s, N_A is copied from the branch to KSTE(s) for process P along with the pointer to branch A.
- If segment A is deleted by any process, the supervisor disconnects the corresponding SDW in process P, if it was connected, and deletes the branch together with N_A .
- If the same directory entry is reused to record the branch information of another segment B, a new unique identifier N_B will be stored in the branch.
- Now, if process P uses the segment number s in order to access segment A, a segment fault will occur; the KST entry contains a pointer to the directory entry which is supposed to be the branch A. But by comparing the N_A of the KST entry and N_B of the directory entry, the supervisor can discover that segment A has been deleted.

Therefore, it is possible to detect the deletion of a branch even though its former directory entry has been reused for another segment.

5. STRUCTURE OF THE SUPERVISOR

Up to now supervisor functions have been described, but supervisor structure has not been discussed. In this section, the different components of the supervisor are covered and the ability of portions of the supervisor to partially utilize the virtual memory is demonstrated.

5.1. Functional Modules

Three functional modules can be identified in the supervisor described; they are called directory control (DC), segment control (SC), and page control (PC).

5.1.1. Directory Control. Directory control is that part of the supervisor which can manipulate all segments in the system. DC identifies a segment by its pathname which uniquely defines a segment in the system. Data bases that are manipulated by DC are the directories and KST's of all processes (see Figure 6). DC provides all primitives to simulate operations on segment attributes; it also provides the assignment of a segment number to a segment within a process.

5.1.2. Segment Control. Segment control is that part of the supervisor which can manipulate only those segments which are known to at least one process. SC identifies a segment by either its segment number within a particular process, which uniquely defines a segment in the system, or by its [PT,ASTE] address which uniquely defines an active segment in the system. Data bases that are manipulated by SC are directories, KST's of all processes, descriptor segments of all processes and [PT,ASTE] pairs of all active segments. SC provides the functions of activation, connection, disconnection and deactivation.

5.1.3. Page Control. Page control is that part of the supervisor which can manipulate only those segments which are active. PC identifies a segment by its [PT,ASTE] address which uniquely defines an active segment in the system. Data bases that are manipulated by PC are [PT,ASTE] pairs of all active segments and the core map. PC provides the mechanism to move pages of active segments between secondary storage and core.




Figure 6. Supervisor Functional Modules and Data Bases

5.2. Use of Segmentation in the Supervisor

Previous to this no assumptions were made about the type of addressing used by the supervisor. It could be written so as not to use segment addressing of course; but organizing the supervisor into procedures and data segments permits one to use in the supervisor the same standard conventions that are used in a user program. For instance, the CALL-SAVE-RETURN conventions made for user programs can be used by the supervisor, the standard way to manufacture pure procedures in a user program can be used in the supervisor, etc. Thus, it seems desirable to use segmentation in the supervisor, and the following (temporary) assumption will be made:

Assumption 1:

- a. The address space of the supervisor is entirely defined by a descriptor segment.
- b. All segments used by the supervisor are always in core.

Assumption 1.b is not realistic, however, since it generally is not possible to dedicate enough core to contain the entire supervisor. It is, therefore, interesting to determine whether there is a way to use the page fault handler to transport supervisor as well as user pages.

5.3. Use of PC in the Supervisor

For the purpose of this paper, let us assume the validity of the following statement. "Page fault handling for a page x must be performed without referencing page x".

It is certainly possible to design a PC module which allows recursive page faults provided that the above condition is always satisfied. Each recursive invocation of the page fault handler should use a set of pages which does not include any of the pages that caused the previous invocations; furthermore, the number of recursive invocations must be guaranteed to be finite. The technique that has been chosen in Multics for page fault handling is to fix this finite number to 1 and thus no recursive page faults will ever occur. This decision has been made for reasons of efficiency and design simplicity. Therefore, it is assumed that all segments used in PC are always in core.

We can observe that the page fault handler need not know if a missing page belongs to a user or to the supervisor; it only expects to find the information it requires in the [PT,ASTE] of the segment to which the missing page belongs. Therefore, if all segments used in SC and DC are always active, then their pages need not be in core since PC can load them when they are referenced.

Thus, assumption 1 can be replaced by the following one (again temporary):

Assumption 2:

- a. The address space of the supervisor is entirely defined by a descriptor segment.
- b. All segments used in PC are always in core.
- c. All segments used in SC and DC must be active and connected.

This convention turned out to be satisfactory in the Multics implementation except for directories. Recall that segments used by SC and DC are: (a) SC and DC procedures, (b) KST's and DS's, and (c) Directories.

The number of segments in class (a) and (b) is relatively small. On the contrary, the number of directory segments may be very large and keeping them always active is not a realistic approach, since a large number of [PT,ASTE] pairs would have to be permanently assigned to them. Therefore, it is desirable to use SC to activate and connect directory segments.

5.4. Use of Segment Control in the Supervisor

A necessary condition for handling a segment fault for segment x in a process is that segment x be known to that process. If SC is to handle segment faults taken by the supervisor for directories, all directories must be known to the supervisor. This means that the address space of the supervisor must be defined not only by its descriptor segment but also by KST, which contains one entry for each directory. After its KST has been so initialized, the supervisor looks like any other process.

Assuming that all directories are known to the supervisor process, but not necessarily active, a supervisor reference to a directory x may cause a segment fault. Recall that when handling this fault, the segment fault handler must reference the parent directory of segment x , where the branch for x is located. This reference to the parent of x could, in turn, cause a recursive invocation of the segment fault handler. Recursive invocations can propagate from directory to parent directory up to the root. If there is a way of stopping the recursion, then any segment fault on directories can be handled.

One way of stopping the recursion is to keep the root active and connected, so that a segment fault never occurs for it.

Assumption 2 can now be replaced by assumption 3, again temporary.

Assumption 3:

- a. The address space of the supervisor process is defined by a descriptor segment and a KST.
- b. All segments used in PC are always in core.
- c. All segments used in SC and DC are always active and connected, except directories.

- d. The root directory is always active and connected.
- e. All directories are known to the supervisor process.

However, it is unsatisfactory to keep all directories known. We would like to keep known only those which may possibly be used in a segment fault handling, provided that other directories can be made known by directory control when needed.

5.5. Use of the Make Known Facility in the Supervisor

Making a segment x known implies searching for its pathname in the KST. If not found, the parent of x must first be made known and so on up to the root. If the root directory is always known to the supervisor, then any directory can be made known to the supervisor by the supervisor itself.

Assumption 3 will now be replaced by the final assumption:

Final Assumption:

- a. The address space of the supervisor process is defined by a descriptor segment and a KST.
- b. All segments used in PC are always in core.
- c. All segments used in SC and DC, except directories, are always active and connected.
- d. The root directory is always active and connected.
- e. If a segment is known to any process, its parent directory must be known to the supervisor process.
- f. The root directory is always known to the supervisor process.

Given the above assumptions, supervisor segments as well as user segments can be stored in the virtual memory that the supervisor provides.

5.6. The Supervisor Address Space

Unlike most supervisors, the Multics supervisor does not operate in a dedicated process or address space. Instead, the supervisor procedure and data segments are shared among all Multics processes. Whenever a new process is created, its descriptor segment is initialized with descriptors for all supervisor segments allowing the process to perform all of the basic supervisory functions for itself. The execution of the supervisor in the address space of each process facilitates communication between user procedures and supervisor procedures. For example, the user can call a supervisor procedure as if he were calling a normal user procedure. Also, the sharing of the Multics supervisor facilitates simultaneous execution, by several processes, of supervisory functions, just as the sharing of user procedures facilitates the simultaneous execution of functions written by users.

Since supervisor segments are in the address space of each process, they must be protected against unauthorized references by user programs. Multics provides the user with a ring protection mechanism which segregates the segments in his address space into several sets with different access privileges. The Multics supervisor takes advantage of the existence of this mechanism and uses it, rather than some other special mechanism, to protect itself.

6. SUMMARY

If only a few points discussed here were to be remembered, they should be those mentioned below. They have been separated into two classes: the point of view of the user of the virtual memory, and the point of view of the supervisor itself.

User Point of View

- The Multics virtual memory is capable of containing a very large number of segments that can be identified by their symbolic names.
- Segment attributes are stored in special segments called directories, which are organized into a tree structure; there is a naming convention, of which the user must be aware, by which a segment name must be the pathname of its branch in the directory tree structure.

- Any operation on directory segments must be done via a call to the supervisor.
- Any operation on a non-directory segment can be done directly in accordance with the access rights that the user has for this segment; any word of any segment which resides in the virtual memory can be referenced with a pair [pathname,i] by the user.
- A process can have only a limited number of segments in its address space. If the programmer wants to overlay a segment A by a segment B in the process address space, he can call the supervisor to do it but he must be aware of the dangers that this operation may present.

Supervisor Point of View

- The supervisor must simulate a large segmented memory directly addressable by segment name such that any access to the memory is submitted to access rights checking.
- It maintains a directory tree where it stores all segment attributes. It can retrieve the attributes of a segment given the pathname of that segment.
- The supervisor itself is organized into segments and runs in the user process address space.
- Any segment, be it a directory or a non-directory segment, is identified by its pathname but can be accessed only using a segment number. For each segment name the supervisor must assign a segment number by which the processor will address the segment in the process.
- The processor accesses a word of a segment through the appropriate SDW and PTW and subject to the access rights recorded in the SDW.
- A segment fault is generated by the processor whenever the page table address or access rights are missing in the SDW. The supervisor then, using the KST entry as a stepping stone, accesses the branch where it finds the needed information. If a PT is to be assigned, the supervisor may have to deactivate another segment.

- A page fault is generated by the processor whenever a PTW does not contain a core address. The supervisor then, using the ASTE associated with the PT, moves the missing page from secondary storage to core. This may require the removal of another page.

Chapter 3

DIRECTORY STRUCTURE

1. INTRODUCTION

A virtual memory system must include some means of storing and retrieving information. A segment is the unit of information in the Multics virtual memory which is so stored and retrieved.

All information about a segment such as its length and its location are called "attributes" of the segment. If the attributes of a segment can be located, then the segment itself can be found.

The attributes of segments are stored in special segments called "directories" and the directories are organized into a tree structure called "the directory hierarchy". All of the attributes of one segment are recorded in one entry in a directory. The entries in a directory can be referenced by literal string names called "entrynames".

The discussion which follows gives some of the details of the directory hierarchy structure, the naming of entries and segments and the contents of directory entries. Segment creation and deletion are also described since those operations are closely related to the creation and deletion of the attributes of a segment.

Other chapters describe the details of the search for a segment and how segments themselves are handled.

2. THE DIRECTORY HIERARCHY AND TERMINOLOGY

2.1. The Structure

A tree structured directory hierarchy is shown in Figure 1. Directory segments are shown as squares and non-directory segments are shown as circles. The lines between segments are branches of the tree structure and in Multics, denote the fact that the attributes of a segment at the lower end of a line are recorded in an entry in the directory at the upper end of the line. Thus the attributes of the segment labeled C in Figure 1, are recorded in the directory labeled B. The directory entries in which the attributes of segments are recorded are called "branch entries" or "branches".

A directory is said to be the parent of a segment if it contains the branch with the attributes of that segment. The parent directory of a segment is said to be "immediately superior to" the segment and the segment is "immediately inferior to" its parent. In Figure 1, the directory at the top of the structure labeled "root" and called "the root directory" or simply "the root" is the parent of or immediately superior to the segment labeled D. The root is superior to the segment labeled E, but not immediately superior to E. The segment labeled E is inferior to the root and the segment labeled D is immediately inferior to the root.

The root is the starting point in the search for segments. Note that the root has no branch. Its attributes, among them its location, are assumed to be known to the modules which perform the search.

There is one and only one branch per segment in the Multics system. This rule arose from the difficulty of finding and updating all the branches of a segment if one of its attributes should be changed. For cases in which it is useful to have a branch in more than one directory a link (see below) can be used.

2.2. Entrynames and Pathnames

An entryname is used to locate an entry in a given directory, but a "pathname" is needed to search the directory hierarchy for a particular entry. In order to uniquely locate a particular entry in a directory, an entryname must be unique in that directory. However, an entry can have several names (synonyms).

A pathname is the concatenation of an ordered sequence of entrynames. The entries must be located in the order they were named in order to follow the path from the root to the desired entry. The entrynames are separated by the character ">".

The name of a segment is the pathname which addresses its branch.




Figure 1. Tree structure directory hierarchy

The names root > A and root > D > E > Z are examples of pathnames for segments in the hierarchy shown in Figure 1. The name root > A > B > C > Z is the name of the segment pointed to by the vertical arrow. Since all pathnames begin with root >, the leading symbol > is used to mean root >. The names above become > A, > D > E > Z and > A > B > C > Z.

An example of the search for the attributes of the segment > B > Z in Figure 2 is instructive. The root is searched for a branch with the name B. The segment > B is accessed using the information in this branch. Next, the directory segment > B is searched for a branch with the name Z. When the branch Z is found in the directory > B then the search is finished.

It must be emphasized that the only way in which the search modules can find a segment is through use of a pathname.

A pathname can have synonyms since the entrynames from which it is constructed can have synonyms. However, a given pathname cannot lead to more than one segment.

2.3. Links

There can exist in a directory a second type of entry in addition to the branch entry. This is called a "link entry" or a "link". A link entry has a name just as branch entry does and like a branch is used to access a segment or its attributes. A link contains no attributes but only the "pathname" of another entry. In Figure 1, the dotted line labeled L is an example of a link. The entryname of the link is L. Loops such as might be generated by two links which reference each other should not be allowed in the directory hierarchy.

Links are addressed by pathname just as branches are. Referencing the link, > A > B > C > L, in Figure 1, will cause accessing of the segment > D > E > X, as that pathname is recorded in the link.

3. DESIGN CONSIDERATIONS

A directory is needed as a place to look up the addresses of other segments. Once a directory exists, there are other advantages to be gained from it. The directory is a

convenient place to store the access rights of a user to a segment so they may be checked at the same time that the address of the segment is located. It may also be useful to reference the attributes of a segment without necessarily accessing the segment body, e.g., to find the length of a segment. For these reasons, all of the attributes of a segment are collected into a list and the lists are stored together in a directory. Why then have more than one directory?




Figure 2. Directories and Attributes

The most important reason for having multiple directories is to avoid the problem of naming conflicts. It is very likely that many users will attempt to use the same names for the different segments they will create. Long and complicated unique names are difficult to remember and inconvenient to look up both for people and computers. If each user is assigned one or more directories then this naming conflict disappears. The key to this simplification is to allow the user to reference the segments of a pre-assigned directory by entryname. Pathnames are constructed by prefixing the directory pathname to the user given entryname and referencing the desired segments via these constructed pathnames.

There are other advantages to be gained with multiple directories. Directories can be used for classification of segments, e.g., all segments of a math library could be accessed through a math library directory, > math-library. Protection is aided since access to directories (and, therefore, complete classes of segments) can be restricted to specific users.

The tree structure carries two advantages. It allows an even better scheme of classification than a linear or limited level structure of directories and it also facilitates an efficient directed search for a given segment in the hierarchy.

4. INTERNAL DIRECTORY AND ENTRY STRUCTURE

Each directory has a small area at its beginning called a header which contains pointers to other areas and items in the directory. There is also in the header a count of the branches in the directory. The header is followed by an array of entries.

Figure 3 shows a directory with a branch entry and a link displayed in some detail. All entries contain the following:

- a name list pointer,
- a unique identifier and
- a branch or link switch.

If the entry is a link, it contains a pathname and no attributes. If the entry is a branch it also contains

- an access list pointer,
- a segment map,
- segment length,
- an active switch,
- a PT-ASTE pointer and
- a directory flag.




Figure 3 - Structure of a Directory

4.1. The Name List Pointer

A branch or link entry can have several names. These are stored in a threaded list. Supervisor primitives exist to add names to the list and to delete names from it. The name list pointer is a pointer to the threaded list of names of the entry.

4.2. The Unique Identifier

The unique identifier, uid, is permanently attached to the entry. It cannot be changed or destroyed while the entry is in use (contains valid branch or link information). All uid's are constructed in part from numbers representing the date and an instant during the creation of the entry, so that a uid can never be repeated. The uniqueness of the uid for each entry in the entire directory hierarchy is thus guaranteed. The uid is used to assure that the correct segment is being accessed when a segment is referenced or activated. Its use will be described in more detail in Chapters 4 and 5.

4.3. The Branch or Link Switch

The branch or link switch is used to tell the search modules the kind of information to be found in the entry.

4.4. Access List Pointer

The access list pointer is a pointer to a threaded list of user names and associated access rights. As an example, user Tom may have the right to execute the contents of a segment and to modify the contents of that segment while user Frank may only be permitted to execute the contents of the segment. The access list contains the names of all users permitted to access the segment, and is described in detail in a companion paper, "Access Control to the Multics Virtual Memory".

4.5. The Segment Map

As the name indicates, the segment map gives the address of the segment in secondary storage. The segment map consists of two parts, the device identifier, did, and a page address list. The did is a number which uniquely identifies the

particular drum, disc or other device on which the segment resides. For each page up to a maximum of sixty four, there is an address in the page address list which locates that page on the device. For unused pages, there is an unassigned address. The order of the addresses in the address list corresponds to the order of the pages in a segment.

4.6. The Segment Length

The segment length is given in pages. It is one plus the number of the highest page accessed counting from zero. There is a maximum of sixty-four pages for any segment.

4.7. The Active Switch

The active switch, when ON, indicates that some of the segment attributes are to be found in the page table, PT, and its associated active segment table entry, ASTE. The segment may be in core or secondary storage or partly in both when this switch is ON. It is always in secondary storage when the active switch is OFF.

4.8. The PT-ASTE Pointer

The PT-ASTE pointer is a pointer to the location of those attributes of the segment in the page and active segment tables. It is only valid when the active switch is ON. One of the attributes found in the active segment table at that time is the segment map.

4.9. The Directory Switch

Finally, the directory switch indicates whether the segment is a directory or not. This information is used when deleting segments and will also be needed for segment handling which is explained in later chapters.

Other attributes are found in the branch. Some of these will be introduced and explained where needed in later chapters.

5. SEGMENT CREATION

A segment is created by establishing for it a branch in a directory. A module is called with the arguments segment name (pathname of the segment to be created), access to the segment and directory switch. The parent directory of the segment to be created is accessed using the pathname argument with the final entryname truncated from it.

To create a segment, information must be written into its parent directory. This requires that a hardware address be assigned to that directory and, therefore, that the directory be assigned a segment number. The assignment of a segment number is called "making a segment known to a process" and is described in Chapter 4. Subsequent addressing is by segment number and the segment control and page control modules handle the problems of activating the segment and bringing its pages into core.

When the parent directory has been accessed in this manner, a free entry is found in it. The entryname to be given this new branch is checked to make sure that it is not already in use in the parent directory. Space is allotted to the name and access lists and they are moved into their allotted places. Pointers to the name and access lists are placed in the entry. A uid is created for the new branch by a special subroutine and the uid is placed in the branch. The segment map is initialized by assigning to the segment a device identifier and setting all of the page addresses to unassigned. Page control will assign addresses to the pages as they are referenced. The branch count of the parent directory is then incremented by one.

At this point all that need be done to create a non-directory segment has been completed. A test is made to see if a directory is being created. If so, the directory being created is assigned a segment number and accessed. The necessary pointers are placed in its header and its branch count is set to zero. Creation of a directory segment is now complete.

An error return with an appropriate comment would have been executed if a free entry had not been found or the entryname had already been in use or the name and access list area had been full. No segment or entry would have been created in any of those cases.

6. SEGMENT DELETION

As in creation, segment deletion is for the most part deletion of a branch entry. However, before a segment may be deleted, several checks must be completed.

The parent directory of the segment to be deleted is accessed as in segment creation. The directory switch is tested to see if the segment to be deleted is a directory. If so, it is accessed and its branch count is checked. A directory cannot be deleted if any branches remain in it since that would break the path to all of its inferior segments. If there are no branches in the directory to be deleted then the execution continues as if it were a non-directory segment.

