

## Lesson 5

## Deadlocks

Ch 6 [Stall 05]

Problem  
Dining Philosophers  
Deadlock occurrence  
Deadlock detection  
Deadlock prevention  
Deadlock avoidance

31.1.2011

Copyright Teemu Kerola 2011

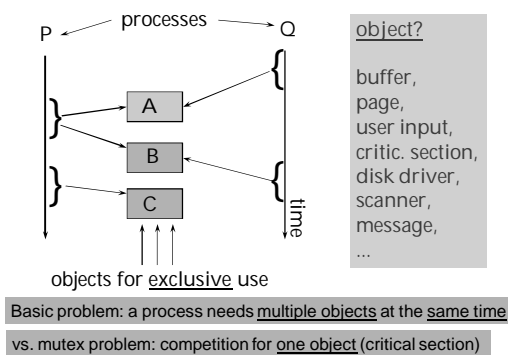
1

## Motivational Example

- New possible laptop for CS dept use
  - Lenovo 400, dual-core, Intel Centrino 2 technology
  - Ubuntu Linux 8.10
- Wakeup from suspend/hibernation, freezes often  
<http://ubuntuforums.org/showthread.php?t=959712>
- Read, study, experiment – some 15 hours?
  - No network?, at home/work?, various units?, ..., ???
  - Problem with Gnome desktop, not with KDE, ..., ???
- Could two processors cause it?
  - Shut down one processor during hibernation/wakeup
  - Wakeup works fine now
- Same problem with many new laptops running Linux
  - All new laptops with Intel Centrino 2 with same Linux driver?
- Concurrency problem in display driver startup?
  - Bug not found yet, use 1-cpu work-around

<http://git.kernel.org/?p=linux/kernel/git/torvalds/linux-2.6.git;a=commitdiff;h=70740d6c3030b339b4ad17fd58ee135dfc13913>  
31.1.2011 Copyright Teemu Kerola 2011 (search "1915\_enable\_vblank" ... ) 2

## Deadlock: Background

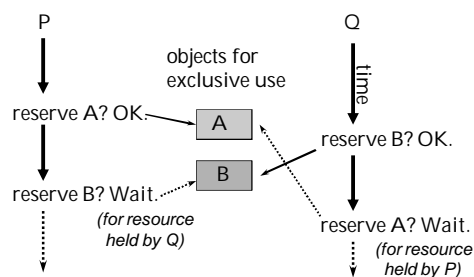


31.1.2011

Copyright Teemu Kerola 2011

3

## Deadlock: an Example (10)

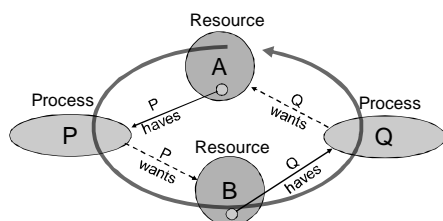


31.1.2011

Copyright Teemu Kerola 2011

4

## Resource Reservation Graph

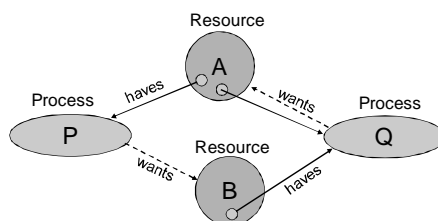
Deadlock cycle in resource reservation graph

