IN5170/IN9170 – Models of concurrency

Fall 2021

Compiled: September 29, 2021

Abstract

This is the "handout" version of the slides for the lecture (i.e., it's a rendering of the content of the slides in a way that does not waste so much paper when printing out). The material is found in [**?**]. Being a handout-version of the slides, some figures and graph overlays may not be rendered in full detail, I remove most of the overlays, especially the long ones, because they don't make sense much on a handout/paper. Scroll through the real slides instead, if one needs the overlays.

This handout version also contains more remarks and footnotes, which would clutter the slides, and which typically contains remarks and elaborations, which may be given orally in the lecture.

Contents

1 Monitors (week 5) **2**

1.1 Semaphores & signalling disciplines 4
1.2 Bounded buffer 6
1.3 Readers/writers problem 7
1.4 Time server 8
1.5 Shortest-job-next scheduling 10
1.6 Sleeping barber 10

1 Monitors (week 5)

23. Sept. 2020

Overview

- Concurrent execution of different processes
- Communication by *shared variables*
- Processes may *interfere* `x := 0; co x := x + 1 || x := x + 2 oc` final value of `x` will be 1, 2, or 3
- **await** language – atomic regions `x := 0; co <x := x + 1> || <x := x + 2> oc` final value of `x` will be 3
- special tools for synchronization: Last week: semaphores Today: monitors

Outline

- Semaphores: review
- Monitors:
  - Main ideas
  - Syntax and semantics
    - Condition variables
    - Signaling disciplines for monitors
  - Synchronization problems:
    - Bounded buffer
    - Readers/writers
    - Interval timer
    - Shortest-job next scheduling
    - Sleeping barber

Semaphores

- Used as "synchronization variables"
- Declaration: $sem\ s = 1;$
- Manipulation: Only two operations, $P(s)$ and $V(s)$
- Advantage: Separation of "business" and synchronization code
- Disadvantage: Programming with semaphores can be tricky:[1]
  - Forgotten P or V operations
  - Too many P or V operations
  - They are shared between processes
    - Global knowledge
    - Need to examine all processes to see how a semaphore is intended

[1] Same may be said about simple locks.

Monitors

Monitor

"Abstract data type + synchronization"

- program *modules* with *more structure* than semaphores
- monitor encapsulates data, which can only be *observed* and *modified* by the monitor's procedures.
  - contains variables that describe the *state*
  - variables can be changed only through the available procedures
- implicit mutex: only 1 procedure may be active at a time.
  - A procedure: mutex access to the data in the monitor
  - 2 procedures in the same monitor: never executed concurrently
- cooperative scheduling
- Condition synchronization:[2] is given by *condition variables*
- At a lower level of abstraction: monitors can be implemented using locks or semaphores (for instance)

Usage

- process = active $\Leftrightarrow$ Monitor: = passive/re-active
- a procedure is *active,* if a statement in the procedure is executed by some process
- all shared variables: inside the monitor
- processes communicate by calling monitor procedures
- processes do not need to know all the implementation details
  - Only the visible effects of public procedures are important
- implementation can be changed, if visible effects remains
- Monitors and processes can be developed relatively independent $\Rightarrow$ Easier to understand and develop parallel programs

Syntax & semantics

```
monitor name {
  monitor variables    # permanent variables
  ## monitor invariant
  initialization
  procedures
}
```

Monitor: a form of abstract data type:

- *only* the procedures' names visible from outside the monitor:

  call *name.opname*(*arguments*)

- statements *inside* a monitor: *no* access to variables *outside* the monitor
- monitor variables: initialized before the monitor is used

Monitor Invariant: describe the monitor's inner states

- can be verified by sequential reasoning inside the monitor

[2] block a process until a particular condition holds.

Condition variables

- monitors contain *special* type of variables: **cond** (condition)
- used for synchronization/to delay processes
- each such variable is associated with a *wait condition*
- "*value*" of a condition variable: *queue* of delayed processes
- value: not directly accessible by programmer
- Instead, manipulated by special operations

```
cond cv;          # declares a condition variable cv
empty(cv);        # asks if the queue on cv is empty
wait(cv);         # causes process to wait in the cv queue
signal(cv);       # wakes up a process in the queue to cv
signal_all(cv);   # wakes up all processes in the cv queue
```