When the directory switch is tested and a non-directory segment is to be deleted, then the active switch in the branch to be deleted is checked. The segment cannot be deleted while this switch is on for that would leave traces of the segment in core and possibly even in other processes. Therefore, segment control is called to deactivate it.

Deactivation removes all traces of the segment from core, in particular from any descriptor segments and from the page table and active segment table entry assigned to this segment. (This forces any subsequent reference of the segment by any process to execute a segment fault thus referencing the branch and seeing that it has been deleted.) Page Control is called to free all secondary storage used by the segment to be deleted. The branch itself is deleted by zeroing its uid. Finally, the branch count of the parent directory is decremented by one and the segment deletion is complete.

Chapter 4

MAKING A SEGMENT KNOWN TO A PROCESS

1. INTRODUCTION

Segments in Multics are identified system-wide, to all users and processes, by their pathnames. However, the hardware references segments by numbers called segment numbers. Therefore, during execution a segment number must be associated with each segment. The segment number associated with any particular segment may differ from one process to another. A segment is said to be "known to a process" (or simply "known") while at least one of its pathnames is associated with a segment number and this association is recorded in a per process segment called the Known Segment Table.

A segment is "unknown to a process" or "unknown" until it has been made known and can again be made unknown to a process by terminating it, that is, by erasing the record of the pathname-segment number association from the Known Segment Table, KST. This breaks the association between pathname and segment number since the record in the KST is the only record of that association.

This chapter describes the way in which segments are made known to a process and the way in which they are made unknown.

2. DATA BASES

There are two major data bases involved in making a segment known. These are the directory hierarchy and the Known Segment Table, KST. Directory structure and contents are discussed in Chapter 3. The KST is discussed in this section.

The KST is a segment with an array of KST entries. Figure 1 is a diagram of the KST with one KST entry, KSTE, displayed in detail. A KSTE contains:

- a name list pointer
- a segment number
- a unique identifier, uid
- a branch pointer
- a directory switch
- an inferior count
- and other information.

The header shown in Figure 1 contains housekeeping information pertinent to the KST such as a pointer to the next free KSTE. The various elements in the KSTE are explained below.

2.1. The Name List Pointer

The name list pointer is a pointer to a threaded list of pathnames. The pathnames are all the different synonyms which have so far been used by this process to refer to the same segment, the one associated with this KSTE.

2.2. The Segment Number

Note that there is no segment number in a KSTE. The index of a KSTE in the KSTE array is the segment number associated with that KSTE. It is, therefore, the segment number of the segment whose name is pointed to by the name list pointer. In Figure 1, "s" is the segment number of the segment whose names are in the pathname list of the expanded KSTE.

2.3. The Unique Identifier, uid

The uid is copied from the branch of the segment when the segment is made known. The uid uniquely identifies a branch and is described in Chapter 3.

2.4. The Branch Pointer

When a segment is made known a pointer to its branch is placed in its KSTE. The presence of the branch pointer in the KSTE is necessary since the name of the segment could be changed while the segment is known. It also allows the segment-fault handler easy access to the attributes of a segment without having to repeat the search for the branch.




Figure 1. Structure of the Known Segment Table (KST)

2.5. The Directory Switch

The directory switch simply tells whether a segment is a directory or not. It is present because of a special rule regarding the handling of directories when making them unknown.

2.6. The Inferior Count

The inferior count is used for directory segments. It is a count of the number of immediately inferior or daughter segments known in the process to which this KST belongs.

2.7. Other Information

There is other information in a KSTE which is put there for use in access control, for example, but which is not pertinent to this discussion. No further mention will be made of it.

3. MAKING A SEGMENT KNOWN

The procedure executed in order to make a segment known is illustrated by the flow chart in Figure 2. Entry is through a call to MAKE-KNOWN. The pathname of the segment to be made known is passed to MAKE-KNOWN and the segment number and a code are returned.

The KST is searched for the pathname and if the pathname is found, MAKE-KNOWN returns a segment number. The third argument, code, is set to inform the caller whether or not the segment just made known is a directory directory. If the pathname passed to MAKE-KNOWN was not found in the KST, then it is tested to determine if it is the root directory pathname. This step is very important as it assures an end to the recursive loop which is described below.

If the pathname is not that of the root directory, then the pathname is parsed and is broken into two parts, a new pathname, and an entry name. The new pathname is the pathname of the directory in which the new entry name can be found. For example, if > A > B > C > D were the original pathname, then the parse would yield the new pathname > A > B > C and the entry name D. MAKE-KNOWN is then called recursively to make the new directory pathname known. This recursive loop is executed until a directory pathname is found in the KST or until the root directory is encountered and made known.

A call to make the root directory known always terminates the recursive loop. Since all pathnames begin with the root pathname it is assured that the recursive loop will always be terminated.

When MAKE-KNOWN returns from a recursive call the code argument is checked to make sure that a directory was found. The directory is searched for the entryname separated from a previous pathname by the parse.

An example is useful at this point. Assume that the segment with pathname > A > B > C is to be made known and that MAKE-KNOWN has been called twice (once recursively) for this segment. On the initial call the pathname > A > B > C was passed to MAKE-KNOWN and on the second call (first recursive call) > A > B was passed to it. Assume that directory segment > A > B was found to be known, then upon return from the last call (the recursive call) the code argument is tested to see if > A > B is a directory. If so, then entry C is looked up in > A > B. We now return to the description of MAKE-KNOWN.

If the entryname is found in the directory whose segment number was just returned, then the entry is tested to determine if it is a branch or a link. If it is a branch, then its uid is looked up in the KST to make sure that the segment is not already known by some other name. If the segment is already known by some other name then the new name is added to its pathname list, the segment number is returned and the code argument is set from the directory switch in its KSTE. If the segment is not known, then a free KSTE is found. The pathname is allocated space and placed in that space. The name list pointer is set to point at the pathname, the uid and the directory switch are copied from the branch into the KSTE and the branch pointer is set in the KSTE. The code argument is set from the directory switch in the branch and the KSTE inferior count of the parent directory is incremented by one. Finally, MAKE-KNOWN returns to its caller (itself or some external caller).

When there is a call to make the root known a special procedure in MAKE-KNOWN is used. Special data is entered into the KSTE for the root without searching the directory hierarchy or attempting to find a branch for it. Its segment number and code are returned in the normal manner.

When a link entry is found, then MAKE-KNOWN is called with the pathname found in the link and the flow continues in the normal manner.

If an error is encountered such as not finding a directory on return from a recursive call to MAKE-KNOWN, then an error code is set into the code argument and a null segment number is returned.




Figure 2. MAKE-KNOWN

4. MAKING UNKNOWN

Normally, a user will experience no difficulty because of the limited size of the KST. However, since there are many more segments in the directory hierarchy than there are KSTE's, a user might wish to free a KSTE (making a segment unknown) so that it can be reused. In any case, certain precautions must be taken before a segment can be made unknown.

First, the directory switch in the KSTE is checked. If the segment to be made unknown is a directory then its inferior count is checked. A directory cannot be made unknown if any of its inferior segments are known. This convention is stated in Chapter 2. It arises from our desire to be able to take segment faults on segments used in the Multics supervisor and not to distinguish between supervisor and user segments.

Second, segment control must be notified that a segment is being made unknown since the KST is prepared for and used by segment control. This is done by a special call to segment control. Upon receiving this call, segment control will disconnect the SDW associated with this segment in this process (see Chapter 5). When segment control has been notified the segment can be made unknown by freeing the area where its name(s) was stored and threading the KSTE onto a free KSTE list. The parent directory's KSTE inferior count is then decremented by one and the operation of making segment unknown is complete.

Finally, a few words must be said about the danger of making a segment unknown and reusing its KSTE. Addresses are prepared using segment numbers. All of these addresses cannot be found when a KSTE is to be freed. If such an address is used after the KSTE has been reused, it will cause information in the corresponding KSTE to be used without further checking. Incorrect segment referencing would result. Therefore, a segment should not be made unknown and its KSTE reused unless it is assured that no address which references that segment will be used again during the existence of the process.

5. INITIAL REQUIREMENTS

The question may be asked, "how much apparatus is required to make a segment known for the first time in a process?"

The data bases used are the directory hierarchy and the KST. The procedures used are MAKE-KNOWN and the procedures called by it. It has been shown that all directories including the root directory may be made known. However, the root directory requires special code. The KST must already have a segment number in order to make a segment known so the KST cannot be made known in this way. Some of the modules called by MAKE-KNOWN could possibly themselves be made known, however, special codes would be necessary to do this. Therefore, MAKE-KNOWN and all of the procedures used by it as well as the KST for a process are assigned segment numbers before the process begins executing as a part of process initialization.

6. OTHER MULTICS CONSIDERATIONS

The fact that some segments must have segment numbers before MAKE-KNOWN can be executed gives a clue to the Multics implementation of segment numbers. There is a group of segments in the Multics supervisor which must have segment numbers assigned in a process before the process can begin execution. These are called hardcore segments. They have no KSTE's. The actual segment numbers assigned to segments when they are made known are the KSTE index plus the highest segment number assigned to a hardcore segment.

Chapter 5

SEGMENT FAULT HANDLING

1. INTRODUCTION

In the Multics Operating System, each process address space is divided into 64K-word items called segments. A segment enters a process address space by being "made known to the Process" (see Chapter 4). In the course of being made known, a segment has a per-process segment number assigned to it. A correspondence is established between this segment number and the segment's pathname and attributes in a per-process table called the "Known Segment Table". Once in a process address space, a segment may be referred to by segment number.

Whenever a process references memory, the 645 hardware references "per process" and "per system" registers. The "per process" information in the hardware accessing path is recorded in a Descriptor Segment which contains one word, called a Segment Descriptor Word (SDW), per segment number. The function of the Nth SDW is to point to the Page Table (see Chapter 6) of the segment which is known to the process as segment #N and to specify the process access rights with respect to that segment.

The Page Table of a segment (and various other data required by the paging mechanism of Multics) must be stored in core. Since there are more segments in Multics than places in core for Page Tables, not all segments can have Page Tables at the same time. When a segment has a Page Table, the segment is called ACTIVE. At other times it is called INACTIVE.

An SDW can, of course, contain the address of a segment's Page Table only if the segment is active. If an SDW contains the address of the segment's Page Table and specifies the process access and the segment's length, then the SDW is called CONNECTED; otherwise, it is called FAULTED or DISCONNECTED. A faulted SDW in fact contains a bit pattern which, when encountered by the 645 addressing hardware, causes the process to "take a Segment Fault" thereby invoking the Segment Fault Handler. See Figure 1.

In view of the above, we may say that:

The function of the Segment Fault Handler is to provide the process with the illusion that all segments known to it are active and all SDW's corresponding to known segments are connected; in short, to render all known segments directly accessible by segment number.




Figure 1. Connected and Disconnected SDW's

2. PREVIEW OF THE SEGMENT FAULT HANDLER (SFH)

2.1. Procedure

A segment fault occurs when a process attempts to access a "target" segment via a faulted SDW. The Segment Fault Handler (SFH), called to repair the faulted SDW, must obtain the address of the Page Table for the "target" segment as well as the process access rights to the segment and store the information in the SDW.

To do this, the SFH must:

- a. Check the validity of the segment number given to the SFH. It may have been incorrectly generated.
- b. Use the segment number of the "target" segment (which the SFH is given) to find a pointer to the segment's branch. This "branch pointer" was put into the segment's entry in the process Known Segment Table (KST) when the segment was made known. (See Chapter 4).
- c. Check the branch to see if it in fact corresponds to the "target" segment. This check is necessary due to the dynamic nature of the File System in which segments can, at any time, be created and destroyed or moved from one directory to another. In checking the branch, a unique identifier (UID) is used which was stored in the segment's KST entry when the segment was made known.

- d. Look in the branch, which contains the segment's attributes, to find the process access to the segment and to locate the segment's Page Table.
- e. Repair the SDW and return.

2.2. Data: Active Segment Table (AST)

We have stated that Page Tables and other data describing active segments must be stored in core, a requirement imposed by the Page Fault Handler (see Chapter 6). Use is, therefore, made of a system-wide table, the Active Segment Table (AST), which resides permanently in core. The AST is a linear array containing one entry per active segment. An active segment's AST Entry (ASTE) contains the segment's Page Table and other paging data and, as we shall see, other data needed by the Segment Fault Handler.

Steps (a) through (e) above and the description of the AST lead to the picture of the data structures used by the Segment Fault Handler and the relations between them shown in Figure 2.




Figure 2. Principal Data Structures Used by the Segment Fault Handler

2.3. Program of Exposition

In order to make segment fault handling more easily comprehensible, it will be useful to discuss the procedures for handling segment faults in three sections corresponding to three successively more likely assumptions about the state of activity of the segments of the system.

These assumptions are:

1. All segments are active simultaneously.
2. All segments can be active simultaneously.
3. All segments cannot be active simultaneously.

3. SEGMENT FAULT HANDLING WHEN ALL SEGMENTS ARE ACTIVE

Let us examine the very unlikely case in which all segments are active. The Segment Fault Handler (SFH) begins by finding the "target" segment's KST entry and branch as explained in Section 2. Two kinds of errors may happen. First, if no segment is "known" by the faulting segment number, then no KST entry exists. This error may happen, for example, if a memory reference is made via a randomly generated segment number. This error is detected through structural details of the KST which do not interest us here.

Second, as mentioned in Section 2, the branch pointer in the KST entry will be incorrect if the "target" segment has been destroyed since this process made it known. To check the correctness of the branch pointer, the SFH compares the unique identifier (UID) in the branch with the one in the KST entry. If they are the same, then the branch pointer is alright; otherwise, the "target" segment has been destroyed and the segment fault cannot be satisfied.

If the branch pointer is valid, the SFH gets the process access rights to "target" segment and the index of the "target" segment's ASTE from the branch. (The branch of an active segment must, of course, contain a pointer to the segment's ASTE. The ASTE index is that pointer.) The address of the "target" segment's Page Table can be calculated easily from the ASTE index. The length of the segment (in pages) can be obtained from the ASTE.

With the table address, the segment length, and the access information, the SFH has enough information to correct the faulted SDW. The SDW is corrected and the SFH returns.

4. SEGMENT ACTIVATION

We wish now to see how segment fault handling must differ if it is possible that the "target" segment may not be active. We assume that there is room in the AST to make an ASTE for any segment. Two new data structures must be introduced.

First, a new piece of information must be added to the branch: an active switch. This switch indicates whether or not the segment is active; i.e., whether the ASTE index in the branch may be used. Second, a new data structure must be added to the AST - a list of available entries.

This structure is called the AST free list. In this section, we assume that the AST free list is never exhausted.

Now let us look at segment fault handling. Once the branch has been accessed, the SFH must inspect the "active switch". If the segment is active, the processing is as described above. If the segment is not active, it is necessary to activate it, that is, to:

1. Find an entry in the AST free list. Remove it from the AST free list.
2. Set the ASTE index in the branch to point to this entry.
3. Copy the paging data from the branch to the ASTE. Initialize the Page Table in the ASTE.
4. Set the branch's "active switch" on.

After activating the segment, the SFH proceeds as before to repair the faulted SDW.

5. SEGMENT DEACTIVATION

Let us now consider the real case in which the size of the AST limits the number of segments which can be active simultaneously. The discussion of this case is sufficiently long that we will divide it into two parts. In this section we will introduce the pieces of the design. In the following section we will put the pieces together.

The assumption that all segments cannot be active simultaneously implies that it may sometimes be necessary to activate a segment at a time when the AST is "full"; i.e., there are no ASTE's in the AST free list. When this happens it becomes necessary to:

- a. Choose an active segment to be deactivated.
- b. Deactivate this segment.
- c. Return the ASTE of the deactivated segment to the AST free list.

Let us defer discussion of how an active segment is chosen for deactivation until the mechanism of deactivation has itself been discussed.

Deactivation of a segment may be characterized as doing all things necessary to disassociate the segment's ASTE from the segment, thus permitting the ASTE to be used to activate another segment. It is important to note at this point that the segment being deactivated is not necessarily "known" to the process performing the deactivation. Many of the details of the design arise from this fact. Section 9 discusses the matter further.

Deactivation is done in the following three steps:

- i page removal - forcing the segment's pages out of core.
- ii disconnection - seeing that all SDW's associated with this segment are faulted.
- iii restoring the branch - moving back to the branch the (presumably altered) values of those attributes of the segment which were moved to the ASTE when the segment was activated.

5.1. Page Removal

A Page Table for a segment can only exist when the segment is active; when the segment is deactivated, the Page Table is destroyed. Thus, one function of deactivation is to remove the segment's pages from core and to inhibit the Page Fault Handler from bringing any of its pages into core during deactivation. To accomplish this, the SFH calls a special entry of the paging module which removes the segment's remaining pages from core and causes page faults (in other processes) on pages of this segment to be transformed into segment faults on this segment.

5.2. Disconnection

Disconnection obviously amounts to faulting all SDW's connected to the segment at the time of deactivation. There may, of course, be many other SDW's associated with the segment, but as they are not connected they must (already)

be faulted. In order to make disconnection possible, a data structure must be associated with each segment which lists all the SDW's connected to the segment. Since this connection list can only be non-empty when the segment is active, it can be stored in the segment's ASTE and need never appear in the segment's branch. When a segment is activated, the connection list is empty. Whenever an SDW is connected to a segment, a pointer to the SDW is added to the segment's connection list. When the segment is deactivated, each of the SDW's on the associated connection list is faulted.

5.3. Restoring the Branch

After page removal and disconnection have been done, all that remains of deactivation is to restore the segment's branch. It is necessary to move back to the segment's branch those attributes which were moved from the branch to the ASTE when the segment was activated (which may have changed during the period of activation) and to reset the "active switch". In order to manipulate the segment's branch in this way, the SFH makes use of another data element in the ASTE of each segment, a pointer to the segment's branch. This branch pointer must be stored in the ASTE when the segment is activated so that it can be used again when the segment is deactivated.

5.4. A Note on Certain Necessarily Active Segments

The reader can easily convince himself that the main path through the SFH (excluding the deactivation path) can only be viable if certain segments are perpetually active: for example, all segments required by the Page Fault Handler, the SFH procedure itself, the per-process KST segment, etc. To permit such segments to be held active, another item is added to the ASTE, the entry hold switch, which may be set to prevent the deactivation of the segment. We will consider these matters further in Section 8 on "Recursion and Initialization." Here we wish to consider a different problem.

During deactivation of a segment the SFH must access all the descriptor segments which contain SDW's connected to the segment being deactivated; it must fault these SDW's. The SFH must also access the segment's branch. In order to assure that these segments may be accessed easily, the following rules are established:

Rule 1 - A Descriptor Segment containing a connected SDW must be active.

Rule 2 - A Directory Segment with an active daughter segment must be active.

Rule 1 may be enforced by use of the entry hold switch. The enforcing of Rule 2 is more complicated. A new data element is associated with each active segment, the inferior count. This is a count of the active daughter segments of the given segment. (For non-directory segments the inferior count is necessarily zero.) When a segment is activated, its own inferior count is set to zero and one is added to its parent's inferior count. When a segment is deactivated, one is subtracted from its parent's inferior count. Needless to say, no segment is chosen for deactivation whose inferior count is non-zero. To enable the deactivation machinery to access the parent directory's inferior count, each ASTE also contains another data element: the parent's ASTE index.

5.5. The Deactivation Algorithm

We have now presented all of the data structures needed to permit deactivation of a segment. Let us now discuss the "deactivation algorithm" which must choose the segment to be deactivated. To facilitate deactivation, another list is introduced, the AST used list, which contains all of the ASTE's corresponding to active segments.

In order to keep deactivation economical, the deactivation algorithm chooses segments with as few pages in core as possible. In practice, it is very often able to choose a segment with no pages in core. The algorithm essentially passes through the AST used list as many times as necessary; on the Nth pass it looks for a segment with fewer than N pages in core whose inferior count and entry hold switch are both zero. When such a segment is found, the search stops and the segment is deactivated.

6. FLOW OF THE COMPLETE SEGMENT FAULT HANDLER

Correct Segment Fault

1. Given segment number, find KST entry.
Check validity of KST entry.
In KST entry, obtain pointer to the branch.
Check validity of the branch pointer by comparing the KST entry's and the branch's version of the unique identifier.
2. If the "active switch" is set, go to Step 8.