31.1.2011

Copyright Teemu Kerola 2011

5

## Resource Reservation Graph

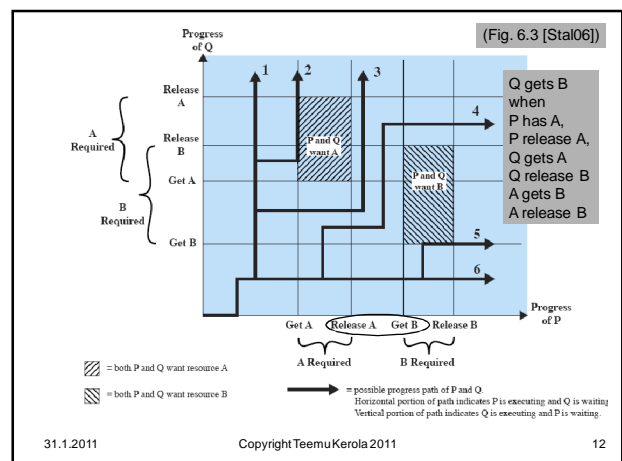
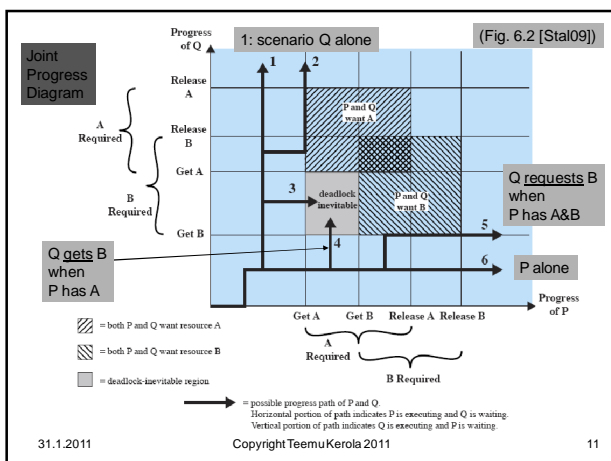
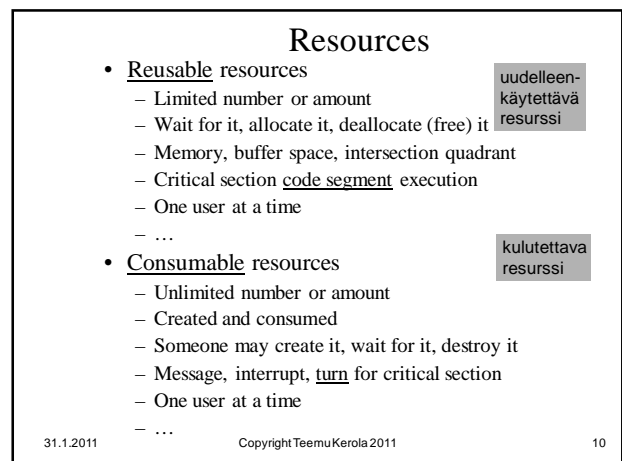
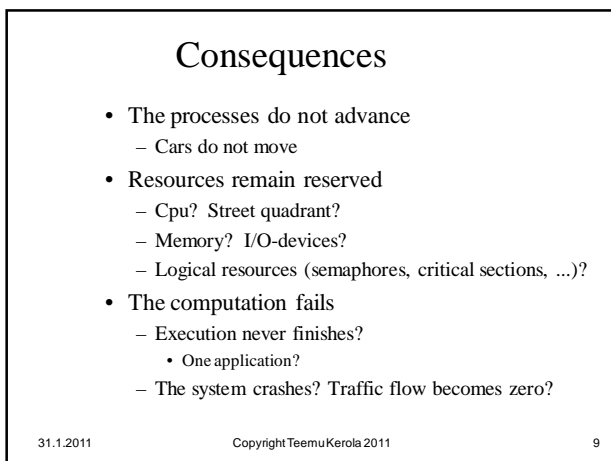
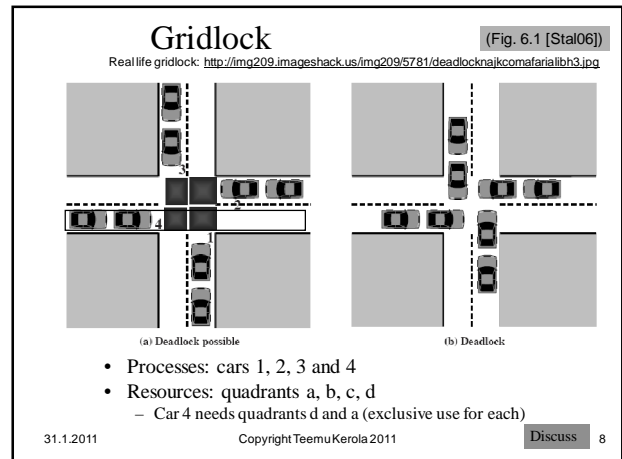
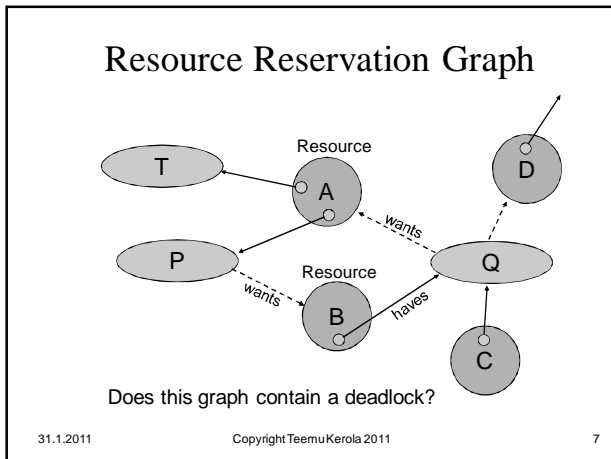


Does this graph contain a deadlock?

