



# Mutual Exclusion

Uwe & Zimmer - The Australian National University

## Mutual Exclusion

### Mutual exclusion: Atomic load & store operations

Assumption 1: every individual base memory cell (word) load and store access is atomic  
Assumption 2: there is no atomic combined load-store access

```
G : Natural := 0; -- assumed to be stored on a f-word cell in memory
task body P1 is
begin
  G := G + 1;
end P1;
task body P2 is
begin
  G := G + 1;
end P2;
task body P3 is
begin
  G := G + 1;
end P3;
```

What is the value of G?

## Mutual Exclusion

### Mutual exclusion: First attempt

```
type Task_Token is mod 2;
Turn : Task_Token := 0;
task body P1 is
begin
  loop
    non-critical_section_1;
    loop
      critical_section_1;
    end loop;
    turn := turn + 1;
  end loop;
end P1;
-- Mutual exclusion!
-- No deadlock!
-- No starvation!
-- Indefinite!
```

scatter  
if turn = n then  
 if turn = n then  
 and if  
 into non-critical sections

## Mutual Exclusion

### Mutual exclusion: Third attempt

```
type Critical_Section_State is (In_CS, Out_CS);
C1, C2 : Critical_Section_State := Out_CS;
task body P1 is
begin
  loop
    non-critical_section_1;
    loop
      In_CS;
    end loop;
    exit when C2 = Out_CS;
    loop
      critical_section_1;
    end loop;
    C1 := Out_CS;
  end loop;
end P1;
-- Mutual exclusion!
-- Potential deadlock!
```

## Mutual Exclusion

### References for this chapter

Robert Arfken,  
M. Bruniati  
Principles of Concurrent and Distributed Programming  
2006, second edition, Prentice-Hall, ISBN 0-13-711827-X

## Mutual Exclusion

### Mutual exclusion: Atomic load & store operations

Assumption 1: every individual base memory cell (word) load and store access is atomic  
Assumption 2: there is no atomic combined load-store access

```
G : Natural := 0; -- assumed to be stored on a f-word cell in memory
task body P1 is
begin
  G := G + 1;
end P1;
task body P2 is
begin
  G := G + 1;
end P2;
task body P3 is
begin
  G := G + 1;
end P3;
```

After the first global initialization, G can have almost any value, between 0 and 24  
After the first global initialization, G will have exactly one value between 0 and 24  
After all tasks terminated, G will have exactly one value between 2 and 24

## Mutual Exclusion

### Mutual exclusion: Second attempt

```
type Critical_Section_State is (In_CS, Out_CS);
C1, C2 : Critical_Section_State := Out_CS;
task body P1 is
begin
  loop
    non-critical_section_1;
    loop
      In_CS;
    end loop;
    exit when C2 = Out_CS;
    loop
      critical_section_1;
    end loop;
    C1 := Out_CS;
  end loop;
end P1;
-- Any better?
```

## Mutual Exclusion

### Mutual exclusion: Fourth attempt

```
type Critical_Section_State is (In_CS, Out_CS);
C1, C2 : Critical_Section_State := Out_CS;
task body P1 is
begin
  loop
    non-critical_section_1;
    loop
      In_CS;
    end loop;
    exit when C2 = Out_CS;
    loop
      critical_section_1;
    end loop;
    C1 := Out_CS;
  end loop;
end P1;
-- Making any progress?
```

## Mutual Exclusion

### Problem specification

### The general mutual exclusion scenario

N processes execute (finite) instruction sequences concurrently. Each instruction belongs to either a critical or non-critical section.  
Safety property 'Mutual exclusion':  
Instructions from critical sections of two or more processes must never be interleaved!

- More required properties:
- No deadlocks, if one or multiple processes try to enter their critical sections then exactly one of them must succeed.
- No starvation, every process which tries to enter one of its critical sections will eventually enter it.
- Efficiency, the fraction of time that processes spend in their critical sections must be as high as possible.

