

# Deadlocks

*Ch 6 [Stall 05]*

Problem

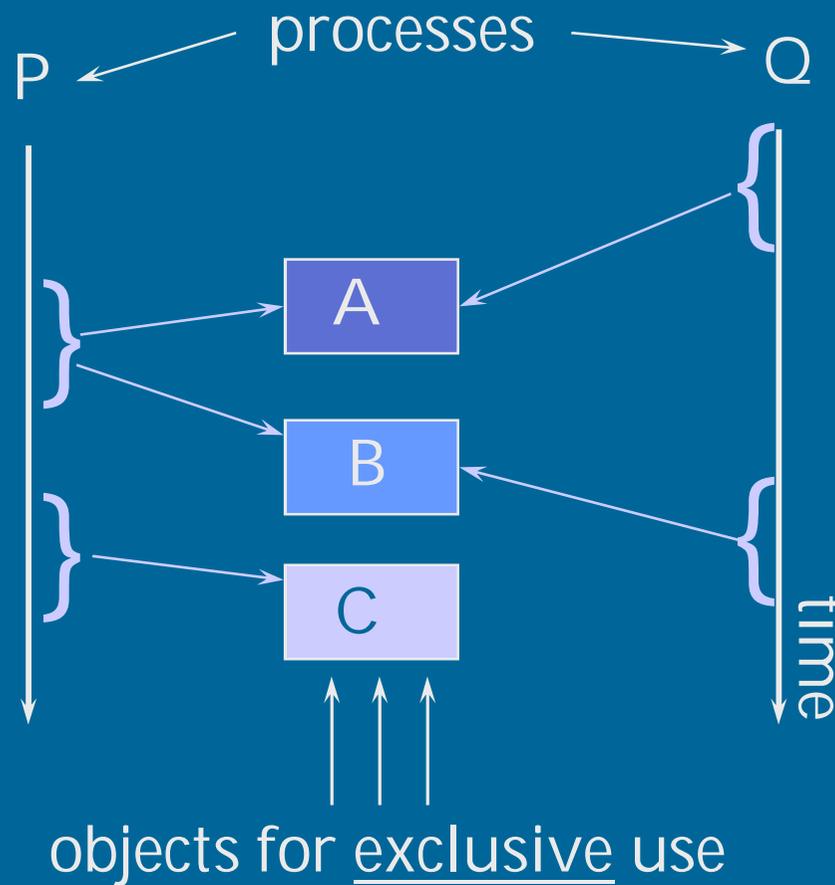
Dining Philosophers  
Deadlock occurrence  
Deadlock detection  
Deadlock prevention  
Deadlock avoidance

# 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=70740d6c93030b339b4ad17fd58ee135dfc13913>  
(search "i915\_enable\_vblank" ...)

# Deadlock: Background



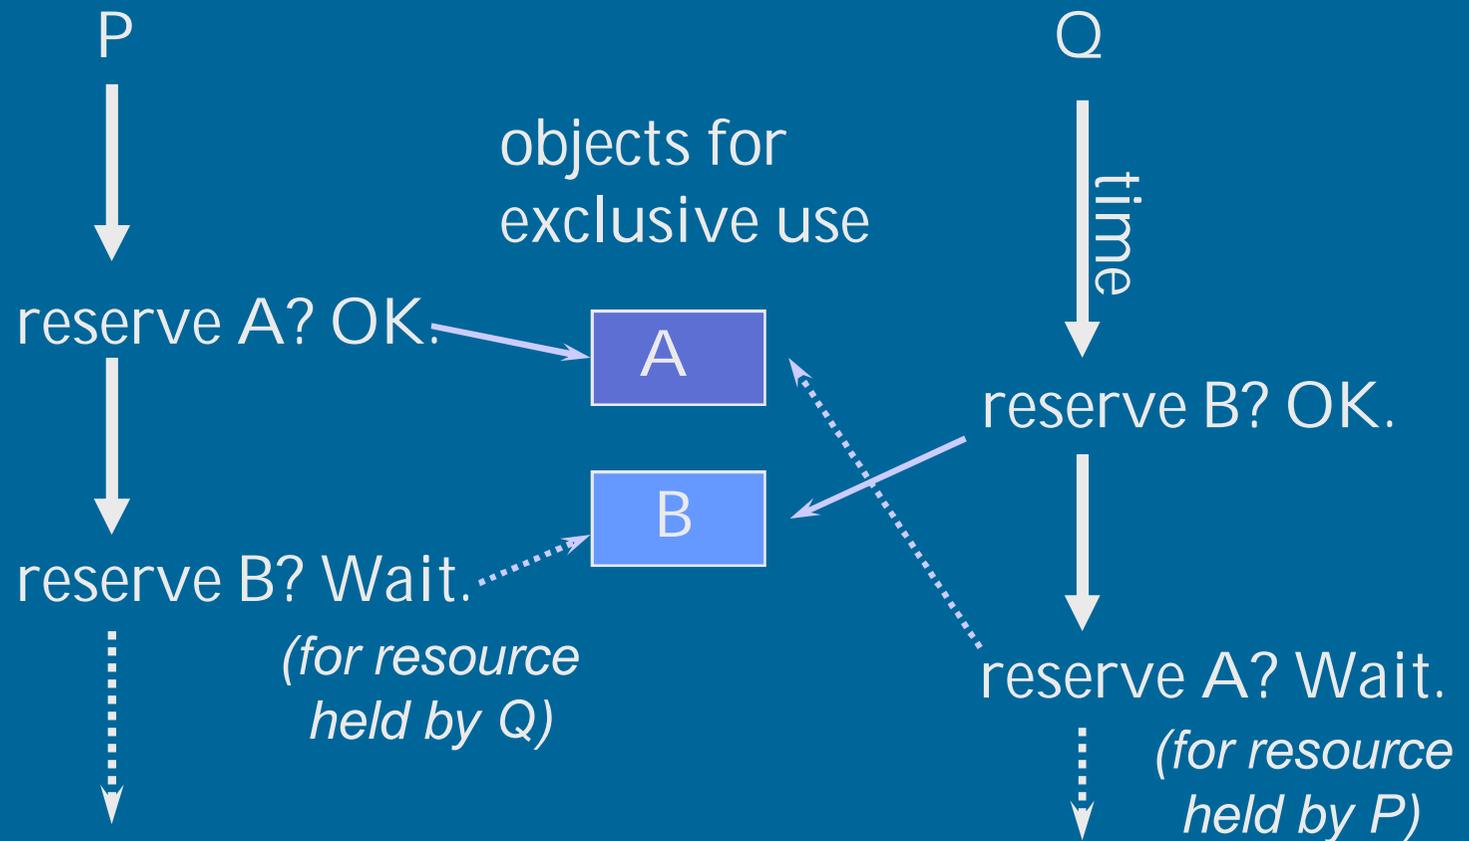
object?

buffer,  
page,  
user input,  
critic. section  
disk driver,  
scanner,  
message,  
...

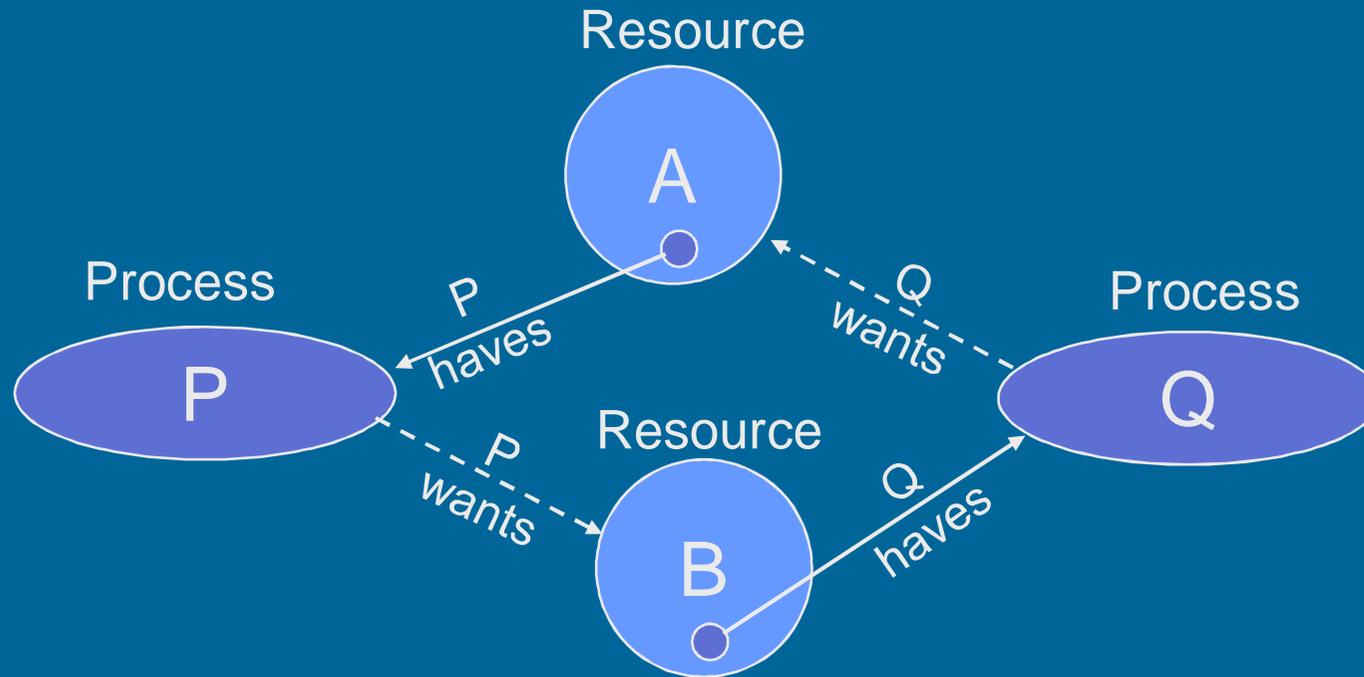
Basic problem: a process needs multiple objects at the same time

