



Virtual Memory

Process Abstraction, Part 2: Private Address Space

Motivation: why not direct physical memory access?


Address translation with pages

Extra benefits: sharing and protection

Memory as a contiguous array of bytes is a lie! Why?


<https://cs.wellesley.edu/~cs240/>

Problems with physical addressing



2

Problem 1: memory management



Also:
Context switches must swap out entire memory contents.
Isn't that expensive?

Problem 2: capacity

64-bit addresses can address several exabytes
(18,446,744,073,709,551,616 bytes)



1 virtual address space per process,
with many processes...

Physical main memory offers ~a few dozen gigabytes (e.g. 8,589,934,592 bytes)

(To scale with 64-bit address space, you can't see it!)



3

4

What does this code print? Why/how?

```
#include <stdio.h>
#include <sys/types.h>
#include <unistd.h>

int main() {
    int s;
    int x = 1;
    int *p;
    int child_pid = fork();
    p = &x;
    if (child_pid == 0) {
        *p = 2;
    } else {
        *p = 3;
    }
    printf("Address 0x%lx holds %d\n", (long)p, *p);
}
```

5

Which is a possible output of this program?

```
#include <stdio.h>
#include <sys/types.h>
#include <unistd.h>

int main() {
    int s;
    int x = 1;
    int *p;
    int child_pid = fork();
    p = &x;
    if (child_pid == 0) {
        *p = 2;
    } else {
        *p = 3;
    }
    printf("Address 0x%lx holds %d\n", (long)p, *p);
}
```

Addr 0x10 holds 2; Addr 0x10 holds 2

Addr 0x10 holds 3; Addr 0x10 holds 3

Addr 0x10 holds 3; Addr 0x10 holds 2

Addr 0x10 holds 3; Addr 0x20 holds 3

Addr 0x10 holds 3; Addr 0x20 holds 2

None of the above

Start the presentation to see live content. For screen share software, share the entire screen. Get help at pollev.com/app


What does this code print? Why/how?

```
#include <stdio.h>
#include <sys/types.h>
#include <unistd.h>


int main() {
    int s;
    int x = 1;
    int *p;
    int child_pid = fork();
    p = &x;
    if (child_pid == 0) {
        *p = 2;
    } else {
        *p = 3;
    }
    printf("Address 0x%lx holds %d\n", (long)p, *p);
}
```

7

Problem 3: protection




Problem 4: sharing




8

Solution: Virtual Memory (address indirection)




Indirection (it's everywhere!)

Direct naming



Indirect naming




What if we move *Thing*?

Tangent: indirection everywhere


- Pointers
- Constants
- Procedural abstraction
- Domain Name Service (DNS)
- Dynamic Host Configuration Protocol (DHCP)
- Phone numbers
- 911
- Call centers
- Snail mail forwarding
- ...

"Any problem in computer science can be solved by adding another level of indirection."
—David Wheeler, inventor of the subroutine, or Butler Lampson


Another Wheeler quote? "Compatibility means deliberately repeating other people's mistakes." 11




Virtual addressing and address translation



Page-based mapping




Some virtual pages do not fit! Where are they stored?




14

Virtual memory: cache for disk?




Example system

Address translation




16

Page table



Virtual memory benefits:
Simple address space allocation

Process needs private *contiguous* address space.



18

Virtual memory benefits:


Protection:

All accesses go through translation.


Impossible to access physical memory not mapped in virtual address space.

Sharing:

Map virtual pages in separate address spaces to same physical page (PP 6).



Virtual memory benefits:
Memory permissions



20

19

Summary: virtual memory

Programmer's view of virtual memory

Each process has its own private linear address space
Cannot be corrupted by other processes

System view of virtual memory

Uses memory efficiently (due to locality) by caching virtual memory pages
Simplifies memory management and sharing
Simplifies protection -- easy to interpose and check permissions
More goodies:

- Memory-mapped files
- Cheap `fork()` with copy-on-write pages (COW)

