CSCI 2021: Virtual Memory

Chris Kauffman

Last Updated: Mon Apr 25 01:06:15 PM CDT 2022

Logistics

${\sf Reading \ Bryant}/O'Hallaron$

- Ch 9: Virtual Memory
- Ch 7: Linking (next)

P5

- 1 Problem
- Implement a small version of malloc() / free()
- Post later today with video, due last day of class

Goals

- Address Spaces, Translation, Paged Memory
- mmap(), Sharing Pages

Date	Event
Fri 4/22	Virtual Mem 1/2
Mon 4/25	Virtual Mem 2/2
Wed 4/27	ELF Files/Linking 1/2 Lab 14 mmap() HW 14 Linking
Fri 4/29	Obj Code/Linking 2/2
Mon 5/2	Last Lecture, Review
	P5 Due

Exercise: The View of Memory Addresses so Far

- Every process (running program) has some memory, divided into roughly 4 areas (which are...?)
- Reference different data/variables through their addresses
- If only a single program could run at time, no trouble: load program into memory and go
- Running multiple programs gets interesting particularly if they both reference the same memory location, e.g. address 1024
 PROGRAM 1 PROGRAM 2

```
## load global from #1024 ## add to global at #1024
movq 1024, %rax addl %esi, 1024
```

. . .

- What conflict exists between these programs?
- What are possible solutions to this conflict?

. . .

. . .

Answers: The View of Memory Addresses so Far

- 4 areas of memory are roughly: (1) Stack (2) Heap (3) Globals (4) Text/Instructions
- Both programs use physical address #1024, behavior depends on order that instructions are interleaved between them

ORDER A: Program 2	1 loads first	ORDER B: Program 2	2 adds first
PROGRAM 1	PROGRAM 2	PROGRAM 1	PROGRAM 2
movq 1024, %rax			addl %esi, 1024
	addl %esi, 1024	movq 1024, %rax	•••

- Solution 1: Never let Programs 1 and 2 run together (bleck!)
- Solution 2: Translate every memory address in every program on loading it, run with physical addresses
 - Tough/impossible as not all addresses are known at compile/load time...
- Solution 3: Translate every memory address/access in every program while it runs (!!!)

Paged Memory

- Physical memory is divided into hunks called pages
- Common page size supported by many OS's (Linux) and hardware is 4KB = 4096 bytes, can be larger with OS config
- CPU models use some # of bits for Virtual Addresses

```
> cat /proc/cpuinfo
vendor_id : GenuineIntel
cpu family : 6
model : 79
model name : Intel(R) Xeon(R) CPU E5-1620 v4 @ 3.50GHz
...
address sizes : 46 bits physical, 48 bits virtual
```

Example of address with page number and offset labelled

xxxxPagenumbrOff : 48 bits used 0x00007ffa0997a428 : 64 bit address

| | +-> Offset 0x428 within page, 12 bits | +-> Page number 0x7ffa0997a, 36 bits +-> Constant bits, not used by processor

Translation happens at the Page Level

- Within a page, addresses are sequential
- Between pages, may be non-sequential

 Page Table:

 |------

 | Virtual Page
 | Size | Physical Page

 |-----

 | 00007ffa0997a000 | 4K | RAM: 0000564955aa1000 |

 | 00007ffa0997b000 | 4K | RAM: 0000321e46937000 |

 | ...

 | ...

Address Space From Page Table:

Virtual Address	Page Offset	Physical Address
00007ffa0997a000 00007ffa0997a001 00007ffa0997a002 00007ffa0997afff	0 1 2 4095	0000564955aa1000 0000564955aa1001 0000564955aa1002 0000564955aa1fff
00007ffa0997b000 00007ffa0997b001 	0 1	0000321e46937000 0000321e46937001

Addresses Translation Hardware

- Translation must be FAST so usually involves hardware
- MMU (Memory Manager Unit) is a hardware element specifically designed for address translation
- Usually contains a special cache, TLB (Translation Lookaside Buffer), which stores recently translated addresses



- OS Kernel interacts with MMU
- Provides location of the Page Table, data structure relating Virtual/Physical Addresses
- Page Fault : MMU couldn't map Virtual to Physical page, runs a Kernel routine to handle the fault

