

Memory Management

CS 111

Operating Systems

Peter Reiher

Outline

- What is memory management about?
- Memory management strategies:
 - Fixed partition strategies
 - Dynamic domains
 - Buffer pools
 - Garbage collection
 - Memory compaction

Memory Management

- Memory is one of the key assets used in computing
- In particular, memory abstractions that are usable from a running program
 - Which, in modern machines, typically means RAM
- We have a limited amount of it
- Lots of processes want to use it
- How do we manage its use?

What Is Memory Used For?

- Anything that a program needs to access
 - Except control and temporary values, which are kept in registers
- The code
 - To allow the process to execute instructions
- The stack
 - To keep track of its state of execution
- The heap
 - To hold dynamically allocated variables

Other Uses of Memory

- The operating system needs memory itself
- For its own code, stack, and dynamic allocations
- For I/O buffers
- To hold per-process control data
- The OS shares the same physical memory that user processes rely on
- The OS provides overall memory management

Aspects of the Memory Management Problem

- Most processes can't perfectly predict how much memory they will use
- The processes expect to find their existing data when they need it where they left it
- The entire amount of data required by all processes may exceed amount of available physical memory
- Switching between processes must be fast
 - Can't afford much delay for copying data
- The cost of memory management itself must not be too high

Memory Management Strategies

- Domains and fixed partition allocations
- Dynamic domains
- Paging
- Virtual memory
- We'll talk about the last two in the next class

Fixed Partition Allocation

- Pre-allocate partitions for n processes
 - Usually one partition per process
 - So n partitions
 - Reserving space for largest possible process
- Partitions come in one or a few set sizes
- Very easy to implement
 - Common in old batch processing systems
 - Allocation/deallocation very cheap and easy
- Well suited to well-known job mix

Memory Protection and Fixed Partitions

- Need to enforce partition boundaries
 - To prevent one process from accessing another's memory
- Could use hardware similar to domain registers for this purpose
- On the flip side, hard to arrange for shared memory
 - Especially if only one segment per process

Problems With Fixed Partition Allocation

- Presumes you know how much memory will be used ahead of time
- Limits the number of processes supported to the total of their memory requirements
- Not great for sharing memory
- *Fragmentation* causes inefficient memory use

Fragmentation

- A problem for all memory management systems
 - Fixed partitions suffer it especially badly
- Based on processes not using all the memory they requested
- As a result, you can't provide memory for as many processes as you theoretically could

Fragmentation Example

Let's say there are three processes, A, B, and C

Their memory requirements: Available partition sizes:

A: 6 MBytes


8 Mbytes

B: 3 MBytes


4 Mbytes

C: 2 MBytes


4 Mbytes



Partition 1
8MB



Partition 2
4MB



Partition 3
4MB

$$\text{Total waste} = 2\text{MB} + 1\text{MB} + 2\text{MB} = 5/16\text{MB} = 31\%$$

Internal Fragmentation

- Fragmentation comes in two kinds:
 - Internal and external
- This is an example of *internal fragmentation*
 - We'll see external fragmentation later
- Wasted space in fixed sized blocks
 - The requestor was given more than he needed
 - The unused part is wasted and can't be used for others
- Internal fragmentation can occur whenever you force allocation in fixed-sized chunks

More on Internal Fragmentation

- Internal fragmentation is caused by a mismatch between
 - The chosen sizes of a fixed-sized blocks
 - The actual sizes that programs use
- Average waste: 50% of each block
- Overall waste reduced by multiple sizes
 - Suppose blocks come in sizes S1 and S2
 - Average waste = $((S1/2) + (S2 - S1)/2)/2$

Multiple Fixed Partitions

- You could allow processes to request multiple partitions
 - Of a single or a few sizes
- Doesn't really help the fragmentation problem
 - Now there are more segments to fragment
 - Even if each contained less memory

Where Was Fixed Partition Allocation Used?

- Old operating systems from the 1960s
 - E.g., IBM's OS 360 and MVT
- Not required until people wanted to do multiprogramming
- Not really great even for those environments, so it didn't last
- A simple model for very basic multiprogramming

Summary of Fixed Partition Allocation

- Very simple
- Inflexible
- Subject to a lot of internal fragmentation
- Not used in many modern systems
 - But a possible option for special purpose systems, like embedded systems
 - Where we know exactly what our memory needs will be


Dynamic Domain Allocation

- A concept covered in a previous lecture
 - We'll just review it here
- Domains are regions of memory made available to a process
 - Variable sized, usually any size requested
 - Each domain is contiguous in memory addresses
 - Domains have access permissions for the process
 - Potentially shared between processes
- Each process could have multiple domains
 - With different sizes and characteristics

