

## Operating Systems Principles

### Deadlock, Prevention, and Avoidance

Mark Kampe  
(markk@cs.ucla.edu)

## Deadlock Prevention and Avoidance

- 10A. Overview
- 10B. Deadlock Avoidance
- 10C. Deadlock Prevention
- 10D. Related Topics
- 10E. Higher level synchronization

Deadlock, Prevention and Avoidance

2

## Why Study Deadlocks?

- A major peril in cooperating parallel processes
  - they are relatively common in complex applications
  - they result in catastrophic system failures
- Finding them through debugging is very difficult
  - they happen intermittently and are hard to diagnose
  - they are much easier to prevent at design time
- Once you understand them, you can avoid them
  - most deadlocks result from careless/ignorant design
  - an ounce of prevention is worth a pound of cure

Deadlock, Prevention and Avoidance


3

## The Dining Philosophers Problem

Five philosophers  
five plates of pasta  
five forks

they eat whenever  
they choose to

one requires two  
forks to eat pasta,  
but must take them  
one at a time



they will not  
negotiate with  
one-another

the problem  
demands an  
absolute solution

Deadlock, Prevention and Avoidance

4

## (The Dining Philosophers Problem)

- the classical illustration of deadlocking
- it was created to illustrate deadlock problems
- it is a very artificial problem
  - it was carefully designed to cause deadlocks
  - changing the rules eliminate deadlocks
  - but then it couldn't be used to illustrate deadlocks

Deadlock, Prevention and Avoidance

5

## Deadlocks May Not Be Obvious

- process resource needs are ever-changing
  - depending on what data they are operating on
  - depending on where in computation they are
  - depending on what errors have happened
- modern software depends on many services
  - most of which are ignorant of one-another
  - each of which requires numerous resources
- services encapsulate much complexity
  - we do not know what resources they require
  - we do not know when/how they are serialized

Deadlock, Prevention and Avoidance

6

## Many Types of Deadlocks

- Different deadlocks require different solutions
- Commodity resource deadlocks
  - e.g. memory, queue space
- General resource deadlocks
  - e.g. files, critical sections
- Heterogeneous multi-resource deadlocks
  - e.g. P1 needs a file, P2 needs memory
- Producer-consumer deadlocks
  - e.g. P1 needs a file, P2 needs a message from P1

Deadlock, Prevention and Avoidance

7

## Approaches

- Avoidance
  - evaluate each proposed action
  - avoid taking actions that would deadlock
- Prevention
  - design system to make deadlock impossible
- Detection and Recovery
  - wait for it to happen
  - try to detect that it has happened
  - take some action to break the deadlock

Deadlock, Prevention and Avoidance

8

## Commodity vs. General Resources

- Commodity Resources
  - clients need an amount of it (e.g. memory)
  - deadlocks result from over-commitment
  - avoidance can be done in resource manager
- General Resources
  - clients need a specific instance of something
    - a particular file or semaphore
    - a particular message or request completion
  - deadlocks result from specific dependency network
  - prevention is usually done at design time

Deadlock, Prevention and Avoidance

9

## Commodity Resource Problems

- memory deadlock
  - we are out of memory
  - we need to swap some processes out
  - we need memory to build the I/O request
- critical resource exhaustion
  - a process has just faulted for a new page
  - there are no free pages in memory
  - there are no free pages on the swap device

Deadlock, Prevention and Avoidance

10

## Avoidance – Advance Reservations

- advance reservations for commodities
  - resource manager tracks outstanding reservations
  - only grants reservations if resources are available
- over-subscriptions are detected early
  - before processes ever get the resources
- client must be prepared to deal with failures
  - but these do not result in deadlocks
- dilemma: over-booking vs. under-utilization

Deadlock, Prevention and Avoidance

11

## Real Commodity Resource Management

- advanced reservation mechanisms are common
  - Unix setbreak system call to allocate more memory
  - disk quotas, Quality of Service contracts
- once granted, reservations are guaranteed
  - allocation failures only happen at reservation time ... hopefully before the new computation has begun
  - failures will not happen at request time
  - system behavior more predictable, easier to handle
- but clients must deal with reservation failures

Deadlock, Prevention and Avoidance

12

## Dealing with Rejection

- resource reservation eliminates difficult failures
  - recovering from a failure in mid-computation
- apps must still deal with reservation failures
  - application design should handle failures gracefully
    - e.g. refuse to perform new request, but continue running
  - app must have a way of reporting failure to requester
    - e.g. error messages or return codes
  - app must be able to continue running
    - all truly critical resources must be reserved at start-up time
- hold resources, wait & try again doesn't help

Deadlock, Prevention and Avoidance

13

## Pre-reserving critical resources

- system services must never deadlock for memory
- potential deadlock: swap manager
  - invoked to swap out processes to free up memory
  - may need to allocate memory to build I/O request
  - If no memory available, unable to swap out processes
- solution
  - pre-allocate and hoard a few request buffers
  - keep reusing the same ones over and over again
  - little bit of hoarded memory is a small price to pay

