

ITP 30002 Operating System

Segmentation

OSTEP Chapters 16 & 17

Shin Hong


Motivation

- **Internal fragmentation problem**

- a process typically uses a small portion of an address space in a sparse manner
- the rest of the address space remains unused and wasted

- **Redundant data**

- The same piece of code may be loaded redundantly if multiple processes use it



Approach

3

- Ideas
 1. split an address space into multiple pieces and manage each separately
 2. allocate the memory to each piece depending on the actual needs
- Define the address space of a process as a set of multiple **segments**
 - a segment is a continuous memory region defined by a pair of base and bounds addresses
 - the available portion of the address space of a process can be represented as a set of segments
 - a segment is initially defined at a loading time
 - a segment can be relocated, extended, or shrunk over time
- Let multiple processes share one segment if they use the same data

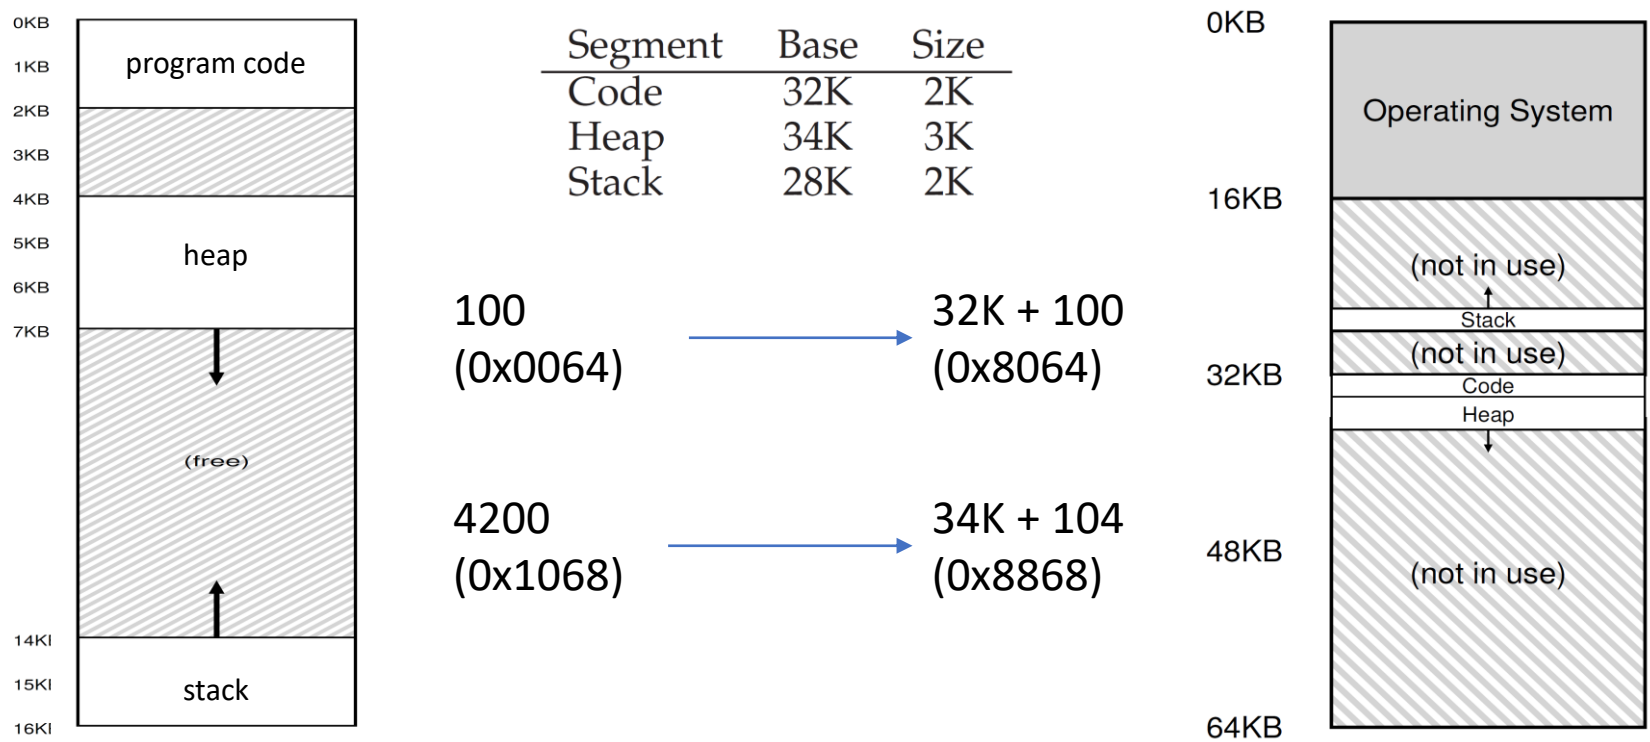
Segmentation

ITP 30002
Operating System

2023-04-06

Segmentation


- A process has three segments, **code**, **heap** and **stack** in an address space
- MMU is required to have three base-bounds pairs, and a mechanism to identify which segment a memory access is on
