

3 UNSW

Test-and-Set

• We can use test-and-set to implement lock() and unlock() primitives

#### • Pros

- Simple (easy to show it's correct)
- Available at user-level
- To any number of processors
- To implement any number of lock variables

#### • Cons

- Busy waits (also termed a spin lock)

  - · Starvation is possible when a process leaves its critical section and more than one process is waiting.

#### to unblock the sleeping process • Waking a ready/running process has no effect.

busy waiting.

Sleep / Wakeup

• The idea

Tackling the Busy-Wait Problem

· When process is waiting for an event, it calls sleep to block, instead of

• The event happens, the event generator (another process) calls wakeup

- Consumes CPU

5 UNSW



4 JUNSW

# The Producer-Consumer Problem

- Also called the bounded buffer problem
- A producer produces data items and stores the items in a buffer
- A consumer takes the items out of the buffer and consumes them.



#### Issues

- We must keep an accurate count of items in buffer • Producer
  - should sleep when the buffer is full,
    - and wakeup when there is empty space in the buffer
  - The consumer can call wakeup when it consumes the first entry of the full buffer Consumer
  - should sleep when the buffer is empty
  - and wake up when there are items available
  - Producer can call wakeup when it adds the first item to the buffer



#### Pseudo-code for producer and consumer

<pre>con() {   while(TRUE) {     if (count == 0)         sleep(con);     remove_item();     count;     if (count == N-1)         wakeup(prod);   } }</pre>
9 🛃 U <u>NS</u> W

# Problems



#### Problems



#### **Proposed Solution**

• Lets use a locking primitive based on test-and-set to protect the concurrent access



# Proposed solution?





<ul> <li>The test for some condition and actually going to sleep needs to be atomic</li> <li>The following does not work:</li> </ul>	The lock is held while asleep ⇒ count will never change	<ul> <li>Dijkstra (1965) introduced two primitives that are more powerful than simple sleep and wakeup alone.</li> <li>P(): proberen, from Dutch to test.</li> <li>V(): verhogen, from Dutch to increment.</li> <li>Also called <i>wait</i> &amp; <i>signal</i>, <i>down</i> &amp; <i>up</i>.</li> </ul>
acquire_lock(buf_lock)	acquire_lock(buf_lock)	
if (count == N)	if (count == N - 1)	
<pre>sleep(prod);</pre>	wakeup(prod);	
<pre> release_lock(buf_lock)</pre>	<pre> release_lock(buf_lock)</pre>	
	15 🤯 UNSW	16

#### How do they work

- If a resource is not available, the corresponding semaphore blocks any process waiting for the resource
- Blocked processes are put into a process queue maintained by the semaphore (avoids busy waiting!)
- When a process releases a resource, it signals this by means of the semaphore
- Signalling resumes a blocked process if there is any, or stores the signal to be read by the next waiting task
- Wait (P) and signal (V) operations cannot be interrupted
- Complex coordination can be implemented by multiple semaphores

#### Semaphore Implementation

Problematic execution sequence

- Define a semaphore as a record
  - typedef struct { int count:

Semaphores

- struct process \*L; } semaphore;
- Assume two simple operations:
  - sleep suspends the process that invokes it.
  - wakeup(P) resumes the execution of a blocked process P.

17 🛃 UNSW



16 🐺 UNSW



# Solving the producer-consumer problem with semaphores

#define N = 4	prod() { while(TRUE) {	con() { while(TRUE) {
<pre>semaphore mutex = 1;</pre>	<pre>item = produce(); wait(full); wait(empty); wait(mutex); wait(cutex);</pre>	<pre>wait(full); wait(mutex); </pre>
/* count empty slots */ semaphore empty = N;	<pre>walt(mutex); insert_item(); signal(mutex);</pre>	<pre>signal(mutex); signal(empty);</pre>
/* count full slots */ semaphore full = 0;	signal(full); } }	}

21 🐺 UNSW

#### Summarising Semaphores

- Semaphores can be used to solve a variety of concurrency problems
- However, programming with them can be error-prone • E.g. must *signal* for every *wait* for mutexes
  - Too many, or too few signals or waits, or signals and waits in the wrong order, can have catastrophic results

#### Monitors

• To ease concurrent programming, Hoare (1974) proposed *monitors*.

Solving the producer-consumer problem

with semaphores

- A higher level synchronisation primitive
- Programming language construct

• Idea

- A set of procedures, variables, data types are grouped in a special kind of module, a *monitor*.
- Variables and data types only accessed from within the monitor
- Only one process/thread can be in the monitor at any one time
  - Mutual exclusion is implemented by the compiler (which should be less error prone)

23 🐺 UNSW



22

UNSW

# Monitor





### Simple example

monitor counter {	Note: "paper"
<pre>int count; procedure inc() {     count = count + 1;</pre>	• Compiler guar one thread ca the monitor a
<pre>} procedure dec() {     count = count -1;</pre>	• Easy to see th mutual exclus • No race cond
} }	• For instance, s methods in Ja

- te: "paper" language ompiler guarantees only ne thread can be active in ne monitor at any one time
- asy to see this provides utual exclusion
- No race condition on **count**.

UNSW

29 🐺 UNSW

27

For instance, **synchronized** methods in Java.

#### How do we block waiting for an event?

- We can use locks to block waiting for an object, held by another task
- We can use semaphores to solve the producer/consumer problem directly
- We would like a mechanism to block waiting for a kind of event (and also respect mutual exclusion)
  - e.g. in the producer-consumer problem
  - A blocked consumer is not waiting on just one producer
- Condition Variables



# **Condition Variable**

• To allow a process to wait within the monitor, a **condition** variable must be declared, as

#### condition x, y;

• Condition variable can only be used with the operations **wait** and **signal**.