Mutex: competition for one object (critical section)

# Deadlock: an Example <sup>(10)</sup>

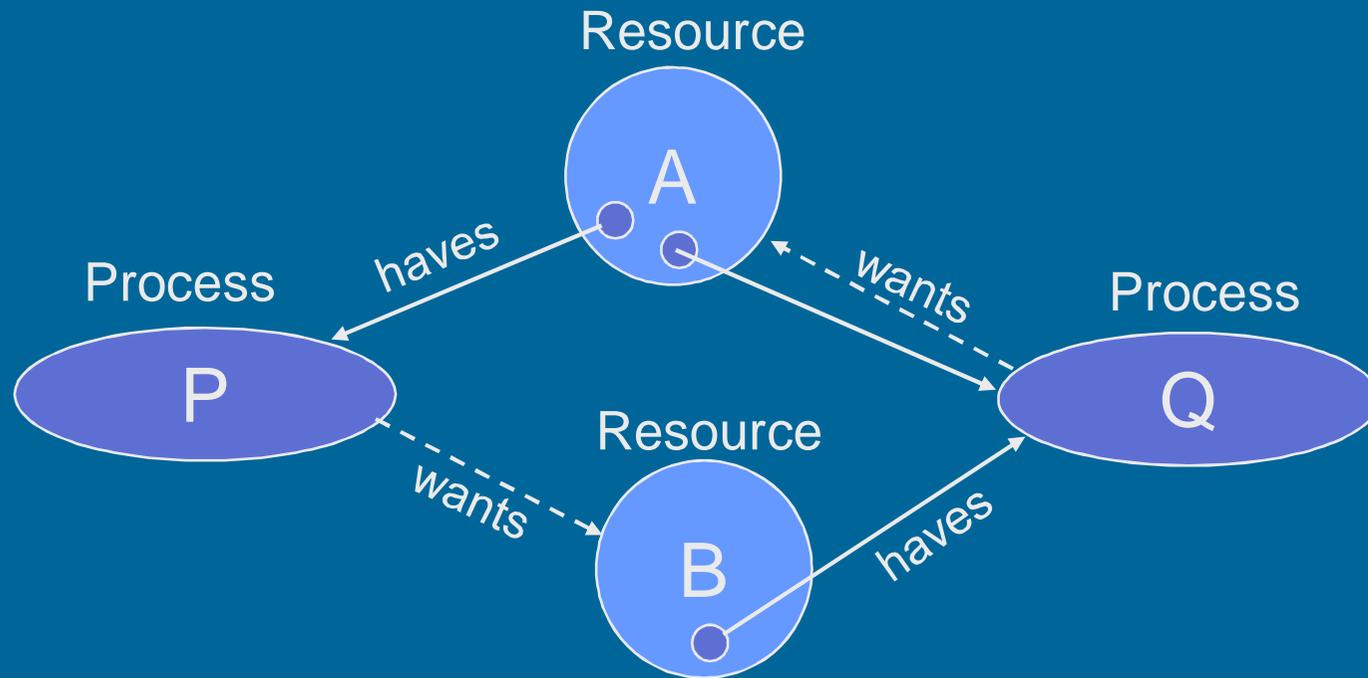


# Resource Reservation Graph



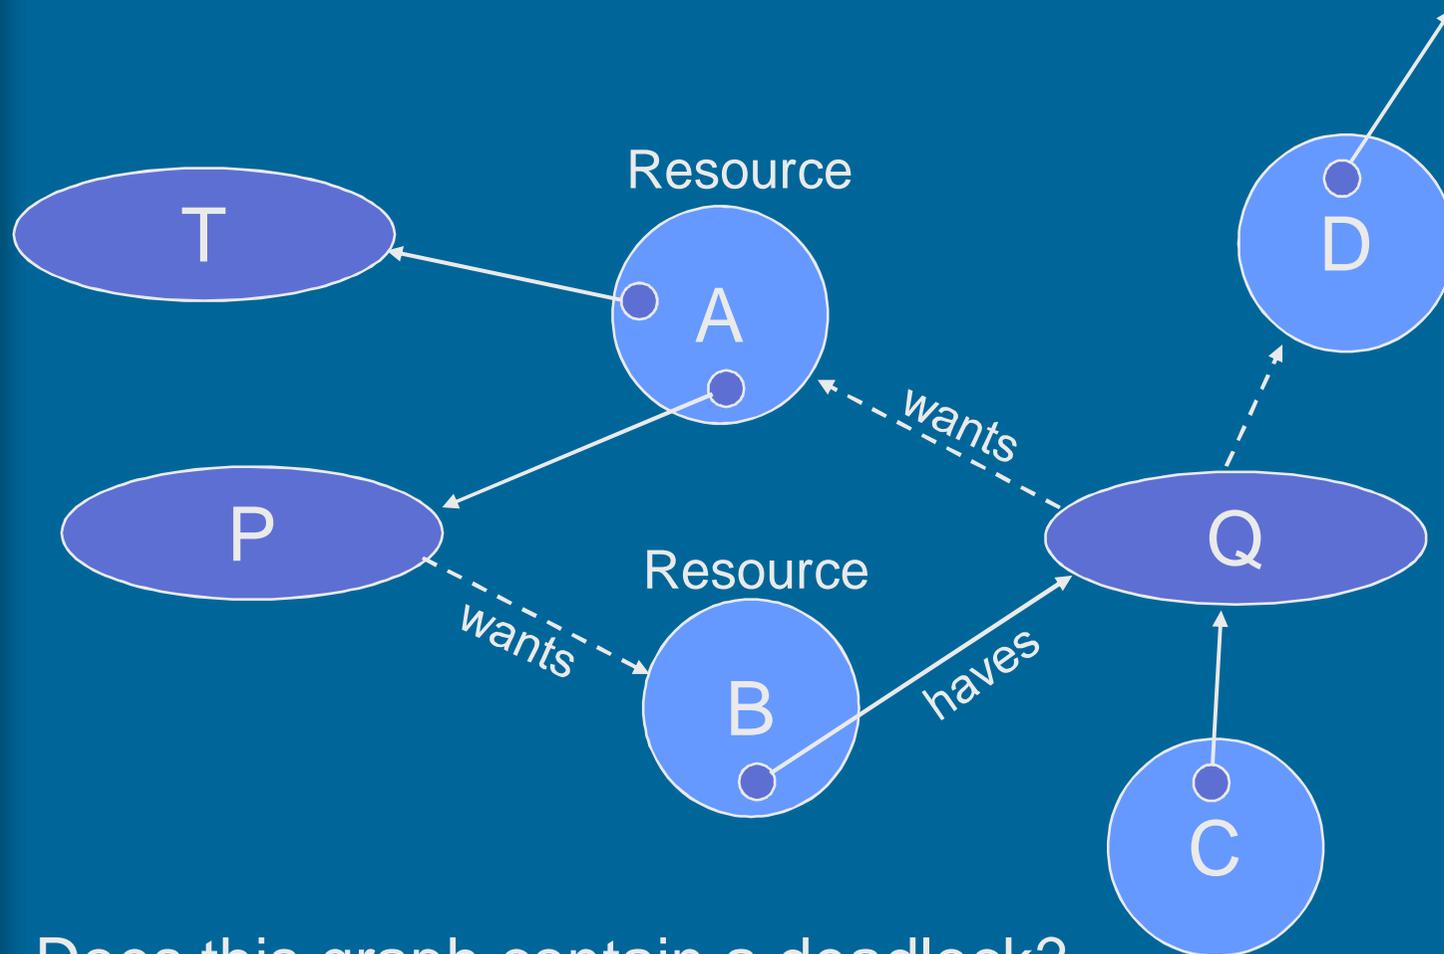
Deadlock cycle in resource reservation graph

# Resource Reservation Graph



Does this graph contain a deadlock?

# Resource Reservation Graph

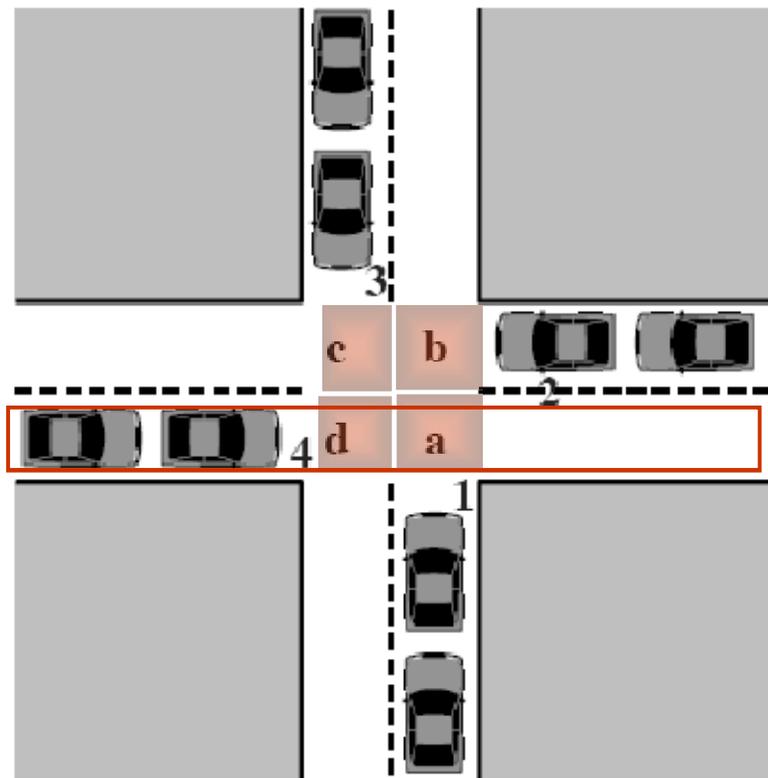


Does this graph contain a deadlock?