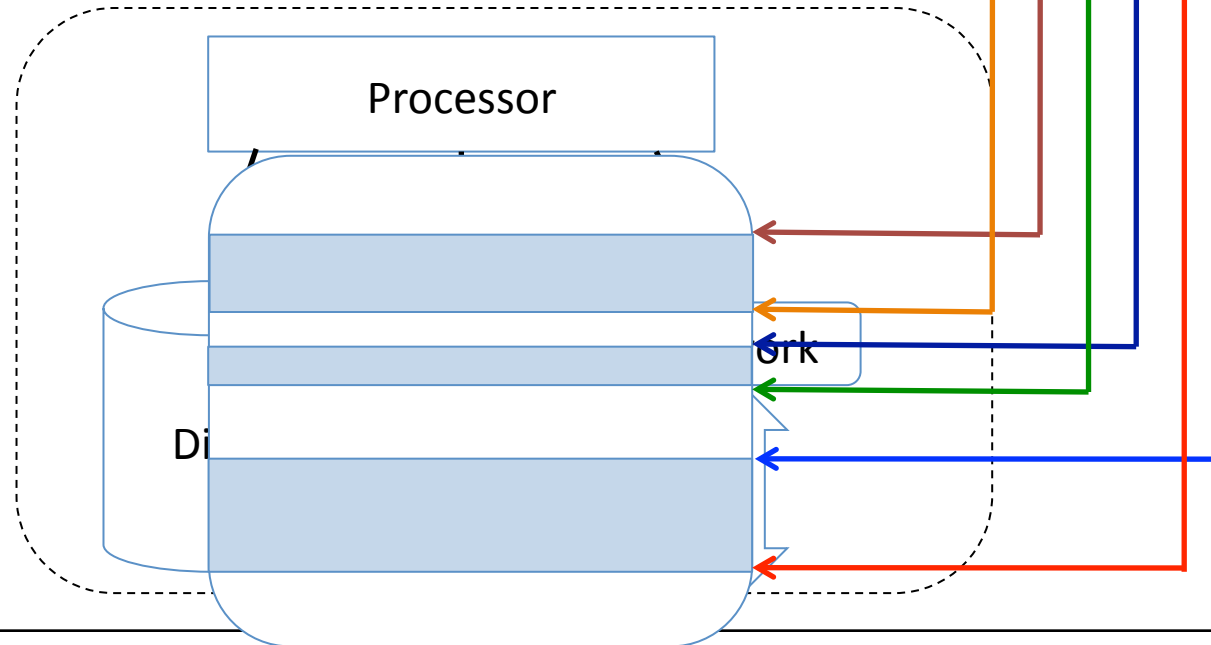

Accessing Domains

- The process issues a memory address
- The address is checked against the domain registers specifying the process' access
 - If address is in one of the process' domains with proper access permissions, allow access
 - Otherwise don't
 - Failures due to access permission problems are permission exceptions
 - Failures due to requesting an address not in your domain are illegal address exceptions

The Domain Concept



Domain Registers



Problems With Domains

- Not relocatable
 - Once a process has a domain, you can't easily move its contents elsewhere
- Not easily expandable
- Impossible to support applications with larger address spaces than physical memory
 - Also can't support several applications whose total needs are greater than physical memory
- Also subject to fragmentation


Relocation and Expansion

- Domains are tied to particular address ranges
 - At least during an execution
- Can't just move the contents of a domain to another set of addresses
 - All the pointers in the contents will be wrong
 - And generally you don't know which memory locations contain pointers
- Hard to expand because there may not be space “nearby”

The Expansion Problem

- Domains are allocated on request
- Processes may ask for new ones later
- But domains that have been given are fixed
 - Can't be moved somewhere else in memory
- Memory management system might have allocated all the space after a given domain
- In which case, it can't be expanded

Illustrating the Problem



Now Process B wants to expand its domain size

But if we do that, Process B steps on Process C's memory

We can't move C's domain out of the way

And we can't move B's domain to a free area

We're stuck, and must deny an expansion request that we have enough memory to handle

Address Spaces Bigger Than Physical Memory

- If a process needs that much memory, how could you possibly support it?
- Two possibilities:
 1. It's not going to use all the memory it's asked for, or at least not all simultaneously
 2. Maybe we can use something other than physical memory to store some of it
- Domains are not friendly to either option

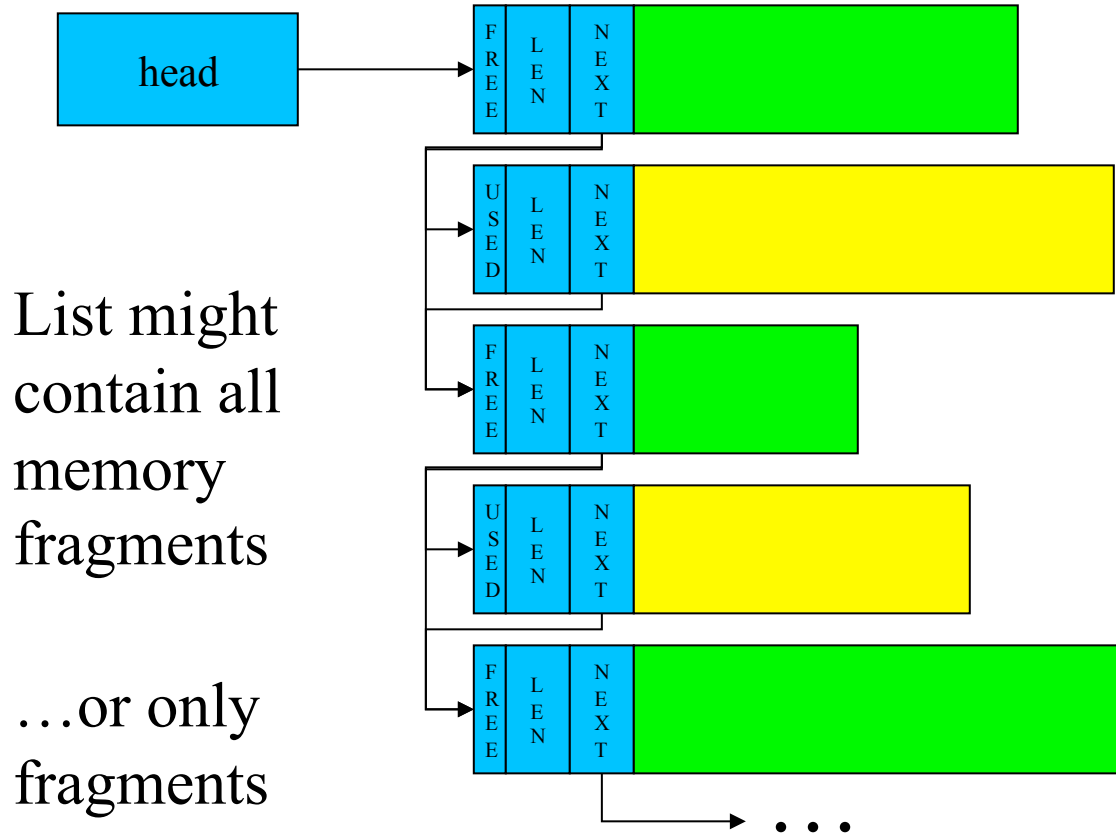
How To Keep Track of Variable Sized Domains?

