

CS370 Operating Systems

Colorado State University

Yashwant K Malaiya

Fall 2025 L15

Deadlocks



Slides based on

- Text by Silberschatz, Galvin, Gagne
- Various sources


Methods for Handling Deadlocks

- Ensure that the system will **never** enter a deadlock state:
 - **Deadlock prevention**
 - ensuring that at least one of the 4 conditions cannot hold
 - **Deadlock avoidance**
 - Dynamically examines the resource-allocation state to ensure that there can never be a circular-wait condition
- Allow the system to enter a deadlock state
 - **Detect and then recover.** Hope is that it happens rarely.
- Ignore the problem and pretend that deadlocks never occur in the system; used by most operating systems, including UNIX.


Single Instance of Each Resource Type

- Maintain **wait-for** graph (based on resource allocation graph)
 - Nodes are processes
 - $P_i \rightarrow P_j$ if P_i is waiting for P_j
 - *Deadlock if cycles*
- Periodically invoke an algorithm that searches for a cycle in the graph. If there is a cycle, there exists a deadlock
- An algorithm to detect a cycle in a graph requires an order of n^2 operations, where n is the number of vertices in the graph

Resource-Allocation Graph and Wait-for Graph



Resource-Allocation Graph



Corresponding wait-for graph

Has cycles. Deadlock.

Several Instances of a Resource Type

Banker's algorithm: Can requests by all process be satisfied?

- **Available:** A vector of length m indicates the number of available (currently free) resources of each type
- **Allocation:** An $n \times m$ matrix defines the number of resources of each type currently allocated to each process
- **Request:** An $n \times m$ matrix indicates the current request of each process. If $\text{Request}[i][j] = k$, then process P_i is requesting k more instances of resource type R_j .

Detection Algorithm

1. Let **Work** and **Finish** be vectors of length m and n , respectively.
Initialize:
 - (a) **Work** = initially available
 - (b) For $i = 1, 2, \dots, n$, if **Allocation** $_i \neq 0$, then
Finish[i] = **false**; otherwise, **Finish**[i] = **true**
2. Find an index i such that both:
 - (a) **Finish**[i] == **false**
 - (b) **Request** $_i \leq \text{Work}$

If no such i exists, go to step 4
3. **Work** = **Work** + **Allocation** $_i$,
Finish[i] = **true**
go to step 2 (find next process)
4. If **Finish**[i] == **false**, for some i , $1 \leq i \leq n$, then the system is in deadlock state. Moreover, if **Finish**[i] == **false**, then P_i is deadlocked

Algorithm requires an order of $O(m \times n^2)$ operations to detect whether the system is in deadlocked state


n = number of processes,
m = number of resources types
Work: res currently free
Finish $_i$: processes finished
Allocation $_i$: allocated to i

Example of Detection Algorithm

- Five processes P_0 through P_4 ; three resource types A (7 instances), B (2 instances), and C (6 instances)
- Sequence $\langle P_0, P_2, P_3, P_1, P_4 \rangle$ will result in $Finish[i] = \text{true}$ for all i. **No deadlock**

Process	Allocation			Request		
	A	B	C	A	B	C
type	A	B	C	A	B	C
available	0	0	0			
P0	0	1	0	0	0	0
P1	2	0	0	2	0	2
P2	3	0	3	0	0	0
P3	2	1	1	1	0	0
P4	0	0	2	0	0	2

After	work		
ini	0	0	0
P0	0	1	0
P2	3	1	3
P3	5	2	4
P1	7	2	4
P4	7	2	6




Example of Detection Algorithm (cont)

- P_2 requests an additional instance of type C

Process	Allocation			Request		
	A	B	C	A	B	C
type	A	B	C	A	B	C
available	0	0	0			
P0	0	1	0	0	0	0
P1	2	0	0	2	0	2
P2	3	0	3	0	0	1
P3	2	1	1	1	0	0
P4	0	0	2	0	0	2

Sequence

After	work		
ini	0	0	0
P0	0	1	0
P2	-	-	-



- State of system?
 - Can reclaim resources held by process P_0 , but insufficient resources to fulfill other processes' requests
 - Deadlock exists, consisting of processes P_1 , P_2 , P_3 , and P_4

Detection-Algorithm Usage

- When, and how often, to invoke depends on:
 - How often a deadlock is likely to occur
 - How many processes will need to be rolled back
 - one for each disjoint cycle
- If detection algorithm is invoked arbitrarily, there may be many cycles in the resource graph and so we would not be able to tell which of the many deadlocked processes “caused” the deadlock.

