



Monitors

Johannes Åman Pohjola
CSE, UNSW
Term 2 2022

Where we are at

Last lecture, we saw a generalisation of *locks* called *semaphores*.

In this lecture, we'll look at another concurrency abstraction, designed to ameliorate some problems with semaphores: *monitors*.

Main Disadvantages of Semaphores

- 1 **Lack of structure:** when building a large system, responsibility is diffused among implementers. Someone forgets to call `signal` \implies possible deadlock.

Main Disadvantages of Semaphores

- 1 **Lack of structure**: when building a large system, responsibility is diffused among implementers. Someone forgets to call `signal` \implies possible deadlock.
- 2 **Global visibility**: when something goes wrong, the whole program must be inspected \implies deadlocks are hard to isolate.

Main Disadvantages of Semaphores

- 1 **Lack of structure**: when building a large system, responsibility is diffused among implementers. Someone forgets to call `signal` \implies possible deadlock.
- 2 **Global visibility**: when something goes wrong, the whole program must be inspected \implies deadlocks are hard to isolate.

Solution

Monitors concentrate one responsibility into a single module and encapsulate critical resources.

They offer more structure than semaphores; more control than `await`.

Monitors

History:

- In the literature: Brinch Hansen (1973) and Hoare (1974)
- languages — Concurrent Pascal (1975)... Java, Pthreads library

Definition

Monitors are a generalisation of **objects** (as in OOP).

- May encapsulate some private data —all fields are private
- Exposes one or more *operations* — akin to methods.
- Implicit mutual exclusion—each operation invocation is **implicitly atomic**.
- Explicit signaling and waiting through *condition variables*.

Our Counting Example

Algorithm 2.1: Atomicity of monitor operations

monitor CS

integer $n \leftarrow 0$

operation increment

integer temp

temp $\leftarrow n$

$n \leftarrow \text{temp} + 1$

p

p1: **loop** ten times

p2: CS.increment

q

q1: **loop** ten times

q2: CS.increment

Program structure

monitor₁ . . . monitor_M
process₁ . . . process_N

- processes interact indirectly by using the same monitor
- processes call monitor procedures
- at most one call active in a monitor at a time — by definition
- explicit signaling using condition variables
- *monitor invariant*: predicate about local state that is true when no call is active

Condition variables

Definition

Condition variables are *named FIFO queues* of blocked processes.

Condition variables

Definition

Condition variables are *named FIFO queues* of blocked processes.

Processes executing a procedure of a monitor with condition variable *cv* can:

- voluntarily suspend themselves using **waitC**(*cv*),
- unblock the first suspended process by calling **signalC**(*cv*), or
- test for emptiness of the queue: **empty**(*cv*).

Condition variables

Definition

Condition variables are *named FIFO queues* of blocked processes.

Processes executing a procedure of a monitor with condition variable *cv* can:

- voluntarily suspend themselves using **waitC**(*cv*),
- unblock the first suspended process by calling **signalC**(*cv*), or
- test for emptiness of the queue: **empty**(*cv*).

Warning

The exact semantics of these differ between implementations!

Algorithm 2.2: Semaphore simulated with a monitor

monitor Sem

integer $s \leftarrow k$

condition notZero

operation wait

if $s = 0$

waitC(notZero)

$s \leftarrow s - 1$

operation signal

$s \leftarrow s + 1$

signalC(notZero)