- Start with one large “heap” of memory
- Maintain a *free list*
 - Systems data structure to keep track of pieces of unallocated memory
- When a process requests more memory:
 - Find a large enough chunk of memory
 - Carve off a piece of the requested size
 - Put the remainder back on a *free list*
- When a process frees memory
 - Put it back on the free list

Managing the Free List

- Fixed sized blocks are easy to track
 - A bit map indicating which blocks are free
- Variable chunks require more information
 - A linked list of descriptors, one per chunk
 - Each descriptor lists the size of the chunk and whether it is free
 - Each has a pointer to the next chunk on list
 - Descriptors often kept at front of each chunk
- Allocated memory may have descriptors too

The Free List



List might contain all memory fragments


...or only fragments that are free

Free Chunk Carving


1. Find a large enough free chunk



2. Reduce its len to requested size





3. Create a new header for residual chunk



4. Insert the new chunk into the list




5. Mark the carved piece as in use



Variable Domain and Fragmentation

- Variable sized domains not as subject to internal fragmentation
 - Unless requestor asked for more than he will use
 - Which is actually pretty common
 - But at least memory manager gave him no more than he requested
- Unlike fixed sized partitions, though, subject to another kind of fragmentation
 - *External fragmentation*

External Fragmentation



We gradually build up small, unusable memory chunks scattered through memory

External Fragmentation: Causes and Effects

- Each allocation creates left-over chunks
 - Over time they become smaller and smaller
- The small left-over fragments are useless
 - They are too small to satisfy any request
 - A second form of fragmentation waste
- Solutions:
 - Try not to create tiny fragments
 - Try to recombine fragments into big chunks

How To Avoid Creating Small Fragments?

- Be smart about which free chunk of memory you use to satisfy a request
- But being smart costs time
- Some choices:
 - Best fit
 - Worst fit
 - First fit
 - Next fit

Best Fit

- Search for the “best fit” chunk
 - Smallest size greater than or equal to requested size
- Advantages:
 - Might find a perfect fit
- Disadvantages:
 - Have to search entire list every time
 - Quickly creates very small fragments

