

Learning Objectives

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

Getting Started

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

```
$ wget http://www.cs.cmu.edu/~213/activities/vm-concepts.tar
$ tar xf vm-concepts.tar
$ cd vm-concepts
```

1 Memory Addresses: A Lie

Examine the file `addr.c` using your editor of choice, or the `less` command.

Notice that this program prints the addresses of several different kinds of variables and one function. Then it uses the `fork` system call to create a new process, which also prints the addresses of all the same things. Finally it waits for the child process to finish, using the `waitpid` system call. If you aren't familiar with these functions, you can use the `man` command to read up on what they do: type `man fork` to display the documentation for `fork`. (This works for *any* C library function.)

Once you're comfortable with the program, compile it and run it:

```
$ make addr
$ ./addr
```

Problem 1. What do you notice about the addresses printed by the two processes?

VM Concepts

Problem 2. Do you think the processes share the same memory? Explain why this either must be or cannot be the case.

Now read the program `large.c`. It uses the `sysinfo` system call to report how much RAM the computer has, and then it uses the `mmap` system call to allocate as much of this memory as possible. It can, optionally, read from or write to each “page” of each allocation.

Problem 3. Compile this program (`make large`) and run it with no arguments (`./large`). Do you notice something odd happening?

Problem 4. Now run this program in “read” mode (`./large read`). What happens? How is that different from what happened before?

Problem 5. Finally, run this program in “write” mode (`./large write`). What happens? How is that different from what happened before?

Problem 6. What do you think is going on here?

2 Memory Addresses: Timings

We just saw that a memory address in one process usually refers to a *different byte of RAM* than the same memory address in a different process. This is called *virtual addressing*. It's an essential safety feature—it guarantees that one program cannot scribble on another program's memory!

We also just saw that the system calls that allocate RAM (`sbrk` and `mmap`) might be “lazy” and not actually allocate any memory until it is used. This means the operating system can tell when freshly-allocated memory is first used.

These features have some performance overhead. Let's try to measure this overhead. Examine the program `timings.c`, which allocates zero-initialized 100 KB memory regions using two different approaches, then performs a series of writes to each, and measures the time these operations take using a helper function from `benchmark.h`. When you are ready, build and run it.

Problem 7. Both `calloc` and `mmap` allocate a block of zero-initialized memory. Which *call* takes less time?

Problem 8. Which memory region is faster for the application to access, the *first* time it does this?

Problem 9. Which memory region is faster for the application to access, the *second* time it does this?

Problem 10. What do these things imply about the work being done by `calloc` versus `mmap`?

Problem 11. Could cache misses alone account for this time difference?

3 Virtual Memory: Page Faults

The memory addresses we've always worked with are known as *virtual addresses*. As we have seen, these are different from the *physical addresses* that actually index the RAM. The operating system maintains a mapping from virtual to physical addresses for each process.

This mapping does not have to be complete. Some virtual addresses might not correspond to any physical address. If a program tries to access a virtual address that has no corresponding physical address, the CPU will trigger a special mechanism called a *page fault* that allows the operating system to intervene. (We'll talk more about this mechanism on Thursday.) One of the things the OS can do when a page fault happens is *create* a mapping for the virtual address and then allow the program to retry the memory access.

Unix allows programs to monitor how many times page faults have happened, and new mappings have been created, because of their memory accesses.

Read `faults.c`, a slightly modified version of `timings.c`. It reports page faults instead of timing, using a different helper function from `benchmark.h`. Build and run the program when you're ready.

Problem 12. Which allocation *call* results in more page faults?

Problem 13. Which memory region incurs more page faults upon initial *access*?

Problem 14. Compare the numbers you get from `faults` with the numbers you get from `timings`. Do the different numbers of page faults explain the different timings?

4 Virtual Memory: Anatomy of an Address


The phenomenon we have just observed—page faults occurring when freshly allocated memory is accessed—is a deliberate operating system design decision known as *demand paging*. Operating systems that implement demand paging wait until freshly allocated memory is first accessed to define a virtual-to-physical mapping for it. This allows them to avoid defining mappings for any memory that is not actually used. Many programs (such as `large.c`) allocate far more memory than they use; demand paging means they only consume the physical RAM they really need.

We keep using the term “page.” A page is a fixed-size set of virtual addresses that are all mapped to physical addresses as a unit. This is similar to how the memory cache divides memory into blocks. Just like how addresses are divided into cache tag, set index, and block offset, they are also divided into a *page number* and a *page offset*.¹ The page number says which page the address belongs to, and therefore which region of physical RAM it is mapped to. The page offset indicates a specific byte within that page.

Because the operating system on the sharks uses demand paging, if we use `mmap` to allocate some memory, and then access it byte by byte and watch for page faults, we can discover the size of a page, and therefore how many bits of an address are the page offset. Familiarize yourself with `bounds.c`, a program that does this, then build and run it.

Problem 15. Looking at the output, how large is each page? How does this compare to the size of a cache line (64 B)?

Problem 16. Below is a diagram of the 64 bits of a virtual memory address. Bit 0 is the least significant bit. Label this diagram to show which bits are part of the *page offset* and which are part of the *page number*.



¹The virtual memory system does not have any equivalent of a set index. As we will see later on, this is because the virtual memory system can be thought of as a *fully associative* cache.

5 Virtual Memory: Program Misbehavior

Sometimes programs access memory addresses they shouldn't. For instance, consider the program `invalid.c`, which has a bug. Compile it and run it—but run it under `gdb`.

Problem 17. What happens when you run it? Which array index is the problem?

Problem 18. What is special about the address of this array (non-)element? (*Hint: `mmap` always returns page-aligned allocations.*)

Problem 19. As you've probably already experienced, the OS cannot always detect out-of-bounds memory accesses. There are at least two ways you could change this program that would make it *appear* to run normally but would not actually fix the bug. Can you think of them?

6 Virtual Memory: Protection

Even if a page of memory is mapped, the program might not be allowed to do just anything to that memory. The hardware maintains *protection bits* for each mapped page which define what operations are allowed on that page. You may have noticed that all the example programs so far passed an argument `PROT_READ|PROT_WRITE` to `mmap`—this asks for memory that is both readable and writable. There are several other possibilities.

Examine the program `protected.c`. It allocates some memory, reads from it, copies a function into it, and tries to execute the code from its new location. You can control how it allocates the memory with command line arguments. Try running it with each of these arguments (one at a time): `""`, `r`, `rw`, and `rx`.

Problem 20. 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