# A benchmark for C program verification

Marko van Eekelen Daniil Frumin Herman Geuvers Léon Gondelman Robbert Krebbers Marc Schoolderman Sjaak Smetsers Freek Verbeek Benoît Viguier Freek Wiedijk

April 1, 2019

#### Abstract

We present twenty-five C programs, as a benchmark for C program verification using formal methods. This benchmark can be used for system demonstration, for comparison of verification effort between systems, and as a friendly competition. For this last purpose, we give a scoring formula that allows a verification system to score up to a hundred points.

# 1 Introduction

Here is a challenge for people who are interested in practical C program verification using formal methods. Show us how to verify the correctness of the twenty-five programs in this paper. The paper and the sources of these programs can be found at:

https://arxiv.org/abs/1904.01009 https://github.com/cverified/cbench/ https://www.cs.ru.nl/~freek/cbench/

If you take this challenge, we will add your work (or a pointer to it, as you prefer) to the repository.

In our Sovereign project [8], which has as its focus the safety and correctness of industrial programs, we *are* interested in practical C program verification. Therefore we would like to have a clearer overview of frameworks for doing this. At first we planned to only have a few small programs for trying out the different systems (e.g., [1], [6], [7], [11], [15]). However, this then grew into this benchmark.

The original idea of this benchmark was that it should be about five programs. It will be completely clear what these programs are when one sees the names:

fac	the basis
cat	I/O
malloc	linked datastructures, memory
qsort	arrays, recursion, function pointers
sqrt	floating point

The few keywords given here indicate why we chose these five programs. (Obvious additions to this list are fib and gcd, but we consider these 'more of the same' with respect to fac.)

For each of these programs we now have several variants, ranging from very small with only a few lines of code to significant programs with reasonable behaviors. We intentionally kept the programs in this benchmark on the spectrum: simplest version

Altogether this gives twenty-five variations on those five programs. Three of these are not written by us, but are GNU implementations [20] that are actually used by many people.

To turn this into a benchmark, then, we award four points for a proper verification of any of the programs. This means that it is possible to score:

 $25 \times 4$  points = 100 points

To get all four points for a program, one should both show that it does not exhibit undefined behavior and prove its functional correctness. Points will be subtracted for the following simplifications:

no verification of defined behavior	-1
no verification of functional correctness	-1
verified modified code, still valid C	-1
verified only a model of the code	-2

For example, if one only verifies functional correctness of a model of the program, one only gets 4 - 1 - 2 = 1 point (-1 for not verifying well-definedness, and -2 for only verifying a model) instead of four. We include the option to modify the code a bit (with the loss of one point), because we expect some frameworks to have restrictions on the code they can verify.

To phrase this more positively, another way to state the way points are awarded is:

verification of defined behavior	+1
verification of functional correctness	+1
verified valid C	+1
verified unmodified code	+1

For the last point it is allowed to add special comments to the code to enable the verification, like in the ACSL language [2].

What it means to correctly verify defined behavior, whether a tool models C accurately (a subject that can lead to heated debate), and how much modification to the code is allowed for the program to still be 'the same', we intentionally leave open. This is a friendly challenge, and should not be taken too seriously: if someone wants to award him- or herself a certain number of points for a verification effort, they should feel free to do so. We primarily hope this 'benchmark' will be useful for demonstrating and comparing verification frameworks on actual C code.

Exactly what it *means* to 'verify' code like this we also intentionally leave open. One is free to decide for oneselves what the specification for the programs should be. In practice, we expect that this means that one has to give:

- Restrictions on the C implementation.
- Pre-conditions for the code, e.g., restrictions on the input.
- Specification of the behavior of the code, e.g., appropriate post-conditions.

The C standard [14] has the notion of *implementation-defined behavior*. This first item is about making explicit under what implementations the program should work.

For example, the last programs in our benchmark presuppose IEEE 754-compliant floating point numbers. While the object representation of floating point numbers is *not* implementation-defined (but unspecified), there is the \_\_STDC\_IEC\_559\_\_ pre-defined constant, which *is* 

implementation-defined. The restriction on the C implementation for these programs then might be that te implementation has to define that constant appropriately, in which case we can assume that the object representation for floating point numbers matches the IEEE 754 standard. Similarly, there might be restrictions on the integer sizes to be 'large enough', or struct layouts to be appropriate to the program.

Here then is a list of our twenty-five programs, with a table of the C features that each of these programs exercises:

$program \ number$	filename	type	lines of code	function calls standard library calls	other library calls recursion	preprocessor: includes	preprocessor: defines	input/output	threads	exit	oreak/continue	switch	labels/goto	arrays/pointers	$function \ pointers$	structs	unions	$defined \ types$	typedefs	floating point	IEEE 754 floating point
1	fac1.c	С	3																		
2	fac2.c	$\mathbf{C}$	3																		
3	fac3.c	$\mathbf{F}$	9																		
4	fac4.c	$\mathbf{F}$	5	×	×																
5	fac5.c	$\mathbf{S}$	20	×	×	×		$\times$										Х			
6	fac6.c	$\mathbf{S}$	70	$\times$ $\times$	×	$\times$	×	×		×								Х			
7	cat1.c	$\mathbf{C}$	2	×				$\times$													
8	cat2.c	$\mathbf{S}$	13	×		×		$\times$						×				Х			
9	cat3.c	$\mathbf{S}$	76	$\times$ $\times$	×	×		$\times$		×	×			×				Х			
10	coreutils/src/cat.c	R	$\sim 800$	$\times$ $\times$	×	$\times$	×	×		×		×	×	$\times$		$\times$		×			
11	malloc1.c	$\mathbf{F}$	16				×							×							
12	malloc2.c	$\mathbf{F}$	28				×							×			$\times$				
13	malloc3.c	$\mathbf{F}$	26				×							×			$\times$				
14	malloc4.c	$\mathbf{F}$	180	$\times$		$\times$	Х							$\times$	$\times$	$\times$		Х	$\times$		
15	glibc/malloc/malloc.c	R	$\sim 6000$	$\times$ $\times$	×	$\times$	Х	$\times$	×			Х	Х	$\times$	$\times$	$\times$					
16	qsort1.c	$\mathbf{F}$	23	$\times$	×									×							
17	qsort2.c	$\mathbf{F}$	50	$\times$	×								×	$\times$							
18	qsort3.c	$\mathbf{F}$	50	$\times$	×		Х				×			$\times$							
19	qsort4.c	$\mathbf{F}$	30	$\times$	×	$\times$	Х				×			$\times$	$\times$						
20	qsort5.c	$\mathbf{S}$	83	$\times$ $\times$	$\times$ $\times$	×	×	×	×		×			$\times$				Х			
21	sqrt1.c	$\mathbf{F}$	14																	$\times$	
22	sqrt2.c	$\mathbf{F}$	12				×							$\times$						$\times$	$\times$
23	sqrt3.c	$\mathbf{F}$	9														Х			$\times$	$\times$
24	sqrt4.c	$\mathbf{F}$	190	×	×		×				×	×				$\times$	Х	×	×	$\times$	$\times$
25	<pre>glibc/soft-fp/sqrtsf2.c</pre>	R	$\sim 500$	×	X	×	$\times$				$\times$	$\times$				$\times$	$\times$	$\times$	$\times$	×	×

There are different types of granularity in this list, ranging from only a fragment of a few isolated lines of code (in which case for example a Hoare triple for that fragment might be proved) to a full multi-file development that spans multiple directories. The four types are indicated with the following letters:

- C code fragment
- F one or more functions
- S single source file
- R real life code, multiple source files

The numbers of lines for the GNU programs are only rough indications. There are *many* files (including many header files) involved in these programs, and we did not establish exactly which files are needed for these programs. (In the list we only list the primary source file, but the verification should include all other relevant source files.) Also, the exact line counts will be dependent on the version of the software. We do not want to restrict the version that should be used for this challenge, but while developing this benchmark we looked at version 8.27 of coreutils, and version 2.24.90 of glibc:

### https://github.com/coreutils/coreutils/tree/v8.27 https://github.com/bminor/glibc/tree/glibc-2.24.90

The listings in this paper only show the relevant parts of the programs. These are the parts that should be verified. However, in case it is useful to take other parts of these programs into account, that is of course allowed. Each of these programs (apart from two of the three GNU ones) is a program that can be compiled on its own, so each program has a main function in its file too.