Activate the Segment

3. Inspect the free list. If there is a free ASTE, then go to Step 7.

Obtain a Free Entry

Choose a Segment to be Deactivated

4. Pass through the AST used list as many times as necessary; on the Nth pass look for a segment having fewer than N pages in core with inferior count and entry hold switch equal to zero. When a segment is found in this way, go on to Step 5.

Deactivate the Segment

5. Remove the pages of the segment from core by a call to the paging module.
Disconnect all of the SDW's listed in the segment's connection list.
Use the ASTE's branch pointer to move the segment's attributes back to the segment's branch.
Reset the "active switch" in the branch.
Use the parent's ASTE index in the ASTE to locate the parent's inferior count, from which subtract one.
6. Remove this entry from the AST used list.
Put it in the AST free list.

Activate the Segment

7. Move the ASTE from the AST free list to the AST used list.
Move certain attributes from the branch to the ASTE.
Set the branch's ASTE index and set the "active switch".
Initialize the Page Table in the ASTE.
Set the segment's inferior count to zero.
Set the segment's entry hold switch to zero.
Find the parent's ASTE; store its index in the segment's ASTE.
Add one to the parent's inferior count.

Connect the SDW

8. Get the page table address and segment length from the ASTE.
Get the access field for the SDW by inspecting the Access Control List in the branch.
Store the page table address, segment length, and access field in the SDW. Reset the Segment Fault Flag.
9. Return.

7. DATA USED IN SEGMENT FAULT HANDLING

7.1. AST (Active Segment Table)

The AST consists of a Header and a linearly indexed array of AST Entries of which there is one per active segment. The AST Header contains:

- a. The head of the AST free list.
- b. The head of the AST used list.
- c. The branch of the root directory.

Each AST Entry contains, in addition to forward and backward pointers used in threading the entries into the various lists,

- a. Paging data
 1. The PAGE TABLE.
 2. The SEGMENT MAP.
 3. Various other paging data.

- b. Data used in choosing a segment for deactivation.
 - 1. The INFERIOR COUNT. If the segment is a directory segment, the inferior count is the number of its active daughter segments. Otherwise it is zero.
 - 2. The NUMBER OF PAGES IN CORE (a counter maintained by the paging mechanism).
 - 3. The ENTRY HOLD SWITCH. When set, this switch makes the segment ineligible for deactivation.
- c. Data required to deactivate the segment.
 - 1. BRANCH-POINTER
 - i ASTE index of the segment's parent directory.
 - ii Offset of the segment's branch within the parent directory.
 - 2. CONNECTION LIST, the list of SDW's which point to the Page Table of this segment and which must be faulted before the segment can be deactivated. Each element of the list consists of:
 - i ASTE index of the descriptor segment containing the SDW.
 - ii Offset within that descriptor segment of the SDW (i.e., the segment number within that process of this segment).

7.2. SDW (Segment Descriptor Word)

A Descriptor Segment is a linear array whose entries are words called SDW's (Segment Descriptor Words). The index of an SDW in a process Descriptor Segment in the segment number, in that process, of the segment associated with the SDW. An SDW consists of:

- a. The ACCESS FIELD
 - 1. If zero, this field causes a segment fault to occur upon an attempted access via this SDW. In this case, items (b) and (c) below are meaningless.

2. If non-zero, this field indicates the accessing permission that this process has toward the segment associated with this SDW (e.g., "read", "write", or "execute" permission).
- b. The PAGE TABLE ADDRESS of the associated segment.
- c. The SEGMENT LENGTH (in pages) of the associated segment.

- 7.3. KST (Known Segment Table) - See Chapter 4.
7.4. Branch - See Chapter 3.
7.5. Page Table, Segment Map - See Chapter 6.

8. RECURSION AND INITIALIZATION

The phrase "recursive segment fault" refers to a segment fault taken by the SFH. In Section 5, we asked the reader to believe that segment faults must not be permitted to occur on certain segments; e.g., the AST. In this section, we shall show why this is so and how such segment faults can be avoided. We shall then consider the recursive segment faults which can be permitted and will conclude with an example.

We may regard each segment fault "taken by the user" as the first in a sequence of segment faults, the rest of which are taken "recursively" by the SFH in handling the first one. It is a design requirement that all such sequences be finite. This implies that the last segment fault in the sequence must be handled entirely by segments which are active and connected in the process handling the faults. Thus, certain segments must be active and connected in a process before the first segment fault is taken.

Which segments must be active? Certainly all segments must be active and connected which are necessarily called in the course of handling every segment fault. In the present design these comprise the AST, KST, SFH procedure and of course all segments required by the paging mechanism (since the SFH references segments which do not reside permanently in core and may take page faults).

Multics initialization must make all of these segments, except the per-process KST, active before the first process is created. Process creation and initialization must create and activate the KST and connect all of these segments before the first segment fault is taken. These segments may be kept active in various ways. Per system segments like the AST and SFH procedure may be kept active by leaving their ASTE's off the AST used list. The per user KST may be kept active by setting the entry hold switch in its ASTE while the process itself is active.

The only segments referenced by the SFH beside those discussed above are directories, the segments which contain branches. The only way to avoid recursive segment faults altogether is to require that every directory, any daughter segment of which is known to any process, be active and appropriately connected. This idea must be rejected as impractical. Hence, recursive segment faults must be reckoned with.

We shall give an example of a sequence of recursive segment faults at the end of this section. We shall show there that all such sequences can be handled if a segment fault on the root directory can be handled. Let us prepare for the example by considering the root directory more closely.

Every segment in Multics except the root directory has a branch in a directory; the branch contains the attributes of the segment. Since no directory is superior to the root, it cannot have a branch in a directory. Nevertheless, if the root is to be accessed, its attributes will have to be recorded somewhere. Since we normally think of a branch as the locus of a segment's attributes, we may provide for the root's accessibility by providing it with a branch. This branch must itself be accessible; since no process can take a segment fault on the AST, it is sufficient that:

The root directory has a branch which resides in the AST segment.

Whenever the root directory is "made known" in a process, a pointer to this branch is placed in its KST entry. Segment faults on the root can obviously be handled in the normal way whether or not the root is active at the time of the segment fault.

NOTE: In practice, a more efficient (if less straightforward) SFH may be obtained by handling segment faults on the root with special code. If this is done, the "branch" of the root disappears from the AST and, by way of trade-off, the SFH procedure segment gets longer.

We have noted that the only segment fault that the SFH can take while handling a segment fault for a segment is a fault on that segment's parent directory. Thus, every sequence of recursive segment faults corresponds to a path up the directory tree structure toward the root directory. Let us now discuss the canonical example of a sequence of recursive segment faults.

8.1. Example of Recursion

Let us assume that a segment fault is taken for a segment with pathname

```
root > d1 > d2 >...> dN > seg
```

Let us further assume that the process SDW's for all the directories "root", "d1", etc., are faulted. Early in handling the fault on "seg", the SFH references the branch of "seg" which causes a segment fault on "dN". The corresponding reference to the branch of "dN" causes a fault on "dN-1" and so it goes until there is a fault on "root". Since the branch of "root" lies in the AST which is active and connected, this segment fault can be handled without another segment fault. With the "root" connected, the SFH can go on with handling the fault on "d1". When "d1" has been connected, "d2" can be handled, and so on. Thus, the only recursive segment fault permitted in the present design is that on a segment's parent; and we have shown that the recursion terminates due to the special treatment of the root directory.

9. SPECIAL ADDRESSING IN DEACTIVATION

9.1. Basic Problem

The deactivation of a segment requires the accessing of a directory (the segment's parent) and of one or more Descriptor Segments (those containing SDW's connected to the segment).

It is very unlikely that these segments are all "known" to the process performing the deactivation. How, then, can the branch pointer and the SDW pointers in the connection list in the ASTE be implemented?

9.2. The Multics Solution

In the present design, the parents of active segments and all descriptor segments connected to active segments are required to be active. This requirement enables the use of the ASTE index as segment specifier for all of the special addresses needed in deactivation. Thus, the branch pointer and SDW pointers are all of the form:

(segment's ASTE index, offset)

Let us look in detail at the trick by which one of these special addresses is used, say the "branch-pointer". A special segment number is reserved in each process for use by the SFH. To access the branch during deactivation, the SFH uses the branch's ASTE index to compute the branch's page table address. This page table address and read and write access permissions are then stored in the SDW which corresponds to the reserved segment number in the deactivating process Descriptor Segment. A pointer using the segment number and the offset given in the "branch-pointer" is then constructed and used to access the branch through the just manufactured SDW.

In effect, the reserved segment number and the corresponding SDW constitute a "window" through which a process can reference a segment not formally known to it.

10. SPECIAL ENTRIES OF THE SEGMENT FAULT HANDLER

There are four functions related to segment fault handling which are accomplished by special entries of the SFH.

10.1. Obtaining a Free ASTE

Some procedures need to be able to obtain a free ASTE in order to activate a segment "by hand". For example, the process creation procedure must make up an active descriptor segment for a new process. A special entry in the SFH is provided which obtains a free ASTE (whether from the AST free list or by deactivation), detaches it from the free list, does not thread it into the AST used list, and returns its index to the caller.

10.2. Setting Faults for All Users of a Segment

Certain procedures wish to fault all SDW's connected to a segment. For example, the Access Control Module will do this if the access control data in the segment's branch has been modified in certain ways. These faults are set by using the part of the deactivation code which disconnects SDW's.

10.3. Deactivating a Segment

Certain procedures must be able to deactivate a segment, for example, the procedure which destroys a segment. Deactivation is accomplished by using the deactivation path in the SFH.

10.4. Disconnecting an SDW from a Segment

Some procedures, for example, the procedure which makes a segment "unknown" to a process, need to be able to disconnect an SDW from a segment. Code in a special entry of the SFH is provided which faults the SDW and removes it from the list of SDW's connected to the segment.

11. CONCLUDING REMARKS

The functions of the Segment Fault Handler (SFH) are:

- to make (the Page Tables of) all segments which are "known" to a process accessible to the process by segment number.
- to multiplex the system's relatively few AST Entries (or, equivalently, Page Tables) among all of the segments "known" to all of the processes executing in the system.

To accomplish these functions, the SFH must establish (activate) and dis-establish (deactivate) the pageability of segments and must establish (connect) and dis-establish (disconnect) the use of such pageability by processes. It is instructive to separate these four functions into two groups.

11.1. Connection and Activation - Process Oriented Functions

The primary tasks in handling a segment fault, connection and activation, are done on a demand basis according to the needs of a process. These functions are performed using only data segments "known" to the process - the per-process KST and Descriptor Segment, the directories superior to the "target" segment, and the AST segment. In this part of handling a segment fault, two types of error may occur. The segment number may (a) not correspond to a segment "known" to the process, or (b) may correspond to a segment which has been destroyed. In either case, the process is trying to access a segment which does not exist and should note the error. For all of these reasons, we say that the connection and activation aspects of segment fault handling are "process oriented functions"; i.e., that the SFH in performing them is acting "for the process".

11.2. Disconnection and Deactivation - System Oriented Functions

The secondary tasks of segment fault handling, disconnection and deactivation, are undertaken by the SFH during activation only when the state of the system demands it, the relevant state of the system being the emptiness of the AST free list. The choosing of the segment to be deactivated is independent of the process in which the SFH is executing. Deactivation and disconnection of a segment require the SFH to access segments not (necessarily) "known" to the process executing the SFH: the parent directory of the segment being deactivated and the Descriptor Segments containing SDW's connected to the segment. The method by which these segments are accessed by the SFH (see Section 9) is clearly independent of the address space of the process in which the SFH is executing.

(The number of directory segments and Descriptor Segments in the system may well exceed 2^{18} . All of these segments are potentially referenceable by any process which is executing in the SFH deactivation path.)

No errors are detected (or caused) during deactivation and disconnection and no per-process errors affect their operation. For these reasons, we say that the deactivation and disconnection aspects of segment fault handling are "system oriented functions", i.e., that the SFH in performing them is acting "for the system".

Chapter 6

PAGE FAULT HANDLING

1. INTRODUCTION

In the Multics Operating System, segments are composed of 1024-word contiguous blocks of data called pages. At a given time, any number of pages of a segment may be located in core memory, but since that memory is limited in size, the hardware blocks of 1024-word core registers must be multiplexed among the many pages of data and procedures which may be referenced. It is the purpose of this chapter to detail the structure of the mechanism which accomplishes (block or) page multiplexing in Multics.

The page multiplexing strategy is similar in broad outline to the page table multiplexing performed by the segment fault handling module (see Chapter 5). However, there are differences in detail which arise in great part due to the 645 hardware used in page fault handling. Page fault handling is closely bound to the various registers and logic functions which the 645 processor can perform; indeed the major purpose of the paging modules is to create the proper environment for hardware access to pages. This access is made through several registers, but the one which uniquely concerns page multiplexing is the page table word (or PTW). This 36-bit register, located in the page table for a segment, contains all of the information used by the 645 processor to deal with a page. The proper maintenance of a PTW is page multiplexing's most basic job.

As a vehicle for carrying the description of page multiplexing, there is a convenient set of machine configurations whose physical capabilities obviate the need for various parts of the page multiplexing function. Considered in order of increasing likelihood they are:

- 1) Infinite core storage and no secondary storage,
- 2) Infinite core storage, with secondary storage, and
- 3) Finite core storage with secondary storage.

Although no multiplexing is required in the first two cases, we shall describe the Multics page fault handling strategy as it would be performed on each of these three configurations, since in this way the real strategy can be described incrementally.

Before discussing the handling of page faults in detail, we shall describe briefly the environment which allows a page fault to occur. It is important to remember that a page fault cannot occur if there is no page table for the segment in question - for a page fault is only generated by hardware reference to a page table. The table is provided for a segment by the segment fault handler when it activates the segment. Also at this time, the Active Segment Table Entry for the segment is initialized by the segment fault handler with all the information needed by the page fault handler. Finally, the page table is set with page faults for each page of the segment, so that the page fault handler will be invoked upon first reference to each page. These actions prepare the environment for page fault handling.

2. PAGE FAULT HANDLING ON A MACHINE WITH INFINITE CORE AND NO SECONDARY STORAGE

When a process attempts to reference a page whose PTW has a fault set, the paging modules are invoked to remove the fault and insert the proper core location of the page into the PTW. The first action, upon receiving notification of the fault, is to locate the page being sought. From the PTW address, the address of the Active Segment Table Entry (ASTE) for the segment can be found. The ASTE contains the segment map (the list of page locations) which yields the actual address of the page (as described in Chapter 7). The segment map is actually split between the ASTE and the page table; however, we will ignore this complication for the moment.

Since we are now assuming a configuration involving only core, the address must be that of a 1024-word block of core. Hence, the paging module need only insert the proper core address into the PTW and fill in the page fault field of the PTW to prevent further page faults on this page. Thereafter, references to the page through the PTW would proceed by hardware without interruption.

The repairing of a page fault in a PTW need be done only once in the environment we are assuming, because all Segment Descriptor Words (SDW's) for a segment point to the same

page table and, therefore, to the same PTW for each page. If, however, a segment were deactivated (its page table destroyed), then before further references could be made to the segment, the page faults set by segment activation would have to be satisfied in the way described above. Page locations being constant assures us that the information stored in a PTW is valid as long as the page table is.

3. THE ADDITION OF SECONDARY STORAGE

If we introduce secondary storage into our machine configuration, we add two problems to the page fault handling mechanism. Since the location of a page can now be outside of core, there is a need to transport that page from its resident device; and also, we must find an appropriate core block into which to put it. The page fault path becomes slightly longer and necessitates referencing a new data base - the core map free list. For, in order to pick an appropriate block of core for the page, we must avoid those blocks currently in use. Since we are postulating infinite core, we need not be concerned with depleting the free page supply. The functions required in this configuration are:

- 1) Receive the fault, go from the PTW to the ASTE to get the segment map and determine the secondary storage location of the page.
- 2) Access the core map free list to obtain a 1024-word block and delete it from the list, and
- 3) call a Device Interface Module (DIM) to retrieve the page and deposit it in the newly acquired block of storage.

The DIM's functions in transporting pages are to queue requests for pages, to make the device perform as efficiently as possible in satisfying the requests, to monitor the device operation, and to notify the page fault handler when input has been completed.

The notification function is performed in such a way as to allow the process which took the page fault to wait for the page without wasting processor time or making unnecessary memory accesses. The page fault handler in the faulting process, after calling the DIM to initiate page transportation, calls the System Traffic Controller to wait for the page. At some later time, when the page has been imported, the Traffic Controller will be called to inform the waiting process that it may continue its computation.

4. RESTRICTION IN CORE SIZE

When we restrict our configuration to have a finite amount of core, we reach the true Multics case where multiplexing is necessary. In addition to the functions previously described, the page multiplexor must also be responsible for finding a free block of core when all blocks are being used. This condition requires a selection algorithm for removing pages from core to secondary storage; and this algorithm requires a new data base - the core map used block list.

The two lists of core blocks - free and used - are implemented by means of an entire core map. The core map consists of one core map entry (CME) for each 1024-word block of core. Each entry serviced by the page fault handler is threaded into one of the two queues - free or used - depending upon its current status, but the entry and its physical block remain associated throughout any change in status. The free list is singly threaded, since only its first element is ever used, but the used list is circularly threaded to allow continuous searching.

The actual information contained in a CME can be deduced from the part it plays in the removal algorithm. Clearly it must allow one to obtain the absolute core location of the beginning of the 1024-word block. But a pointer to the PTW for the page currently residing in the block is also necessary if the block is being used, since the removal of the page means that access to it should be inhibited by setting a page fault in the PTW which controls it.

The sequence of actions for handling a page fault in the limited core environment is:

- 1) As before, get the fault, find the device address from the segment map in the ASTE.
- 2) Access the core map free list to obtain a free block: if successful, place a pointer to the PTW in the CME selected, unthread the CME from the free list, flag it as being used for I/O and thread it into the used list. (Continue at Step 6).
- 3) If the free list is empty, perform the replenishment algorithm to find a block in which the referenced page can be put.

The replenishment algorithm is driven by a bit in the PTW called the "Page has been used" or PHU bit. This bit is set by the 645 hardware whenever the page is accessed. When a page must be removed, the used block list of the core map is accessed and the entries are examined in the order of their threading, starting where the last search stopped. Each entry is examined to determine whether the PHU bit has been turned on. Each page's PHU bit is turned on by the software when the page is brought to core, but after each examination during the removal algorithm the page fault handler turns it off. Therefore, the effective criterion for removal is whether the page has been used since it was last examined for removal. (For further information about the philosophy of this algorithm, see "A paging experiment with the Multics System", F. J. Corbato, Multics Repository Document M0104.)

- 4) Having found a candidate for removal, set a page fault in the PTW, determine the device address from the ASTE and call the DIM to transport the page to secondary storage.

In order to avoid unnecessary page transportation, the 645 hardware maintains a "page has been modified" bit in each PTW which is turned on only if the page has been written into. Unless this bit is on, the page need not be returned to secondary storage.

- 5) In contrast to the page important strategy, do not wait for this page to be moved, but continue to select candidates for removal and call the DIM to transport them until a block appears on the free list. This block will have been placed on the free list by the DIM when a page has been completely transported.
- 6) Using the device address developed in step 1), call the DIM to import the desired page into the newly acquired block of core.
- 7) Call the Traffic Controller to wait for the page to arrive.
- 8) Since the PTW's page fault switch and CME's I/O busy switch are reset by the process which called the Traffic Controller to awaken the waiting process, the page fault has been entirely repaired and the page fault handler may return.

Figure 1 is a gross flow chart of the page fault handling strategy as described in Section 4. Special points of interest are:

- 1) Any call to the DIM, whether for reading or writing, causes all transactions to be observed and the completed ones to be "posted". The posting process for a write operation consists of threading the CME out of the used list and into the free list. For a read, posting requires that the page fault switch be reset and that any processes which might be waiting for the page to be imported be informed of its arrival (through the Traffic Controller).
- 2) The call to the Traffic Controller to wait for a page to be read is only made for reasons of efficiency and plays no logical part in the page fault handling strategy.




