#### Operating Systems "Memory Management"

Mathieu Delalandre University of Tours, Tours city, France <u>mathieu.delalandre@univ-tours.fr</u>

Lecture available at http://mathieu.delalandre.free.fr/teachings/operating2.html

# Operating Systems "Memory Management"

1. Introduction

2. Contiguous memory allocation

- 2.1. Partitioning and placement algorithms
- 2.2. Memory fragmentation and compaction

2.3. Process swapping

- 2.4. Loading, address binding and protection
- 3. Simple paging and segmentation
  - 3.1. Paging, basic method
  - 3.2. Segmentation

#### Introduction (1)

**Memory hierarchy:** memory is a major component in any computer. Ideally, memory should be extremely fast (faster than executing an instruction on CPU), abundantly large and dirt chip. No current technology satisfies all these goals, so a different approach is taken. The memory system is constructed as a hierarchy of layers.

	Access time (4KB)	Capacity
Registers	0.25 - 0.5 ns	< 1 KB
Cache	0.5 - 25 ns	>16 MB
Main memory	80 - 250 ns	> 16 GB
Disk storage	30 µs - plus	> 100 GB

As one goes down in the hierarchy, the following occurs:

- a. decreasing cost per bit,
- b. increasing capacity,
- c. increasing access time,
- d. decreasing frequency of access to the memory by the processor.

i.e. from 130.2 Mb.s<sup>-1</sup> to 15.6 Gb.s<sup>-1</sup>

#### Introduction (2)

Memory hierarchy: memory is a major component in any computer. Ideally, memory should be extremely fast (faster than executing an instruction on CPU), abundantly large and dirt chip. No current technology satisfies all these goals, so a different approach is taken. The memory system is constructed as a hierarchy of layers.

	Access time (4KB)	Capacity
Registers	0.25 - 0.5 ns	< 1 KB
Cache	0.5 - 25 ns	>16 MB
Main memory	80 - 250 ns	> 16 GB
Disk storage	30 µs - plus	> 100 GB

i.e. from 130.2 Mb.s<sup>-1</sup> to 15.6 Gb.s<sup>-1</sup>

The strategy of using a memory hierarchy works in principle, but only if conditions (a) through (d) in the preceding list apply.

e.g. with a two-level memory hierarchy

His the hit ratio, the faction of all memory accesses that are found in the faster memory.

- is the access time to level 1.  $T_{I}$
- is the access time to level 2, with  $T_2 >> T_1$ .  $\frac{T_2}{T}$
- is the average access time, computed as:



#### Introduction (3)

**Memory hierarchy:** memory is a major component in any computer. Ideally, memory should be extremely fast (faster than executing an instruction on CPU), abundantly large and dirt chip. No current technology satisfies all these goals, so a different approach is taken. The memory system is constructed as a hierarchy of layers.

	Access time (4KB)	Capacity
Registers	0.25 - 0.5 ns	< 1 KB
Cache	0.5 - 25 ns	>16 MB
Main memory	80 - 250 ns	> 16 GB
Disk storage	30 µs - plus	> 100 GB

i.e. from 130.2 Mb.s<sup>-1</sup> to 15.6 Gb.s<sup>-1</sup>

The strategy of using a memory hierarchy works in principle, but only if conditions (a) through (d) in the preceding list apply.

e.g. with a two-level memory hierarchy Considering time accesses of 200 to the main memory T<sub>1</sub> and 40 µs to the main plus disk memories (T<sub>1</sub>+T<sub>2</sub>),  $\overline{T} = H \times 200 + (1 - H) \times 4 \times 10^4$ 

For a performance degradation less than 10 percent in main memory,  $220 = H \times 200 + (1 - H) \times 4 \times 10^4$ 

H = 0,999497

then, 1 fault access out of 1990.

additional constraints must be considered, a typical hard disk has:

Latency	3 ms
Seek time	5 ms
Access time	0.05 ms
Total	≈8 ms

#### Introduction (4)

**Memory management:** managing the lowest level of cache memory is normally done by hardware, the focus of memory management is on the programmer's model of main memory and how it can be managed well.

**Memory management without memory abstraction:** the simplest memory management is without abstraction. Main memory is generally divided in two parts, one part for the operating system and one part for the program currently executed. The model of memory presented to the programmer was physical memory, a set of addresses belonging to the user's space.

e.g. three simple ways to organize memory with an operating system and user programs:



When a program executed an instruction like

MOV REGISTER1, 80

the computer just moved the content of physical memory location 80 to REGISTER1.

#### Introduction (5)

**Memory management:** managing the lowest level of cache memory is normally done by hardware, the focus of memory management is on the programmer's model of main memory and how it can be managed well.

**Memory management without memory abstraction:** the simplest memory management is without abstraction. Main memory is generally divided in two parts, one part for the operating system and one part for the program currently executed. The model of memory presented to the programmer was physical memory, a set of addresses belonging to the user's space.

e.g.

programming version of C,D



loadable version of C,D, physical addresses must be directly specified by the programmer in the program itself



#### Introduction (6)

**Memory management:** managing the lowest level of cache memory is normally done by hardware, the focus of memory management is on the programmer's model of main memory and how it can be managed well.