Deadlock, Prevention and Avoidance

14

## Deadlock Prevention

- necessary condition #1: mutual exclusion
  - P1 cannot use a resource until P2 releases it
- necessary condition #2: hold and wait
  - process already has R1 blocks to wait for R2
- necessary condition #3: no preemption
  - R1 cannot be taken away from P1
- necessary condition #4: circular dependency
  - P1 has R1, and needs R2
  - P2 has R2, and needs R1

Deadlock, Prevention and Avoidance

15

## Attack #1 – Mutual Exclusion

deadlock requires mutual exclusion

- P1 having the resource precludes P2 from getting it
- you can't deadlock over a shareable resource
  - perhaps maintained with atomic instructions
  - even reader/writer locking can help
    - readers can share, writers may be attacked in other ways
- you can't deadlock if you have private resources
  - can we give each process its own private resource?

Deadlock, Prevention and Avoidance

16

## Attack #2: hold and block

deadlock requires you to block holding resources

1. allocate all resources in a single operation
  - you hold nothing while blocked
  - when you return, you have all or nothing
2. disallow blocking while holding resources
  - you must release all held locks prior to blocking
  - reacquire them again after you return
3. non-blocking requests
  - a request that can't be satisfied immediately will fail

Deadlock, Prevention and Avoidance

17

## Attack #3: no preemption

- deadlock can be broken by resource confiscation
  - resource "leases" with time-outs and "lock breaking"
  - resource can be seized & reallocated to new client
- revocation must be enforced
  - invalidate previous owner's resource handle
  - if revocation is not possible, kill previous owner
- resources may be damaged by lock breaking
  - previous owner was in the middle of critical section
  - may need mechanisms to audit/repair resource
- resources must be designed for revocation

Deadlock, Prevention and Avoidance

18


### Attack #4: circular dependencies

- total resource ordering
  - all requesters allocate resources in same order
  - first allocate R1 and then R2 afterwards
  - someone else may have R2 but he doesn't need R1
- assumes we know how to order the resources
  - order by resource type (e.g. groups before members)
  - order by relationship (e.g. parents before children)
- may require a lock dance
  - release R2, allocate R1, reacquire R2

Deadlock, Prevention and Avoidance

19

### “Lock Dances” to preserve ordering



list head must be locked for searching, adding & deleting

individual buffers must be locked to perform I/O & other operations

To avoid deadlock, we must always lock the list head before we lock an individual buffer.

To find a desired buffer:

- read lock list head
- search for desired buffer
- lock desired buffer
- unlock list head
- return (locked) buffer

To delete a (locked) buffer from list


- unlock buffer
- write lock list head
- search for desired buffer
- lock desired buffer
- remove from list
- unlock list head

Deadlock, Prevention and Avoidance



### Deadlock – Practical Examples

- the problem – urban gridlock
  - resource: being in the intersection
  - deadlock: nobody can get through



Deadlock, Prevention and Avoidance

21

### Deadlocks: divide and conquer!

- There is no one universal solution to all deadlocks
  - fortunately, we don't need a universal solution
  - we only need a solution for each resource
- Solve each individual problem any way you can
  - make resources sharable wherever possible
  - use reservations for commodity resources
  - ordered locking or no hold-and-block where possible
  - as a last resort, leases and lock breaking
- OS must prevent deadlocks in all system services
  - applications are responsible for their own behavior

Deadlock, Prevention and Avoidance

22

### Closely related forms of "hangs"

- live-lock
  - process is running, but won't free R1 until it gets msg
  - process that will send the message is blocked for R1
- Sleeping Beauty, waiting for “Prince Charming”
  - a process is blocked, awaiting some completion
  - but, for some reason, it will never happen
- neither of these is a true deadlock
  - wouldn't be found by deadlock detection algorithm
  - both leave the system just as hung as a deadlock

Deadlock, Prevention and Avoidance

23

### Deadlock vs. "hang" detection

- deadlock detection seldom makes sense
  - it is extremely complex to implement
  - only detects true deadlocks for known resources
- service/application "health monitoring" does
  - monitor application progress/submit test transactions
  - if response takes too long, declare service "hung"
- health monitoring is easy to implement
- it can detect a wide range of problems
  - deadlocks, live-locks, infinite loops & waits, crashes

Deadlock, Prevention and Avoidance

24

## Hang/Failure Detection Methodology

- look for obvious failures
  - process exits or core dumps
- passive observation to detect hangs
  - is process consuming CPU time, or is it blocked
  - is process doing network and/or disk I/O
- external health monitoring
  - “pings”, null requests, standard test requests
- internal instrumentation
  - white box audits, exercisers, and monitoring

