


CSE 451: Operating Systems
Autumn 2008

Deadlock

H. Levy



10/17/2008

(Is Google the greatest, or what?)

2


Definition

- A thread is deadlocked when it's waiting for an event that can never occur
 - I'm waiting for you to clear the intersection, so I can proceed
 - but you can't move until he moves, and he can't move until she moves, and she can't move until I move
 - thread A is in critical section 1, waiting for access to critical section 2; thread B is in critical section 2, waiting for access to critical section 1
 - I'm trying to book a vacation package to Tahiti – air transportation, ground transportation, hotel, side-trips. It's all-or-nothing – one high-level transaction – with the four databases locked in that order. You're trying to do the same thing in the opposite order.

10/17/2008

3

Resource graph



- A deadlock exists if there is an *irreducible cycle* in the resource graph (such as the one above)

10/17/2008

4


Graph reduction

- A graph can be *reduced* by a thread if all of that thread's requests can be granted
 - in this case, the thread eventually will terminate – all resources are freed – all arcs (allocations) to it in the graph are deleted
- Miscellaneous theorems (Holt, Havender):
 - There are no deadlocked threads iff the graph is completely reducible
 - The order of reductions is irrelevant
- (Detail: resources with multiple units)

10/17/2008

5

Resource allocation graph with no cycle




What would cause a deadlock?

10/17/2008

Silberschatz, Galvin and Gagne ©2002

6

Resource allocation graph with a deadlock




10/17/2008

Silberschatz, Galvin and Gagne ©2002

7

Resource allocation graph with a cycle but no deadlock



10/17/2008

Silberschatz, Galvin and Gagne ©2002

8

Approaches to deadlock

- Prevention – don't let deadlock occur
 1. each thread obtains all resources at the beginning; blocks until all are available
 - drawback?
 2. resources are numbered; each thread obtains them in sequence (which means acquiring some before they are actually needed)
 - why does this work?
 - pros and cons?
 3. each thread states its maximum claim for every resource type; system runs the Banker's algorithm at each allocation request
 - if I were to allocate you that resource, and then everyone were to request their maximum claim for every resource, would there be a deadlock?
 - how do I tell if there would be a deadlock?

10/17/2008

9

Approaches (cont'd.)

- Detection and correction
 - every once in a while, check to see if there's a deadlock
 - how?
 - if so, eliminate it
 - how?

