

Memory Management

summary

- goals and requirements
- techniques that do not involve virtual memory

memory management

- tracking used and free memory
- primitives
 - allocation of a certain amount of memory
 - de-allocation of what allocated or garbage collectors, permit reuse of de-allocated memory
- memory used for...
 - data structures (e.g. array, objects, ecc.)
 - kernel data structures (allocator implemented in the kernel)
 - process data structures (allocator implemented by language runtime libraries, e.g. C/C++ malloc)
 - processes (within an O.S.)

memory management within the o.s.

- subdividing memory to accommodate multiple processes
- memory needs to be allocated to ensure a reasonable supply of ready processes to consume available processor time
- needs for memory are hard to know before run
 - usually dynamically allocated to processes upon request (requires a complex management system)

summary and applicability

- many techniques and concepts in memory management equally apply to memory allocation for processes and for data
 - fixed partitioning, dynamic compaction, fragmentation, placement algorithms, buddy system
 - we talk about a “process” but it may be any kind of data
- hardware supported techniques apply only to processes
 - virtual memory, paging, segmentation

memory management
techniques that do not involves
virtual memory



Fixed Partitioning

- Equal-size partitions
 - Any process or data whose size is less than or equal to the partition size can be loaded into an available partition
 - If all partitions are full, the operating system can swap a process out of a partition
 - A process/data may not fit in a partition.
 - For processes, the programmer must design the program with overlays
 - still used in hard disk partitioning
 - LVM overcome such limitation (linux)

Fixed Partitioning

- Main memory use is inefficient. Any program, no matter how small, occupies an entire partition. This is called **internal fragmentation**.


partitions size




Placement Algorithm with Partitions

- Equal-size partitions
 - Because all partitions are of equal size, it does not matter which partition is used
- Unequal-size partitions
 - Can assign each process to the smallest partition it will fit into
 - Queue for each partition
 - Processes/data are assigned in such a way as to minimize wasted memory within a partition