p

q

loop forever

non-critical section

p1: Sem.wait

critical section

p2: Sem.signal

loop forever

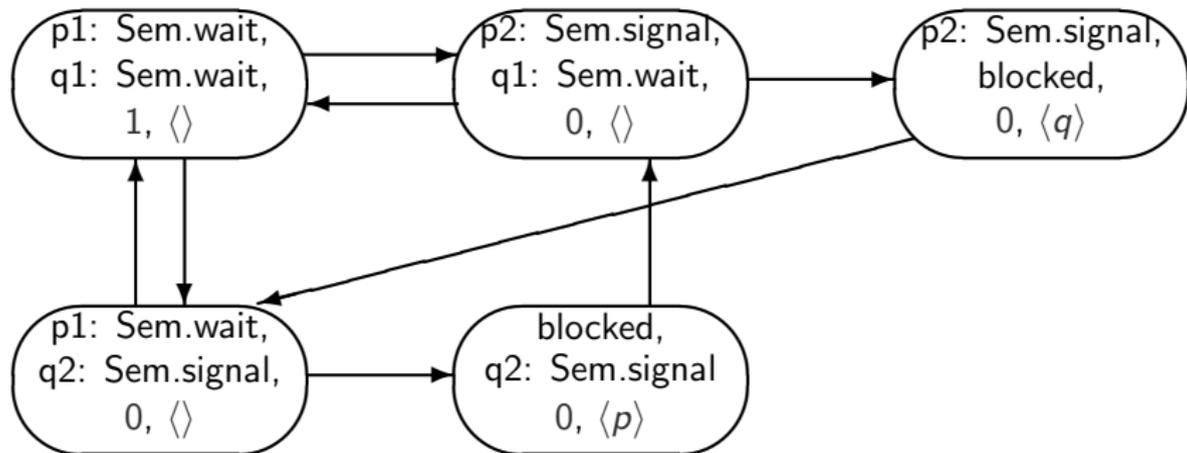
non-critical section

q1: Sem.wait

critical section

q2: Sem.signal

State Diagram for the Semaphore Simulation



Algorithm 2.3: Producer-consumer (finite buffer, monitor)**monitor** PCbufferType buffer \leftarrow empty**condition** notEmpty**condition** notFull**operation** append(datatype V) **if** buffer is full **waitC**(notFull)

append(V, buffer)

signalC(notEmpty)**operation** take()

datatype W

if buffer is empty **waitC**(notEmpty) W \leftarrow head(buffer) **signalC**(notFull) **return** W

Algorithm 2.3: Producer-consumer ... (continued)**producer**

datatype D

loop forever

p1: D ← produce

p2: PC.append(D)

consumer

datatype D

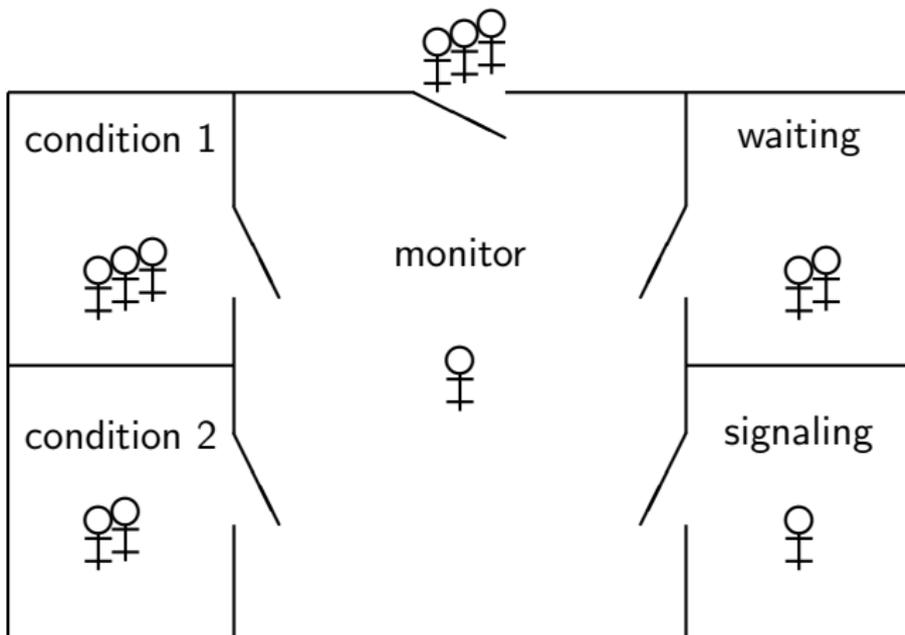
loop forever

q1: D ← PC.take

q2: consume(D)

The Immediate Resumption Requirement

Question: When a condition variable is signalled, who executes next? **It depends!**



Signaling disciplines

Precedences:

