

Lesson 3

# Critical Section Problem

Ch 3 [BenA 06]

Critical Section Problem  
Solutions without HW Support  
State Diagrams for Algorithms  
Busy-Wait Solutions with HW Support

3.11.2008 Copyright Teemu Kerola 2008 1

## Mutual Exclusion Real World Example





Fig. Pesutuvan varaus

- How to reserve a laundry room?
  - Housing corporation with many tenants
- Reliable
  - No one else can reserve, once one reservation for given time slot is done
  - One can not remove other's reservations
- Reservation method
  - One can make decision independently (without discussing with others) on whether laundry room is available or not
  - One can have reservation for at most one time slot at a time
- People not needing the laundry room are not bothered
- One should not leave reservation on when moving out
- One should not lose reservation tokens/keys

3.11.2008 Copyright Teemu Kerola 2008 2



PESUTUVAN VARAUS

Taloyhtiön pesutuvan varaus toimii laittamalla varauslukko talle sopivan päivän ja kellonajan kohdalle varaustauluun.

Varauslukko tulee poistaa varauksen jälkeen tai mikäli ette käytä varaamaanne aikaa

Terveisin  
Isännöitsijä

Photo P. Niklander

3.11.2008 Copyright Teemu Kerola 2008 3

## Concurrent indivisible operations

- Echo
 

	Process P1	Process P2
char out, in; //globals	...	...
procedure echo {	input (in,..);	...
input (in, keyboard);	out = in;	input(in,..);
out = in;	output (out, display);	out = in;
output (out, display);	...	output (out,..);
}	output(out,..);	

  - What if *out* and/or *in* local variables?
- Data base update
  - Name, id, address, salary, annual salary, ...
- How/when/by whom to define granularity for indivisible operations?

3.11.2008 Copyright Teemu Kerola 2008 4

## Critical Section (CS)

- Mutex (mutual exclusion) solved **poissulkemisongelma ratk.**
- No deadlock: someone will succeed **ei lukkiutumista**
- No starvation (and no unnecessary delay) **ei nälkiintymistä**
  - Everyone succeeds eventually
- Protocol does not use common variables with CS actual work
  - Can use it's own local or shared variables

**Algorithm 3.1: Critical section problem**

global variables	
p	q
local variables	local variables
loop forever	loop forever
non-critical section	non-critical section
preprotocol	preprotocol
<b>critical section</b>	<b>critical section</b>
postprotocol	postprotocol

3.11.2008 Copyright Teemu Kerola 2008 5

## Critical Section Assumptions

**Algorithm 3.1: Critical section problem**

global variables	
p	q
local variables	local variables
loop forever	loop forever
non-critical section	non-critical section
preprotocol	preprotocol
<b>critical section</b>	<b>critical section</b>
postprotocol	postprotocol

← "unsafe zone" →

← "safe zone" →

- Preprotocol and postprotocol have **no common** local/global variables with critical/non-critical sections
  - They do not disturb/affect each other
- Non-critical section may stall or terminate
  - Can not assume it to complete
- Critical section will complete (will not terminate)
  - Postprotocol eventually executed once critical section is entered
- Process will not terminate in preprotocol or postprotocol (!!!)

3.11.2008 Copyright Teemu Kerola 2008 6

### Critical Section Solution

Algorithm 3.2: First attempt

integer turn ← 1	
p	q
loop forever	loop forever
p1: non-critical section	q1: non-critical section
p2: await turn = 1	q2: await turn = 2
p3: critical section	q3: critical section
p4: turn ← 2	q4: turn ← 1

- How to prove correct? (or incorrect?)
  - Mutex? (functional correct)
  - No deadlock? (eventually someone from many will get in)
  - No starvation? (eventually specific one will get in)


### Correctness Proofs

- Prove incorrect
  - Come up with one scenario that does not work
    - Two processes execute in sync? **often non-trivial**
    - Some other unlikely scenario?
- Prove correct
  - Heuristics: "I did not come up with any proofs (counterexample) for incorrectness and I am smart"
    - I can not prove incorrectness **"easy", unreliable**
    - It must be correct...
  - State diagrams **difficult, reliable**
    - Describe algorithm with states:
      - { relevant control pointer (cp) values, relevant local/global variable values }
    - Analyze state diagrams to prove correctness

### State Diagram for Alg. 3.2

Algorithm 3.2

- State {p<sub>i</sub>, q<sub>i</sub>, turn}
  - Control pointer p<sub>i</sub>
  - Control pointer q<sub>i</sub>
  - Global variable turn
  - 1<sup>st</sup> four states →
- Mutex ok
  - State {p<sub>3</sub>, q<sub>3</sub>, turn} **not accessible** in state diagram?
- No deadlock?
  - When many processes try concurrently, one will succeed
- No starvation?
  - Whenever any (one) process tries, it will eventually succeed




How to prove it?

