Lecture 26: End Semester Review

Vivek Kumar Computer Science and Engineering IIIT Delhi vivekk@iiitd.ac.in



- Segments of memory can become unusable due to the result of allocation scheme
- Two types of fragmentation
 - External fragmentation
 - Memory remains unallocated
 - Internal fragmentation
 - Memory is allocated but unused
 - Fixed allocation sizes



Malloc using Implicit List



- Each memory block has a header/footer
- Blocks connected as linked list
- Block splitting during malloc
 - E.g., p=malloc(1)
- Block coalescing during free
 E.g., free(p)

© Vivek Kumar

Address Spaces

- Using some mapping technique to translate the address seen by the user to the actual address on the RAM gives two different views of addresses
 - Virtual address space (as seen by the process)
 - Physical address space (as seen by the OS depending on DRAM size)
- Virtual address space: Set of N = 2ⁿ virtual addresses {0, 1, 2, 3, ..., N-1}
- Physical address space: Set of M = 2^m physical addresses {0, 1, 2, 3, ..., M-1}

Logical Address to Linear Address



Paging

Page #0
Page #1
Page #2
Page #3
Page #4
Page #5
Page #6
Page #7

Physical Memory

- Segmentation is easy to implement but it has major drawbacks because segments could have variable sized memory allocation
 - Can lead to external fragmentation
 - To allocate space for a new process, segments of some other process has to be swapped to the disk (huge overhead)
- Paging solves both the above issues by dividing the available memory into small and fixed size blocks (pages)
 - The memory can then be viewed as an array of fixed sized slots
- Default page size on Linux is 4096 bytes (4KB)
 - Total number of pages in 32-bit addressing mode is 2³²/4KB, i.e., 2²⁰
 - Total number of pages in 64-bit addressing mode is 2⁶⁴/4KB, i.e., 2⁵²

Segmentation Along with Paging



- Segmentation cannot be disabled in x86 processors, but address translation is complex using the segmentation-based approach
 - Hence, to simplify the address translation, segment registers are set to have the same base (0) and bound (2³²-1) values (Lecture #16)
- Segmentation converts the logical address to linear address, which is then converted into virtual address to support paging
- However, as the base and bound of each segment registers maps to entire 2³² memory size (4GB), logical address will be the same as virtual address

CSE231: Operating Systems

Two Components of a Virtual Address

- Dedicate few bits to identify VPN and the remaining bits to identify the offset into that VPN
- How many bits required to represent a VPN for pages of N bytes in a virtual memory of M bytes (total M/N pages)?
 - \circ M=4GB (2³²) and N=4KB
 - Pages = $2^{32}/4096 = 2^{20}$ and bits required to identify 2^{20} pages are 20



A Simple Page Table in Action



A Simple Page Table in Action (Using TLB)



Lecture 26: End Semester Review



- Recall, process address space is not contiguous
 - Inter-segment memory is not contiguous
 - Only intra-segment virtual memory is contiguous (the physical memory may/may-not be contiguous)
- Page table until now is an **array** where the length of the array will be the total number of VPN being used in the process address space
 - Not all the VPNs will be used as the segments are not contiguous (there could be gaps between them)

Two-Level Page Table in x86 VPN



- VA stores page directory index, page table index, and address offset in PP
- CR3 register stores the starting address of the Page Directory
- Each PDE contains the Page Frame Number and a status bit
 - Status=0 if there is not a single valid entry in the PFN, otherwise status =1 (valid PFN only if status=1)
- Page Table Entry (PTE) is found from the PDE
 - Physical Page Frame Number (PFN) is found from the PTE



Demand Paging



- It the mechanism to provide the illusion of infinite memory
- OS keeps some amount of DRAM always free by deciding some upper cap
 - Whenever the upper cap memory level is breached, unused pages are moved out to disk Ο
- OS brings the pages into memory back into DRAM from disk when it is accessed
 - Happens during page fault
 - PTE makes demand paging implementable using the Present bit
 - $\begin{array}{l} \mathsf{P}=\mathsf{Valid} \Rightarrow \mathsf{Page in memory, PTE} \\ \mathsf{points at physical page} \\ \mathsf{P}=\mathsf{Not Valid} \Rightarrow \mathsf{Page not in memory} \end{array}$ \bigcirc
 - \bigcirc



IPC Within a Multicore Processor

- Inter-process communication in shared memory
 - Transfer of control from user space to kernel space and vice-versa
- Complicated IPC mechanism for communication
- OS has to reserve extra memory / resources
 - Separate heap, stack, .text segment, etc. for each process
 - Same copy of .text segment in each process
- Separate page table for each process
- Cost of IPC may exceed the cost of actual computation!

Array Sum using Pthread

```
int main(int argc, char *argv[]) {
                                                     int result:
#include <pthread.h>
                                                     if (SIZE < 1024) {
#include <stdio.h>
                                                       result = array_sum(0, SIZE);
#include <stdlib.h>
                                                     } else {
int A[SIZE]; // Initialization code elided
                                                       pthread_t tid[NTHREADS];
int array_sum(int low, int high) {
                                                       thread_args args[NTHREADS];
  int sum = 0:
                                                       int chunk = SIZE/NTHREADS;
  for (int i=low; i<high; i++) {</pre>
                                                       for (int i=0; i<NTHREADS; i++) {</pre>
    sum += A[i];
                                                           args[i].low=i*chunk; args[i].high=(i+1)*chunk;
  }
                                                           pthread_create(&tid[i],
  return sum;
                                                                             NULL.
}
                                                                             thread_func,
typedef struct {
                                                                             (void*) &args[i]);
  int low:
  int high;
                                                       for (int i=0; i<NTHREADS; i++) {</pre>
  int sum:
                                                           pthread_join(tid[i] , NULL);
} thread_args:
                                                           result += args[i].sum;
                                                       }
void *thread_func(void *ptr) {
                                                     }
  thread_args * t = ((thread_args *) ptr);
                                                     printf("Total Sum is %d\n", result);
  t->sum = array_sum(t->low, t->high);
                                                     return 0:
  return NULL;
}
    CSE231: Operating Systems
                                                    © Vivek Kumar
                                                                                                            20
```

Producer Consumer using Pthreads

- 1. pthread_mutex_lock(&mutex);
- while(task_queue_size() == 0)
- 3. pthread_cond_wait(&cond, &mutex);
- 4.
- 5. task = pop_task_queue();
- pthread_mutex_unlock(&mutex);
- 7. execute_task (task);

- 1. pthread_mutex_lock(&mutex);
- 2. int queue_size = task_queue_size();
- 3. push_task_queue(&task);
- 4. if(queue_size == 0) {
- 5. pthread_cond_broadcast(&cond);
- 6. }
- pthread_mutex_unlock(&mutex);



- pthread_cond_wait causes the current thread to relinquish the CPU and wait until another thread invokes the signal or the broadcast
- Upon call for wait, the thread releases ownership of the mutex and waits until another thread signals the waiting threads

Consumer(s)

Deadlock





- Deadlock occurs when multiple threads need the same locks but obtain them in different order
 - It could be avoided by using lock ordering
 - Ensure that all locks are taken in same order by any thread



© Vivek Kumar



Inode Structure



I hope you enjoyed the course..

All the best for your exams !