#### • The operation

- x.wait();
- means that the process invoking this operation is suspended until another process invokes
   Another thread can enter the monitor while original is suspended
- x.signal(); • The x.signal operation resumes exactly one suspended process. If no process is

#### suspended, then the signal operation has no effect.

# **Condition Variables**



#### Monitors

#### onitor ProducerConsumer procedure producer; begin while true do condition full, empty; integer count; begin item = produce\_item; ProducerConsumer.insert(item) insert item(item): end count := count + 1; if count = 1 then signal(empty) end procedure consumer; end: begin function remove: integer; while true do begin if count = 0 then wait(empty); begin item = ProducerConsumer.remove; remove = $remove\_item;$ count := count - 1;if count = N - 1 then signal(full) consume\_item(item) end end; end count := 0;end monitor • Outline of producer-consumer problem with monitors • only one monitor procedure active at one time

31

• buffer has N slots

• Locks

- Semaphores
- Condition Variables

OS/161 Provided Synchronisation Primitives

32 🛃 UNSW

34 🛃 UNSW

36 🐺 UNSW

#### Locks

#### • Functions to create and destroy locks

struct lock \*lock\_create(const char \*name); void lock\_destroy(struct lock \*);

• Functions to acquire and release them

void	lock_acquire(struct	lock *);
void	lock_release(struct	lock *);

#### Example use of locks

int count; struct lock \*count\_lock main() { count = 0; count\_lock = lock\_create("count lock"); if (count\_lock == NULL) panic("I'm dead"); stuff(); }

```
lock_acquire(count_lock);
count = count + 1;
lock_release(count_lock);
}
procedure dec() {
lock_acquire(count_lock);
count = count -1;
lock_release(count_lock);
}
```

procedure inc() {

33 **UNSW** 

#### Semaphores

struct semaphore	<pre>*sem_create(const char *name, int initial_count);</pre>		
void	<pre>sem_destroy(struct semaphore *);</pre>		
void void	P(struct semaphore *); V(struct semaphore *);		

#### Example use of Semaphores

procedure inc() { int count; struct semaphore
 \*count\_mutex; P(count\_mutex); count = count + 1; V(count\_mutex); main() { } count = 0;procedure dec() { count\_mutex = P(count\_mutex); sem\_create("count", count = count -1; 1); V(count\_mutex); if (count\_mutex == NULL) } panic("I'm dead"); stuff(); }

35 **UNSW** 

# **Condition Variables**



#### **Dining Philosophers**

#define N #define LEFT #define RIGHT	5 (i+N–1)%N (i+1)%N	/* number of philosophers */ /* number of i's left neighbor */ /* number of i's right neighbor */		#define N 5	/* number of philosophers */
#define THINKING #define HUNGRY #define EATING typedef int semapho int state[N]; semaphore mutex = semaphore s[N]:	0 1 2 re; 1;	/* philosopher is thinking */ /* philosopher is trying to get forks */ /* philosopher is eating */ /* semaphores are a special kind of int */ /* array to keep track of everyone's state */ /* mutual exclusion for critical regions */ /* one semaphore per philosopher */		<pre>void philosopher(int i) {     while (TRUE) {         think();         take_fork(i);         take_fork(ii,1) % N);     } }</pre>	/* i: philosopher number, from 0 to 4 */ /* philosopher is thinking */ /* take left fork */
void philosopher(int { while (TRUE) { think(); take_forks eat(); put_forks(	i) (i); i);	/* i: philosopher number, from 0 to N-1 */ /* repeat forever */ /* philosopher is thinking */ /* acquire two forks or block */ /* yum-yum, spaghetti */ /* put both forks back on table */		eat(); put_fork((i+1) % N); put_fork((i+1) % N); }	/* yum-yum, spaghetti */ /* put left fork back on the table */ /* put right fork back on the table */
}				A <u>non</u> solution to t	he dining philosophers problem
Solution t	o dining pl	nilosophers problem (part 1)	<mark>/  </mark>		42 🗸 UNSW

**Dining Philosophers** 

## **Dining Philosophers**





# • Models access to a database

- E.g. airline reservation system
  Can have more than one concurrent reader
- To check schedules and reservations
- Writers must have exclusive access
- To book a ticket or update a schedule

#### The Readers and Writers Problem



**NSW** 

45 🛃 UNSW

# The Sleeping Barber Problem



#### The Sleeping Barber Problem



#### • Counting semaphores versus binary semaphores:

- In a counting semaphore, count can take arbitrary integer values
- In a binary semaphore, *count* can only be 0 or 1
- Can be easier to implement
- Counting semaphores can be implemented in terms of binary semaphores (how?)
- Strong semaphores versus weak semaphores:
  - In a strong semaphore, the *queue* adheres to the FIFO policy
  - In a weak semaphore, any process may be taken from the queue
  - Strong semaphores can be implemented in terms of weak semaphores (how?)



#### Two processes want to access shared memory at same time



# Making Single-Threaded Code Multithreaded



Conflicts between threads over the use of a global variable

51 UNSW

49 🐺 UNSW

#### A Producer-Consumer Solution Using OS/161 CVs

<pre>int count = 0; #define N 4 /* buf size */ prod() { while(TRUE) { item = produce() lock_aquire(l) while (count == N) cv_wait(full,l); insert_item(item); count++; if (count == 1) cv_signal(empty,l); lock_release(l) } }</pre>	<pre>con() {   while(TRUE) {     lock_acquire(l)     while (count == 0)         cv_wait(empty,l);     item = remove_item();     count;     if (count == N-1)         cv_signal(full,l);     lock_release(l);     consume(item);   } }</pre>
	52 🛃 UNSW

#### Peterson's Solution

- For the curious
- Avoids strict alternation • see the textbook
  - or Internet

FYI