Systemprogrammering 2007 Föreläsning 6 Dynamic Memory Allocation

Topics

- Simple explicit allocators
 - Data structures
 - Mechanisms
 - Policies
- Explicit doubly-linked free list
- Segregated free lists
- Garbage collection
- Memory-related perils and pitfalls

Dynamic Memory Allocation

Application Dynamic Memory Allocator

Heap Memory

Explicit vs. Implicit Memory Allocator Explicit: application allocates and frees space

- E.g., malloc and free in C
- Implicit: application allocates, but does not free space
- E.g. garbage collection in Java, ML or Lisp

Allocation

- In both cases the memory allocator provides an abstraction of memory as a set of blocks
- Doles out free memory blocks to application

F6 – 3 –

Systemprogrammering 2007

Harsh Reality

Memory Matters!

Memory is not unbounded

- It must be allocated and managed
- Many applications are memory dominated
 - Especially those based on complex, graph algorithms

Memory referencing bugs especially pernicious

Effects are distant in both time and space

Memory performance is not uniform

- Cache and virtual memory effects can greatly affect program performance
- Adapting program to characteristics of memory system can lead to major speed improvements

F6 – 2 –

Systemprogrammering 2007

Process Memory Image



Malloc Package

#include <stdlib.h>

void *malloc(size_t size)

- If successful:
 - Returns a pointer to a memory block of at least size bytes, (typically) aligned to 8-byte boundary.
 - If size == 0, returns NULL
- If unsuccessful: returns NULL (0) and sets errno.

void free(void *p)

- Returns the block pointed at by p to pool of available memory
- p must come from a previous call to malloc or realloc.

void *realloc(void *p, size_t size)

- Changes size of block p and returns pointer to new block.
- Contents of new block unchanged up to min of old and new size.

F6 – 5 –

Systemprogrammering 2007

Assumptions

Assumptions made in this lecture

Memory is word addressed (each word can hold a pointer)



Malloc Example

```
void foo(int n, int m) {
 int i, *p;
 /* allocate a block of n ints */
 if ((p = (int *) malloc(n * sizeof(int))) == NULL) {
   perror("malloc");
   exit(0);
  for (i=0; i<n; i++)</pre>
   p[i] = i;
  /* add m ints to end of p block */
  if ((p = (int *) realloc(p, (n+m) * sizeof(int))) == NULL) {
   perror("realloc");
   exit(0);
 for (i=n; i < n+m; i++)
   p[i] = i;
 /* print new array */
 for (i=0; i<n+m; i++)</pre>
   printf("%d\n", p[i]);
 free(p); /* return p to available memory pool */
```

F6 - 6 -

F6-8-

Systemprogrammering 2007

Constraints

Applications:

- Can issue arbitrary sequence of allocation and free requests
- Free requests must correspond to an allocated block

Allocators

- Can't control number or size of allocated blocks
- Must respond immediately to all allocation requests
 - i.e., can't reorder or buffer requests
- Must allocate blocks from free memory
 - i.e., can only place allocated blocks in free memory
- Must align blocks so they satisfy all alignment requirements
 8 byte alignment for GNU malloc (libc malloc) on Linux boxes
- Can only manipulate and modify free memory
- Can't move the allocated blocks once they are allocated
 - i.e., compaction is not allowed

Systemprogrammering 2007

Performance Goals: Throughput

Given some sequence of malloc and free requests:

R₀, R₁, ..., R_k, ..., R_{n-1}

Want to maximize throughput and peak memory utilization.

These goals are often conflicting

Throughput:

- Number of completed requests per unit time
- Example:
 - 5,000 malloc calls and 5,000 free calls in 10 seconds
 - Throughput is 1,000 operations/second.

Goals of Good malloc/free

Primary goals

- Good time performance for malloc and free
 - Ideally should take constant time (not always possible)
 - Should certainly not take linear time in the number of blocks
- Good space utilization
 - User allocated structures should be large fraction of the heap.
 - Want to minimize "fragmentation".