# 2 Factorial

# 2.1 A simple for loop

The very first program fragment is a for loop, multiplying the numbers from 1 to n:

```
6 f = 1;
for (i = 1; i <= n; i++)
8 f = f*i;
```

#### Program 1: fac1.c

This program has a straight-forward counterpart in most programming languages, so it should be easy to model it in any system. The declarations of f, i and n are not part of this fragment, but of course should be taken into account when verifying this.

In practice, this program will exhibit integer overflow when n is too large. With 32-bit integers this happens for  $n \ge 13$ . The exact way this is specified in the pre-condition is left open. For example both n < 13 and  $n! \le INT_MAX$  might be reasonable choices for this.

# 2.2 Turing's program

8

The second program fragment is very similar to the previous one, but is a bit more involved. This is a C rendering of a program by Alan Turing from his 1949 paper [21] on program verification:

```
for (u = r = 1; v = u, r < n; r++)
for (s = 1; u += v, s++ < r; )
.</pre>
```

Program 2: fac2.c

This calculates factorial without using multiplication. The original version of this program was written for the Manchester Mark 1 computer [22], but as far as we know no C compiler for that computer exists.

# 2.3 A function

We now wrap the code in a function. This time the verification has to be about pre- and post-conditions of C functions, not of isolated statements:

```
int
fac(int n)
{
    int f = 1;
    while (n)
    f *= n--;
    return f;
}
```

#### Program 3: fac3.c

We used more 'C-like' idiom in this example. For example, the update of the counter and the multiplication now happen together in a single statement, making this example slightly less trivial.

# 2.4 Recursion

5

We claimed that the for loop has a counterpart in most programming languages, but this might not hold for functional languages. Here is the obvious recursive definition of the factorial function in C:

```
int
fac(int n)
{
    return n == 0 ? 1 : n*fac(n - 1);
}
```

Program 4: fac4.c

When called with too large an n this program can overflow its stack (for example when compiled with gcc -00 and called with 1000000 this happens). The C standard does not address that issue, but the verification of this code might take this into account in its pre-condition.

# 2.5 Big numbers

Until now the output of the code is meaningless for most inputs because of overflow. The program might as well just index in a small array to get the appropriate output. Here is a version of the program that uses the GMP library [10] to compute factorial for larger inputs:

```
1 #include <stdio.h>
#include <stdlib.h>
#include "gmp.h"
5 int
```

```
main(int argc, char **argv)
     {
       int n, i;
       mpz_t f;
10
       n = atoi(argv[1]);
       mpz_init(f);
       mpz_set_ui(f, 1);
       for (i = 1; i <= n ; i++)</pre>
         mpz_mul_ui(f, f, i);
15
       mpz_out_str(stdout, 10, f);
       putchar('\n');
       mpz_clear(f);
       return 0;
20
     }
```

Program 5: fac5.c

A specification of this program should also contain appropriate specifications for the GMP types and functions being used here. This program computes the factorial of one million (a number of 5,565,709 decimal digits) in a few minutes.

# 2.6 Big numbers without a library

Here is a simple implementation of factorial with big number arithmetic without using an external library:

```
#include <stdio.h>
     #include <stdlib.h>
    #include <stdint.h>
     #define M 100000000
 5
     #define D 9
     int h = 0;
10
     int
     calc_fac(uint32_t *a, int L, uint64_t n)
     {
       uint64_t b, c;
       int i, l;
15
       a[0] = 1;
       for (i = 1; i < L; i++)</pre>
        a[i] = 0;
       1 = 1;
       while (n) {
20
         c = 0;
         for (i = 0; i < 1 || c; i++) {
           b = n*a[i] + c;
           c = b/M;
           a[i] = b - c*M;
25
         }
         if (i > 1)
           l = i;
```

```
}
30
      return 1;
    }
    void
    print_digits(int n, int c)
35
    {
      int d;
      if (c <= 0)
40
        return;
      print_digits(n/10, c - 1);
      d = n\%10;
      if (h || d) {
        putchar('0' + d);
45
         h = 1;
      }
    }
    int
    main(int argc, char **argv)
50
    {
      uint32_t *a;
      int n, m, L, i;
55
      if (argc != 2) {
         fprintf(stderr, "usage: %s input\n", *argv);
         exit(2);
      }
      n = atoi(argv[1]);
      if (n < 0) {
60
         fprintf(stderr, "%s: input is negative\n", *argv);
         exit(2);
       }
      L = 0;
      for (m = n; m; m = m/10)
65
        L++;
      L = L*n/D + 1;
       a = malloc(L*sizeof(*a));
       if (!a) {
         fprintf(stderr, "%s: out of memory\n", *argv);
70
         exit(1);
      }
      L = calc_fac(a, L, n);
      for (i = L - 1; i >= 0; i--) {
75
         print_digits(a[i], D);
       }
      putchar(' n');
       return 0;
    }
79
```

n--;

### Program 6: fac6.c

The output of the digits does not use printf, but prints the digits individually using putchar.

However, for the error handling fprintf is used.

This program is about ten times slower than the GMP version.

# 3 Cat

## 3.1 The simplest input/output program

Input/output is underrated in C verification. For example, in the  $CH_2O$  semantics from Robbert Krebbers' very thorough PhD thesis [17] this is not properly modelled. However, we think that input/output (the way that software interacts with its environment) is a fundamental property of a program and should be treated as such.

The simplest program with both input and output just copies its input to its output. In Unix this program is called **cat**, and the simplest C implementation of this behavior is:

```
while ((c = getchar()) != EOF)
putchar(c);
```

Program 7: cat1.c

How the functional behavior of this program should be specified is not addressed in this paper.

### 3.2 Buffering

In practice the previous program is slow, even though getchar internally uses a buffer. By putting an explicit buffer in the program, it gets as fast as the GNU implementation of cat:

```
1 #include <stdio.h>
char buf[128*1024];
5 int
main()
{
    size_t count;
10 while ((count = fread(buf, 1, sizeof(buf), stdin)) > 0)
    fwrite(buf, 1, count, stdout);
    return 0;
13 }
```

This program uses the size\_t type, because that is what the fread function returns and the fwrite function expects. Therefore the verification framework should know about this type, in order to be able to verify this example.

This code is similar to the example program from [12].