Figure 1. The Page Fault Handler

5. RECAPITULATION OF PAGE FAULT DATA BASE USE

Having followed the basic paths through the page fault handler, we have seen all of the data bases used by page multiplexing. However, not all of the items in these data bases have been specified, only those portions significant to a general understanding of page fault handling. This section describes all of the data items referenced by the page multiplexor and indicates their uses. (See also Figures 2 and 3).

The most basic page multiplexing data base for a segment is the page table, which contains the page table words for each page. Page tables in Multics are allocated at system initialization time in a permanently core-resident system-wide segment, the System Segment Table (see Figure 3). Hence, references to a page table are made through the normal segment addressing mechanism - although no page faults are ever taken on such references. Each PTW has six types of data used by the page fault handler. These types can be divided further into hardware referenced data and solely software referenced. Both types are important to page multiplexing.

The set of hardware referenced data in a PTW consists of three items - the page fault switch, the core address field and the page has been modified/used bits. Whenever the core address is not meaningful, the page fault switch is set to inhibit processor attempts to access through that address. Since the address field is not used when the page fault switch is on, the page multiplexor makes use of the storage thus provided by placing part of the device address in that field. Logically, the entire address can be considered to be in the Active Segment Table Entry as mentioned earlier, but to save storage space, the "segment map entry" for a page is stored in the core address field of its PTW when the page is not in core. The PHM and PHU bits, set by the hardware and reset by the software when appropriate, have already been described.

The data in the PTW which is only referenced by software consists of two flags used by the I/O handlers and a bit which designates "wired down" status. The latter is interpreted by the page fault handler during the page removal algorithm to mean that this page may not be removed. The two former flags are: one to signify that there is I/O pending for this page (and whether read or write) and another which is set if an I/O error is encountered while transporting the page.




Figure 2. Data Base Inter-relationships

Notes for Figure 2.

- A. To bring in a page from secondary storage, start with the PTW which took the fault. Go to the ASTE to get the Device I.D. Use the core map free list to get a suitable CME. Fill in the pointer to the PTW and transfer the file map to the CME, then read in the page and remove the fault.
- B. To remove a page from core, start with the CME picked from the core map used list. Go to the PTW by the pointer and set a fault; thence to the ASTE for the Device I.D. to complete the secondary address. Move the file map back to the PTW. If the page has been modified, write it out.

NOTE: It may seem that Multics has two copies of a data page when the data is in core. Logically there is only one and we could easily free the storage used by the old page each time it was read in. There are three reasons of efficiency why we do not do so.


- 1) Assigning and reassigning secondary storage blocks takes processor time.
- 2) If the page to be removed from core has not been modified and we have retained the old copy, we do not need to write it out.
- 3) If the system were to crash, losing core but not secondary storage contents, we would still have a (possibly obsolete) copy of the data.

Core MAP Free List Pointer

1

Core MAP Used List Pointer

5



The n^{th} CME refers to n^{th} 1024-word block of core.

Figure 3. The System Segment Table

A fourth item of information, the address of the segment's ASTE, although not contained in the PTW, is implicit in its address; the ASTE's (which are fixed in length) are stored in an array parallel to the page tables, thereby yielding the ASTE address from the page table (or PTW) address. This information is necessary to provide the rest of the device address for a page when reading it in from secondary storage.

Another data base important to page fault handling is the Active Segment Table Entry (ASTE) for a segment. The ASTE contains more data than is used by the page multiplexor, but that portion which is used consists of the device I.D., the page fault count, a "no page fault" switch, the number of pages in core for the segment and the current segment length, in pages. Of these, only the device I.D. (see Chapter 7) is crucial to page fault handling - the others are measurements for tuning purposes and aids to Segment Activation and Deactivation which can be best provided by the page multiplexor. Additional comment is made on these items in Section 6 of this paper.

The third, and last data base used by the page fault handler is the core map. The core map is also allocated in the System Segment Table whose header contains pointers to the head of the free list and to the next entry to be examined on the used list. Each core map entry (CME) has two threading pointers: the entries in the free list use only the "forward" pointer, since entries are removed from the top and added there, while the used list employs both "forward" and "backward" pointers to allow insertion of an entry between any two other entries. A CME on the free list has no other useful information contained in it. The core location of the block controlled by a CME is implicit in its position within the core map, since as many CME's are allocated as there are 1024-word blocks of core. We note that not all CME's are put on the free or used threads - only those blocks to be serviced by the page fault handler. All permanently core-resident pages of core are represented by CME's which are not pointed to by entries on one of the threads. In this way, the removal algorithm need not explicitly check entries which could never be removed.

Two other items are kept in the CME for a page only if the CME is on the used list. First, the segment map entry for the page, which was kept in the PTW while the page resided in secondary storage, is transferred to the CME before the core address is inserted into the PTW. Clearly the core address could be kept in this space in the CME when it was threaded on the free list if the information were not implicitly available. Second, a pointer to the PTW is maintained to permit the segment map entry to be replaced and the page fault switch to be reset when the page is removed from core.

We are now in a position to understand the initial requirements of the page fault handler in order that it be able to function. First, the core map entries for multiplexible blocks must all be threaded onto the free list. Then, for each segment which is to be referenced, each page table word must be filled in with a segment map entry and page fault switch setting; and also the Active Segment Table Entry must be initialized in its Device I.D. and other entries. This work would allow a page fault to be serviced for non-page-fault-handling procedures.

But what of the page multiplexor itself? May it use the same page fault mechanism which it provides? Not entirely. While selected parts of the data and procedures of the multiplexor could be transportable, at least one path through the multiplexor must be guaranteed page-fault free to prevent infinite recursion. In Multics, the choice has been made to prevent any page faults whatever from occurring during the handling of a page fault. To this end, all of the page fault handling procedures are permanently core-resident, and for this reason the ASTE is used by the page multiplexor. For although all of the necessary information about the pages of a segment can be found in the branch for that segment, branches are not wired down and are, therefore, susceptible to page faults. Hence, that part of the information kept in the branch which must be referenced by the page fault handler is transferred to the ASTE at segment activation time, to keep it in a permanently core-resident data base. This choice permits Multics to follow a more predictable and shorter path while handling page faults.

6. ADDITIONAL REMARKS

Several important functions of the page multiplexor have been excluded from the previous discussion since they are not necessary to allow page multiplexing. Some of these functions receive additional coverage in chapters written to explicate the areas in which they are used, but their appearance in this paper is germane to a better understanding of the page fault handler's importance in Multics.

An especially complex area in any multiplexed computing system is that of synchronization of processes. Empirical evidence has led Multics to the path of least complexity where possible. An example is the synchronization of multiple processes, all of which may desire the handling of a page fault at the same time. To prevent interference between the various processes in handling common data, only one process is allowed to execute in the page multiplexor at a time. Other processes wishing to deal with a page fault are forced to wait their turn.

It is also necessary to synchronize the physical devices which transport pages and the processes which have requested the transportation. As we have seen, there is no problem of synchrony when writing pages. Any process which needs to have a write operation completed in order to continue its computation simply loops on the core map free list and calls to the DIM until a block appears in the free list. When reading pages, synchrony is established through communication with the Traffic Controller - both to wait and to notify. The only remaining problem is to ensure that the DIM is called subsequent to the completion of every read so that notification can be performed.

The DIM is normally called by the next process to take a page fault. But if the DIM is not called in a sufficiently long time (this could happen if each process were waiting for a page!), the secondary storage devices cause interrupts which are directed to a special process whose sole purpose is to wait for these interrupts and call the DIM in response to them (see Chapter 8).

Another function performed by the page multiplexor is the assignment of blocks of secondary storage (see Chapter 7). No secondary storage is assigned to a page until it has been referenced. A special device address (the "null" address) is assigned to all pages of a segment which have never been referenced. When such an address is encountered by the page fault handler, it creates a page of zeroes in core rather than reading in data. Ideally, then, the page need not be written out until the page-has-been-modified (PHM) bit has been turned on by the hardware. In fact, the Multics page fault handler causes the assignment of a secondary storage address at first reference and sets the PHM bit to force write-out of the page. This sometimes results in storing a page of zeroes in secondary storage but eliminates a check for "null" addresses in the page removal path.

Actual device storage handling is incorporated into a separate module whose algorithm for assigning storage on a particular device can be easily changed to accommodate any system discipline (such as directory segments on Drum and non-directory segments on Disk).

The movement of data from one device to another is also accomplished in the page multiplexor. Chapter 7 describes the function in detail, but the basic mechanism used is an additional item in the ASTE for a segment which can specify a "move device I.D." and a bit in the segment map entry which specifies whether or not the page (which must be in core) has already been moved. Using these items, a segment can be moved by setting the move device to I.D. when activating the segment. Then, whenever a page is brought to and subsequently removed from core, it is rewritten onto the new device.

The final functional area incorporated into the page multiplexing modules is that of services to the segment activation and deactivation module. A special set of entry-points allows individual page manipulation on demand. Specifically, the functions are:

- 1) To read or write a page from/to secondary storage.
- 2) To "wire" or "unwire" a page by setting the "wired down" bit in the PTW which allows a page to be skipped by the removal algorithm (this function is

used only for temporary wiring). Pages are "wired" (made permanently core-resident) by leaving their CME's off the core-map used list.

- 3) To truncate a segment by destroying all pages beyond a certain page and returning their secondary storage to the free pool.
- 4) To cleanup all traces of a segment in preparation for its deactivation. This function consists of exporting the entire segment (removing all of its pages) and waiting until the pages have all been exported to their resident secondary storage.

7. THE HISTORY OF THE MULTICS PAGE MULTIPLEXOR

The entire Multics file system has gone through two incarnations. The original version carried the Multics penchant for elaborate original design to some lengths and was successfully implemented. Its performance and amenability to debugging left something to be desired. Therefore, this second attempt was made, using the knowledge gained from the first to avoid areas of difficult implementation and slow execution.

There were two principal changes respecting paging. First, a single page size of 1024-words was chosen, replacing the previous strategy in which pages of two sizes, 1024-words and 64-words, were allowed. This simplification resulted in the elimination of elegant but time-consuming algorithms for page removal and for "change-making" and coalescing free blocks in core and in secondary storage. Second, a single segment size of 64 pages (implying a single page table size of 64 words) was chosen, replacing the previous strategy in which segments could vary in size from 64 to 256 pages (always in units of 64 pages). This simplification resulted in a greatly simplified Page Table-Active Segment Table Entry arrangement to the benefit of all the modules involved: page control, segment control, core control. As a result of these changes, the present implementation has avoided several lengthy computations in the most frequently used path in page fault handling, achieving a great advantage in average execution time over the former implementation.

Chapter 7

SECONDARY STORAGE MANAGEMENT

1. INTRODUCTION

Secondary storage management cuts across many parts of the Multics virtual memory system. In this chapter, we shall try to minimize repetition by discussing only those points which have not been discussed elsewhere.

We shall discuss the assignment of a segment to a secondary storage device and the assignment to its pages of blocks of secondary storage. We shall also discuss "moving" a segment from one secondary storage device to another, i.e., changing a segment's assigned device.

2. ORGANIZATION OF SECONDARY STORAGE

The physical devices used for storing information in Multics--core, disks, drums--are divided into 1024-word "blocks" corresponding to the division of segments into 1024-word "pages". Addresses of blocks in secondary storage are given as pairs:

block address = (device identifier, block number)

For the most part, the "device identifier" specifies a particular physical device. It is possible, however, that one device identifier specifies part of a large device or a collection of small devices. We should, therefore, use the phrase "logical device identifier."

Each (logical) device of secondary storage has an associated "Device Map" which records which of its blocks are assigned to pages of segments and which are free. The "Device Map" contains one bit per block of the device. This bit is set to "1" to indicate that the block is free and to "0" to indicate that the block is assigned to a page.

To assign a block on a device to a page, it suffices to search the appropriate Device Map for a bit set to "1", note the corresponding block number, and reset the "1" to "0". To free a block (when a page is destroyed, for instance) it suffices to reset the corresponding bit in the Device Map from "0" to "1".

3. SECONDARY STORAGE OF SEGMENTS AND PAGES

3.1. Strategy for Secondary Storage of Segments and Pages

3.1.1. Segment Assignment. Devices of secondary storage are not equivalent. Due to differences of latency and transmission timings and to differences in the accessing code, some devices prove to be faster than others. For example, in the present Multics configuration, the drum is faster than the disk. Because of this non-uniformity, each segment is assigned to a single secondary device: segments expected to be used often are assigned to fast devices, segments expected to be used more rarely are relegated to slower devices. (When we say that a segment is assigned to a device, we mean that the pages of the segment are to be stored in blocks of that device.)

A segment is assigned to a device of secondary storage when the segment is created. The algorithm by which segments are initially assigned to devices of secondary storage is called the "Multi-Level Storage Algorithm". The phrase "Multi-Level" emphasizes the differences in device characteristics. A discussion of the Multi-Level Storage Algorithm is beyond the scope of this paper. We shall examine only the mechanisms used to execute this algorithm's decisions.

3.1.2. Page Assignment. When a segment is created, its 64 pages are also, in a sense, created. But before a page is referenced for the first time, it cannot contain any information.

Although a few segments may ultimately contain 64 information-filled pages, many segments never contain more than a few such pages. It would be wasteful to tie up blocks of secondary storage for pages that contain no information and may never be referenced. Therefore, pages are assigned blocks in secondary storage only after they have been referenced.

3.2. Data Relating to Secondary Storage of Segments and Pages

3.2.1. Segment Map. We may now specify the "Segment Map", that attribute of a segment which tells where the pages of the segment are stored in secondary storage. The Segment Map consists of:

- device identifier
- 64 block numbers

3.2.2. The "Null" Block. A special block number, called the "null" block number, is used to indicate that a page has not been assigned a block of secondary storage. We often say of such a page that it is assigned the "the null block". The null block may be regarded as a page of zeros.

3.3. Procedural Implications of Secondary Storage Strategy

3.3.1. Creating a Segment. When a segment is created, the Multi-Level Storage Algorithm is used to assign the segment to a device of secondary storage. The segment's 64 pages, which have not been referenced, are all assigned the "null" block. The device identifier and the 64 "null" block numbers are all recorded in the segment's Segment Map.

3.3.2. Bringing a Page to Core. When a page is referenced for the first time, the Page Fault Handler (PFH) is asked to "bring to core" a page which is "stored" in the "null" block. The PFH handles such a request by:

- assigning a block to the page from the segment's assigned device.
- zeroing out the block in core which is to contain the new page.
- setting the "page has been modified" switch in the page's PTW to make sure that the page will ultimately be moved to its newly assigned block.

3.3.3. Removing a Page from Core. When the PFH wishes to use a block of core presently occupied by a page, it inspects that page's "page has been modified" switch. If the page has been modified, then it must be written into its assigned block in secondary storage. If the page has not been modified, then it may be overwritten directly since it is equivalent to the information in the assigned block of secondary storage.

4. MOVING A SEGMENT FROM ONE SECONDARY STORAGE DEVICE TO ANOTHER

4.1. Strategy for Moving a Segment

It is occasionally necessary to "move" a segment from one secondary storage device to another. A "move" is necessary, for instance, if a secondary storage device becomes full or if a segment's usage changes substantially. The decision to re-assign a segment to a new device (to "move" a segment) is made by the same Multi-Level Storage Algorithm which assigned the segment to a device at segment creation time.

4.2. The Segment Map Revised to Permit "Moves"

The Segment Map must be revised if it is to be possible to move a segment from one device to another. The Segment Map must indicate not only the device to which the segment is assigned (in the case of a move, the segment is assigned to the new device) but also the device to which the segment was assigned. During the move, information must also be stored to show to which of these two devices the segment's 64 pages are assigned. Last, the Segment Map must show whether or not a "move" is in process.

The revised Segment Map has the form:

- device identifier (or, in case of a move, "old device identifier")
- new device identifier (non-zero only if a move is in progress; thus acts also as a "move in progress" switch)
- 64 block numbers
- 64 "moved" switches (a page's "moved" switch shows to which of the two devices the pages are assigned; the "moved" switch is meaningful only when a "move" is in progress)

4.3. Procedural Implications of "Moves"

4.3.1. Bringing a Page to Core. When a page is brought to core, it must be brought from its presently assigned block in secondary storage. The Page Fault Handler (PFH), by inspecting the "new device identifier", "old device identifier", and the page's "moved" switch, can determine the device to which the page is assigned. The page's secondary storage address then consists of the device identifier so calculated and the page's block number.

In the case of a "null" block assignment, the page referenced for the first time is assigned to a block of the new device.

4.3.2. Removing a Page from Core. When the PFH removes a page from core it must see that the page is removed to a block in the correct device. This means that:

- (a) If the segment is being moved from one secondary storage device to another, and
- (b) if the page's "moved" switch shows that the page has NOT been moved, then the PFH must
- (c) release the block assigned to the page on the "old" device,
- (d) assign a block to the page on the "new" device,
- (e) record the new assignment in the Segment Map, setting the "moved" switch, and
- (f) move the page from core to its newly assigned block in secondary storage.

4.3.3. Deactivating a Segment - Completing a Change of Devices. We know that when a segment is deactivated, the Segment Fault Handler calls a special entry of the paging module to force the segment's remaining pages out of core. Before it removes any pages from core, this procedure checks to see if the segment being deactivated is being moved from one device to another. If the segment is being moved, code is executed which brings to core those pages of the segment which are

stored in blocks of the "old" device (that is, pages which have non-"null" block numbers and whose "moved" switches say "not moved".) When this has been done, the job of removing the segment's pages from core is performed. The pages are removed from core as described in the previous paragraph. Thus, at the end of the page removal, all of the segment's pages are necessarily assigned to blocks of the "new" device.

When the page removal procedure returns to the Segment Fault Handler, the latter updates the Segment Map to show the correct device identifier, a zero new device identifier, and all "moved" switches showing "not moved". With this, the move is complete; we see that deactivation completes a "move".

4.3.4. Performing a "Move". We may now describe the procedure which the Multi-Level Storage Algorithm uses to move a segment from one device to another. The "new-device" identifier is placed in the Segment Map of the segment (in the branch if the segment is not active, in the ASTE if the segment is active). If the segment is not active, it is made active. Finally, the segment is deactivated by means of a call to a special entry of the Segment Fault Handler. By supplying a "new-device" and activating the segment, the "move" is initiated. By forcing the deactivation of the segment, the "move" is terminated, as described in the previous paragraph.

Chapter 8

DEVICE INTERFACE MODULES

1. INTRODUCTION

The responsibility of the Device Interface Module (DIM) respecting paging is (a) to initiate transfers of pages between blocks of core and blocks on secondary storage devices as requested by the Page Control Module, and (b) to notify the Page Control Module upon the completion of these transfers. In general, there is a considerable time lapse between the performance of these two functions.

Page Control uses the DIM by (a) initiating a transfer and then (b) waiting to be notified by the DIM of the completion of the transfer. The DIM must, therefore, perform its notification function with respect to a given transfer without being called by the process which requested that transfer.

Notification of completed transfers is usually performed by the DIM just after the latest transfer is initiated. The DIM inspects its data bases, determines which transfers (by whatever process requested) have been completed, and performs the necessary notifications. It then returns to its caller which may in turn wait for notification.

It may happen that all processes are waiting and, thus, that no process will (by taking a page fault) invoke the DIM. Therefore, as a precaution, the DIM arranges to have an interrupt sent to the processor by the secondary storage device sometime after it completes its last pending transfer. This interrupt will cause a special process to be wakened which will call the DIM and cause the required notifications to be performed. The DCW (see below) which causes this interrupt also disconnects the controller; we call it the "last" DCW.

