

Section 6: Synchronization, Scheduling

July 13, 2020

Contents

1	Vocabulary	2
2	Synchronization	3
2.1	test_and_set	3
3	Scheduling	5
3.1	Round Robin Scheduling	5
3.2	Life Ain't Fair	5
3.3	Delivery Service	7

1 Vocabulary

- **Scheduler** - Routine in the kernel that picks which thread to run next given a vacant CPU and a ready queue of unblocked threads.
- **Preemption** - The process of interrupting a running thread to allow for the scheduler to decide which thread runs next.
- **FIFO Scheduling** - First-In-First-Out (aka First-Come-First-Serve) scheduling runs jobs as they arrive. Turnaround time can degrade if short jobs get stuck behind long ones (convoy effect);
- **round-robin Scheduling** - Round-Robin scheduling runs each job in fixed-length time slices (quanta). The scheduler preempts a job that exceeds its quantum and moves on, cycling through the jobs. It avoids starvation and is good for short jobs, but context switching overhead can become important depending on quanta length;
- **Shortest Time Remaining First Scheduling** - A scheduling algorithm where the thread that runs is the one with the least time remaining. This is ideal for throughput but also must be approximated in practice.
- **Linux CFS** - Linux scheduling algorithm designed to optimize for fairness. It gives each thread a weighted share of some target latency and then ensures that each thread receives that much virtual CPU time in its scheduling decisions.
- **Earliest Deadline First** - Scheduling algorithm used in real time systems. It attempts to meet deadlines of threads that must be processed in real time by selecting the thread that has the closest deadline to complete first.
- **Multi-Level Feedback Queue Scheduling** - MLFQS uses multiple queues with priorities, dropping CPU-bound jobs that consume their entire quanta into lower-priority queues.

2 Synchronization

2.1 test_and_set

In the following code, we use `test_and_set` to emulate locks.

```
int value = 0;
int hello = 0;

void print_hello() {
    while (test_and_set(&value));
    hello += 1;
    printf("Child thread: %d\n", hello);
    value = 0;
    pthread_exit(0);
}

void main() {
    pthread_t thread1;
    pthread_t thread2;
    pthread_create(&thread1, NULL, (void *) &print_hello, NULL);
    pthread_create(&thread2, NULL, (void *) &print_hello, NULL);
    while (test_and_set(&value));
    printf("Parent thread: %d\n", hello);
    value = 0;
}
```

Assume the following sequence of events:

1. Main starts running and creates both threads and is then context switched right after
2. Thread2 is scheduled and run until after it increments hello and is context switched
3. Thread1 runs until it is context switched
4. The thread running main resumes and runs until it get context switched
5. Thread2 runs to completion
6. The thread running main runs to completion (but doesn't exit yet)
7. Thread1 runs to completion

Is this sequence of events possible? Why or why not?

At each step where `test_and_set(&value)` is called, what value(s) does it return?

Given this sequence of events, what will C print?

Is this implementation better than using locks? Explain your rationale.

3 Scheduling

3.1 Round Robin Scheduling

Which of the following are true about Round Robin Scheduling?

1. The average wait time is less than that of FCFS for the same workload.
2. It requires pre-emption to maintain uniform quanta.
3. If quanta is constantly updated to become the # of cpu ticks since boot, Round Robin becomes FIFO.
4. If all threads in the system have the same priority, Priority Schedulers **must** behave like round robin.
5. Cache performance is likely to improve relative to FCFS.
6. If no new threads are entering the system all threads will get a chance to run in the cpu every $QUANTA * SECONDS_PER_TICK * NUMTHREADS$ seconds. (Assuming $QUANTA$ is in ticks).
7. It is the fairest scheduler

3.2 Life Ain't Fair

Suppose the following threads denoted by `THREADNAME : PRIORITY` pairs arrive in the ready queue at the clock ticks shown. Assume all threads arrive unblocked and that each takes 5 clock ticks to finish executing. Assume threads arrive in the queue at the beginning of the time slices shown and are ready to be scheduled in that same clock tick. (This means you update the ready queue with the arrival before you schedule/execute that clock tick.) Assume you only have one physical CPU.

```
0 Sam : 7
1
2 Kevin : 1
3 Bobby: 3
4
5 Jonathan : 5
6
7 William: 11
8
9 Alex: 14
```

Determine the order and time allocations of execution for the following scheduler scenarios:

- Round Robin with time slice 3
- Shortest Time Remaining First (SRTF/SJF) WITH preemptions
- Preemptive priority (higher is more important)

Write answers in the form of vertical columns with one name per row, each denoting one clock tick of execution. For example, allowing Sam 3 units at first looks like:

```
0 Sam
1 Sam
2 Sam
```

It will probably help you to draw a diagram of the ready queue at each tick for this problem.



3.3 Delivery Service

Assume each numbered line of code takes 1 CPU cycle to run, and that a context switch takes 2 CPU cycles. Hardware preemption occurs every 50 CPU cycles and takes 1 CPU cycle. The scheduler is run after every hardware preemption and takes 0 time. Finally, the currently running thread does not change until the end of a context switch.

Lock lock_a, lock_b; // Assume these locks are already initialized and unlocked.

```
int a = 0;
```

```
int b = 1;
```

```
bool run = true;
```

```

    Kiki() {
1.     bool cond = run;
2.     while (cond) {
3.         int x = a;
4.         int y = b;
5.         int sum = x + y;
6.         lock_a.acquire();
7.         lock_b.acquire();
8.         a = y;
9.         b = sum;
10.        lock_a.release();
11.        lock_b.release();
12.        cond = run;
    }
}

```

```

    Jiji() {
1.     bool cond = run;
2.     while (cond) {
3.         int x = b;
4.         int sum = a + b;
5.         lock_b.acquire();
6.         lock_a.acquire();
7.         b = sum;
8.         a = x;
9.         lock_b.release();
10.        lock_a.release();
11.        cond = run;
    }
}

```

```

    Tombo() {
1.     while (true) {
2.         lock_a.acquire()
3.         a = 0;
4.         lock_a.release()
5.         lock_b.acquire()
6.         b = 1;
7.         lock_b.release()
    }
}

```

Thread 1 runs `Kiki`, Thread 2 runs `Jiji`, and Thread 3 runs `Tombo`. Assuming round robin scheduling (threads are initially scheduled in numerical order):

1. What are the values of `a` and `b` after 50 CPU cycles?

2. What are the values of `a` and `b` after 200 CPU cycles? What line is the program counter on for each thread?

Now assume we use the same round robin scheduler but we just run 2 instances of `Kiki`.

3. What are the values of `a` and `b` after 100 cycles?

Now we replace our scheduler with a CFS-like scheduler. In particular, this scheduler is invoked on every call to `lock_acquire` or `lock_release`. We still have 2 threads running `Kiki`, but this time thread 1 runs for 50 cycles before thread 2 starts.

4. What are the values of `a` and `b` at the end of the 73rd cycle?