**Memory management with memory abstraction** provides a different view of a memory location depending on the execution context in which the memory access is made. The memory abstractions makes the task of programming much easier, the programmer no longer needs to worry about the memory organization, he can concentrate instead on the problem to be programmed. The memory abstraction covers:

Memory partitioning is interested for managing the available memory into partitions.

contiguous / noncontiguous allocation	assigns a process to consecutive / separated memory blocks.
fixed / dynamic partitioning	manages the available memory into regions with fixed / deformable boundaries.
complete / partial loading	refers to the ability to execute a program that is only fully or partially in memory.

#### Introduction (7)

**Memory management:** managing the lowest level of cache memory is normally done by hardware, the focus of memory management is on the programmer's model of main memory and how it can be managed well.

**Memory management with memory abstraction** provides a different view of a memory location depending on the execution context in which the memory access is made. The memory abstractions makes the task of programming much easier, the programmer no longer needs to worry about the memory organization, he can concentrate instead on the problem to be programmed. The memory abstraction covers:

**Placement algorithms:** when it is time to load a process into main memory, the OS must decide which memory blocks to allocate.

**Fragmentation** / **compaction:** is a phenomenon in which storage space is used inefficiently, reducing capacity or performance and often both. compaction can eliminate, in part, the fragmentation.

**Process swapping** is a strategy to deal with memory overload, it consists in bringing each process in its entirely, running if for a while, then putting it back on the disk.

Address protection determines the range of legal addresses that the process may access and to ensure that this process can access only these legal addresses.

Address binding: the addresses may be represented in a different way between the disk and main memory spaces. Address binding is a mapping from one address space to another.

# Introduction (8)

Methods	Partitioning	Placement	Fragmentation	Swapping	Address binding	Layer
		algorithms	/ compaction		& protection	

Fixed partitioning	contiguous / fixed / complete		yes / no		no / yes (MMU)	OS kernel
Memory management with bitmap	contiguous / dynamic / complete					
Memory management with linked lists	contiguous / dynamic / complete	searching algorithms	yes / yes	yes		
Buddy memory allocation	contiguous / hybrid / complete	urgoriumis .	yes / no			
Simple paging and segmentation	noncontiguous / dynamic / complete		yes / no		yes (TLB)	programs / services

# Operating Systems "Memory Management"

1. Introduction

2. Contiguous memory allocation

- 2.1. Partitioning and placement algorithms
- 2.2. Memory fragmentation and compaction

2.3. Process swapping

- 2.4. Loading, address binding and protection
- 3. Simple paging and segmentation
  - 3.1. Paging, basic method
  - 3.2. Segmentation

Methods	Partitioning	Placement	Fragmentation	Swapping	Address binding	Layer
		algorithms	/ compaction		& protection	

Fixed partitioning	contiguous / fixed / complete	searching algorithms	yes / no	yes	no / yes (MMU)	OS kernel
Memory management with bitmap	contiguous / dynamic / complete		yes / yes			
Memory management with linked lists	contiguous / dynamic / complete					
Buddy memory allocation	contiguous / hybrid / complete		yes / no			
Simple paging and segmentation	noncontiguous / dynamic / complete		yes / no		yes (TLB)	programs / services

# Partitioning and placement algorithms "Fixed partitioning"

**Fixed partitioning:** the simple scheme for managing available memory is to partition into regions with fixed boundaries. There are two alternatives for fixed partitioning, with equal-size or unequal-size partitions.

	Size of partitions	Max loading size	Memory fragmentation	Degree of multi- programming	Placement algorithms
equal-size partitions	М	<m< td=""><td>High</td><td>High</td><td>Single FIFO queue</td></m<>	High	High	Single FIFO queue
unequal-size partitions	[M-N] with M <n< td=""><td><n< td=""><td>Medium</td><td>Medium</td><td>Single / multiple FIFO queue(s)</td></n<></td></n<>	<n< td=""><td>Medium</td><td>Medium</td><td>Single / multiple FIFO queue(s)</td></n<>	Medium	Medium	Single / multiple FIFO queue(s)



13

Methods	Partitioning	Placement	Fragmentation	Swapping	Address binding	Layer
		algorithms	/ compaction		& protection	

Fixed partitioning	contiguous / fixed / complete		yes / no		no / yes (MMU)	OS kernel
Memory management with bitmap	contiguous / dynamic / complete		vos / vos			
Memory management with linked lists	contiguous / dynamic / complete	searching algorithms	yes / yes	yes		
Buddy memory allocation	contiguous / hybrid / complete	uigonninis	yes / no			
Simple paging and segmentation	noncontiguous / dynamic / complete		yes / no		yes (TLB)	programs / services

# Partitioning and placement algorithms "Memory management with bitmaps"

**Memory management with bitmaps:** with a bitmap, memory is divided into allocation units as small as a few words and as large as several kilobytes. Corresponding to each allocation units is a bit in the bitmap, which is 0 if the unit is free and 1 if it is occupied (or vice versa).

e.g. a part of memory with five processes and three holes, with the corresponding bitmap.



Another problem is the placement algorithm. When it has been decided to bring k unit process into memory, we must search k consecutive 0 bits in the map, that is a slow operation.

