



1B11 Operating Systems


Processes

Prof. Steve R Wilbur
s.wilbur@cs.ucl.ac.uk



Lecture Objectives


- 1 How do we make best use of the processor cycles?
- 1 How do we do Multiprogramming?
- 1 What is a Process?
- 1 How are they managed?



1B11-2 1999 Slide 2

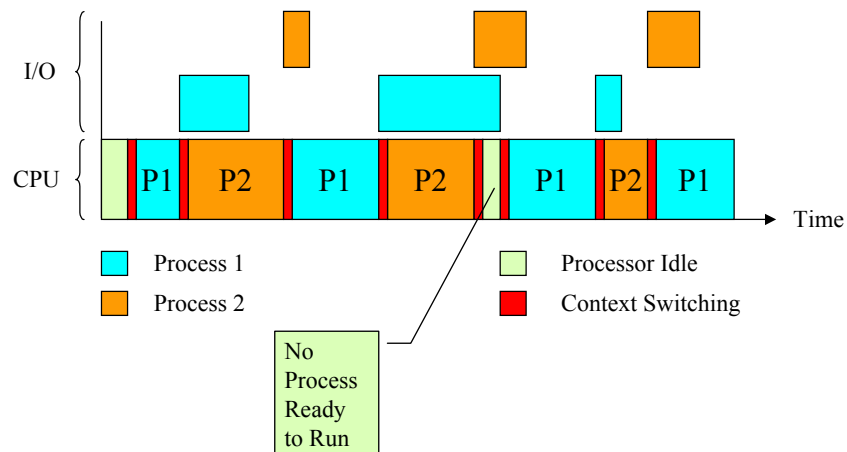
Why have Processes?

- 1 Computers (were) expensive
- 1 CPU idle while waiting for I/O (polled)
- 1 Could have several programs in memory
- 1 Switch to another as they engage in I/O
- 1 Now important management concept



1B11-2 1999 Slide 3


In More Detail



1B11-2 1999 Slide 4


A Process

- | A Process is a running program - the program that the CPU is currently fetching and executing instructions for
- | A process exists in a virtual machine consisting of:
 - o memory
 - o a schedule
 - o a set of I/O channels that it is allowed to use


1B11-2 1999 Slide 5

Process Life Cycle

- | To start a program, the OS must:
 - o Find the binary/executable on disk
 - o Allocate storage for the code and data
 - o Allocate swap space on disk
 - o Map the pages in memory and swap space
 - o Copy the code into the physical pages
 - o Start execution by saving current context (see later) and jumping (loading the PC) to the start address of the code!

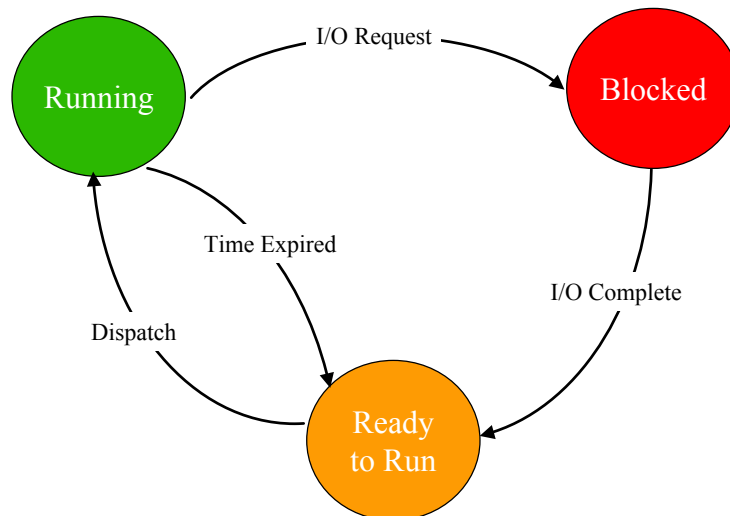

1B11-2 1999 Slide 6

Important Features

- 1 Code and data for process must be in memory
- 1 When process is suspended from running need to save its *current state*
- 1 Saving, restoring this state and deciding which process to run is known as *context switching* - "*wasted*" time or *overhead*
- 1 Needs to be *policy* for choosing next process to run - high level *scheduler*

1B11-2 1999 Slide 7

Process State Transitions



1B11-2 1999 Slide 8

Process Representation

- 1 OS keeps *Process Control Block (PCB)* for each process
- 1 Contains management information
 - o Location of P in memory
 - o P's state
 - o P's owner (user)
 - o etc.

