

Programmieren in C

Fabian Klötzl

MetaNook 2017

- Master in Informatik 2015
- PhD im MPI für Evolutionsbiologie, Plön
- 29 GitHub Repos
- Debian Med Team
- Folien unter `kloetzl.info#downloads`

Section 1

Eine Einführung in C

Hello World

```
1 #include <stdio.h>
2
3 int main(void)
4 {
5     printf(" Hello, World!\n");
6     return 0;
7 }
```

Warum C im Jahre 2017?

Warum C im Jahre 2017?

Performance

Es gibt genügend gute andere Optionen: C++, D, Rust, Go.
Die bieten gleichzeitig bessere Abstraktionen, mehr Features, mehr Sicherheit, ...

Warum C im Jahre 2017?

Performance

Es gibt genügend gute andere Optionen: C++, D, Rust, Go.
Die bieten gleichzeitig bessere Abstraktionen, mehr Features, mehr Sicherheit, ...

Systems Programming

Kernel, Treiber, Glue-Code

- Statisches Typsystem
- Strukturen
- Zeiger
- Compiliert schnell
- Minimal

Warum nicht C++?

```
1 auto numbers = std::vector<int>{0,1,2,3,4,5,6,7,8,9};
2
3 auto count = 0;
4 auto odd = [=] (int n) mutable {
5     if (n & 1) {
6         printf("%i_0%i\n", n, count);
7         count++;
8         return true;
9     }
10    return false;
11 };
12
13 std::remove_if(std::begin(numbers), std::end(numbers), odd);
```

C++ is a beast

Expected Output

```
1 0  
3 1  
5 2  
7 3  
9 4
```

https://gcc.gnu.org/bugzilla/show_bug.cgi?id=81482

C++ is a beast

Expected Output

```
1 0  
3 1  
5 2  
7 3  
9 4
```

Actual Output

```
1 0  
3 0  
5 1  
7 2  
9 3
```

https://gcc.gnu.org/bugzilla/show_bug.cgi?id=81482

Wie groß ist ein int?

- a) 16 bit
- b) 32 bit
- c) 40 bit
- d) 64 bit

Wie groß ist ein int?

- a) 16 bit
- b) 32 bit
- c) 40 bit
- d) 64 bit

Lösung

Alles ist möglich!

2 byte \leq short \leq int,

int \leq long,

4 byte \leq long

Vorzeichen

Signed oder Unsigned, das ist hier die Frage.

- char
- short
- int
- size_t

Vorzeichen

Signed oder Unsigned, das ist hier die Frage.

- char
- short
- int
- size_t

Lösung

- char: je nach Encoding
- short: signed
- int: signed
- size_t: unsigned (ssize_t)

Alle Typen die ihr je braucht

- `char`: Byte, Zeichen
- `int`: Bool, File Descriptor
- `size_t`: Indices, Größen
- `double`: Reelle Zahlen
- `arrays`: Arrays
- `struct`: Zusammengesetzte Datentypen
- `pointer`: Zeiger


```
1  int add(int a, int b);  
2  void fn(void);  
3  
4  int add(int a, int b)  
5  {  
6      return a + b;  
7  }
```

```
1 struct range {  
2     size_t lower, upper,  
3     double weight;  
4 };  
5  
6 struct range r = {0, 100, 0.5};  
7 double area = (r.upper - r.lower) * r.weight;  
8 struct range s = r; // copy
```

Statische Arrays

```
1 int buffer[1024] = {0};
2
3 buffer[0] = 23;
4 memset(buffer, 1, sizeof(buffer));
5 printf("%i\n", buffer[0]); // 16843009 = 0x1010101
```

```
1 void fn(size_t count)
2 {
3     int arr[count] = {0};
4
5     do_something(arr);
6
7     // automatic deallocation
8 }
```