# 3.3 A simplified version of the real program

Here is a simplified version of the GNU implementation of cat. The main missing functionality is that there are no command line options ('flags'). Also, we simplified the code quite a bit:

Program 8: cat2.c

```
#include <stdio.h>
1
    #include <string.h>
    #include <stdlib.h>
    #include <stdint.h>
5
    #include <unistd.h>
    #include <fcntl.h>
    char *program_name, *infile;
10
    int input_desc, ok = 1;
    void
    error(int status, int errnum, char *message)
    {
15
      fprintf(stderr, "%s: ", program_name);
      if (errnum)
        perror(message);
       else
        fprintf(stderr, "%s\n", message);
      fflush(stderr);
20
      if (status)
         exit(status);
      ok = 0;
    }
25
    int
    main(int argc, char **argv)
    ł
      size_t insize = 128*1024, page_size = getpagesize(), count, n_rw;
30
       char *inbuf, *buf, *ptr;
      int argind = 0;
       program_name = argv[argind++];
       infile = "-";
       inbuf = malloc(insize + page_size - 1);
35
       if (!inbuf)
        error(EXIT_FAILURE, 0, "memory exhausted");
       buf = inbuf + page_size - 1;
       buf = buf - (uintptr_t) buf % page_size;
40
       do {
         if (argind < argc)</pre>
           infile = argv[argind];
         if (*infile == '-' && infile[1])
           error(0, 0, "options not implemented");
45
         input_desc =
           strcmp(infile, "-") == 0 ? STDIN_FILENO : open(infile, O_RDONLY);
         if (input_desc < 0) {</pre>
           error(0, 1, infile);
           continue;
         }
50
         while (1) {
           count = read(input_desc, buf, insize);
           if (count == (size_t) -1) {
             error(0, 1, infile);
```

```
break;
55
            }
            if (count == 0)
             break:
           ptr = buf;
            while (count > 0) {
60
             n_rw = write(STDOUT_FILENO, ptr, count);
              if (n_rw == (size_t) -1)
                break;
              if (n_rw == 0)
65
                break:
              count -= n_rw;
             ptr += n_rw;
            }
            if (count != 0)
              error(EXIT_FAILURE, 1, "write error");
70
         }
         if (strcmp(infile, "-") != 0 && close(input_desc) < 0)</pre>
            error(0, 1, infile);
       } while (++argind < argc);</pre>
       return ok ? EXIT_SUCCESS : EXIT_FAILURE;
75
     }
76
```

Program 9: cat3.c

This code does not use the standard C library functions fread and fwrite, but instead uses the basic Unix system calls read and write. Therefore, the code includes the unistd.h header file, and should be verified with a respect to a specification of a Unix environment.

### 3.4 The GNU version

This is the real program that the previous program was a simplified version of. When on a Linux system one runs cat, this is the code that is executed.

The file coreutils-8.27/src/cat.c contains 768 lines of code, but that code includes eight GNU header files, links to a function for handling command line options, and uses the GNU function xmalloc, a wrapper around the C standard library function malloc which is defined in a different file.

# 4 Allocate memory

## 4.1 The simplest allocator

Here is the simplest memory allocator for C possible. It just cuts off pieces of memory and does not support free:

```
1 #define HEAPSIZE 1000
unsigned char heap[HEAPSIZE];
unsigned char *limit = heap;
5
unsigned char *
malloc(int n)
{
```

```
unsigned char *ptr;

if (n <= 0 || n > heap + HEAPSIZE - limit)

return 0;

ptr = limit;

limit += n;

15 return ptr;

16 }
```

### Program 11: malloc1.c

This code does not satisfy the alignment restrictions on the real C malloc. All that is guaranteed is that a block with the appropriate number of bytes is returned.

This code is very similar to the memory allocator verified in [3, 4].

### 4.2 A typed allocator

If one does not need arrays, a memory allocator for C is simple to write. Here is an allocator that allocates objects of a given type T. (The definition of this type is intentionally not given in the code fragment shown, as it should be irrelevant for the verification.)

The code both has an array (to allocate from) and a linked list (for keeping track of the freed objects), which makes it a nice target for separation logic:

```
6 #define HEAPSIZE 1000
```

```
union cell { T allocated; union cell *next_free; } heap[HEAPSIZE];
    union cell *limit = heap, *first_free = 0;
10
    T *
    T_alloc()
    {
       union cell *ptr = first_free;
15
       if (ptr)
        first_free = first_free->next_free;
       else {
         if (limit >= heap + HEAPSIZE)
           return 0;
20
         ptr = limit++;
       }
       return &(ptr->allocated);
    }
25
    void
    T_free(T *T_ptr)
    {
       union cell *ptr = (union cell *) T_ptr;
30
       ptr->next_free = first_free;
       first_free = ptr;
    }
33
```

Program 12: malloc2.c

### 4.3 A typed allocator that does not need pointer casts

The previous program needs to get a pointer to a union cell from a pointer to the allocated T object that is in that cell. According to the C standard [14] this is allowed, as it says in 6.7.2.1/15:

[...] A pointer to a union object, suitably converted, points to each of its members (or if a member is a bit-field, then to the unit in which it resides), and vice versa. [...]

However, in some formal versions of the C semantics [18] this is not allowed. Therefore, here is a variant of this code where the user needs to keep track of the union cell for him- or herself, and that problem is not present:

```
#define HEAPSIZE 1000
 6
     union cell { T allocated; union cell *next_free; } heap[HEAPSIZE];
     union cell *limit = heap, *first_free = 0;
10
     union cell *
     T_alloc()
     {
       union cell *ptr = first_free;
15
       if (ptr)
        first_free = first_free->next_free;
       else {
         if (limit >= heap + HEAPSIZE)
20
           return 0;
         ptr = limit++;
       7
       return ptr;
     }
25
     void
     T_free(union cell *ptr)
     {
       ptr->next_free = first_free;
       first_free = ptr;
30
     }
31
```