10/17/2008

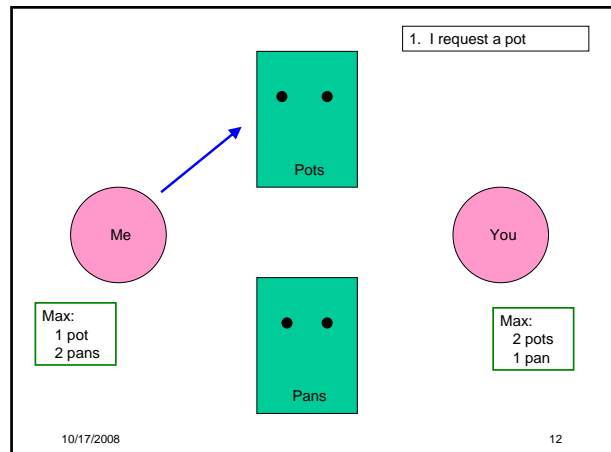
10

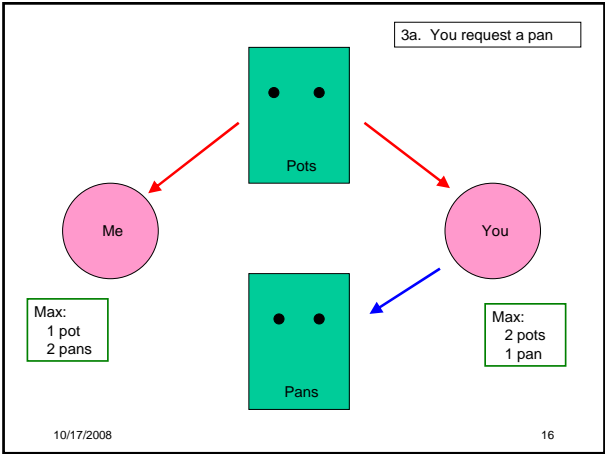
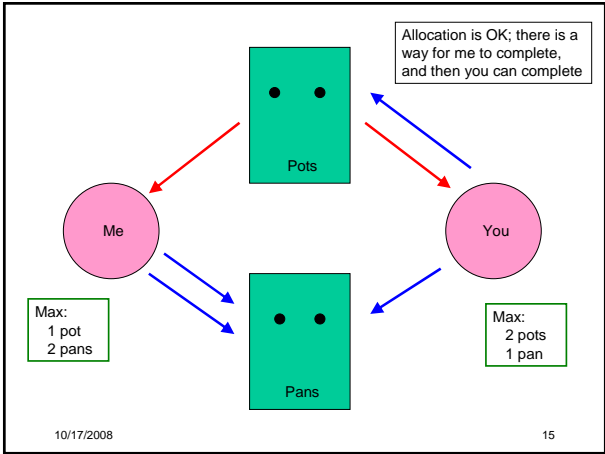
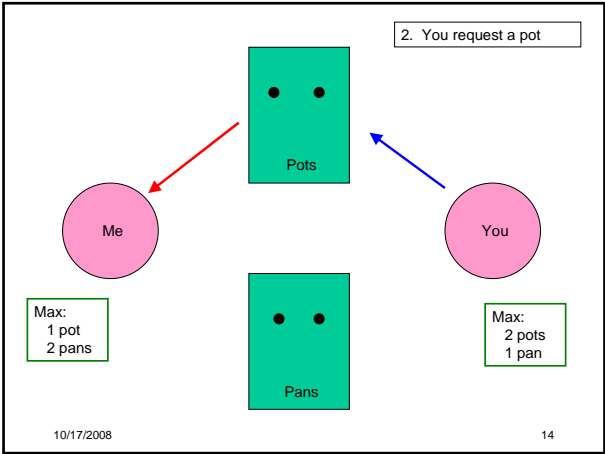
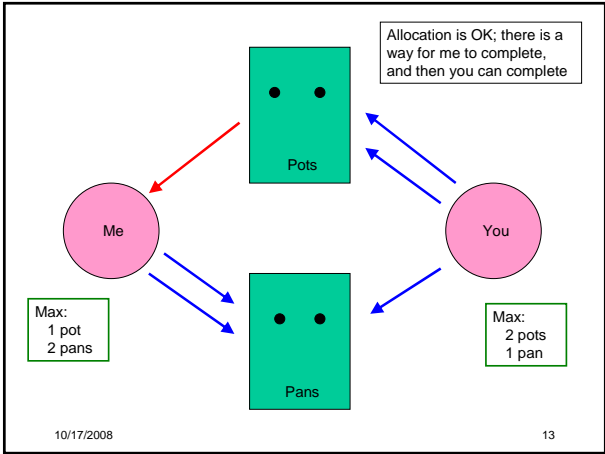
Banker's Algorithm example