31.1.2011

Copyright Teemu Kerola 2011

6



## Definitions

- **Deadlock**
  - Eternal wait in blocked state
  - Does not block processor (unless one resource is processor)
- **Livelock**
  - Two or more processes continuously change their state (execute/wait) as response to the other process(es), but never advance to real work
  - E.g., ping-pong "you first – no, you first – ..."
    - two processes alternate offering the turn to each other - no useful work is started
  - Consumes processor time
- **Starvation**
  - the process will never get its turn
  - E.g., in ready-to-run queue, but never scheduled

31.1.2011

Copyright Teemu Kerola 2011

13

## Deadlock Problems

- How to know if deadlock exists?
  - How to locate deadlocked processes?
- How to prevent deadlocks?
- How to know if deadlock might occur?
- How to break deadlocks?
  - Without too much damage?
  - Automatically?
- How to prove that your solution is free of deadlocks?

31.1.2011

Copyright Teemu Kerola 2011

14

## Good Deadlock Solution

- Prevents deadlocks in advance, or detects them, breaks them, and fixes the system
- Small overhead
- Smallest possible waiting times
- Does not slow down computations when no danger exists
- Does not block unnecessarily any process when the resource wanted is available

31.1.2011

Copyright Teemu Kerola 2011

15

### Conditions for Possible Deadlock

- Three policy conditions
  - S1. Resource mutual exclusion yksi käyttäjä
    - one user of any resource at a time (not just code)
  - S2. Hold and wait pitää ja odota
    - a process may hold allocated resources while waiting for others
  - S3. No preemption ei keskeytettävissä
    - resource can not be forcibly removed from a process holding it
- A dynamic (execution time) condition takes place
  - D1. Circular wait: a closed chain of processes exists, each process holds at least one resource needed by the next process in chain

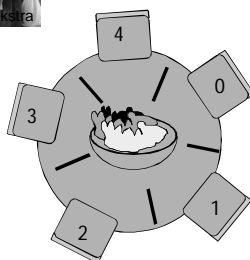
<http://portal.acm.org/citation.cfm?id=356588&coll=GUIDE&dl=GUIDE&CFID=4442763&CFTOKEN=75849639&ret=1#Fulltext>

31.1.2011

Copyright Teemu Kerola 2011

16

## Dining Philosophers (Dijkstra)



See philosopher art in web

<http://images.google.fi/images?q=dinig%20philosophers&ie=UTF-8&oe=utf-8&rs=org.mozilla.en-US:official&client=firefox-a&um=1&sa=N&tab=wi>  
31 1 2011 Copyright Teemu Kerkola 2011

<http://images.g>  
31.1.2011

Copyright Teemu Kerola 2011

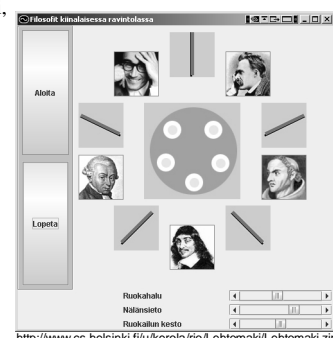
17

Philosopher:  
think  
take two forks ...  
... one from each side  
eat rice until satisfied  
return the forks

Problem:  
how to reserve the forks  
without causing  
- deadlock  
- starvation  
and everybody may be  
present

## Dining Philosophers in Java

- Tapio Lehtomäki, MikroBitti
- Load program from course schedule page
- Modify paths in script philosophers.bat and run it
- Modify program for homework?
  - Next year?



[http://www.cs.helsinki.fi/u/kerola/rip/l\\_ehtomaki/l\\_ehtomaki.zip](http://www.cs.helsinki.fi/u/kerola/rip/l_ehtomaki/l_ehtomaki.zip)

31.1.2011

Copyright Teemu Kerola 2011

18

```

/* program diningphilosophers */
semaphore fork [5] = {1}; /* mutex, one at a time */
int i;
void philosopher (int i)
{
    while (true)
    {
        think();
        wait (fork[i]); /* left fork */
        wait (fork [(i+1) mod 5]); /* right fork */
        eat();
        signal(fork [(i+1) mod 5]);
        signal(fork[i]);
    }
}
void main()
{
    parbegin (philosopher (0), philosopher (1), philosopher (2),
              philosopher (3), philosopher (4));
}

```

(Fig. 6.12 [Stal09])

**Trivial Solution #1**

- Possible deadlock scenario – not good
  - All 5 grab left fork “at the same time”