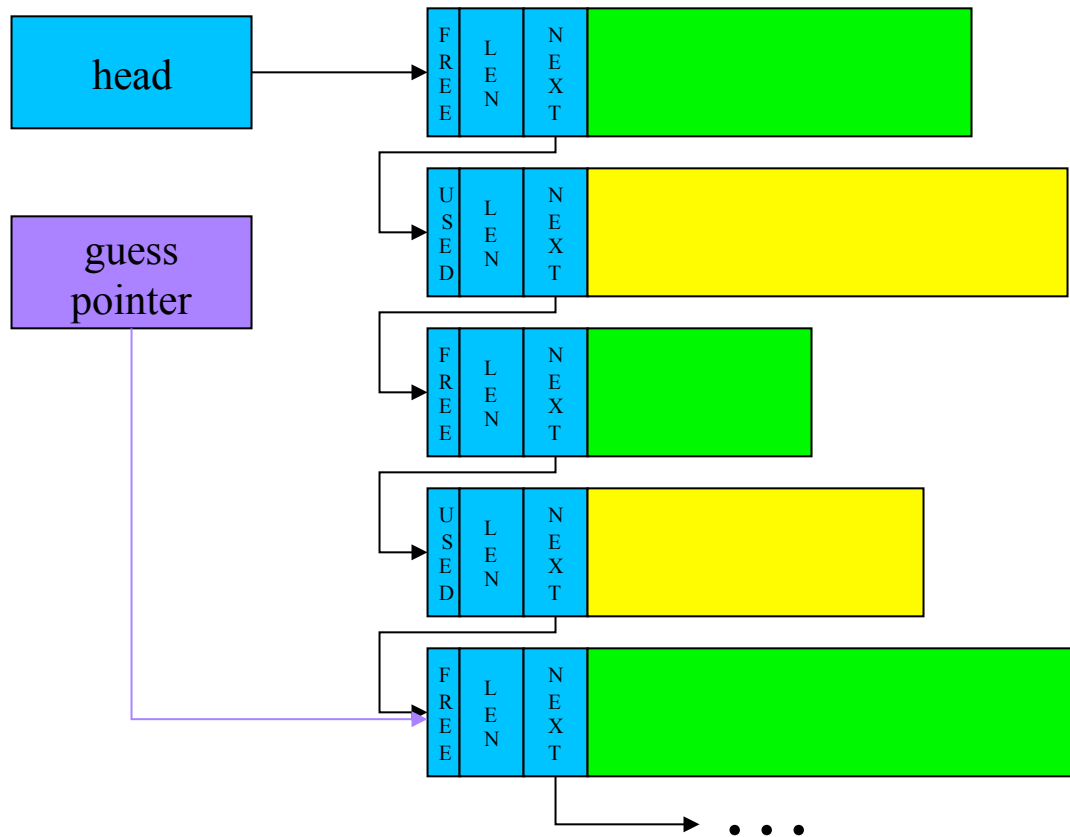
Worst Fit

- Search for the “worst fit” chunk
 - Largest size greater than or equal to requested size
- Advantages:
 - Tends to create very large fragments
... for a while at least
- Disadvantages:
 - Still have to search entire list every time

First Fit

- Take first chunk you find that is big enough
- Advantages:
 - Very short searches
 - Creates random sized fragments
- Disadvantages:
 - The first chunks quickly fragment
 - Searches become longer
 - Ultimately it fragments as badly as best fit

Next Fit



After each search, set guess pointer to chunk after the one we chose.

That is the point at which we will begin our next search.