Deadlock, Prevention and Avoidance

25

## Automated Recovery

- kill and restart “all of the affected software”
- how will this affect service/clients
  - design services to automatically fail-over
  - components can warm-start, fall back to last check-point, or cold start
- which, and how many processes to kill?
  - define service failure/recovery zones
  - processes to be started/killed as a group
  - progressive levels of increasingly scope/severity

Deadlock, Prevention and Avoidance

26

## Synchronization is Difficult

- recognizing potential critical sections
  - potential combinations of events
  - interactions with other pieces of code
- choosing the mutual exclusion method
  - there are many different mechanisms
  - with different costs, benefits, weaknesses
- correctly implementing the strategy
  - correct code, in all of the required places
  - maintainers may not understand the rules

Deadlock, Prevention and Avoidance

27

## We need a “Magic Bullet”

- We identify shared resources
  - objects whose methods may require serialization
- We write code to operate on those objects
  - just write the code
  - assume all critical sections will be serialized
- Compiler generates the serialization
  - automatically generated locks and releases
  - using appropriate mechanisms
  - correct code in all required places

Deadlock, Prevention and Avoidance

28

## Monitors – Protected Classes

- each monitor class has a semaphore
  - automatically acquired on method invocation
  - automatically released on method return
  - automatically released/acquired around CV waits
- good encapsulation
  - developers need not identify critical sections
  - clients need not be concerned with locking
  - protection is completely automatic
- high confidence of adequate protection

Deadlock, Prevention and Avoidance

29

## Monitors: use

```
monitor CheckBook {
    // class is locked when any method is invoked
    private int balance;
    public int balance() {
        return(balance);
    }
    public int debit(int amount) {
        balance -= amount;
        return( balance)
    }
}
```

Deadlock, Prevention and Avoidance

30

## Evaluating: Monitors

- correctness
  - complete mutual exclusion is assured
- fairness
  - semaphore queue prevents starvation
- progress
  - inter-class dependencies can cause deadlocks
- performance
  - coarse grained locking is not scalable

Deadlock, Prevention and Avoidance

31

## Java Synchronized Methods

- each object has an associated mutex
  - acquired before calling a synchronized method
  - nested calls (by same thread) do not reacquire
  - automatically released upon final return
- static synchronized methods lock class mutex
- advantages
  - finer lock granularity, reduced deadlock risk
- costs
  - developer must identify serialized methods

Deadlock, Prevention and Avoidance

32

## Java Synchronized: use

```
class CheckBook {
    private int balance;
    public int balance() {
        return(balance);
    }
    // object is locked when this method is invoked
    public synchronized int debit(int amount) {
        balance -= amount;
        return( balance)
    }
}
```

Deadlock, Prevention and Avoidance

33

## Evaluating Java Synchronized Methods

- correctness
  - correct if developer chose the right methods
- fairness
  - priority thread scheduling (potential starvation)
- progress
  - safe from single thread deadlocks
- performance
  - fine grained (per object) locking
  - selecting which methods to synchronize

Deadlock, Prevention and Avoidance

34

## Encapsulated Locking

- opaquely encapsulate implementation details
  - make class easier to use for clients
  - preserve the freedom to change it later
- locking is entirely internal to class
  - search/update races within the methods
  - critical sections involve only class resources
  - critical sections do not span multiple operations
  - no possible interactions with external resources

Deadlock, Prevention and Avoidance

35

## Client Locking

- Class cannot correctly synchronize all uses
- critical section spans multiple class operations
  - updates in a higher level transaction
- client-dependent synchronization needs
  - locking needs depend on how object is used
  - client may control access to protected objects
  - client may select best serialization method
- potential interactions with other resources
  - deadlock prevention must be at higher level

Deadlock, Prevention and Avoidance

36

## Assignments


- for the next lecture:
  - load and stress testing

Deadlock, Prevention and Avoidance

37

## Discussion Slides

### nested monitors – example



Deadlock, Prevention and Avoidance



### (nested monitors – simpler isn't safer)

- consider two monitors:
  - QUEUE with methods: enqueue, dequeue
  - ADAPTOR with methods: process, receive
    - where ADAPTORS are implemented with QUEUES
- possible static deadlocks:
  - QUEUE.enqueue adds entry, calls ADAPTOR.process
  - ADAPTOR.process calls QUEUE.dequeue
- possible dynamic deadlocks:
  - thread 1 calls QUEUE.enqueue, calls ADAPTOR.process
  - thread 2 calls ADAPTOR.receive, calls QUEUE.enqueue

Deadlock, Prevention and Avoidance

40

### Monitors: simplicity vs. performance

- monitor locking is very conservative
  - lock the entire class (not merely a specific object) ●
  - lock for entire duration of any method invocations
- this can create performance problems
  - they eliminate conflicts by eliminating parallelism
  - if a thread blocks in a monitor a convoy can form
- There Ain't No Such Thing As A Free Lunch
  - fine-grained locking is difficult and error prone
  - coarse-grained locking creates bottle-necks

Deadlock, Prevention and Avoidance

41