S the signaling process

W waiting on a condition variable

E waiting on entry

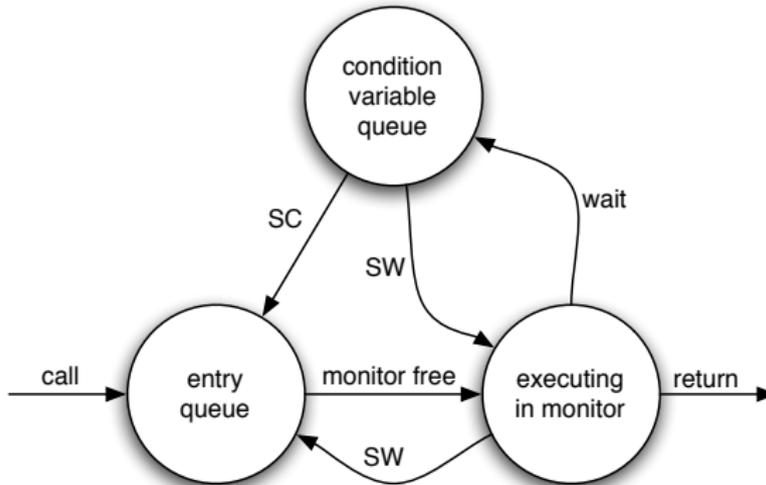
Signal and Urgent Wait

In Hoare's paper, $E < S < W$. This is also called the *immediate resumption requirement* (IRR). That is, a signalling process must wait for the **signalled** process to exit the monitor (or wait on a condition variable) before resuming. Signalling gives up control!

Signal and Continue

In Java, pthreads, and many other implementations, $E = W < S$. This means that signalling processes continue executing, and signalled processes await entry to the monitor at the same priority as everyone else.

Diagram for monitors



Simulating Monitors in Promela 1

```
1  bool lock = false;
2
3  typedef Condition {
4      bool gate;
5      byte waiting;
6  }
7  inline enterMon() {
8      atomic {
9          !lock;
10         lock = true;
11     }
12 }
13 inline leaveMon() {
14     lock = false;
15 }
```

Simulating Monitors in Promela 2

```
1  inline waitC(C) {
2      atomic {
3          C.waiting++;
4          lock = false; /* Exit monitor */
5          C.gate; /* Wait for gate */
6          lock = true; /* IRR */
7          C.gate = false; /* Reset gate */
8          C.waiting--;
9      }
10 }
```

Simulating Monitors in Promela 3

```

1  inline signalC(C) {
2      atomic {
3          if
4              /* Signal only if waiting */
5              :: (C.waiting > 0) ->
6                  C.gate = true;
7                  !lock; /* IRR - wait for released lock */
8                  lock = true; /* Take lock again */
9              :: else
10             fi;
11         }
12     }
13
14     #define emptyC(C) (C.waiting == 0)

```

Monitors in Java

An object in Java can be made to approximate a monitor with one waitset (i.e. unfair) condition variable and no immediate resumption:

- A method is made mutually exclusive using the `synchronized` keyword.

Monitors in Java

An object in Java can be made to approximate a monitor with one waitset (i.e. unfair) condition variable and no immediate resumption:

- A method is made mutually exclusive using the `synchronized` keyword.
- Synchronized methods of an object may call their `wait()` to suspend until `notify()` is called, analogous to condition variables.

Monitors in Java

An object in Java can be made to approximate a monitor with one waitset (i.e. unfair) condition variable and no immediate resumption:

- A method is made mutually exclusive using the `synchronized` keyword.
- Synchronized methods of an object may call their `wait()` to suspend until `notify()` is called, analogous to condition variables.
- **No immediate resumption requirement** means that waiting processes need to **re-check** their conditions!

Monitors in Java

An object in Java can be made to approximate a monitor with one waitset (i.e. unfair) condition variable and no immediate resumption:

- A method is made mutually exclusive using the `synchronized` keyword.
- Synchronized methods of an object may call their `wait()` to suspend until `notify()` is called, analogous to condition variables.
- **No immediate resumption requirement** means that waiting processes need to **re-check** their conditions!
- **No strong fairness guarantee** about wait lists, meaning any arbitrary waiting process is awoken by `notify()`.

Monitors in Java

An object in Java can be made to approximate a monitor with one waitset (i.e. unfair) condition variable and no immediate resumption:

- A method is made mutually exclusive using the `synchronized` keyword.
- Synchronized methods of an object may call their `wait()` to suspend until `notify()` is called, analogous to condition variables.
- **No immediate resumption requirement** means that waiting processes need to **re-check** their conditions!
- **No strong fairness guarantee** about wait lists, meaning any arbitrary waiting process is awoken by `notify()`.

