Virtual Memory Preview

- RAM is broken up into blocks called pages.
- Address translation allows pages of physical memory to be mapped into a process' address space.

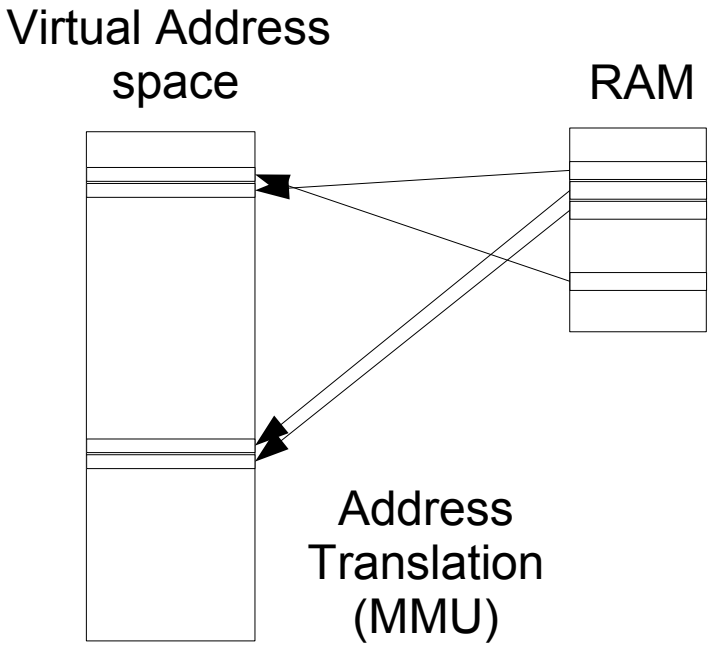

Cache/VM Interactions

- Caching was designed to take advantage of spatial and temporal locality, assuming physical memory addresses are used.
- What problems will this cause with virtual memory?

Cache/VM Interactions

- Which address to use?
- Virtual Address
  - On a process switch(virtual address space switch), caches need to be flushed because they no longer contain the right data.
- Physical Address
  - Can break spatial locality if non-contiguous memory pages are mapped to contiguous ranges of virtual memory
  - Even worse if virtually contiguous pages conflict in the cache

Solution?

- Physical Addresses + smart allocation
  - Allocate pages that will show up in contiguous cache lines to contiguous addresses in the virtual memory.
  - Example: 4 line, 4k per line cache
    - If we allocate virtual addresses 4k->7k (line 1, 0 based)to physical mem 16k->19k, for virtual addresses 8k->11k we would try to allocate 4k-7k or 20k->23k or 36k->39k .....
    - definitely avoid allocating 0k->3k, 32k->35k, etc. to virtual addresses 8k->11k
  - This is called page or cache coloring and is done by the OS when allocating memory to a process.