Schuhe

```
1 struct range normalize(struct range r) {
2     double inv = 1.0 / (r.upper - r.lower);
3     r.weight = inv;
4     return r;
5 }
6
7 void normalize2(struct range *range_ptr) {
8     double inv = 1.0 / (( *range_ptr).upper - range_ptr->lower);
9     range_ptr->weight = inv;
10 }
11
12 struct range rr;
13 rr = normalize(rr);
14 normalize2(&rr);
```

```
1 void normalize2(struct range *range_ptr) {
2     double inv = 1.0 / (( *range_ptr).upper - range_ptr->lower);
3     range_ptr->weight = inv;
4 }
5
6 normalize2(NULL); // SEGFAULT
7 int foo;
8 normalize2(&foo); // int* != struct range*
```

```
1 void swap_int(int *a, int *b) {  
2     int temp = *a;  
3     *a = *b;  
4     *b = temp;  
5 }  
6  
7 int x = 4;  
8 int y = 5;  
9 swap_int(&x, &y);
```

Definition

Ein *String* ist eine Folge von Zeichen, die mit NUL (einem Null-Byte) endet.

```
1  const char *str = "Hallo";
2  // str[0] = 'H'
3  // str[1] = 'a'
4  // str[2] = 'l'
5  // str[3] = 'l'
6  // str[4] = 'o'
7  // str[5] = '\0'
8  strlen(str) // 5
9  printf("%s\n", str) // Hallo
```


Strings

```
1  const char *str = "Hallo";  
2  str[1] = 'e';  
3  str[4] = '\0';  
4  strlen(str) // 4  
5  printf("%s\n", str) // Hell
```

Strings

strdup

```
1  const char *str = "Hallo";  
2  char *mut = strdup(str);  
3  mut[1] = 'e';  
4  mut[4] = '\\0';  
5  strlen(str) // 5  
6  strlen(mut) // 4
```

```
1 int *arr = malloc(count * sizeof(int));  
2  
3 do_something(arr);  
4  
5 free(arr);
```

Daumenregel

```
grep -c 'free' ≈ grep -c 'malloc'
```

gilt auch für strdup!

```
1 int *arr = malloc(count * sizeof(int));
2 if (arr == NULL) {
3     err(errno, "out_of_memory");
4 }
5
6 do_something(arr);
7
8 free(arr);
```

```
1 int *arr = reallocarray(NULL, count, sizeof( *arr));
2 if (arr == NULL) {
3     err(errno, "out_of_memory");
4 }
5
6 do_something(arr);
7
8 free(arr);
```

Section 2

Übung

Eigenschaften:

- kontinuierlicher Speicher
- dynamisch verlängerbar

Operationen:

- Erstellen
- Ein Element anhängen
- Anzahl der gespeicherten Elemente abfragen
- Auf Daten zugreifen
- Speicher freigeben

Elementtyp: `long`

```
1 struct vector_long {  
2     long *data;  
3     size_t size;  
4     size_t capacity;  
5 };
```



```
1 // int vector_long_init(struct vector_long*);  
2 void vector_long_free(struct vector_long*);  
3 int vector_long_push(struct vector_long*, long);  
4 size_t vector_long_size(struct vector_long*);  
5 long* vector_long_data(struct vector_long*);
```

```
1  int vector_long_init(struct vector_long* vec) {
2      if (!vec) return 1;
3
4      vec->data = reallocarray(NULL, 4, sizeof(*vec->data));
5      vec->size = 0;
6      vec->capacity = 4;
7
8      // return 0 iff successful
9      return vec->data == NULL ? 1 : 0;
10 }
```

```
1 void vector_long_free(struct vector_long* vec) {  
2     if (!vec) return;  
3     free(vec->data);  
4     *vec = (struct vector_long){0};  
5 }
```

```
1 long *vector_long_data(struct vector_long* vec) {
2     if (!vec) return NULL;
3     return vec->data;
4 }
5
6 size_t vector_long_size(struct vector_long *vec) {
7     if (!vec) return 0;
8     return vec->size;
9 }
```

```
1 int vector_long_push(struct vector_long* vec, long element) {
2     if (!vec) return 1;
3
4     if (vec->size < vec->capacity) {
5         vec->data[vec->size++] = element;
6         return 0;
7     }
8
9     long *ptr = reallocarray(vec->data, vec->size / 2, 3 * sizeof(long));
10    if (!ptr) return 1;
11    vec->data = ptr;
12    vec->capacity = (vec->capacity / 2) * 3;
13    vec->data[vec->size++] = element;
14
15    return 0;
16 }
```