Placement Algorithm with Partitions



(a) One process queue per partition

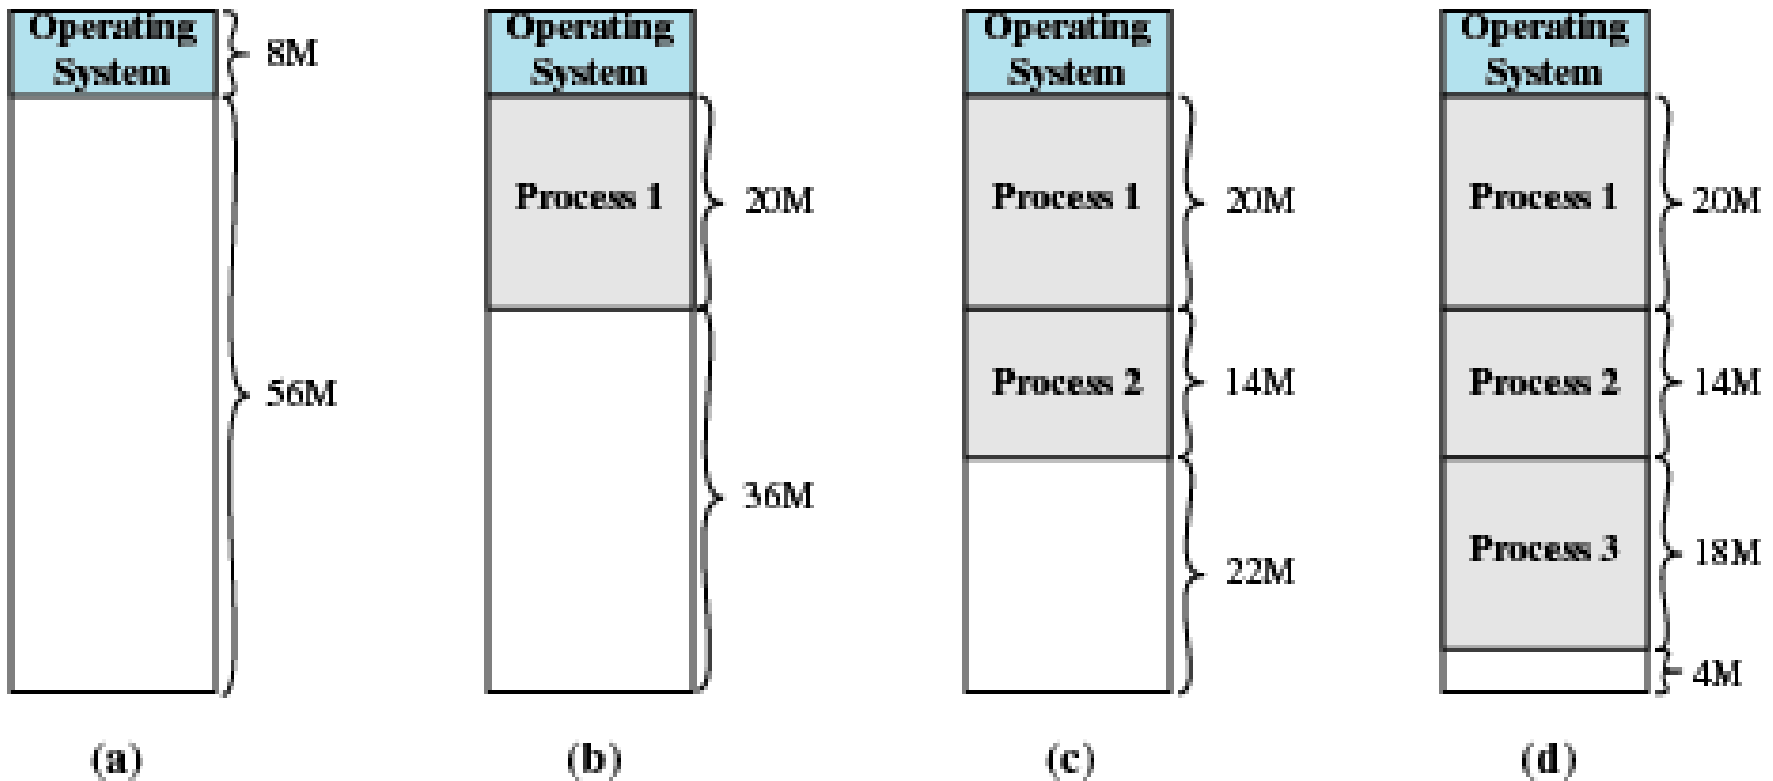


(b) Single queue


dynamic partitioning and compaction

- Partitions are of variable length and number
- Process/data is allocated exactly as much memory as required
- Eventually get holes in the memory. This is called **external fragmentation**
- Must use **compaction** to shift processes so they are contiguous and all free memory is in one block
 - in the general case compaction may be unfeasible (e.g. for C/C++ memory allocators)


external fragmentation




external fragmentation



(e)



(f)



(g)

Dynamic Partitioning Placement Algorithm

- Operating system must decide which free block to allocate to a process
- Best-fit algorithm
 - Chooses the block that is closest in size to the request
 - Worst performer overall
 - since smallest block is found for process, the smallest amount of fragmentation is left
 - Memory compaction must be done more often

Dynamic Partitioning Placement Algorithm



- **First-fit algorithm**
 - Scans memory from the beginning and chooses the first available block that is large enough
 - Fastest
 - May have many process loaded in the front end of memory that must be searched over when trying to find a free block

Dynamic Partitioning Placement Algorithm

- Next-fit
 - Scans memory from the location of the last placement
 - More often allocate a block of memory at the end of memory where the largest block is found
 - The largest block of memory is broken up into smaller blocks
 - Compaction is required to obtain a large block at the end of memory

examples

- allocation of a block of 16MB



Buddy System

- simple but powerful allocator
- widely used in O.S. to allocate large chunks of fixed size
 - e.g. 4KB pages in today architectures (x86_32)
 - no fragmentation in this case
- it can be used as a base for a more fine grained allocator
 - which is called slab allocator in Linux and Sun Solaris and is used for kernel data structures