Discuss

31.1.2011 Copyright Teemu Kerola 2011 19

## Resource Allocation (Dijkstra's)

- Processes  $P_i \in P_1..P_n$
- Resources (or objects)  $R_j \in R_1..R_m$
- Number of resources of type  $R_j$ 
  - total amount of resources  $R = (r_1, \dots, r_m)$
  - currently free resources  $V = (v_1, \dots, v_m)$
- Allocated resources (allocation matrix)
  - $A = [a_{ij}]$ , “process  $P_i$  has  $a_{ij}$  units of resource  $R_j$ ”
- Outstanding requests (request matrix)
  - $Q = [q_{ij}]$ , “process  $P_i$  requests  $q_{ij}$  units of resource  $R_j$ ”

31.1.2011

Copyright Teemu Kerola 2011

20

How many R4 resources exists?

	R1	R2	R3	R4	R5
P1	1	0	1	1	0
P2	1	1	0	0	0
P3	0	0	0	1	0
P4	0	0	0	0	0

Resource vector  $R$

	R1	R2	R3	R4	R5
$R$	0	0	0	0	1

Available vector  $V$

	R1	R2	R3	R4	R5
P1	0	1	0	0	1
P2	0	0	0	0	1
P3	0	0	0	0	1
P4	1	0	1	0	1

Request matrix  $Q$

Which resources are now free?

Who has now R4?

P2 has now R1 and R2, P2 wants now R3 and R5

Is there now a deadlock or not?

(Fig. 6.10 [Stal09])

Discuss

31.1.2011 Copyright Teemu Kerola 2011 21

## DDA- Deadlock Detection Algorithm (Dijkstra)

- Find a (any) process that could terminate
  - All of its current resource requests can be satisfied
- Assume now that
  - This process terminates, and
  - It releases all of its resources
- Repeat 1&2 until can not find any more such processes
- If any processes still exist, they are deadlocked
  - They all each need something
  - The process holding that something is waiting for something else
    - That process can not advance and release it



31.1.2011

Copyright Teemu Kerola 2011

22

## Deadlock Detection Algorithm (DDA)

- DL1. [Remove the processes with no resources] Mark all processes with null rows in  $A$ .
- DL2. [Initialize counters for available objects] Initialize a working vector  $W = V$
- DL3. [Search for a process  $P_i$  which could get all resources it requires] Search for an unmarked row  $i$  such that  $q_{ij} \leq w_j$   $j = 1..n$ . If none is found terminate the algorithm.
- DL4. [Increase  $W$  with the resources of the chosen process] Set  $W = W + A_i$ , i.e.  $w_j = w_j + a_{ij}$  when  $j = 1..n$ . Mark process  $P_i$  and return to step DL3.

When the algorithm terminates, unmarked processes correspond to deadlocked processes. Why?

31.1.2011

Copyright Teemu Kerola 2011

23

## Example: Initial state

	allocation matrix $A$	request matrix $Q$	
row 1:	1 0 1 1 0	0 1 0 0 1	
2:	1 1 0 0 0	0 0 1 0 1	E.g., "process 2 has resources 1 & 2, and it wants resources 3 & 5"
3:	0 0 0 1 0	0 0 0 0 1	
4:	0 0 0 0 0	1 0 1 0 1	
	all resources $R$	2 1 1 2 1	Who holds resource 4?
	free resources $V$	0 0 0 0 1	Which resources are free?

(Fig. 6.10 [Stal09])

Deadlock or not?

What now?

31.1.2011

Copyright Teemu Kerola 2011

24

Example: Deadlock Detection (6)

10110

11000

00010

00000

DL4: mark

DL1: mark

all resources

free resources

may become free

DL4: new W

01001

00101

00001

10101

DL3: no request can be satisfied:  
 $\exists i \forall j: q_{ij} \leq w_{ij} \rightarrow \text{Deadlock}$

DL3: this request can be satisfied:  
 $q_{3j} \leq w_{ij} \forall j$

DL2: copy

DL4: new W

R: 21121

V: 00001

W: 00001

00011

Discuss

Example: Deadlock Detection (phases)

10110

11000

00010

00000

all resources

free resources

may become free

01001

00101

00001

10101

R: 21121

V: 00001

W:

Example: Deadlock Detection (phases)

10110

11000

00010

00000

DL1: mark

all resources

free resources

may become free

01001

00101

00001

10101

R: 21121

V: 00001

W:

Example: Deadlock Detection (phases)

10110

11000

00010

00000

DL1: mark

all resources

free resources

may become free

01001

00101

00001

10101

R: 21121

V: 00001

W: 00001

DL2: copy

Example: Deadlock Detection (phases)