2. GENERALITIES

2.1. DCW's

We are concerned with the I/O transactions of transferring a page from a block of core to a block of secondary storage (or vice versa) involving the GE-645 computer and its peripheral units. The Page Control Module "requests" such a

transfer by a call to the DIM of the appropriate device. The DIM passes the request along to the device controller via a special word, called a DCW - Data Control Word. A DCW contains the following information of interest to us in this paper:

- op-code read, write, or no-op (on some devices)
- core address
- device address
- disconnect bit causes the device to disconnect after satisfying the request specified in the op-code
- interrupt bit causes the controller to send an interrupt to the processor after the request has been satisfied

2.2. DCW Lists

Each device controller is driven by a list of DCW's which it runs through, consecutively, interpreting the DCW's, until it encounters a DCW with a disconnect bit set after satisfying which the controller disconnects itself and waits to be reconnected. The DCW lists are all regarded as circular in the sense that the controller accesses the DCW's consecutively, modulo some N. The DCW lists are finite in the sense that the DIM's always store a DCW with a disconnect bit somewhere in each DCW list. This DCW, whose interrupt bit is also set, is called the "last" DCW.

2.3. Status Queues

While the device controllers obtain their instructions from DCW lists, they record the status of requested transfers in special "status queues". A word in the status queue is associated with a DCW and written into by the controller after the transfer specified by the DCW is begun. The status word will show whether the action begun was completed, whether there was a parity error, etc.

2.4. Function of the DIM

The DIM for a particular device interfaces with the device through the two data bases discussed above, the DCW lists and the status queues. The DIM interfaces with the user (the Page Control Module, in our case) as follows:

- on being called, the DIM sets up appropriate DCW's,
- makes sure that the interrupt and disconnect bits are set in the proper DCW, and
- connects the controller if necessary.

After acting for the particular user, as above, the DIM

- inspects the status queue,
- "posts" all completed transfers in the associated Page Table Words,
- "notifies" the processes waiting for the completed transfers, and
- cleans up the SDW list and status queue as required.

2.5 Normal Operation

It is expected that enough page faults will occur that "requests" for page transfers will, in general, occur while each DCW list is non-empty. This means that the controller has not yet reached the "last" DCW with its disconnect and interrupt bits set. The DIM accordingly establishes new DCW's (as indicated above) and then advances the "last" DCW, that is, resets the disconnect and interrupt bits in the DCW in which they are presently set and sets them instead in the now appropriate DCW, further along on the circular DCW list. (We will discuss in detail below just which DCW is "last".) In this way, it is expected that the controller will run for a long time without reaching the "end" of the DCW list and will consequently remain connected and will not have to send interrupts. Since the purpose of the interrupts would be to force the invocation of the DIM to "notice" the completion of transfers, and since the DIM will be called quite often as part of paging, there is no need for interrupts. The interrupt associated with the "last" DCW takes care of the unlikely case that all processes are waiting for page I/O and that no more calls to the DIM from the "user" will occur. In this case, the interrupt will force a special process to be wakened which will call the DIM and so enable the noticing of completed transfers, their "posting", and the associated notification.

3. DRUM

3.1. Drum Configuration

The drum contains $M \times N$ blocks of 1024 words as shown in Figure 1. As the m th row of blocks comes under the read/write head, any of the N blocks in the row can become part of a transfer.




Figure 1. Drum Configuration

3.2. Drum's DCW List Configuration

The drum DIM maintains a circular DCW list for the drum of length $L \times M$ where $L \gg N$. Each DCW contains the following information:

(read, write, or no-op), core address, m , n , ("last" or "not-last")

where "last" means the disconnect and interrupt bits are set (in one DCW only), m is the row number, and n is the number of the block in the row. The drum's DCW list is initialized with all of the m 's set, in order, so that the DCW list consists of L one-entry-per-row coverings of the drum. As the drum rotates and as the controller advances through the DCW list, the number of the drum's "presenting" row equals the row number of the DCW then pointed at by the controller. For this reason, the drum DCW may have a no-op

op-code to specify that no transfer is to occur involving a block on the mth row at this time. See Figure 2. It should be emphasized that this L-fold "covering" of the drum by the DCW list and the consequent parallelism between the drum's physical position and the words of the DCW list is not required by the hardware but is a software construct (of some beauty.)




Figure 2. Drum DCW List Configuration

3.3. The "Last" DCW

Whenever it is called, the drum DIM resets the "last" DCW to be the DCW last passed by the controller, that is, the DCW for which the drum controller has just completed the recording of final status. This guarantees that the drum controller will never generate an interrupt (and disconnect) until L revolutions of the drum after the last call to the drum DIM. The drum DIM maintains a pointer to the "last" DCW and, upon being called, erases the "last" information from the presently "last" DCW and writes it into the DCW just passed by (as explained above), resetting the "last" pointer.

3.4. DCW List Management for the Drum

When a transfer request is sent to the drum DIM, a pair (m,n) is sent as drum block address. The DIM examines the DCW list and finds that unique DCW in the L*M DCW's of the list which (a) will first be in the controller's path, (b) has a no-op op-code, and (c) has the given value of m in its row number slot. There are L DCW's with row number m and it is expected that at least one of them is no-op. (If all L of them are in real use, the DIM loops, waiting for one of them to come free.) When such a DCW is found, the DIM writes the appropriate value of n into the DCW and sets the appropriate op-code (read or write). The DIM then goes through the usual steps of setting the "last" position correctly, observing, posting, and notifying for completed transactions. One part of cleaning up after completed transactions is to reset to "no-op" the op-code of DCW's whose requests have been serviced.

It is expected that the drum controller will, in general, continuously run through the circular DCW list, the "last" DCW running along L revolutions behind it, with a band of non-null DCW's just in front of the current controller position in the list. See Figure 3.




Figure 3. Density of Non-Null DCW's in Drum's DCW List

4. DISK

4.1. Disk Configuration

The "disk" consists of l platters each with one movable I/O head. The I/O head for any platter may be moved to any of J tracks. Each track has K blocks which come under the I/O head as the disk rotates. A disk DCW specified:

(read or write), core address, i , j , k , ("last" or "not-last")

4.2. Disk DCW List Configuration

The disk's DCW list is circular only in the sense that its DCW's are accessed by the controller consecutively modulo the list length. DCW's are put on the list in the order received; there are no no-op DCW's, and the last DCW to be put on the list is always the "last" DCW in the sense that it contains the disconnect and interrupt bits.

4.3. Expected Operation

The Disk DCW list is ordered randomly in the sense that requests are put on the list as received and hence without regard for the position of the disk (k) or of the l arm positions $j(i)$. This technique is chosen because of the large amount of processing that would otherwise have to be done to keep track of the arm positions, the value of k , and the list of unsatisfied DCW's; and because of the low probability that such processing would pay off.

No interrupt and disconnect will occur as long as there is at least one request in the queue. Since no reuse of old DCW's is made (as in the drum's case), there is no need, in cleaning up, to erase the contents of the DCW's of completed requests.

B. Access Control to the Multics Virtual Memory

CONTENTS

<u>Chapter</u>	<u>Title</u>
	Introduction
1	Access Control Philosophy
2	Multics Ring Structure Philosophy
3	Software Functions in Ring Changing
4	Simulation of Rings Using the 645
	Bibliography

INTRODUCTION

An important trend in the design of large computer systems is the inclusion of hardware and software for the sharing of information, both procedure and data. Thus, the concept of pure, re-entrant procedure has lost its novelty and the sharing of data, as in multiuser information retrieval systems, has become commonplace. The introduction of sharing into large systems has, however, brought the difficult problem of access control into the realm of the computer system. The comparatively easy problem of protecting the supervisor in a batch environment has grown into the complex task of permitting the flexible sharing of information between system and user and between user and user.

The Multics access control system has been described in a number of places with a number of purposes. Graham³ discusses the fundamental reasoning behind the choice of the Multics ringed access control system; Organick⁴ discusses the details of the implementation and use of this system; and the Multics System Programmers' Manual goes into even greater detail on implementation. The purpose of the present paper is not to duplicate any of the excellent material already available but rather to highlight certain aspects of the Multics ringed access control system which are thought to be of particular interest to system programmers.

In this paper we shall develop the ideas of the ringed access control system as an approximation of access control conditioned on the identity of the procedure in execution, as suggested by Evans and Leclerc²; we shall describe "ringed hardware" to support the ringed access control system; we shall show how this "ringed hardware" is simulated on the 645 processor; and we shall discuss at some length the software mechanisms which are implied by the concept of "ring".

This paper was written in conjunction with, and logically follows, another paper, The Multics Virtual Memory¹.

Chapter 1

ACCESS CONTROL PHILOSOPHY

In the Multics virtual memory, the segment is the unit of information to which access is controlled. In fact, the possibility of controlling access to shared information was a principal justification for designing a segmented memory system. In Multics, every segment is directly addressable and it is, therefore, necessary, upon each attempted memory access, for the accessing hardware to answer the question:

Shall this attempted access be permitted?

The answer to this question, with respect to a given segment, is obtained by interpreting a data base associated with the segment, the segment's "access control attributes". It is the purpose of this chapter to discuss the basis on which the hardware might go about answering this question, hence, to specify the content of a segment's access control attributes.

We feel that, at the least, a segment's access control attributes should indicate:

1. who may access the segment; a segment may be accessible to a single user only or shared by a number of users.
2. how each of these users may access the segment; distinct users may have distinct access rights.
3. in what circumstances each user may exercise his access rights to a segment; a user's rights may be made to depend, in some way, on what he is doing.

User-Name. Let us look at these points in turn. To begin with, it is a process executing on a processor which attempts to access memory, not a user. For this reason, every process has associated with it the name ("user-name") of the user on whose behalf it is executing; all access rights of the process derive from the process' user-name.

We note that the security of an access control mechanism depending in this way on the user-name depends strongly on the technique by which a process is assigned a user-name.

A simple and perhaps sufficient technique for assigning user-names to processes is to require each user, when he "logs in", to specify a user-name and then give a secret password which validates his right to use the given user-name. All processes which subsequently act for him as a result of this "login" will then do so with the authority of the given, validated user-name.

Access-Mode. The basic types of memory access are READ, WRITE, and EXECUTE. We use the term "access-mode" to refer to any combination of these types including the null combination. It is clear that a process' access rights respecting a segment are at any moment characterizable in an access-mode.

If a process' access rights were to be independent of its activities, then a segment's access control attributes could be recorded in a list of user-name/access-mode pairs. A process' right to access a segment in a given way would then be determined by (a) whether the process' user-name appeared in the segment's list and (b) whether the given access type appeared in the corresponding access-mode. This system of access control is illustrated in Figure 1. The access control mechanism takes as arguments the process' user-name, the type of the attempted access, and the name of the target segment. It then searches the target segment's access control attributes list for the given user-name. If it is found, the corresponding access-mode is then searched for the given access type. The access is permitted only if the user-name and access type are found.

The system of access control just described is already quite powerful. It permits a user who has created a segment to grant himself READ-and WRITE-access to it, to store information in the segment, and then to give a number of other users READ-access to it. He and these others may now read the segment whereas he retains for himself the right to change it.




Figure 1. Illustrating an Access Control Mechanism Depending on User-name and Access-type

Circumstance-Dependence of Access Rights. There are two reasons why a process' access rights should somehow be made to depend on the process' current business. First, the problem of error may suggest that access, particularly WRITE-access, to a segment should be limited to debugged procedures or groups of procedures. Second, the problem of intentional misuse of a segment may suggest that semi-trusted users be forced to access the segment via procedures or groups of procedures specifically designed for their use.

As an example of the latter, consider a Management Information System with a data base including individual salary information. This data base would generally be "readable" by all users of the system; but the less privileged users would have to "read" the segment via procedures designed not to disclose individual salaries.

In the remainder of this chapter we shall take for granted the dependence of access rights on user-name and shall concentrate on finding a good and workable way to make a process' access rights to a segment circumstance-dependent.

1. ACCESS CONTROL BY PROCEDURE

The most obvious way to achieve circumstance dependence in access control would be to condition a process' access rights on the procedure by which the access is attempted. A segment's access control attributes would then be recorded, in effect, in a user-name versus procedure table whose entries are access-modes.

Figure 2 illustrates this system of access control. The access control mechanism takes as arguments the process' user-name, the type of attempted access, the name of the procedure in execution, and the name of the target segment. The target segment's access control attribute table is then searched for an entry corresponding to the given user-name and procedure. If the entry is found, the corresponding access-mode is then searched for the given access type. The access is permitted only if an entry with a suitable access-mode is found.

This type of control would permit access control as illustrated in the following example:

user-1 "owns" data segment D and procedure P and gives user-2 WRITE-access to D only when executing P and EXECUTE-access to P only when executing in procedure P (and, of course, P itself).

This implies that user-2 can only write in D by calling P from P (to which user-2 has access from some source other than user-1).




Figure 2. Illustrating an Access Control Mechanism Depending on User-name, Access-type, and Procedure

2. ACCESS CONTROL BY SET

The idea of conditioning access rights on the procedure-in-execution has been proposed by Evans and Leclerc² and is an idea that occurs to many system programmers at some point when they are struggling with difficult access control problems. We would recommend this technique were it not for some difficulties which render it infeasible. The principal hindrances to the conditioning of access rights on the procedure-in-execution are:

- no hardware presently exists which would permit this type of access control to be practiced in any but an interpretive mode
- too much effort and space must be expended in constructing and updating each segment's table of access control attributes
- too much must be foreseen: This technique requires knowledge of all of the uses to which each segment may legitimately be put.

A natural idea for approximating the procedure-in-execution strategy is based on grouping related procedure segments into "sets" and basing access rights to segments on the identity of the set to which the procedure attempting the access belongs.

There is no reason, by the way, to suppose that these sets of procedures would be disjoint; indeed, service procedures such as PL/1 runtime routines would probably be included in every set.

Access control by procedure-set appears to have two advantages over access control by procedure. First, each segment would have a somewhat smaller table of access control attributes, a practical system presumably having fewer sets than procedures in sets. Second, updating the per-segment access control attributes tables should be easier, since adding another procedure to a set would mean revising the definition of the set, not amending the access control attribute tables of a large number of segments.

Figure 3 illustrates access control based on set-in-execution. The form of a segment's access control attributes and the interpretation of these attributes by the access control mechanism are just as in access control by procedure, as described in Section 1 above, except that "set" replaces "procedure" wherever it occurs.

The concept of "sets" introduces a number of interesting problems. Given that each procedure is potentially an element of several sets, and stipulating that a change of set can only occur upon a change of procedure (i.e., upon a call or return), how shall the access control mechanism determine to which set to change (if any) upon each transition between procedures? How shall the composition (membership) of sets be initially defined and by what mechanism shall the composition of sets be changed? Shall the composition of sets be determined in a system-wide way or per-user or per-project? And so on.

We have introduced this concept of "sets" of procedures in order to make the definition of "rings" (see below) less abrupt and also to put the concept of "rings" in perspective.

3. ACCESS CONTROL BY RING

The implementation of an access control strategy based on sets, as described above, is judged infeasible due to the difficulty of defining sets, of unambiguously defining all transitions between sets, etc. It is useful to define a restricted theory which produces the more useful concept of "rings".

We use the term "rings" to refer to "sets" (as described above) to which access rights to all segments are assigned in such a way that the sets can unambiguously be ordered by increasing power or privilege. Precisely, we say that a collection of sets is a collection of rings if the sets can be numbered 0, 1, 2, ... in such a way that the possession by a particular user of an access right to a segment in set k implies the possession by that user of that segment for all sets j, $j < k$.




Figure 3. Illustrating an Access Control Mechanism Depending on User-name, Access-type, and Set

A corollary of this definition is that a user's access rights to a segment can in part be expressed as an access-mode and a triple of ring numbers - $r(\text{READ})$, $r(\text{WRITE})$, and $r(\text{EXECUTE})$ - indicating that a given access type, X , if present in the access-mode, is to be available to the user in the rings 0 - $r(X)$, inclusive. We shall defer to the next chapter consideration of how a process changes from one ring to another.

Figure 4 illustrates access control base on the ring in which the process is executing. The essential point to notice is that a segment's access control attributes can be very concisely recorded. The interpretation of the access control attributes is as discussed in the preceding paragraph.

A few comments about rings may be in order. First, the introduction of rings greatly simplifies the recording of a segment's access control attributes, as indicated above. Second, the fact that rings are ordered removes the ambiguity about the changing of power that was inherent in the idea of a transition between sets: when the processor changes from ring i to ring j , $j > i$ implies an increase (or at least no decrease) of power or privilege with respect to all segments; and $j < i$ implies a decrease. This homogeneous and evident change of power with the change of ring makes it much easier to think about the problems of changing rings than it could ever have been to think about the changing of sets. As we shall see, and notwithstanding the previous remark, most of the difficulty in the fully worked out strategy of access control by ring nevertheless resides in the mechanics of changing rings.

The following points seem to be necessary elements of any access control strategy based on the idea of rings:

- Attempts by the processor to pass control from one ring to another must be supervised by the access control mechanism.




Figure 4. Illustrating an Access Control Mechanism Depending on User-name, Access-type, and Ring

- The transition from one ring to another can only occur upon a call or return; the transition (if any) associated with every call and return must be unambiguously defined.
- The concept of a "return" from a call must be extended to imply returning to the ring from which the call was made.

3.1 Multics Terminology; Gates

In Multics the term "inward" is used to characterize a transition from one ring to a more privileged ring, the term "outward" to characterize a transition in the other direction. Procedures which may be called by "inward" calls are called "gate" procedures since they are, in effect, gates through which the processor may enter the more privileged ring. We shall see that the major difficulties in the design of a ringed access control system relate to allowing outer ring procedures to use inner ring procedures without allowing them to defeat the protection purposes responsible for the existence of rings.

Chapter 2

MULTICS RING STRUCTURE PHILOSOPHY

In Chapter 1 we discussed access control and developed the idea of rings in a general (i.e., non-Multics) way. We now turn to Multics itself. In this chapter we shall enlarge on the subjects of rings and gates, state and justify the full Multics ring structure strategy, and show how this strategy can be implemented with the aid of suitable hardware.

In this and the following chapters, the emphasis in the definition of "ring" will shift slightly. We will think of a ring not as a process state defined by a set of procedures, but rather as an abstract process state in which, by virtue of the access control rules of the system, a particular set of procedures may be permitted to execute. There are 64 rings in Multics which are conventionally numbered, in order of decreasing power, from 0 to 63.

1. THE PRELIMINARY STRATEGY

A preliminary (and conceptually useful) idea for the use of rings is based on specifying a user's access rights to a given segment with an access-mode, a single ring number "r", and a gate-switch.

The Rules. The ring number r, gate-switch and access-mode are interpreted as follows. (Note that all ring intervals are inclusive).

- a. If the user's access-mode contains WRITE, the user may, in rings (0,r), write in the segment.
- b. If the user's access-mode contains READ, the user may, in rings (0,r), read the segment.

- c. If the user's access-mode contains EXECUTE, the user may,
 - 1. in ring r , call and execute the segment
 - 2. in ring $R < r$, call the segment, switching to ring r to execute it
 - 3. in ring $R < r$, but only if the gate-switch is set, call the segment, switching to ring r to execute it

Every attempt by the process to switch to a lower numbered ring in this way must pass a legitimacy test imposed by the access control mechanism and by the procedure being entered.

- d. All ring switching must be done under the supervision of the access control mechanism.
- e. The concept of "return from a call" must be extended to imply a return to caller's ring.

The Need for "Gates". Since an "inward" call (i.e., a call through a "gate") increases the processor's power, it is necessary that a test be made to verify that the process has attempted to enter the more powerful ring on a legitimate errand. For, if the process could freely change its ring so as to increase its power, the protection offered by the ring aspect of the access control mechanism would be wholly illusory. The kind of testing is occasioned by an attempted inward ring change will be discussed in detail in Chapter 3. As an obvious example, we note that a call to a gate segment should be permitted only if the target address is in fact an entry point of the segment.

2. THE "RING BRACKET" STRATEGY

The principal difficulty with the "preliminary" strategy described above is that procedure segments may be executed in one ring only. This means that a procedure likely to be called in several rings will often be called from a ring other than its ring of execution, occasioning a great deal of ring changing, an expensive business as we shall later see. A second difficulty

with the "preliminary" strategy is that users with both READ- and WRITE-access rights for a segment have these rights equally in all of the rings from 0 to r. Since the ability to write in a segment is intrinsically more powerful than the ability to read it, it would be desirable to be able to grant write permission to a user in a (relatively privileged) subset of the rings in which he may read. As a result of these and other considerations, Multics has rejected the "preliminary" strategy for a "ring bracket" strategy.