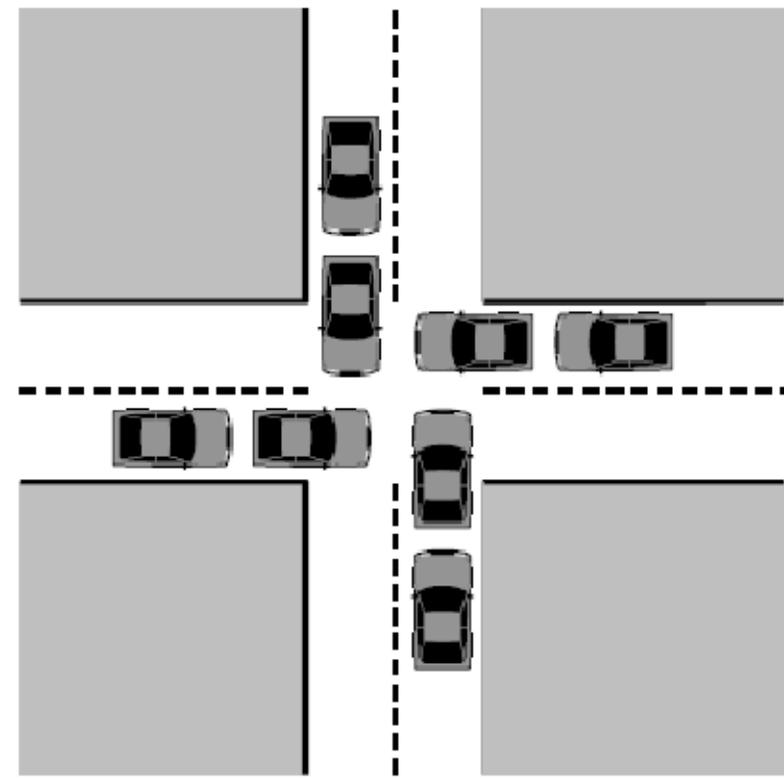
# Gridlock

(Fig. 6.1 [Stal06])

Real life gridlock: <http://img209.imageshack.us/img209/5781/deadlocknajokcomafarialibh3.jpg>



(a) Deadlock possible



(b) Deadlock

- Processes: cars 1, 2, 3 and 4
- Resources: quadrants a, b, c, d
  - Car 4 needs quadrants d and a (exclusive use for each)

# Consequences

- The processes do not advance
  - Cars do not move
- Resources remain reserved
  - Cpu? Street quadrant?
  - Memory? I/O-devices?
  - Logical resources (semaphores, critical sections, ...)?
- The computation fails
  - Execution never finishes?
    - One application?
  - The system crashes? Traffic flow becomes zero?

# Resources

- Reusable resources

- Limited number or amount
- Wait for it, allocate it, deallocate (free) it
- Memory, buffer space, intersection quadrant
- Critical section code segment execution
- ...