Exercise: Translating Virtual Addresses

Nearby diagram illustrates relation of Virtual Pages to Physical Pages

- 1. How many page tables are there?
- 2. Where can a page table entry refer to?
- 3. Count the number of Virtual pages, compare to the number of physical pages which his larger?
- 4. What happens if PID #123 accesses its Virtual Page #2
- 5. What happens if PID #456 accesses its Virtual Page #2



Translating Virtual Addresses 1/2

- On using a Virtual Memory address, MMU will search TLB for physical DRAM address,
- If found in TLB, Hit, use physical DRAM address
- If not found, MMU will search Page Table, if found and in DRAM, cache in TLB
- Else Miss = Page fault, OS decides..
 - Page is swapped to Disk, move to DRAM, potentially evicting another page
 - 2. Page not in page table = Segmentation Fault



Translating Virtual Addresses 2/2

- Each process has its own page table, OS maintains mapping of Virtual to Physical addresses
- Processes "compete" for RAM
- OS gives each process impression it owns all of RAM
- OS may not have enough memory to back up all or even 1 process
- Disk used to supplement ram as Swap Space
- Thrashing may occur when too many processes want too much RAM, "constantly swapping"



Trade-offs of Address Translation

Wins of Virtual Memory

- 1. Avoids processes each referencing the same address, conflicting
- Allows each Process (running program) to believe it has entire memory to itself
- 3. Gives OS tons of flexibility and control over memory layout
 - Present a continuous Virtual chunk which is spread out in Physical memory
 - Use Disk Space as memory
 - Check for out of bounds memory references

Losses of Virtual Memory

- Address translation is not constant O(1), has an impact on performance of real algorithms*
- 2. Requires special hardware to make translation fast enough: MMU/TLB
- Not needed if only a single program is running on a machine

Wins outweigh Losses in most systems so Virtual Memory is used widely, a *great idea* in CS

The Many Other Advantages of Virtual Memory

- 1. Caching: Seen that VirtMem can treat main memory as a cache for larger memory
- 2. Security: Translation allows OS to check memory addresses for validity, segfault on out-of bounds access
- 3. Debugging: Valgrind checks addresses for validity
- 4. Sharing Data: Processes can share data with one another; request OS to map virtual addresses to same physical addresses
- 5. **Sharing Libraries**: Can share same program text between programs by mapping address space to same shared library
- Convenient I/O: Map internal OS data structures for files to virtual addresses to make working with files free of read()/write()

Virtual Memory and mmap()

- Normally programs interact indirectly with Virtual Memory system
 - Stack/Heap/Globals/Text are mapped automatically to regions in Virtual Memory System
 - Maps are adjusted as Stack/Heap Grow/Shrink
- mmap() / munmap() directly manipulate page tables
 - mmap() creates new entries in page table
 - munmap() deletes entries in the page table
 - Can map arbitrary or specific addresses into memory
- mmap() is used to initially set up Stack / Heap / Globals / Text when a program is loaded by the program loader
- While a program is running can also use mmap() to interact with virtual memory
- A convenient way to do File I/O via Memory Mapped Files

Exercise: Printing Contents of file

Examine the two programs below which print the contents of a file

- Identify differences between them
- Which has a higher memory requirement?

```
// print_file.c
                                               // mmap_print_file.c
1
                                            1
   int main(int argc, char *argv[]){
                                            2 int main(int argc, char *argv[]){
2
      FILE *fin = fopen(argv[1], "r");
3
                                            3
                                                 int fd = open(argv[1], O_RDONLY);
4
      char inbuf[256]:
                                            4
5
      while(1){
                                            5
                                                 struct stat stat buf:
6
        int nread =
                                            6
                                                 fstat(fd. &stat buf):
7
          fread(inbuf, sizeof(char),
                                            7
                                                 int size = stat buf.st size:
8
                                            8
                256. fin):
9
        if(nread == 0){
                                            9
                                                 char *file chars =
10
                                           10
                                                   mmap(NULL, size,
          break:
        }
11
                                           11
                                                        PROT READ, MAP SHARED,
12
        fwrite(inbuf,sizeof(char),
                                           12
                                                        fd, 0);
13
               nread.stdout);
                                           13
14
      }
                                           14
                                                 for(int i=0; i<size; i++){</pre>
15
                                           15
                                                   printf("%c",file chars[i]);
16
      fclose(fin):
                                           16
                                                 r
17
      return 0;
                                           17
                                                 printf("\n");
18
   }
                                           18
                                           19
                                                 munmap(file_chars, size);
                                           20
                                                 close(fd):
                                           21
                                                 return 0;
```