Under the "ring bracket" strategy, a user's access rights respecting a given segment are encoded in an access-mode and a triple of ring numbers, (r1, r2, r3), called the user's "ring brackets" for the given segment.

The Rules. The ring brackets, (r1, r2, r3), which must satisfy the relations $r1 \leq r2 \leq r3$, are interpreted as follows. (Note that all ring intervals are inclusive).

- a. If the user's access-mode contains WRITE the user may, in rings (0,r1), write in the segment.
- b. If the user's access-mode contains READ the user may, in rings (0,r2), read the segment.
- c. If the user's access-mode contains EXECUTE the user may,
 1. in rings (r1,r2) call the segment without changing ring
 2. in rings (0,r1-1), call the segment, switching to ring r1
 3. in rings (r2+1,r3), call the segment, switching to ring r2

Every attempt by the process to switch to a lower numbered ring in this way must pass a legitimacy test imposed by the access control mechanism and by the procedure being entered.

- d. All ring switching must be done under the supervision of the access control mechanism.
- e. The concept of "return from a call" must be extended to imply a return to the caller's ring.

Under these rules we see that a utility routine may be given ring-brackets (0,63,63) and so be callable in all rings, but never occasion a change of rings upon being called. On the other hand, a critical system procedure might have ring brackets (0,0,0) and so be callable and executable only in ring 0.

We also see that a user who has read and write permission for a data segment may be given ring brackets (a,b,b) with $a < b$ so that the domain in which he has write permission, rings (0,a) is a relatively privileged subset of the domain in which he has read permission, rings (0,b). These comments show how the ring bracket strategy corrects the defects which we noticed in the preliminary strategy.

Ring Changing Calls. Let us now discuss inward and outward calls. The "rules" provide that every procedure segment for which $0 < r_1$ may be entered via an outward call (from ring 0, for instance) and that those procedure segments for which $r_2 < r_3$ are "gate" segments and may, therefore, be entered via inward calls (from ring r_3 , for instance). What is the nature of such calls?

An inward call is made when a procedure in an outer ring wants to increase the power of its process temporarily in order to do a job requiring such increased power. For example, a user procedure may call a system procedure in ring 0. The notion of "inward call" brings to mind "the tail wagging the dog", since lesser power directs the use of greater power. The only segments which can be entered via inward calls are, therefore, the "gate" segments. The duty of a gate segment, as a gate segment, is to perform a test of the legitimacy of the inward call, that is, to see that the caller has not, by accident or design, asked the gate segment to behave irresponsibly. Whether or not a segment is a "gate" for a particular user depends on that user's ring brackets and access-mode respecting that segment.

An outward call is made when a procedure executing in an inner ring wants a job done which can (and perhaps must) be accomplished with the comparatively feebler power of an outer ring. For example, a process in Multics initializes itself (a system function) in ring 0 but calls out to a user ring when ready to do the user's work. In this case, the process must call out since a Multics convention forbids user work to be done in ring 0. For another example, a programmer with a collection of more or less debugged procedures may use several rings, keeping the more debugged procedures and their data in the inner rings so that damage from the other procedures will be isolated in the outer rings. If these procedures call each other freely, outward calls will presumably occur.

3. RECORDING AND RECOVERING ACCESS CONTROL RIGHTS

In "The Multics Virtual Memory"¹, we find that all of a segment's attributes of interest to the system are stored in the segment's "branch" in a "directory" segment. The access control attributes of a segment are stored in its branch in a variable length table called the access control list (ACL). Each entry of a segment's ACL specifies a particular user's access rights respecting the segment and is of the form:

user-name, access-mode, ring brackets

The procedure responsible for determining a user's access rights for a given segment searches that segment's ACL for the user's user-name. If it is not found, then the user has no rights. If it is found, then the user's access rights are determined by the associated access-mode and ring brackets.

4. "RINGED HARDWARE"

In "The Multics Virtual Memory"¹ we discussed the use of the 645's descriptor segment and Segment Descriptor Word (SDW) in providing the Virtual Memory. It was noted that part of each SDW was reserved for an access control field. In this section we shall discuss hardware similar to the 645's which is consistent with the description given in "The Multics Virtual Memory" and which permits a simply described implementation of the Multics ringed access control strategy. In Chapter 4, we shall describe the actual 645 hardware and discuss the software modifications needed to provide for the differences from the hardware described here.

We propose "ringed hardware" with the following features:

1. The processor has a ring register whose value defines the process' ring. This register may be changed by instructions only in ring 0, that is, when its value is 0.
2. The SDW's access control field contains the process' access-mode and ring brackets.
3. The processor has an access control mechanism which checks attempted memory accesses according to the rules stated in Section 2 and causes the processor to fault (trap) to an appropriate procedure in ring 0 in cases where the attempted access cannot be (or cannot be directly) performed. Such a fault causes the hardware to set the ring register to 0.

It should be clear that a procedure executing in ring n should not be able to change the value of the ring register to m , $m < n$, simply by executing an instruction. It might, however, seem that ring changing should be accomplished by the hardware itself, during the execution of a transfer instruction, by a simple change of the contents of the ring register. We have avoided specifying such hardware, for, as we shall see in Chapter 3, changing rings is quite complex and requires considerable software assistance. In light of this fact and considering the hardware organization described above, we may describe the functioning of this system as follows:

When a memory access is attempted, the type of access (read, write, or execute) and the processor's ring register are compared, by the processor's access checking mechanism, with the access-mode and ring brackets fields of the target segment's SDW. As a result of the comparison, three actions may be taken:

1. The memory access is performed and the ring register is unchanged.
2. The memory access is a ring changing transfer; the processor faults to the ring changing fault handler, in ring 0.
3. The access attempted is illegal; the processor faults to a suitable access violation fault handling procedure, in ring 0.

Note that the fault handling mechanism must have the power to change the ring register. This is achieved by making the fault handling procedure executable in ring-0 only, making the hardware enter ring 0 upon taking a fault, and making the ring register changeable (by instruction) in ring 0 only.

Chapter 3

SOFTWARE FUNCTIONS IN RING CHANGING

We indicated, in Chapter 2, that ring changing is a complex activity requiring considerable software assistance. In this chapter we will discuss various aspects of an operating system imposed by a ringed access control system and will discuss the software functions consequently required in ring changing.

We will structure our exposition by separately describing the four types of ring changes, inward and outward calls and returns, attending to points of interest as they arise. That done, we will conclude with a discussion of important facts and concepts and a quick once-over of the ring changing software.

Many of the functions to be described below might be performed, at least in part, in the inner-ring procedure involved in the change of ring rather than in the procedures of the ring changing mechanism, and some of the functions might more naturally be performed there. We take the point of view, however, that the code required to perform ring changes should be concentrated in a single place and we give the ring changing mechanism responsibility for performing all of these functions. In order to handle ring changing in this way, it is necessary to establish certain conventions between the ring changing mechanism and the inner-ring procedures involved in ring changes, as we shall see below.

1. INWARD CALLS

Detection. An inward ring changing call is detected when an inward ring changing fault occurs as the result of a "call" (rather than of a "return") type of instruction. The fault handler obtains the number of the target ring from the process' ring brackets for the target segment; according to the rules of Chapter 3, Section 2, the target ring is "r2".

Gate Address Validation. The handler's first business is to verify that the address to which the caller wishes to transfer is indeed a gate entry point for the process. This verification is based on a "gate-list" (i.e., a list of gate entry points) associated with the target segment; this gate-list may be system-wide in the sense that all users who may use the segment as a gate may use the same gate entry points or may be per-user in the sense that each user who may use the segment as a gate has a private gate-list. In today's Multics, the gate-list is system wide and is stored in the procedure segment (rather than, for instance, in the segment's access control attributes in the segment's branch).

Per Ring Data. We must now consider the nature of the "workspaces" of the calling and of the called procedures. In the ringed environment, all data must be "ring bracketed", including workspace data, e.g., the PL/1 static and automatic data. Since a procedure executing in ring r may freely copy into the (ring r) workspace any data readable in ring r , including all such data not readable in ring $r+1$, it is clear that ring r must use a workspace with ring brackets (r,r,r) . Thus, assuming that any workspace segment has an access-mode implying read and write permission, the workspace for ring r is readable and writeable in rings 0 to r and cannot be accessed at all in the rings $r+1$ to 63. The above considerations imply that the process needs distinct workspace segments corresponding to the rings in which the process executes. Hence, the inward ring changing fault handler will have to provide the proper workspace for the called procedure.

The Environment. We may generalize from the idea of a workspace segment to the idea of an environment. Object procedures expect to execute and, therefore, to be transferred to, in a conventional environment defined by various appropriately valued hardware registers and data structures. Among other things, the environment specifies the workspace to be used by the procedure. (Thus, for example, Multics procedures expect certain processor base registers to be pointing to appropriate "stack" and "linkage" segments). Since this conventional environment is assumed, it is obviously a duty of the ring changing mechanism to create the environment which the procedure being entered will expect to use.

Argument Validation. Now let us consider the arguments which may be passed by the caller to the called (gate) procedure. To begin with, providing a suitably initialized environment for the called procedure involves copying the address of the argument-list (or copying the argument-list itself) into the environment of the called procedure. Certain precautionary measures then become necessary which relate to the need for a "gate" to act responsibly, as discussed in Sections 1 and 2 of Chapter 2. Let us motivate the discussion of these precautions by considering two examples of inward calls which should be aborted by a careful ring changing mechanism.

1. A ring-50 procedure calls a gate procedure in ring-32 and specifies a return argument in the workspace of ring-40. If the call is not aborted, the gate procedure may write in the ring-40 segment at the explicit request of the ring-50 procedure. The gate procedure would thus in effect permit the ring-50 procedure to overwrite the ring-40 segment, a clear violation of the access control philosophy.
2. A ring-50 procedure calls a gate procedure in ring-32, specifying return arguments in ring-50 segments and input arguments in ring-40 segments. If the call is not aborted, the gate procedure may copy ring-40 data into ring-50 segments. The gate procedure would thus in effect permit the ring-50 procedure to read ring-40 data, another violation of access control philosophy.

The responsibility of a gate procedure may be characterized as avoiding the improper use, on behalf of an outer ring procedure, of that part of its accessing power which exceeds that of its caller. To fulfill this responsibility, a gate procedure must, before accessing memory via an address obtained from its caller (or from any other outer ring source), verify that the intended type of access could have been performed by the caller. We shall refer to this as "validating the address". Once all addresses obtained from outer rings have been validated, the gate procedure may freely proceed, since it is clearly safe to use all other addresses.

Although this "address validation" can all be done by the gate procedure itself, our point of view suggests that as much of it as reasonably possible be done by the ring changing mechanism. Since most of the addresses supplied to a gate procedure by its caller are the addresses of arguments, we assign the business of validating these addresses to the ring changing mechanism and we leave it to the gate procedure itself to validate all other suspect addresses. Checking the addresses of the arguments is called "argument validation". Argument validation should include checking that the caller has READ-access for all of the arguments being passed and WRITE-access for those arguments, including "return" arguments, in which the called procedure may write. Argument validation implies a further step in inward ring changing: argument-list copying. For, if a pointer is checked to see that its value may safely be used, then the pointer may not safely be left in a segment where it may be changed by a process executing in a ring less privileged than the gate's. Therefore, all addresses to be checked must be copied into the gate segment's workspace prior to such checking.

It is clear that the argument validation mechanism must make use of an argument-list-descriptor, presumably coded as data, associated with the gate entry point. This descriptor tells how many arguments are expected and how they are to be used (i.e., whether they will be read and/or written in).

The argument-list-descriptor(s) for a gate segment may be implemented in many ways, for example, as part of the gate segment's gate-list. In any case, it is clear that the argument-list-descriptor, like the gate-list, must be supplied to the ring changing mechanism by the gate (inner-ring) procedure rather than by the calling procedure.

Figure 5 illustrates a gate procedure segment, with its gate-list and argument-list-descriptors, in a straightforward implementation. When this segment is called from an outer ring, the ring changing mechanism validates the attempted transfer address by finding it on the gate-list and validates arguments by checking that the caller has the access rights toward the arguments which are specified in the argument-list-descriptor associated with the transfer address.




Figure 5. A Gate Segment Showing Gate Entry Points, Gate-List and Argument-List Descriptors

2. OUTWARD RETURNS

The Detection Problem. The detection of an outward return is not straightforward. Since the procedure to be returned to may well have ring brackets permitting it to execute in the returner's ring and, indeed, in a number of other rings, one may well wonder how the ring changing fault associated with the return is generated and how the ring changing mechanism decides which ring to return to.

An example may make these remarks clearer. Consider a call from ring-20 to a procedure P with ring brackets 5-10-20 and a call by it to a procedure Q with ring brackets 3-7-12. The first call takes the process into ring-10 and the second takes it into ring-7. It is clear that an ordinary return from procedure Q would not cause a ring changing fault. It is also clear that if it did cause a fault, the fault handler would have to choose a ring to return the process to from the interval ring-5 to ring-10.

Forcing A Fault. We see that in the case of a ring changing return, the ring bracket mechanism cannot by itself be dependent upon to cause the necessary ring changing fault or to provide the information required to identify the caller's ring. A special trick is, therefore, used to cause the fault. The normal return pointer in the returner's workspace is over-written with a conventional replacement so that when the process attempts to return via this "return pointer", a fault will occur which is associated with the ring changing return fault handler. This device for forcing a return fault applies equally to inward returns and is also used in that case.

The Return Stack. When the artificial ring changing return fault occurs, as a result of a "return" type of instruction, the ring changing return mechanism is invoked. It must not look in the returner's workspace to find the information that it needs to perform the return - caller's ring number and the return pointer - for these items could be manufactured by the "returner" to imply a "return" to a more privileged ring. The ring changing mechanism, therefore, maintains a ring-0 data base called the "return stack" in which it records all the information needed to perform all uncompleted ring changing returns, both inward and outward. At any time, the last entry on this stack specifies the return from the ring in which the process is then executing. We may now say that an outward ring changing return is a return which causes a ring changing return fault and whose entry in the return stack indicates a return to an outer ring.

Address and Argument Validation. There is no need, from an access control viewpoint, to validate a return address for an outward return since an inner ring procedure may in any case freely enter an outer ring at any point. However, to protect against error, the return pointer recorded in the return stack may be compared against a "validation return pointer" stored in the returner's workspace. Both the validation return pointer and the return pointer in the return stack would be recorded at the time of the corresponding call by the ring changing mechanism. If these return pointers disagree, then the ring changing return can be regarded as an error and treated accordingly.

There is no need for argument validation at the time of an outward return; the work was done at the time of the corresponding call.

The Restoration of the Environment. Finally, let us note that it is necessary, in servicing an outward return, to re-establish the environment that existed at the time of the corresponding call. The caller's workspace must be re-established, base registers must be restored, the entry on the return stack must be removed, etc. Any information which may be needed for this work must be found in the return stack entry for this return and must thus have been stored there at the time of the call.

3. OUTWARD CALLS

An outward ring changing call is detected when an outward ring changing fault occurs as the result of a "call" type of instruction. The ring to be entered is determined from the target procedure's ring brackets; according to the rules of Chapter 2, Section 2, the target ring is "r1". There is no need to validate the target address of the call, for gates govern inward calls only. As with the inward call, there is a need to establish the environment required by the called procedure.

Argument Copying. If, as is usual, the caller's arguments are stored in the caller's workspace, the arguments will be inaccessible to the called procedure in its outer ring. It is, therefore, insufficient to copy only a pointer to the argument-list or the argument-list itself into the workspace of the called procedure. It is necessary to copy the arguments themselves. This in turn implies that a new argument-list must be fabricated in the workspace of the called procedure which contains the addresses of the local copies of the arguments.

There is no need to perform access validation on the arguments. The inner ring procedure may judge for itself what data to pass to the outer ring. The copying of arguments is done, of course, with the authority of the calling procedure's ring, if not by the calling procedure itself. If the copying is actually done in a ring more privileged than the caller's, e.g., by the ring changing fault handler (which executes in ring-0), then the arguments must be access validated to make sure that no data are copied into the workspace of the called procedure to which the caller itself does not have access.

Note that argument copying depends on information, represented as a "copying descriptor", associated with the outward call (see Figure 6). The copying descriptor tells how many arguments there are, how they are to be used (i.e., whether or not they are to be written into), and what their lengths are (so that they can be correctly copied). We will discuss the question of arguments which are to be written into by the called procedure in the following section.




Figure 6. Example of a Procedure which makes two Outward Calls Showing Copying Descriptors

4. INWARD RETURNS

Inward returns are detected when the artificial ring changing fault occurs (see the discussion of outward returns in Section 2) and the return stack entry indicates an inward return. The return pointer in the return stack entry may be compared with a validation return pointer in the returner's workspace in order to avoid erroneous ring changing returns.

The arguments which the outer ring procedure may have written in, as identified in the copying descriptor, are then copied from the returner's workspace into the locations specified for them in the caller's (original) argument-list. Validation of these addresses is only necessary if the copying is done in a ring more privileged than that being returned to, e.g., by the ring changing fault handler which executes in ring 0.

Once these arguments have been copied, the ring changing mechanism re-establishes the environment of the calling procedure and returns to it.

5. REVIEW AND DISCUSSION

Detection. A ring changing transfer is detected when the ring changing mechanism is invoked in response to a suitable fault. The ring bracket mechanism (i.e., a mechanism respecting the rules set forth in Chapter 2, Section 2) will produce such a fault in the case of inward and outward calls; such calls are in fact so defined. Ring changing returns, though, are defined as returns from ring changing calls and the ring bracket mechanism cannot be depended on to detect these returns. The strategy of the "artificial ring changing return fault" was introduced (see Section 2) to guarantee that these returns would always be detected.

Transfer Address Validation. The basic fact about a ringed access control system is that a process' power depends on the ring in which it executes. This is meaningful only because of the rules which govern inter-ring transfers. The basic rule is that outward (power decreasing) transfers may be made at the procedure's discretion whereas inward (power increasing) transfers may be made only with "the permission of the ring to be entered". Transfer address validation, which consists of making sure that the target address is an address at which the target ring will permit entry, thus applies only to inward transfers.

In the case of an inward call, the target address is validated by finding it on the target procedure's gate-list, that is, finding it to be the address of a gate entry point. In the case of an inward return, the target address (which is obtained from the return stack) is validated implicitly by virtue of the fact that it was earlier supplied to the ring changing mechanism by the outward calling procedure, the very procedure being returned to.

The Return Stack. The return stack was introduced (see Section 2) as the data base in which the ring changing mechanism stores the ring number and return address of a caller so that the ring changing return mechanism can subsequently validate the return. The return stack must thus be accessed by the ring changing mechanism upon every ring changing call and return, being "pushed" at each such call and "popped" at each such return. To the extent that the ring changing mechanism may profit from storing other information from the time of a call to the time of the corresponding return, the return stack is evidently the "right" data base to use. Without going into detail we suggested, for example, that the return stack was a good place to record information needed for the restoration of the calling procedure's environment.

Argument Validation. Whenever an inner ring procedure accesses memory via an address obtained from an outer ring source, there is the danger that the supplier of the address is "using" the more privileged procedure to "get around" access control restrictions. Addresses obtained from outer rings are, therefore, suspect and must be used with discretion.

The most outstanding examples of suspect addresses are the addresses of arguments associated with inward calls. "Argument validation" is a technique by which the ring changing mechanism, acting on behalf of the class of gate entry points, does a standard and generally sufficient job of checking these addresses.

Argument validation is not only used in the case of inward calls but in the case of those outward calls where arguments are copied as well. When the arguments for an outward call are copied into the workspace of the called procedure and later, when the return arguments are copied back into the workspace of the calling procedure, the copier of these arguments, being part of the ring-0 ring changing mechanism, obtains its arguments from the rings of the calling and called procedures and must validate these addresses.

Although argument validation doesn't handle all cases of "suspect" addresses, the existence of argument validation does have the useful effect of isolating the cases which aren't covered, making life easier for the programmer of a gate procedure. For, if he can make sure that all of the suspect addresses to be used by the gate procedure and its dynamic descendants are the addresses of arguments, he may be assured that he has written a proper gate procedure. And if there are a few other addresses requiring checking, he can handle them on a case by case basis.