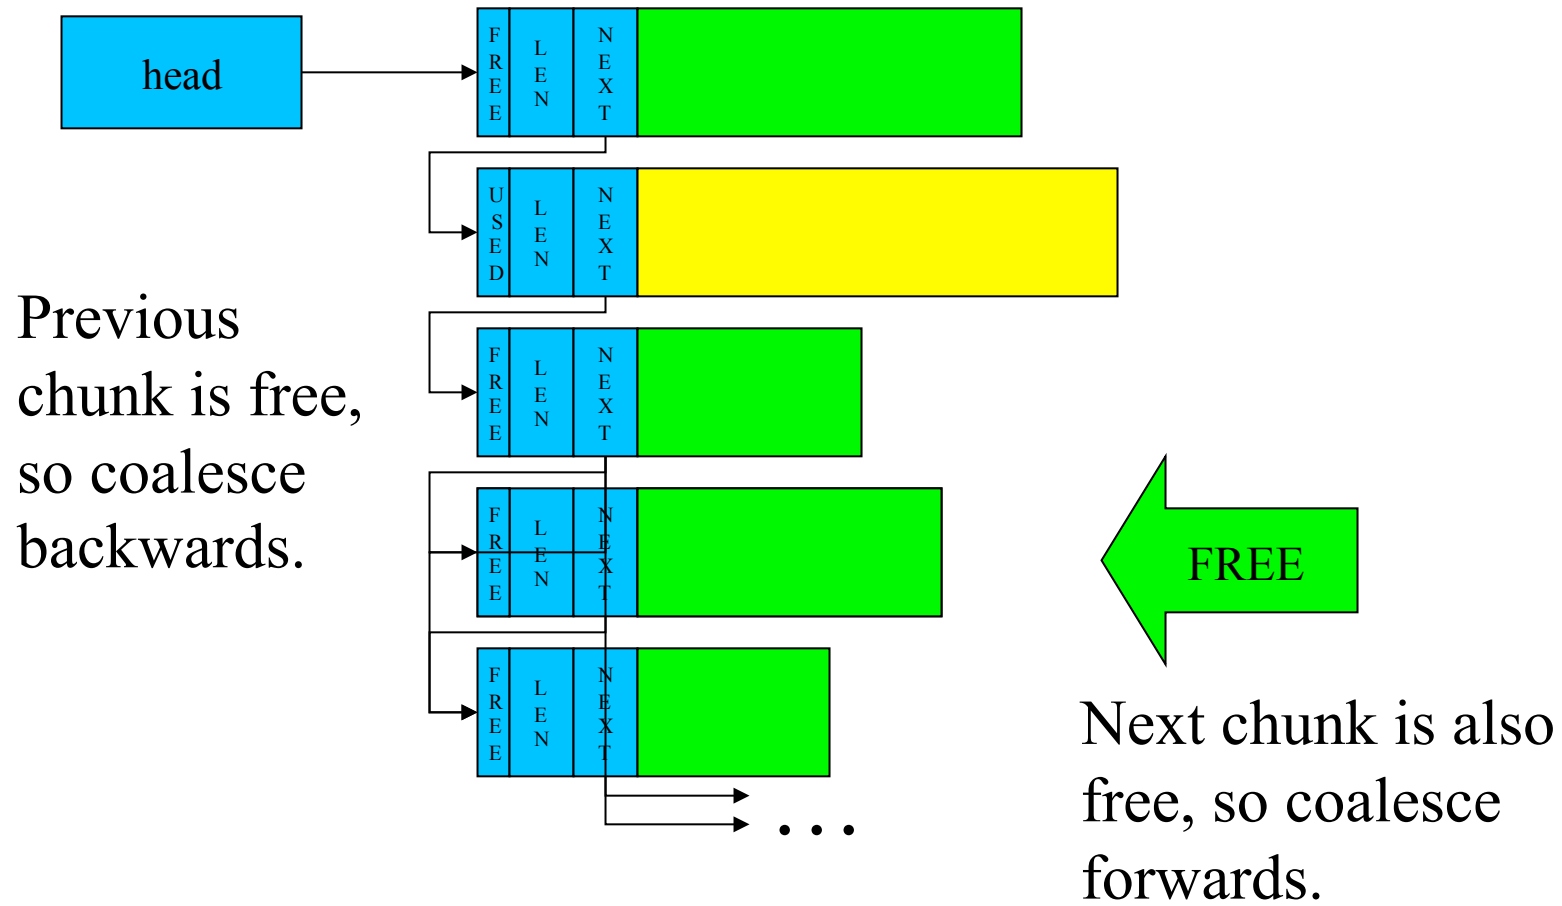
Next Fit Properties

- Tries to get advantages of both first and worst fit
 - Short searches (maybe shorter than first fit)
 - Spreads out fragmentation (like worst fit)
- Guess pointers are a general technique
 - Think of them as a lazy (non-coherent) cache
 - If they are right, they save a lot of time
 - If they are wrong, the algorithm still works
 - They can be used in a wide range of problems

Coalescing Domains

- All variable sized domain allocation algorithms have external fragmentation
 - Some get it faster, some spread it out
- We need a way to reassemble fragments
 - Check neighbors whenever a chunk is freed
 - Recombine free neighbors whenever possible
 - Free list can be designed to make this easier
 - E.g., where are the neighbors of this chunk?
- Counters forces of external fragmentation

Free Chunk Coalescing



Fragmentation and Coalescing

- Opposing processes that operate in parallel
 - Which of the two processes will dominate?
- What fraction of space is typically allocated?
 - Coalescing works better with more free space
- How fast is allocated memory turned over?
 - Chunks held for long time cannot be coalesced
- How variable are requested chunk sizes?
 - High variability increases fragmentation rate
- How long will the program execute?
 - Fragmentation, like rust, gets worse with time

Coalescing and Free List Implementation

- To coalesce, we must know whether the previous and next chunks are also free
- If the neighbors are guaranteed to be in the free list, we can look at them and see if they are free
- If allocated chunks are not in the free list, we must look at the free chunks before and after us
 - And see if they are our contiguous neighbors
 - This suggests that the free list must be maintained in address order

Variable Sized Domain Summary

- Eliminates internal fragmentation
 - Each chunk is custom-made for requestor
- Implementation is more expensive
 - Long searches of complex free lists
 - Carving and coalescing
- External fragmentation is inevitable
 - Coalescing can counteract the fragmentation
- Must we choose the lesser of two evils?

Another Option


- Fixed partition allocations result in internal fragmentation
 - Processes don't use all of the fixed partition
- Dynamic domain allocations result in external fragmentation
 - The elements on the memory free list get smaller and less useful
- Can we strike a balance in between?

A Special Case for Fixed Allocations

frequency

Internal fragmentation results from mismatches between chunk sizes and request sizes (which we have assumed to be randomly distributed)

But if we look at what actually happens, it turns out that memory allocation requests aren't random at all.



Why Aren't Memory Request Sizes Randomly Distributed?

- In real systems, some sizes are requested much more often than others
- Many key services use fixed-size buffers
 - File systems (for disk I/O)
 - Network protocols (for packet assembly)
 - Standard request descriptors
- These account for much transient use
 - They are continuously allocated and freed
- OS might want to handle them specially

Buffer Pools

- If there are popular sizes,
 - Reserve special pools of fixed size buffers
 - Satisfy matching requests from those pools
- **Benefit: improved efficiency**
 - Much simpler than variable domain allocation
 - Eliminates searching, carving, coalescing
 - Reduces (or eliminates) external fragmentation
- **But we must know how much to reserve**
 - Too little, and the buffer pool will become a bottleneck
 - Too much, and we will have a lot of unused buffer space
- **Only satisfy perfectly matching requests**
 - Otherwise, back to internal fragmentation

How Are Buffer Pools Used?

- Process requests a piece of memory for a special purpose
 - E.g., to send a message
- System supplies one element from buffer pool
- Process uses it, completes, frees memory
 - Maybe explicitly
 - Maybe implicitly, based on how such buffers are used
 - E.g., sending the message will free the buffer “behind the process’ back” once the message is gone

Dynamically Sizing Buffer Pools

- If we run low on fixed sized buffers
 - Get more memory from the free list
 - Carve it up into more fixed sized buffers
- If our free buffer list gets too large
 - Return some buffers to the free list
- If the free list gets dangerously low
 - Ask each major service with a buffer pool to return space
- This can be tuned by a few parameters:
 - Low space (need more) threshold
 - High space (have too much) threshold
 - Nominal allocation (what we free down to)
- Resulting system is highly adaptive to changing loads

Lost Memory

- One problem with buffer pools is memory leaks
 - The process is done with the memory
 - But doesn't free it
- Also a problem when a process manages its own memory space
 - E.g., it allocates a big area and maintains its own free list
- Long running processes with memory leaks can waste huge amounts of memory

Garbage Collection

- One solution to memory leaks
- Don't count on processes to release memory
- Monitor how much free memory we've got
- When we run low, start garbage collection
 - Search data space finding every object pointer
 - Note address/size of all accessible objects
 - Compute the compliment (what is inaccessible)
 - Add all inaccessible memory to the free list

How Do We Find All Accessible Memory?

- Object oriented languages often enable this
 - All object references are tagged
 - All object descriptors include size information
- It is often possible for system resources
 - Where all possible references are known
 - (E.g., we know who has which files open)
- How about for the general case?