22 }

Answers: Printing Contents of file

- 1. Write a simple program to print all characters in a file. What are key features of this program?
 - Open file
 - Read up to 256 characters into memory using fread()/fscanf()
 - Print those characters with printf()
 - Read more characters and print
 - Stop when end of file is reached
 - Close file
- Examine mmap_print_file.c: does it contain all of these key features? Which ones are missing?
 - Missing the fread()/fscanf() portion
 - Uses mmap() to get direct access to the bytes of the file
 - Treat bytes as an array of characters and print them directly

mmap(): Mapping Addresses is Amazing

- ptr = mmap(NULL, size,...,fd,0) arranges backing entity of fd to be mapped to be mapped to ptr
- fd often a file opened with open() system call

```
int fd = open("gettysburg.txt", O_RDONLY);
// open file to get file descriptor
```

```
printf("%c",file_chars[0]); // print 0th file char
printf("%c",file_chars[5]); // print 5th file char
```

OS usually Caches Files in RAM

- ► For efficiency, part of files are stored in RAM by the OS
- OS manages internal data structures to track which parts of a file are in RAM, whether they need to be written to disk
- mmap() alters a process Page Table to translate addresses to the cached file page
- OS tracks whether page is changed, either by file write or mmap() manipulation
- Automatically writes back to disk when needed
- Changes by one process to cached file page will be seen by other processes
- See diagram on next slide

Diagram of Kernel Structures for mmap()



Changing Files

- > mmap() exposes several capabilities from the OS
 char *file_chars =
 mmap(NULL, size,
 PROT_READ | PROT_WRITE, // map allowing read + write
 MAP_SHARED, // share changes with original file
 fd, 0); // file to map + offset from start
- Assign new value to memory, OS writes changes into the file
- Example: mmap_tr.c to transform one character to another

Mapping things that aren't characters

mmap() just gives a pointer: can assert type of what it points at

- Example int *: treat file as array of binary ints
- Notice changing array will write to file

// mmap_increment.c: demonstrate working with mmap()'d binary data

```
int fd = open("binary_nums.dat", O_RDWR);
// open file descriptor, like a FILE *
```

```
int *file_ints = mmap(NULL, size, PROT_READ | PROT_WRITE, MAP_SHARED, fd, 0);
// get pointer to file bytes through mmap,
// treat as array of binary ints
```

```
int len = size / sizeof(int);
// how many ints in file
```

```
for(int i=0; i<len; i++){
    printf("%d\n",file_ints[i]); // print all ints
}</pre>
```

```
for(int i=0; i<len; i++){
   file_ints[i] += 1; // increment each file int, writes back to disk
}</pre>
```

mmap() Compared to Traditional fread()/fwrite() I/O

Advantages of mmap()

- Avoid following cycle
 - fread()/fscanf() file contents into memory
 - Analyze/Change data
 - fwrite()/fscanf() write memory back into file
- Saves memory and time
- Many Linux mechanisms backed by mmap() like processes sharing memory

Drawbacks of mmap()

- Always maps pages of memory: multiple of 4096b (4K)
- ▶ For small maps, lots of wasted space
- Cannot change size of files with mmap(): must used fwrite() to extend or other calls to shrink
- No bounds checking, just like everything else in C

Page Table Size

- Page tables map a virtual page to physical location
- Page tables maintained by operating system in Kernel Memory
- A direct page table has one entry per virtual page
- Each page is 4K = 2¹² bytes, so 12 bits for offset of address into a page
- Virtual Address Space is 2⁴⁸ bytes
- So, 2³⁶ virtual pages mapped in the page table...
 - 68,719,476,736 pages
 - At 8 bytes per page entry...
 - 1 Terabyte for a page table



How big does the page table mapping virtual to physical pages need to be?

Page "Tables" are Multi-Level Sparse Trees



"What Every Programmer Should Know About Memory" by Ulrich Drepper, Red Hat, Inc.