Some other goals

- Good locality properties
 - Structures allocated close in time should be close in space
 - "Similar" objects should be allocated close in space
- Robust
 - Can check that free(p1) is on a valid allocated object p1
 - Can check that memory references are to allocated space

F6 – 10 –

F6 – 12 –

Systemprogrammering 2007

Performance Goals: Peak Memory Utilization

Given some sequence of malloc and free requests:

• $R_0, R_1, ..., R_k, ..., R_{n-1}$

Def: Aggregate payload P_k:

- malloc(p) results in a block with a payload of p bytes..
- After request R_k has completed, the aggregate payload P_k is the sum of currently allocated payloads.

Def: Current heap size is denoted by H_k

Assume that H_k is monotonically nondecreasing

Def: Peak memory utilization:

- After k requests, peak memory utilization is:
- $U_k = (max_{i < k} P_i) / H_k$

Internal Fragmentation

Poor memory utilization caused by fragmentation.

Comes in two forms: internal and external fragmentation

Internal fragmentation

For some block, internal fragmentation is the difference between the block size and the payload size.



- Caused by overhead of maintaining heap data structures, padding for alignment purposes, or explicit policy decisions (e.g., not to split the block).
- Depends only on the pattern of *previous* requests, and thus is easy to measure.

F6 – 13 –

F6 – 15 –

Systemprogrammering 2007

Implementation Issues

- How do we know how much memory to free just given a pointer?
- How do we keep track of the free blocks?
- What do we do with the extra space when allocating a structure that is smaller than the free block it is placed in?
- How do we pick a block to use for allocation -- many might fit?
- How do we reinsert freed block?



External Fragmentation

Occurs when there is enough aggregate heap memory, but no single free block is large enough



Knowing How Much to Free

Standard method

Keep the length of a block in the word preceding the block.

• This word is often called the header field or header

Requires an extra word for every allocated block



Keeping Track of Free Blocks

<u>Method 1</u>: Implicit list using lengths -- links all blocks

<u>Method 2</u>: Explicit list among the free blocks using pointers within the free blocks

6



Method 3: Segregated free list

Different free lists for different size classes

Method 4: Blocks sorted by size

Can use a balanced tree (e.g. Red-Black tree) with pointers within each free block, and the length used as a key

F6 – 17 –

Systemprogrammering 2007

Implicit List: Finding a Free Block

First fit:

Search list from beginning, choose first free block that fits

- Can take linear time in total number of blocks (allocated and free)
- In practice it can cause "splinters" at beginning of list

Next fit:

- Like first-fit, but search list from location of end of previous search
- Research suggests that fragmentation is worse

Best fit:

- Search the list, choose the free block with the closest size that fits
- Keeps fragments small --- usually helps fragmentation
- Will typically run slower than first-fit

Systemprogrammering 2007

Method 1: Implicit List

Need to identify whether each block is free or allocated

- Can use extra bit
- Bit can be put in the same word as the size if block sizes are always multiples of two (mask out low order bit when reading size).



Implicit List: Allocating in Free Block

Allocating in a free block - splitting





Implicit List: Bidirectional Coalescing

Boundary tags [Knuth73]

- Replicate size/allocated word at bottom of free blocks
- Allows us to traverse the "list" backwards, but requires extra space
- Important and general technique!



Implicit List: Coalescing

Join (coelesce) with next and/or previous block if they are free

Coalescing with next block





Constant Time Coalescing (Case 1)





Constant Time Coalescing (Case 3)



Constant Time Coalescing (Case 2)



F6 – 26 –

Constant Time Coalescing (Case 4)



Systemprogrammering 2007

Summary of Key Allocator Policies

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 adjacent blocks each time free is called
- Deferred coalescing: try to improve performance of free by deferring coalescing until needed. e.g.,
 - Coalesce as you scan the free list for malloc.
 - Coalesce when the amount of external fragmentation reaches some threshold.