10110

11000

00010

00000

DL1: mark

all resources

free resources

may become free

01001

00101

00001

10101

DL3: this request can be satisfied:  
 $q_{3j} \leq w_{ij} \forall j$

DL2: copy

R: 21121

V: 00001

W: 00001

Example: Deadlock Detection (phases)

10110

11000

00010

00000

DL1: mark

all resources

free resources

may become free

01001

00101

00001

10101

DL3: this request can be satisfied:  
 $q_{3j} \leq w_{ij} \forall j$

DL2: copy

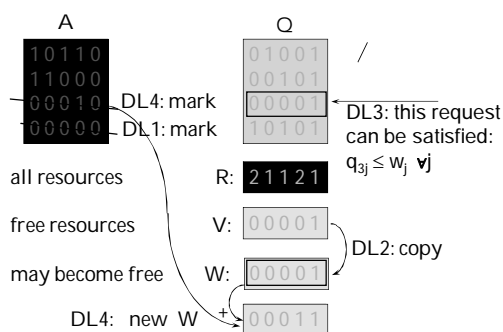
R: 21121

V: 00001

W: 00001

DL4: new W

## Example: Deadlock Detection (phases)

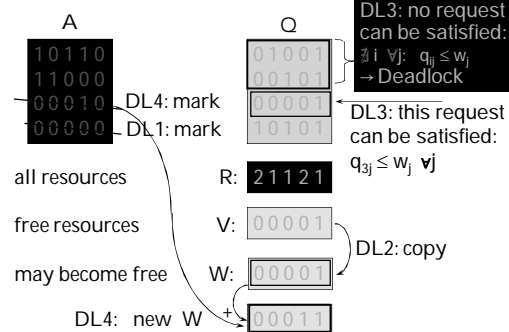


31.1.2011

Copyright Teemu Kerola 2011

31

## Example: Deadlock Detection (phases)



31.1.2011

Copyright Teemu Kerola 2011

32

## Example: Breaking Deadlocks

- Processes P1 and P2 are in deadlock
  - What next?
- Abort P1 and P2?
  - Most common solution
- Rollback P1 and P2 to previous safe state, and try again
  - Rollback states must exist
  - May deadlock again (or may not!)
- Abort /Rollback P1 because it is less important
  - Must have some basis for selection
  - Who makes the decision? Automatic?
- Preempt R3 from P1
  - Must be able to preempt (easy if R3 is CPU?)
  - Must know what to preempt from whom
  - How many resources need preemption?

31.1.2011

Copyright Teemu Kerola 2011

33

## Deadlock Prevention

- How to prevent deadlock occurrence in advance?
- Deadlock possible only when all 4 conditions are met:
  - S1. Mutual exclusion: Yksi käyttäjä resurssilla
  - S2. Hold and wait: pidä ja odota
  - S3. No preemption: ei saa ottaa pois kesken kaiken
  - D1. Circular wait: kehäodotus
- Solution: disallow any one of the conditions
  - S1, S2, S3, or D1?
  - Which is possible to disallow?
  - Which is easiest to disallow?

31.1.2011

Copyright Teemu Kerola 2011

34

## Disallow S1 (mutual exclusion)

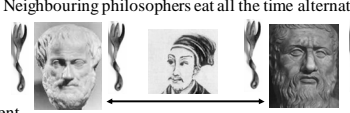
- Can not do always
  - There are reasons for mutual exclusion!
    - Can not split philosophers fork into 2 resources
- Can do sometimes
  - Too high granularity blocks too much
    - Resource room in trivial solution #2
  - Finer granularity allows parallelism
    - Smaller areas, parallel usage, more locks
    - More administration to manage more locks
    - Too fine granularity may cause too much administration work
  - Normal design approach in data bases, for example
- Get more resources, avoid mutex competition?
  - Buy another fork for each philosopher?

31.1.2011

Copyright Teemu Kerola 2011

35

## Disallow S2 (hold and wait)

- Request all needed resources at one time
  - Wait until all can be granted simultaneously
    - Can lead to starvation
      - Reserve both forks at once (simultaneous wait!)
      - Neighbouring philosophers eat all the time alternating
- 
- Inefficient
    - long wait for resources (to be used much later?)
    - worst case reservation (long wait period for resources which are possibly needed - who knows?)
  - Difficult/impossible to implement?
    - advance knowledge: resources of all possible execution paths of all related modules ...

