

Section 8: Banker's Algorithm and Address Translation

March 8, 2019

Contents

1	Vocabulary	2
2	Problems	3
2.1	Conceptual Questions	3
2.2	Banker's Algorithm	5
2.3	Page Allocation	6
2.4	Address Translation	7
2.5	Inverted Page Tables	9

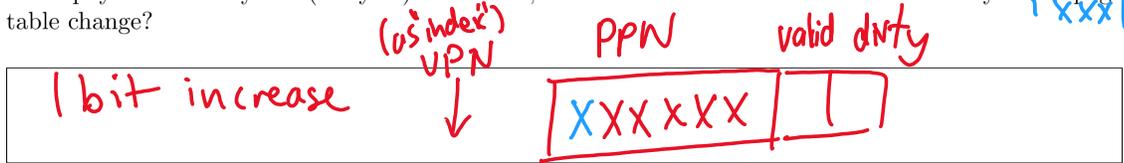
1 Vocabulary

- **Deadlock** - Situation in which two computer programs sharing the same resource are effectively preventing each other from accessing the resource, resulting in both programs ceasing to function.
- **Banker's Algorithm** - A resource allocation and deadlock avoidance algorithm that tests for safety by simulating the allocation for predetermined maximum possible amounts of all resources, before deciding whether allocation should be allowed to continue.
- **Virtual Memory** - Virtual Memory is a memory management technique in which every process operates in its own address space, under the assumption that it has the entire address space to itself. A virtual address requires translation into a physical address to actually access the system's memory.
- **Memory Management Unit** - The memory management unit (MMU) is responsible for translating a process' virtual addresses into the corresponding physical address for accessing physical memory. It does all the calculation associating with mapping virtual address to physical addresses, and then populates the address translation structures.
- **Address Translation Structures** - There are two kinds you learned about in lecture: segmentation and page tables. Segments are linearly addressed chunks of memory that typically contain logically-related information, such as program code, data, stack of a single process. They are of the form (s,i) where memory addresses must be within an offset of i from base segment s. A page table is the data structure used by a virtual memory system in a computer operating system to store the mapping between virtual addresses and physical addresses. Virtual addresses are used by the accessing process, while physical addresses are used by the hardware or more specifically to the RAM.
- **Inverted Page Table** - The inverted page table scheme uses a page table that contains an entry for each physical frame, not for each logical page. This ensures that the table occupies a fixed fraction of memory. The size is proportional to physical memory, not the virtual address space. The inverted page table is a global structure – there is only one in the entire system. It stores reverse mappings for all processes. Each entry in the inverted table contains has a tag containing the task id and the virtual address for each page. These mappings are usually stored in associative memory (remember fully associative caches from 61C?). Associatively addressed memory compares input search data (tag) against a table of stored data, and returns the address of matching data. They can also use actual hash maps.
- **translation lookaside buffer (TLB)** - A translation lookaside buffer (TLB) is a cache that memory management hardware uses to improve virtual address translation speed. It stores virtual address to physical address mappings, so that the MMU can store recently used address mappings instead of having to retrieve them multiple times through page table accesses.

2 Problems

2.1 Conceptual Questions

If the physical memory size (in bytes) is doubled, how does the number of bits in each entry of the page table change?



If the physical memory size (in bytes) is doubled, how does the number of entries in the page map change?

Page Table based on Virtual Addr space
 → 1 entry per virtual page

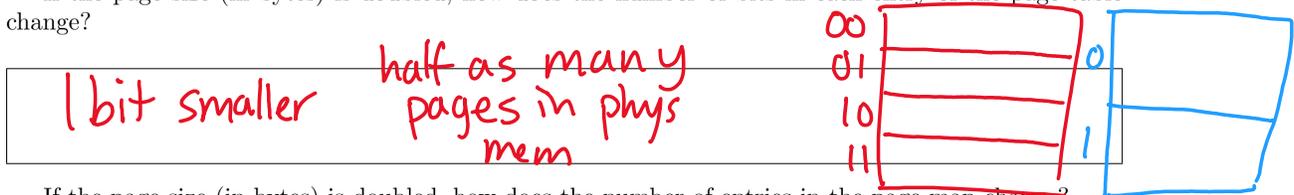
If the virtual memory size (in bytes) is doubled, how does the number of bits in each entry of the page table change?

same per entry

If the virtual memory size (in bytes) is doubled, how does the number of entries in the page map change?

twice as many entries (twice as many indices)

If the page size (in bytes) is doubled, how does the number of bits in each entry of the page table change?



If the page size (in bytes) is doubled, how does the number of entries in the page map change?

Half as many entries
 - half as many virtual pages ⇒ 1/2 as many entries (1 entry per page)

The following table shows the first 8 entries in the page map. Recall that the valid bit is 1 if the page is resident in physical memory and 0 if the page is on disk or hasn't been allocated.

3.2 Banker's Algorithm

Covered in detail next week
- will do high level explanation

Suppose we have the following resources: A, B, C and threads T1, T2, T3 and T4. The total number of each resource as well as the current/max allocations for each thread are as follows:

Bankers Algo
- check if we are in a safe state
safe meaning deadlock can be avoided for sure

		Total					
		A	B	C			
		7	8	9			
		Current			Max		
T/R		A	B	C	A	B	C
T1		0	2	2	4	3	3
T2		2	2	1	3	6	9
T3		3	0	4	3	1	5
T4		1	3	1	3	3	4

Bankers is conservative
- non safe state ≠ deadlock (mostly like a warning)
- assumes: need max resources to complete, no preemption of resources

Is the system in a safe state? If so, show a non-blocking sequence of thread executions.