F6 – 29 –

Systemprogrammering 2007

Systemprogrammering 2007

Keeping Track of Free Blocks

Method 1: Implicit list using lengths -- links all blocks



- <u>Method 2</u>: Explicit list among the free blocks using pointers within the free blocks
- Method 3: Segregated free lists

5

- Different free lists for different size classes
- <u>Method 4</u>: Blocks sorted by size (not discussed)
 - Can use a balanced tree (e.g. Red-Black tree) with pointers within each free block, and the length used as a key

Implicit Lists: Summary

- Implementation: very simple
- Allocate: linear time worst case
- Free: 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 allocate. Used in many special purpose applications.

However, the concepts of splitting and boundary tag coalescing are general to *all* allocators.

F6 – 30 –



It is important to realize that links are not necessarily in the sam order as the blocks



Explicit List Summary

Comparison to implicit list:

- Allocate is linear time in number of free blocks instead of total blocks -- much faster allocates when most of the memory is full
- Slightly more complicated allocate and free since needs to splice blocks in and out of the list
- Some extra space for the links (2 extra words needed for each block). This means that the smallest allocated block gets bigger which increases internal fragmentation.

Main use of linked lists is in conjunction with segregated free lists

Keep multiple linked lists of different size classes, or possibly for different types of objects

Freeing With Explicit Free Lists

Insertion policy: Where in the free list do you put a newly freed block?

- LIFO (last-in-first-out) policy
 - Insert freed block at the beginning of the free list
 - Pro: simple and constant time
 - Con: studies suggest fragmentation is worse than address ordered.
- Address-ordered policy
 - Insert freed blocks so that free list blocks are always in address order
 - » i.e. addr(pred) < addr(curr) < addr(succ)
 - Con: requires search
 - Pro: studies suggest fragmentation is better than LIFO

F6 – 34 –

Systemprogrammering 2007

Keeping Track of Free Blocks

Method 1: Implicit list using lengths -- links all blocks



<u>Method 2</u>: Explicit list among the free blocks using pointers within the free blocks



Method 3: Segregated free list

Different free lists for different size classes

Method 4: Blocks sorted by size

Can use a balanced tree (e.g. Red-Black tree) with pointers within each free block, and the length used as a key

Segregated Storage

Each size class has its own collection of blocks



Often have separate size class for every small size (2,3,4,...)

For larger sizes typically have a size class for each power of 2

F6 – 37 –

Systemprogrammering 2007

Segregated Fits

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

To free a block:

Coalesce and place on appropriate list (optional)

Tradeoffs

- Faster search than sequential fits (i.e., log time for power of two size classes)
- Controls fragmentation of simple segregated storage
- Coalescing can increase search times
 - Deferred coalescing can help

Systemprogrammering 2007

Simple Segregated Storage

Separate heap and free list for each size class

No splitting

To allocate a block of size n:

- If free list for size n is not empty,
 - allocate first block on list (note, list can be implicit or explicit)
- If free list is empty,
 - get a new page
 - create new free list from all blocks in page
 - allocate first block on list
- Constant time

To free a block:

- Add to free list
- If page is empty, return the page for use by another size (optional)

Tradeoffs:

Fast, but can fragment badly

F6 – 38 –

Systemprogrammering 2007

For More Info on Allocators

D. Knuth, "The Art of Computer Programming, Second Edition", Addison Wesley, 1973

The classic reference on dynamic storage allocation

Wilson et al, "Dynamic Storage Allocation: A Survey and Critical Review", Proc. 1995 Int'l Workshop on Memory Management, Kinross, Scotland, Sept, 1995.