- Fix this absurdity with **multi-level page tables**: a sparse tree
- Virtual address divided into sections which indicate which PTE to access at different table levels
- 3-4 level page table is common in modern architectures
- Programs typically use only small amounts of virtual memory: most entries in different levels are NULL (not mapped) leading to much smaller page tables than a direct (array) map

Direct Page Table vs Sparse Tree Page Table



Textbook Example: Two-level Page Table

Space savings gained via NULL portions of the page table/tree



Pages and Mapping

- Memory is segmented into hunks called pages, 4Kb is common (use page-size.c to see your system's page size)
- OS maintains tables of which pages of memory exist in RAM, which are on disk
- OS maintains tables per process that translate process virtual addresses to physical pages
- Shared Memory can be arranged by mapping virtual addresses for two processes to the same memory page



Shared Libraries: *.so Files

 Code for libraries can be shared

> libc.so: shared library with malloc(), printf() etc in it

 OS puts into one page, maps all linked procs to it



Source: John T. Bell Operating Systems Course Notes

pmap: show virtual address space of running process

- > ./memory parts 0x5575555a71e9 : main() 0x5575555aa0c0 : global arr 0x557555b482a0 : heap arr 0x60000000000 : mmap'd block1 0x60000001000 : mmap'd block2 0x7f2244dc4000 : mmap'd file 0x7ffff0133b70 : stack arr my pid is 496605 press any key to continue
 - Determine process id of running program
 - pmap reports its virtual address space
 - More details of pmap output in this article from Andreas Fester
 - His diagram is awesome

> pmap 496605 ./memory parts 496605: 00005575555a6000 00005575555a7000 00005575555a8000 00005575555a9000 00005575555aa000 00005575555ab000 0000557555b48000 00006000000000000 00007f2244bca000 00007f2244bcc000 00007f2244d3f000 00007f2244d8e000 00007f2244d91000 00007f2244dc4000 00007f2244dc5000 00007f2244dc7000 00007f2244de8000 00007f2244df2000 00007ffff0114000 00007ffff014d000 total 2352K

4K r---- memory parts 4K r-x-- memory parts TEXT 4K r---- memory parts 4K r---- memory parts 4K rw--- memory parts GLOBALS 4K rw---[anon] [anon] HEAP 132K rw---8K rw---[anon] 8K rw---[anon] 152K r---- libc-2.32.so 00007f2244bf2000 1332K r-x-- libc-2.32.so 304K r---- libc-2.32.so 12K rw--- libc-2.32.so 24K rw---[anon] 4K r---- gettysburg.txt 8K r---- 1d-2.32.so 132K r-x-- 1d-2.32.so 36K r---- 1d-2.32.so 8K rw--- 1d-2.32.so 132K rw---[stack] STACK 12K r----[anon]

Memory Protection

- Output of pmap indicates another feature of virtual memory: protection
- OS marks pages of memory with Read/Write/Execute/Share permissions like files
- Attempt to violate these and get segmentation violations (segfault)
- Ex: Executable page (instructions) usually marked as r-x: no write permission.
- Ensures program don't accidentally write over their instructions and change them
- Ex: By default, pages are not shared (no 's' permission) but can make it so with the right calls

Exercise: Quick Review

- 1. While running a program, memory address #1024 always refers to a physical location in DRAM (True/False: why?)
- 2. Two programs which both use the address #1024 cannot be simultaneously run (True/False: why?)
- 3. What do MMU and TLB stand for and what do they do?
- 4. What is a memory page? How big is it usually?
- 5. What is a Page Table and what is it good for?

Answers: Quick Review

- 1. While running a program, memory address #1024 always refers to a physical location in DRAM (True/False: why?)
 - False: #1024 is usually a virtual address which is translated by the OS/Hardware to a physical location which may be in DRAM but may instead be paged out to disk
- 2. Two programs which both use the address #1024 cannot be simultaneously run (True/False: why?)
 - False: The OS/Hardware will likely translate these identical virtual addresses to different physical locations so that the programs doe not clobber each other's data
- 3. What do MMU and TLB stand for and what do they do?
 - Memory Management Unit: a piece of hardware involved in translating Virtual Addresses to Physical Addresses/Locations
 - Translation Lookaside Buffer: a special cache used by the MMU to make address translation fast
- 4. What is a memory page? How big is it usually?

A discrete hunk of memory usually 4Kb (4096 bytes) big

- 5. What is a Page Table and what is it good for?
 - A table maintained by the operating system that is used to map Virtual Addresses to Physical addresses for each page

Additional Review Questions

- What OS data structure facilitates the Virtual Memory system? What kind of data structure is it?
- What does pmap do?
- What does the mmap() system call do that enables easier I/O? How does this look in a C program?
- Describe at least 3 benefits a Virtual Memory system provides to a computing system