# Partitioning and placement algorithms "Memory management with linked lists" (1)

**Memory management with linked lists:** another way of keeping track of memory is to maintain a linked list of allocated and free memory segments, where a segment either contains a process or is an empty hole between two processes.

e.g. a part of memory with five processes and three holes, with the corresponding linked list



Here, the segment list is sorted by address. Sorting this way has the advantage that when a process terminates, updating the list is straightforward.

# Partitioning and placement algorithms "Memory management with linked lists" (2)

**Memory management with linked lists:** another way of keeping track of memory is to maintain a linked list of allocated and free memory segments, where a segment either contains a process or is an empty hole between two processes.

A terminating process has two neighbors (except when it is at the very top or bottom of the memory). These may be either processes or holes, leading to four combinations.



Since the process table slot for the terminating process will normally point to the list entry for the process itself, it may be more convenient to have the list as a double-linked list.

# Partitioning and placement algorithms "Memory management with linked lists" (3)

**Memory management with linked lists:** another way of keeping track of memory is to maintain a linked list of allocated and free memory segments, where a segment either contains a process or is an empty hole between two processes.

Several algorithms can be used to allocate memory for a created process.

Algorithms	5 Descriptions
------------	----------------

first-fit	It scans along the list of segments from beginning until it finds a hole that is big enough.
next-fit	It works the same way as first fit, except that it keeps track of where it is whenever it finds a suitable hole.
best-fit	It searches the entire list, from beginning to end, and takes the smallest hole that is adequate rather than breaking up a big hole that might be needed later.

worst-fit	To get around the problem, one could think about worst fit, that is, always take the largest available hole.
quick-fit	It maintains separate lists for some of the most common sizes requested, to speed up the best-fist.

Statistical analysis reveals that the first-fit algorithm is not only the simplest one but usually the best and fastest as well. To support the next-fit, circularly linked list with a tail pointer must be modeled.

## Partitioning and placement algorithms "Memory management with linked lists" (4)

Memory management with linked lists: another way of keeping track of memory is to maintain a linked list of allocated and free memory segments, where a segment either contains a process or is an empty hole between two processes.

	initial state			first fit		next fit		best fit		
	Туре	Size (MB)		Туре	Size (MB)		Туре	Size (MB)	Туре	Size (MB)
	Н	8		Н	8		Н	8	Н	8
	P4		1	P4			P4		P4	
	Н	12	1	Н	12		Н	12	Н	12
	P1		1	P1			P1		P1	
	Н	22	1	P9	16		Н	22	Н	22
the last allocated	P3			Н	6		P3		P3	
hlook	Н	18		P3			Н	18	P9	16
UIUCK	P5		1	Н	18		P5		Н	2
	P2			P5			P2		P5	
	Н	8		P2			Н	8	P2	
	P7			Н	8		P7		Н	8
	Н	6		P7			Н	6	P7	
	P6			Н	6		P6		Н	6
	Н	14		P6			Н	14	P6	
	P8			Н	14		P8		Н	14
	Н	36		P8			P9	16	P8	
		•	•	Н	36		Н	20	Н	36

e.g. a memory with 8 processes (P1 to P8), P2 is the last allocated. What about P9 (16 MB) with the placement algorithms.

Methods	Partitioning	Placement	Fragmentation	Swapping	Address binding	Layer
		algorithms	/ compaction		& protection	

Fixed partitioning	contiguous / fixed / complete	searching algorithms	yes / no			OS kernel
Memory management with bitmap	contiguous / dynamic / complete				no / yes (MMU)	
Memory management with linked lists	contiguous / dynamic / complete		yes / yes	yes		
Buddy memory allocation	contiguous / hybrid / complete		yes / no			
Simple paging and segmentation	noncontiguous / dynamic / complete		yes / no		yes (TLB)	programs / services

# Partitioning and placement algorithms "Buddy memory allocation" (1)

**The Buddy memory allocation:** fixed and dynamic partitioning schemes have drawbacks. A fixed partitioning limits the number of process and may use space inefficiently. The dynamic partitioning is more complex to maintain and includes the overhead of compaction. An interesting compromise is the Buddy memory allocation, employing an hybrid partitioning.

*blocks* in a buddy system, memory blocks of holes are available of size  $2^{K}$ , with  $L \le K \le U(L, U)$  are the smallest and largest sizes blocks respectively).

*init* to begin, the entire space is treated as a single block hole of size  $2^{K=U}$ .

*request* at any time, the buddy system maintains a list of holes (unallocated blocks) of each size  $2^i$ : (1) a hole may be removed from the i+1 list by splitting it in half to create two buddies of size  $2^i$  in the *i* list.

(2) whenever a pair of buddies on the *i* list becomes unallocated, they are removed from that list and coalesced into a single block on the i+1 list of size  $2^{i+1}$ .

(3) presented a request for an allocation of size *s* such that  $2^{i-1} \le s \le 2^i$ , the following recursive algorithm applies.

get <i>hole</i> of size <i>i</i>	
if <i>i</i> equals <i>U</i> +1 or <i>L</i> -1	failura casa
failure	
if the <i>list i</i> is empty	
get <i>hole</i> of size <i>i</i> +1	recursive search
split <i>hole</i> into <i>buddies</i>	
put <i>buddies</i> on the <i>i list</i>	
take first <i>hole</i> on the <i>i list</i>	access case

# Partitioning and placement algorithms "Buddy memory allocation" (2)

**The Buddy memory allocation:** fixed and dynamic partitioning schemes have drawbacks. A fixed partitioning limits the number of process and may use space inefficiently. The dynamic partitioning is more complex to maintain and includes the overhead of compaction. An interesting compromise is the Buddy memory allocation, employing an hybrid partitioning.