6. OUTLINE OF RING CHANGING SOFTWARE

A. Inward Calls

1. Check that the specified address is a gate entry point.
2. Store information in the "return stack" specifying the caller's environment, including caller's ring number and the return address specified by caller.
3. Determine the ring (NEW-RING) to be entered; that is the value r2 from the called procedure's ring brackets.
4. Create an environment for the called procedure in NEW-RING.
5. Copy the addresses of the arguments into the environment of the called procedure and perform "argument validation".
6. Associate a ring-changing-return fault with the normal return from the called procedure.
7. Set the ring register to NEW-RING.
8. Perform the call.

B. Outward Returns

1. Check that this return corresponds to the last entry in the "return stack".
2. Clean up the environment of the returning procedure (undo A-4).
3. Determine the ring to be returned to, OLD-RING, from the "return stack".
4. Restore the caller's environment, as specified in the "return stack".
5. Set the ring register to OLD-RING.
6. Return to the caller at the address specified in the return stack.

C. Outward Calls

1. Store information in the "return stack" describing the caller's environment.
2. Determine the ring, NEW-RING, to be entered; this is the value r1 from the called procedure's ring brackets.
3. Create an environment for the called procedure in NEW-RING.
4. Copy the caller's arguments into the new environment and create an argument-list pointing to the copied values, also in the new environment.
5. Associate a ring-changing-return fault with the normal return from the called procedure.
6. Set the ring register to the value NEW-RING.
7. Perform the call.

D. Inward Returns

1. Check that this return corresponds to the last entry in the "return stack".
2. Determine the ring, OLD-RING, to be returned to from the "return stack".
3. Copy the return arguments back into the caller's environment (in OLD-RING).
4. Clean up the returning procedure's environment.
5. Restore the caller's environment, as specified in the "return stack".
6. Set the ring register to OLD-RING.
7. Return to caller at the address specified in the return stack.

Chapter 4

SIMULATION OF RINGS USING THE 645

The 645 differs from the "ringed hardware" described in Chapter 2 in several respects which, taken together, add up to the fact that the 645 is a 2-ring rather than a 64-ring machine. In this chapter we shall discuss the relevant aspects of the 645 hardware and show how the ringed access control strategy described in Chapters 2 and 3 can be simulated on the 645.

1. FEATURES OF THE 645 NEEDED FOR THE SIMULATION

1.1 The 645 does not have a "ring register" but does have two states, called master mode and slave mode. The processor has greater power when in master mode than when in slave mode; in particular, (a) certain instructions can only be executed when the processor is in master mode and (b) the access control field of the 645's SDW permits the specification, in addition to the access-mode, of a limiting descriptor - "accessible in master mode only."

1.2 The access control field of the 645's SDW contains no information about rings; in particular it does not contain ring brackets. It does, however, contain either:

- a. access-mode information possibly including either of the two descriptors:
 - accessible in master mode only
 - master mode procedure
- b. the specification of one of eight special "directed" faults (traps) which is to occur whenever the SDW is accessed.

The processor is only "in master mode" when executing a procedure whose SDW indicates a "master mode procedure." The processor may enter master mode while executing a slave mode procedure by:

- faulting
- taking an interrupt

There is another way of switching from slave mode to master mode; it will be discussed later since it invokes a hardware feature that is not needed to simulate a ringed machine.

1.3 The 645 processor's access control machinery interprets the SDW during the addressing cycle and causes an appropriate action to occur depending on the SDW and (usually) on the attempted access, as follows:

- a. If the SDW implies a particular "directed fault", then that fault occurs.
- b. Otherwise, if the SDW does not permit the attempted access, the appropriate access violation fault occurs.
- c. Otherwise, the SDW permits the attempted access and the access is performed.

When a fault occurs, the 645 enters master mode and transfers control to the appropriate master mode fault handling procedure.

1.4 Among the instructions which are "master mode only" are those which access the processor's DBR (the Descriptor Base Register, which contains the absolute address of the descriptor segment currently in effect) and all I/O connect instructions.

2. SIMULATING THE "RINGED HARDWARE" ON THE 645

The technique of simulating the "ringed hardware" on the 645 can practically be deduced from the requirements of that simulation:

1. It must be possible to simulate being in a given ring.
2. It must be possible to simulate changing from one ring to another.

To simulate being in a ring, it must be possible to set up a 645 descriptor segment to define the same set of potential actions in response to potential attempted accesses as is defined by any given "ringed descriptor segment, ring register" pair. The potential actions will be the same in the 645 as on the "ringed hardware" if (a) permitted accesses are performed by the 645 without causing a fault and (b) if accesses which would cause a fault on the "ringed hardware" cause a fault on the 645.

To simulate changing from one ring to another it is obviously necessary to be able to change the 645 descriptor segment. This may be done in two ways. If space is felt to be at a premium, the 645's master mode fault handlers may "change rings" by over-writing the existing descriptor segment with the values appropriate to the other ring. On the other hand, if processing time is felt to be more important than space, the fault handler in master mode may "change rings" by altering the DBR to point to that descriptor segment (waiting in the wings, so to speak) which corresponds to the ring being entered. This second technique, used in Multics, requires one descriptor segment per-ring for each process. The per-ring descriptor segment thus becomes part of the "environment" which pertains to each ring, and switching descriptor segments becomes part of the job of the ring changing mechanism.

3. AN ADDITIONAL FEATURE OF THE 645

The 645 processor has the ability of switching from slave mode to master mode without invoking the trap mechanism, as follows:

A slave mode procedure can transfer to a master mode procedure M provided that:

- a) the segment descriptor of M contains the "accessible in slave mode" attribute, and
- b) the transfer be directed to location zero of M.

This technique for increasing a process' power differs from ring changing in the sense that no fault is generated. However, the philosophy remains the same. By checking that conditions (a) and (b) are true, the hardware performs the "gate validation". The fact that the transfer is guaranteed to be into location zero permits one to code explicitly any type of subsequent validation in the procedure M and to guarantee that the validation code will be executed. (The only system responsibility is to make sure that the transfer is directed to a gate; the gate procedure must take care of the rest.) This feature is used in Multics as explained below.

4. MASTER MODE AND SLAVE MODE IN RING ZERO

Master mode is the most powerful state of the 645 processor; ring zero is the most powerful state of the ringed processor simulated on the 645. It should follow that executing in ring zero means executing in master mode, and it would so follow, were it not for the 645 feature discussed in Section 3 above. That feature is used in order to permit the ring zero supervisor to execute partly in master mode and partly in slave mode, easily switching from one mode to the other.

The Multics ring zero can be regarded as being itself composed of two concentric rings. The more powerful or "master ring 0" contains all master procedures and also all data accessible only in master mode. The less powerful or "slave ring 0" contains all slave procedures and also all data accessible in slave mode. Going from slave ring 0 to master ring 0 can be done through gates provided by master ring 0; these gates are in fact master procedures accessible in slave mode, with the entry point at location zero of the segment. This technique permits efficient switching between slave and master mode in the supervisor and this is the motivation for this additional hardware feature in the GE 645.

Two questions are raised by this discussion. First, why don't "ring zero" and "master mode" coincide? And second, why isn't the special mechanism for entering master mode more generally used?

The supervisor should use master mode only for jobs requiring its special power. To use it for other purposes would increase the chance of disastrous errors due to hardware and software bugs since, for example, all I/O connect instructions are executable in master mode only. But since the supervisor must use master mode fairly frequently, it is desirable that the supervisor have a way of entering master mode which involves just enough validation to prevent accidental entry. Thus the special mechanism. And thus the restriction of the use of the special mechanism to "slave ring zero."

BIBLIOGRAPHY

1. Bensoussan, A. et al. The Multics Virtual Memory T.I.S. R69LSD3.
2. Evans, D.C. and Leclerc, J.Y. Address Mapping and the Control of Access in an Interactive Computer SJCC 1967
3. Graham, R.M. Protection in an Information Processing Utility, CACM (May 1968)
4. Organick, E.I. A Guide to Multics for Sub-system Writers Chapter IV, Access Control and Protection Multics Repository Document MO106.

C. Series 6000 Features for the Multics Virtual Memory

CONTENTS

	<u>Page</u>
Introduction	167
Hardware Compatibility	167
Modifications of Paging and Segmentation Hardware	167
Hardware Implementation of Multics Ring Protection	168
Instructions for String Manipulation and Decimal Arithmetic	169
Segmentation and Paging in New Processor	169
Segment Descriptor Word	169
Page Table Word	171
Modifying the Page Size	173
Absolute Address Formation	173
Descriptor Segment Base Register	174
Rings and Ring Brackets	174
Read/Write Bracket (Rings 0-R1)	176
Read/Execute Bracket (Rings R1-R2)	176
Call Bracket (Rings (R2 +1) - R3)	177
Summary	177
Processor Address Registers	177
Instruction Pointer Register	178
Temporary Pointer Register	178
Eight Pointer Registers	179
Access Control Mechanism	181
CALL Instruction	182
Associative Memory	190

ILLUSTRATIONS

Figure 1. Appending Cycle	175
Figure 2. Instruction Fetch and Initial Address Calculation	183
Figure 3. Indirect Addressing	184
Figure 4. Access Checks for Nontransfer Instructions	185
Figure 5. Access Checks for Transfer Instructions, Except CALLn	186
Figure 6. Access Checks for CALL Instruction	187
Figure 7. Execution of CALL Instruction	189

INTRODUCTION

Experience to date with the Model 645 hardware and the Multics software has uncovered many areas in which system performance and maintainability could be substantially improved by certain modifications of the 645 specifications. Four areas have been investigated and shown to be of major importance to the performance of Multics. The Multics extensions of Series 6000 processors both allow execution of Multics on these processors and provide considerable improvements in system performance. These features include:

1. Hardware aids for improved compatibility with other product line software.
2. Refinements of the paging and segmentation hardware to improve the performance of the system software.
3. Implementation of Multics ring protection mechanism entirely in hardware to improve system performance and reduce software complexity.
4. Addition of instructions for string manipulation and decimal arithmetic.

HARDWARE COMPATIBILITY

The issue of compatibility with product line software is really two issues stemming from two distinctly different motivations. One issue is that of "stand-alone compatibility," which allows the running of standard product line software (e.g., GCOS, T and D monitor) on a stand-alone machine. The other issue is that of "slave program compatibility," which allows for the efficient execution of Series 600/6000 slave programs under the control of the Multics system. A compatibility switch on the processor is used to handle the problem of stand-alone compatibility, while a program-settable mode is used as a software aid in handling the problem of slave program compatibility. The issues of stand-alone and slave program compatibility are treated separately in the following discussion.

MODIFICATIONS OF PAGING AND SEGMENTATION HARDWARE

This paper describes a number of modifications of the 645 paging and segmentation hardware which improve the performance of the Multics software and simplify some areas now felt to have been overdesigned in the 645. The changes in 645 specifications are summarized below (and are described in detail later in the paper):

1. The address field in the segment descriptor word is extended to a full 24-bit absolute address to allow page tables (and unpagged segments) to begin on any legal memory address. This modification allows the software a great deal more flexibility in the allocation of page table space and greatly reduces the amount of wired-down core storage reserved for page tables.

2. The processor supports only a single page size rather than the two page sizes supported by the 645. Provision is made to allow the page size to be modified by field engineering in an orderly and well understood fashion.
3. Each of the eight pointer registers of the processor is extended to contain both a segment number and a word number portion. The 645 concept of internal and external base registers and control fields is dropped. Each of the eight pointer registers on the processor behaves as a 645 base register pair. In addition, each of the new pointer registers contains a bit offset field for use by new instructions for string manipulation and decimal arithmetic.
4. Some minor changes have been made in the definition of the 645 master and slave modes, which are renamed respectively as the privileged and unprivileged modes to avoid confusion with existing terminology. A minor change has also been made in the treatment of execute-only segments, to allow entry at locations other than zero.
5. The access control information contained in the 645 page table word (not used in Multics) has been removed. In the new processor, all access control is implemented in the segment descriptor word.

HARDWARE IMPLEMENTATION OF MULTICS RING PROTECTION

The Multics concept of protection rings, ring crossing, and argument validation has been implemented as an integral part of the paging and segmentation hardware on the processor. The hardware implementation of rings is really a further modification of the 645 paging and segmentation hardware. However, the modification is introduced separately at this point since it involves perhaps the most significant deviation from the 645 specification and, as such, deserves somewhat more motivation. In the current version of Multics running on the 645, the ring protection mechanism is, of necessity, completely simulated by the Multics software. The current system maintains, in parallel, separate descriptor segments for each ring of each process. The ring crossing is simulated by a rather costly and complex fault processing mechanism which includes the copying and validation of argument pointers and the switching of descriptor segments to simulate the effect of switching rings. The cost of the current simulation amounts to approximately 10-20 percent of the useful chargeable CPU time and contributes substantially to the overall complexity of the system. In the new processor, ring crossing and argument validation are handled directly by the hardware without costly software intervention. As a result, a call to an inner ring will require no more CPU time than a call to a procedure in the same ring.

INSTRUCTIONS FOR STRING MANIPULATION AND DECIMAL ARITHMETIC


Extension of the 645 instruction set to include instructions for string manipulation and decimal arithmetic allows considerable simplification of both supervisor and user programs. The Series 6000 Extended Instruction Set (EIS) provides commands to directly process bytes, BCD characters, packed decimal data, and strings. The supervisor will take full advantage of the savings allowed by these new instructions. The language compilers will also take advantage of these space and time saving instructions.

SEGMENTATION AND PAGING IN NEW PROCESSOR

This section describes in detail the segmentation and paging hardware for the new processor. In most respects, the mechanism is quite similar to the 645 appending hardware, with the addition of some refinements to improve the performance of the system software. The single substantial deviation from the 645 specification is the addition of hardware to implement the Multics ring protection mechanism.

Segment Descriptor Word

In order to accommodate the hardware-implemented ring crossing and argument validation and other changes, the Segment Descriptor Word (SDW) has been extended to a 72-bit double precision word to be interpreted as described below.



- ADDR (0-23) - Is a full 24-bit absolute address and specifies the core address of either a page table (for a paged segment) or the first location of an unpagged segment.
- R1 (24-26) Specifies the highest ring number of the read/write bracket for this segment (0-R1)¹.
- R2 (27-29) Specifies the highest ring number of the read/execute bracket of this segment (R1-R2)¹.
- R3 (30-32) Specifies the highest ring number of the call bracket of this segment ((R2 + 1) - R3)¹.

¹ See following section on Rings and Ring Brackets.

- F (33) Is a directed fault indicator and if off (=0) specifies that the processor is to execute the directed fault specified in the FC field (see below).
- FC (34-35) Indicates (if F is off) which of the four directed faults (DF0-DF3) the processor is to execute.
- BOUND (1-14) Is the boundary field and indicates the highest 16-word block of the segment which can be addressed without causing an out-of-bounds fault. If the high order 14 bits of an address to this segment is greater than the value of the boundary field, an out-of-bounds fault is generated. A boundary field of 14 bits is chosen because some instructions (e.g., the new version of STB) reference up to 16 contiguous words. (The boundary field could be maintained to the nearest word, but special checks would have to be made for instructions which reference two or more contiguous words.) A further implication is that the software is expected to allocate unpagged segments in a zero mode 16-word boundary.
- R (15) Is the read-permit indicator. Data fetches by other segments are permitted to this segment only if this indicator is on (=1) and if the processor is executing in a ring less than or equal to R2 (i.e., within the read/write or read/execute bracket).¹
- E (16) Is the execute-permit indicator. Instruction fetches from this segment are permitted only if this indicator is on (=1) and if the processor is executing in a ring greater than or equal to R1 and less than or equal to R2 (i.e., within the read/execute bracket; see below). Note that when the E indicator is on and the R indicator is off, the segment is to be treated as an "execute-only" procedure segment. An execute-only procedure segment is permitted to reference data within itself (i.e., within the same segment) in spite of the lack of the read indicator. However, read permission is denied to other segments.
- W (17) Is the write-permit indicator. Attempts to store into this segment are honored only if this indicator is on (=1) and if the processor is executing in a ring less than or equal to R1 (i.e., within the read/write bracket; see below).

¹See following section on Rings and Ring Brackets.

- P (18) Is the privileged mode indicator. If this indicator is on (=1) and if the processor is executing in ring 0, the procedure segment is permitted to execute privileged instructions and inhibit interrupts under control of bit 28. Privileged procedures need no further powers and are subject to all other access checking (read, write permission bounds checking, etc.). Since privileged procedures can be executed only in ring 0, it is no longer necessary to limit calls to privileged procedures to enter via word 0 of the segment.
- U (19) Indicates whether the segment is paged (=0) or unpaged (=1). If the segment is unpaged, ADDR is the full absolute address of the first word (word 0) of the segment. If the segment is paged, ADDR is the full absolute address of the beginning of the page table for the segment.
- G (20) Is the gate indicator and if off (=0) any call to this segment from a different segment must be directed to an address value less than the value of the CL field (see below).
- CL (22-35) Is the call limiter. If G is on, any external transfer to this segment via the new CALL instruction (described below) must be directed to a word number less than the value of this field.

Page Table Word

The format of the page table word (PTW) has been somewhat simplified from the 645 version in that no access control is performed at the PTW level. The PTW format is described below.



- ADDR (0-17) Is the high order 18 bits of the 24-bit absolute address of the first word of the page. The hardware assumes that all pages begin on addresses which are zero modulo the page size. For example, if the page size is set to 1024 words, the hardware assumes that each page begins on a zero modulo 1024 address and that the low order 10 bits of the 24-bit absolute address are zero.

- S (18-23) Is reserved for the use of the system software for the maintenance of page status information and is never modified (or used) by the hardware.
- U (26) Indicates whether (=1) or not (=0) the page has been used since the last time this bit was interrogated (and reset) by the system software. Whenever this bit is zero and the processor addresses any word within the page (corresponding to this PTW), the processor sets this bit to 1 using a 3-bit store-by-zone command for bits 24-26. The store-by-zone is used to avoid a race condition with another processor attempting to set the "modified" bit (see below) for the same page. This technique is necessitated by the lack of a read-alter-rewrite command in the memory controller. A further implication of the lack of read-alter-rewrite is that the software must reset this bit via a store to the third 9-bit field (character) in the PTW in order not to disturb the modified bit. Note that any usage of the page between the time the used bit is read by the software and then reset (if on) is not noticed by the software. Since the used bit is used only for maintaining the core-usage statistics, this race between hardware and software has no effect, inasmuch as (1) if the page is heavily used (i.e., needed in core) it will be used again turning the used bit back on, and (2) the software does not reset the bit if it is already zero.
- M (29) Indicates whether (=1) or not (=0) the page has been modified since the last time this bit was interrogated (and reset) by the system software. Whenever this bit is zero and the processor modifies any word within the page, the processor sets this bit and the usage bit to 1 using a 6-bit store-by-zone command for bits 24-29. The software uses this bit to determine whether or not the contents of a page must be written on secondary storage before the core is released for other usage. As a result, the software is expected to store a directed fault in the PTW and clear the associative memories of all processors before the modified bit may be safely tested and then reset. In addition, the directed fault must be stored using a store to the sixth 6-bit field (character) of the PTW to compensate for the lack of read-alter-rewrite.
- F (33) Is the directed fault indicator and if off (=0) indicates that the directed fault indicated in the FC field is to be executed by the processor.

FC (34-35) Indicates (if F is off) which of the four directed faults (DFO-DF3) is to be executed by the processor.

Modifying the Page Size

As indicated by the format of the SDW, the new processor supports only a single page size (the 645 allows two page sizes). However, it is extremely desirable to have the ability to change the page size in order to allow system optimization with respect to core "breakage" and storage device access times. For example, replacing the current highspeed drum with a bulk core would most likely give even better performance with a page size smaller than 1024 words.

Since a decision to change the page size is not a casual one and should not be made very often, the page size is changeable by field modification to any power of 2 from 64 words to 4096 words.