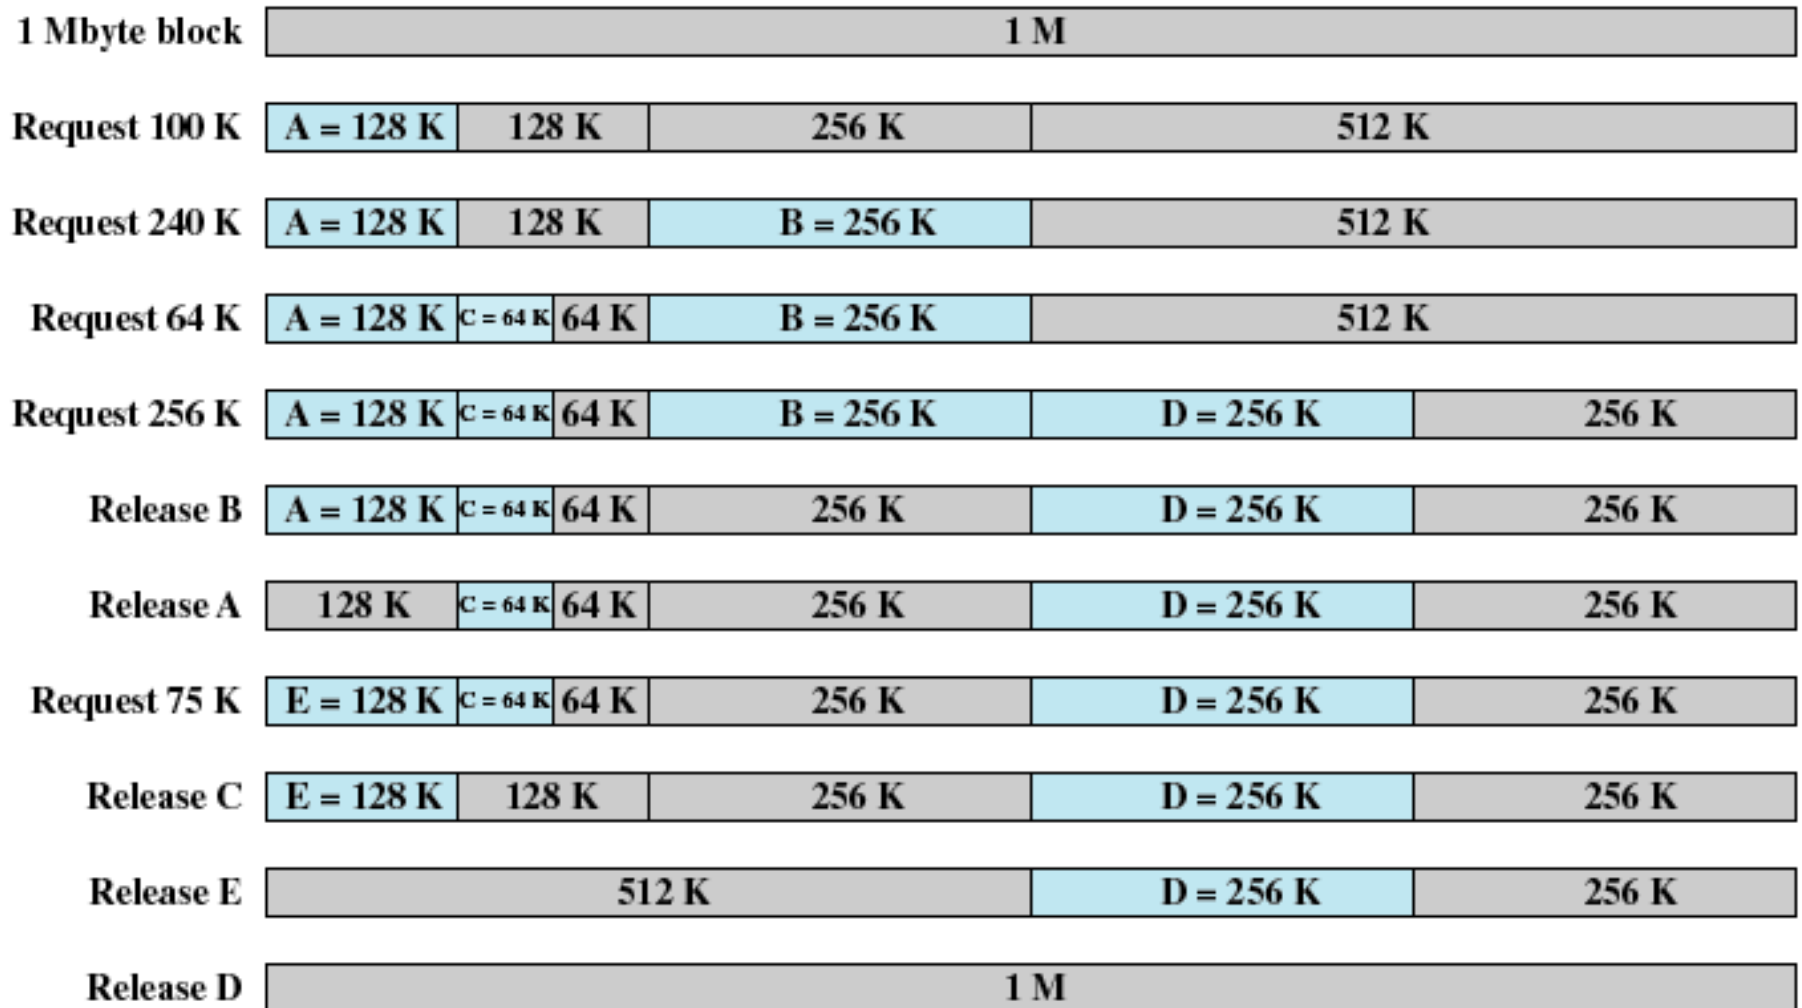
Buddy System

- entire space available is treated as a single block of 2^U
- a request of s bytes returns a block of $\text{ceil}(\log_2 s)$ bytes
 - if a request of size s such that $2^{i-1} < s \leq 2^i$, a block of length 2^i is allocated
 - a 2^i block can be split into two equal **buddies** of 2^{i-1} bytes
 - for each request a “big” block is found and split until the smallest block greater than or equal to s is generated

