Week 9: Scheduling and Concurrency

CS211: Programming and Operating Systems

Niall Madden (`Niall.Madden@NUIGalway.ie`)

Wednesday and Thursday, 14+15 April 2021

CS211, this week, will be all about . . .

CS211
Week 9: Scheduling and Concurrency

Start of ...

PART 1: Part 1: Scheduling Processes - examples

Part 1: Scheduling - examples

For more, see Chapter 7 of the Textbook.

Last week we started studying ***SCHEDULING***: algorithms by which the Operating System decides which of the available processes will be given access to the CPU, i.e., set **running**.

First: recall the **states of a process** and how they relate to each other.

We studied four **Scheduling Algorithms**:

1. First-Come-First-Served (FSFS)
2. **Shortest-Job-First** (SJF)
3. Shortest Time-to-Completion First (STCF)
4. Round-Robin (RR)

For each of these, we consider a few examples of process mix; for each example we'll assumed that processes have a single CPU burst, measured in seconds (though this unit is not important).

We compared algorithms according to the following **METRICS**:

1. Turnaround time – the time that elapses between when process arrives in the system, and when it finally completes.
2. Wait time – the amount of time between when a process arrives, and when it completes, that it spends doing nothing.
3. Response time – the time that elapses between when process arrives in the system, and when it executes for the first time.

Each process gets a small unit of CPU time called a *time quantum* or *time slice* –usually 10-100 milliseconds. After this time has elapsed, the process is preempted and added to the end of the ready queue.

If there are n processes in the ready queue and the time quantum is q, then each process gets $1/n$ of the CPU time in chunks of at most q time units at once. No process waits more than $(n-1)q$ time units. (*Why?*)

The size of the *quantum* is of central importance to the **RR** algorithm. If it is too large, then its is just the FCFS model. If it is too low, them too much time is spent on context switching.

Suppose the following arrive in the following order:

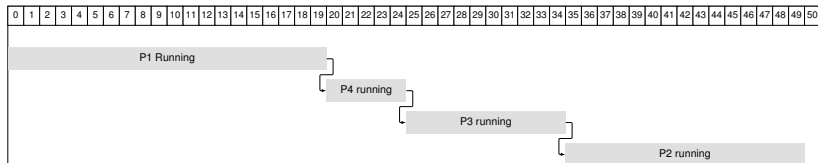| Proc	Arrive Time	Burst Time
P_1	0	20
P_2	2	15
P_3	4	10
P_4	6	5

Calculate the

- [a] ***Average Turnaround Time***,
- [b] ***Average Wait Time***, and
- [c] ***Average Response Time*** for

1. FCFS
2. SJF
3. STCF
4. RR with $q = 10$

FCFS
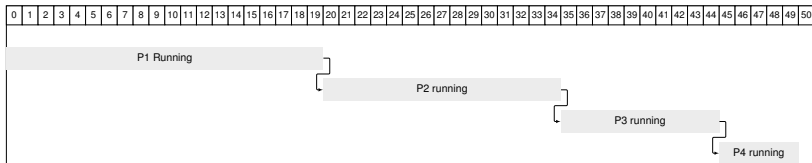
Shortest Job First (SJF)

| 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 | 21 | 22 | 23 | 24 | 25 | 26 | 27 | 28 | 29 | 30 | 31 | 32 | 33 | 34 | 35 | 36 | 37 | 38 | 39 | 40 | 41 | 42 | 43 | 44 | 45 | 46 | 47 | 48 | 49 | 50 |

P1 Running

P4 running

P3 running

P2 running

Shortest Time to Completion First (STCF)

| 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 | 21 | 22 | 23 | 24 | 25 | 26 | 27 | 28 | 29 | 30 | 31 | 32 | 33 | 34 | 35 | 36 | 37 | 38 | 39 | 40 | 41 | 42 | 43 | 44 | 45 | 46 | 47 | 48 | 49 | 50 |

P1 runs   P1 runs

    P2 runs   P2 runs

        P3 runs   P3 runs

            P4 runs

Round Robin with $q = 10$

| 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 | 21 | 22 | 23 | 24 | 25 | 26 | 27 | 28 | 29 | 30 | 31 | 32 | 33 | 34 | 35 | 36 | 37 | 38 | 39 | 40 | 41 | 42 | 43 | 44 | 45 | 46 | 47 | 48 | 49 | 50 |

P1 runs

P2 runs

P3 runs

P4 runs

P1 runs

P2 runs

CS211
Week 9: Scheduling and Concurrency

END OF PART 1

CS211
Week 9: Scheduling and Concurrency

Start of ...