## Mutual Exclusion

### Mutual exclusion: First attempt

```
type Task_Token is mod 2;
Turn : Task_Token := 0;
task body P1 is
begin
  loop
    non-critical_section_1;
    loop
      critical_section_1;
    end loop;
    turn := turn + 1;
  end loop;
end P1;
-- Mutual exclusion!
-- No deadlock!
-- Starvation?
-- Work without contention?
```

## Mutual Exclusion

### Mutual exclusion: Second attempt

```
type Critical_Section_State is (In_CS, Out_CS);
C1, C2 : Critical_Section_State := Out_CS;
task body P1 is
begin
  loop
    non-critical_section_1;
    loop
      In_CS;
    end loop;
    exit when C2 = Out_CS;
    loop
      critical_section_1;
    end loop;
    C1 := Out_CS;
  end loop;
end P1;
-- No mutual exclusion!
```

## Mutual Exclusion

### Mutual exclusion: Fourth attempt

```
type Critical_Section_State is (In_CS, Out_CS);
C1, C2 : Critical_Section_State := Out_CS;
task body P1 is
begin
  loop
    non-critical_section_1;
    loop
      In_CS;
    end loop;
    exit when C2 = Out_CS;
    loop
      critical_section_1;
    end loop;
    C1 := Out_CS;
  end loop;
end P1;
-- Mutual exclusion!
-- Potential starvation!
-- Potential global deadlock!
```

## Mutual Exclusion

### Problem specification

### The general mutual exclusion scenario

N processes execute (finite) instruction sequences concurrently. Each instruction belongs to either a critical or non-critical section.  
Safety property 'Mutual exclusion':  
Instructions from critical sections of two or more processes must never be interleaved!

- Further assumptions:
- Pre- and post-protocols can be executed before and after each critical section.
- Processes may delay infinitely in non-critical sections.
- Processes do not delay infinitely in critical sections.

## Mutual Exclusion

### Mutual exclusion: First attempt

```
type Task_Token is mod 2;
Turn : Task_Token := 0;
task body P1 is
begin
  loop
    non-critical_section_1;
    loop
      critical_section_1;
    end loop;
    turn := turn + 1;
  end loop;
end P1;
-- Mutual exclusion!
-- No deadlock!
-- No starvation!
-- Locks up, if there is no contention!
```

## Mutual Exclusion

### Mutual exclusion: Third attempt

```
type Critical_Section_State is (In_CS, Out_CS);
C1, C2 : Critical_Section_State := Out_CS;
task body P1 is
begin
  loop
    non-critical_section_1;
    loop
      In_CS;
    end loop;
    exit when C2 = Out_CS;
    loop
      critical_section_1;
    end loop;
    C1 := Out_CS;
  end loop;
end P1;
-- Any better?
```

## Mutual Exclusion

### Mutual exclusion: Deekers' algorithm

```
type Task_Range is mod 2;
type Critical_Section_State is (In_CS, Out_CS);
Turn : Task_Range := Task_Range_First;
task type One_Of_Two_Tasks (this_Task : Task_Range);
CS : One_Of_Two_Tasks := Out_CS;
task body One_Of_Two_Tasks is
begin
  other_Task := this_Task + 1;
  begin
    non-critical_section
  end if;
  CS (this_Task) := In_CS;
  loop
    other_Task := Task_Range;
  end loop;
end if;
end One_Of_Two_Tasks;
turn := One_Of_Two_Tasks;
```

### Mutual Exclusion

**Mutual exclusion: Decker's Algorithm** \* Two tasks only!

```

type TaskRange is range 0..1;
type CriticalSectionState is (In_CS, Out_CS);
CSS : array (TaskRange) of CriticalSectionState := (others => Out_CS);
Turn : TaskRange := TaskRange.First;
loop
  exit when Turn = this_Task;
  CSS (this_Task) := In_CS;
  if Turn = other_Task then
    loop
      exit when Turn = this_Task;
    end loop;
  end if;
  CSS (this_Task) := Out_CS;
  Turn := other_Task;
end One_of_Two_Tasks;
  
```

task body one\_of\_two\_tasks is  
 other\_Task := TaskRange.First;  
 begin  
 ----- non-critical-section  
 \* Mutual exclusion! \* No starvation!  
 \* No deadlock! \* No livelock!  
 end One\_of\_Two\_Tasks;

### Mutual Exclusion

**Mutual exclusion: Peterson's Algorithm**

```

type TaskRange is range 0..1;
type CriticalSectionState is (In_CS, Out_CS);
CSS : array (TaskRange) of CriticalSectionState := (others => Out_CS);
Last : TaskRange := TaskRange.First;
task type One_of_Two_Tasks (this_Task : TaskRange);
task body one_of_two_tasks is
  other_Task := this_Task + 1;
  begin
  ----- non-critical-section
  or else last /= this_Task;
  end loop;
  CSS (this_Task) := Out_CS;
  Last := this_Task;
end One_of_Two_Tasks;
  