### State Diagram for Algorithm 3.2

Algorithm 3.2

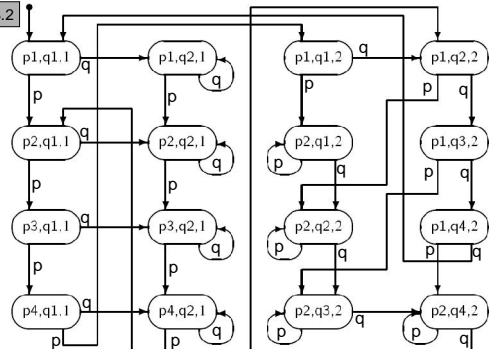
- Create complete diagram with all accessible states
- No states
  - {p<sub>3</sub>, q<sub>3</sub>, 1}
  - {p<sub>3</sub>, p<sub>3</sub>, 2}
- I.e., mutex secured **proof!**
- Problem:
  - Too many states?
  - Difficult to create
  - Difficult to analyze



(Fig. 3.1)

### Alternate Layout for Full State Diagram


Alg. 3.2



### Correctness (3)

Algorithm 3.2

- Mutex?
  - Ok, no state {p<sub>3</sub>, q<sub>3</sub>, ??}
- No deadlock?
  - many try, one can always get in? (into a state with p<sub>3</sub> or q<sub>3</sub>)
  - {p<sub>2</sub>, q<sub>1</sub>, 1}: P can get in
  - {p<sub>2</sub>, q<sub>2</sub>, 1}: P can get in
  - {p<sub>2</sub>, q<sub>1</sub> tai q<sub>2</sub>, 2}:
    - Q can get in
    - P can get in eventually
  - {p<sub>i</sub>, q<sub>2</sub>, ?} similarly
- No starvation?
  - One tries, it will eventually get in
  - {p<sub>2</sub>, q<sub>1</sub>, 2}
    - Q dies (ok to die in q<sub>1</sub>).
    - P will starve! **Not good!**



(Fig. 3.1)

### Reduced Algorithm for Easier Analysis

**Algorithm 3.2: First attempt**  
integer turn ← 1

p	q
loop forever	loop forever
p1: non-critical section	q1: non-critical section
p2: await turn = 1	q2: await turn = 2
p3: critical section	q3: critical section
p4: turn ← 2	q4: turn ← 1

- Reduce algorithm to reduce number of states of state diagrams: leave irrelevant code out
  - Nothing relevant (for mutex) left out?

**Algorithm 3.5: First attempt (abbreviated)**  
integer turn ← 1

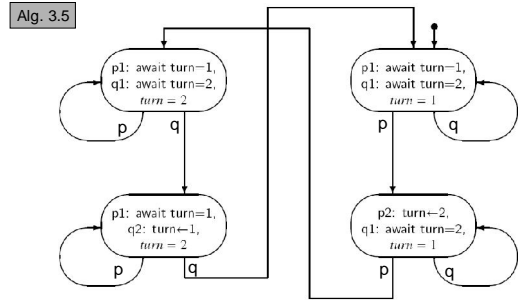
p	q
loop forever	loop forever
p1: await turn = 1	q1: await turn = 2
p2: turn ← 2	q2: turn ← 1

3.11.2008

Copyright Teemu Kerola 2008

13

### State Diagram for Reduced Algorithm



- Much fewer states!