Program 13: malloc3.c

### 4.4 A simplified version of the real program

Here is a very much simplified version of the GNU implementation of malloc and free in glibc:

```
struct malloc_chunk {
15
      size_t mchunk_prev_size;
       size_t mchunk_size;
      struct malloc_chunk *fd;
      struct malloc_chunk *bk;
      struct malloc_chunk *fd_nextsize;
20
      struct malloc_chunk *bk_nextsize;
    };
    typedef struct malloc_chunk *mchunkptr;
    typedef struct malloc_chunk *mbinptr;
25
    mchunkptr av_top;
    struct malloc_chunk av_bin;
    #define chunk2mem(p) ((void*) ((char*) (p) + 2*SIZE_SZ))
30
    #define mem2chunk(mem) ((mchunkptr) ((char*) (mem) - 2*SIZE_SZ))
     #define SIZE_SZ (sizeof(size_t))
    #define MALLOC_ALIGNMENT (2*SIZE_SZ < _Alignof(max_align_t) \</pre>
       ? _Alignof(max_align_t) : 2*SIZE_SZ)
35
     #define MALLOC_ALIGN_MASK (MALLOC_ALIGNMENT - 1)
     #define MIN_CHUNK_SIZE (sizeof(struct malloc_chunk))
     #define MINSIZE (unsigned long) \
       (((MIN_CHUNK_SIZE + MALLOC_ALIGN_MASK) & ~MALLOC_ALIGN_MASK))
40
     #define PREV_INUSE 0x1
    #define SIZE_BITS (PREV_INUSE)
    #define chunksize(p) (chunksize_nomask(p) & ~(SIZE_BITS))
45
    #define prev_inuse(p) ((p)->mchunk_size & PREV_INUSE)
    #define chunksize_nomask(p) ((p)->mchunk_size)
    #define prev_size(p) ((p)->mchunk_prev_size)
    #define chunk_at_offset(p, s) ((mchunkptr) (((char *) (p)) + (s)))
50
    #define set_head(p, s) ((p)->mchunk_size = (s))
    #define set_foot(p, s) \
       (((mchunkptr) ((char *) (p) + (s)))->mchunk_prev_size = (s))
    #define inuse_bit_at_offset(p, s) \
55
       (((mchunkptr) (((char *) (p)) + (s)))->mchunk_size & PREV_INUSE)
    #define set_inuse_bit_at_offset(p, s) \
       (((mchunkptr) (((char *) (p)) + (s)))->mchunk_size |= PREV_INUSE)
    #define clear_inuse_bit_at_offset(p, s) \
       (((mchunkptr) (((char *) (p)) + (s)))->mchunk_size &= ~PREV_INUSE)
60
    #define unlink(P, BK, FD) \
       { FD = P->fd; BK = P->bk; FD->bk = BK; BK->fd = FD; }
    static void
65
    free_unsorted_chunk(mchunkptr p, size_t size)
    ł
      mchunkptr fwd, bck;
       set_head(p, size | PREV_INUSE);
70
       set_foot(p, size);
```

```
bck = &av_bin;
       fwd = bck->fd;
       if (fwd != bck) {
         size |= PREV_INUSE;
75
         while ((unsigned long) size < chunksize_nomask(fwd))</pre>
           fwd = fwd->fd_nextsize;
         if ((unsigned long) size ==
              (unsigned long) chunksize_nomask(fwd))
         {
80
           p->fd_nextsize = p->bk_nextsize = NULL;
           fwd = fwd->fd;
         } else {
           p->fd_nextsize = fwd;
           p->bk_nextsize = fwd->bk_nextsize;
85
           fwd->bk_nextsize = p;
           p->bk_nextsize->fd_nextsize = p;
         }
         bck = fwd->bk;
       } else
90
         p->fd_nextsize = p->bk_nextsize = p;
       p->bk = bck;
       p->fd = fwd;
       fwd->bk = p;
       bck \rightarrow fd = p;
95
     }
     static void
     malloc_init_state()
100
     {
       size_t size;
       av_bin.fd = av_bin.bk = &av_bin;
       size = (HEAPSIZE - MIN_CHUNK_SIZE + 2*SIZE_SZ) & ~MALLOC_ALIGN_MASK;
105
       size -= MALLOC_ALIGNMENT;
       av_top = (mchunkptr) (heap.bytes + MALLOC_ALIGNMENT - 2*SIZE_SZ);
       set_inuse_bit_at_offset(av_top, size);
       free_unsorted_chunk(av_top, size);
     }
110
     void *
     malloc(size_t bytes)
     {
       size_t nb, size;
       mbinptr bin;
115
       mchunkptr victim, remainder, fwd, bck;
       unsigned long remainder_size;
       nb = bytes + SIZE_SZ + MALLOC_ALIGN_MASK;
       nb = nb < MINSIZE ? MINSIZE : nb & ~MALLOC_ALIGN_MASK;</pre>
120
       if (!(bin = &av_bin))
         malloc_init_state();
       if ((victim = bin->fd) != bin
           && (unsigned long) chunksize_nomask(victim)
               >= (unsigned long) (nb)) {
125
         victim = victim->bk_nextsize;
```

```
while (((unsigned long) (chunksize(victim))
                < (unsigned long) (nb)))
           victim = victim->bk_nextsize;
         if (victim != bin->bk
130
             && chunksize_nomask(victim) == chunksize_nomask(victim->fd))
           victim = victim->fd;
         size = chunksize(victim);
         remainder_size = size - nb;
         unlink(victim, bck, fwd);
135
         if (remainder_size < MINSIZE)</pre>
           set_inuse_bit_at_offset(victim, size);
         else {
           remainder = chunk_at_offset(victim, nb);
           set_head(victim, nb | PREV_INUSE);
140
           free_unsorted_chunk(remainder, remainder_size);
         }
         return chunk2mem(victim);
       }
       victim = av_top;
145
       size = chunksize(victim);
       if ((unsigned long) (size) >= (unsigned long) (nb + MINSIZE)) {
         remainder_size = size - nb;
         remainder = chunk_at_offset(victim, nb);
         av_top = remainder;
150
         set_head(victim, nb | PREV_INUSE);
         set_head(remainder, remainder_size | PREV_INUSE);
         return chunk2mem(victim);
       }
155
       return 0;
     }
     void
     free(void *mem)
160
     {
       mchunkptr p, nextchunk, bck, fwd;
       size_t size, nextsize, prevsize;
       p = mem2chunk(mem);
       size = chunksize(p);
165
       nextchunk = chunk_at_offset(p, size);
       nextsize = chunksize(nextchunk);
       if (!prev_inuse(p)) {
         prevsize = prev_size(p);
170
         size += prevsize;
         p = chunk_at_offset(p, -((long) prevsize));
         unlink(p, bck, fwd);
       }
       if (nextchunk != av_top) {
175
         if (!inuse_bit_at_offset(nextchunk, nextsize)) {
           unlink(nextchunk, bck, fwd);
           size += nextsize;
         } else
           clear_inuse_bit_at_offset(nextchunk, 0);
180
         free_unsorted_chunk(p, size);
```

```
} else {
    size += nextsize;
    set_head(p, size | PREV_INUSE);
    av_top = p;
    }
    }
```

#### Program 14: malloc4.c

Some features of the real GNU malloc that are missing from this version:

- $\bullet\,$  no fast bins
- just a single regular bin
- no bitmap to quickly find non-empty bins (not needed, as there is just one bin)
- no sbrk to get more memory (just one static array)
- memory never returned to the system (the av\_top pointer never decreases)
- no memory mapped blocks
- no thread-awareness

However, the free list in the one bin is still kept in sorted order, and still uses the fd\_nextsize and bk\_nextsize pointers to search for an appropriate block.

### 4.5 The GNU version

And again, the last program of this section is the full GNU implementation of malloc and free in glibc. The main file glibc/malloc/malloc.c in version 2.24.90 of glibc has 5,289 lines of code, but this file depends on many other files.

The file contains the sentence:

The implementation is in straight, hand-tuned ANSI C.

Exactly what it means to say that this code is 'ANSI C', and whether this is actually true, is an interesting question.

# 5 Quicksort

## 5.1 Straight-forward quicksort

Here is a very straight-forward version of a quicksort function for an array of doubles:

```
4 double a[N];
5 void
quicksort(int m, int n)
{
int i, j;
10 double pivot, tmp;
if (m < n) {</pre>
```

```
pivot = a[n];
          i = m; j = n;
          while (i <= j) {</pre>
15
            while (a[i] < pivot) i++;</pre>
            while (a[j] > pivot) j--;
            if (i <= j) {
              tmp = a[i]; a[i] = a[j]; a[j] = tmp;
20
              i++; j--;
            }
         }
          quicksort(m, j);
          quicksort(i, n);
       7
25
     }
26
```

#### Program 16: qsort1.c

This function of course has to be called with quicksort(0, N - 1).

We tried to make this code as simple as possible. For example it contains no return or break and only has a single swap.

There is a variant of this in which there only is a single comparison between i and j (with the bounds in the recursive call different), but that really makes use of the fact that we used a[n] for the pivot. For that reason we decided not to take that variant here, even though it is still shorter. The version here also works with different choices for the pivot, like taking the median or mean of some elements in the array.

This program has the property that if the array is already sorted, then it needs stack space for the recursion proportional to the length of the array. On a Linux system with a (default) stacksize of 8MB, the full program, which has N equal to 666,666, fails after about three minutes with the error:

#### segmentation fault

It would be nice if the verification of this program would be compatible with this behavior, i.e., that it would not be possible to prove this program to be correct for too large N.