Recovery from Deadlock: Process Termination

Choices

- Abort all deadlocked processes
- Abort one process at a time until the deadlock cycle is eliminated

In which order should we choose to abort?

1. Priority of the process
2. How long process has computed, and how much longer to completion
3. Resources the process has used
4. Resources process needs to complete
5. How many processes will need to be terminated
6. Is process interactive or batch?

- **Selecting a victim** – minimize cost
- **Rollback** – return to some safe state, restart process for that state
- **Starvation** – same process may always be picked as victim, include number of rollbacks in cost factor

- **Checkpoint** process periodically
 - Contains memory image and resource state
- Deadlock detection tells us *which* resources are needed
- Process owning a needed resource
 - **Rolled back** to before it acquired needed resource
 - Work done since rolled back checkpoint discarded
 - **Assign** resource to deadlocked process

Livelocks

In a livelock two processes need each other's resource

- Both run and make no progress, but neither process blocks
- Use CPU quantum over and over without making progress

Ex: If fork fails because process table is full



- Wait for some time and try again
- But there could be a collection of processes each trying to do the same thing
- Avoided by ensuring that only one process (chosen randomly or by priority) takes action

Two people meet in a narrow corridor, and each tries to be polite by moving aside to let the other pass. But they end up swaying from side to side without making any progress because they both repeatedly move the same way at the same time.

Welcome to CS370 Second Half

- Topics: Memory, Storage, File System, Virtualization
- Class rules: See [Syllabus](#)
 - Class, Canvas, Teams
 - participation
 - Final
 - Sec 001, local 801: in class.
 - Sec 801 non-local: on-line.
 - SDC: Sec 001, Sec 801: must be taken at SDC
 - Project, deadlines, Plagiarism


Some OS History Lessons 1: UNIX



- **Unix** 1969 at AT&T's Bell Laboratories by Ken Thompson, Dennis Ritchie, Douglas McIlroy, and Joe Ossanna. Initially released in 1971 and written in assembly language, Unix was re-written in C in 1973 by Dennis Ritchie, making it more portable across platforms.
- **BSD** (Berkeley Software Distribution) is a Unix variant developed at UC Berkeley. Derivatives like FreeBSD, OpenBSD, and NetBSD have emerged from BSD. OS X (macOS) and PS4 also have roots in BSD.
- **Linux**, released by Linus Torvalds in 1991, open-source operating system. Initially built for Intel x86 PCs, Linux has since been ported to more platforms than any other OS. It is now widely used on servers, supercomputers, mobile phones (Android), and gaming consoles like the Nintendo Switch.

Some OS History Lessons 2: Windows

- 1974: CP/M Intel 8080, Gary Kildall, Digital Research
 - 8-bit, min 16 kB RAM, floppy
- 1980: 86-DOS, Intel 8086, Tim Paterson, Seattle Computer Products
 - Inspired by CP/M?
- 1981: PC DOS, Bill Gates, Microsoft
 - 86-DOS licensed for \$25,000, hired Paterson
- 1985: Windows, Bill Gates, Microsoft
 - GUI inspired by MAC? Xerox PARC Star?



Gary Kildall net worth \$1.9 Million at death
Tim Paterson Net Worth: \$250,000

CS370 Operating Systems

Colorado State University

Yashwant K Malaiya

Fall 2025



Main Memory

Slides based on

- Text by Silberschatz, Galvin, Gagne
- Various sources



Colorado State University
Yashwant K Malaiya
Fall 2025

Main Memory

Slides based on
• Text by Silberschatz, Galvin, Gagne
• Various sources

Chapter 8: Main Memory

Objectives:

- Organizing memory for multiprogramming environment
 - Partitioned vs separate address spaces
- Memory-management techniques
 - Virtual vs physical addresses
 - Chunks
 - segmentation
 - Paging: page tables, caching (“TLBs”)
- Examples: the Intel (old/new) and ARM architectures

What we want

- Memory capacities have been increasing
 - But programs are getting bigger faster
 - Parkinson's Law*: Programs expand to fill the memory available to hold
- What we would like
 - Memory that is
 - infinitely large, infinitely fast
 - Non-volatile
 - Inexpensive too
- Unfortunately, no such memory exists as of now