1 MB block	1 MB							
request A = 100K	A=100K	12	8K	256K	512K			
request $B = 240 K$	A=100K	12	8K	B=240K	512	2K		
request $C = 64 K$	A=100K	C=64K	64K	B=240K	512	2K		
request $D = 256K$	A=100K	C=64K	64K	B=240K	D=256K	256K		
release B	A=100K	C=64K	64K	256K	D=256K	256K		
release A	128K	C=64K	64K	256K	D=256K	256K		
request $E = 75K$	E=75K	C=64K	64K	256K	D=256K	256K		
release C	E=75K	128K		256K	D=256K	256K		
release E		512	2K		D=256K	256K		
release D				1MB				

e.g. a 1MB memory with  $L=6 \le K \le U=10$ , considered blocks are 64K, 128K, 256K, 512K, 1MB.

# Partitioning and placement algorithms "Buddy memory allocation" (3)

**The Buddy memory allocation:** fixed and dynamic partitioning schemes have drawbacks. A fixed partitioning limits the number of process and may use space inefficiently. The dynamic partitioning is more complex to maintain and includes the overhead of compaction. An interesting compromise is the Buddy memory allocation, employing an hybrid partitioning.

e.g. a 1MB memory with  $L=6 \le K \le U=10$ , considered blocks are 64K, 128K, 256K, 512K, 1MB.



# Operating Systems "Memory Management"

1. Introduction

2. Contiguous memory allocation

- 2.1. Partitioning and placement algorithms
- 2.2. Memory fragmentation and compaction

2.3. Process swapping

- 2.4. Loading, address binding and protection
- 3. Simple paging and segmentation
  - 3.1. Paging, basic method
  - 3.2. Segmentation

Methods	Partitioning	Placement	Fragmentation	Swapping	Address binding	Layer
		algorithms	/ compaction		& protection	

Fixed partitioning	contiguous / fixed / complete		yes / no			OS kernel
Memory management with bitmap	contiguous / dynamic / complete				no / yes (MMU)	
Memory management with linked lists	contiguous / dynamic / complete	searching algorithms	yes / yes	yes		
Buddy memory allocation	contiguous / hybrid / complete		yes / no			
Simple paging and segmentation	noncontiguous / dynamic / complete		yes / no		yes (TLB)	programs / services

#### Memory fragmentation and compaction (1)

**External fragmentation:** with dynamic partitioning, as processes are loaded and removed from memory, the free memory space could be broken into little pieces. This phenomenon is known as external fragmentation of memory.



#### Memory fragmentation and compaction (2)

**External fragmentation:** with dynamic partitioning, as processes are loaded and removed from memory, the free memory space could be broken into little pieces. This phenomenon is known as external fragmentation of memory.

**50-percent rule:** memory fragmentation depends of the exact sequence of incoming processes, the considered system and placement algorithms (e.g. first, next and best fit). Statistical analysis reveals that even with some optimization, given N allocated blocks 0,5N block will be lost to fragmentation. That is, one third of memory  $\frac{0,5N}{1N+0,5N}$  may be unusable.

## Memory fragmentation and compaction (3)

**External fragmentation:** with dynamic partitioning, as processes are loaded and removed from memory, the free memory space could be broken into little pieces. This phenomenon is known as external fragmentation of memory.

**Memory compaction:** when internal fragmentation results in multiple holes in memory, it is possible to combine them all into one big one by moving all the processes downward as far as possible. This technique is knows as memory compaction. It requires lot of CPU time.

e.g. on a 1GB machine that can copy 4 bytes in 20 ns, it would take 5 s to compact all memory.

#### Memory fragmentation and compaction (4)

Internal fragmentation: memory allocated to a process may be larger than the request memory, two cases occur.

**Track overhead:** the general approach is to break the physical memory into fixed size blocks and allocate memory in units based on block size. The difference between these two numbers is internal fragmentation of memory.

e.g.



#### Memory fragmentation and compaction (5)

Internal fragmentation: memory allocated to a process may be larger than the request memory, two cases occur.

**Extra memory allocation:** when it is expected that the most processes will grow as they run, it is probably a good idea to allocate a little extra memory.

e.g. a process with two growing segments, a data segment to being used as a heap for variables that are dynamically allocated and released, and a stack segment for the local normal variables and return addresses. An arrangement is to place the stack at the top that is growing downward, and the data segment just beyond the program text that is growing upward.



# Operating Systems "Memory Management"

1. Introduction

2. Contiguous memory allocation

- 2.1. Partitioning and placement algorithms
- 2.2. Memory fragmentation and compaction

2.3. Process swapping

- 2.4. Loading, address binding and protection
- 3. Simple paging and segmentation
  - 3.1. Paging, basic method
  - 3.2. Segmentation

Methods	Partitioning	Placement	Fragmentation	Swapping	Address binding	Layer
		algorithms	/ compaction		& protection	

Fixed partitioning	contiguous / fixed / complete		yes / no			
Memory management with bitmap	contiguous / dynamic / complete		wastwas		no / yes (MMU)	OS kernel
Memory management with linked lists	contiguous / dynamic / complete	searching algorithms	yes / yes	yes		
Buddy memory allocation	contiguous / hybrid / complete		yes / no			
Simple paging and segmentation	noncontiguous / dynamic / complete		yes / no		yes (TLB)	programs / services

#### Process swapping (1)

**Swapping** is a strategy to deal with memory overload, it consists in bringing each process in its entirely, running if for a while, then putting it back on the disk.

e.g. a 88MB memory system, we consider the following events with swapping.

