Slide 1

## Administrivia

- Homework 4 on the Web. Due next Friday.

Slide 2

## Minute Essay From Last Lecture

- Answer I had in mind can be gleaned from the reading (and many people did, or figured it out).
- There were some mentions that keeping the table in registers needed more memory. ?? (In context “memory” means RAM and is distinct from registers.)
- There were also mentions that keeping the table in registers seemed like it might let one process interfere with another. But remember that each process has a “virtual CPU” with registers.

Slide 3

## Paging — Review/Recap

- Simple schemes for memory management, in which each process's address space is mapped to a single single contiguous block of physical memory, are simple but not very flexible. Paging is one way to do better.
- Idea — divide both address spaces and memory into fixed-size blocks ("pages" and "page frames"), allow non-contiguous allocation.
- Consider tradeoffs yet again — complexity versus flexibility, efficient use of memory.

Slide 4

## Page Sizes and Other Details

- How big to make pages? compare extreme cases (really big, really small).
- If you know how big addresses are, what does that tell you about (maximum) sizes of physical/virtual memory?
- How big are page tables . . .

Slide 5

## Page Table Size — Example

- Given a page size of 64K ($2^{16}$), 64-bit addresses, and 4G ($2^{32}$) of main memory, at least how much space is required for a page table? Assume that you want to allow each process to have the maximum address space possible with 64-bit addresses, i.e., $2^{64}$ bytes.
- (Hints: How many entries? How much space for each one? and no, this is not a very realistic system.)

Slide 6

## Page Table Size — Example Continued

- Number of entries is $2^{64}/2^{16}$, i.e., $2^{48}$.
- Size of each entry — at least enough for page frame number. There are $2^{16}$ of them, so we need 16 bits for that. Probably should also include a valid/invalid bit, for a total of 17 bits. Rounding up to a multiple of 8 bits (one byte), that's 3 bytes per entry.
- Total space is $2^{48} \times 3$ — bigger than main memory!! so, not realistic.

Slide 7

## Performance / Feasibility Concerns

- Clearly page tables can be impractically large. What to do?
- Also, every memory reference involves page-table access — how to make this feasible/fast?

Slide 8

## Sidebar: Memory Management Within Processes

- What if we don't know before the program starts how much memory it will want? with very old languages, maybe not an issue, but with more modern ones it is.

  I.e., we might want to manage memory within a process's "address space" (range of possible program/virtual addresses).
- Typical scheme involves
  - Fixed-size allocation for code and any static data.
  - Two variable-size pieces ("heap" and "stack") for dynamically allocated data.
  - Notice — combined sizes of these pieces might be less than size of address space, maybe a lot less.

Slide 9

## Paging, Continued — Performance / Large Address Spaces

- Even with good choice of page size, serious performance implications — page table can still be big, and every memory reference involves page-table access — how to make this feasible/fast?
- (Remember that the MMU is hardware, and a bit about registers — local to the CPU, faster to access than memory but limited in number, can be general-purpose or dedicated to a particular use (e.g., the program counter).)

Slide 10

## Page Tables — Performance Issues

- One possibility is to keep the whole page table for the current process in registers. Could possibly use general-purpose registers for this but likely would not. Should make for fast translation of addresses, but — is this really feasible for a large table? and what about context switches?
- Another possibility is to keep the process table in memory and just have one register (probably a special-purpose one) point to it. Cost/benefit tradeoffs here seem like the opposite of the first scheme, no?

  The big downside is slow lookup, though, and that can be improved with a "translation lookaside buffer" (TLB) — special-purpose cache.

Slide 11

## Large Address Spaces

- Clearly page tables can be big. How to make this feasible?
- One approach — multilevel page tables.
- Another approach — inverted page tables (one entry per page frame).

Slide 12

## Paging and Virtual Memory

- Idea — if we don't have room for all pages of all processes in main memory, keep some on disk ("pretend we have more memory than we really do").
- Or a simpler view: All address spaces live in secondary memory / swap space / backing store, and we "page in" as needed (demand paging).
- Making this work requires help from both hardware (MMU) and software (operating system).

Slide 13

## Page Fault Interrupts

- We said MMU should generate a "page fault" interrupt for a page that's not present in real memory. What happens then? It's an interrupt, so . . .
- Control goes to an interrupt handler. What should it do? (Are there different possibilities for what caused the page faults?)

Slide 14

## Page Fault Interrupts, Continued

- One possible cause — an address that's not valid. You know (sort of) what happens then . . .
- Another cause — an address that's valid, but the page is on disk rather than in real memory. So — do I/O to read it in. Where to put it? If there's a free page frame, choice is easy. What if there's not?

Slide 15

## Finding A Free Frame — Page Replacement Algorithms

- Processing a page fault can involve finding a free page frame. Would be easy if the current set of processes aren't taking up all of main memory, but what if they are? Must steal a page frame from someone. How to choose one?
- Several ways to make choice (as with CPU scheduling) — "page replacement algorithms".
- "Good" algorithms are those that result in few page faults. (What happens if there are many page faults?)
- Choice usually constrained by what MMU provides (though that is influenced by what would help o/s designers).
- Many choices . . . (To be continued.)

Slide 16

## Minute Essay

- Why is a "good" page replacement algorithm one that generates as few page faults as possible? (I.e., what happens if there are a lot of page faults?)

## Minute Essay Answer

- The usual result of lots of page faults is that the computer spends more time doing “paging” (moving data back and forth between memory and disk), sometimes to the point where it isn’t doing much else.