*work expands so as to fill the time available for its completion. 1955

Background


- Program must be brought (from disk) into memory and run as a process
- Main memory and registers are only storage CPU can access directly
- Memory unit only sees a stream of
 - addresses + read requests, or
 - address + data and write requests
- n-bit address: address space of size 2^n bytes.
 - Ex: 32 bits: addresses 0 to $(2^{32}-1)$ bytes
 - Addressable unit is always 1 byte.
- Access times:
 - Register access in one CPU clock (or less)
 - Main memory can take many cycles, causing a **stall**
 - **Cache** sits between main memory and CPU registers making main memory appear much faster
- **Protection** of memory required to ensure correct operation

$$2^{10} = 1,024 \approx \text{K}$$

$$2^{20} = 1,048,576 \approx \text{M}$$

$$2^{30} \approx \text{G}$$

Hierarchy




Ch 11,13,14,16: Disk, file system

Cache: CS470

Memory Technology

somewhat inaccurate




Protection: Making sure each process has separate memory spaces

- OS must be protected from accesses by user processes
- User processes must be protected from one another
 - Determine range of legal addresses for each process
 - Ensure that process can access only those
- Approaches:
 - Partitioning address space (early system)
 - Separate address spaces (modern practice)


Partitioning: Base and Limit Registers

- Base and Limit for a process
 - **Base**: Smallest legal physical address
 - **Limit**: Size of the range of physical address
- A pair of **base** and **limit registers** define the logical address space for a process
- CPU must check every memory access generated in user mode to be sure it is between base and limit for that user
- Base: **Smallest** legal physical address
- Limit: Size of the **range** of physical address
- Eg: Base = 300040 and limit = 120900
- Legal: 300040 to $(300040 + 120900 - 1) = 420939$




Addresses: decimal, hex/binary

Hardware Address Protection



Legal addresses: **Base address to Base address + limit -1**

Multistep Processing of a User Program



Address Binding Questions

- Programs on disk, ready to be brought into memory to execute from an **input queue**
 - Without support, must be loaded into address 0000
- Inconvenient to have first user process physical address always at 0000
 - How can it not be?
- Addresses represented in different ways at different stages of a program's life
 - **Source code** addresses are symbolic
 - **Compiled code** addresses **bind** to relocatable addresses
 - i.e., “14 bytes from beginning of this module”
 - **Linker or loader** will bind relocatable addresses to absolute addresses
 - i.e., 74014
 - Each binding maps one address space to another

Binding of Instructions and Data to Memory

- Address binding of instructions and data to memory addresses can happen at three different stages
 - **Compile time:** If memory location known a priori, **absolute code** can be generated; must recompile code if starting location changes
 - **Load time:** Must generate **relocatable code** if memory location is not known at compile time
 - **Execution time:** Binding delayed until run time if the process can be moved during its execution from one memory segment to another
 - Need hardware support for address maps (e.g., base and limit registers)

Separate Address Spaces

Modern

- Each process has its own private address space.
 - **Logical address space** is the set of all logical addresses used by a process.
- However, the physical memory has just one address space.
 - **Physical address space** is the set of all physical addresses
- Need to map one to the other.


Logical vs. Physical Address Space

- The concept of a logical address space that is bound to a separate **physical address space** is central to proper memory management
 - **Logical address** – generated by the CPU; also referred to as **virtual address**
 - **Physical address** – address seen by the memory unit
- **Logical address space** is the set of all logical addresses generated by a program
- **Physical address space** is the set of all physical addresses

Memory-Management Unit (MMU)

- Hardware device that at run time maps virtual to physical address
 - Many methods possible, we will see them soon
- Consider simple scheme where the value in the relocation register is added to every address generated by a user process at the time it is sent to memory
 - Base register now called **relocation register**
 - MS-DOS on Intel 80x86 used 4 relocation registers
- The user program deals with *logical addresses*; it never sees the *real physical addresses*
 - Execution-time binding occurs when reference is made to location in memory
 - Logical address bound to physical addresses

Dynamic relocation using a relocation register



Linking: Static vs Dynamic

- **Linking**
 - Takes some smaller executables and joins them together as a single larger executable.
- **Static linking** – system libraries and program code combined by the loader into the binary image
 - Every program includes library: wastes memory
- **Dynamic linking** –linking postponed until execution time
 - Operating system locates and links the routine at run time

Dynamic Linking

- **Dynamic linking** –linking postponed until execution time
- Small piece of code, **stub**, used to locate the appropriate memory-resident library routine
- Stub replaces itself with the address of the routine, and executes the routine
- Operating system checks if routine is in processes' memory address
 - If not in address space, add to address space
- Dynamic linking is particularly useful for
 - **shared libraries**