PART 2: Concurrency

Part 2: Concurrency

Please read Chapter 26 (Concurrency) of the textbook for much more detail on threads.

A *cooperating* process is one that can affect or be affected by another process that is executing on the system.

Threads are prime examples of this: we can think of them as a single process with multiple points of execution. They share program code and, crucially, data.

In this section, we consider the problems that occue then one or more threads try to access the same data, and we look at potential solutions.

A classic data inconsistency problem is the so-called

"*Race Condition*",

which we'll study in Lab 6 (a more complicated version is discussed in Sections 26.3-26.4).

A **Race Condition** (also called a **data race**) is one where the result depends on the order in which instructions are executed.

For a single-thread process, this is predetermined.

But for multi-threaded processes, we do not have control over the order in which individual threads execute their instructions.

Consider the following example: two cooperating process called P_1 and P_2 share the variable `count`. At various times during execution either may increment or decrement `count`.

The machine usually implements an **increment** as follows:

1. load the value of `count` into a register: REG_1 = count
2. add 1 to the contents of the register: REG_1 = REG_1 + 1
3. overwrite the contents of `count` with the contents of the register: `count` = REG_1.

A **decrement** would be implemented as

1. load the value of `count` into a register: REG_2 = count
2. subtract 1 from the contents of the reg: REG_2 = REG_2 - 1
3. save the contents of the register as `count`: `count` = REG_2

Suppose the value of `count` is 5. If P_1 executes an increment and P_2 executes a decrement, then the value of count should still be 5. Unless the individual operations happen in the following order...

P_1 executes	REG_1	=	$count$	$REG_1 = 5$
P_1 executes	REG_1	=	$REG_1 + 1$	$REG_1 = 6$
P_2 executes	REG_2	=	$count$	$REG_2 = 5$
P_2 executes	REG_2	=	$REG_2 - 1$	$REG_2 = 4$
P_1 executes	$count$	=	REG_1	$count = 6$
P_2 executes	$count$	=	REG_2	$count = 4$

We arrive at the wrong state because we allowed both threads to manipulate the variable `count` at the same time.

Since the outcome depends on the order in which each operation takes place, we have a ***race condition***.

CS211
Week 9: Scheduling and Concurrency

END OF PART 2

CS211
Week 9: Scheduling and Concurrency

Start of ...

PART 3: Critical sections

Part 3: Critical sections

Critical Section

(From Section 26.4 of the text-book). A **critical section** is a piece of code that accesses a shared variable (or more generally, a shared resource) and must not be concurrently executed by more than one thread.

The example given on Slides 17 and 18 shows that multiple threads executing the same code can result in a race condition, that is an example of a **critical section.**

To resolve this, we would like to enforce **mutual exclusion**: This property guarantees that if one thread is executing within the critical section, the others will be prevented from doing so.

One possible solution is to make the operation **"atomic"** (or ***indivisible***. This is, the critical section is executed as though it were a single operation, and so impossible to interrupt.

In a realistic setting, that is not possible for all race conditions. But, as we will see, the use of some atomic operations can help us solve the larger problem, by creating **locks**.

See Section 28.1 of the textbook

So now we know we would like to execute a series of instructions atomically. But, in general, on a multiprocessor system, we can't.

But what we can do is create a **lock** which we put around critical sections, and thus ensure that any such critical section executes as if it were a single atomic instruction.

For the **lock** approach to work, the following 3 conditions must be satisfied:

- **Mutual Exclusion:** If process T_i is executing in its critical section, then no other processes can be executing in their critical sections.
- **Fairness/Progress:** If there are some procs that wish to enter the critical section, then the selection of the process that will enter the critical section next cannot be postponed indefinitely.
- **Performance/Bounded Waiting:** after a process has made a request to enter its critical section and before that request is granted, then must be a **bound** (i.e., a limit) on the number of times other processes are allowed to enter their critical sections.