31.1.2011

Copyright Teemu Kerola 2011

36

### Disallow S3 (no preemption)

- Allow preemption in crisis
- Release of resources => fallback to some earlier state
  - Initial reservation of these resources
  - Fall back to specific checkpoint
  - Checkpoint must have been saved earlier
  - Must know when to fall back!
- OK, if the system has been designed for this
  - Practical, if saving the state is cheap and the chance of deadlock is to be considered
  - Standard procedure for transaction processing

```
wait (fork[i]);
if "all forks taken" then
  "remove fork" from philosopher [i+1]
wait (fork[i+1])
```

- What will philosopher  $i+1$  do now? Think? Eat? Die?

31.1.2011

Copyright Teemu Kerola 2011

37

### Disallow D1 (circular wait)

- Linear ordering of resources
  - Make reservations in this order only – no loops!
- Pessimistic approach – prevent “loops” in advance
  - Advance knowledge of resource requirements needed
  - Reserve all at once in given order
  - Prepare for “worst case” behavior

Forks in global ascending order  
philosophers 0, 1, 2, 3:

```
wait (fork[i]);
wait (fork[i+1]);
```

last philosopher 4:

```
wait (fork[0]);
wait (fork[4]);
```

- Optimistic approach – worry only at the last moment
  - Reservation dynamically as needed (but in order)
  - Reservation conflict => restart from some earlier stage
    - Must have earlier state saved somewhere

31.1.2011

Copyright Teemu Kerola 2011

38

```
/* program diningphilosophers */
semaphore fork[5] = {1};
semaphore room = {4}; /* only 4 at a time, 5th waits */
int i;
void philosopher (int I)
{
  while (true)
  {
    think();
    wait (room);
    wait (fork[i]);
    wait (fork [(i+1) mod 5]);
    eat();
    signal (fork [(i+1) mod 5]);
    signal (fork[i]);
    signal (room);
  }
}
void main()
{
  parbegin (philosopher (0), philosopher (1), philosopher (2),
            philosopher (3), philosopher (4));
}
```

(Fig. 6.13 [Stal09])

- No deadlock, no starvation - circular wait not possible

31.1.2011

Copyright Teemu Kerola 2011

39

### Deadlock Detection and Recovery

- Let the system run until deadlock problem occurs
  - “Detect deadlock existence”
  - “Locate deadlock and fix the system”
- Detection is not trivial:
  - Blocked group of processes is deadlocked? or
  - Blocked group is just waiting for an external event?
- Recovery
  - Detection is first needed
  - Fallback to a previous state (does it exist?)
  - Killing one or more members of the deadlocked group
    - Must be able to do it without overall system damage
- Needed: information about resource allocation
  - In a form suitable for deadlock detection!

31.1.2011

Copyright Teemu Kerola 2011

40

### Banker's Algorithm: Deadlock Avoidance with DDA

- Use Dijkstra's algorithm to avoid deadlocks in advance?
- Banker's Algorithm Pankkiirin algoritmi
  - Originally for one resource (money)
  - Why “Banker's”?
  - “Ensure that a bank never allocates its available cash so that it can no longer satisfy the needs of all its customers”

31.1.2011

Copyright Teemu Kerola 2011

41

### Banker's Algorithm (Dijkstra, 1977?)



- Keep state information on resources allocated to each process
- Keep state information on number of resources each process might still allocate
- For each resource allocation, first find an ordering which allows processes to terminate, if that allocation is made
  - Assume that allocation is made and then use DDA to find out if the system remains in a safe state even in the worst case
  - If deadlock is possible, reject resource request
  - If deadlock is not possible, grant resource request

Discuss

31.1.2011

Copyright Teemu Kerola 2011

42

## Deadlock Avoidance with Banker's Algorithm

Matrices as before, and some more

- For each process: the maximum needs of resources
  - $C = [c_{ij}]$ , "Pi may request  $c_{ij}$  units of  $R_j$ "
- The current hypothesis of resources in use
  - $A' = [a'_{ij}]$ , "if this allocation is made, Pi would have  $a'_{ij}$  units of  $R_j$ "
- The current hypothesis of future maximum demands
  - $Q' = [q'_{ij}]$ , "Pi could still request  $q'_{ij}$  units of  $R_j$ "
  - $Q' = C - A'$
- Apply DDA to  $A'$  and  $Q'$ 
  - If no deadlock possible, grant resource request

31.1.2011

Copyright Teemu Kerola 2011

43

## Banker's Algorithm Example

Allocation A						Requests Q						Max allocation C					
P1	R1	R2	R3	R4	R5	P1	R1	R2	R3	R4	R5	P1	R1	R2	R3	R4	R5
0	1	0	0	0	0	1	0	0	0	0	0	2	1	0	1	0	0
1	1	0	0	0	0	0	0	0	0	1	0	1	1	0	0	0	1
0	0	1	0	1	0	0	0	0	1	0	0	1	0	1	1	1	1
0	0	1	1	0	0	0	0	0	0	1	0	0	2	1	1	1	1

Resources R						Available V					
R1	R2	R3	R4	R5		R1	R2	R3	R4	R5	
2	3	2	1	2		1	1	0	0	1	

(Fig. 16.11, Bacon, Concurrent Systems, 1993)

P1 requests R1. Is request granted?  
Could system deadlock, if R1 is granted?

