

Memory (Stack and Heap)

Praktikum „C-Programmierung“

Nathanael Hübbe, Eugen Betke, Michael Kuhn, Jakob Lüttgau, Jannek Squar

Wissenschaftliches Rechnen
Fachbereich Informatik
Universität Hamburg

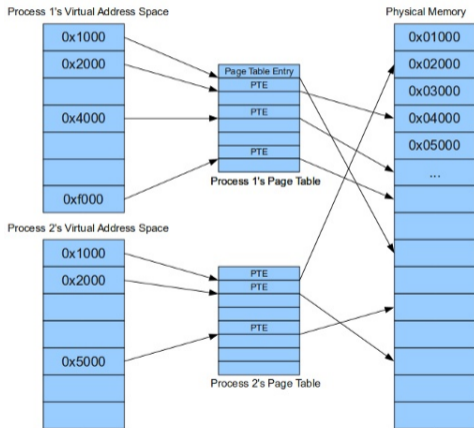
2018-12-03



Universität Hamburg
DER FORSCHUNG | DER LEHRE | DER BILDUNG

- 1 Introduction
- 2 Heap vs. Stack
- 3 Memory management in C
 - Library
 - Common techniques
- 4 Pitfalls
- 5 Quellen

Virtual address space



- Modern operating systems manage memory allocation
- A page table maps physical memory to virtual address space
- Each process see contiguous memory space

Memory layout of C programs

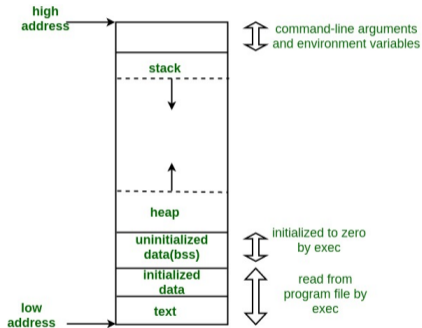


Abbildung: Memory Layout [1]

- Text segment
 - Contains executable instructions
 - Often read only and shared by processes
- DATA:
 - Data segment, initialized by programmer
- BSS:
 - Data segment, uninitialized by programmer
 - Initialized to arithmetic 0
 - Stores global and static variables
- Stack
 - Stores automatic variables
 - Typically grows from higher addresses towards zero
- Heap
 - Dynamically allocated space
 - Begins at the end of BSS segment

1 Introduction

2 Heap vs. Stack

3 Memory management in C

■ Library

■ Common techniques

4 Pitfalls

5 Quellen

Stack vs. Heap [2]

■ Stack

- very fast access
- don't have to explicitly de-allocate variables
- space is managed efficiently by CPU, memory will not become fragmented
- local variables only
- limit on stack size (OS-dependent)
- variables cannot be resized

■ Heap

- variables can be accessed globally
- no limit on memory size
- (relatively) slower access
- no guaranteed efficient use of space, memory may become fragmented over time as blocks of memory are allocated, then freed
- you must manage memory (you're in charge of allocating and freeing variables)
- variables can be resized using `realloc()`

- 1 Introduction
- 2 Heap vs. Stack
- 3 Memory management in C**
 - Library
 - Common techniques
- 4 Pitfalls
- 5 Quellen

C memory management functions

<code>malloc</code>	allocates memory
<code>calloc</code>	allocates and zeroes memory
<code>realloc</code>	expands previously allocated memory block
<code>free</code>	deallocates previously allocated memory

- The memory management functions
 - allocate and deallocated memory on the heap
 - are defined in header `<stdlib.h>`

malloc

```
1 void *malloc(size_t size);
```

- Allocates size bytes and returns a pointer to the allocated memory.
- The memory is not initialized.
- If size is 0, then malloc() returns either NULL, or a unique pointer value that can later be successfully passed to free().

calloc

```
1 void *calloc(size_t nmemb, size_t size);
```

- Allocates memory for an array of `nmemb` elements of `size` bytes each and returns a pointer to the allocated memory
- The memory is set to zero
- If `nmemb` or `size` is 0, then `calloc()` returns either `NULL`, or a unique pointer value that can later be successfully passed to `free()`

realloc

```
1 void *realloc(void *ptr, size_t size);
```

- Changes the size of the memory block pointed to by ptr to size bytes
- The contents will be unchanged in the range from the start of the region up to the minimum of the old and new sizes.
- If the new size is larger than the old size, the added memory will not be initialized.

```
1 p = realloc(NULL, size);  
2 // the same as  
3 p = malloc(size);
```

```
1 realloc(ptr, 0);  
2 // the same as  
3 free(ptr)
```

free

```
1 void free(void *ptr);
```

- Frees the memory space pointed to by ptr
- ptr must have been returned by a previous call to malloc(), calloc(), or realloc().
- Otherwise, or if free(ptr) has already been called before, undefined behavior occurs.
- If ptr is NULL, no operation is performed.