(a) P1 (32MB) (b) P2 (16MB) loaded loaded		(c) P3 ( loa	(c) P3 (32MB) loaded		wapped m disk	(e) P4 (16MB) loaded		(f) P2 (16MB) leaves		(g) P1 swapped in at a different location			
Туре	Size (MB)	Туре	Size (MB)	Туре	Size (MB)	Туре	Size (MB)	Туре	Size (MB)	Туре	Size (MB)	Туре	Size (MB)
P1	32	P1	32	P1	32	Н	32	P4	16	P4	16	P4	16
Н	56	P2	16	P2	16	P2	16	Н	16	Н	32	P1	32
Total	88	Н	40	P3	32	P3	32	P2	16	P3	32	P3	32
		Total	88	Н	8	Н	8	P3	32	Н	8	Н	8
				Total	88	Total	88	H	8	Total	88	Total	88
								Total	88				

#### Process swapping (2)

**Swapping** is a strategy to deal with memory overload, it consists in bringing each process in its entirely, running if for a while, then putting it back on the disk.



**Backing store** is a fast disk that must be large enough to accommodate copies o all memory images. The system maintains suspend queues consisting in all processes whose are on the backing store. The context switch time in such a swapping system is fairly high.

#### Process swapping (3)

**Swapping** is a strategy to deal with memory overload, it consists in bringing each process in its entirely, running if for a while, then putting it back on the disk.



The context switch time in such a swapping system is fairly high.

e.g. a user process of 10 MB and backing store with a transfer rate of 40 MB.s<sup>-1</sup>

- The standard swap operation is  $10/40 = \frac{1}{4}$  s = 250 ms.
- Assuming no head seeks are necessary, and a latency of 8 ms, the swap time is 258 ms.
- The total swap time (in and out) is then 516 ms.

# Operating Systems "Memory Management"

1. Introduction

2. Contiguous memory allocation

- 2.1. Partitioning and placement algorithms
- 2.2. Memory fragmentation and compaction

2.3. Process swapping

- 2.4. Loading, address binding and protection
- 3. Simple paging and segmentation
  - 3.1. Paging, basic method
  - 3.2. Segmentation

Methods	Partitioning	Placement	Fragmentation	Swapping	Address binding	Layer
		algorithms	/ compaction		& protection	

Fixed partitioning	contiguous / fixed / complete	complete				
Memory management with bitmap	contiguous / dynamic / complete		wastwas		no / yes (MMU)	OS kornal
Memory management with linked lists	contiguous / dynamic / complete	searching algorithms	yes / yes	yes		05 kenner
Buddy memory allocation	mory contiguous / hybrid / complete		yes / no			
Simple paging and segmentation	noncontiguous / dynamic / complete		yes / no		yes (TLB)	programs / services

# Loading, address binding and protection "Loading and address binding" (1)



**Loading:** the first step is to load a program into the main memory and to create a process image. Any program consists in compiled modules in an object-code form linked together or to library routines. The library routines can be incorporated into the program or referenced as a shared code to be call at the run time.

Address space is the set of addresses that a process can use to address memory. Each process has its own address space, independent of those belonging to other processes.

Address binding: addresses may be represented in different way between the disk and main memory spaces. Address binding is a mapping from one address space to another.

## Loading, address binding and protection "Loading and address binding" (2)



# Loading, address binding and protection "Loading and address binding" (3)

**Absolute loading:** an absolute loader requires that a given load module always be loaded into the same location in main memory. Thus, in the load module presented to the loader, all address references must be specific to main memory addresses.

Binding time	Function
programming	All actual physical addresses are directly specified
time	by the programmer in the program itself.



# Loading, address binding and protection "Loading and address binding" (4)

**Absolute loading:** an absolute loader requires that a given load module always be loaded into the same location in main memory. Thus, in the load module presented to the loader, all address references must be specific to main memory addresses.

Binding time	Function
compile or assembly time	The program contains symbolic address references, and these are converted to actual physical addresses by compiler or assembler.

	Module C a	at lo	oad module	С	c main					
ľ	orogrammir	ıg aft	tion	memory	256					
						230				
	ADD		ADD	92						
	MOV	Х	MOV	88						
				84						
				80		96				
				76	С					
				72	Operating	64				
				68	System in					
	JMP X		JMP 88	64	RAM	0				
	L	1				0				

# Loading, address binding and protection "Loading and address binding" (5)

**Relocatable loading:** when many programs share a main memory, it may not desirable to decide ahead of time into which region of memory a particular module should be loaded. It is better to make that decision at load time, thus we need a load module that can be located anywhere in main memory.