# 5.2 Hoare's version of quicksort

Here is a variation on the previous program. This is a C rendering of the programs in the original publication on quicksort by Tony Hoare [13]:

```
17
     void
     partition(double *a, int m, int n, int *_i, int *_j)
     {
       double x;
20
       int f, i, j;
       f = random(m, n);
       x = a[f];
       i = m;
25
       j = n;
     up:
       for (i = i; i <= n; i++)</pre>
         if (x < a[i])
30
           goto down;
```

```
i = n;
     down:
       for (j = j; j >= m; j--)
         if (a[j] < x)
           goto change;
35
       j = m;
     change:
       if (i < j) {</pre>
         exchange(&a[i], &a[j]);
         i = i + 1;
40
         j = j - 1;
         goto up;
       } else
       if (i < f) {</pre>
         exchange(&a[i], &a[f]);
45
         i = i + 1;
       } else
       if (f < j) {
         exchange(&a[f], &a[j]);
50
         j = j - 1;
       7
       *_i = i;
       *_j = j;
     }
55
     void
     quicksort(double *a, int m, int n)
     {
       int i, j;
60
       if (m < n) {
         partition(a, m, n, &i, &j);
         quicksort(a, m, j);
         quicksort(a, i, n);
65
       }
     }
66
```

#### Program 17: qsort2.c

This program uses goto instead of break. The original Algol 60 version of this code uses the call-by-name feature of Algol 60 (which is more like macro expansion), which is modeled here by giving the partition function pointers to the variables to be modified.

### 5.3 The algorithm in the real program

If we take the GNU implementation of quicksort, and simplify it and also make it only apply to arrays of doubles, this is what one gets:

```
#define SWAP(a, b) do { tmp = *a; *a = *b; *b = tmp; } while(0)
void
quicksort(double *base_ptr, int total_elems)
{
    double tmp, *lo, *hi, *left_ptr, *right_ptr, *mid;
```

```
if (total_elems == 0)
         return:
       lo = base_ptr;
10
       hi = lo + total_elems - 1;
       while (lo < hi) {</pre>
         mid = lo + ((hi - lo) >> 1);
         if (*mid < *lo)</pre>
            SWAP(mid, lo);
15
         if (*hi < *mid) {</pre>
            SWAP(mid, hi);
            if (*mid < *lo)</pre>
              SWAP(mid, lo);
         }
20
         left_ptr = lo + 1;
         right_ptr = hi - 1;
         do {
            while (*left_ptr < *mid)</pre>
              left_ptr++;
25
            while (*mid < *right_ptr)</pre>
              right_ptr--;
            if (left_ptr < right_ptr) {</pre>
              SWAP(left_ptr, right_ptr);
              if (mid == left_ptr)
30
                mid = right_ptr;
              else if (mid == right_ptr)
                mid = left_ptr;
              left_ptr++;
35
              right_ptr--;
            } else if (left_ptr == right_ptr) {
              left_ptr++;
              right_ptr--;
              break;
           }
40
         } while (left_ptr <= right_ptr);</pre>
         if (right_ptr - lo > hi - left_ptr) {
            quicksort(left_ptr, hi - left_ptr + 1);
           hi = right_ptr;
         } else {
45
            quicksort(lo, right_ptr - lo + 1);
            lo = left_ptr;
         }
       }
     }
50
```

Program 18: qsort3.c

This is just a variation on the previous programs, but it has a few nice features, like using a macro for swapping elements in the array. One probably wants to prove this *macro* correct, which means that just verifying the macro-expanded version of the program will not be too attractive in this case.

Apart from the fact that this function only sorts doubles instead of sorting arrays with elements of an arbitrary type, it omits two optimizations from the GNU version of this code. This uses insertion sort when the size of the subarray drops below a certain threshold, and it manages a (short) explicit stack for the recursive calls, instead of using the recursion of C.

### 5.4 Quicksort on arbitrary types

Here is a version of the simple version of quicksort that we gave in 5.1, but this time for arbitrary types, so this is an implementation of the interface of the qsort function in the C standard library:

```
#include <string.h>
     #define a(i) (((unsigned char *) base) + i*size)
    void
    qsort(void *base, size_t nmemb, size_t size,
       int (*compar)(const void *, const void *))
     {
       int m, n, i, j;
       unsigned char pivot[size], tmp[size];
10
       m = 0;
       n = nmemb - 1;
       if (m < n) {
         memcpy(pivot, a(n), size);
15
         i = m; j = n;
         while (i <= j) {
           while ((*compar)(a(i), pivot) < 0) i++;</pre>
           while ((*compar)(a(j), pivot) > 0) j--;
           if (i <= j) {
20
             memcpy(tmp, a(i), size);
             memcpy(a(i), a(j), size);
             memcpy(a(j), tmp, size);
             i++; j--;
           }
25
         }
         qsort(a(m), j - m + 1, size, compar);
         qsort(a(i), n - i + 1, size, compar);
       }
    }
30
```

Program 19: qsort4.c

Notable features of this program are function pointers, (void \*) pointers, and the use of memcpy to copy C objects.

## 5.5 Parallel quicksort

Finally, here is a version of quicksort that uses threads. The basic idea of this program is that at most THREADS threads will be running. Once a recursive call gets a piece of the array of which the size is below THRESHOLD (which means the data will fit in the cache of one of the cores), a separate thread will be used for this if one is available. The count of free threads is kept in the variable available for which a mutex is used when updating it.

We did not use C11 threads for this, because we did not have a good implementation of it available. Instead this uses the pthreads library. The means the program might need to be compiled with -lpthread.

Again, algorithmically this is just the simple algorithm from 5.1:

```
#include <stdio.h>
     #include <stdlib.h>
     #include <pthread.h>
     #define N 666666
5
     double a[N];
     #define THRESHOLD 0x2000
     #define THREADS 16
10
     pthread_mutex_t mutex;
     int active = 1;
     void quicksort(int, int, int);
    struct quicksort_args {int m; int n;};
15
     void *
     run_quicksort(void *void_args)
     {
20
       struct quicksort_args *args = void_args;
       quicksort(args->m, args->n, 1);
      return 0;
     }
25
     void
     quicksort(int m, int n, int threaded)
     {
       int i, j;
       double pivot, tmp;
30
       pthread_t thread;
       struct quicksort_args args;
       if (m < n) {
        pivot = a[n];
35
         i = m; j = n;
         while (i <= j) {</pre>
           while (a[i] < pivot) i++;</pre>
           while (a[j] > pivot) j--;
40
           if (i <= j) {
             tmp = a[i]; a[i] = a[j]; a[j] = tmp;
             i++; j--;
           }
         }
         if (threaded) {
45
           if (n - m < THRESHOLD)</pre>
             threaded = 0;
           else {
             pthread_mutex_lock(&mutex);
50
             if (active < THREADS) {</pre>
               active++;
               pthread_mutex_unlock(&mutex);
               args.m = m;
               args.n = j;
               pthread_create(&thread, 0, &run_quicksort, &args);
55
```

```
quicksort(i, n, 1);
                pthread_mutex_lock(&mutex);
                active--;
                pthread_mutex_unlock(&mutex);
               pthread_join(thread, 0);
60
               return;
             }
             pthread_mutex_unlock(&mutex);
           }
         }
65
         quicksort(m, j, threaded);
         quicksort(i, n, threaded);
       }
     }
70
     int
     main()
     {
       int i;
75
       pthread_mutex_init(&mutex, 0);
       for (i = 0; i < N; i++)</pre>
         a[i] = random();
       quicksort(0, N - 1, 1);
       for (i = 0; i < N; i++)</pre>
80
         printf("%f\n", a[i]);
       return 0;
     }
83
```

Program 20: qsort5.c

# 6 Square root

# 6.1 Newton's approximation

The last class of examples in this benchmark are implementations of the square root function on single precision floating point numbers. Here is a simple version of an implementation that uses the Newton method:

```
float
    sqrt_newton(float x)
    {
      float y, z;
5
       if (x <= 0)
         return 0;
      y = x >= 1 ? x : 1;
       do {
10
         z = y;
         y = (z + x/z)/2;
      } while (y < z);
       return y;
    }
14
```

This will not always give exactly the right answer according to the IEEE 754 standard, but it will be close. If a float has finitely many different representations (which one might expect the C standard would enforce, because there are only finitely many different object representations), then this program will always terminate.