Resources for Java Programming

See also Vladimir's videos introducing concurrent programming in Java, available on the course website.

Shared Data

Consider the *Readers and Writers* problem, common in any database:

Problem

We have a **large data structure** which cannot be updated in one atomic step. It is shared between many writers and many readers.

Shared Data

Consider the *Readers and Writers* problem, common in any database:

Problem

We have a **large data structure** which cannot be updated in one atomic step. It is shared between many writers and many readers.

Desiderata:

- *Atomicity*. An update should happen in one go, and updates-in-progress or partial updates are not observable.

Shared Data

Consider the *Readers and Writers* problem, common in any database:

Problem

We have a **large data structure** which cannot be updated in one atomic step. It is shared between many writers and many readers.

Desiderata:

- *Atomicity*. An update should happen in one go, and updates-in-progress or partial updates are not observable.
- *Consistency*. Any reader that starts after an update finishes will see that update.

Shared Data

Consider the *Readers and Writers* problem, common in any database:

Problem

We have a **large data structure** which cannot be updated in one atomic step. It is shared between many writers and many readers.

Desiderata:

- *Atomicity*. An update should happen in one go, and updates-in-progress or partial updates are not observable.
- *Consistency*. Any reader that starts after an update finishes will see that update.
- Minimal *waiting*.

A Crappy Solution

Treat both reads and updates as critical sections — use any old critical section solution to sequentialise all reads and writes to the data structure.

A Crappy Solution

Treat both reads and updates as critical sections — use any old critical section solution to sequentialise all reads and writes to the data structure.

Observation

Updates are *atomic* and reads are *consistent* — but reads can't happen concurrently, which leads to unnecessary *contention*.

A Better Solution

A *monitor* with two condition variables (à la Ben-Ari chapter 7).

Requirements

- 1 *Atomicity* and *consistency* (still)
- 2 Multiple reads can execute concurrently.
- 3 If someone writes: no concurrent reads or writes.

Algorithm 2.4: Readers and writers with a monitor**monitor** RWinteger readers \leftarrow 0integer writers \leftarrow 0**condition** OKtoRead, OKtoWrite**operation** StartRead **if** writers \neq 0 or not **empty**(OKtoWrite) **waitC**(OKtoRead) readers \leftarrow readers + 1 **signalC**(OKtoRead)**operation** EndRead readers \leftarrow readers - 1 **if** readers = 0 **signalC**(OKtoWrite)

Algorithm 2.4: Readers and writers with a monitor (continued)**operation** StartWrite**if** writers \neq 0 or readers \neq 0**waitC**(OKtoWrite)writers \leftarrow writers + 1**operation** EndWritewriters \leftarrow writers - 1**if empty**(OKtoRead)**then signalC**(OKtoWrite)**else signalC**(OKtoRead)**reader****writer**

p1: RW.StartRead

p2: *read the database*

p3: RW.EndRead

q1: RW.StartWrite

q2: *write to the database*

q3: RW.EndWrite

Proving Atomicity

Essentially we desire **mutual exclusion** of writers with any other process (writer or reader).

Proving Atomicity

Essentially we desire **mutual exclusion** of writers with any other process (writer or reader).

Like any safety property, we can prove it by **gathering invariants**. Let R be the number of active readers and W be the number of active writers.

Proving Atomicity

Essentially we desire **mutual exclusion** of writers with any other process (writer or reader).

Like any safety property, we can prove it by **gathering invariants**.

Let R be the number of active readers and W be the number of active writers.

- readers = $R \geq 0$ and writers = $W \geq 0$, trivially.

Proving Atomicity

Essentially we desire **mutual exclusion** of writers with any other process (writer or reader).

Like any safety property, we can prove it by **gathering invariants**.

Let R be the number of active readers and W be the number of active writers.

- readers = $R \geq 0$ and writers = $W \geq 0$, trivially.
- $(R > 0 \Rightarrow W = 0) \wedge (W \leq 1) \wedge (W = 1 \Rightarrow R = 0)$.

This is preserved across the **eight** possible transitions in this system: the four monitor operations running unhindered, and the four partial operations resulting from a signal. See Ben-Ari p159 for details.

Proving Atomicity

Essentially we desire **mutual exclusion** of writers with any other process (writer or reader).