Binding time	Function	m	odule befo	Main				
load time	The compiler or assembler produces relative	t	to be loade	ed load	ded with X	=128	memory	_
	addresses. The loader translates these to absolute	[		]		]		
	addresses at the time of program loading.		ADD		ADD	92		
		[	MOV	Y+X	MOV	88		132
	The load module must include information about					84	С	128
	that tells the loader where the addresses references			-		80		120
	are. This set of information is prepared by the					76		
	compiler and referred as the relocation directory.			1		72	Operating	1
		Ī				68	System in	

JMP X+Y

Х

0

RAM

64

JMP 88

# Loading, address binding and protection "Loading and address binding" (6)

**Dynamic runtime loading:** to maximize memory utilization, we would like to be able to swap the process image back into different locations at different times, that involves to update the load module at every swap. The relocatable loaders cannot support this scheme. The alternative is to defer the calculation of the absolute address until it is actually needed at run time.

Binding time	Function
run time	The loaded program retains relative addresses, these are converted dynamically to absolute addresses by processor hardware.



**Logical address space** is the set of all logical addresses generated by a program, in the range [0,max].

**Relocation register:** the value *R* in the relocation register is added to every logical address to obtain the corresponding physical address.

**Physical address space:** is the set of all physical addresses corresponding to the logical addresses, in the range [R+0,R+max].

**Memory Management Unit (MMU):** the run time mapping from logical to physical addresses is done by an hardware device called MMU.

# Loading, address binding and protection "Address protection" (1)

Address protection: we need to make sure that every process has a separate memory space. To do this, we need the ability to determine the range of legal addresses that the process may access and to ensure that the process can access only these legal addresses. Protection of memory space is accomplished by having a special CPU hardware.

#### Non-dynamic address protection (absolute and relocatable loading)



**Base register** holds the smallest legal physical memory address.

**Limit (or bound) register** specifies the size of the range e.g. if the base register holds 300 040 and limit register is 120 900, then the program can legally access all addresses from 300 040 to 420 940.

# Loading, address binding and protection "Address protection" (2)

**Address protection:** we need to make sure that every process has a separate memory space. To do this, we need the ability to determine the range of legal addresses that the process may access and to ensure that the process can access only these legal addresses. Protection of memory space is accomplished by having a special CPU hardware.

#### Dynamic address protection (dynamic runtime loading)



The protection must be tuned with the relocation register in the case of MMU.

# Operating Systems "Memory Management"

1. Introduction

2. Contiguous memory allocation

- 2.1. Partitioning and placement algorithms
- 2.2. Memory fragmentation and compaction

2.3. Process swapping

- 2.4. Loading, address binding and protection
- 3. Simple paging and segmentation
  - 3.1. Paging, basic method
  - 3.2. Segmentation

Methods	Partitioning	Placement	Fragmentation	Swapping	Address binding	Layer
		algorithms	/ compaction		& protection	

Fixed partitioning	contiguous / fixed / complete	ntiguous / fixed / complete					
Memory management with bitmap	contiguous / dynamic / complete					OS korral	
Memory management with linked lists	contiguous / dynamic / complete	searching algorithms	yes / yes	yes	no / yes (wiwic)	05 kenner	
Buddy memory allocation	ry contiguous / hybrid / complete		yes / no				
Simple paging and segmentation	noncontiguous / dynamic / complete		yes / no		yes (TLB)	programs / services	

## Simple paging and segmentation "Paging, basic method" (1)

**Paging:** is a memory-management scheme that permits the physical address space of a process to be noncontiguous. The basic method for implementing paging involves breaking logical memory into fixed-sized blocks called **pages** and breaking physical memory into blocks of the same size called **frame**. A **page table** contains the base address of each page in physical memory. When a process is to be executed, its pages are loaded into any available memory frames from the backing store.



#### Simple paging and segmentation "Paging, basic method" (2)

**Paging:** is a memory-management scheme that permits the physical address space of a process to be noncontiguous. The basic method for implementing paging involves breaking logical memory into fixed-sized blocks called **pages** and breaking physical memory into blocks of the same size called **frame**. A **page table** contains the base address of each page in physical memory. When a process is to be executed, its pages are loaded into any available memory frames from the backing store.

	$(t_0)$ initial stat	te (t	(1) A loaded	(t	2) B loaded	(t	3) C loaded	(t <sub>4</sub> ) l	B swapped	out (t	5) D loaded
	0	0	A.0	0	A.0	0	A.0	0	A.0	0	A.0
	1	1	A.1	1	A.1	1	A.1	1	A.1	1	A.1
	2	2	A.2	2	A.2	2	A.2	2	A.2	2	A.2
	3	3	A.3	3	A.3	3	A.3	3	A.3	3	A.3
	4	4		4	B.0	4	B.0	4		4	D.0
	5	5		5	B.1	5	B.1	5		5	D.1
	6	6		6	B.2	6	B.2	6		6	D.2
frame	7	7		7		7	C.0	7	C.0	7	C.0
number	8	8		8		8	C.1	8	C.1	8	C.1
	9	9		9		9	C.2	9	C.2	9	C.2
	10	10		10		10	C.3	10	C.3	10	C.3
	11	11		11		11		11		11	D.3
	12	12		12		12		12		12	D.4
	13	13		13		13		13		13	
	14	14		14		14		14		14	

e.g. 4 processes are loaded in main memory using paging: A, C (4 pages), B (3 pages) and D (5 pages).