 - MMU will find the base value and add it to the offset part of a virtual address
- OS can change bounds on demand of a process to extend/shrink a segment



Which Segment A Memory Access is On?

5

- By virtual address
 - use top few bits as a segmentation indicator
 - ex. virtual address 4200



- By instruction type
 - referring the code segment if the address is derived from PC
 - referring the stack segment if the address is derived from stack pointer
 - referring the heap segment otherwise

Segmentation

ITP 30002
Operating System

2023-04-06

Segment Attributes

6

- A bit to indicate whether a segment grows forward or backward

- ex.

Segment	Base	Size (max 4K)	Grows Positive?
Code ₀₀	32K	2K	1
Heap ₀₁	34K	3K	1
Stack ₁₁	28K	2K	0

- A bit to indicate whether a segment is for read-write or read-only

- ex.

Segment	Base	Size (max 4K)	Grows Positive?	Protection
Code ₀₀	32K	2K	1	Read-Execute
Heap ₀₁	34K	3K	1	Read-Write
Stack ₁₁	28K	2K	0	Read-Write

Segmentation

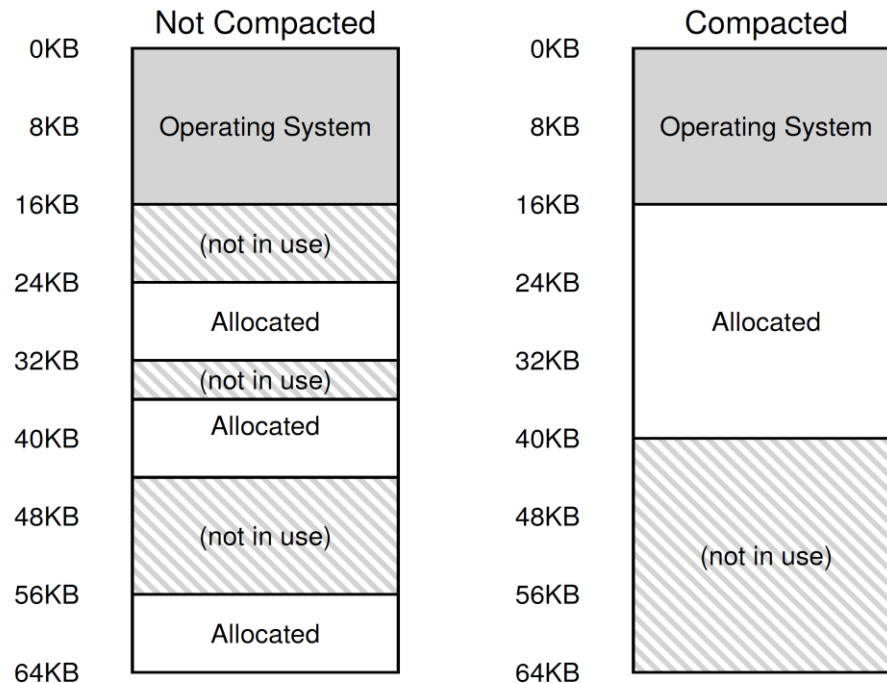
ITP 30002
Operating System

2023-04-06

Operating System Supports

7

- context switching
- serving requests to grow/shrink a segment
- managing free space in physical memory to deal with external fragmentation problem
 - compaction
 - free-list management algorithms are applied



