

15213 Lecture 17: Virtual Memory Concepts

Learning Objectives

- Understand the distinction between virtual and physical addresses, and which are visible to processes.
- Describe page faults and give at least two situations where one might occur.
- Compare and contrast pages with cache lines.
- Identify what the OS does to programs that perform illegal memory accesses.

Getting Started

The directions for today's activity are on this sheet, but refer to accompanying programs that you'll need to download. To get set up, run these commands on a shark machine:

1. `$ wget http://www.cs.cmu.edu/~213/activities/lec17.tar`
2. `$ tar xf lec17.tar`
3. `$ cd lec17`

1 Memory Addresses: A Lie

Examine the `addr.c` file using your editor of choice or the `cat` command. Notice that the program prints example addresses from a number of its sections, then forks a child process that does the same. Once you're comfortable with the program, build (`$ make addr`) and run (`$./addr`) it.

1. What do you notice about the addresses printed by the two processes?
2. Do you think the processes share the same memory? Explain why this either must be or cannot be the case.
3. Now consider the `large.c` program, which performs a number of 1-GB memory allocations. Building and running the program, do you notice anything about its output?

2 Memory Addresses: Timings

We just saw that memory addresses are *remappable*, and that the system has been silently managing the mappings for us all along! Let's try to observe the performance overhead of this management. Examine the `timings.c` program, which allocates zero-initialized 100-KB memory regions using two different approaches, then performs a series of writes to each. When ready, build and run it.

4. Both the `calloc()` and `mmap()` library calls allocate a block of zero-initialized memory. Which *call* takes less time? `calloc()` `mmap()`

5. Which memory region exhibits faster initial *access* times? What does this suggest about how much setup work each of the allocation calls does?
6. Do you suspect that cache misses alone account for the access time difference? Why or why not?
7. Overall, which of the library interfaces was faster to use?

3 Virtual Memory: Page Faults

The memory addresses we've always worked with are known as **virtual memory addresses**, and usually differ from their corresponding **physical memory address** in DRAM. Any time the system attempts to translate a virtual address, there's a possibility that no valid mapping is present: such a situation is called a **page fault**, and handling it is one of the operating system's main jobs. One common OS response is to simply map the address to some unused portion of physical memory, then allow the program to retry; we'll focus on this for now and look at the other possibilities later.

Unix exposes counters that reveal how many page faults the OS has handled behind the scenes. Take a look at `faults.c`, a slightly modified version of the last program that uses a helper function from `benchmark.h` to report these counts. Build and run the program when you're ready.

8. Which allocation *call* results in more page faults? Which memory region incurs more page faults upon initial *access*? Now compare against the results from the previous section.

4 Virtual Memory: Anatomy of an Address

The difference between the two regions' page fault patterns is the result of an approach known as **demand paging**, whereby the OS defers mapping virtual addresses until they are first accessed (faulted). Whereas `mmap()` merely informs the OS that the locations may be accessed at some point, `calloc()` goes further by **prefaulting** all requested pages before it returns¹².

Much like the CPU cache, the virtual memory system breaks memory into a huge number of fixed-sized regions analogous to blocks; however, this time they are called **pages**. Also like cache, a portion of the address (the **page number**) determines which page it occupies, and a different portion (the **page offset**) describes the location within that page. Thanks to demand paging, we can reveal the system page size using `mmap()` by observing *when* within our traversal loop the page faults occur. Familiarize yourself with `bounds.c`, a program that does this, then build and run it.

¹Technically, `calloc()` performs an additional task before returning: The OS only zeros a page if a page fault actually occurs. However, because `calloc()` allocates from the process's heap, it sometimes reuses pages that the process has previously faulted (e.g., allocated then `free()`'d). Thus, `calloc()` must explicitly zero such pages.

²As an optimization, the GNU versions of `malloc()` and `calloc()` perform allocations larger than a certain size threshold (typically a little larger than the 100 KB we allocated) using `mmap()` instead of allocating heap space.

9. Looking at the output, how large is each page? How does this compare to the size of a cache line (64 B)?
10. Given the following depiction of a virtual memory address, draw two vertical lines to split it into the *page number*, *page offset*, and *unused bits*, labeling each part and the bits between.



5 Virtual Memory: Program Misbehavior

Of course, sometimes an application is wrong to attempt a certain memory access. For instance, consider the program `invalid.c`, which has a bug.

11. What happens when you run it? Which array index is the problem?
12. Thinking about the address of this array element, what distinguishes it from all the others? (*Hint: `mmap()` always returns page-aligned allocations.*)
13. As you've probably already experienced, the OS cannot always detect out-of-bounds memory accesses. In fact, there are at least two trivial changes to this program that would cause it to (apparently) run normally without actually fixing the bug; can you think of them?

6 Virtual Memory: Protection

The hardware stores some metadata for each mapped page. One important such entry is a **valid bit** to indicate whether that page is mapped; however, there are also **protection bits** that control which operations are valid on that page, useful for preventing some of the attacks you saw earlier in the course. Examine `protected.c`, a program that allocates some memory, reads from it, copies a function into it, and tries to execute the code from its new location. Try modifying the memory mapping by invoking it once with each of the arguments "", `r`, `rw`, and `rx` (one at a time).

14. Now that you've seen the reasons the OS might have to intervene in the middle of a memory access, complete this summary table by marking which categories of memory access (*not* allocation call) cause each of the listed outcomes. The first row has been done for you.

Valid access		Invalid access		Outcome
<code>calloc()</code> 'd	<code>mmap()</code> 'd	Unallocated page	Protection bits	
		x	x	Segfault Page fault No OS involvement OS maps page