```

### Mutual Exclusion

**Mutual exclusion: Peterson's Algorithm** \* Two tasks only!

```

type TaskRange is range 0..1;
type CriticalSectionState is (In_CS, Out_CS);
CSS : array (TaskRange) of CriticalSectionState := (others => Out_CS);
Last : TaskRange := TaskRange.First;
task type One_of_Two_Tasks (this_Task : TaskRange);
task body one_of_two_tasks is
  other_Task := this_Task + 1;
  begin
  ----- non-critical-section
  or else last /= this_Task;
  end loop;
  CSS (this_Task) := Out_CS;
  Last := this_Task;
end One_of_Two_Tasks;
  
```

### Mutual Exclusion

**Problem specification**

**The general mutual exclusion scenario**

- Processes execute (finite) instruction sequences concurrently.
- Each instruction belongs to either a **critical** or **non-critical** section.
- Safety property** **Mutual exclusion**: Instructions from **critical** sections of two or more processes must never be interleaved!
- More required properties:
  - No deadlocks: If one or multiple processes try to enter their critical sections then exactly one of them must succeed.
  - No starvation: Every process which tries to enter one of its critical sections must eventually succeed.
  - Efficiency: The decision which process may enter the critical section must be made **efficiently** in all cases, i.e. also when there is no contention.

### Mutual Exclusion

**Mutual exclusion: Bakery Algorithm**

**The idea of the Bakery Algorithm**

A set of  $N$  Processes  $P_1, \dots, P_N$ , competing for mutually exclusive execution of their critical regions. Every process  $P_i$  (for  $i = 1, \dots, N$ ) supplies a globally readable number  $l_i$  (ticket) (initialized to 0).

- Before a process  $P_i$  enters a critical section:
  - $P_i$  draws a new number  $l_i > l_j \forall j \neq i$
  - $P_i$  is allowed to enter the critical section iff  $\forall j \neq i: l_j < l_i$  or  $j = 0$
  - After a process left a critical section:
    - $P_i$  resets its  $l_i = 0$

Issues

- Can you ensure that processes won't read each others ticket numbers while still calculating?
- Can you ensure that no two processes draw the same number?

### Mutual Exclusion

**Mutual exclusion: Bakery Algorithm**

```

No_of_Tasks : constant Positive := 1;
type TaskRange is range 0..No_of_Tasks;
Choosing : array (TaskRange) of Boolean := (others => false);
Ticket : array (TaskRange) of Natural := (others => 0);
task type P (this_Id : TaskRange);
task body P is
  begin
  ----- non-critical-section
  Ticket (this_Id) := 0;
  or else
  or else Ticket (this_Id) < Ticket (Id);
  end loop;
  Choosing (this_Id) := true;
  Choosing (this_Id) := false;
  end loop;
  if Id in this_Id then
    ----- critical-section;
    end loop;
  end P;
  
```

### Mutual Exclusion

**Beyond atomic memory access**

**Realistic hardware support**

Atomic test-and-set operations:

- $l := G \wedge C := 1$

Atomic exchange operations:

- $Temp := l; l := G; C := Temp$

Memory cell reservations:

- $l, C, C_{old}$  - read by using a **special instruction**, which puts a reservation on C
- $C_{old} := C$  - calculate a new value for C...
- $C := C_{old} \wedge new\_value$
- succeeds iff C was not manipulated by other processors or devices since the reservation

### Mutual Exclusion

**Mutual exclusion: atomic test-and-set operation**

```

type Flag is Natural range 0..1; C : Flag := 0;
task body P1 is
  L : Flag := 1;
  begin
  loop
    loop
      exit when L = 0;
    end loop;
    ----- change process
    end loop;
    C := 0;
  end loop;
  end P1;
  
```

Does that work?

### Mutual Exclusion

**Mutual exclusion: atomic test-and-set operation**

```

type Flag is Natural range 0..1; C : Flag := 0;
task body P1 is
  L : Flag := 1;
  begin
  loop
    loop
      exit when L = 0;
    end loop;
    ----- change process
    end loop;
    C := 0;
  end loop;
  end P1;
  
```

Does that work?