Segmentation


ITP 30002
Operating System

2023-04-06

Free-Space Management

8

- External fragmentation hurts memory utilization
 - commonly happens for dynamic allocation of variable-length memory units
 - example. with a 30-bytes address space
 - alloc 15 bytes as A
 - alloc 10 bytes as B
 - free A (15 bytes)
 - alloc 5 bytes as C
 - alloc 15 bytes as D



Segmentation


ITP 30002
Operating System

2023-04-06

Free List

9

- maintains free space as a linked list of free chunks (i.e., continuous unused memory regions)
- to allocate a chunk of size m
 - find a node of a free chunk whose size is greater than or equal to m
 - split the free chunk if its size is greater than m
- to free an allocated memory chunk,
 - add a node of the newly freed chunk to the free list
 - merge adjacent chunks into a single larger chunk (coalescing)




Allocation Strategies

10

- The ideal allocator should be fast and minimize fragmentation.
- **Best fit:** search through the free list and find the smallest chunks that are large enough to afford the memory request (i.e., closest to what the user asks)
- **Worst fit:** find the largest chunk in order to keep large chunks remain in the free list
- **First fit:** find the first chunk that is large enough to afford the requested memory
- **Next fit:** conduct the first fit search from the node where the last search was stopped (not from the beginning of the free list)