49

## Simple paging and segmentation "Paging, basic method" (3)

**Paging:** is a memory-management scheme that permits the physical address space of a process to be noncontiguous. The basic method for implementing paging involves breaking logical memory into fixed-sized blocks called **pages** and breaking physical memory into blocks of the same size called **frame**. A **page table** contains the base address of each page in physical memory. When a process is to be executed, its pages are loaded into any available memory frames from the backing store.

(t	5) D loaded	l					data	atura	<b>t</b>					
0	A.0						data	struc	ture					
1	A.1													٦
2	A.2		nage	table	nage	table	nage	table		nage	table	free	fram	
3	A.3		page A	A	page ]	B	page (	C		page		nee	list	U
4	D.0		n	f	n	f	n	f		n	f	Γ	f	
5	D.1		<b>P</b>	1	<b>P</b>	1	<u>Р</u>	1		P 0	1	-		
6	D.2		0	0	0	-	0	1		0	4		13	
7	C.0		1	1	1	-	1	8		1	5		14	
8	C.1		2	2	2	-	2	9		2	6			
9	C.2		3	3			3	10		3	11			
10	C.3								•	4	12			
11	D.3								1					
12	D.4													
13														
14														

e.g. 4 processes are loaded in main memory using paging with A, C (4 pages), B (3 pages) and D (5 pages).

# Simple paging and segmentation "Paging, basic method" (4)

Address binding: every address generated by the CPU is dived in two parts, a page number (p) and a page offset (d). The page number is used as index into the page table, and the page table contains the base address of each page in physical memory. This base address is combined with the page offset (d) to define the physical memory address that is sent to the memory unit.

The size of a page is typically a power of 2 that makes the translation of a logical address into a page number and page offset particularly easy e.g. if the size of logical address space is  $2^m$ , and a page size is  $2^n$  addressing units, then the high order *m*-*n* bits designates the page number. As a result, the addressing scheme is transparent to programmer, assembler and linker.



- number of bits to encode the logical address
- number of bit to encode the page offset
- *m-n* number of bits to encode the page number

# Simple paging and segmentation "Paging, basic method" (5)

Address binding: every address generated by the CPU is dived in two parts, a page number (p) and a page offset (d). The page number is used as index into the page table, and the page table contains the base address of each page in physical memory. This base address is combined with the page offset (d) to define the physical memory address that is sent to the memory unit.



#### Simple paging and segmentation "Paging, basic method" (6)

Address binding: every address generated by the CPU is dived in two parts, a page number (p) and a page offset (d). The page number is used as index into the page table, and the page table contains the base address of each page in physical memory. This base address is combined with the page offset (d) to define the physical memory address that is sent to the memory unit.

physical address logical address address d data page table f d f р р address space space а b с d e f g h i the logical address i 0101 (page 1, offset 1) corresponds to the physical address k 1101 (frame 3, offset 1) m n р

e.g. a memory of 16 bytes with a page size of 4 bytes, consider the logical address 0101.

data

i

k

а

b

с

d

m

n

р

e

f

g

h

# Simple paging and segmentation "Paging, basic method" (7)

**Transaction look-aside buffer (TLB):** in principle every binding from logical address space to physical address space using paging causes two physical memory accesses (1) to fetch the appropriate page table entry (2) to fetch the desired data. This will have the effect of doubling the memory access time. To overcome this problem, we use a special high speed cache for page table entries called Transaction look-aside buffer (TLB).



**TLB** is an associative high speed memory. Each entry in the TLB consists of two parts {key; value}.

When an **item** (i.e. a page number) is presented, it is compared with all the **keys** simultaneously. This technique is referred as **associative mapping**. If **item == key**, the corresponding **value** (i.e. offset) is returned.

The search is fast and hardware, however, is expensive. Typically, the number of entries in a TLB is small e.g. 64 to 1024 entries.

# Simple paging and segmentation "Paging, basic method" (8)

**Transaction look-aside buffer (TLB):** in principle every binding from logical address space to physical address space using paging causes two physical memory accesses (1) to fetch the appropriate page table entry (2) to fetch the desired data. This will have the effect of doubling the memory access time. To overcome this problem, we use a special high speed cache for page table entries called Transaction look-aside buffer (TLB).



**TLB hit:** it occurs when the page number is in the TLB.

**TLB miss:** if the page number is not in the TLB, we add the page number and frame so that they will be found quickly on the next reference. If the TLB is already full of entries, OS must select one entry for replacement (LRU or FIFO policy, etc).

**wired down:** some TLBs allow entries to be wired down, meaning that they cannot be removed from the TLB.

# Simple paging and segmentation "Paging, basic method" (9)

**Transaction look-aside buffer (TLB):** in principle every binding from logical address space to physical address space using paging causes two physical memory accesses (1) to fetch the appropriate page table entry (2) to fetch the desired data. This will have the effect of doubling the memory access time. To overcome this problem, we use a special high speed cache for page table entries called Transaction look-aside buffer (TLB).



**TLB hit:** it occurs when the page number is in the TLB.

**TLB miss:** if the page number is not in the TLB, we add the page number and frame so that they will be found quickly on the next reference. If the TLB is already full of entries, OS must select one entry for replacement (LRU or FIFO policy, etc).

wired down: some TLBs allow entries to be wired down, meaning that they cannot be removed from the TLB.