(Fig. 3.2)

3.11.2008

Copyright Teemu Kerola 2008

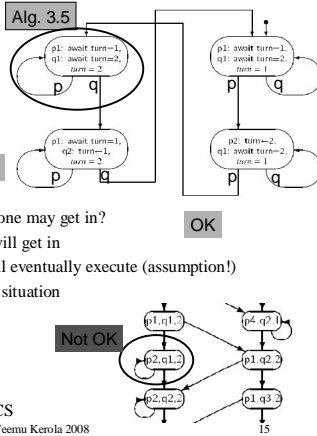
14

### Correctness of Reduced Algorithm (2)

- Mutex?
  - No state {p2, q2, turn}
- No deadlock: Some are trying, one may get in?
  - Top left (p & q trying): q will get in
  - Bottom left (p trying): q will eventually execute (assumption!)

should be OK to die in NCS, but not OK to die in protocol

- Top & bottom right: mirror situation
- starvation?
  - Tricky, reduced too much!
    - NCS combined with await
  - Look at original diagram
    - Problem if Q dies in NCS



3.11.2008

Copyright Teemu Kerola 2008

15

### Critical Section Solution #2

**Algorithm 3.6: Second attempt**  
boolean wantp ← false, wantq ← false

p	q
loop forever	loop forever
p1: non-critical section	q1: non-critical section
p2: await wantq = false	q2: await wantp = false
p3: wantp ← true	q3: wantq ← true
p4: critical section	q4: critical section
p5: wantp ← false	q5: wantq ← false

- Each have their own global variable *wantp* and *wantq*
  - True when process is in critical section
- Process dies in NCS?
  - Starvation problem ok, because it's *want*-variable is false
- Mutex? Deadlock?

3.11.2008

Copyright Teemu Kerola 2008

16


### Attempt #2 Reduced

**Algorithm 3.7: Second attempt (abbreviated)**  
boolean wantp ← false, wantq ← false

p	q
loop forever	loop forever
p1: await wantq = false	q1: await wantp = false
p2: wantp ← true	q2: wantq ← true
p3: wantp ← false	q3: wantq ← false

- No mutex! {p3, q3, ?} reachable
  - Problem: p2 should be part of critical section (but is not!)

proto-col



3.11.2008

Copyright Teemu Kerola 2008

17

### Critical Section Solution #3

**Algorithm 3.8: Third attempt**  
boolean wantp ← false, wantq ← false

p	q
loop forever	loop forever
p1: non-critical section	q1: non-critical section
p2: wantp ← true	q2: wantq ← true
p3: await wantq = false	q3: await wantp = false
p4: critical section	q4: critical section
p5: wantp ← false	q5: wantq ← false

- Avoid previous problem, **mutex ok**
- Deadlock possible: {p3, q3, wantp=true, wantq=true}
- Problem: cyclic wait possible, both insist their turn next
  - No preemption

3.11.2008

Copyright Teemu Kerola 2008

18

**Algorithm 3.9: Fourth attempt**

boolean wantp ← false, wantq ← false

P	Q
loop forever p1: non-critical section p2: wantp ← true p3: while wantq p4:     wantp ← false p5:     wantp ← true p6: critical section p7: wantp ← false	loop forever q1: non-critical section q2: wantq ← true q3: while wantp q4:     wantq ← false q5:     wantq ← true q6: critical section q7: wantq ← false

- Avoid deadlock by giving away your turn if needed
- Mutex ok: P in p6 only if !wantq (Q is not in q6)
- Deadlock (livelock) possible:
  - {p3, q3, ...} → {p4, q4, ...} → {p5, q5, ...}
  - Unlikely but possible!
  - **Livelock**: both executing all the time, not waiting suspended
    - Neither one advances

eloukko

3.11.2008 Copyright Teemu Kerola 2008 19

**Algorithm 3.10: Dekker's algorithm**

boolean wantp ← false, wantq ← false  
integer turn ← 1

P	Q
loop forever p1: non-critical section p2: wantp ← true p3: while wantq p4:     if turn = 2 p5:         wantp ← false p6:         await turn = 1 p7:     wantp ← true p8: critical section p9: turn ← 2 p10: wantp ← false	loop forever q1: non-critical section q2: wantq ← true q3: while wantp q4:     if turn = 1 q5:         wantq ← false q6:         await turn = 2 q7:     wantq ← true q8: critical section q9: turn ← 1 q10: wantq ← false

- Combine 1st and 4th attempt
- 3 global (mutex ctr) variables: shared *turn*, semi-private *want*'s
  - only one process writes to *wantp* or *wantq* (= semi-private)
- *turn* gives you the right to insist, i.e., priority
  - Used only when both want CS at the same time