Dynamic loading of routines

- Routine is not loaded until it is called
- Better memory-space utilization; unused routine is never loaded
- All routines kept on disk in relocatable load format
- Useful when large amounts of code are needed to handle infrequently occurring cases
- OS can help by providing libraries to implement dynamic loading
- Static library
 - Linux .a (archive)
 - Windows .lib (Library)
- Dynamic Library
 - Linux .so (Shared object)
 - Windows .dll (Dynamic link library)

Swapping a process

- A process can be **swapped** temporarily out of memory to a backing store, and then brought back into memory for continued execution
 - Total physical memory space of processes can exceed physical memory
- **Backing store** – fast disk large enough to accommodate copies of all memory images for all users; must provide direct access to these memory images
- Major part of swap time is transfer time; total transfer time is directly proportional to the amount of memory swapped
- System maintains a **ready queue** of ready-to-run processes which have memory images on disk

Schematic View of Swapping



Do we really need to keep the entire process in the main memory? Stay tuned.

Context Switch Time including Swapping

- If next processes to be put on CPU is not in memory, need to swap out a process and swap in target process
- Context switch time can then be very high
- 100MB process swapping to hard disk with transfer rate of 50MB/sec
 - Swap out time of $100\text{MB}/50\text{MB/s} = 2\text{ seconds}$
 - Plus swap in of same sized process
 - Total context switch swapping component time of 4 seconds + some latency
- Can reduce if reduce size of memory swapped – by knowing how much memory really being used by a process

Context Switch Time and Swapping (Cont.)

- Standard swapping not used in modern operating systems
 - But modified version common
 - Swap only when free memory extremely low

Memory Allocation

Memory Allocation Approaches

- **Contiguous allocation**: entire memory for a program in a single contiguous memory block. Find where a program will “fit”. earliest approach
- **Segmentation**: program divided into logically divided “segments” such as main program, functions, stack etc.
 - Need table to track segments.
- **Paging**: program divided into fixed size “pages”, each placed in a fixed size “frame”.
 - Need table to track pages.


Contiguous Allocation

- Main memory must support both OS and user processes
- Limited resource, must allocate efficiently
- Contiguous allocation is one **early** method
- Main memory usually into two **partitions**:
 - Resident operating system, usually held in low memory with interrupt vectors
 - User processes then held in high memory
 - Each process contained in **single contiguous section of memory**

Contiguous Allocation (Cont.)

- Registers used to protect user processes from each other, and from changing operating-system code and data
 - Relocation (Base) register contains value of smallest physical address
 - Limit register contains range of logical addresses – each logical address must be less than the limit register
- MMU maps logical address *dynamically*


Hardware Support for Relocation and Limit Registers



MMU maps logical address *dynamically*
 $\text{Physical address} = \text{relocation reg} + \text{valid logical address}$

Multiple-partition allocation

- Multiple-partition allocation
 - Degree of multiprogramming limited by number of partitions
 - **Variable-partition** sizes for efficiency (sized to a given process' needs)
 - **Hole** – block of available memory; holes of various size are scattered throughout memory
 - When a process arrives, it is allocated memory from a hole large enough to accommodate it
 - Process exiting frees its partition, adjacent free partitions combined
 - Operating system maintains information about:
 - allocated partitions
 - free partitions (hole)



Dynamic Storage-Allocation Problem

How to satisfy a request of size n from a list of free holes?

- **First-fit**: Allocate the *first* hole that is big enough
- **Best-fit**: Allocate the *smallest* hole that is big enough; must search entire list, unless ordered by size
 - Produces the smallest leftover hole
- **Worst-fit**: Allocate the *largest* hole; must also search entire list
 - Produces the largest leftover hole

Simulation studies:

- First-fit and best-fit better than worst-fit in terms of speed and storage utilization
- Best fit is **slower** than first fit . Surprisingly, it also results in more **wasted memory** than first fit
 - Tends to fill up memory with tiny, useless holes

Fragmentation

- **External Fragmentation** – External fragmentation: memory wasted due to small chunks of free memory interspersed among allocated regions
- **Internal Fragmentation** – allocated memory may be slightly larger than requested memory; this size difference is memory internal to a partition, but not being used
- Simulation analysis reveals that given N blocks allocated, $0.5 N$ blocks lost to fragmentation
 - 1/3 may be unusable -> **50-percent rule**

Fragmentation (Cont.)

- Reduce external fragmentation by **compaction**
 - Shuffle memory contents to place all free memory together in one large block
 - Compaction is possible *only* if relocation is dynamic, and is done at execution time
 - I/O problem
 - Latch job in memory while it is involved in I/O
 - Do I/O only into OS buffers