- Comprehensive survey
- Available from CS:APP student site (csapp.cs.cmu.edu)

Implicit Memory Management: Garbage Collection

Garbage collection: automatic reclamation of heap-allocated storage -- application never has to free

```
void foo() {
    int *p = malloc(128);
    return; /* p block is now garbage */
```

Common in functional languages, scripting languages, and modern object oriented languages:

Lisp, ML, Java, Perl, Mathematica,

Variants (conservative garbage collectors) exist for C and

C++

Cannot collect all garbage

```
F6-41-
```

t collect all garbage

Systemprogrammering 2007

Dereferencing Bad Pointers

The classic scanf bug

int val;
scanf("%d", val);

Using a pointer that doesn't point at allocated memory



Memory-Related Bugs

Dereferencing bad pointers

Using unallocated memory

Reading uninitialized memory

Overwriting memory

Referencing nonexistent variables

Freeing blocks multiple times

Referencing freed blocks

Failing to free blocks

F6 – 42 –

Systemprogrammering 2007

Reading Uninitialized Memory

Assuming that heap data is initialized to zero

/* return y = Ax */
int *matvec(int **A, int *x) {
 int *y = malloc(N*sizeof(int));
 int i, j;

for (i=0; i<N; i++)
 for (j=0; j<N; j++)
 y[i] += A[i][j]*x[j];
return y;</pre>

F6 – 44 –



<section-header><section-header><section-header><section-header><section-header><text><text><page-footer>

Overwriting Memory

Referencing a pointer instead of the object it points to

<pre>int *BinheapDelete(int **binheap, int *size) {</pre>
<pre>int *packet;</pre>
<pre>packet = binheap[0];</pre>
<pre>binheap[0] = binheap[*size - 1];</pre>
*size;
<pre>Heapify(binheap, *size, 0);</pre>
return(packet);

F6 – 48 –

Misunderstand	Overwriting Memory ing pointer arithmetic int *search(int *p, int val) { while (*p && *p != val) p += sizeof(int); return p; }		Referen Forgetting tha returns	tint *foo () { int *foo () { int val; return &val }	les
F6-49- Free	ing Blocks Multiple Ti	Systemprogrammering 2007	F6-50-	oferencing Freed Blocks	Systemprogrammering 2007
Nasty!	<pre>x = malloc(N*sizeof(int)); <manipulate x=""> free(x); y = malloc(M*sizeof(int)); <manipulate y=""> free(x);</manipulate></manipulate></pre>		Evil!	<pre>x = malloc(N*sizeof(int)); <manipulate x=""> free(x); y = malloc(M*sizeof(int)); for (i=0; i<m; i++)<br="">y[i] = x[i]++;</m;></manipulate></pre>	
F6 – 51 –		Systemprogrammering 2007	F6 – 52 –		Systemprogrammering 200

Failing to Free Blocks (Memory Leaks)

Slow, long-term killer!

foo() {
 int *x = malloc(N*sizeof(int));
 ...
 return;

F6 – 53 –

Systemprogrammering 2007

Dealing With Memory Bugs

Conventional debugger (gdb)

Good for finding bad pointer dereferences

Hard to detect the other memory bugs

Debugging malloc (CSRI UToronto malloc)

- Wrapper around conventional malloc
- Detects memory bugs at malloc and free boundaries
 - Memory overwrites that corrupt heap structures
 - Some instances of freeing blocks multiple times
 - Memory leaks
- Cannot detect all memory bugs
 - Overwrites into the middle of allocated blocks
 - Freeing block twice that has been reallocated in the interim
 - Referencing freed blocks

Failing to Free Blocks (Memory Leaks)

Freeing only part of a data structure

struct list {		
inc var,		
struct list	*next;	
};		
foo() {		
struct list	*head =	
	<pre>malloc(sizeof(struct list));</pre>	
head->val =	0;	
head->next	= NULL;	
<create and<="" td=""><td>manipulate the rest of the list></td><td></td></create>	manipulate the rest of the list>	
•••		
<pre>free(head);</pre>		
return;		
}		

F6 – 54 –

Systemprogrammering 2007

Dealing With Memory Bugs (cont.)

Binary translator

- Powerful debugging and analysis technique
- Rewrites text section of executable object file
- Can detect all errors as debugging malloc
- Can also check each individual reference at runtime
 - Bad pointers
 - Overwriting
 - Referencing outside of allocated block

Garbage collection (Boehm-Weiser Conservative GC)

Let the system free blocks instead of the programmer.

Systemprogrammering 2007

F6 – 56 –