Buddy System

- it maintains a lists L_i ($i=1..U$) of unallocated blocks (*holes*) of size 2^i
 - splitting: remove a hole from L_{i+1} split it, and put the two buddies it into L_i
 - coalescing: remove two unallocated buddies from L_i and put it into L_{i+1}

buddy system: example



Buddy System

procedure **get_hole**

- input: i (precondition: $i \leq U$)
- output: a block c of size 2^i (postcondition: L_i does not contain c)

if (L_i is empty)

$b = \text{get_hole}(i+1);$

$\langle \text{split } b \text{ into two buddies } b_1 \text{ and } b_2 \rangle$


$\langle \text{put } b_1 \text{ and } b_2 \text{ into } L_i \rangle$

$c = \langle \text{first hole in } L_i \rangle$

$\langle \text{remove } c \text{ from } L_i \rangle$

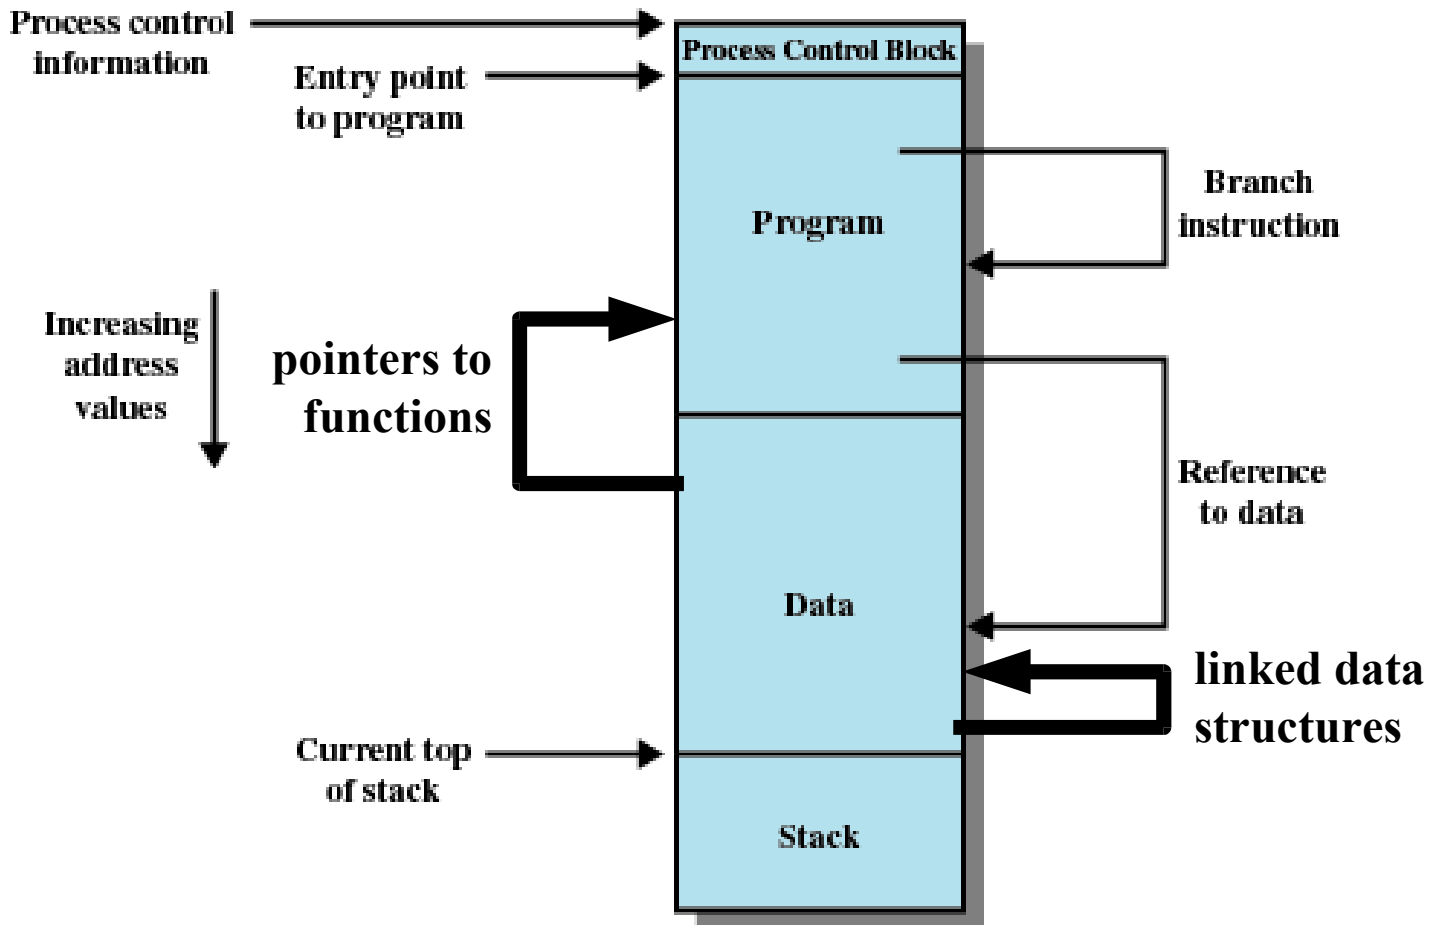
return c

buddy system: tree representation



memory management requirements for processes

- relocation



memory management requirements for processes

- relocation for processes
 - programmer does not know where the program will be placed in memory when it is executed
 - while the program is executing, it may be swapped to disk and returned to main memory at a different location (relocated)
 - memory references in the code must be translated to actual physical memory address
- relocation for data is hard
 - data created by process, relocation performed by OS, need for a standard, never done

memory management

requirements for processes

- protection
 - processes should not be able to reference memory locations in another process without permission
 - references **must** be checked at run time
 - impossible to check memory references at compile time (may directly depend on the input)
 - exercise: given a generic input and program prove that reference check is not computable! (reduce stopping problem to it)