Ex. the user asks 18



Segmentation

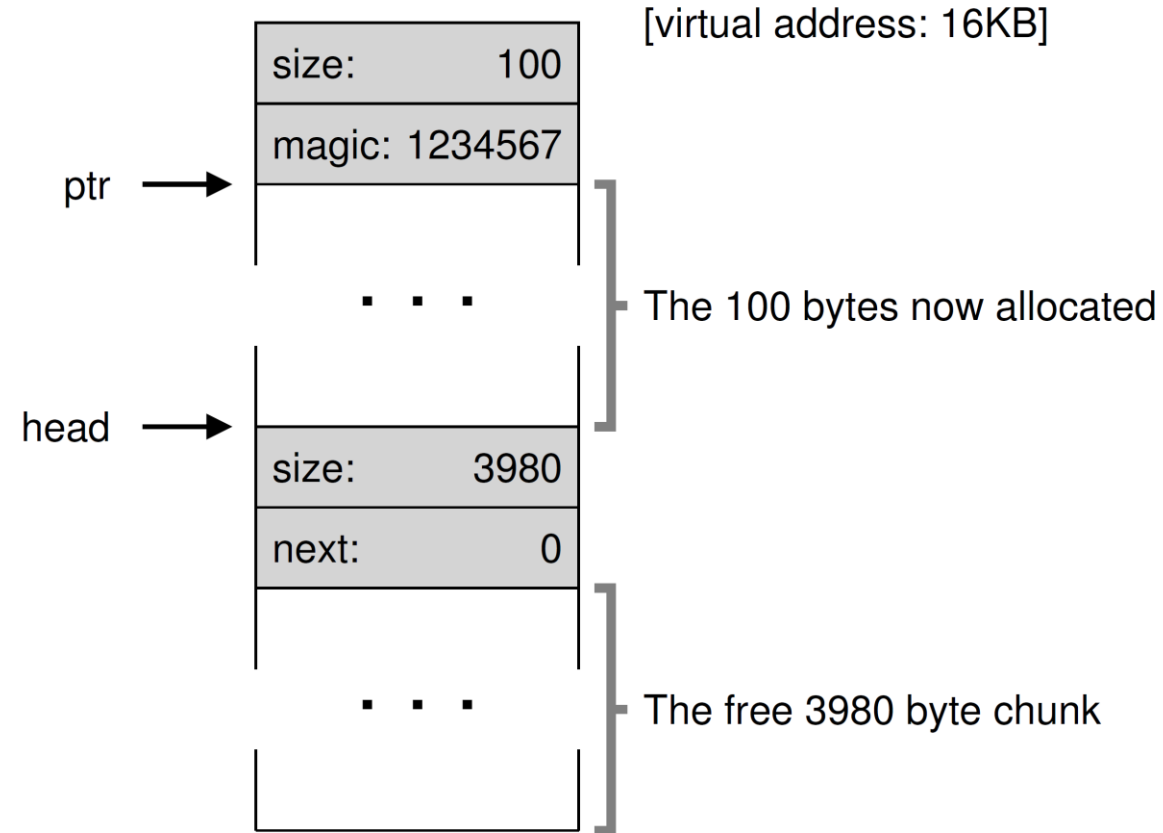

ITP 30002
Operating System

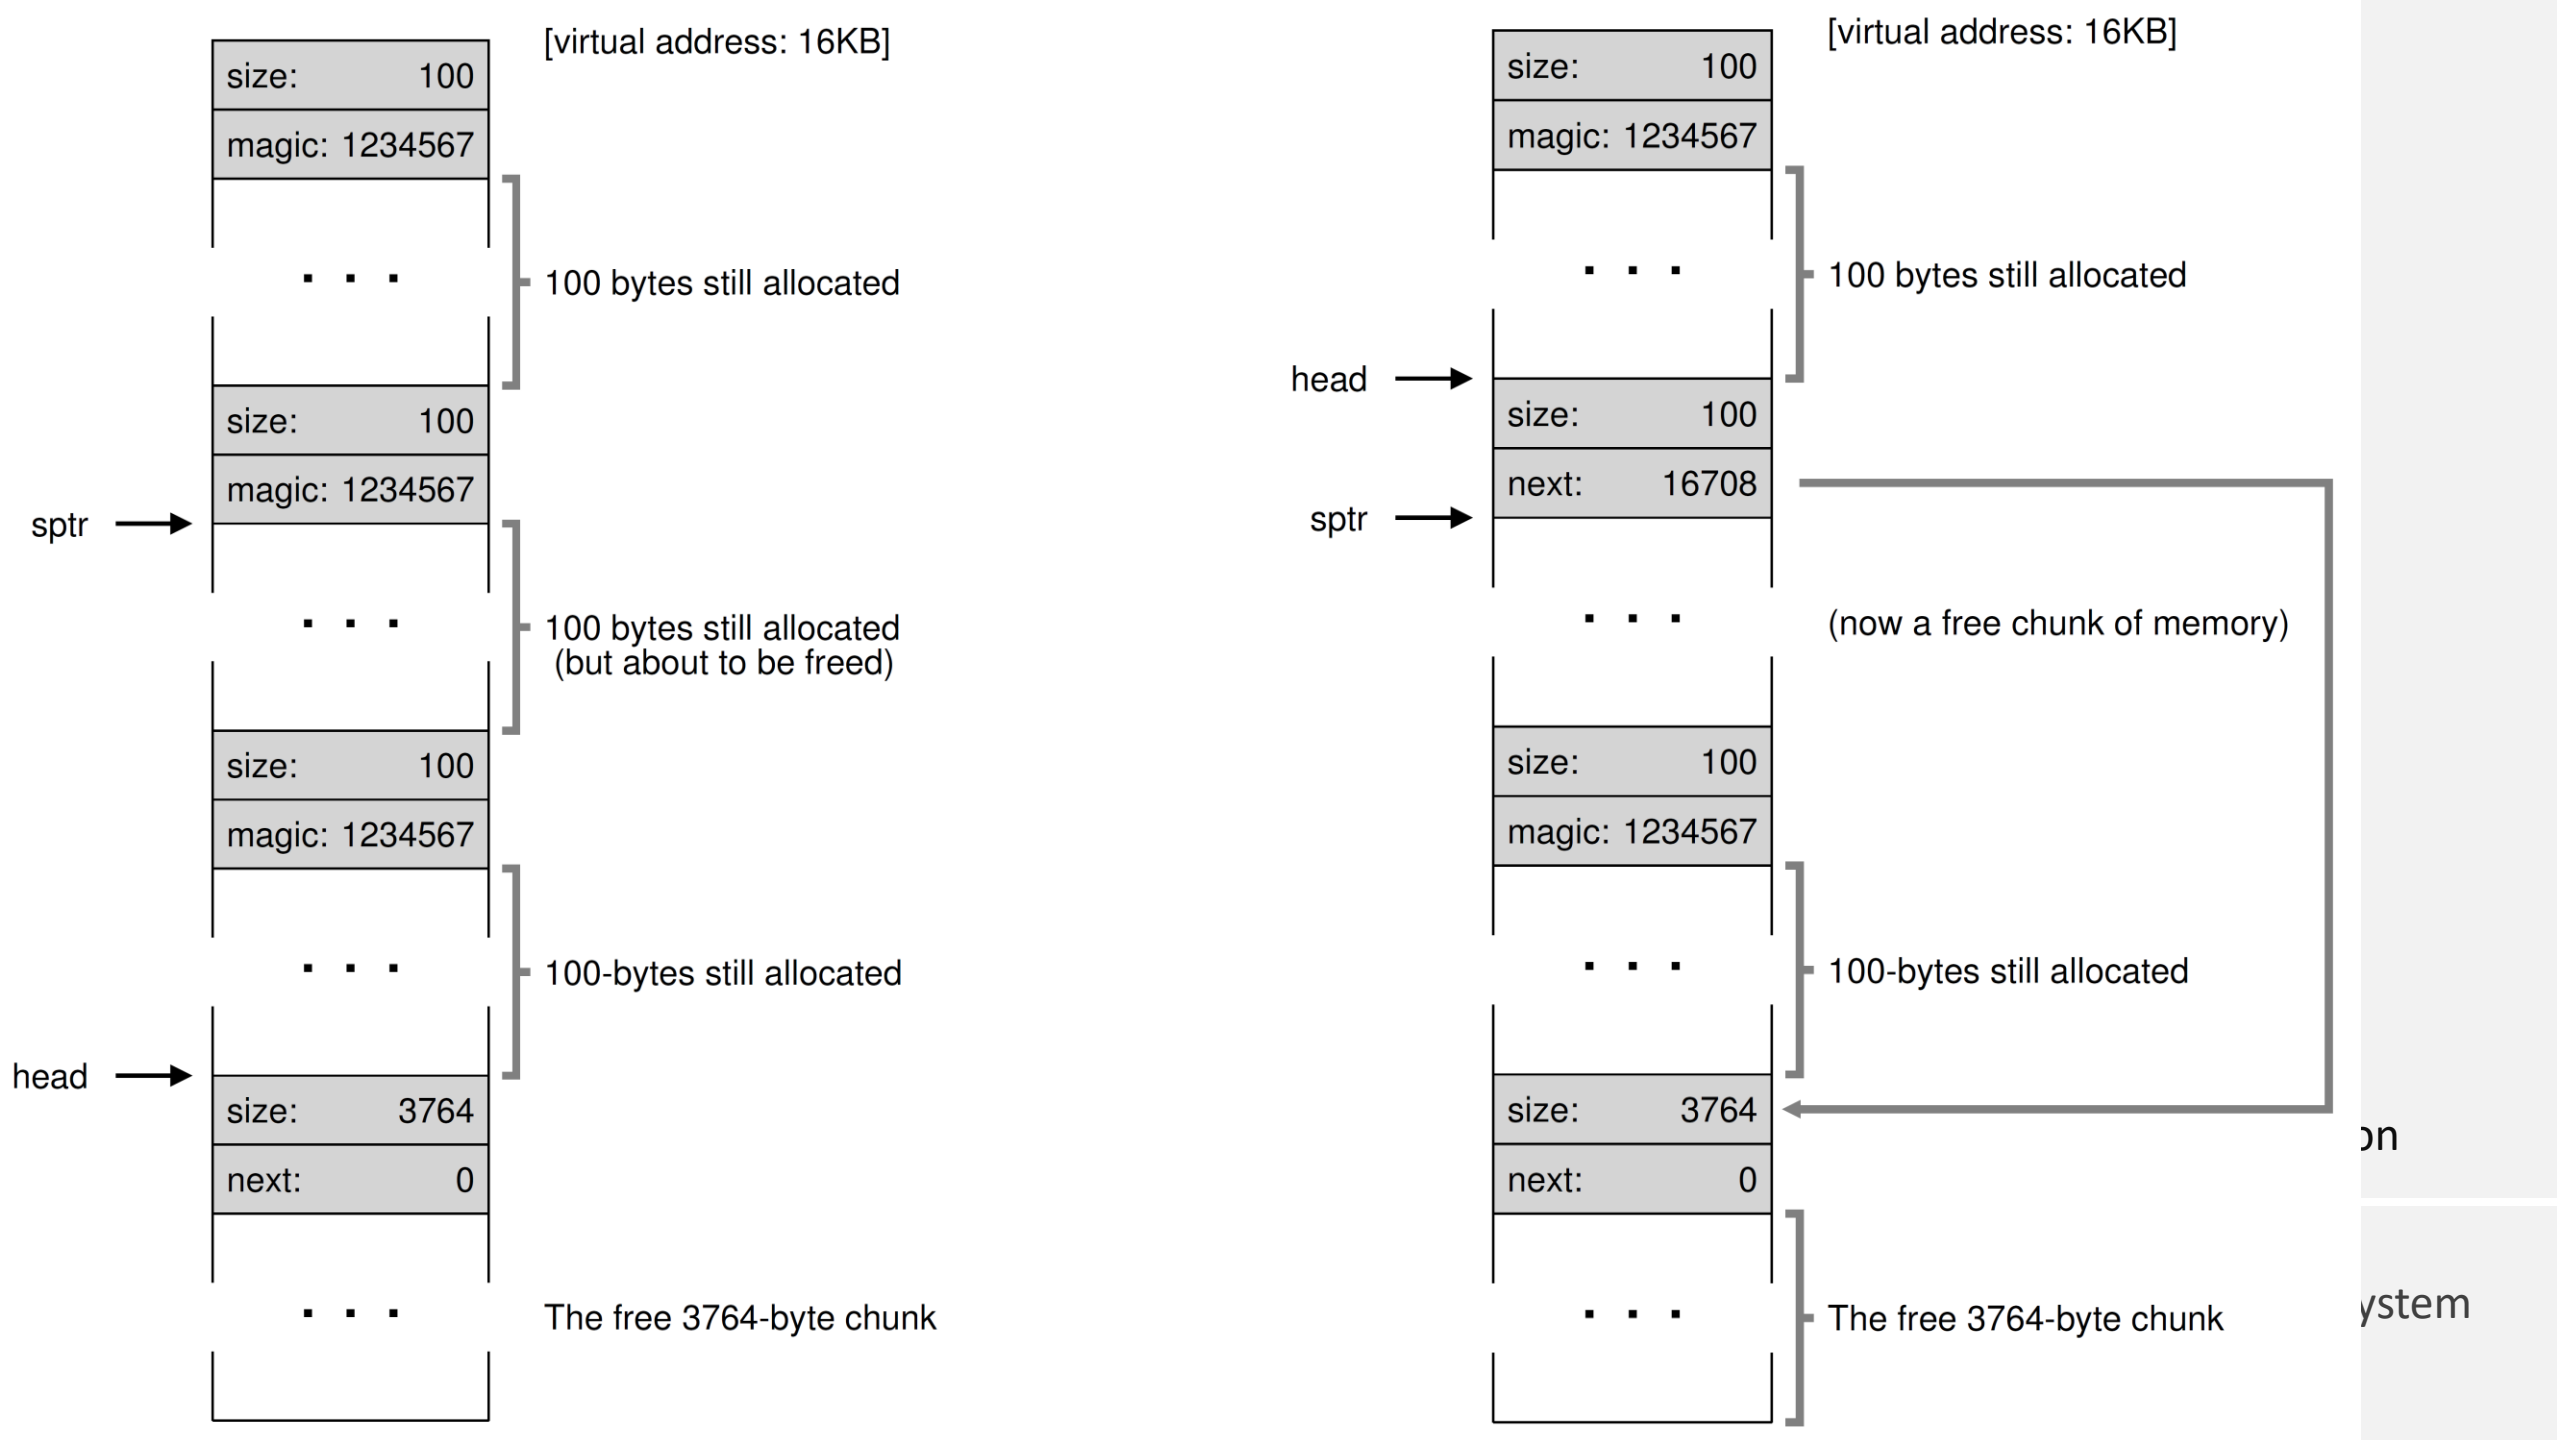
2023-04-06

Embedding A Free List

11

- Example: 4096-byte memory






Slab Allocator

- Alternate strategy
- **Slab** is one or more physically contiguous pages
- If slab is full of used objects, next object allocated from empty slab
 - If no empty slabs, new slab allocated
- Benefits include no fragmentation, fast memory request satisfaction

