

Today

- Dynamic Memory Allocation (Ch 9.9)
 - Tracking Free Blocks
 - Implicit Lists
 - Explicit Lists
 - Segregated Lists
- Exceptional Control Flow
- Exceptions

Administrative Details

- Read CSAPP 9.9, 8.1-8.2
- Quiz open today at 2:30 until Friday at 2:30pm
- Lab #6 due Friday at 5pm
- Review session on Sun or Tues?
 - https://www.when2meet.com/?27868617-GeNrF
- Final exam
 - Wednesday, December 11, 9:30am 11:30am (2 hours)
 - Clark Hall 105
- Colloquium talk on Friday at 2:35pm in Wege
 - "Systems research to address societal problems"
 - Aruna Balasubramanian, Stony Brook University

2

Optimization: No Boundary Tag for Allocated Blocks

- Boundary tag is only needed for free blocks
- Insight: when sizes are multiples of 4 or more, have 2+ spare bits











- Placement policy:
 - First-fit, next-fit, best-fit, etc.
 - Trades off lower throughput for less fragmentation
 - Interesting observation: segregated free lists approximate a best fit placement policy without having to search entire free list
- Splitting policy:
 - When do we go ahead and split free blocks?
 - How much internal fragmentation are we willing to tolerate?
- Coalescing policy:
 - Immediate coalescing: coalesce each time free is called
 - Deferred coalescing: try to improve performance of free by deferring coalescing until needed. Examples:
 - Coalesce as you scan the free list for malloc
 - Coalesce when the amount of external fragmentation reaches some threshold

9

Practice on Your Own

Consider a system where

- The memory is byte addressable.
- Memory accesses are to 1-byte words
- Virtual addresses are 16 bits wide.
- Physical addresses are 14 bits wide.
- The page size is 1024 bytes = 2¹⁰
- The TLB is 4-way set associative with 16 total entries

For the following virtual addresses, fill out the table above

- 0x76BD
- 0x57EB

	Parameter	Value						
	VPN	0x						
	TLB Index	0x						
	TLB Tag	0x						
	TLB Hit?							
	(Y/N)							
	Page Fault?							
	(Y/N)							
	PPN	0x						
	PA							
лÈ	1 1 1							

Implicit Lists: Summary

- Implementation: very simple
- Allocate cost:
 - linear time worst case
- Free cost:
 - constant time worst case
- even with coalescing
- Memory usage:
 - will depend on placement policy
 - First-fit, next-fit or best-fit
- Not used in practice for malloc/free because of linear-time allocation
 - used in many special purpose applications
- However, the concepts of splitting and boundary tag coalescing are general to all allocators

TLB				Page Table					
Index	Tag	PPN	Valid	VPN	PPN	Valid	VPN	PPN	Valid
0	А	С	0	00	7	0	10	9	1
	4	9	1	01	5	0	11	0	0
	3	D	0	02	В	1	12	0	0
	9	3	0	03	D	1	13	5	0
1	2	0	1	04	9	0	14	3	0
	F	2	0	05	8	0	15	6	1
	7	4	0	06	5	1	16	Е	0
	5	6	1	07	2	0	17	А	0
2	7	7	0	08	ш	1	18	1	0
	2	9	0	09	0	1	19	0	0
	0	В	1	0A	F	0	1A	F	1
	6	F	1	0B	D	0	1B	9	0
3	6	A	0	0C	С	0	1C	7	0
	7	4	1	0D	4	0	1D	2	1
	С	1	0	0E	3	0	1E	8	0
	1	8	0	0F	1	0	1F	4	1

Practice on Your Own

- Assuming double-word alignment is used and the implicit free list format discussed in slides is used (i.e., 4 byte header). Block sizes are rounded up to nearest multiple of 8 bytes to maintain alignment.
- What would the block size be in bytes (including payload, header, and padding)? What would be stored in the block header (in hex)?
 - malloc(2)
 - malloc(9)
- malloc(16)

13

Today

- Dynamic Memory Allocation (Ch 9.9)
 - Tracking Free Blocks
 - Implicit Lists
 - Explicit Lists
 - Segregated Lists
- Exceptional Control Flow
- Exceptions

Practice on Your Own

- When using implicit lists that only have a header, what is the smallest size of an allocated block (assuming double word alignment)?
- When using implicit lists that use headers and footers in all blocks, what is the smallest size of an allocated block (assuming double word alignment)?
- When using implicit lists that use headers in all blocks but footers only in free blocks, what is the smallest size of an allocated block (assuming double word alignment)?











Splitting Tasks Suppose we split a free block B into two: block A, and block F: Allocated block A needs to be updated and removed from the list Free block B needs to be updated so F is kept in the list sizeAndTags for both A and F needs to be adjusted Shrink A to its allocation size Create a sizeAndTags field in F to reflect new block's size (B – A) The next/previous pointers from A need to be added to F The next/previous blocks need their previous/next pointers to be updated to point to F Useful macros: #define UNSCALED_POINTER_ADD(p,x) ((void*)((char*)(p) + (x))) #define UNSCALED_POINTER_SUB(p,x) ((void*)((char*)(p) - (x)))



Freeing With a LIFO Policy (Case 1)









Today Dynamic Memory Allocation (Ch 9.9) Tracking Free Blocks Implicit Lists Explicit Lists Segregated Lists Exceptional Control Flow Exceptions

Seglist Allocator (cont.)

- To free a block:
 - Coalesce and place on appropriate list
- Advantages of seglist allocators
 - Higher throughput
 - Better memory utilization
 - First-fit search of segregated free list approximates a bestfit search of entire heap.
 - Extreme case: Giving each block its own size class is equivalent to best-fit.
- Seglist is used by GNU malloc provided in C std library

Seglist Allocator

Given an array of free lists, each one for some size class

To allocate a block of size *n*:

- Search appropriate free list for block of size m >= n
- If an appropriate block is found:
- Split block and place fragment on appropriate list (optional)
- If no block is found, try next larger class
- Repeat until block is found

If no block is found:

- Request additional heap memory from OS (using sbrk())
- Allocate block of n bytes from this new memory
- Place remainder as a single free block in largest size class.
- 30

Today

- Dynamic Memory Allocation (Ch 9.9)
 - Tracking Free Blocks
 - Implicit Lists
 - Explicit Lists
 - Segregated Lists
- Exceptional Control Flow
- Exceptions

- Exists at all levels of a computer system
- Low level mechanisms
 - 1. Exceptions
 - Change in control flow in response to a system event (i.e., change in system state)
 - Implemented using combination of hardware and OS software
- Higher level mechanisms
 - 2. Process context switch
 - Implemented by OS software and hardware timer
 - 3. Signals
 - Implemented by OS software
 - 4. Nonlocal jumps: setjmp() and longjmp()
 - Implemented by C runtime library