Section 3

Tipps und Tools

Commandline Arguments

```
1  int main(int argc, char **argv) {
2      int c;
3      int flag = 0;
4      int iterations = 10;
5      while((c = getopt(argc, argv, "hfi:")) != -1) {
6          if (c == 'h') usage(EXIT_SUCCESS);
7          if (c == 'f') flag = 1;
8          if (c == 'i') iterations = atoi(optarg);
9          if (c == '?') usage(EXIT_FAILURE);
10     }
11     ...
12 }
```

Konvertierung

string to int

```
1 int iterations;
2 const char *errstr;
3
4 iterations = strtoum(optarg, 1, 64, &errstr);
5 if (errstr)
6     errx(1, "number_of_iterations_is_%s:_%s", errstr, optarg);
```


Konvertierung

int to string

```
1 int iterations = 10;
2 char *str;
3
4 int check = asprintf(&str, "%i", iterations);
5 if (check < 0)
6     errx(1, "critical_error");
```

Man Pages

- `man malloc`
- `man 3 free`
- `man ascii`
- `help2man`

- Compiler: clang, gcc
- Warnings: -Wall -Wextra
- Optimierung: -O0, -O2, -O3
- -march=native
- Debugging: -ggdb -fno-omit-frame-pointer

perf stat -d tool file.ext

```
Performance counter stats for 'andi /home/kloetzl/Projects/Arbeit/data/eco29.fasta':

 19813,809593      task-clock:u (msec)          #    2,443 CPUs utilized
          0        context-switches:u          #    0,000 K/sec
          0        cpu-migrations:u           #    0,000 K/sec
      34.279      page-faults:u                #    0,002 M/sec
48.182.898.185    cycles:u                     #    2,432 GHz              (62,39%)
25.241.341.345    instructions:u               #    0,52 insn per cycle   (74,88%)
 5.546.686.564    branches:u                   # 279,940 M/sec            (74,66%)
 339.623.906     branch-misses:u             #    6,12% of all branches (74,95%)
 5.563.169.715    L1-dcache-loads:u           # 280,772 M/sec            (75,11%)
 620.570.988     L1-dcache-load-misses:u     # 11,15% of all L1-dcache hits (75,02%)
 306.180.557     LLC-loads:u                  # 15,453 M/sec            (50,24%)
 254.924.849     LLC-load-misses:u           #  83,26% of all LL-cache hits (50,13%)

 8,110930991 seconds time elapsed
```

perf record

perf record tool file.ext
perf report

```
23,11%  andi    andi    [.] get_interval
19,14%  andi    andi    [.] esa_init
10,89%  andi    libdivsufsort.so.3.0.0 [.] divsufsort
6,54%   andi    andi    [.] get_match_cached
4,85%   andi    andi    [.] get_match_from
3,61%   andi    andi    [.] esa_init_FVC
2,53%   andi    andi    [.] dist_anchor
2,01%   andi    libdivsufsort.so.3.0.0 [.] 0x000000000000061ec
1,87%   andi    libdivsufsort.so.3.0.0 [.] 0x00000000000006220
1,54%   andi    andi    [.] model_count
1,01%   andi    andi    [.] revcomp
0,83%   andi    andi    [.] pfasta_read_seq
0,58%   andi    andi    [.] esa_init_cache_fill
0,56%   andi    libdivsufsort.so.3.0.0 [.] 0x00000000000003e20
```

perf record tool file.ext

perf report

```
0,06      mov     (%rdi),%r8
0,31      ↓ je     d8
           }

           int m = ij.m;
           int l = ij.l;

           char c = S[SA[i] + l];
27,81     mov     0x0(%r13,%r15,4),%eax
           const saidx_t *LCP = self->LCP;
1,11     mov     0x10(%rdi),%rbx
           const saidx_t *CLD = self->CLD;
0,06     mov     0x30(%rdi),%r12
           char c = S[SA[i] + l];
0,09     add     %esi,%eax
0,26     cltq
```

Bad Code

```
1  name[M] = (char*) calloc(kseq->name.l + 1, sizeof(char));  
2  strcpy(name[M], kseq->name.s);
```