Slab Allocation



Buddy System

15

- Allocates memory from fixed-size segment consisting of physically-contiguous pages
- Memory allocated using **power-of-2 allocator**
 - Satisfies requests in units sized as power of 2
 - Request rounded up to next highest power of 2
 - When smaller allocation needed than is available, current chunk split into two buddies of next-lower power of 2
 - Continue until an appropriately-sized chunk is available
- E.g., assume 256KB chunk available, there is a request of 21 KB
 - Split into A_L and A_R of 128 KB each
 - One further divided into B_L and B_R of 64 KB
 - One further into C_L and C_R of 32 KB each – one used to satisfy request
- Advantage – quickly **coalesce** unused chunks into larger chunk
- Disadvantage - fragmentation

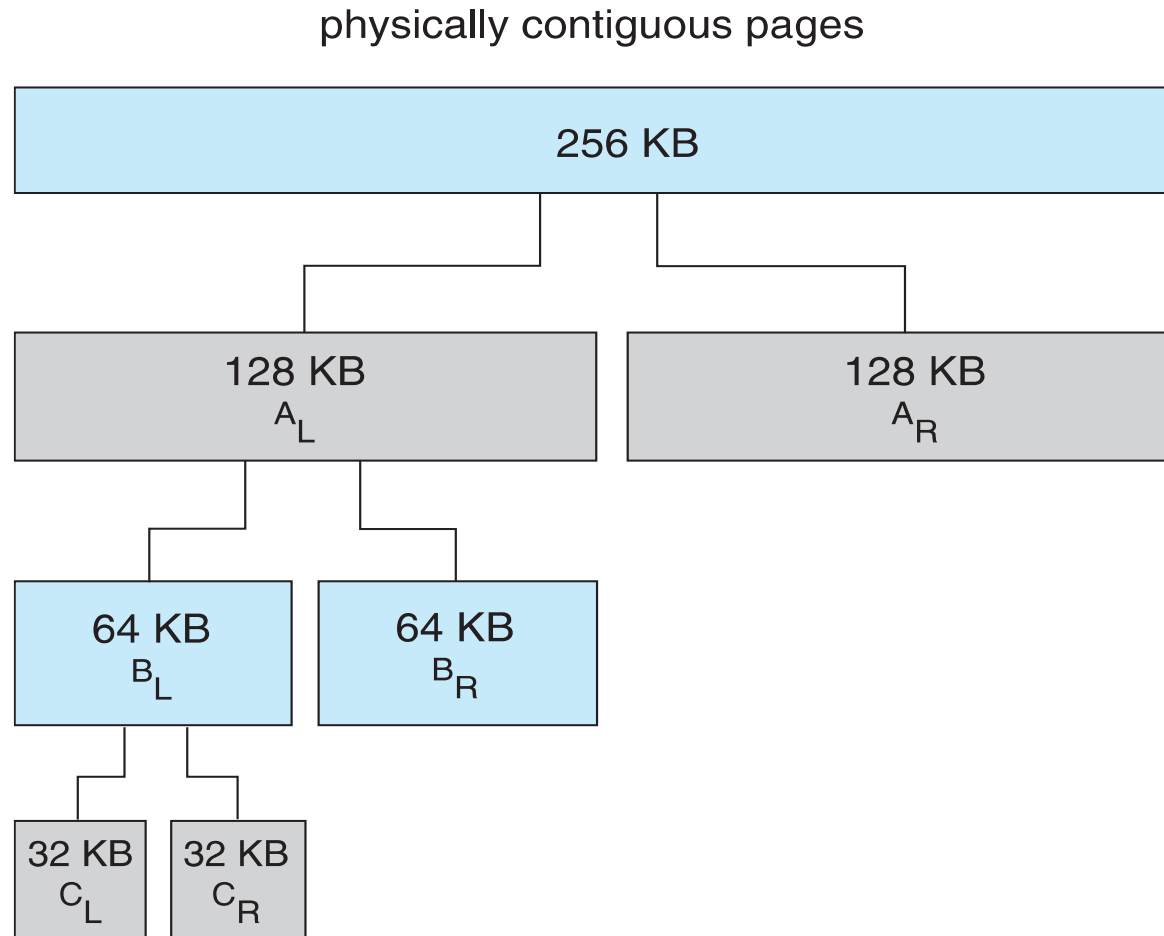
Segmentation

ITP 30002
Operating System

2023-04-06

Buddy System Allocator

16



Segmentation

ITP 30002
Operating System

2023-04-06