General Garbage Collection

- Well, what would you need to do?
- Find all the pointers in allocated memory
- Determine “how much” each points to
- Determine what is and is not still pointed to
- Free what isn’t pointed to
- Why might that be difficult?



Problems With General Garbage Collection

- A location in the data or stack segments might seem to contain addresses, but ...
 - Are they truly pointers, or might they be other data types whose values happen to resemble addresses?
 - If pointers, are they themselves still accessible?
 - We might be able to infer this (recursively) for pointers in dynamically allocated structures ...
 - But what about pointers in statically allocated (potentially global) areas?
- And how much is “pointed to,” one word or a million?

Compaction and Relocation

- Garbage collection is just another way to free memory
 - Doesn't greatly help or hurt fragmentation
- Ongoing activity can starve coalescing
 - Chunks reallocated before neighbors become free
- We could stop accepting new allocations
 - But resulting convoy on memory manager would trash throughput
- We need a way to rearrange active memory
 - Re-pack all processes in one end of memory
 - Create one big chunk of free space at other end

Memory Compaction



Now let's compact!

An obvious improvement!

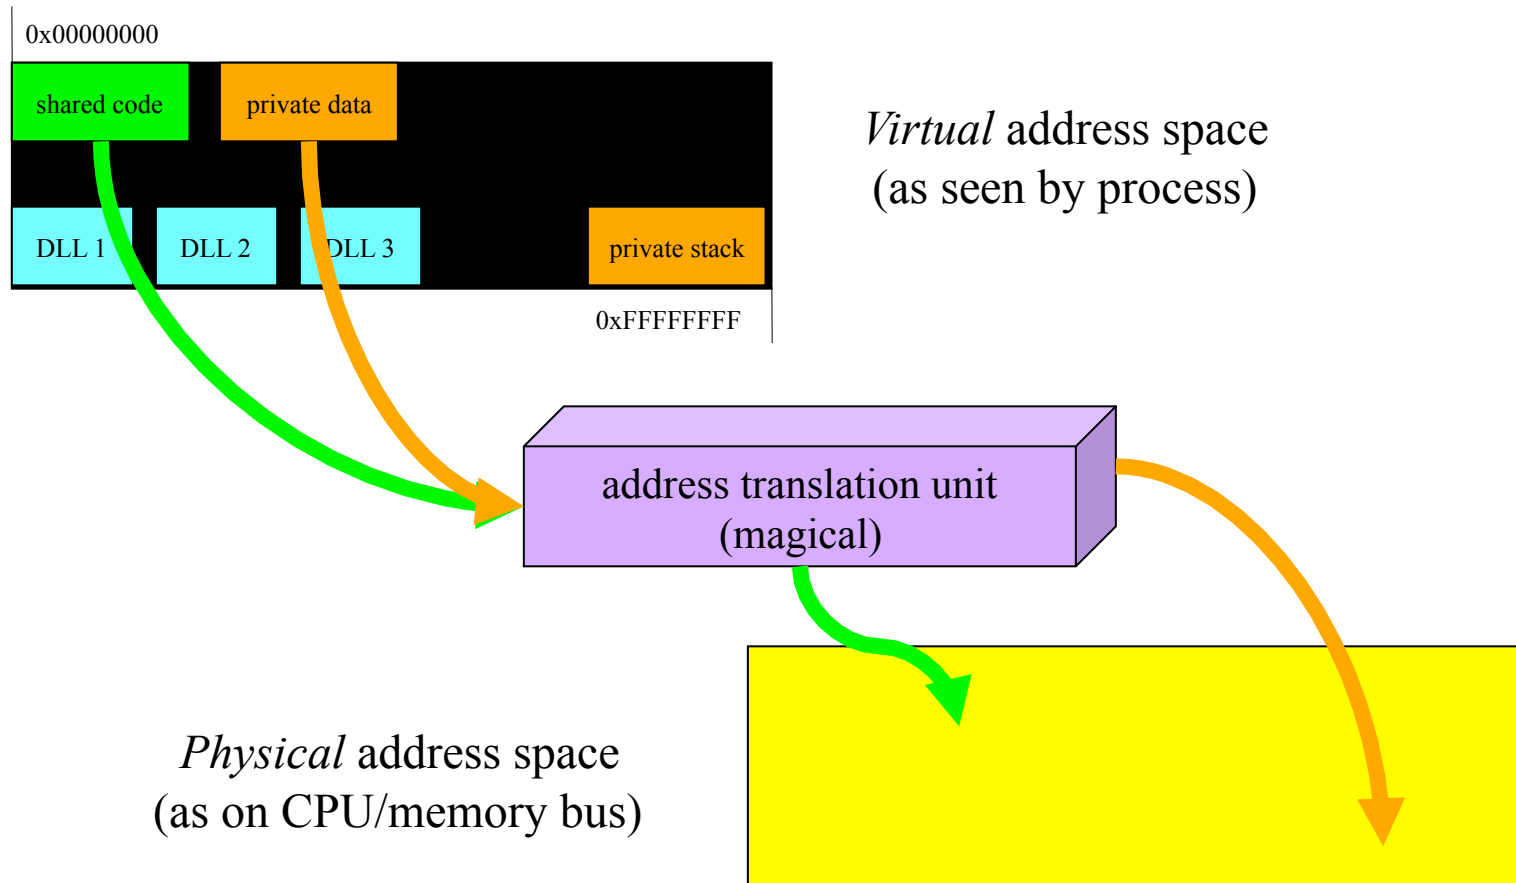
All This Requires Is Relocation . . .