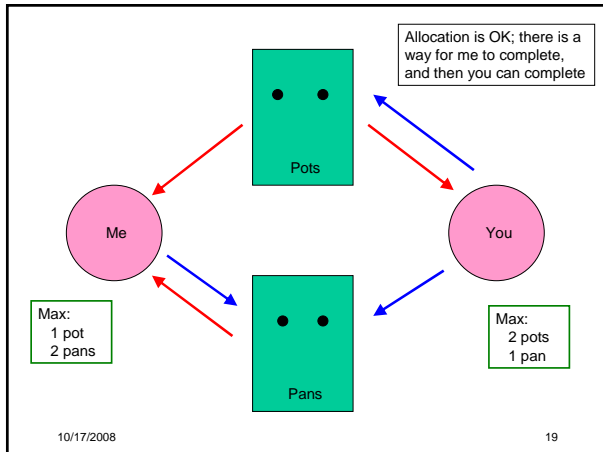
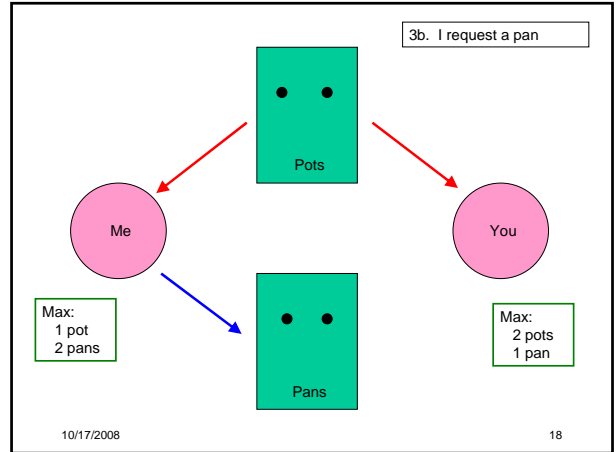
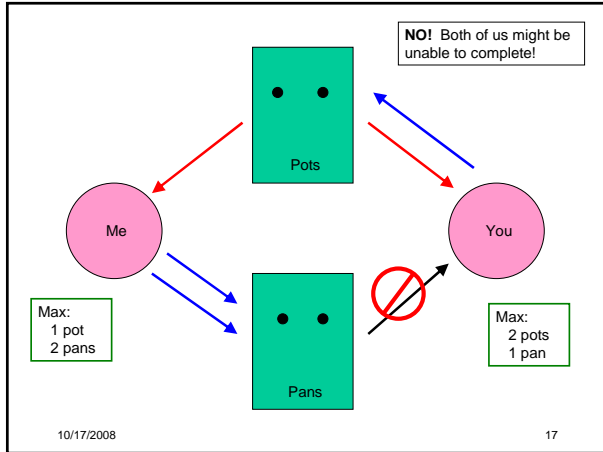
- When a request is made
 - pretend you granted it
 - pretend all other legal requests were made
 - can the graph be reduced?
 - if so, allocate the requested resource
 - if not, block the thread

10/17/2008

11







Current practice

- Microsoft SQL Server
 - "The SQL Server Database Engine automatically detects deadlock cycles within SQL Server. The Database Engine chooses one of the sessions as a deadlock victim and the current transaction is terminated with an error to break the deadlock."
- Oracle
 - As Microsoft SQL Server, plus "Multitable deadlocks can usually be avoided if transactions accessing the same tables lock those tables in the same order... For example, all application developers might follow the rule that when both a master and detail table are updated, the master table is locked first and then the detail table."

10/17/2008 20

More current practice....

- Windows internals (Linux no different)
 - "... the NT kernel architecture is a deadlock minefield. With the multi-threaded re-entrant kernel there is plenty of deadlock potential."
 - "Lock ordering is great in theory, and NT was originally designed with mutex levels, but they had to be abandoned. Inside the NT kernel there is a lot of interaction between memory management, the cache manager, and the file systems, and plenty of situations where memory management (maybe under the guise of its modified page writer) acquires its lock and then calls the cache manager. This happens while the file system calls the cache manager to fill the cache which in turn goes through the memory manager to fault in its page. And the list goes on."
 - Enhancements to semaphores for file systems:
 - Reacquire by owning thread
 - Test for ownership before acquiring

10/17/2008

21

Summary

- Deadlock is bad!
- We can deal with it either statically (prevention) or dynamically (avoidance and detection)
- In practice, you'll encounter lock ordering, periodic deadlock detection/correction, and minefields

10/17/2008

22