
High Performance Computing

Lecture 16

Matthew Jacob

Indian Institute of Science

Operating System

Software that manages the resources of a computer system

- ❑ Main memory
- ❑ CPU time
- ❑ I/O devices
- ❑ Software resources

Process Management

- What is a Process?
 - A program in execution ...
 - but some programs run as multiple processes ...
 - and the same program can be running multiply as more than one process at the same time
- We need a clearer idea of what a process is

Process: A New Perspective

- We could view process as a data structure
- What is a data structure?
 - Example: Stack
 - Defined by the operations that can be done on the associated data
 - In the case of the stack: Push and Pop
 - And a few more
 - Create, Destroy
- We will discuss the idea of process in terms of a set of operations and the associated data

Process as a Data Structure


- What are the operations on processes?
 - Process related **system calls**
 - `fork()`, `exec()`, `wait()`, `exit()`, ...
 - Recall: When we talked about the stack operations Push and Pop, we thought in terms of explicit data
 1. A contiguous region of memory (a 1-dimensional array)
 2. A Top of Stack Pointer

Process as a Data Structure.


- What is the data manipulated by these process operations?
 - Think about the `fork()` operation
 - Create a new process

fork() and exec()

Parent Process



Child Process



```
retval = fork();  
if (retval == 0) {  
    /* child */  
  
} else {  
    /* parent */  
}
```

Process as a Data Structure.

- What is the data manipulated by these process operations?
 - Think about the `fork()` operation
 - Create a new process
 - Text, data, stack, heap copied from its parent process

Process as a Data Structure.

- What is the data manipulated by these process operations?
 1. Text, Data, Stack, Heap
 2. Data stored in hardware
 - PC value, register values (R1-R31, etc)

Process as a Data Structure.

- What is the data manipulated by these process operations?
 1. Text, Data, Stack, Heap
 2. Data stored in hardware
 3. Other information maintained by the OS
 - Process identifiers
 - A unique name/integer for each process
 - With each process, also remember its parent process identifier

Process as a Data Structure.

- What is the data manipulated by these process operations?
 1. Text, Data, Stack, Heap
 2. Data stored in hardware
 3. Other information maintained by the OS
 - Process and parent identifiers
 - User identifier
 - The identity of the user who is running the program as a result of which the process was created

Process as a Data Structure.

- What is the data manipulated by these process operations?
 1. Text, Data, Stack, Heap
 2. Data stored in hardware
 3. Other information maintained by the OS
 - Process, parent and user identifiers
 - Memory management information: Page table
 - CPU time used by the process, in user/system
 - File related info: Open files, file pointers

Process : A New Perspective

1. We can view a process as a data structure
2. We can view a process as an OS created abstraction
 - Like virtual memory, it does not really (physically) exist
 - It is just an OS interface for program execution
 - It is the unit of resource management by OS
 - There is a separate page table for each process
 - Unit of sharing of CPU time

Process vs Program

Some adjectives to describe these concepts

- Program: static passive dead
- Process: dynamic active living
- A process changes state with time
 - Example: After executing LW R5, -8(R29), the value in R5 changes
 - as well as the values in PC, IR
 - and many other of the data items we listed
- The process is said to be **running**

A Running Process

- Question: Can 2 processes be running at the same time?
- Answer: Not if there is only one PC, one IR, one set of general purpose registers
 - i.e., not if there is only one CPU of the kind that we have been talking about
- So, if there are 100 processes on a computer system and process P_1 is running, what state are the other 99 processes in?