Need: T1 4 1 1
T2 1 4 8
T3 0 1 1
T4 2 0 3

0	Available: $7-0-2-3-1=1$ $8-2-2-3=1$ $9-2-1-4-1=1$
1	Run: T3 → Available: $1+3=4$ $1+0=1$ $1+4=5$
2	Run: T1 → Available: $4+0=4$ $1+2=3$ $5+2=7$
3	Run: T4 → Available: $4+1=5$ $3+3=6$ $7+1=8$ → run T2

Repeat the previous question if the total number of C instances is 8 instead of 9.

0	Available: 1 1 $8-2-1-4-1=0$
→ no thread can complete → NOT SAFE	

↑
deadlock? maybe, maybe not
↑
what if a thread gives up resource before completing? ⇒ state does not necessarily lead to deadlock.

2.3 Page Allocation

Suppose that you have a system with 8-bit virtual memory addresses, 8 pages of virtual memory, and 4 pages of physical memory.

How large is each page? Assume memory is byte addressed.

$$\text{total VA size} = 2^8 \text{ B} \quad 2^8 \text{ B} / 2^3 \text{ pages} = 2^5 \text{ B/page} = 32 \text{ B}$$

Suppose that a program has the following memory allocation and page table.

Memory Segment	Virtual Page Number	Physical Page Number
N/A	000	NULL
Code Segment	001	10
Heap	010	11
N/A	011	NULL
N/A	100	NULL
N/A	101	NULL
N/A	110	NULL
Stack	111	01

1)
2)
I don't really like this Q
Not going to finish

What will the page table look like if the program runs the following function? Page out the least recently used page of memory if a page needs to be allocated when physical memory is full. Assume that the stack will never exceed one page of memory.

What happens when the system runs out of physical memory? What if the program tries to access an address that isn't in physical memory? Describe what happens in the user program, the operating system, and the hardware in these situations.

```
#define PAGE_SIZE 1024; // replace with actual page size
```

```
void helper(void) {
    char *args[5];
    int i;
    for (i = 0; i < 5; i++) {
        // Assume malloc allocates an entire page every time
        args[i] = (char*) malloc(PAGE_SIZE);
    }
    printf("%s", args[0]);
}
```

1) runs out → kicks out oldest pages
 2) if valid → fetch from disk (stored somewhere)
 user prog requests, OS realizes not in phys mem
 → reads from disk → brings into phys mem
 ↑ PageFile / swap partition

or we can think of it as bytes reached
 $2^{11} \cdot 2^{11} \cdot 2^{11} \cdot 2^{13} = 2^{46}$
 PTE1 PTE2 PTE3 page

2.4 Address Translation

Consider a machine with a physical memory of 8 GB, a page size of 8 KB, and a page table entry size of 4 bytes. How many levels of page tables would be required to map a 46-bit virtual address space if every page table fits into a single page?

$2^3 \cdot 2^{10} = 2^{13}$ B pages, $46 - 13 = 33$ VPN bits
 4 bytes/PTE $\rightarrow 2^{13} / 2^2 = 2^{11}$ entries
 $2^{11} \cdot 2^{11} \cdot 2^{11} = 2^{33}$ VPN addresses \Rightarrow 3 page tables

List the fields of a Page Table Entry (PTE) in your scheme.

pointer (PPN) to page, status/permission bits like read, write, execute, valid, dirty, etc.

Without a cache or TLB, how many memory operations are required to read or write a single 32-bit word?

4 operations: 3 lookup + final access

With a TLB, how many memory operations can this be reduced to? Best-case scenario? Worst-case scenario?

Best case: 2 (TLB hit \rightarrow data)
 Worst case: 5 (TLB miss \rightarrow 3 lookups \rightarrow data)

The pagemap is moved to main memory and accessed via a TLB. Each main memory access takes 50 ns and each TLB access takes 10 ns. Each virtual memory access involves:

- mapping VPN to PPN using TLB (10 ns)
- if TLB miss: mapping VPN to PPN using page map in main memory (50 ns)
- accessing main memory at appropriate physical address (50 ns)

Assuming no page faults (i.e. all virtual memory is resident) what TLB hit rate is required for an average virtual memory access time of 61ns.

see solns

Assuming a TLB hit rate of .50, how does the average virtual memory access time of this scenario compare to no TLB?

see solns

2.5 Inverted Page Tables

Why IPTs? Consider the following case:

- 64-bit virtual address space
- 4 KB page size
- 512 MB physical memory

See solns

How much space (memory) needed for a single level page table? Hint: how many entries are there? 1 per virtual page. What is the size of a page table entry? access control bits + physical page #.

How about multi level page tables? Do they serve us any better here?

What is the number of levels needed to ensure that any page table requires only a single page (4 KB)?

Linear Inverted Page Table

What is the size of of the hashtable? What is the runtime of finding a particular entry?

Assume the following:

- 16 bits for process ID
- 52 bit virtual page number (same as calculated above)
- 12 bits of access information

Hashed Inverted Page Table

What is the size of of the hashtable? What is the runtime of finding a particular entry?

Assume the following:

- 16 bits for process ID
- 52 bit virtual page number (same as calculated above)
- 12 bits of access information