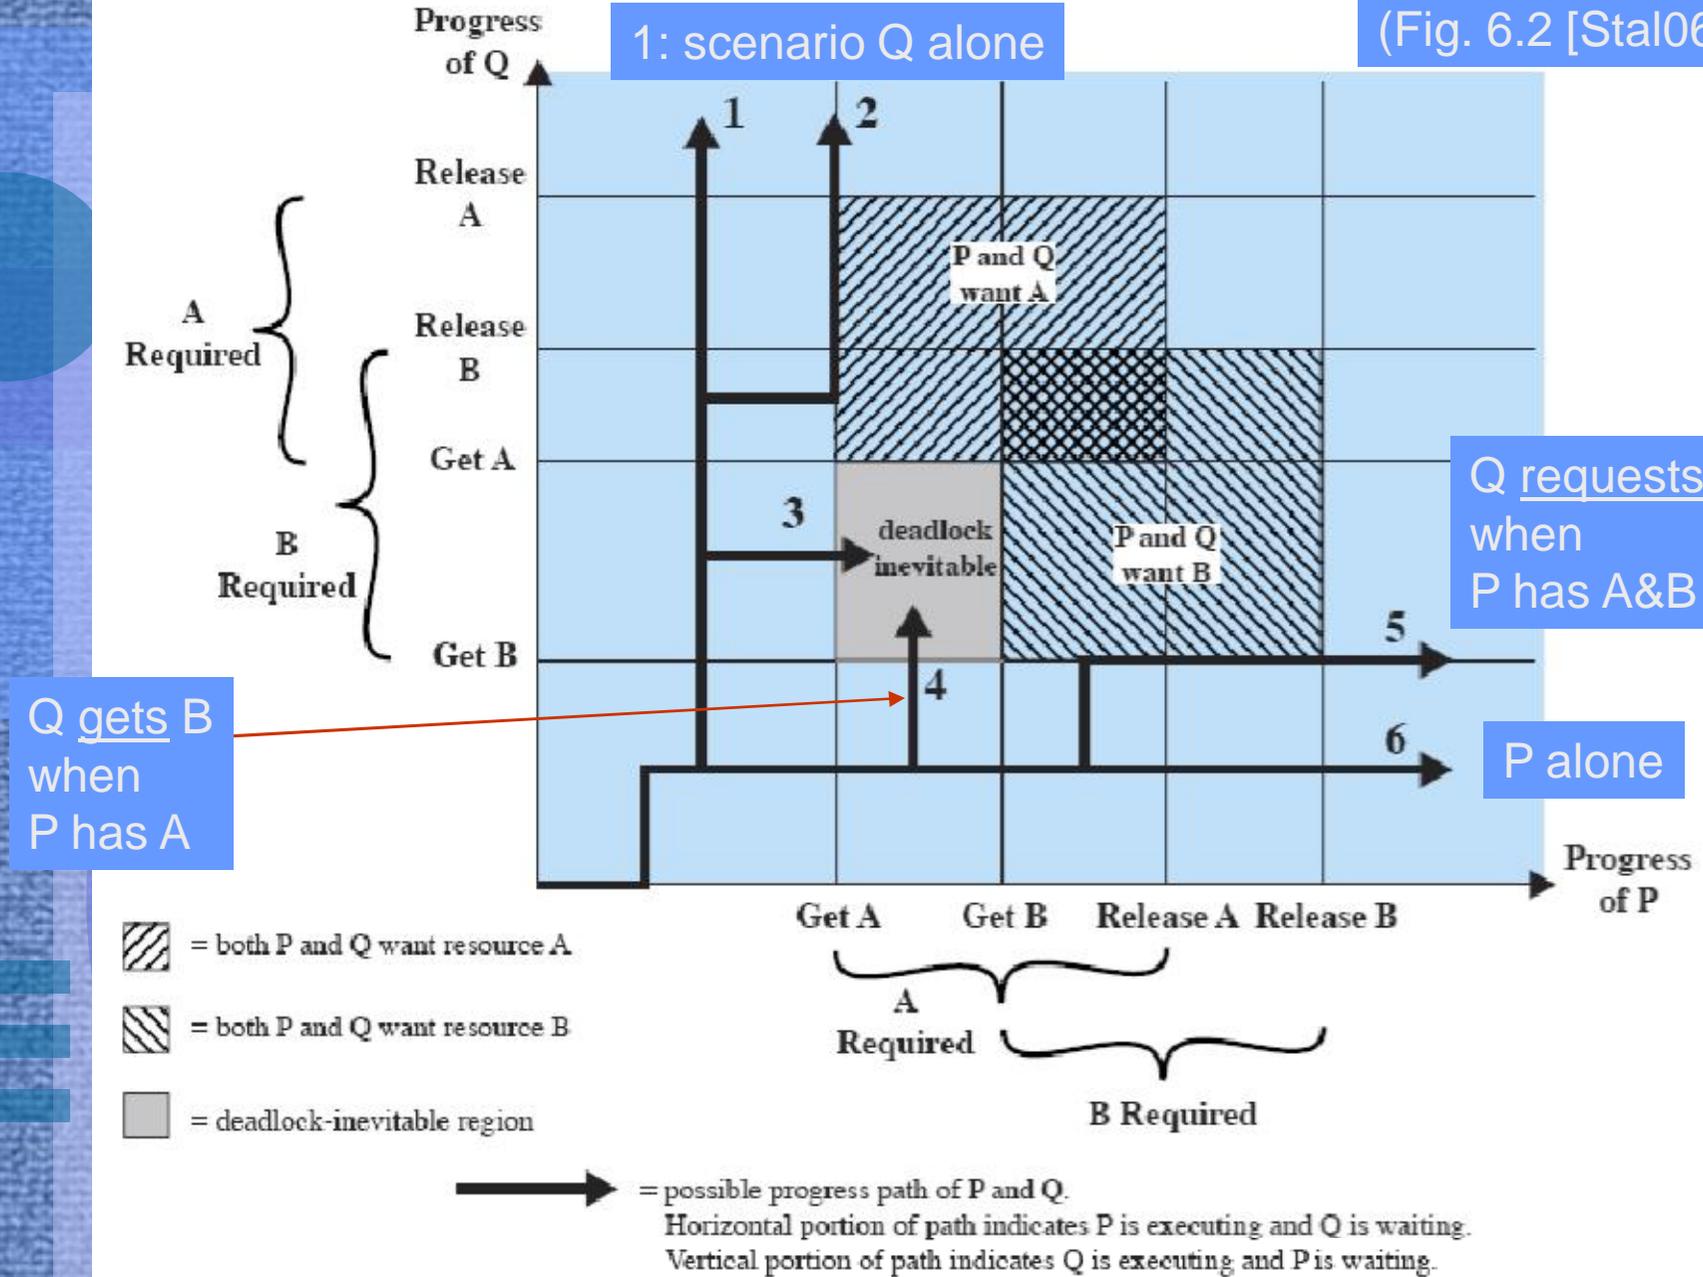
uudelleen-  
käytettävä  
resurssi

- Consumable resources

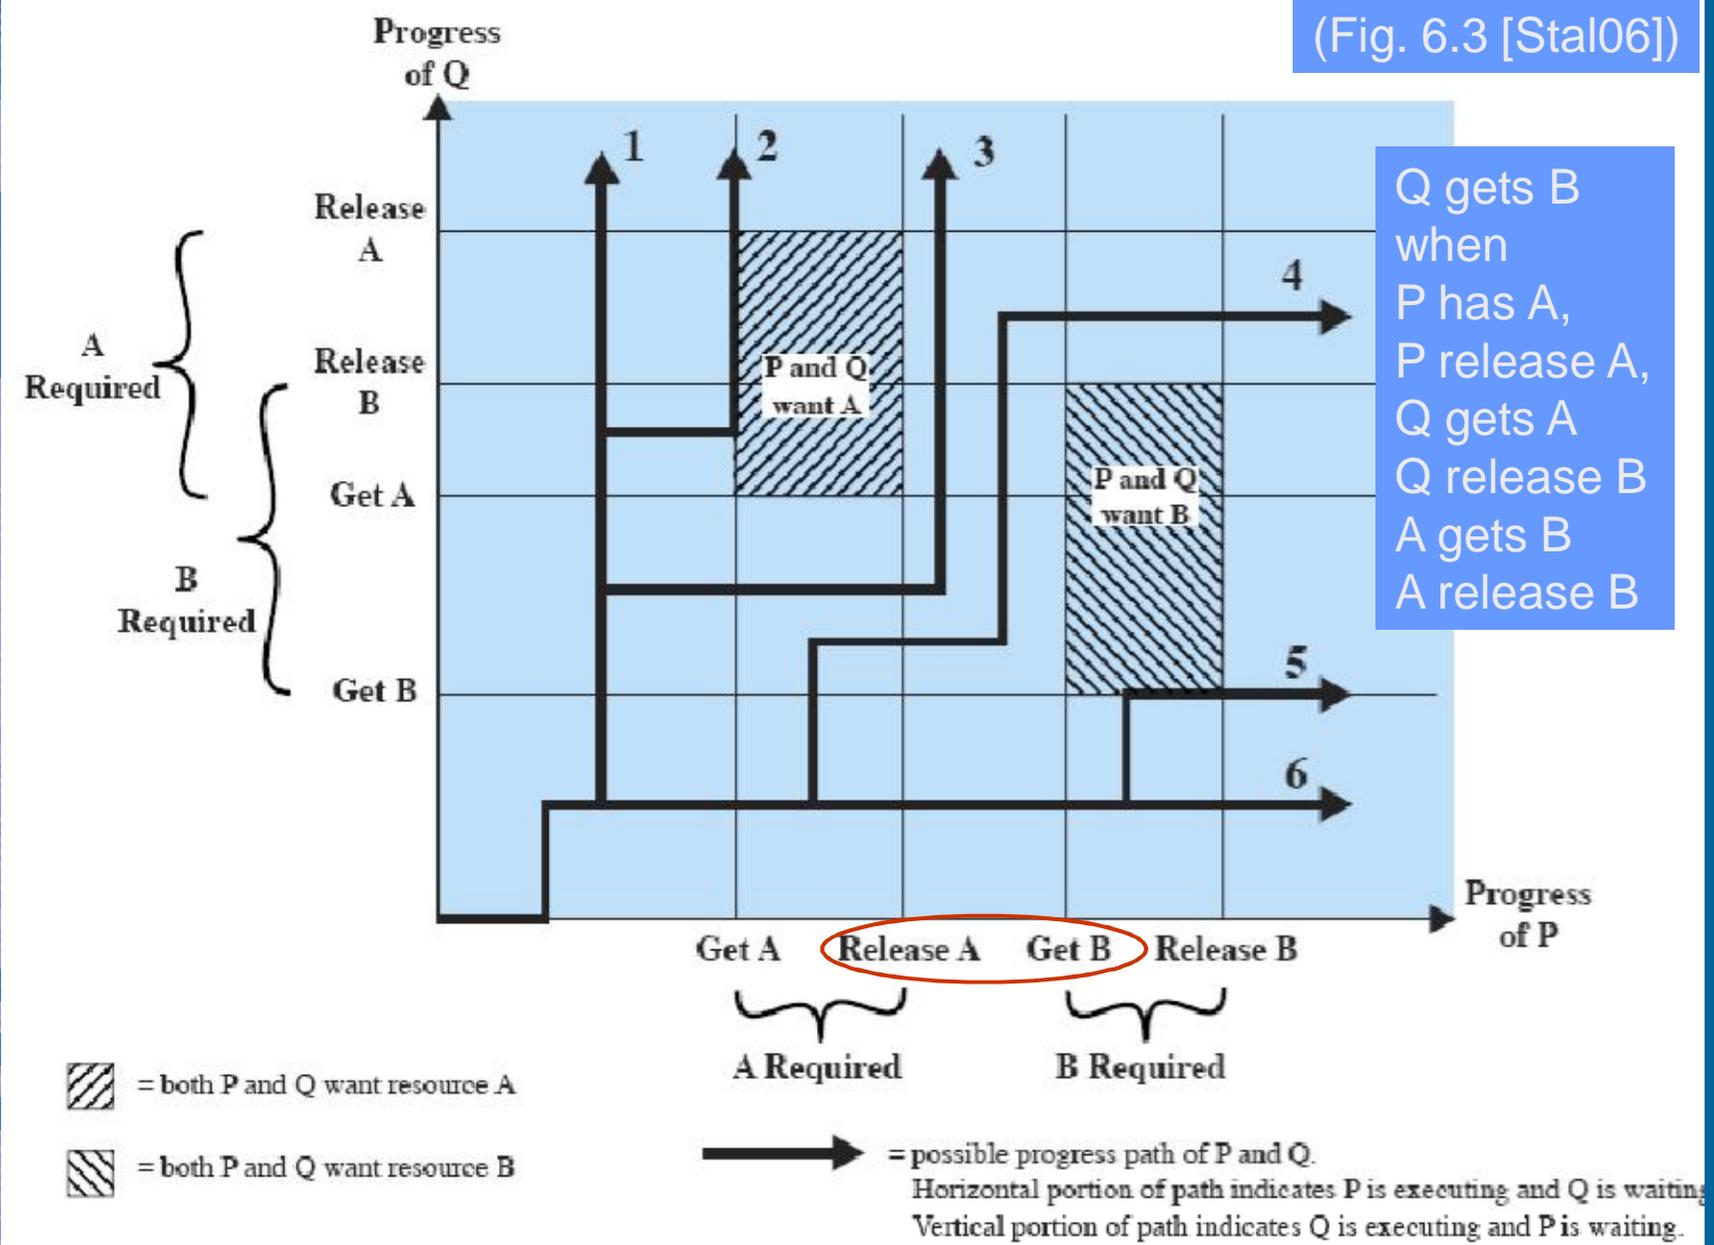
- Unlimited number or amount
- Created and consumed
- Someone may create it, wait for it, destroy it
- Message, interrupt, turn for critical section
- ...

kulutettava  
resurssi

(Fig. 6.2 [Stal06])



(Fig. 6.3 [Stal06])



# Definitions

- Deadlock

lukkiintuminen

- Eternal wait in blocked state
- Does not block processor (unless one resource is processor)

- Livelock

”elolukko”

- 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

nälkiintyminen

- the process will never get its turn
- E.g., in ready-to-run queue, but never scheduled

# 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?

# 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

# Conditions for Deadlock (6)

- Three policy conditions

Coffman, 1971

- S1. Mutual exclusion

yksi käyttäjä

- one user of any resource at a time (not just code)

- S2. Hold and wait

pidä 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



E.G. Coffman

- A dynamic (execution time) condition takes place

kehäodotus

- D1. Circular wait: a closed chain of processes exists, each process holds at least one resource needed by the next process in chain