Usage example

```
1  /* Allocate space for an array with ten elements of type int. */
2  int *ptr = (int *) malloc(10 * sizeof (int));
3  if (ptr == NULL) {
4      /* Memory could not be allocated, the program should
5       handle the error here as appropriate. */
6  } else {
7      /* Allocation succeeded. Do something. */
8      free(ptr);
9      ptr = NULL;
10 }
```

Allocation idiom

```
1 int *ptr1 = malloc(10 * sizeof(int));
```

Listing 1: Straight forward allocation

```
1 int *ptr2 = malloc(10 * sizeof(*ptr));
```

Listing 2: Allocation idiom

- Instead of `sizeof(int)` we used `sizeof(*ptr)`
- `sizeof()` automatically determines the correct size of `*ptr`
- If type changes, the `malloc` still provides right amount of memory

Single dynamic allocation with ownership semantics

```
1 {  
2     struct obj otmp;  
3     /* do stuff with otmp */  
4 }  
5  
6 {  
7     struct obj *otmp;  
8     otmp = malloc(sizeof *otmp);  
9     /* do stuff with otmp */  
10    free(otmp);  
11 }
```

- User is responsible for memory allocation and deallocation
- Memory is allocated and released in the same scope

Dynamic memory allocation in functions

```
1 int *my_func(void) {  
2     int * x;  
3     x = (int *) malloc(25);  
4     return x;  
5 }  
6  
7 int main(int argc, char** argv) {  
8     int *pi = my_func();  
9     /* do something with pi */  
10    free(pi);  
11    return 0;  
12 }
```

- Memory is allocated by `malloc` in a function
- The user is responsible for deallocation

- 1 Introduction
- 2 Heap vs. Stack
- 3 Memory management in C
 - Library
 - Common techniques
- 4 Pitfalls**
- 5 Quellen

Common pitfalls

Memory management is a source for bugs [4], e.g.

- Memory leaks
- Use after free
- Freeing not dynamically allocated memory
- Accessing not allocated memory
- Multiple free

Memory leaks

```
1 int memory_leak() {  
2     int* ptr = malloc(sizeof(*ptr));  
3     return 0;  
4 }
```

- Pointer get lost after function returns
- Memory stays allocated, until program exits

Use after free

```
1 int *ptr = malloc(sizeof (int));  
2 free(ptr);  
3 *ptr = 7; /* Undefined behavior */
```

- Memory is used after it was freed
- Results in undefined behaviour

Freeing not dynamically allocated memory

```
1 char *msg = "Default message";  
2 int tbl[100];  
3 free(msg);  
4 free(tbl);
```

- Freeing memory that was not allocated by `malloc`, `calloc` or `realloc`
- Results in undefined behaviour

Accessing not allocated memory

```
1 int *my_func(void) {  
2     int x[25];  
3     return x;  
4 }  
5  
6 int main(int argc, char** argv) {  
7     int *pi = my_func();  
8     *(pi + 1) = 5;  
9     free(pi);  
10    return 0;  
11 }
```

- Not allocated memory is used after freed

Multiple free

```
1 void main(){
2     int *p;
3     p = (int *)malloc(10 * sizeof(int));
4     f(p);
5     free(p);
6 }
7 void f(int *g){
8     printf("%d", g);
9     free(g);
10 }
```

- The same allocation is freed several times
- Results in undefined behaviour

Summary

- From application's perspective, memory is a contiguous space partitioned in segments (Text, DATA, BSS, Stack and Heap)
- Memory management
 - On stack memory is managed for you
 - On heap the user is responsible for memory management
 - Memory management functions are defined in `<stdlib.h>`
- Common bugs
 - Memory leaks
 - Use after free
 - Freeing not dynamically allocated memory
 - Accessing not allocated memory
 - Multiple free

- 1 Introduction
- 2 Heap vs. Stack
- 3 Memory management in C
 - Library
 - Common techniques
- 4 Pitfalls
- 5 Quellen**

Quellen I

- [1] **GeeksforGeeks**. Memory Layout of C Programs.
<https://www.geeksforgeeks.org/memory-layout-of-c-program/>.
Accessed on 03.12.2014.
- [2] **Gribblelab**. Memory: Stack vs Heap. https://www.gribblelab.org/CBootCamp/7_Memory_Stack_vs_Heap.html.
Accessed on 03.12.2014.
- [3] **Turkeyland**. Buffer Overflows and You.
<https://turkeyland.net/projects/overflow/intro.php>. Accessed on 03.12.2014.

Quellen II

- [4] Wikibooks. C Programming/stdlib.h/malloc. https://en.wikibooks.org/wiki/C_Programming/stdlib.h/malloc. Accessed on 03.12.2014.