- The ability to move a process
 - From region where it was initially loaded
 - Into a new and different region of memory
- What's so hard about that?
- All addresses in the program will be wrong
 - References in the code segment
 - Calls and branches to other parts of the code
 - References to variables in the data segment
 - Plus new pointers created during execution
 - That point into data and stack segments

The Relocation Problem

- It is not generally feasible to re-relocate a process
 - Maybe we could relocate references to code
 - If we kept the relocation information around
 - But how can we relocate references to data?
 - Pointer values may have been changed
 - New pointers may have been created
- We could never find/fix all address references
 - Like the general case of garbage collection
- Can we make processes location independent?

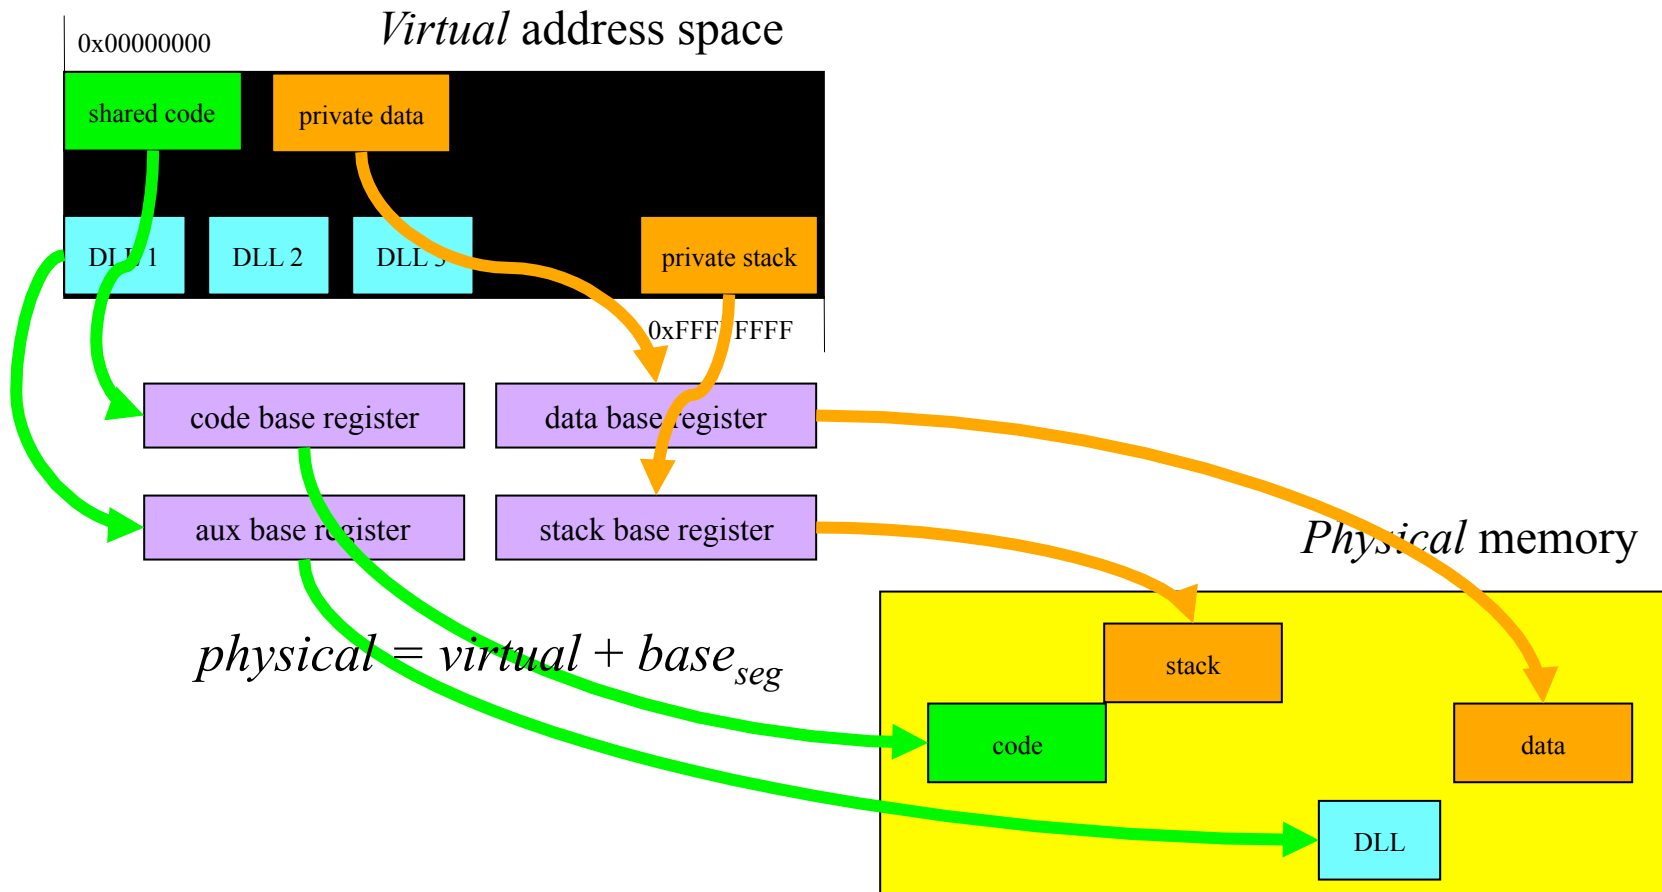
Virtual Address Spaces



Memory Segment Relocation

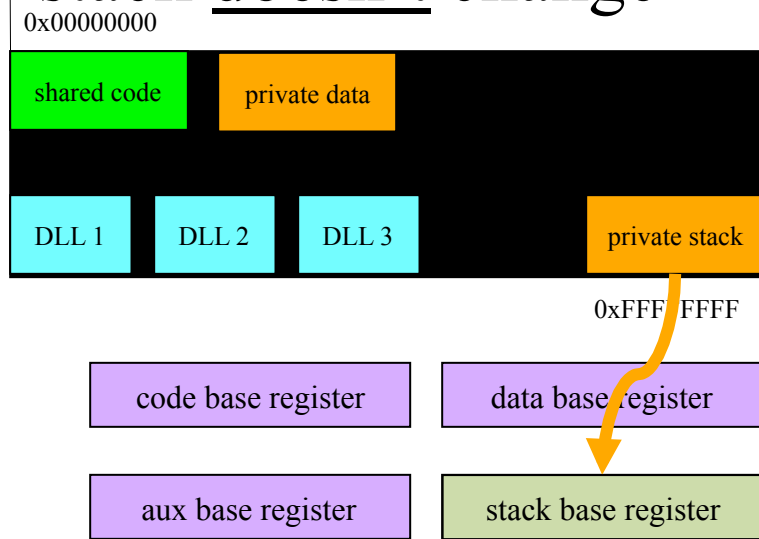
- A natural model
 - Process address space is made up of multiple segments
 - Use the segment as the unit of relocation
 - Long tradition, from the IBM system 360 to Intel x86 architecture
- Computer has special relocation registers
 - They are called segment base registers
 - They point to the start (in physical memory) of each segment
 - CPU automatically adds base register to every address
- OS uses these to perform virtual address translation
 - Set base register to start of region where program is loaded
 - If program is moved, reset base registers to new location
- Program works no matter where its segments are loaded

How Does Segment Relocation Work?



Relocating a Segment

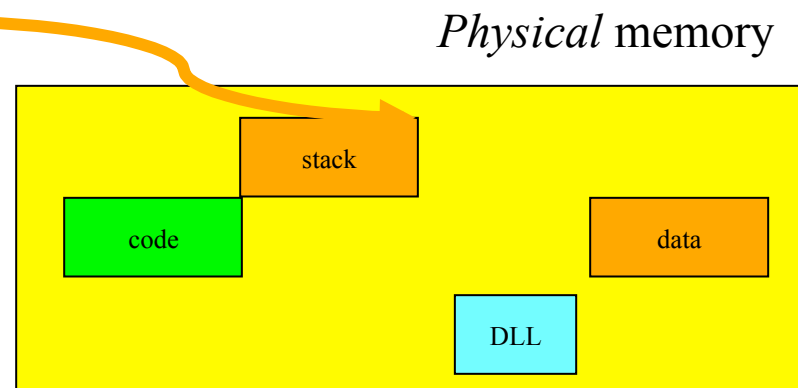
The virtual address of the stack doesn't change



Let's say we need to move the stack in physical memory

$$physical = virtual + base_{seg}$$

We just change the value in the stack base register



Relocation and Safety

- A relocation mechanism (like base registers) is good
 - It solves the relocation problem
 - Enables us to move process segments in physical memory
 - Such relocation turns out to be insufficient