User Space

Kernel Space (OS)

Memory

1B11-2 1999 Slide 9

Process Representation - 2

Current (Running) P: ● → PCB5

Ready to Run: ● → PCB3 ● → PCB1 ● → PCB0

Blocked: ● → PCB2 ● → PCB4

Note: Idle "Process" → PCB0

1B11-2 1999 Slide 10

Process State Transitions - 2

Running-> Ready | Time slot expires
| Scheduler determines time slot

Ready->Running | Current P-> Blocked or *pre-empted*
| Highest priority runnable-> running
| Null/Idle process if Q empty

Running->Blocked | P makes *synchronous I/O request*
| or waits for earlier *asynchronous I/O request*

Blocked->Ready | I/O or events complete

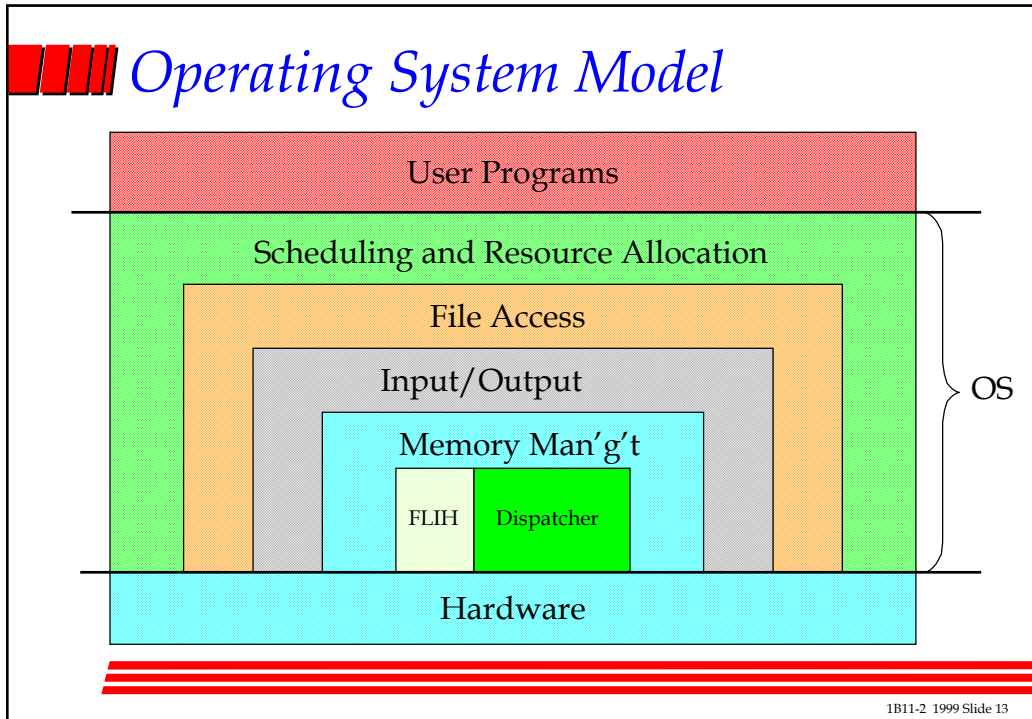
1B11-2 1999 Slide 11

Process State Transitions - 3

- | Null/Idle Process
 - o May do something useful: machine diagnostics, compute prime numbers, etc
 - o or just wait in kernel loop for event (I/O)

- | Scheduling vs Dispatching
 - o Dispatcher just takes next P from Q and sets it running
 - o Scheduler decides on ordering of Ps on Q on basis of policy, eg. shortest job, premium payment, real-time activity, etc.

1B11-2 1999 Slide 12



- ### *Context...*
- | Processes are defined by their code, data, and state
 - | What does the state consist of:
 - o Program Counter
 - o Stack Pointer
 - o Register Contents
 - o Page/MMU register Contents
 - o Anything else?
 - | How do we suspend one program, to start another (i.e. give the illusion of concurrency)?
- 1B11-2 1999 Slide 14

Process Control Block

1 Contains **state information** such as:

o Process State	Ready, etc
o Process ID	to relate I/O, events inter-process communications (IPC), etc
o Priority	or time slot, time to run etc.
o Memory pointers	location of process code and data
o Resources Allocated	terminals, devices
o Register Save Area	processor registers, stack pointer etc. User ID
o Owner	Parent process ID
o Parent	

1B11-2 1999 Slide 15

Process Control Block - 2

1 When P **swapped out**:

- o Not all of PCB needs to be retained in kernel
- o Need to also store swap device address of Process image in PCB

1B11-2 1999 Slide 16

■■■■ Cost of Context Switch?