We consider two basic approaches to this:

1. Interrupt suspension
2. Automatic *test-and-set()* and *swap()* instructions.

1. Interrupt suspension (Section 28.5)

Suppose a process is in its critical section. If it cannot be preempted then data consistency should be maintained. On a single processor system, the problem could be solved by disabling interrupts while a shared variable is begin modified. However, such a method is not feasible on a multiproc system: large over-heads would be incurred informing all procs that interrupts are dis-allowed.

2 Automatic instructions

The processor has facilities to swap the contents of two words (in memory), or test and change the contents of a word *automatically* – i.e., as a single instruction.

There are other hardware and software solutions to synchronization problems. The most important, perhaps, is a tool known as a *semaphore* (Chapter 31).

CS211
Week 9: Scheduling and Concurrency

END OF PART 3

CS211
Week 9: Scheduling and Concurrency

Start of ...

PART 4: Semaphores

Part 4: Semaphores

A Semaphore *S* is an integer variable that can only be accessed via one of two operations:

....

(Test/*sem_wait*) $P(S)$:

```
while (S ≤ 0)
      { wait(); }
 S--;
```

....

(**Increment/*sem_post***) $V(S)$:   `S++;`

....

(Historically, these functions were called *P*robern (Dutch for "test") and *V*erhogen (increment)).

These operations must be ***indivisible*** (or "atomic"). This is, when one process (or thread) modifies a semaphore value, no other process can modify it at the same time.

Part 4: Semaphores

There are two types of semaphore:

1. **Binary semaphores (locks):** These are used to control access to a single resource, such as a memory location. If the resources is available then $S = 1$. Otherwise $S = 0$. When a process wants to access it,

   (i) it calls the function $P(S)$
   (ii) enters it critical section
   (iii) calls $V(S)$ when it exits the critical section.

2. **General (or counting) semaphores:** These are used to control access to a pool consisting of a finite number of identical resources. Say there are 5 units available. The S is initialised to 5. Whenever a process requests the resource, it calls $P(S)$ and decrements the value of S. If S reaches 0 then the next proc that requests that resource must wait until another frees it by running $V(S)$.

Lab 6 will be all about coding a senario when a **race condition** can occur, and then solving it.

For consider the example in `adder.c`.

- a (parent) process creates a subprocess (i.e., "child process").
- Then the subprocess tries to calculate the sum of 4 numbers by placing them in a pipe, from where the parent will read them.
- The parent sums the four numbers, and places the result in another pipe for the child to read.
- The child reads this solution and prints it.

Since there are no competing processes, nothing should go wrong (not does it).

adder.c (main)

```
   int main(void )
30 {
     int ParentsPID, ans;
32   pipe(inpipe);
     pipe(outpipe);
34   ParentsPID = getpid(); // now I'll always know who I am
     fork(); // Now have 2 procs. Subprocess will have differnt pid
36   if ( getpid() == ParentsPID )
       Adder();   // The parent will be the adder
38   else
     {
40     ans = SubProc(1,2,3,4);
       printf("SubProc (%d): 1+2+3+4= %d\n", getpid(), ans);
42   }
     return(0);
44 }
```

`adder.c`: `adder()`, run by parent

```
46  void Adder(void ) // run by parent
    {
48    int i, number, sum=0;

50    for (i=0; i<4; i++)
      {
52      read(inpipe[0], &number, sizeof(int));
        sum += number;
54    }
      write(outpipe[1], &sum, sizeof(int));
56  }
```

adder.c: SubProc(), run by subprocess ("child")

```c
58  int SubProc ( int a , int b , int c , int d )
    {
60     int ans ;
       printf ( "SubProc (%d) writes four numbers to the pipe ()\n" ,
62             getpid ());
       write ( inpipe [1] , &a , sizeof ( int ));
64     write ( inpipe [1] , &b , sizeof ( int ));
       sleep (1);  // Pause for a second to encourage race condition
66     write ( inpipe [1] , &c , sizeof ( int ));
       write ( inpipe [1] , &d , sizeof ( int ));

       printf ( "SubProc (%d) reads the answer from a pipe ()\n" ,
70             getpid ());
       read ( outpipe [0] , &ans , sizeof ( int ));
72     return ( ans );
    }
```

The output I get when I run this is:

```
SubProc (4285) writes four numbers to the pipe()
SubProc (4285) reads the answer from a pipe()
SubProc (4285): 1+2+3+4= 10
```

So - no problem!