31.1.2011

Copyright Teemu Kerola 2011

44

## Banker's Algorithm Example (7)

If P1 request for R1 approved, can deadlock occur?

Possible allocation A'						Possible requests Q'						Max allocation C					
P1	R1	R2	R3	R4	R5	P1	R1	R2	R3	R4	R5	P1	R1	R2	R3	R4	R5
1	1	0	0	0	0	1	0	0	1	0	0	2	1	0	1	0	0
1	1	0	0	0	0	0	0	0	0	1	0	1	1	0	0	0	1
0	0	1	0	1	0	1	0	0	1	0	0	1	0	1	1	1	1
0	0	1	1	0	0	0	2	0	0	1	0	0	2	1	1	1	1

Resources R						Available V					
R1	R2	R3	R4	R5		R1	R2	R3	R4	R5	
0	1	0	0	1		1	1	0	0	1	

Resources R						Possibly available V'					
R1	R2	R3	R4	R5		R1	R2	R3	R4	R5	
0	1	0	0	1		0	1	0	0	1	

31.1.2011

Copyright Teemu Kerola 2011

45

## Banker's Algorithm Example (7)

If P1 request for R1 approved, can deadlock occur?

Possible allocation A'						Possible requests Q'						Max allocation C					
P1	R1	R2	R3	R4	R5	P1	R1	R2	R3	R4	R5	P1	R1	R2	R3	R4	R5
1	1	0	0	0	0	1	0	0	1	0	0	2	1	0	1	0	0
1	1	0	0	0	0	0	0	0	0	1	0	1	1	0	0	0	1
0	0	1	0	1	0	1	0	0	1	0	0	1	0	1	1	1	1
0	0	1	1	0	0	0	2	0	0	1	0	0	2	1	1	1	1

Resources R						Available V					
R1	R2	R3	R4	R5		R1	R2	R3	R4	R5	
0	1	0	0	1		1	1	0	0	1	

Resources R						Possibly available V'					
R1	R2	R3	R4	R5		R1	R2	R3	R4	R5	
0	1	0	0	1		0	1	0	0	1	

31.1.2011

Copyright Teemu Kerola 2011

46

## Banker's Algorithm Example (7)

If P1 request for R1 approved, can deadlock occur?

Possible allocation A'						Possible requests Q'						Max allocation C					
P1	R1	R2	R3	R4	R5	P1	R1	R2	R3	R4	R5	P1	R1	R2	R3	R4	R5
1	1	0	0	0	0	1	0	0	1	0	0	2	1	0	1	0	0
1	1	0	0	0	0	0	0	0	0	1	0	1	1	0	0	0	1
0	0	1	0	1	0	1	0	0	1	0	0	1	0	1	1	1	1
0	0	1	1	0	0	0	2	0	0	1	0	0	2	1	1	1	1

Resources R						Available V					
R1	R2	R3	R4	R5		R1	R2	R3	R4	R5	
0	1	0	0	1		1	1	0	0	1	

Resources R						Possibly available V'					
R1	R2	R3	R4	R5		R1	R2	R3	R4	R5	
0	1	0	0	1		0	1	0	0	1	

31.1.2011

Copyright Teemu Kerola 2011

47

## Banker's Algorithm Example (7)

If P1 request for R1 approved, can deadlock occur?

Possible allocation A'						Possible requests Q'						Max allocation C					
P1	R1	R2	R3	R4	R5	P1	R1	R2	R3	R4	R5	P1	R1	R2	R3	R4	R5
1	1	0	0	0	0	1	0	0	1	0	0	2	1	0	1	0	0
1	1	0	0	0	0	0	0	0	0	1	0	1	1	0	0	0	1
0	0	1	0	1	0	1	0	0	1	0	0	1	0	1	1	1	1
0	0	1	1	0	0	0	2	0	0	1	0	0	2	1	1	1	1

Resources R						Available V					
R1	R2	R3	R4	R5		R1	R2	R3	R4	R5	
0	1	0	0	1		1	1	0	0	1	