It is interesting exactly what is provable about this program in a C implementation where one does not know how the floating point implementation works in detail. Maybe if there are some guarantees about relative error bounds of the arithmetic operations, an error bound on the result of this function can be proved.

#### 6.2 An approximation of square root by bit shifting

The IEEE 754 representation of floating point numbers allows for an approximation to the square root function that can be calculated extremely fast. The counterpart for the inverse square root function (both square root and inverse square root use exactly the same trick) is known as 'the 0x5f3759df method', after the magic constant that appears in it. (Comments in the first program that had this constant in it were *evil floating point bit level hacking* and *what the fuck?*)

The implementation just consists of considering the bits of the number by casting it to an integer, shifting those bits one place to the right, and then adding a magic constant to that integer representation:

```
1 #define BITS(x) (*((unsigned *) &xx))
#define FLOAT(x) (*((float *) &xx))

5 sqrt_approx(float x)
{
    float one = 1;
    unsigned i;
10 i = (BITS(x) >> 1) + BITS(one) - (BITS(one) >> 1);
    return FLOAT(i);
12 }
```

#### Program 22: sqrt2.c

It is surprising that this gives a sensible result. One is shifting one of the bits of the exponent into a bit of the mantissa, how is that going to make sense? However, when analyzing it more closely [5], it turns out that this *does* make sense, and that for sufficiently large inputs, it gives a continuous piecewise linear function that approximates the square root function quite well (it is exact on powers of four, see the graph in Fig. 1.) The maximum error occurs when the argument is twice a power of four, and the relative error on those values is  $(1\frac{1}{2} - \sqrt{2})/\sqrt{2} = \frac{3}{4}\sqrt{2} - 1 \approx 6.0660\% \approx 2^{-4.0431}$ , which means that generally the first three bits of the fractional part will be correct.

Unlike the code in the next section, this function also makes sense if floats do not have 32 bits single precision. All that is needed is that the bits are laid out in the style of a IEEE 754 representation. As was noted before, this can be established by the implementation-defined property of \_\_STDC\_IEC\_559\_\_ having an appropriate definition.

We had some discussion whether to include this program in our benchmark. We wanted our examples to be 'realistic', and this seems more of a cute hack than something resembling realistic



Figure 1: The graph of sqrt\_approx.

code. Also, nowadays square root is supported by hardware, and getting a good approximation as a start for Newton's method is not really important anymore. However, this program is about the bit representation of floating point numbers, and about using casts to reinterpret one type as another (which is a common thing in C). For this reason we decided to still include this as one of our examples.

# 6.3 An approximation of square root by bit shifting that does not violate effective types

The previous program uses type casts to get at the bits of a floating point number, in order to be able to shift them. However, that violates the *effective type* rules, as described in 6.5/7 of the C11 standard. In fact, when compiled using GCC with the -Wall and -O3 flags, the compiler emits the warning:

```
warning: dereferencing type-punned pointer will break strict-
aliasing rules [-Wstrict-aliasing]
```

We now give a version of the same program that uses a union to get at the bits of the float, which means it does not suffer from this problem:

```
1 float
sqrt_approx(float x)
{
        union { float x; unsigned i; } u;
5
        u.x = x;
        u.i = (u.i >> 1) + 0x1fc00000;
        return u.x;
```

#### Program 23: sqrt3.c

Using a union to get at the bits of an object is allowed by the C11 standard, as in footnote 95 of 6.5.2.3/3 it reads:

If the member used to read the contents of a union object is not the same as the member last used to store a value in the object, the appropriate part of the object representation of the value is reinterpreted as an object representation in the new type as described in 6.2.6 (a process sometimes called "type punning"). This might be a trap representation.

When using unions, it is harder to write the constant to be added to the expression in a way that it will work whatever the number of bits in the floating point representation, and in such a way that the compiler will insert it as a constant in the code at compile time. To keep the program small, we therefore expanded this constant to its value for 32-bit floating point numbers: 0x1fc00000. This means that to work correctly, this program needs an implementation with 32-bit IEEE 754 floats.

# 6.4 A simplified version of a serious implementation

The GNU library glibc contains software implementations of the IEEE 754 operations in the subdirectory soft-fp. In the case of square root, the algorithm used is the square root analog of long division. The GNU code for this consists of a whole stack of macro definitions. When expanding all these definitions, and simplifying everything quite a bit, one gets:

```
typedef float SFtype;
                                       24
     #define FRACBITS
    #define EXPBITS
                                       8
    #define EXPBIAS
                                       127
5
                                       255
    #define EXPMAX
     #define W_TYPE
                                       unsigned long
    #define I_TYPE
                                       long
    #define W_TYPE_SIZE
                                       32
10
    #define CLS_NORMAL
                                       0
     #define CLS_ZERO
                                       1
     #define CLS_INF
                                       2
    #define CLS_NAN
                                       3
15
    union UNION
    {
       SFtype flt;
       struct STRUCT_LAYOUT
20
       ł
         unsigned frac:FRACBITS - 1;
         unsigned exp:EXPBITS;
         unsigned sign:1;
25
       } bits;
    };
    SFtype
```