Like any safety property, we can prove it by **gathering invariants**.

Let R be the number of active readers and W be the number of active writers.

- readers = $R \geq 0$ and writers = $W \geq 0$, trivially.
- $(R > 0 \Rightarrow W = 0) \wedge (W \leq 1) \wedge (W = 1 \Rightarrow R = 0)$.

This is preserved across the **eight** possible transitions in this system: the four monitor operations running unhindered, and the four partial operations resulting from a signal. See Ben-Ari p159 for details.

Liveness Properties

We may also wish to prove some analogue of starvation freedom as Ben-Ari does on p160. This gets a bit handwavy. Without a concrete monitor implementation, it's hard to know whether starvation is possible!

Reading and Writing

Complication

Now suppose we don't want readers to wait (much) while an update is performed. Instead, we'd rather they get an *older version* of the data structure.

Reading and Writing

Complication

Now suppose we don't want readers to wait (much) while an update is performed. Instead, we'd rather they get an *older version* of the data structure.

Trick: A writer creates *their own local copy* of the data structure, and then updates the (shared) *pointer* to the data structure to point to their copy.

Johannes: Draw on the board

Reading and Writing

Complication

Now suppose we don't want readers to wait *(much)* while an update is performed. Instead, we'd rather they get an *older version* of the data structure.

Trick: A writer creates *their own local copy* of the data structure, and then updates the (shared) *pointer* to the data structure to point to their copy.

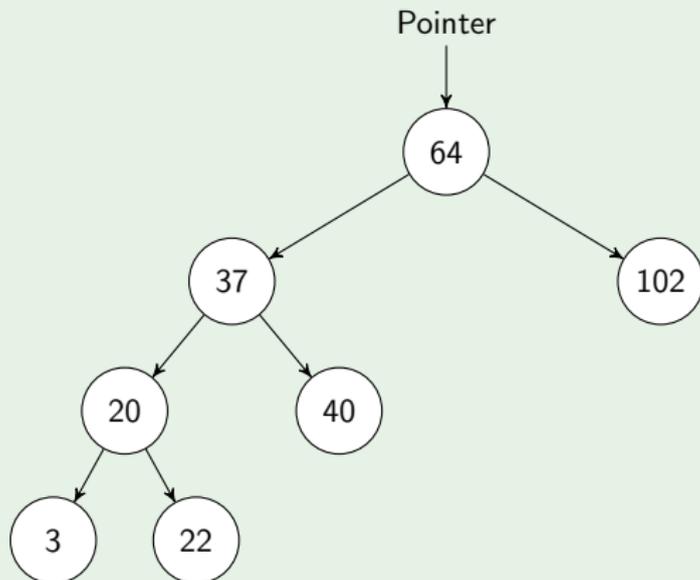
Johannes: Draw on the board

- Atomicity** The only shared write is now just to one pointer.
- Consistency** Reads that start before the pointer update get the older version, but reads that start after get the latest.

Persistent Data Structures

Copying is $\mathcal{O}(n)$ in the worst case, but we can do better for many tree-like types of data structure.

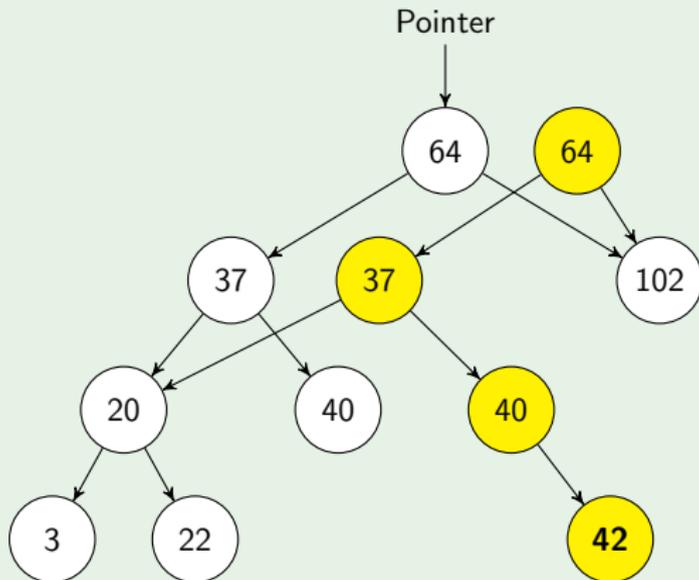
Example (Binary Search Tree)



Persistent Data Structures

Copying is $\mathcal{O}(n)$ in the worst case, but we can do better for many tree-like types of data structure.