But in the next version, the parent has two children both doing the same thing. See `adder_race_condition.c`

Now the output is

```
SubProc (4485) writes four numbers to the pipe()
SubProc (4486) writes four numbers to the pipe()
SubProc (4485) reads the answer from a pipe()
SubProc (4485): 1+2+3+4= 6
SubProc (4486) reads the answer from a pipe()
SubProc (4486): 1+2+3+4= 14
```

In Lab 6 we'll design a semaphore solution to this problem

Recall that a **Semaphore** S is an integer variable that can only be accessed via one of two operations: **Test**/***sem_wait*** $P(S)$, and **Increment**/***sem_post*** $V(S)$.

A more detailed explanation of semaphores

Implementing semaphores with `pipe()` in C.
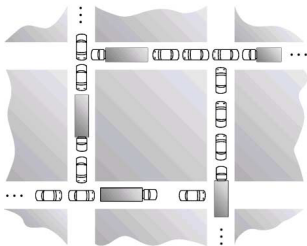
CS211
Week 9: Scheduling and Concurrency

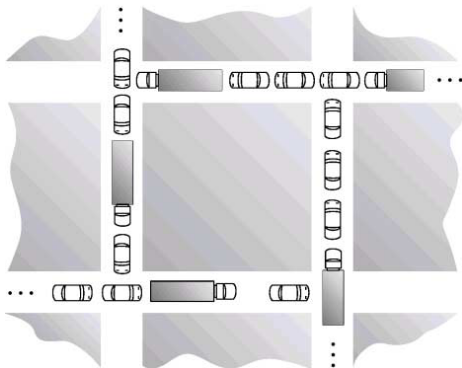**END OF PART 4**

CS211
Week 9: Scheduling and Concurrency

Start of ...

PART 5: Deadlock and Starvation

Part 5: Deadlock and Starvation

Deadlock is when two or more procs are waiting indefinitely for an event that can only be caused by one of the waiting processes. E.g., all are stuck in the `wait()` loop of the *P*() function.

Part 5: Deadlock and Starvation

Deadlock can arise if four conditions hold simultaneously

1. **_Mutual exclusion_**: only one process can have access to a particular resource at any given time.

2. **_Hold and wait_**: a process holding at least one resource is waiting to acquire additional resources held by other processes.

3. **_No preemption_**: a resource can only be released voluntarily by the process holding it, after that process has completed its task.

4. **_Circular wait_**: there exists a set $\{P_0, P_1, ..., P_n, P_0\}$ of waiting processes such that
   - P_0 is waiting for a resource that is held by P_1,
   - P_1 is waiting for a resource that is held by P_2, ...,
   - P_{n-1}, is waiting for a resource that is held by P_n which in turn is waiting for P_0.

Part 5: Deadlock and Starvation

CS211
Week 9: Scheduling and Concurrency

END OF PART 5

CS211
Week 9: Scheduling and Concurrency

Start of ...

PART 6: Resource Allocation Graphs

Part 6: Resource Allocation Graphs

Deadlocks may be described in using a directed graph called a ***resource allocation graph***.

This graph has two sets of **vertices:**

- **Processes:** P_0, P_1, ..., P_n and
- **Resources:** R_0, R_1, ..., R_m

and ***Edges***

- from P_j to R_k if process j as requested resource k but not yet been allocated it,
- from R_k to P_k ig process j as been allocated resource k and not yet released it.

Example (Example 1)

A system has $m = 2$ resources, R_1 and R_2,
and $n = 2$ processes, P_1 and P_2.