9 }

```
__sqrtsf2(SFtype a)
    {
30
      SFtype r;
      union UNION RAW_flo;
      I_TYPE shift;
      I_TYPE A_c;
35
      I_TYPE A_s;
      I_TYPE A_e;
      W_TYPE A_f;
      I_TYPE R_c;
      I_TYPE R_s;
      I_TYPE R_e;
40
      W_TYPE R_f;
      W_TYPE S_f;
      W_TYPE T_f;
      W_TYPE q;
45
      RAW_flo.flt = a;
      A_s = RAW_flo.bits.sign;
      A_e = RAW_flo.bits.exp;
      A_f = RAW_flo.bits.frac;
      switch (A_e) {
50
      default:
        A_c = CLS_NORMAL;
        A_f |= ((W_TYPE) 1 << (FRACBITS - 1));
        A_f <<= 3;
55
        A_e -= EXPBIAS;
        break;
       case 0:
        if (A_f == 0)
          A_c = CLS_ZERO;
         else {
60
           shift = __builtin_clzl(A_f);
           shift -= (W_TYPE_SIZE - FRACBITS);
          A_c = CLS_NORMAL;
          A_f <<= (shift + 3);
          A_e -= (EXPBIAS - 1 + shift);
65
         }
        break;
       case EXPMAX:
         if (A_f == 0)
          A_c = CLS_INF;
70
         else
           A_c = CLS_NAN;
        break;
      }
75
       switch (A_c) {
       case CLS_NAN:
        R_c = CLS_NAN;
        R_s = A_s;
        R_f = A_f;
        break;
80
      case CLS_INF:
        if (A_s) {
```

```
R_c = CLS_NAN;
            R_s = 1;
           R_f = ((W_TYPE) 1 \iff (FRACBITS - 2));
85
         } else {
           R_c = CLS_INF;
           R_s = 0;
         }
90
         break;
       case CLS_ZERO:
         R_c = CLS_ZERO;
         R_s = A_s;
         break;
       case CLS_NORMAL:
95
         if (A_s) {
           R_c = CLS_NAN;
           R_s = 1;
           R_f = ((W_TYPE) 1 \iff (FRACBITS - 2));
           break;
100
         }
         R_c = CLS_NORMAL;
         R_s = 0;
         R_e = A_e >> 1;
105
         if (A_e & 1)
           A_f += A_f;
         S_f = 0;
         R_{f} = 0;
         q = ((W_TYPE) 1 << (3 + FRACBITS)) >> 1;
110
         while (q != ((W_TYPE) 1 << 2)) {</pre>
           T_f = S_f + q;
            if (T_f <= A_f) {
              S_f = T_f + q;
              A_f -= T_f;
115
             R_f += q;
            }
            A_f += A_f;
           q >>= 1;
         }
120
         if (A_f) {
            if (S_f < A_f)
             R_f |= ((W_TYPE) 1 << 2);</pre>
           R_f |= ((W_TYPE) 1 << 0);</pre>
         }
125
       }
       switch (R_c) {
       case CLS_NORMAL:
         R_e += EXPBIAS;
130
         if (R_e > 0) {
            if (R_f & 7) {
              if ((R_f & 15) != ((W_TYPE) 1 << 2))
                R_f += ((W_TYPE) 1 << 2);</pre>
135
            }
            if (R_f & ((W_TYPE) 1 << (3 + FRACBITS))) {</pre>
```

```
R_f &= ~((W_TYPE) 1 << (3 + FRACBITS));</pre>
              R_e++;
            }
            R_f >>= 3;
140
            if (R_e >= EXPMAX) {
              R_c = CLS_INF;
              R_e = EXPMAX;
              R_{f} = 0;
            }
145
          } else {
            R_e = -R_e + 1;
            if (R_e <= (3 + FRACBITS)) {</pre>
              R_f = (R_f >> R_e);
              if (R_f & 7) {
150
                if ((R_f & 15) != ((W_TYPE) 1 << 2))</pre>
                  R_f += ((W_TYPE) 1 << 2);</pre>
              }
              if (R_f & (((W_TYPE) 1 << (3 + FRACBITS)) >> 1)) {
155
                R_e = 1;
                R_{f} = 0;
              } else {
                R_{e} = 0;
                R_f >>= 3;
              }
160
            } else {
              R_{e} = 0;
              if (R_f != 0) {
                R_f = 1;
                if ((R_f & 7) && ((R_f & 15) != ((W_TYPE) 1 << 2)))
165
                  R_f += ((W_TYPE) 1 << 2);</pre>
                R_f >>= 3;
              }
            }
          }
170
         break;
        case CLS_ZERO:
          R_{e} = 0;
          R_{f} = 0;
175
          break;
        case CLS_INF:
          R_e = EXPMAX;
          R_f = 0;
          break;
        case CLS_NAN:
180
          R_e = EXPMAX;
          R_f |= ((W_TYPE) 1 << (FRACBITS - 2));</pre>
         break;
        }
185
       RAW_flo.bits.frac = R_f;
       RAW_flo.bits.exp = R_e;
       RAW_flo.bits.sign = R_s;
       r = RAW_flo.flt;
       return r;
```

Program 24: sqrt4.c

This code implements IEEE 754 behavior, but (as the primary simplification) only implements the default rounding mode. The full GNU version supports the other rounding modes as well. Also, the full version implements more possible variant behaviors on not-a-numbers.

Most of this code is about getting at the bits at the start and reassembling the floating point number at the end, and about handling various special cases like infinities, not-a-numbers, and denormalized numbers. The actual algorithm for the 'long division' analog is in lines 103–125.

### 6.5 The GNU version

And the last program in our verification benchmark is the actual implementation of square root in libc/soft-fp. That code is a bit obfuscated, because many of the macros implement the operations for different floating point sizes at the same time. The function that should be verified is \_\_sqrtsf2, implemented in the file sqrtsf2.c:

```
31
     #include "soft-fp.h"
     #include "single.h"
     SFtvpe
     __sqrtsf2 (SFtype a)
35
     {
       FP_DECL_EX;
       FP_DECL_S (A);
       FP_DECL_S (R);
40
       SFtype r;
       FP_INIT_ROUNDMODE;
       FP_UNPACK_S (A, a);
       FP_SQRT_S (R, A);
       FP_PACK_S (r, R);
45
       FP_HANDLE_EXCEPTIONS;
       return r;
    }
49
```

Program 25: sqrtsf2.c

Of course the macros defined through the **#include** lines pull in lots of other code from other files, and expand this program significantly.

In the latest versions of glibc the source code of square root in libc/soft-sp has been moved to a larger file that contains many arithmetical operations together. However, we prefer to have the 25th program in our benchmark be about a single small source file that we essentially can include in this paper. For this reason, this benchmark stays linked to version 2.24.90 of glibc.

# 7 Versions

The benchmark presented in this paper will go through various versions. We will use dates for versioning. The initial version is dated April 1, 2019, and might be referred to as the *April fool's* version.

190 }

Every time a problem with one of the programs in this benchmark is discovered, we will modify the program to remove the problem, and give the benchmark a new version date. In that case we will add a short description to this section of the problem, how it was discovered and how the program was changed to remove the problem. This means that this section will be some kind of high level changelog.

Note that bad behavior outside the intended domain of behavior of a program does not count as a problem. For example, integer overflow in the earlier factorial programs in Section 2 is not a problem. This should be handled by a precondition of the correctness statement of the code, and does not merit modification of the code in the benchmark.

# 8 Discussion

We can imagine different reactions to the collection of programs that we present here:

- These programs are too easy.
- These programs are too hard.

#define fac(n) EVAL3(PRIM(fac)(n))

• These programs are not too easy or too hard, but not interesting.

In the first case we would like to see how to verify them. If they are easy, showing how to do that should not be much work, right? In the second case we would like to know how far the technology currently reaches. We intentionally provided a spectrum from very small fragments to larger pieces of code, so if you have the second reaction then just show us how to do the easier ones.

However, in the third case we disagree. This is based on actual programs from the GNU project, and as such are representative of real C programming. Why is it not interesting to be able to verify what to us are characteristic C programs?

It was tempting to also include programs in our benchmark that are cute and/or dazzling. (The square root approximation by bit shifting in Section 6.2 was borderline in this respect.) For instance, consider two possible additions to the factorial section, the first being a surprising application of the preprocessor:

```
#include <stdio.h>
     /* after all the macro violence, PRIM(fac) will be left in the output;
        one solution is to define these as ... existing nowhere */
5
     extern int fac();
     extern int (*PRIM())();
     #define NOTHING
    #define PRIM(f) f##_prim
10
     /* NOTHING between PRIM and (f) prevens expansion of PRIM */
     #define REC PRIM NOTHING
15
     /* this simply forces the re-scanning of x (for macros) many times */
    #define EVAL1(x) x
    #define EVAL2(x) EVAL1(EVAL1(EVAL1(EVAL1(x))))
    #define EVAL3(x) EVAL2(EVAL2(EVAL2(EVAL2(x))))
    #define fac_prim(n) (n)<=0? 1 : (n)*(REC(fac)(n-1))</pre>
20
```

```
int main()
{
25 printf("%ld\n", fac(20L));
26 }
```