<https://github.com/tseemann/snp-dists>

```
1 name[M] = (char*) calloc(kseq->name.l + 1, sizeof(char));  
2 strcpy(name[M], kseq->name.s);
```

<https://github.com/tseemann/snp-dists>

- 1 `sizeof(char)` ist immer 1
- 2 warum `calloc`, wenn wir gleich wieder überschreiben?
- 3 `malloc` nicht casten
- 4 wir kennen die Länge → `memcpy`
- 5 memory leak (not shown)


```
1 name[M] = (char*) calloc(kseq->name.l + 1, sizeof(char));  
2 strcpy(name[M], kseq->name.s);
```

```
1 name[M] = malloc(kseq->name.l + 1);  
2 memcpy(name[M], kseq->name.s, kseq->name.l + 1);
```

```
1 name[M] = strdup(kseq->name.s);
```

Bad Code 2

```
1  int ** array;
2  array = malloc(nrows * sizeof(double * ));
3
4  for(i = 0; i < nrows; i++){
5      array[i] = malloc(ncolumns * sizeof(double));
6  }
7
8  array[0][1] = 3.141;
```

<http://moreisdifferent.com/2015/07/16/why-physicsts-still-use-fortran/>

Bad Code 2

```
1  double **array = malloc(nrows * sizeof(double *));
2
3  for(i = 0; i < nrows; i++){
4      array[i] = malloc(ncolumns * sizeof(double));
5  }
6
7  array[0][1] = 3.141;
```

Bad Code 2

```
1  double **array = malloc(nrows * sizeof(double *));
2  double *base = malloc(ncolumns * nrows * sizeof(double));
3
4  for(i = 0; i < nrows; i++){
5      array[i] = base + i * ncolumns;
6  }
7
8  array[0][1] = 3.141;
```

Bad Code 2

```
1 double *base = malloc(ncolumns * nrows * sizeof( *base));
2
3 #define B(X,Y) (base[(X)*ncolumns + (Y)])
4
5 B(0,1) = 3.141;
```

- BSD Userland
- libc
- Debian Code Search
- GitHub

Optionale und Benamte Argumente

```
1 struct opts { size_t size; };
2
3 int vector_long_init(struct vector_long* vec, struct opts o) {
4     vec->data = malloc(o.size * sizeof(long));
5     ...
6 }
7
8 #define vector_long_init(VEC, ...) \
9     vector_long_init((VEC), (struct opts){.size = 4, __VA_ARGS__})
10
11
12 vector_long_init(&vec);
13 vector_long_init(&vec, .size = 10);
```

Bad Code 3

```
1 char *path_name(const struct name_path *path, const char *name)
2 {
3     const struct name_path *p;
4     int nlen = strlen(name);
5     int len = nlen + 1;
6
7     for (p = path; p; p = p->up) {
8         if (p->elem_len)
9             len += p->elem_len + 1;
10    }
11    char *n = xmalloc(len);
12    ...
13    return n;
14 }
```

<https://github.com/git/git/blob/v1.7.0/revision.c>

Bad Code 3

Verbessert

```
1 char *path_name(const struct name_path *path, const char *name)
2 {
3     const struct name_path *p;
4     size_t nlen = strlen(name);
5     size_t len = nlen + 1;
6
7     for (p = path; p; p = p->up) {
8         if (p->elem_len)
9             len += p->elem_len + 1;
10    }
11    char *n = xmalloc(len);
12    ...
13    return n;
14 }
```

Bad Code 4

```
1 char *line = fd_read_line (fd);
2 char *endl;
3 if (line == NULL)
4     break;
5
6 remaining_chunk_size = strtol (line, &endl, 16);
7 xfree (line);
8
9 if (remaining_chunk_size == 0) {
10     line = fd_read_line (fd);
11     xfree (line);
12     break;
13 }
14 fd_read (fd, dlbuf, remaining_chunk_size, -1);
```

