### Course 02158

## **Barriers and Memory Models**

Hans Henrik Løvengreen

DTU Compute

# 

### **Barrier with Shared Variables**

```
Double barrier
```

```
var count<sub>1</sub>, count<sub>2</sub> : integer := 0;
process Worker[i : 1..n]
loop
do task;
  (count<sub>1</sub> := count<sub>1</sub> + 1);
  if count<sub>1</sub> = n then count<sub>1</sub> := 0;
  await count<sub>1</sub> = 0
  (count<sub>2</sub> := count<sub>2</sub> + 1);
  if count<sub>2</sub> = n then count<sub>2</sub> := 0;
  await count<sub>2</sub> = 0
end loop
```

#### **Barrier with Shared Variables** Coordinator-based • **var** arrive[1..n] : int := 0; continue[1..n] : bool := 0; • process Worker[i : 1...n] process Coordinator loop loop **for** *i* : 1..*n* **do** do task<sub>i</sub> { **await** arrive[i] = 1; arrive[i] := 0 } arrive[i] := 1;await continue[i] = 1;**for** *i* : 1..*n* **do** continue[i] := 0;continue[i] := 1;end loop end loop

• General principle: Process awaiting flag resets it



### **Caveats of Low-level Sharing**

- Consider **var** x, y : int; x := 1; y := 2
- x and y may not be assigned as expected:
  - The *compiler* may reorder the assignments
  - The processor may reorder the store instructions
  - ► The *memory system* may reorder the write operations
  - The compiler/processor may keep values in registers
- Proper ordring must be ensured by *synchronization instructions*

#### Java low-level memory model

- Access to variables (except long and double) is *atomic*
- Joining a thread ensures that its memory operations are completed
- Monitor entry/exit synchronize memory operations
- Access to volatile variables is always synchronized

### **Linux Spin Locks**

```
Linux 2.6.22
```

```
• static inline void __raw_spin_lock(raw_spinlock_t *lock)
{
    asm volatile("\n1:\t"
        LOCK_PREFIX " ; decb %0\n\t"
        "jns 3f\n"
        "2:\t"
        "rep;nop\n\t"
        "cmpb $0,%0\n\t"
        "jle 2b\n\t"
        "jmp 1b\n"
        "3:\n\t"
        : "+m" (lock->slock) : : "memory");
}
```



## Linux Spin Locks

```
Linux 2.6.22
```

```
• /*
 * Your basic SMP spinlocks, allowing only a single CPU anywhere
 *
 * Simple spin lock operations. There are two variants, one clears IRQ's
 * on the local processor, one does not.
 *
 * We make no fairness assumptions. They have a cost.
 *
 * (the type definitions are in asm/spinlock_types.h)
 */
```



### **Critical Region with Shared Variables**

#### **Ticket Algorithm**

# Linux Spin Locks

```
Linux 3.11
```

```
• static __always_inline void __ticket_spin_lock(arch_spinlock_t *lock)
{
    register struct __raw_tickets inc = { .tail = 1 };
    inc = xadd(&lock->tickets, inc);
    for (;;) {
        if (inc.head == inc.tail)
            break;
        cpu_relax();
        inc.head = ACCESS_ONCE(lock->tickets.head);
    }
    barrier(); /* make sure nothing creeps before the lock is taken */
}
```