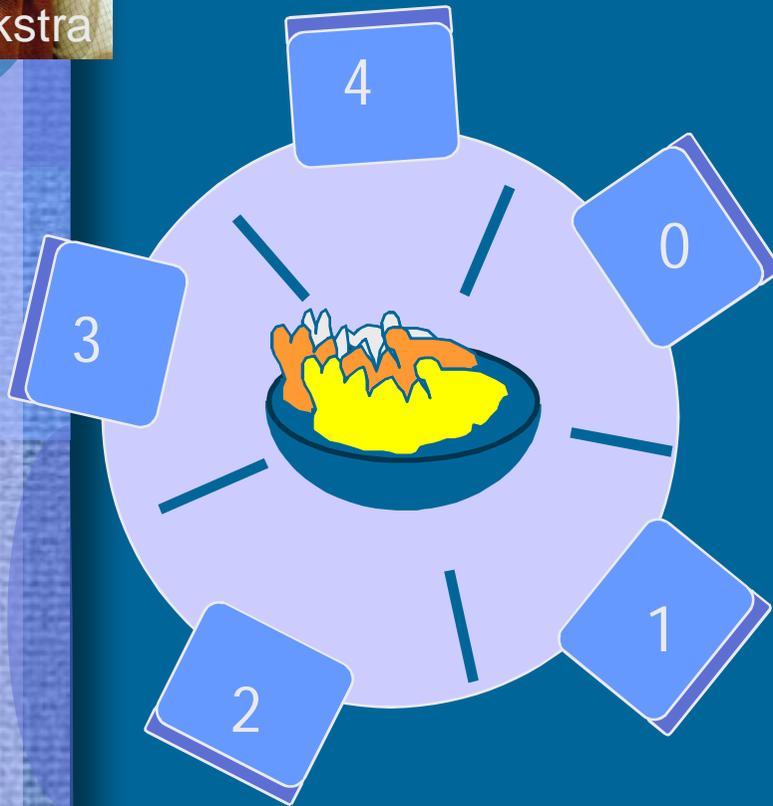
E.g., slide 5

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



Dijkstra

# Dining Philosophers (Dijkstra)



See philosopher art in web

## Philosopher:

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

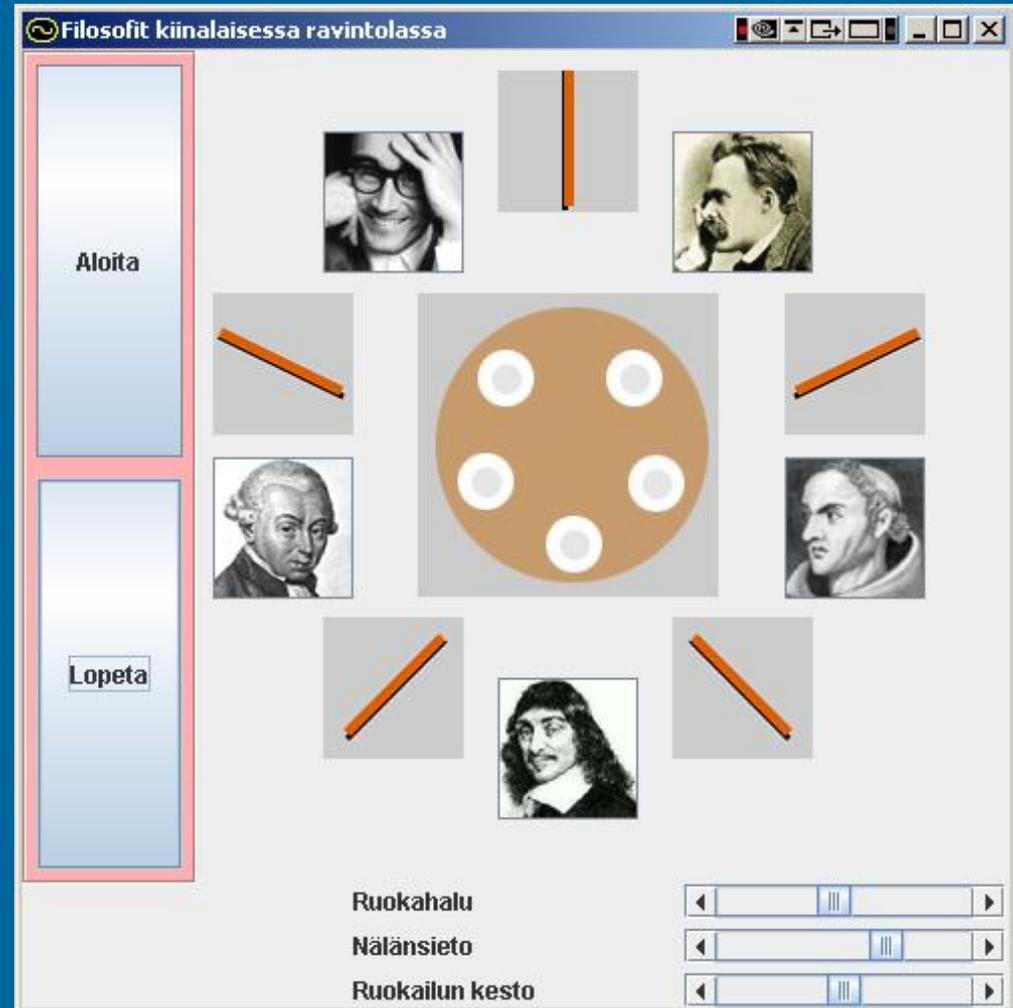
## Problem:

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?

[Lehtomaki.zip](#)



<http://www.cs.helsinki.fi/u/kerola/rio/Lehtomaki/Lehtomaki.zip>

```
/* 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));
}
```

# Trivial Solution #1

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

(Fig. 6.13 [Stal06])

```
/* 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));
}
```

- No deadlock, no starvation, and no company while eating – not good
- Waiting when resources are available – not good

which scenario?

# Deadlock Prevention

- How to prevent deadlock occurrence in advance?
- Deadlock possible only when all 4 conditions are met:
  - S1. Mutual exclusion poissulkemistarve
  - 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?

# 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?

# 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?)  $\frac{A}{B}$
  - 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 ...

# 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  $i \oplus 1$  do now? Think? Eat? Die?

# 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

# Deadlock Detection and Recovery (4)

- 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!

# Resource Allocation

- 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  $R_j$ "

How many R4 resources exists?

R1	R2	R3	R4	R5
2	1	1	2	1

Resource vector **R**

R1	R2	R3	R4	R5
0	0	0	0	1

Available vector **V**

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

Request matrix **Q**

	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

Allocation matrix **A**

Which resources are now free?

Who has now R4?

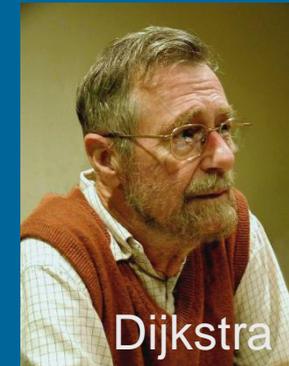
P2 has now R1 and R2,  
P2 wants now R3 and R5

Is there now a deadlock or not?

(Fig. 6.10 [Stal06])

# Deadlock Detection (Dijkstra) (4)

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



# Deadlock Detection Algorithm (DDA)

DL1. [*Remove the processes with no resources*]  
Mark all processes with null rows in  $\mathbf{A}$ .

DL2. [*Initialize counters for available objects*]  
Initialize a working vector  $\mathbf{W} = \mathbf{V}$

DL3. [*Search for a process  $P_i$  which could get all resources it requires*]  
Search for an unmarked row  $i$  such that  
$$q_{ij} \leq w_j \quad j = 1..n$$
  
If none is found terminate the algorithm.

DL4. [*Increase  $\mathbf{W}$  with the resources of the chosen process*]  
Set  $\mathbf{W} = \mathbf{W} + \mathbf{A}_{i^*}$  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?

# Example: Initial state

allocation matrix

A

row 1: 1 0 1 1 0  
2: 1 1 0 0 0  
3: 0 0 0 1 0  
4: 0 0 0 0 0

request matrix

Q

0 1 0 0 1  
0 0 1 0 1  
0 0 0 0 1  
1 0 1 0 1

E.g.,  
"process 2 has  
resources 1 & 2,  
and it wants  
resources 3 & 5"

all resources R 2 1 1 2 1

Who holds  
resource 4?