1.1 Semaphores & signalling disciplines

Implementation of semaphores

A monitor with P and V operations:

```
monitor Semaphore { # monitor invariant: s ≥ 0
  int s := 0        # value of the semaphore
  cond pos;         # wait condition

  procedure Psem() {
    while (s=0)  { wait (pos) };
    s := s - 1     }

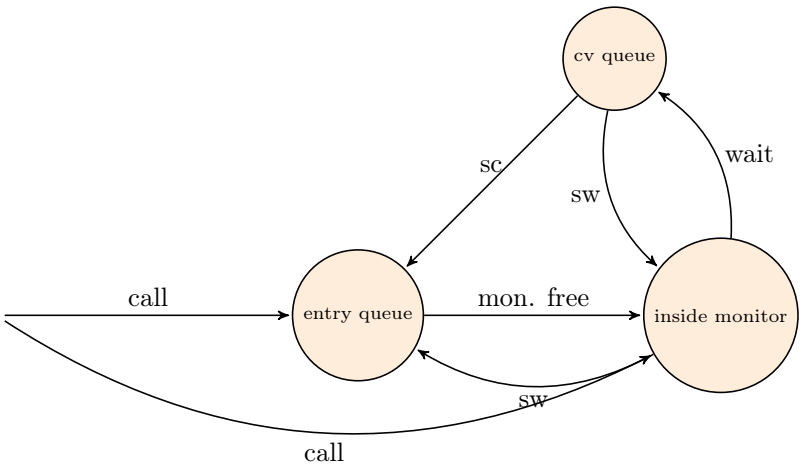
  procedure Vsem() {
    s := s+1;
    signal (pos);  }
}
```

Signaling disciplines

- signal on a condition variable `cv` roughly has the following effect:
  - empty queue: no effect
  - the process at the head of the queue to `cv` is woken up
- `wait` and `signal`: *FIFO signaling strategy*
- When a process executes `signal(cv)`, then it is inside the monitor. If a waiting process is woken up: *two active processes* in the monitor?

2 disciplines to provide mutex:

- Signal and Wait (SW): the signaller waits, and the signalled process gets to execute immediately
- Signal and Continue (SC): the signaller continues, and the signalled process executes later



Note: The figure is *schematic* and combines the "transitions" of signal-and-wait and signal-and-continue in a single diagram. The corresponding transition, here labelled SW and SC are the state changes caused by being *signalled* in the corresponding discipline.

Signalling disciplines

Is this a FIFO semaphore assuming SW or SC?

```
monitor Semaphore { # monitor invariant: s ≥ 0
  int s := 0          # value of the semaphore
  cond pos;           # wait condition

  procedure Psem() {
    while (s=0)  { wait (pos) };
    s := s - 1     }

  procedure Vsem() {
    s := s+1;
    signal (pos);  }
}
```

Signalling disciplines: FIFO semaphore for SW

```
monitor Semaphore { # monitor invariant: s ≥ 0
  int s := 0          # value of the semaphore
  cond pos;           # wait condition

  procedure Psem() {
    while (s=0)  { wait (pos) };
    s := s - 1     }

  procedure Vsem() {
    s := s+1;
    signal (pos);  }
}
```

```
monitor Semaphore { # monitor invariant: s ≥ 0
  int s := 0          # value of the semaphore
  cond pos;           # wait condition

  procedure Psem() {
    if   (s=0)  { wait (pos) };
    s := s - 1     }

  procedure Vsem() {
    s := s+1;
    signal (pos);  }
}
```

FIFO semaphore

FIFO semaphore with SC: can be achieved by explicit transfer of control inside the monitor (forward the condition).

```
monitor Semaphore        { # monitor invariant: s ≥ 0
  int s := 0;              # value of the semaphore
  cond pos;                # wait condition

  procedure Psem() {
    if   (s=0)
         wait (pos);
    else
         s := s - 1
  }

  procedure Vsem() {
    if   empty(pos)
         s := s + 1
    else
         signal(pos);
  }
}
```

1.2 Bounded buffer

Bounded buffer synchronization (1)

- buffer of size n ("channel", "pipe")
- producer: performs `put` operations on the buffer.
- consumer: performs `get` operations on the buffer.
- `count`: number of items in the buffer
- two access operations ("methods")
  - `put` operations must wait if buffer full
  - `get` operations must wait if buffer empty
- assume SC discipline[3]

Bounded buffer synchronization (2)

- When a process is woken up, it goes back to the monitor's entry queue
  - Competes with other processes for entry to the monitor
  - Arbitrary delay between awakening and start of execution

  $\Longrightarrow$ *re-test* the wait condition, when execution starts
  - E.g.: `put` process wakes up when the buffer is not full
    - Other processes can perform `put` operations before the awakened process starts up
    - Must therefore re-check that the buffer is not full

Bounded buffer synchronization monitors (3)

```
monitor Bounded_Buffer {
   typeT buf[n]; int count := 0;
   cond not_full, not_empty;

   procedure put(typeT data){
      while (count = n) wait(not_full);
      # Put element into buf
      count := count + 1; signal(not_empty);
   }

   procedure get(typeT &result) {
      while (count = 0) wait(not_empty);
      # Get element from buf
      count := count - 1; signal(not_full);
   }
}
```

[3] It's the commonly used one in practical languages/OS.