- 1 A context switch entails at least 1 trap (call to scheduler or I/O or clock interrupt)
- 1 Then it involves a *full save of registers and page table*
- 1 What if we knew we were only going to do a bit of work? Couldn't we do something quicker?
- 1 Yes: threads.....

1B11-2 1999 Slide 17

■■■■ Threads - bare processes...

- 1 Threads are also known as "*lightweight processes*"
- 1 Assume we have a "safe" language, then we can have "concurrency" and scheduler, but avoid saving all registers and so on - let threads run in *same virtual address space and so on*
- 1 All the same except that scheduler is part of library not a system call (cheaper) and much less state saved and restored in a context switch (minimum is stack, register and PC)

1B11-2 1999 Slide 18

Scheduling Algorithms

- | What is the order in which a scheduler chooses the next process when another one suspends or stops?
 - o Random, round robin (fair) priority, mix
- | Can try to balance between I/O and CPU intensive based on past behaviour
- | Can have a real-time requirement (e.g. must respond before the reactor blows up).
- | Complex- especially in newer multiprocessor or distributed systems!

1B11-2 1999 Slide 19

Communicating Processes

- | Normally in a time-sharing multi-tasking system processes are unrelated. But, sometimes...we want them to communicate with each other.
- | Communicating processes in different machines or address spaces use messages and protocols.
- | Communicating processes in the same address space can share memory, and therefore can communicate by sharing variables (memory)

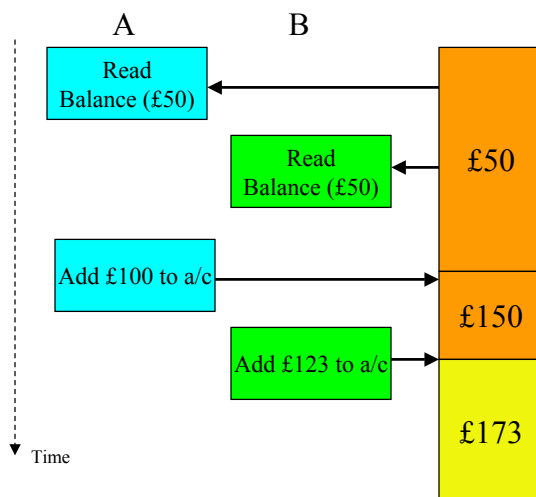
1B11-2 1999 Slide 20

Shared Variable Comms

- There are *interesting* problems with concurrent access to shared variables:
 - Run process A, which reads bank account v
 - Switch to process B, which reads same value
 - Switch to process A which adds £100 to it and writes it
 - Switch to process B which adds £100 to it and writes it
 - Result is wrong (only 1 update had an effect)
 - "Lost update"* problem

1B11-2 1999 Slide 21

Lost Update Problem




- Result should be £273
- One update lost because balance read a second time before first update complete
- Need to *bind* two steps of update together
- A could *lock* data until it has finished with it

1B11-2 1999 Slide 22


Shared Variable Comms - 2

- 1 Solutions include:
 - o Locks
 - o Queues
 - o Monitors - lots of others - will be covered later in course - a difficult subject


1B11-2 1999 Slide 23


Problems with message passing

- 1 Concurrent systems that pass messages or share variables can *deadlock*
 - o Process a is waiting for process c
 - o Process c is waiting for b
 - o Process b is waiting for a
 - o (consider problems that can occur at a road junction when each car blocks the exit of another)
- 1 or *livelock* or be unfair and so on


1B11-2 1999 Slide 24


Summary

- | Process - running program
- | Scheduler picks between processes to run
- | Scheduler saves and restores contexts in process table
- | Virtual Machine hides processes from each other
- | Shared memory and messages can allow communication between processes
- | Concurrency causes problems


1B11-2 1999 Slide 25

Further Reading

- | Ritchie C, "Operating Systems inc. UNIX and Windows", 3rd ed, 1997, Letts, 1 85805 302 1, £13.95
- | and many more


1B11-2 1999 Slide 26