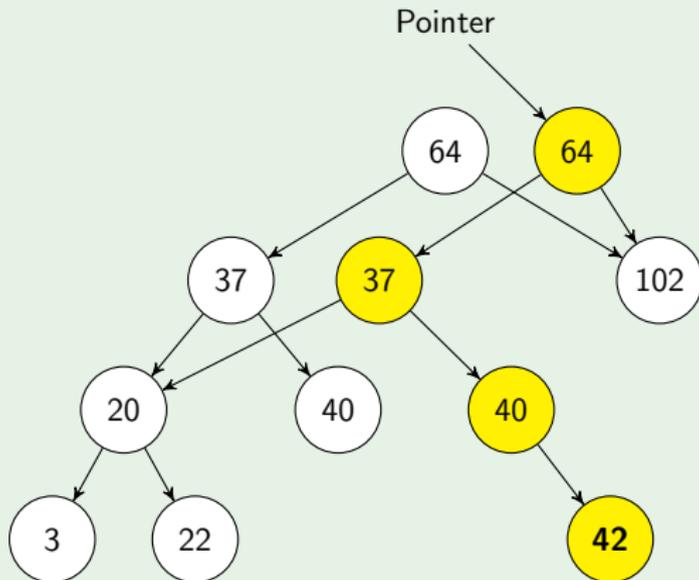
Example (Binary Search Tree)



Persistent Data Structures

Copying is $\mathcal{O}(n)$ in the worst case, but we can do better for many tree-like types of data structure.

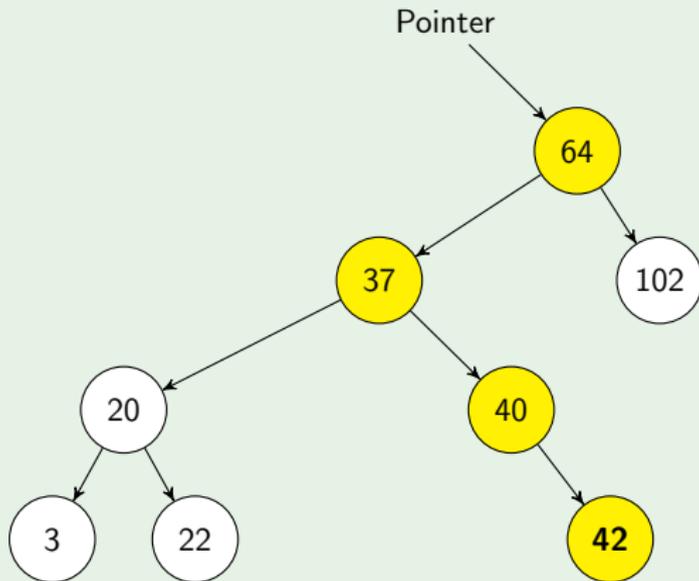
Example (Binary Search Tree)



Persistent Data Structures

Copying is $\mathcal{O}(n)$ in the worst case, but we can do better for many tree-like types of data structure.

Example (Binary Search Tree)



Purely Functional Data Structures

Persistent data structures that exclusively make use of copying over mutation are called *purely functional* data structures. They are so called because operations on them are best expressed in the form of mathematical functions that, given an input structure, return a *new* output structure:

$$\begin{aligned}
 \textit{insert } v \text{ Leaf} &= \text{Branch } v \text{ Leaf Leaf} \\
 \textit{insert } v (\text{Branch } x \text{ } l \text{ } r) &= \text{if } v \leq x \text{ then} \\
 &\quad \text{Branch } x (\textit{insert } v \text{ } l) \text{ } r \\
 &\quad \text{else} \\
 &\quad \text{Branch } x \text{ } l (\textit{insert } v \text{ } r)
 \end{aligned}$$

Purely functional programming languages like **Haskell** are designed to facilitate programming in this way.

What Now?

Next lecture, we'll be looking at **message-passing**, the foundation of distributed concurrency.

This homework involves Java programming. There are some resources to assist you on the course website.

Assignment 1 is out this week, hopefully tonight.