 - memory protection requirement must be satisfied by the processor (hardware) rather than the operating system (software)
 - Operating system cannot anticipate all of the memory references a process will perform

memory management requirements for processes

- sharing
 - allow several processes to access the same portion of memory
 - better to allow each process access to the same copy of the program rather than have their own separate copy

memory management requirements for processes

- logical organization
 - programs are written in modules
 - sw engineering reasons: divide the responsibility for development, maintenance, testing, ecc
 - modules can be written and compiled independently
 - different degrees of protection given to modules (read-only, execute-only)
 - share modules among processes

memory management requirements for processes

- physical organization
 - memory available for a program plus its data may be insufficient
 - overlaying allows various modules to be assigned the same region of memory
 - programmer does not know how much memory will be available

relocation

- a process may occupy different partitions, which means different absolute memory locations
- when a program is loaded into memory the absolute memory locations are determined
 - different execution may lead to different locations
- on-the-fly relocation during execution
 - swap out and swap in
 - compaction of allocated partitions

addresses in the program

- Physical
 - The absolute address or actual location in main memory
- Logical
 - Reference to a location in a “logical” memory independent of the current assignment of data to memory
 - Translation must be made to the physical address by the hardware (MMU)
- Relative (logical or physical)
 - Address expressed as a location relative to some known point