free resources V 0 0 0 0 1

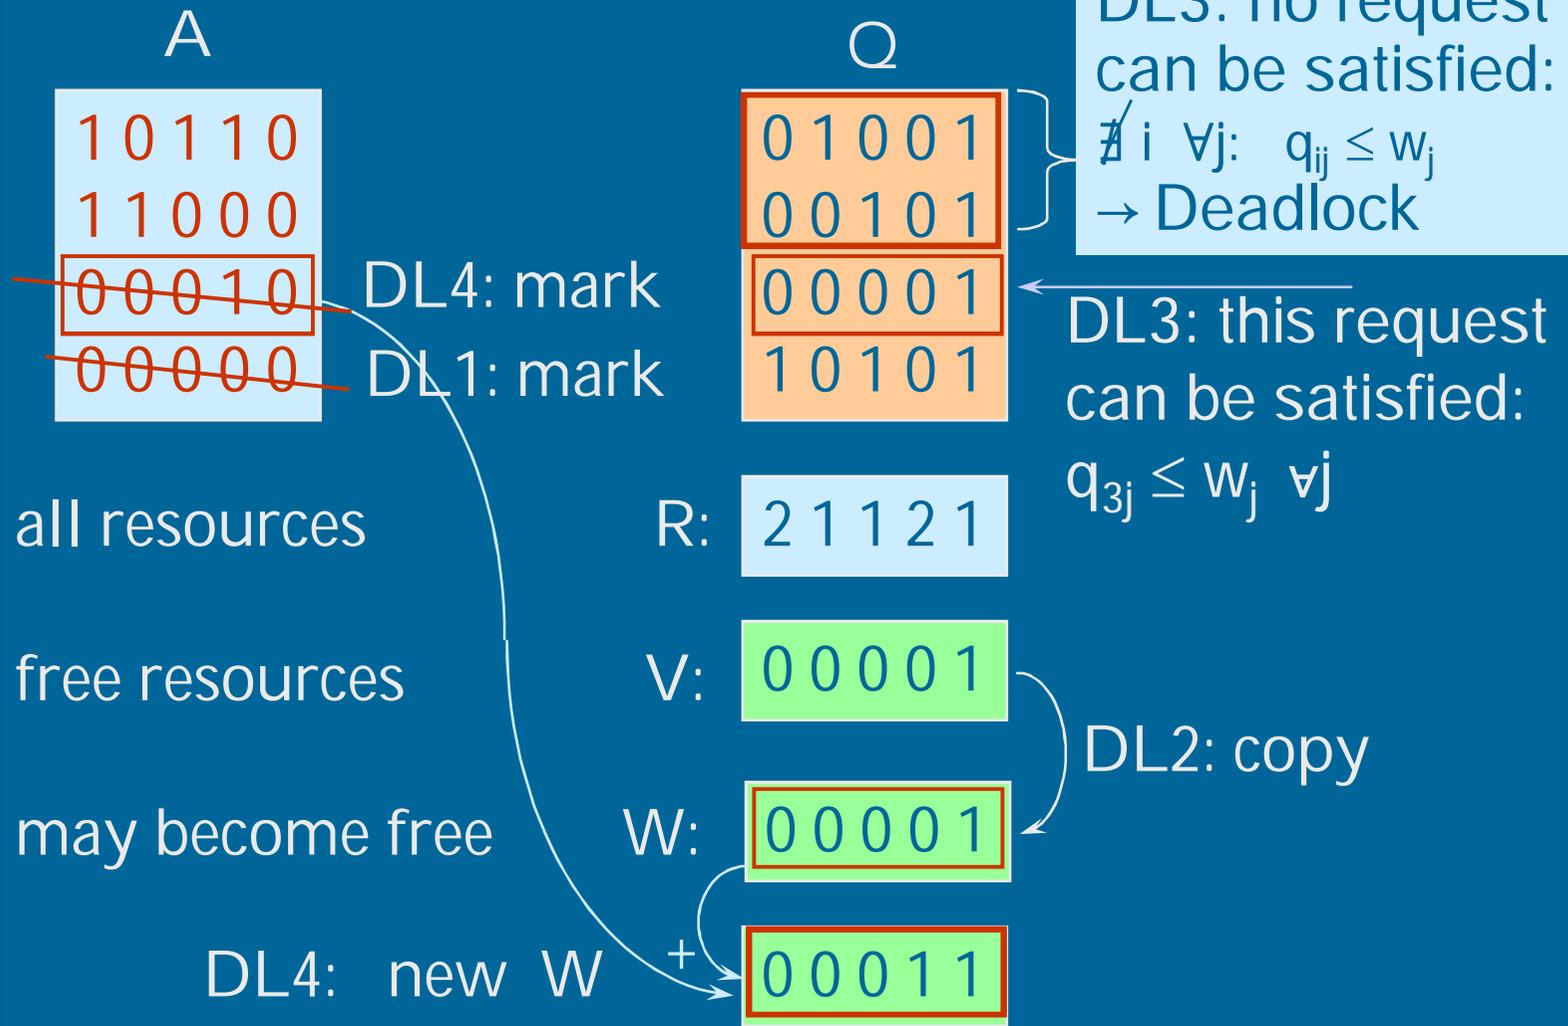
Which resources  
are free?

(Fig. 6.10 [Stal06])

Deadlock or not?

What now?

# Example: Deadlock Detection



# Example: Deadlock Detection (phases)

A

1	0	1	1	0
1	1	0	0	0
0	0	0	1	0
0	0	0	0	0

Q

0	1	0	0	1
0	0	1	0	1
0	0	0	0	1
1	0	1	0	1

all resources

R: 2 1 1 2 1

free resources

V: 0 0 0 0 1

may become free

W:

# Example: Deadlock Detection (phases)

A

```
1 0 1 1 0
1 1 0 0 0
0 0 0 1 0
0 0 0 0 0
```

DL1: mark

Q

```
0 1 0 0 1
0 0 1 0 1
0 0 0 0 1
1 0 1 0 1
```

all resources

R: 2 1 1 2 1

free resources

V: 0 0 0 0 1

may become free

W:

# Example: Deadlock Detection (phases)

A

```
1 0 1 1 0
1 1 0 0 0
0 0 0 1 0
0 0 0 0 0
```

DL1: mark

Q

```
0 1 0 0 1
0 0 1 0 1
0 0 0 0 1
1 0 1 0 1
```

all resources

R: 2 1 1 2 1

free resources

V: 0 0 0 0 1

may become free

W: 0 0 0 0 1

DL2: copy

# Example: Deadlock Detection (phases)

A

1	0	1	1	0
1	1	0	0	0
0	0	0	1	0
<del>0</del>	<del>0</del>	<del>0</del>	<del>0</del>	<del>0</del>

DL1: mark

all resources

R: 2 1 1 2 1

free resources

V: 0 0 0 0 1

may become free

W: 0 0 0 0 1

Q

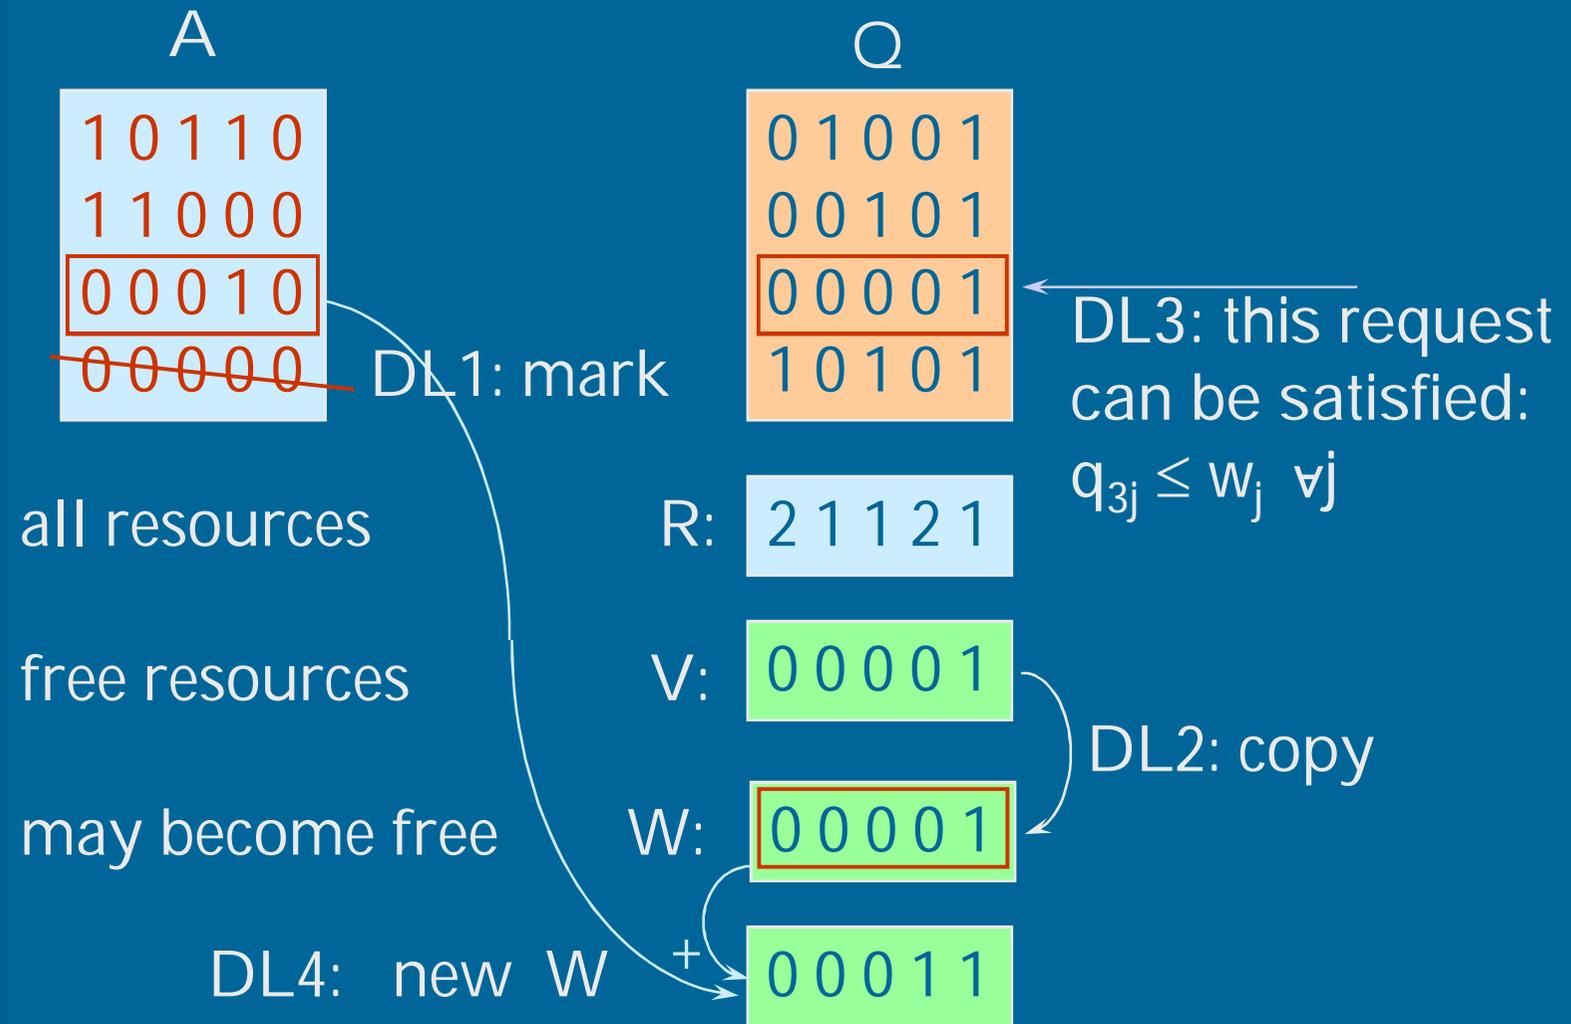
0	1	0	0	1
0	0	1	0	1
0	0	0	0	1
1	0	1	0	1

DL3: this request can be satisfied:

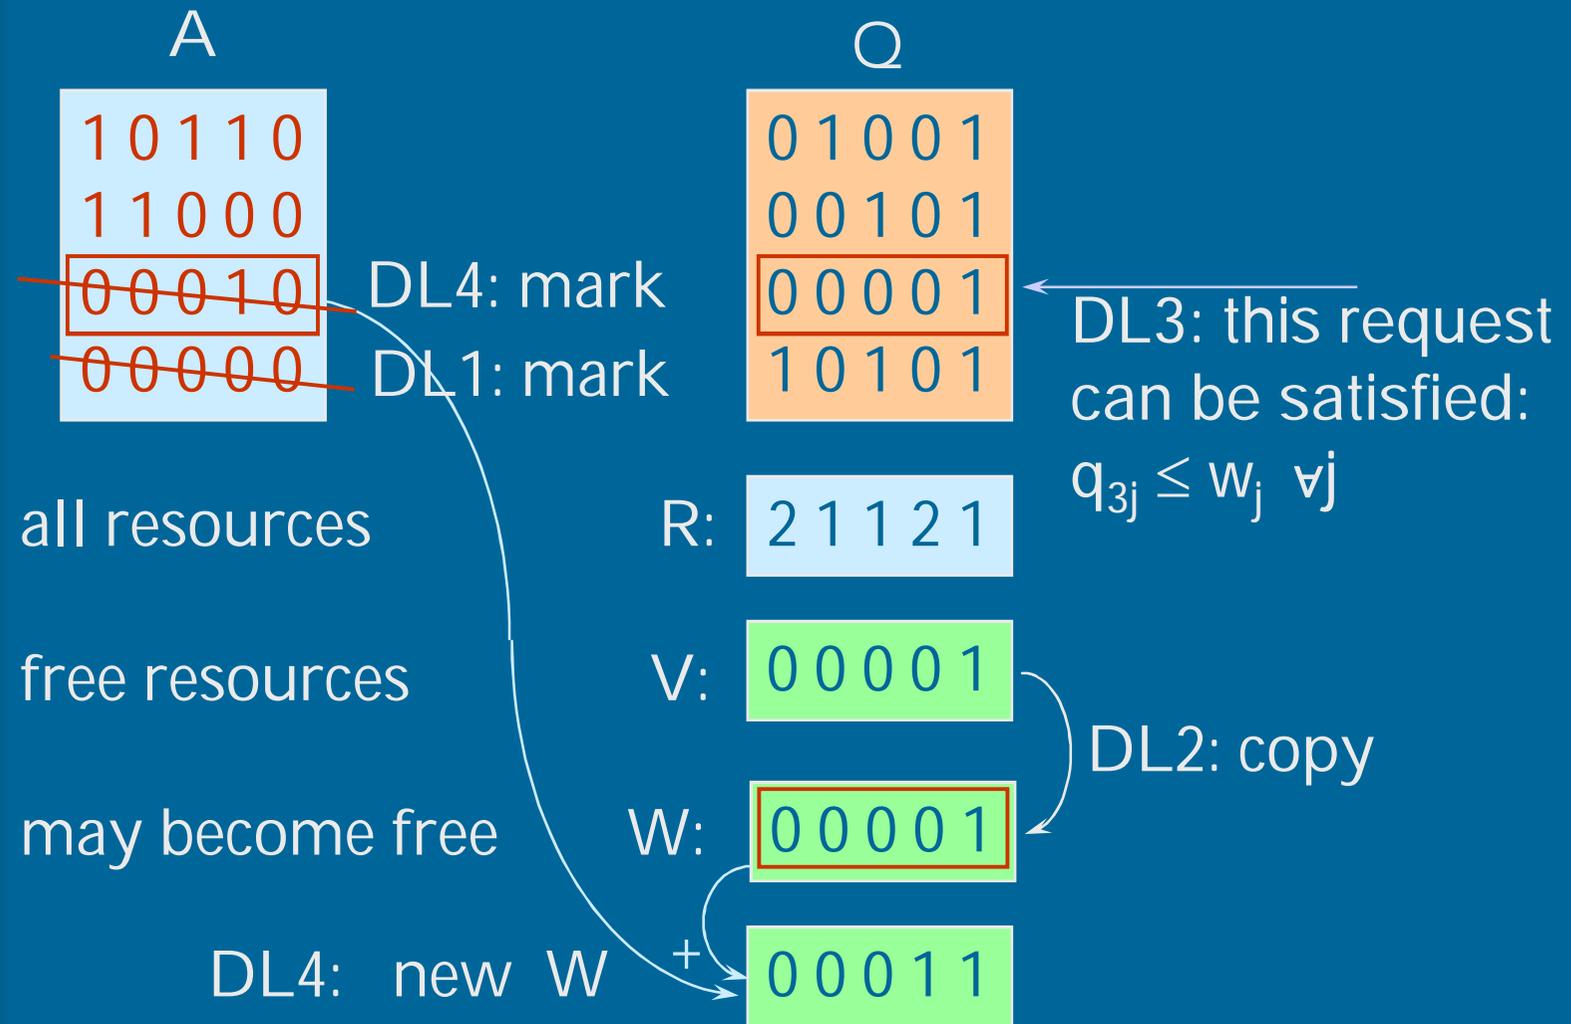
$$q_{3j} \leq w_j \quad \forall j$$

DL2: copy

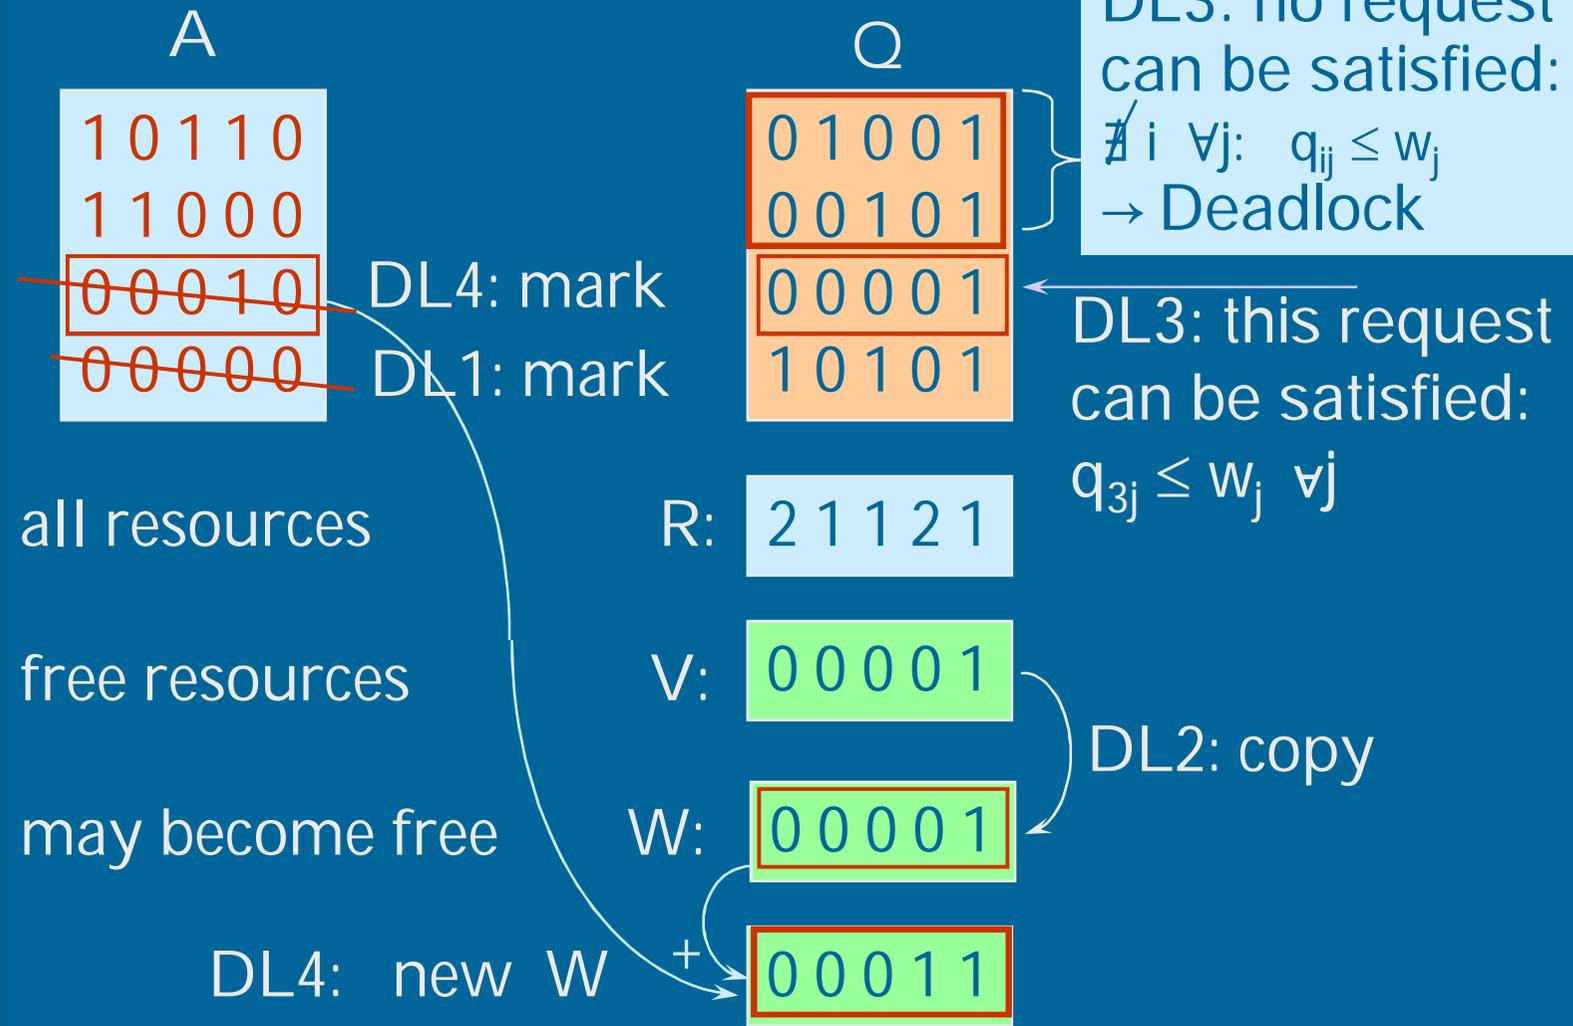
# Example: Deadlock Detection (phases)