3.11.2008 Copyright Teemu Kerola 2008 20

**Algorithm 3.10: Dekker's algorithm**

boolean wantp ← false, wantq ← false  
integer turn ← 1

P	Q
loop forever p1: non-critical section p2: wantp ← true p3: while wantq p4:     if turn = 2 p5:         wantp ← false p6:         await turn = 1 p7:     wantp ← true p8: critical section p9: turn ← 2 p10: wantp ← false	loop forever q1: non-critical section q2: wantq ← true q3: while wantp q4:     if turn = 1 q5:         wantq ← false q6:         await turn = 2 q7:     wantq ← true q8: critical section q9: turn ← 1 q10: wantq ← false

**Proof**

- Mutex ok: P in p8 only if !wantq (Q can not be in q8)
- No deadlock, because P or Q can continue to CS from {p3, q3, ...}
- No starvation, because
  - If in {p6, ...}, then eventually {p6, q9, ...} and {..., q10, ...}
  - Next time {p3, ...} or {p4, ...} will lead to {p8, ...}

3.11.2008 Copyright Teemu Kerola 2008 21

**Algorithm 3.10: Dekker's algorithm**

boolean wantp ← false, wantq ← false  
integer turn ← 1

P	Q
loop forever p1: non-critical section p2: wantp ← true p3: while wantq p4:     if turn = 2 p5:         wantp ← false p6:         await turn = 1 p7:     wantp ← true p8: critical section p9: turn ← 2 p10: wantp ← false	loop forever q1: non-critical section q2: wantq ← true q3: while wantp q4:     if turn = 1 q5:         wantq ← false q6:         await turn = 2 q7:     wantq ← true q8: critical section q9: turn ← 1 q10: wantq ← false

- mutex with no HW-support needed, need only shared memory
- Bad: complex, many instructions
  - Must execute each instruction at a time, in this order
    - Will not work, if compiler optimizes code too much!
  - In simple systems, can do better with HW support
    - Special machine instructions to help with this problem

3.11.2008 Copyright Teemu Kerola 2008 22

## Mutex with HW Support

- Specific machine instructions for this purpose
  - Suitable for many situations
  - Not suitable for all situations
- Interrupt disable/enable instructions
- Test-and-set instructions
  - Other similar instructions
- Specific memory areas
  - Reserved for concurrency control solutions
  - Lock variables (for test-and-set) in their own cache?
    - Different cache protocol for lock variables?
    - Busy-wait without memory bus use?

Disable  
 -- Critical Section --  
 Enable

Lock (L)  
 -- Critical Section --  
 Unlock (L)

3.11.2008 Copyright Teemu Kerola 2008 23

## Disable Interrupts

- Environment
  - All (competing) processes on same processor
  - Not for multiprocessor systems
    - Disabling interrupts does it only for the processor executing that instruction
- Disable/enable interrupts
  - Prevent process switching during critical sections
    - Good for only very short time
    - Prevents also (other) operating system work while in CS

Disable  
 -- CS --  
 Enable

Disable  
 -- CS --  
 Enable

3.11.2008 Copyright Teemu Kerola 2008 24

### Test-and-set locking variables

- Environment
  - All processes with shared memory
  - Should have multiple processors
  - Not very good for uniprocessor systems (or synchronizing processes running on the same processor)
    - Wait (**busy-wait**) while holding the processor!
- Test-and-set *machine instruction*
  - Indivisibly read old value and write new value (complex mem-op)

Test-and-set (common, local)  
 local ← common ; read state  
 common ← 1 ; mark reserved

shared

Test-and-set (shLock, locked);  
 while (locked)  
 Test-and-set (shLock, locked);  
 -- CS --  
 shLock = 0;

local

Test-and-set (shLock, locked);  
 while (locked)  
 Test-and-set (shLock, locked);  
 -- CS --  
 shLock = 0;

3.11.2008 Copyright Teemu Kerola 2008 25

### Other Machine Instructions for Synchronization Problem Busy-Wait Solutions

- Test-and-set
  - Test-and-set (common, local)  
 local ← common ; read state  
 common ← 1 ; mark reserved
- Exchange
  - Exchange (common, local)  
 local ↔ common ; swap values
- Fetch-and-add
  - Fetch-and-add (common, local, x)  
 local ← common ; read state  
 common ← common+x ; add x
- Compare-and-swap
  - int Compare-and-swap (common, old, new)  
 return\_val ← common  
 if (common == old)  
 common ← new

Use all in busy-wait loops

"read-modify-write" memory bus transaction (local in HW register)

"read-after-write" memory bus transaction may also be used

3.11.2008 Copyright Teemu Kerola 2008 26