- We also need protection
 - Prevent process from reaching outside its allocated memory
 - E.g., by overrunning the end of a mapped segment
- Segments also need a length (or limit) register
 - Specifies maximum legal offset (from start of segment)
 - Any address greater than this is illegal (in the hole)
 - CPU should report it via a segmentation exception (trap)

How Much of Our Problem Does Relocation Solve?

- We can use variable sized domains
 - Cutting down on internal fragmentation
- We can move domains around
 - Which helps coalescing be more effective
 - But still requires contiguous chunks of data for segments
 - So external fragmentation is still a problem
- We need to get rid of the requirement of contiguous segments

Overlays

- Another problem not yet addressed is limiting a process' memory to the amount of RAM
- Even relocatable segments doesn't solve that
 - Since it just moves them from place to place
- One solution is overlays
 - Define parts of the address space not needed simultaneously
 - Have such overlaid segments “share” a portion of the physical address space

Using Overlays

- Typically defined at the program level
 - By program authors
- For example, consider a two pass assembler
 - Pass 1 produces a symbol table
 - Pass 2 generates machine code
 - The instructions for the two passes need not be in memory simultaneously
 - So we could
 - Run pass 1
 - Load the code for pass 2 on top of pass 1's code
 - Run pass 2

Making Overlays Work

- Need a special loader that is overlay aware
- Explicit instructions are used to invoke the loader
 - Which knows which segments can be used as overlays
 - Also need to fiddle with linking
- Can be done without OS support
- But very complicated and not widely used today

But Overlays Suggest Something

- Overlays are bad because they require users to think about what memory segments are disjoint
- Why force the users to think about it?
- Why not have the system do it for them?
- Could be done statically (e.g., program analysis), but why not dynamically?
- Why not have the system determine at run time what needs to be in memory?