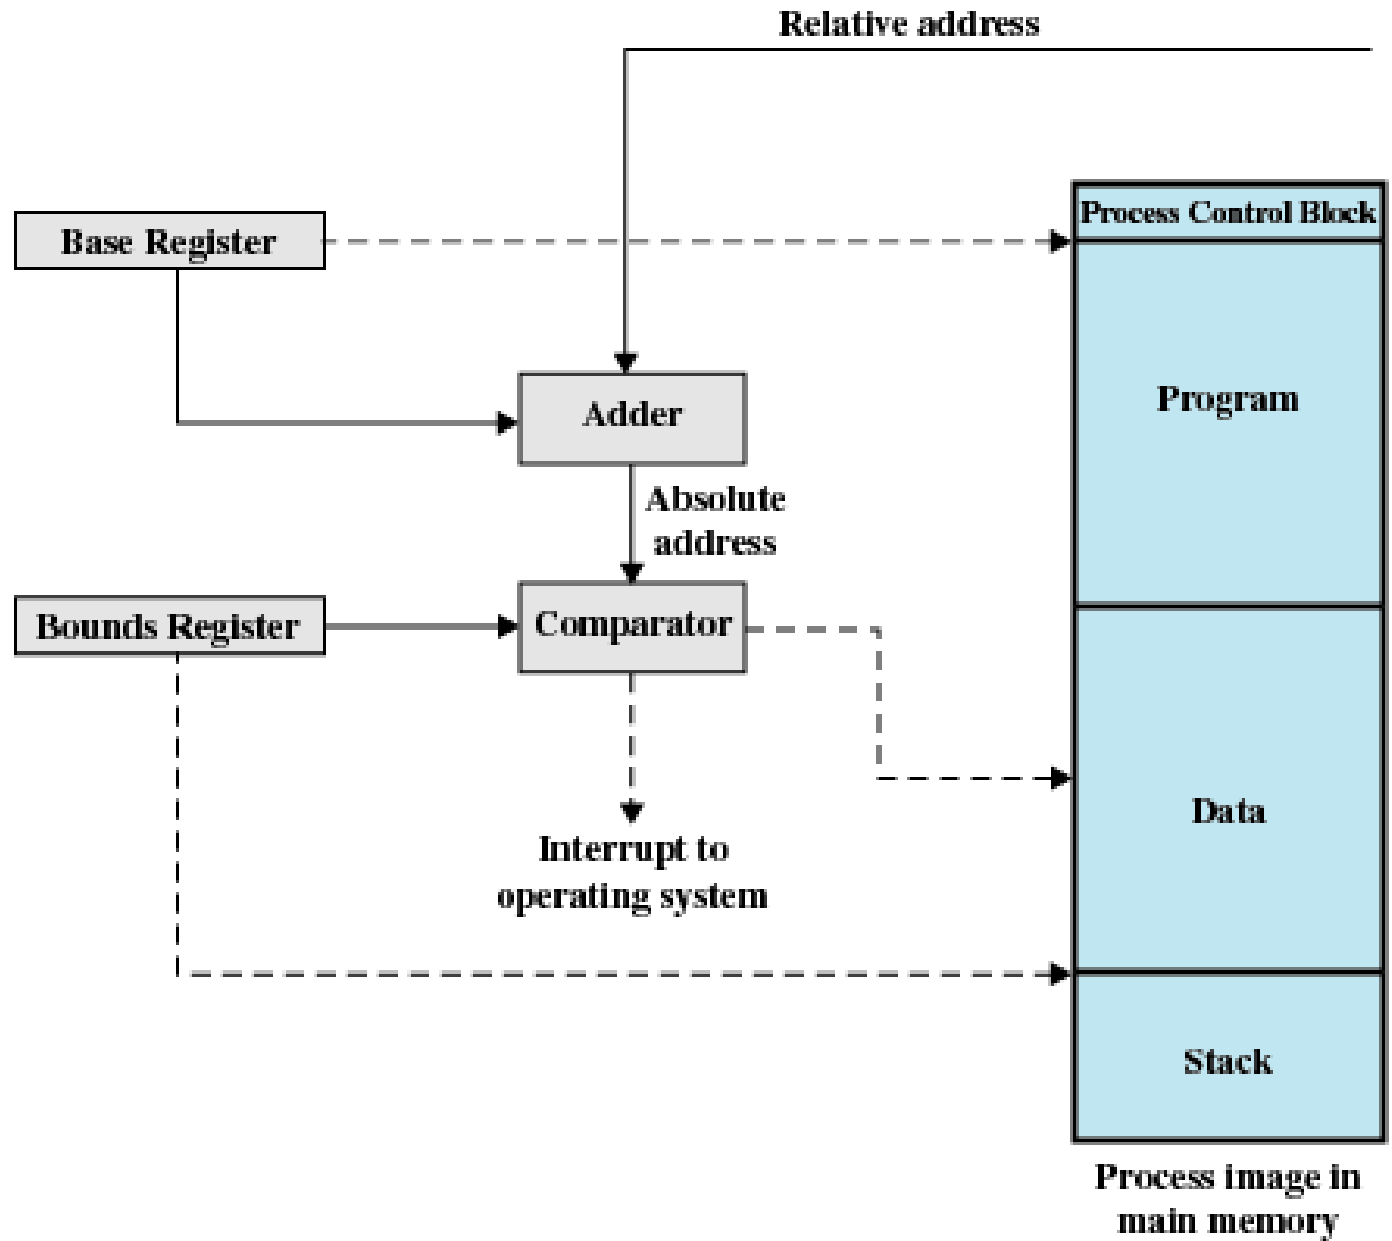
hardware support for relocation

- Base register
 - Starting address for the process
- Bounds register
 - Ending location of the process
- These values are set when the process is loaded or when the process is swapped in

hardware support for relocation

- The value of the base register is added to a relative address to produce an absolute address
- The resulting address is compared with the value in the bounds register
- If the address is not within bounds, an interrupt is generated to the operating system
- If the address is ok it is used to access memory
- relocation is performed by setting appropriate value in the registers


hardware support for relocation




Paging

- Partition memory into small equal fixed-size chunks and divide each process into the same size chunks
- The chunks of a process are called **pages** and chunks of memory are called **frames**
- Operating system maintains a **page table** for each process
 - Contains the **frame location for each page** in the process
 - **Memory address consist of a page number and offset within the page**