- P_1 has been allocated R_1, and is requesting R_2.
- P_2 has been allocated R_2, and is requesting R_1.

Draw the resource allocation graph for the scenario.
Does the system reach deadlock?

Example (Example 2)

A system has $m = 2$ resources, R_1 and R_2,
and $n = 2$ processes, P_1 and P_2.

- P_1 is requesting R_2.

- P_2 has been allocated R_2, and is requesting R_1.

Draw the resource allocation graph for the scenario.
Does the system reach deadlock?

- If there are no cycles, there is no deadlock,
- If there is deadlock, there must be a cycle
- If there is a cycle, there *may* be deadlock
- If each resource has only one instance, and there is a cycle, then there is deadlock

**CS211
Week 9: Scheduling and Concurrency**

END OF PART 6

CS211
Week 9: Scheduling and Concurrency

Start of ...

PART 7: The Dining Philosophers Problem

Part 7: The Dining Philosophers

(See also Section 31.6 of the textbook)

Starvation occurs when a particular process is always waiting for a particular semaphore to become available.

The ideas of **Deadlock** and **Starvation** are exhibited in the classic synchronization problem: ***The Dining Philosophers Problem***.

- There are five philosophers seated at a round table.
- Each has a bowl of rice in front of them and a chopstick to their left and right.
- They spend their day alternating between eating and thinking.
- However there are only five chopsticks...

Part 7: The Dining Philosophers
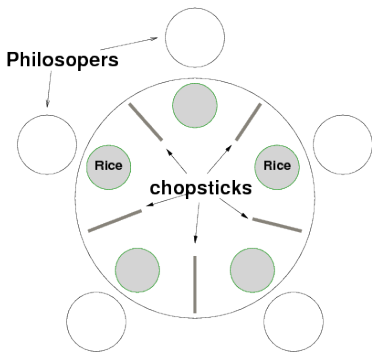
If a philosopher is hungry, she will try to pick up the chopstick to the left and then the chopstick to the right. If she manages to do this they will eat for a while before putting down both and thinking for a while. However, if she it picks up one, she will not let go of it until she can picks up the second and eat.

Suppose each of the picks up the fork to their left. No forks remain on the table so we reach a state of deadlock.

The challenge is to find a solution so that

- **Deadlock** does not occur.
- neither does **starvation** where one philosopher never gets to eat.

Part 7: The Dining Philosophers

Example (Taken from 1718-CS211 Exam)

In the dining philosophers problem, each philosopher wants to pick up the 2 forks beside him/her so that they can eat. Suppose we have 5 such philosophers and

- at time $t = 0$: nobody has picked up any fork
- at time $t = 1$: Philosophers 1, 2 and 3 have picked up the forks to their left, philosopher 4 has not picked up any fork and philosopher 5 has picked up two forks.

In general operating systems take one of three approaches to deal with deadlock:

1. Ensure that the system will never enter a deadlock state.
2. Allow the system to enter a deadlock state and then recover.

In Case 1, there are two possibilities:

- **_Prevention:_** we ensure at least one of the four necessary conditions never hold.
- **_Deadlock avoidance:_** where the OS uses **_a priori_** information about the procs the devise an algorithm to circumvent deadlock.

In Case 2, the OS must have mechanisms for first detecting deadlock and then dealing with it.

Exercises

Exercise (9.1)

(This is taken from the CS211 Semester 2 from 2017/2018)
Given the data (all time in seconds)

Process	Arrival time	Process duration	
P1	3	5	
P2	1	3	*for four processes,*
P3	0	8	
P4	4	6	

determine the scheduling result for the policies of

1. *Round Robin (with time quantum 4)*
2. *First Come First Served*

(c) Calculate the average turnaround time and average waiting time for these examples.

Note: The "wait time" of a process is the length of time it spends doing nothing.

Exercises

Exercise (9.2)

Draw the resource allocation graph for the scenario where all philosophers pick up the fork/chopstick to their left.

Exercise (9.3)

A system has $m = 4$ identical resources, and $n = 3$ processes, P_1, P_2 and P_3, which make a request for 1, 2, and 3 resources, respectively. Draw the resource allocation graph for the scenario. Can the system reach deadlock?