Bounded buffer synchronization: client-sides

```
process Producer[i = 1 to M]{
        while (true){
            ...
            call Bounded_Buffer.put(data);
        }
}
process Consumer[i = 1 to N]{
        while (true){
            ...
            call Bounded_Buffer.get(result);
        }
}
```

1.3 Readers/writers problem

Readers/writers problem

- Reader and writer processes share a common resource ("database")
- Reader's transactions can read data from the DB
- Write transactions can read and update data in the DB
- Assume:
  - DB is initially consistent and that
  - Each transaction, seen in isolation, maintains consistency
- To avoid interference between transactions, we require that
  - writers: exclusive access to the DB.
  - No writer: an arbitrary number of readers can access simultaneously

Monitor solution to the reader/writer problem (2)

- database should not be encapsulated in a monitor, as the readers will not get shared access
- monitor instead regulates access of the processes
- processes don't enter the critical section (DB) until they have passed the `RW_Controller` monitor

Monitor procedures:

- `request_read`: requests read access
- `release_read`: reader leaves DB
- `request_write`: requests write access
- `release_write`: writer leaves DB

Invariants and signalling

Assume that we have two counters as local variables in the monitor:

`nr` — number of readers
`nw` — number of writers

Invariant

$$\texttt{RW: (nr = 0 or nw = 0) and nw} \leq 1$$

We want `RW` to be a *monitor invariant*

- chose carefully condition variables for "communication" (waiting/signaling)

Let two condition variables `oktoread` and `oktowrite` regulate waiting readers and waiting writers, respectively.

```
monitor RW_Controller {  # RW  (nr = 0 or nw = 0) and nw ≤ 1
  int nr:=0, nw:=0
  cond oktoread ;    # signalled when nw = 0
  cond oktowrite;    # sig'ed when nr = 0 and nw = 0

  procedure request_read() {
    while (nw > 0) wait(oktoread);
    nr := nr + 1;
  }
  procedure release_read() {
    nr := nr - 1;
    if nr = 0 signal (oktowrite);
  }

  procedure request_write() {
    while (nr > 0 or nw > 0) wait(oktowrite);
    nw := nw + 1;
  }

  procedure release_write() {
    nw := nw -1;
    signal(oktowrite);      # wake up 1 writer
    signal_all(oktoread);   # wake up all readers
  }
}
```

Invariant

- monitor invariant I: describe the monitor's inner state
- expresses relationship between monitor variables
- maintained by execution of procedures:
  - must hold: after initialization
  - must hold: when a procedure terminates
  - must hold: when we suspend execution due to a call to `wait`

  ⇒ can assume that the invariant holds *after* `wait` and when a procedure starts
- Should be as *strong* as possible

Monitor solution to reader/writer problem (6)

$$RW\text{: (nr} = 0 \text{ or nw} = 0) \text{ and nw} \leq 1$$

```
procedure request_read() {
      # May assume that the invariant holds here
      while (nw > 0) {
            # the invariant holds here
            wait(oktoread);
            # May assume that the invariant holds here
      }
      # Here, we know that nw = 0...
      nr := nr + 1;
      # ...thus:  invariant also holds after increasing nr
}
```

Do we need $nr \geq 0$ and $nw \geq 0$?

Do we need to protect release_read and release_write?

1.4 Time server

Time server

- Monitor that enables sleeping for a given amount of time
- Resource: a logical clock (`tod`)
- Provides two operations:

  - `delay(interval)`: caller wishes to sleep for `interval` time
  - `tick` increments the logical clock with one tick Called by the hardware, preferably with high execution priority
- Each process which calls `delay` computes its own time for wakeup: `wake_time := tod + interval;`
- Waits as long as `tod < wake_time`
  - Wait condition is dependent on local variables

Covering condition:

- all processes are woken up when it is possible for some to continue
- Each process checks its condition and sleeps again if this does not hold

Time server: covering condition

Invariant: $CLOCK$: `tod` $\geq 0 \wedge$ `tod` increases monotonically by 1

```
monitor Timer { int tod = 0; # Time Of Day
    cond check; # signalled when tod is increased

    procedure delay(int interval) {
        int wake_time;
        wake_time = tod + interval;
        while (wake_time > tod) wait(check);
    }

    procedure tick() {
        tod = tod + 1;
        signal_all(check);
    }
}
```

- Not very efficient if many processes will wait for a long time
- Can give many false alarms