### Mutual Exclusion

**Mutual exclusion: atomic exchange operation**

```

type Flag is Natural range 0..1; C : Flag := 0;
task body P1 is
  L : Flag := 1;
  begin
  loop
    loop
      exit when L = 0;
    end loop;
    ----- change process
    end loop;
    L := 0;
  end loop;
  end P1;
  
```

Does that work?

### Mutual Exclusion

**Mutual exclusion: atomic exchange operation**

```

type Flag is Natural range 0..1; C : Flag := 0;
task body P1 is
  L : Flag := 1;
  begin
  loop
    loop
      exit when L = 0;
    end loop;
    ----- change process
    end loop;
    L := 0;
  end loop;
  end P1;
  
```

Mutual exclusion! No deadlock! No global livelock!  
 Works for any dynamic number of processes.  
 Individual starvation possible: busy waiting loops!

### Mutual Exclusion

**Mutual exclusion: memory cell reservation**

```

type Flag is Natural range 0..1; C : Flag := 0;
task body P1 is
  L : Flag := 0;
  begin
  loop
    loop
      exit when L = 0;
    end loop;
    ----- change process
    end loop;
    C := 0;
  end loop;
  end P1;
  
```

Does that work?

### Mutual Exclusion

**Mutual exclusion: memory cell reservation**

```

type Flag is Natural range 0..1; C : Flag := 0;
task body P1 is
  L : Flag := 0;
  begin
  loop
    loop
      exit when L = 0;
    end loop;
    ----- change process
    end loop;
    C := 0;
  end loop;
  end P1;
  
```

Does that work?

### Mutual Exclusion

**Mutual exclusion: memory cell reservation**

```

type Flag is Natural range 0..1; C : Flag := 0;
task body P1 is
  L : Flag := 0;
  begin
  loop
    loop
      exit when L = 0;
    end loop;
    ----- change process
    end loop;
    C := 0;
  end loop;
  end P1;
  
```

Does that work?

### Mutual Exclusion

**Mutual exclusion: memory cell reservation**

```

type Flag is Natural range 0..1; C : Flag := 0;
task body P1 is
  L : Flag := 0;
  begin
  loop
    loop
      exit when L = 0;
    end loop;
    ----- change process
    end loop;
    C := 0;
  end loop;
  end P1;
  
```

Does that work?



## Mutual Exclusion

### Semaphores

```

S1, S2 : Semaphore := 1;

task body P1 is
begin
loop
----- non_critical_section_1;
wait (S1);
----- critical_section_1;
signal (S2);
signal (S1);
end loop;
end P1;

task body P2 is
begin
loop
----- non_critical_section_2;
wait (S2);
----- critical_section_2;
signal (S1);
signal (S2);
end loop;
end P2;

```

⇨ Works too?

© 2003 Lee & Zilles, The Australian National University page 21 of 18 chapter 2: "Mutual Exclusion" up to page 251

## Mutual Exclusion

### Semaphores

```

S1, S2 : Semaphore := 1;

task body P1 is
begin
loop
----- non_critical_section_1;
wait (S1);
----- critical_section_1;
signal (S2);
signal (S1);
end loop;
end P1;

task body P2 is
begin
loop
----- non_critical_section_2;
wait (S2);
----- critical_section_2;
signal (S1);
signal (S2);
end loop;
end P2;

```

⇨ Mutual exclusion!, No global live-lock!

⇨ Works for any dynamic number of processes.

⇨ Individual starvation possible!

⇨ Deadlock possible!

© 2003 Lee & Zilles, The Australian National University page 22 of 18 chapter 2: "Mutual Exclusion" up to page 251

## Mutual Exclusion

### Summary

### Mutual Exclusion

- **Definition of mutual exclusion**
- **Atomic load and atomic store operations**
  - ... some classical errors
  - Decker's algorithm, Peterson's algorithm
  - Bakery algorithm
- **Realistic hardware support**
  - Atomic test-and-set, Atomic exchanges, Memory cell reservations
- **Semaphores**
  - Basic semaphore definition
  - Operating systems style semaphores

© 2003 Lee & Zilles, The Australian National University page 23 of 18 chapter 2: "Mutual Exclusion" up to page 251