Assignment of Process Pages to Free Frames



(a) Fifteen Available Frames




(b) Load Process A




(c) Load Process B


Assignment of Process Pages to Free Frames



(d) Load Process C



(e) Swap out B



(f) Load Process D

Page Tables

0	0
1	1
2	2
3	3

Process A
page table

0	N
1	N
2	N

Process B
page table

0	7
1	8
2	9
3	10

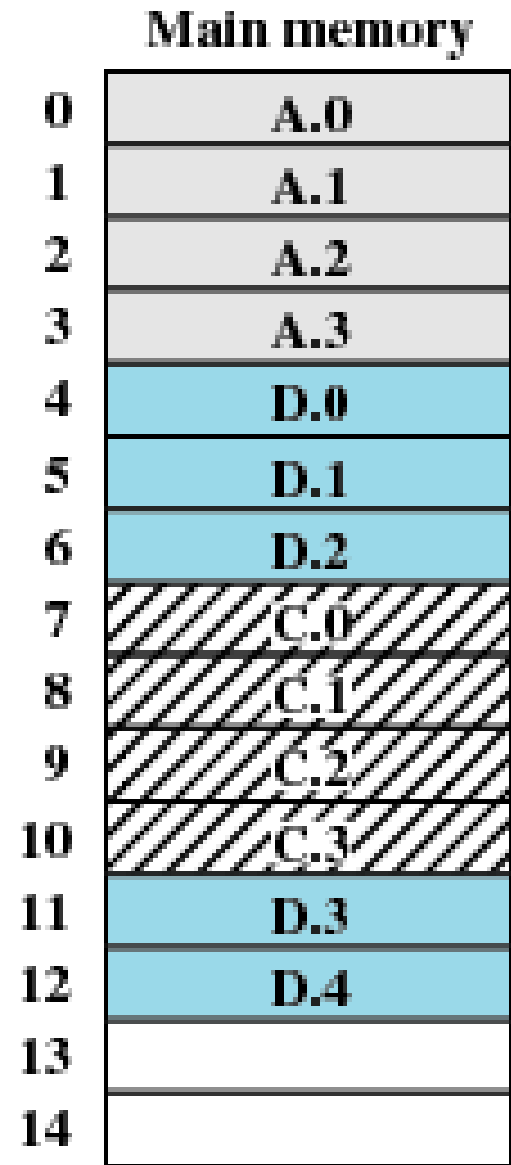
Process C
page table

0	4
1	5
2	6
3	11
4	12

Process D
page table

13
14

Free frame
list




(f) Load Process D

Segmentation


- All segments of all programs do **not** have to be of the **same length**
- There is a maximum segment length
- Addressing consist of two parts - a **segment number and an offset**
- Since segments are not equal, segmentation is similar to dynamic partitioning