## Simple paging and segmentation "Paging, basic method" (10)

**Shared pages:** an advantage of paging is the possibility of sharing common code. This can appear with reentrant code (or pure code) that is a non-self-modifying code i.e. it never changes during execution.

e.g. consider a system that supports three users, each of whom executes a text editor. It the text editor consists of 150 KB of code and 50 KB of data space, we need 600 KB to support the three users. If the code is reentrant code, it can be shared as:



## Simple paging and segmentation "Paging, basic method" (11)

Address protection: address (or memory) protection in a paged environment is accomplished by protection bits associated to each frame. Normally, these bits are kept in the page table.

valid bit (v): when this bit is set to valid, the associated page in the process's logical address space is thus a valid page. When the bit is set to invalid then the page is not in the process's logical address space.

e.g. in a system with m=14 bits address space (0 to 16383), using a page size of 2 KB, then n=11, m-n=3 (i.e. 8 pages). We have a process P of size 10438 bytes.

logical address space		physical address space		page table P				
р	data	f	data		р	f	v	
0	page 0	0	page 4		0	1	true	valid
1	page 1	1	page 0		1	4	true	bit
2	page 2	2			2	3	true	
3	page 3	3	page 2		3	7	true	
4	page 4	4	page 1		4	0	true	
5	page 5	5			5	8	true	
		6			6	×	false	
		7	page 3		7	×	false	
		8	page 5				•	

Accesses to addresses up to  $12287 (6 \times 2^{11})$  are valid, only the addresses from 12288 to 16383 are not valid.

Because the program extends to the address from 10468 to 12288, only the references beyond that address are illegal. This problem results of the 2KB page size and reflects the internal fragmentation of paging.

# Simple paging and segmentation "Paging, basic method" (12)

Address protection: address (or memory) protection in a paged environment is accomplished by protection bits associated to each frame. Normally, these bits are kept in the page table.

**r/w bit:** defines a page to be read-write or read only; we can easily expand this approach to provide a finer level of protection by considering execute-only sate.

# Operating Systems "Memory Management"

1. Introduction

2. Contiguous memory allocation

- 2.1. Partitioning and placement algorithms
- 2.2. Memory fragmentation and compaction

2.3. Process swapping

- 2.4. Loading, address binding and protection
- 3. Simple paging and segmentation
  - 3.1. Paging, basic method
  - 3.2. Segmentation

Methods	Partitioning	Placement	Fragmentation	Swapping	Address binding	Layer
		algorithms	/ compaction		& protection	

Fixed partitioning	contiguous / fixed / complete		yes / no	yes	no / yes (MMU)	OS kernel
Memory management with bitmap	contiguous / dynamic / complete	searching algorithms	yes / yes			
Memory management with linked lists	contiguous / dynamic / complete					
Buddy memory allocation	contiguous / hybrid / complete		yes / no			
Simple paging and segmentation	ple paging and noncontiguous / nentation dynamic / complete		yes / no		yes (TLB)	programs / services

# Simple paging and segmentation "Segmentation" (1)

**One vs. separated address spaces:** logical addressing discussed so far is one-dimensional because the logical addresses go from zero to some maximum. For many problems, having separate logical address space may be much more better.

e.g. a compiler has tables that are built up as compilation proceeds, and each of them grows continuously. In a one dimensional memory, these tables are allocated contiguous chunks of logical addresses. Consider what happens if a program has a much larger than usual number of variables; the chunk of address space allocated for a table may fill up.



Some approaches to deal with are:

(1) the compiler could simply issue a message saying that the compilation cannot continue.

(2) to play "Robin Hood", tacking space from the tables with an excess of room and giving it to the tables with little room. This is a nuisance at best and a great deal of tedious, unrewarding work, at worst.

# Simple paging and segmentation "Segmentation" (2)

Segmentation is a memory management scheme that supports independent address spaces called segment.

-Each segment consists in linear sequence of addresses, from 0 to some maximum (usually very large).

-The length of each segment may be anything from 0 to the maximum allowed.

-Different segments may have different lengths.

-Segment lengths may change during execution.

-With segments, a logical address consists of a two tuple <segment number, offset>.

e.g. a compiler has tables that are built up as compilation proceeds, and each of them grows continuously. Because each segment constitutes a separate address space, the different segments can grow independently without affecting each other.



Logical address space

### Simple paging and segmentation "Segmentation" (3)

Segmentation: is a memory management scheme that supports independent address spaces called segment.

- -Each segment consists in linear sequence of addresses, from 0 to some maximum (usually very large).
- -The length of each segment may be anything from 0 to the maximum allowed.
- -Different segments may have different lengths.
- -Segment lengths may change during execution.
- -With segments, a logical address consists of a two tuple <segment number, offset>.



## Simple paging and segmentation "Segmentation" (4)

**Combined paging and segmentation:** both paging and segmentation have their strengths. In combined paging / segmentation system, the user's address space is broken up into a number of segments, at the discretion of the programmer. Each segment is, in turn, broken up into a number of fixed-size page.

-From the programmer's point of view, a logical address still consists of a segment number and a segment offset.

-From the system's point of view, the segment offset is viewed as a page number and page offset for a page within the specified segment.