Absolute Address Formation



For each memory reference we assume the program presents the processor with the following address:



SEGNO (15 bits) Specifies the desired segment (i.e., index into the descriptor segment). The segment number is constrained to 15 bits (rather than 18) by considerations described later in this document.

WORDNO (18 bits) Specifies the desired word address (18 bits) within the specified segment.

We also assume (for discussion only) that the processor has the following two internal registers:



PN (12 bits) Is used to hold the page number (i.e., index into page table) when forming an absolute address within a paged segment.

PO (12 bits) Is used to hold the offset within the page when forming an absolute address within a paged segment.

PN is initialized with the high-order portion of WORDNO to obtain an index relative to the base of the page table of the segment. PO is initialized with the remaining portion of WORDNO and is augmented by the ADDR field of

the PTW to form the absolute address. The "break" in WORDNO is determined by the current page size. For example, if the page size is 1024 words (initial setting), the high order 8 bits of WORDNO are used to initialize PN and the low order 10 bits of WORDNO are used to initialize PO.

Figure 1 summarizes the major steps necessary to transform a program generated address (SEGNO/WORDNO) into an absolute address (ABSADDR). In most respects, the address formation is simpler than the 645 mechanism, in that there is only one page size to consider and that no access control is specified at the PTW level.

Note that Figure 1 and all the flow charts in this paper make use of the following PL/1 notations:

1. A.B is used to denote the quantity B contained in A. For example, PTW.ADDR denotes the ADDR field within the PTW.
2. The double vertical bar (||) is used to denote concatenation (e.g., PTW.ADDR || 000000).
3. The single vertical bar (|) is used to denote the logical inclusive OR.

Descriptor Segment Base Register

The Descriptor Segment Base Register (DSBR) is an internal processor register used to locate the current descriptor segment. In the new processor, the DSBR has been extended to 51 bits to accommodate the longer address and bound fields and to contain a stack offset. The DSBR is loaded from and stored into a doubleword having the same format as a Segment Descriptor Word (SDW) with the exception that unused fields are ignored during loading of the DSBR and are set to zero when the DSBR is stored. Only the following four SDW fields have meaning when loaded into a DSBR.

1. ADDR (24 bits)
2. BOUND (14 bits)
3. U (page/ unpage switch; 1 bit)
4. STACK (12 bits)

The STACK field specifies the upper 12 bits of the 15-bit stack segment number. This register is used only during the execution of a CALL instruction.

RINGS AND RING BRACKETS

A Multics process consists of procedure and data segments which are all directly addressable through the descriptor segment of that process. However, a process may access a segment only when the process is running at an appropriate level of privilege. For example, all the segments of the




Figure 1. Appending Cycle

hardcore supervisor are shared and accessible to all Multics processes but only when executing at the highest level of privilege.

The Multics system allows segments to be grouped into an ordered set of collections called rings in which segments requiring the highest level of privilege to reference can be accessed only from within the innermost ring of the set. Each ring is identified with a ring number designating the required level of privilege necessary to access segments in that ring. In Multics, the ring with the highest privilege is ring 0, which contains the procedures and data bases of the hardcore supervisor. Each user process has at least two rings, one for the hardcore supervisor and one for user programs and data. The user process may generate more rings (levels of lesser privilege) if desired.

Frequently, it is useful to allow a segment to be accessible in more than one ring. For example, it is often useful for a data base which is writeable in an inner ring to be readable in an outer ring. For this reason, the concept of ring brackets was introduced.

The access of a user to a specific segment is controlled by two quantities: the access attributes (e.g., read, execute, write) and the ring brackets. The ring brackets of a segment are specified by three integers (R1, R2, and R3) each of which must be greater than or equal to the preceding number. The first number (R1) specifies the top (highest ring number) of the read/write bracket, the second number (R2) specifies the top of the read/execute bracket, and the last number (R3) specifies the top of the call bracket.

Read/Write Bracket (Rings 0-R1)

Attempts to read or write a segment by a procedure executing in a ring within the read/write bracket are allowed if the appropriate (read or write) access indicators are on for the segment being referenced. Execution of a procedure in a ring within this bracket is permitted only at the top of the read/write bracket (R1), which is also the bottom of the read/execute bracket. Note that the highest ring from which a segment can be written is specified by R1. As a result, the data in the segment is no more reliable than the procedure segments which operate in that ring.

Read/Execute Bracket (Rings R1-R2)

Attempts to read or execute (transfer to) a segment by a procedure executing in a ring within the read/execute bracket are allowed if the appropriate (read and execute) indicators are on for the segment being referenced. Writing of a segment within its read/execute bracket is permitted only from the ring at the bottom of the bracket (R1), which is also the top of the read/write bracket. If R2 is equal to R1, the read/execute bracket specifies a single ring (R1).

Call Bracket (Rings (R2 + 1) - R3)

Attempts to call a (procedure) segment from a segment executing in a ring above the read/execute bracket but within the call bracket of the procedure to be called are allowed if the execute indicator of the procedure to be called is on and if the new CALL instruction (described below) is used. When the CALL instruction is directed to a procedure in an inner ring which has the appropriate execute access and call bracket, the processor automatically switches to the ring specified as the R2 of the procedure being called. The call bracket and the CALL instruction are the only means (except for faults) by which control can be passed from an outer ring to an inner (more privileged) ring. If R3 is equal to R2, the call bracket is null and the procedure cannot be called from an outer ring.

Summary

Assuming that the appropriate (read, execute, or write) indicators are on, the following list summarizes the effects of the three ring brackets:

1. Writing is permitted from a ring within the read/write brackets only (i.e., if ring \leq R1).
2. Reading is permitted from a ring within the read/write bracket or the read/execute bracket (i.e., ring \leq R2).
3. Execution (or transfer of control) is permitted only from a ring within the read/execute bracket (i.e., R1 \leq ring \leq R2).
4. Calling (via CALL only) is permitted from a ring within the read/execute or call brackets (i.e., R1 \leq ring \leq R3).
5. The CALL instruction is the only instruction which may be used to access a segment in a ring within its call bracket (i.e., R2 < ring \leq R3).
6. No access is permitted to a segment from a ring higher than the call bracket (i.e., ring > R3).

PROCESSOR ADDRESS REGISTERS

Like the 645, the new processor has 10 address or pointer registers (PRs). Eight of these pointer registers can be directly accessed and modified by the software, one is used to locate the current instruction,

and one is used exclusively by the processor for effective address calculations. Unlike the 645, each of the eight program addressable pointer registers specifies a full segmented address including the segment number and the word number in a single pointer register. These registers have also been extended to include a bit number.

Instruction Pointer Register

The instruction pointer register (IPR) is used by the processor to locate the current instruction and may be modified by the software to effect a transfer of control. The IPR is actually an extension of the PBR and IC of the 645. The contents of the 36-bit IPR are outlined below.

PRR	PSR	IC
-----	-----	----

PRR (3 bits) Is the procedure ring register and specifies the ring (level of privilege) in which the processor is currently executing. PRR may be set to a higher value only by an RTCD or RCU instruction. It may be set to a lower value only by a CALL instruction (see below) or by a fault or interrupt.

PSR (15 bits) Is the procedure segment register (same as the PBR in the 645) and specifies the segment number of the current procedure segment.

IC (18 bits) Is the instruction counter (same as in the 645).

Temporary Pointer Register

The temporary pointer register (TPR) is used exclusively by the processor for operand address calculations and serves the same general purpose as the TBR and computed address (CA) of the 645.

TRR	TSR	CA	BITNO
-----	-----	----	-------

TRR (3 bits) Is the temporary ring register and is used to maintain the lowest level of privilege (i.e., highest ring number) encountered during operand address calculation. The TRR is initialized with the value of the PRR field of the IPR at the beginning of each instruction. During the operand address calculation, TRR is used to record the

highest value of SDW.RI (the top of the read/write bracket) of any segment used, in the address calculation. For example, if an indirect address is fetched from segment X, TRR is set to the larger of TRR (its current value) and the RI field in the SDW for segment X. Note that during the operand address calculation, the value of TRR may get larger but never smaller.

- TSR (15 bits) Is the temporary segment register (same as the TBR in the 645) and is initialized with the value of the PSR field of the IPR at the beginning of each instruction. During operand address calculation, TSR contains the segment number portion of the current address calculation.
- CA (18 bits) Is the computed address and serves the same function as the 645 register of the same name. The computed address is initialized at the beginning of each instruction with the contents of the instruction counter of the IPR. During operand address calculation, CA contains the word number portion of the current address calculation.
- BITNO (6 bits) Is a bit-offset relative to the first bit in the word specified by CA. This field is ignored by all instructions except the new instructions specifically designed for string manipulation or decimal arithmetic.

Once an operand address calculation is complete, the value of TRR is compared with the ring brackets of the segment containing the operand address to determine whether the operation is to be allowed. For example, if the instruction intends to store into this operand, the value of TRR must be less than or equal to the RI (in the SDW) of the segment to be modified.


Eight Pointer Registers

The new processor contains eight program accessible pointer registers, which replace the eight address base registers (ABRs) of the 645. The PRs of the new processor differ from the ABRs in that each PR contains both a segment number and a word number portion. In effect, each PR of the new processor behaves as a 645 base register pair. The contents of each 42-bit PR are outlined below.

RN	SEGNO	WORDNO	BITNO
----	-------	--------	-------

- RN (3 bits) Is used by the software to specify the level of privilege (i.e., ring number) at which the processor is to treat the address contained in the address register. When the processor uses the contents of a PR for address modification (e.g., bit 29 in the instruction word is on) the value of TRR is set to the larger of TRR (its current value) and the RN field of the specified PR. The use of the RN field of a PR allows the software to save the TRR of an operand address calculation — e.g., through the use of an EAP (effective address to pointer) instruction.
- SEGNO (15 bits) Specifies the segment number portion of the segmented address.
- WORDNO (18 bits) Specifies the word number portion of the segmented address.
- BITNO (6 bits) Specifies a bit-offset relative to WORDNO and is ignored by all instructions except those designed specifically for string manipulation or decimal arithmetic.

The software may store the contents of a PR into an ITS (indirect to segment) word pair with the use of an STP (store pointer) instruction. The software may then address indirectly through the ITS indirect word rather than using the original PR. Alternatively, the software may reload another PR from the ITS word pair through the use of the EAP instruction. In either case, it is necessary to save the value of the RN of the PR in the ITS word pair so that the privilege level of the original operand address calculation is not lost. As a result, the ITS word pair is modified to include a ring number field as outlined below.



- SEGNO (3-17) Is the segment number field (as in the 645). Note that bits 0-2 of the ITS word pair are set to zero for compatibility with 645 programs expecting an 18-bit segment number in the upper half of the first word.
- RN (18-20) Is set to the value of the RN field of the PR during the STP instruction. If the processor attempts to indirect through an ITS word pair, TRR is set to the larger of TRR, RN (of the ITS), and R1 of the segment containing the ITS. Note that an improper value of RN in an ITS word pair has no ill effect, since the processor always takes

the maximum of TRR and RN. In other words, it is impossible for an ITS word pair to specify a higher privilege than the segment in which it resides.

- ITS (30-35) Specifies the modifier code (octal 43) for the ITS modifier (same as in the 645).
- WORDNO (0-17) Is the word number portion of the saved PR.
- BITNO (21-26) Is the bit-offset of the PR saved by the STP instruction. The strange placement of the BITNO field is necessary to remain compatible with the current PL/1 software implementation of bit-offsets.
- MOD (30-35) Is set to zero by the STP instruction but may be set by the software to specify further address modification (same as in the 645).

Since most Multics compilers (notably PL/1) calculate addresses via an EAP instruction, it can be expected that compiler generated code can take full advantage of the hardware protection mechanism with little modification. If all addresses of all input parameters are calculated and saved (for use as outgoing argument pointers) via the use of the EAP and STP instructions, it will be possible for a procedure operating in ring 1 to pass to ring 0 a parameter given to the procedure from ring 2, without checking the address of the parameter. The access checking is fully automatic as long as the TRR of the original address calculation continues to be maintained and passed along as the RN field of a PR or ITS word pair.

The STCD (store control double) instruction is modified to store the PRR in the same manner as STP stores the RN field of a PR. The PRR is stored by the STCD to allow an RTCD (return control double) instruction to return to the proper ring.

ACCESS CONTROL MECHANISM

Figures 2 through 6 attempt to flow chart the entire access control mechanism from the instruction fetch up to actual execution of the instruction. In order to concentrate on the access control mechanism, many details have been left out of the flow charts (indexing, IT modifiers, etc.). If all the access control checks are successfully met, control will end up in a circle marked "done." The contents of the flow charts are summarized below.

Figure 2 begins with the instruction fetch and continues through the initial address calculation.

Figure 3 shows how indirect addressing affects the access computation. (The notation "Rl (ITS)" is used to denote the Rl of the segment containing the ITS word pair.)

Figure 4 shows the access checks made for all instructions except for transfer of control.

Figure 5 shows the access checks performed on all transfer instructions with the exception of the CALL instruction. Note that the PRR cannot be changed by a normal transfer instruction (even to a higher value). However, it is possible to set the PRR to a higher value with the modified RTCD (described below).

Figure 6 shows the access checks performed by the CALL instruction (the only slave instruction permitted to set the PRR to a lower value).

CALL INSTRUCTION

The CALL instruction is provided as the only means by which a procedure segment may call a procedure in an inner ring (i.e., set PRR to a lower value). The CALL instruction is to be used in all standard interprocedure calls and is intended to replace the transfer instruction as the last instruction of the standard Multics calling sequence.

The CALL instruction uses two PRs: PR_n and PR_{n+1}, where n is even. The value of n is wired into the processor and is currently 6. It is possible for a field engineer to change this value to 4, 2, or 0 by an orderly procedure. It must not be possible, however, to change this value under user program control. (This use of a pair of PRs involves two full PRs (RN, SEGNO, WORDNO, and BITNO) and should not be confused with a 645 base register pair.) When the CALL instruction is used to transfer control to another ring, the assumption is made (by convention) that the stack segment of the target ring has a segment number equal to the ring number of the target ring (i.e., the stack segment for ring X is a segment number X). The CALL instruction behaves as a TRA (transfer) instruction with the following exceptions:

1. The access checking for a CALL instruction allows PRR to be set to a lower value, provided that the call is made from a ring within the call bracket of the target segment (see Figure 6).
2. If an attempt is made to call a procedure in an outer ring (a relatively rare case), an access violation occurs. Because of the necessity of copying all arguments, the standard call, save, and return sequences cannot handle calls to an outer ring without excessive software overhead. Therefore, calls to outer ring procedures will continue to cause a fault to allow the system software to interpret the call.




Figure 2. Instruction Fetch and Initial Address Calculation




Figure 3. Indirect Addressing




Figure 4. Access Checks for Nontransfer Instructions




Figure 5. Access Checks for Transfer Instructions, Except CALLn




Figure 6. Access Checks for CALL Instruction

3. At the beginning of the CALL, the contents of PR_n are assumed to point to a location (i.e., the beginning of the current stack frame) within the stack segment of the calling ring. During the execution of the CALL, the processor sets the contents of PR_{n+1} to point to word 0 of the stack segment of the target ring in one of two ways:
 - a. If control is to remain in the current ring (i.e., TRR = PRR), the SEGNO portion of PR_{n+1} is set to the SEGNO of PR_n and the WORDNO portion of PR_{n+1} is set to zero.
 - b. If control is to be passed to an inner ring (i.e., TRR < PRR), the SEGNO portion of PR_{n+1} is set to the value of the target ring number (TRR) and the WORDNO of PR_{n+1} is set to zero.

If an attempt is made to call to an outer ring (i.e., TRR > PRR), an access violation is generated as indicated in Figure 6.

Figure 7 details the operation of the CALL instruction after the effective address computation has been completed (i.e., TPR has been computed), and TRR is set to the target ring number (see Figure 6).

The software stores a pointer to the end of the current (or last used) stack frame in the beginning of that stack. The standard call and save sequences might then be modified as follows:

Calling Sequence:	ZERO	ARGLIST	ZERO points to ARGLIST
	STCD	6 20	Set return location
	CALL	ENTRYPOINT	Call external procedure
Save Sequence:	EAP1	7 NEXT,*	Load pointer with base of new stack frame
	STP6	1 16	Save pointer to old frame
	EAP6	1 0	Switch to new frame
	EAP1	6 TEMP	Compute pointer to next frame (allocate new frame)
	STP1	6 18	Save pointer to next frame
	STP1	7 NEXT	Update stack base
	STP0	6 26	Save ARGLIST PR




Figure 7. Execution of CALL Instruction

ASSOCIATIVE MEMORY

as in the 645, the new processor requires a small associative memory in order to avoid memory fetches of frequently used SDWs and PTWs. A series of measurements and experiments has determined the effectiveness and behavior of the 645 associative memory. The experiments were conducted and measurements taken during normal Multics operation, under varying user loads.

The experiments indicate that the current 645 associative memory is quite effective. It appears that the most significant aspect of the associative memory is in the speed of the search, since it is not possible to overlap completely the associative lookup with other work. This aspect suggests that a one-pass lookup would be a desirable objective. There are at least three ways in which the effect of a one-pass lookup can be achieved:

1. One approach is derived from the fact that the "hit rate" on SDWs for paged segments on the 645 is extremely low (about 0.21 percent). This fact suggests a one-pass search of an associative memory containing only PTWs and SDWs for unpagged segments. The search would look for an unpagged SDW for the referenced page within the segment. The copy of the PTW in the associative memory must be extended to include access control information from the original SDW for the segment. This approach has the drawback that any change in the operating environment (e.g., the use of smaller page sizes) which causes SDWs for paged segments to be in higher demand would begin to degrade system performance.
2. Another approach is to achieve the effect of a one-pass search using a two-pass search and overlapping the first pass during address preparation. In this approach, the single associative memory contains both SDWs (paged and unpagged) and PTWs extended with access control information. During address preparation, the associative memory is searched for the SDW of the segment to be referenced. After address preparation, a second pass is made to locate the PTW for the page to be referenced. Only if the second pass fails are the results of the first pass interrogated. If the first pass had succeeded, only the PTW must be fetched from core memory. Otherwise, both the SDW and the PTW must be fetched from core memory.
3. A third approach is to search two associative memories in parallel, one for SDWs and the other for PTWs. If either the SDW or PTW is not found in its respective associative memory, it is retrieved from core memory and updated into the appropriate associative memory. Although this approach requires duplicate circuitry, it is appealing in its logical simplicity and is the method chosen.

D. Abbreviations and Acronyms

ADDR	Address portion of PTW
AST	Active Segment Table
ASTE	Active Segment Table Entry
CA	Computed Address
CM	Core Map
CME	Core Map Entry
DBR	Descriptor Base Register
DC	Directory Control
DCW	Data Control Word
DID	Device Identifier
DIM	Device Interface Module
DS	Descriptor Segment
DSBR	Descriptor Segment Base Register
IC	Instruction Counter
IPR	Instruction Pointer Register
ITP	Indirect to Pointer Register
ITS	Indirect To Segment
KST	Known Segment Table
KSTE	Known Segment Table Entry
MC	Memory Controller
PC	Page Control
PHM	Page Has Been Modified
PHU	Page Has Been Used
PN	Page Number
PO	Page Offset

PR	Pointer Register
PRR	Procedure Ring Register
PSR	Procedure Segment Register
PT	Page Table
PTW	Page Table Word
RA	Ring Alarm register
RN	Ring Number
SC	Segment Control
SDW	Segment Descriptor Word
SFH	Segment Fault Handler
SST	Systems Segment Table
TPR	Temporary Pointer Register
TRR	Temporary Ring Register
TSR	Temporary Segment Register
UID	Unique Identifier

Honeywell

Honeywell Information Systems

In the U.S.A.: 200 Smith Street, MS 486, Waltham, Massachusetts 02154
In Canada: 155 Gordon Baker Road, Willowdale, Ontario M2H 3N7
In the U.K.: Great West Road, Brentford, Middlesex TW8 9DH
In Australia: 124 Walker Street, North Sydney, N.S.W. 2060
In Mexico: Avenida Nuevo Leon 250, Mexico 11, D.F.

33111, 2C1181, Printed in U.S.A.

AG95, Rev. 0