

Memory Allocation for Kernel

- Kernel memory must be allocated from a separate pool, different from the pool used user processes
- Most kernel data structures must stay **resident** in RAM all the time (should not be paged out)
- Most kernel data structures must be allocated **contiguously** in RAM (must be accessed without going through the virtual memory interface)
 - Pages = frames

Strategies for Kernel Mem Allocation

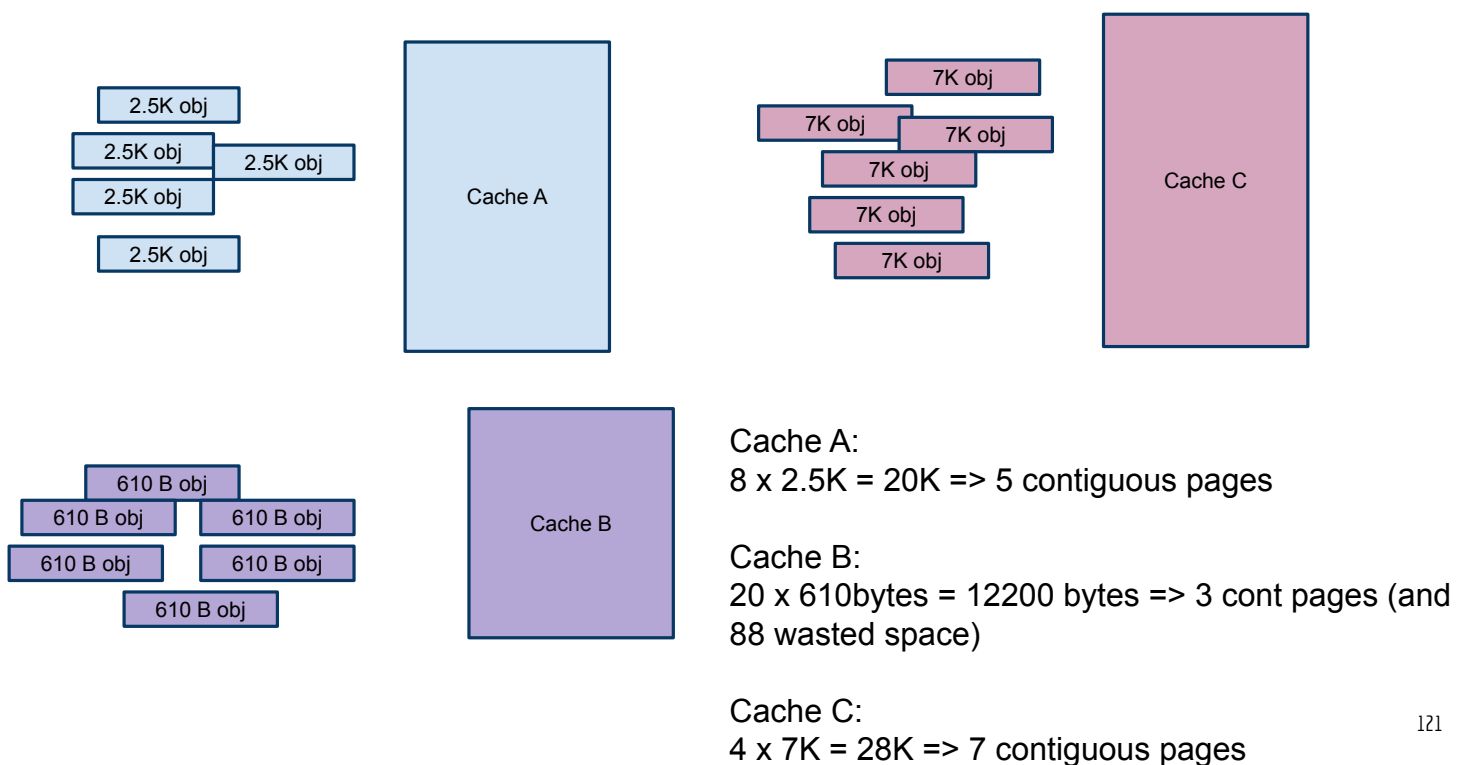
- Buddy System (power-of-two allocator)
 - Always allocate to the nearest (larger) **power of two** memory size
 - Request 15 MB => Allocate 16MB
 - Request 17MB => Allocate 32MB (too much wasted space)
 - May suffer large internal fragmentation
- Slab Allocation (better than Buddy)

Slab States

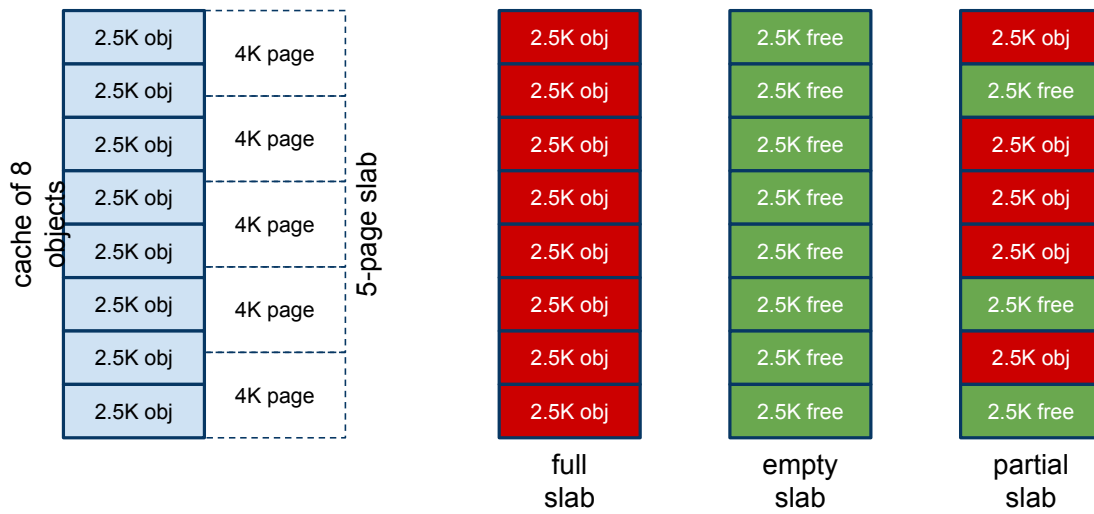
- Full: all of the kernel objects within the slab are assigned,
- Empty: all of the kernel objects are free
- Partial: some free, some assigned

Slab Allocation: One Cache per Kernel Object Type

Assume page size 4K



Slab Allocation into Contiguous Pages



$$8 \times 2.5K = 20K = 5 \times 4K$$

For the above example: cache grows/shrinks 5 pages at a time

OS Examples

Linux Page Replacement

- LRU 2Q (LRU clock approximation with two queues)
 - Inactive_list: pages which are not referenced (ref-bit ZERO)
 - Active_list: pages which are referenced (ref-bit ONE)
- Working Set
- Global replacement strategy

Windows Page Replacement

- Working Set
- LRU with local replacement strategy