# Example: Deadlock Detection (phases)



# Example: Deadlock Detection (phases)



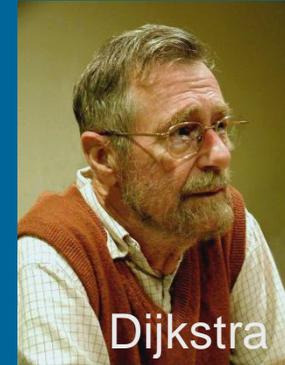
# Example: Breaking Deadlocks

- 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 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?

# 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"

# Banker's Algorithm (6)



- 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

# Deadlock Avoidance with Banker's Algorithm (6)

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

Possible allocation

Possible request

# Banker's Algorithm Example

## Allocation A

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

## Requests Q

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

## Max allocation C

	R1	R2	R3	R4	R5
P1	2	1	0	1	0
P2	1	1	0	0	1
P3	1	0	1	1	1
P4	0	2	1	1	1

## Resources R

2	3	2	1	2
---	---	---	---	---

R1 R2 R3 R4 R5

## Available V

1	1	0	0	1
---	---	---	---	---

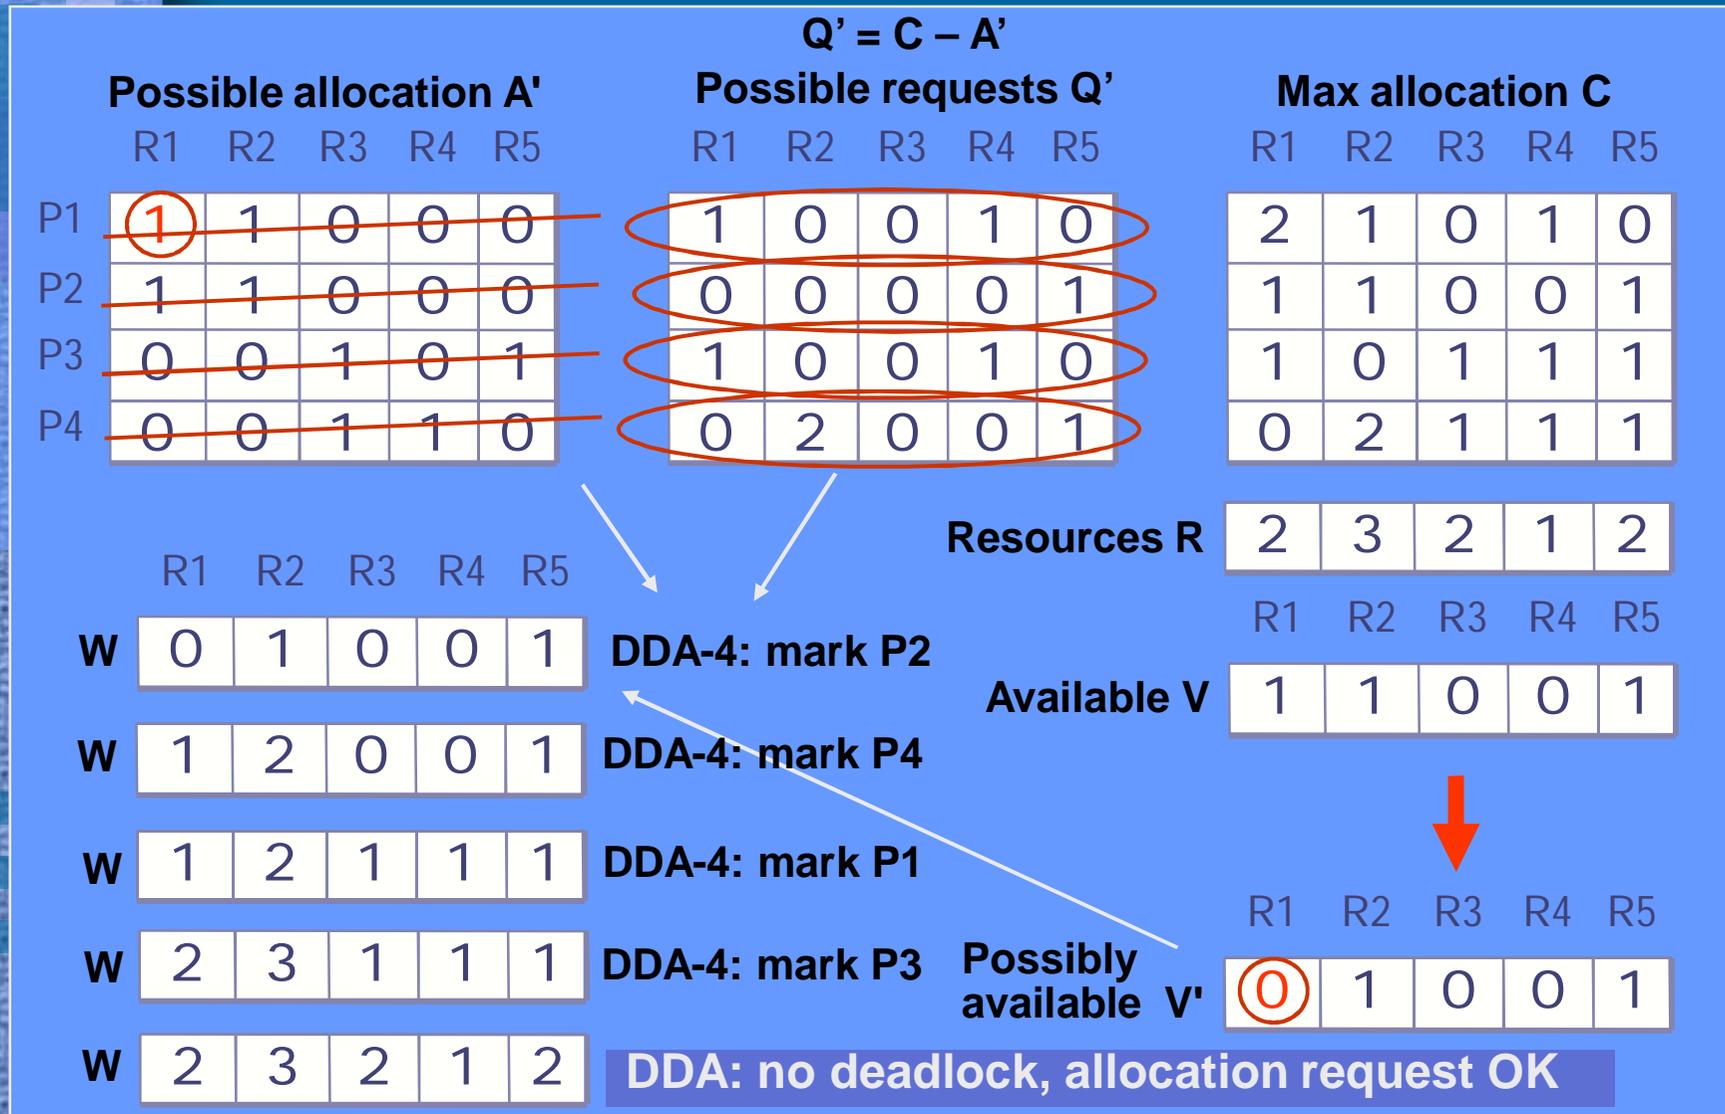
(Fig. 16.11, Bacon, Concurrent Systems, 1993)

P1 requests R1. Is request granted?

Could system deadlock, if R1 is granted?

# Banker's Algorithm Example (7)

If P1 request for R1 approved, can deadlock occur?



# 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>

# 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 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