paging vs. segmentation



logical to physical translation

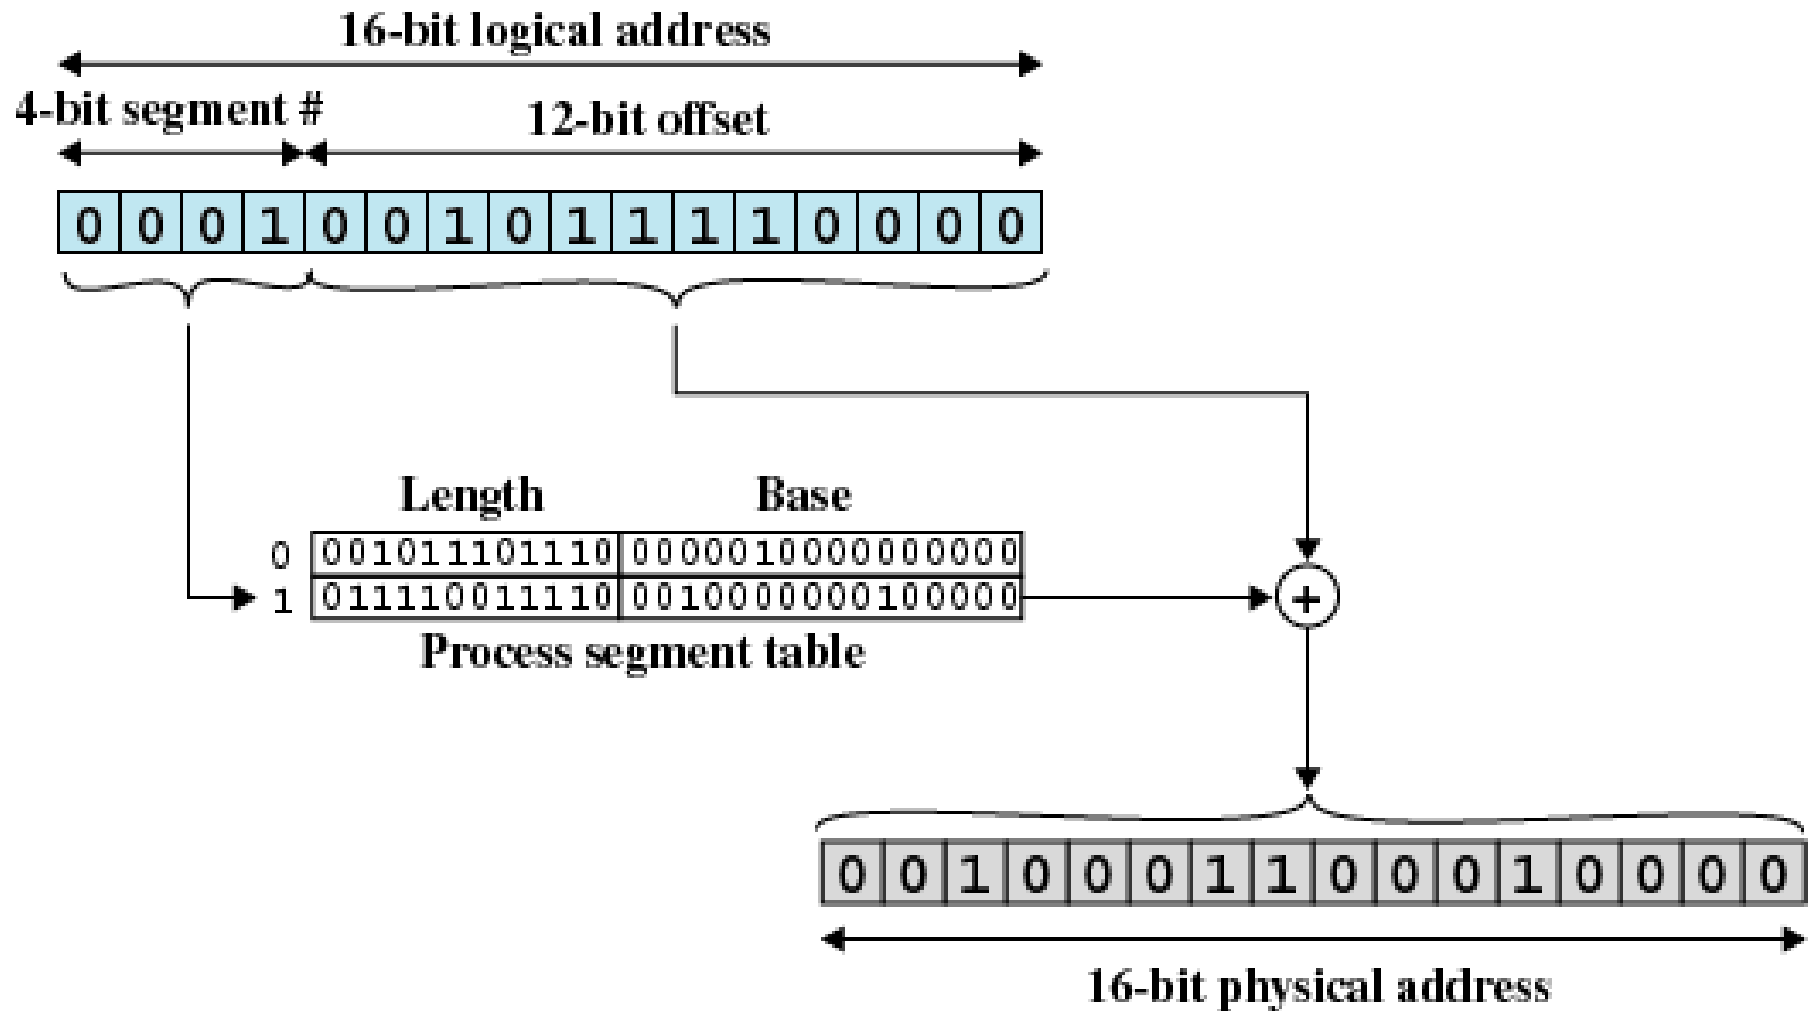
paging



(a) Paging

logical to physical translation

segmentation



(b) Segmentation