Resources R						Possibly available V'					
R1	R2	R3	R4	R5		R1	R2	R3	R4	R5	
0	1	0	0	1		0	1	0	0	1	

31.1.2011

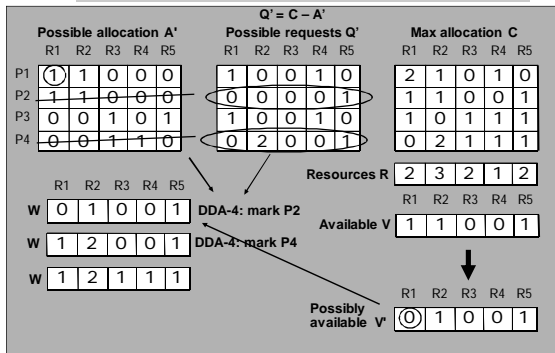
Copyright Teemu Kerola 2011

48



## Banker's Algorithm Example (7)

If P1 request for R1 approved, can deadlock occur?



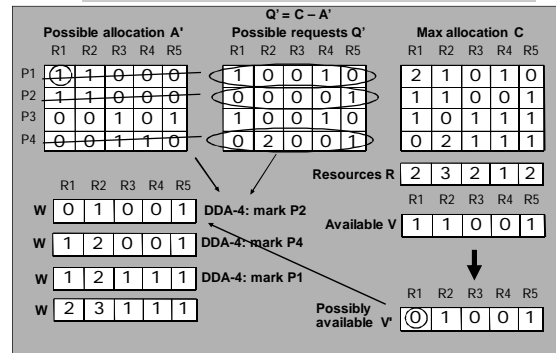
31.1.2011

Copyright Teemu Kerola 2011

49

## Banker's Algorithm Example (7)

If P1 request for R1 approved, can deadlock occur?



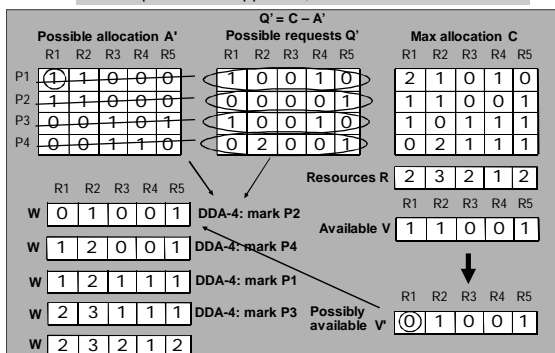
31.1.2011

Copyright Teemu Kerola 2011

50

## Banker's Algorithm Example (7)

If P1 request for R1 approved, can deadlock occur?



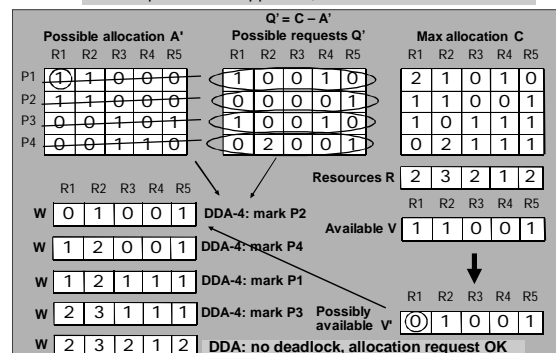
31.1.2011

Copyright Teemu Kerola 2011

51

## Banker's Algorithm Example (7)

If P1 request for R1 approved, can deadlock occur?



31.1.2011

Copyright Teemu Kerola 2011

52

## Deadlock Avoidance Problems

- Each allocation: a considerable overhead
  - Run Banker's algorithm for 20 processes and 100 resources?
- Knowledge of maximum needs
  - In advance?
    - An educated guess? Worst case?
  - Dynamically?
    - Even more overhead
- A safe allocation does not always exist
  - An unsafe state does not always lead to deadlock
  - You may want to take a risk!

Another Banker's Algorithm example: B. Gray, Univ. of Idaho  
<http://www.if.uidaho.edu/~bgray/classes/cs341/doc/banker.html>

31.1.2011

Copyright Teemu Kerola 2011

53

## Summary

- Difficult real problem
- Can detect deadlocks Dijkstra's DDA
  - Need specific data on resource usage
- Difficult to break deadlocks
  - How will killing processes affect the system?
- Can prevent deadlocks e.g., Bankers
  - Prevent any one of those four conditions
    - E.g., reserve resources always in given order
  - Can analyze system at resource reservation time to see whether deadlock might result
    - Complex and expensive

31.1.2011

Copyright Teemu Kerola 2011

54