Prioritized waiting

- Can also give additional argument to `wait`: `wait(cv, rank)`
  - Process waits in the queue to `cv` in ordered by the argument `rank`.
  - At `signal`: Process with lowest `rank` is awakened first
- Call to `minrank(cv)` returns the value of `rank` to the first process in the queue (with the lowest rank)
  - The queue is not modified (no process is awakened)
- Allows more efficient implementation of `Timer`

Time server: Prioritized wait

- Uses prioritized waiting to order processes by `check`
- The process is awakened only when `tod` $\geq$ `wake_time`
- Thus we do not need a `while` loop for `delay`

```
monitor Timer {
    int tod = 0; # Invariant: CLOCK
    cond check; # signalled when minrank(check) ≤ tod

    procedure delay(int interval) {
        int wake_time;
        wake_time := tod + interval;
        if (wake_time > tod) wait(check, wake_time);
    }

    procedure tick() {
        tod := tod + 1;
        while (!empty(check) && minrank(check) ≤ tod)
        signal(check);
    }
}
```

1.5 Shortest-job-next scheduling

Shortest-Job-Next allocation

- Competition for a shared resource
- A monitor administrates access to the resource
- Call to `request(time)`
  - Caller needs access for time interval `time`
  - If the resource is free: caller gets access directly
- Call to `release`
  - The resource is released
  - If waiting processes: The resource is allocated to the waiting process with lowest value of `time`
- Implemented by prioritized wait

Shortest-Job-Next allocation (2)

```
monitor Shortest_Job_Next {
 bool free = true;
 cond turn;

 procedure request(int time) {
   if (free)
     free := false
   else
     wait(turn,time)
   }

  procedure release() {
    if (empty(turn))
      free := true;
    else
      signal(turn);
}
```

1.6 Sleeping barber



The story of the sleeping barber

- barbershop: with two doors and some chairs.
- customers: come in through one door and leave through the other. Only one customer sits in the barber chair at a time.
- Without customers: barber sleeps in one of the chairs.
- When a customer arrives and the barber sleeps ⇒ barber is woken up and the customer takes a seat.
- barber busy ⇒ the customer takes a nap
- Once served, barber lets customer out the exit door.
- If there are waiting customers, one of these is woken up. Otherwise the barber sleeps again.

Interface

Assume the following *monitor procedures*

Client: `get_haircut`: called by the customer, returns when haircut is done

Server: barber calls:

- `get_next_customer`: called by the barber to serve a customer
- `finish_haircut`: called by the barber to let a customer out of the barbershop

Rendezvous

Similar to a two-process barrier: *Both* parties must arrive before either can continue.[4]

- The barber must wait for a customer
- Customer must wait until the barber is available

The barber can have rendezvous with an arbitrary customer.

Organize the sync.: Identify the synchronization needs

1. barber must wait until
   (a) customer sits in chair
   (b) customer left barbershop
2. customer must wait until
   (a) the barber is available
   (b) the barber opens the exit door

client perspective:

- two phases (during `get_haircut`)
  1. "entering"
     - trying to get hold of barber,
     - sleep otherwise
  2. "leaving":
- between the phases: suspended

Processes signal when one of the wait conditions is satisfied.

Organize the synchronization: state

3 var's to synchronize the processes:
`barber, chair` and `open` (initially 0)

binary variables, alternating between 0 and 1:

- for entry-rendezvous
  1. `barber = 1` : the barber is ready for a new customer
  2. `chair = 1`: the customer sits in a chair, the barber hasn't begun to work
- for exit-sync
  3. `open = 1`: exit door is open, the customer has not yet left

[4]Later, in the context of message passing, will have a closer look at making rendezvous synchronization (using channels), but the pattern "2 partners must be present at a point at the same time" is analogous.

Sleeping barber

```
monitor Barber_Shop {
  int barber := 0, chair := 0, open := 0;
  cond barber_available;          # signalled when barber > 0
  cond chair_occupied;            # signalled when chair > 0
  cond door_open;                 # signalled when open > 0
  cond customer_left;             # signalled when open = 0

procedure get haircut() {
  while (barber = 0) wait(barber_available); # RV with barber
  barber := barber - 1;
  chair := chair + 1; signal(chair_occupied);

  while (open = 0) wait(door_open);          # leave shop
  open := open - 1; signal(customer_left);
}
procedure get next customer() {              # RV with client
  barber := barber + 1; signal (barber available);
  while (chair = 0) wait(chair_occupied);
  chair := chair - 1;
}
 procedure finished_cut() {
   open := open + 1; signal(door_open);      # get rid of customer
   while (open > 0) wait(customer_left);
}
```