<https://access.redhat.com/security/cve/cve-2017-13089>

Bad Code 4

Verbessert

```
1 char *line = fd_read_line (fd);
2 char *endl;
3 if (line == NULL)
4     break;
5
6 remaining_chunk_size = strtol (line, &endl, 16);
7 xfree (line);
8
9 if (remaining_chunk_size < 0)
10     return false;
11
12 if (remaining_chunk_size == 0) {
13     line = fd_read_line (fd);
14     xfree (line);
15     break;
16 }
17 fd_read (fd, dlbuf, remaining_chunk_size - 1);
```

```
1 struct range {
2     size_t lower, upper;
3     double weight;
4 };
5
6 struct colored_range {
7     size_t lower, upper;
8     double weight;
9     int color;
10 };
```

```
1 struct colored_range {  
2     struct range base;  
3     int color;  
4 };  
5  
6 struct colored_range cr;  
7 cr.base.left = 0;  
8 normalize(&cr.base);
```

```
1 struct colored_range {  
2     struct range; // gcc -fms-extensions  
3     int color;  
4 };  
5  
6 struct colored_range cr;  
7 cr.left = 0;  
8 normalize(&cr); // gcc -fplan9-extensions
```

```
1 struct colored_range {
2     union {
3         struct range;
4         struct range as_range;
5     };
6     int color;
7 };
8
9 struct colored_range cr;
10 cr.left = 0;
11 normalize(&cr.as_range);
```

Bad Code 5

```
1 struct Score
2 {
3     long int value;
4     char used;
5 };
6
7 struct Node
8 {
9     Score S[3];
10    int m_max;
11 };
12
13 struct Node huge_vector[1000]; // conceptually
```


Bad Code 5

Verbessert

```
1 struct Node
2 {
3     long int values[3];
4     char used[3];
5     int m_max;
6 };
7
8 struct Node huge_vector[1000]; // conceptually
```

- *gdb --args !!*
- *run* Starten
- *list* Aktuelle Codezeile ausgeben
- *p* Variable ausgeben

- `-fsanitize=address`
- Address: Erkennt Pointerfehler
- Integer: Erkennt Over/Underflows
- Undefined: UB
- Thread: Race-Conditions
- Dynamisch (zur Laufzeit), gute Unittests!

- *scan-build gcc*
- *scan-build make*
- Findet Bugs wzb. Null-Pointer-Deref
- Gibt Pfad zum Bug an.
- Bei Compilierung (statisch)

- `clang-format -i *.c`
- Formatiert euren Code automatisch passend zu (euren Vorlieben) den Vorgaben des Projektes
- Braces, Zeilenumbrüche, Includes sortieren, ...
- Clang-Format-Datei:

- 1 `BasedOnStyle: LLVM`
- 2 `IndentWidth: 4`
- 3 `TabWidth: 4`
- 4 `UseTab: Always`
- 5 `AllowShortIfStatementsOnASingleLine: true`
- 6 `AllowShortFunctionsOnASingleLine: false`

```
1 int func(char *bar) __attribute__((target_clones("AVX,default")))  
2 {  
3     // code goes here  
4     return 0;  
5 }
```

- 1 Dynamische Auflösung zur Laufzeit
- 2 Gute Lösung für Paketverwaltung

- `/usr/bin/time -f "%e%M%P" tool args`
- `%e`: elapsed time (wall clock)
- `%M`: memory peak
- `%P`: durchschnittliche CPU-Auslastung
- $\hat{t} = t + t_{clock} + t_{sched} + t_{caching}$
- $\hat{t}_1, \hat{t}_2, \hat{t}_3 \rightarrow t?$

- `/usr/bin/time -f "%e%M%P" tool args`
- `%e`: elapsed time (wall clock)
- `%M`: memory peak
- `%P`: durchschnittliche CPU-Auslastung
- $\hat{t} = t + t_{clock} + t_{sched} + t_{caching}$
- $\hat{t}_1, \hat{t}_2, \hat{t}_3 \rightarrow t?$
- $t_x \geq 0$
- $t = \min(\hat{t}_1, \hat{t}_2, \hat{t}_3)$

Wofür ich keine Zeite hatte

- unions, enums
- undefined behaviour → CppCon Talks
`unsigned long file_size = 5 << 30; //5gb`
- const, inline
- typedef
- valgrind
- Google Benchmark
- <https://matt.sh/howto-c>
- fuzzing
- strict aliasing

Zusammenfassung

`kloetzl.info#downloads`