Program 26: facx.c

1	<pre>#include<stdarg.h>/*Y=\f*/</stdarg.h></pre>
	<pre>#include<stdio.h>/*.(\x.f(*/</stdio.h></pre>
	<pre>#define z return/*xx))(\x.*/</pre>
	<pre>#include<stdlib.h>/*f(xx))*/</stdlib.h></pre>
5	typedef union w w;union w{
	<pre>#define L(f,y,x,t) w _ ##\</pre>
	f(w x,){va_list argh;w\
	y; va_start(argh,x);y= *\
	<pre>(va_arg(argh,w*));va_end\</pre>
10	(argh);z(t); w f (w fw)
	{r((w){.c= &_## f}, fw);}
	<pre>int _; void*p;w(*c)(w,</pre>
	<pre>#define r(x,y)w*k=malloc\</pre>
	(sizeof(w) << 1); k[0] = x
15	; k[1]=y; z((w) {.p= k})
	<pre>#define l(f,x,t)w _##f(\</pre>
	w x,){z(t);}w f={.p\
	=(w[1]){{.c= &_##f}}};
	);};w a(w f,w x){w*d=f
20	.p;z((*d->c)(x,d+1)
	);}
	$w a_(w f, w fw) \{r(f,$
25	fw);}w _(w f){w*k=f.p,r=
	a(*k,k[1]); free(k) ;z(r);}
	L(F,f,x,(w){=x?x*a(_(f),
	(w){=x1}):1})1(F_,f,F(f)
	)L( W,(f),/*=\*/x,a(f,a_(x,x))
30	)l( $Y,/*=\*/f,a(W(f),W(f))$
	)int main(){printf("%.""f\n",
	a(a(Y,F_),(w){=10})
	/60./60/24/*d */
34	);}

# Program 27: facy.c

(it turns out that the factorial is calculated here at compile time), and the second being a price winning submission to the international obfuscated C code contest (it implements a functional language and then does the recursion by evaluating Church's Y-combinator.) However, we do not consider these two programs representative of real life C programming, so we resisted the temptation to include them in our benchmark.

There are twenty-five programs in this paper, and there are multiple frameworks for formal C verification out there, like the ones described in [1], [6], [7], [11] and [15]. We think this could be a great source for student projects. Just say, for example: 'prove the correctness of program number three in framework number one' (actually something close to that one already has been

done, see the slides of [23]). That way we will get a good comparison of what are the strengths and weaknesses of the different current approaches for formal C verification.

# 9 Related work

Benchmarks/competitions very similar to the one presented in this paper already exist:

- VACID-0 [19]: https://archive.codeplex.com/?p=vacid
- Verfified Software Competition [9, 16].
- VerifyThis Competition: http://www.pm.inf.ethz.ch/research/verifythis.html

The main difference from the benchmark presented here is that they do not target C explicitly. Yet, they sometimes use C code in their statements, and sometimes participants are using C for their solutions.

There are also benchmarks that *are* about specific programming languages, which often have a C section, like for example:

• Collection of Verification Tasks of the SoSy-Lab at the Ludwig-Maximilian University in Munich:

https://github.com/sosy-lab/sv-benchmarks/tree/master/c

However, in that case the focus is not so much on formally proving that programs are fully correct, but on benchmarking the performance of non-interactive tools for discovering problems in the programs.

# Acknowledgments

Thanks to Jean-Christophe Filliâtre, Andrew Appel and Sandrine Blazy for many helpful comments.

# References

- Andrew W. Appel. Verified software toolchain. In Proceedings of the 20th European Symposium on Programming, ESOP'11/ETAPS'11, pages 1–17, 2011.
- [2] Patrick Baudin, Pascal Cuoq, Jean-Christophe Filliâtre, Claude Marché, Benjamin Monate, Yannick Moy, and Virgile Prevosto. ACSL: ANSI/ISO C Specification Language, version 1.12, 2016. https://frama-c.com/acsl.html.
- [3] Nick Benton. Abstracting Allocation, The New new thing. In Computer Science Logic, 20th International Workshop, CSL 2006, 15th Annual Conference of the EACSL, Szeged, Hungary, September 25-29, 2006, Proceedings, pages 182–196, 2006.
- [4] Nick Benton. Machine Obstructed Proof: How many months can it take to verify 30 assembly instructions? In Workshop on Mechanizing Metatheory, 2006.
- [5] James F. Blinn. Floating-point tricks. IEEE Comput. Graph. Appl., 17(4):80–84, July 1997.

- [6] Loïc Correnson, Pascal Cuoq, Florent Kirchner, André Maroneze, Virgile Prevosto, Armand Puccetti, Julien Signoles, and Boris Yakobowski. Frama-C User Manual, 2018.
- [7] Markus Dahlweid, Michal Moskal, Thomas Santen, Stephan Tobies, and Wolfram Schulte. VCC: Contract-based modular verification of concurrent C. In 31st International Conference on Software Engineering, ICSE 2009, May 16-24, 2009, Vancouver, Canada, Companion Volume, pages 429–430. IEEE, 2009.
- [8] M.C.J.D. van Eekelen, J.H. Geuvers, J.E.W. Smetsers, and F. Wiedijk. SOVEREIGN: A Framework for Modular Formal Verification of Safety Critical Software, STW OTP-project proposal, 2014.
- [9] Jean-Christophe Filliâtre, Andrei Paskevich, and Aaron Stump. The 2nd Verified Software Competition: Experience Report. In Vladimir Klebanov, Bernhard Beckert, Armin Biere, and Geoff Sutcliffe, editors, COMPARE 2012, Comparative Empirical Evaluation of Reasoning Systems, 1st Intl. Workshop, volume 873 of CEUR Workshop Proceedings, pages 36–49, Manchester, United Kingdom, June 2012.
- [10] Torbjörn Granlund and the GMP development team. GNU MP, The GNU Multiple Precision Arithmetic Library, Edition 6.1.1, 2016.
- [11] David Greenaway, June Andronick, and Gerwin Klein. Bridging the Gap: Automatic Verified Abstraction of C. In Lennart Beringer and Amy Felty, editors, *Interactive Theorem Proving: Third International Conference, ITP 2012, Princeton, NJ, USA, August 13-15,* 2012. Proceedings, pages 99–115, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg.
- [12] Armaël Guéneau, Magnus Myreen, Ramana Kumar, and Michael Norrish. Verified characteristic formulae for CakeML. In *European Symposium on Programming*, pages 584–610, Uppsala, Sweden, April 2017.
- [13] C. A. R. Hoare. Algorithm 63: Partition and Algorithm 64: Quicksort. Communications of the ACM, 4(7):321, 1961.
- [14] ISO. ISO/IEC 9899:2011 Information technology Programming languages C. International Organization for Standardization, Geneva, Switzerland, December 2011.
- [15] Bart Jacobs, Jan Smans, Pieter Philippaerts, Frédéric Vogels, Willem Penninckx, and Frank Piessens. VeriFast: A Powerful, Sound, Predictable, Fast Verifier for C and Java. In Proceedings of the Third International Conference on NASA Formal Methods, NFM'11, pages 41–55, Berlin, Heidelberg, 2011. Springer-Verlag.
- [16] Vladimir Klebanov, Peter Müller, Natarajan Shankar, Gary T. Leavens, Valentin Wüstholz, Eyad Alkassar, Rob Arthan, Derek Bronish, Rod Chapman, Ernie Cohen, Mark Hillebrand, Bart Jacobs, K. Rustan M. Leino, Rosemary Monahan, Frank Piessens, Nadia Polikarpova, Tom Ridge, Jan Smans, Stephan Tobies, Thomas Tuerk, Mattias Ulbrich, and Benjamin Weiß. The 1st Verified Software Competition: Experience report. In Michael Butler and Wolfram Schulte, editors, Proceedings, 17th International Symposium on Formal Methods (FM 2011), volume 6664 of LNCS. Springer, 2011.
- [17] Robbert Krebbers. The C standard formalized in Coq. PhD thesis, Radboud University Nijmegen, 2015.
- [18] Robbert Krebbers. A formal C memory model for separation logic. J. Autom. Reasoning, 57(4):319–387, 2016.

- [19] K. Rustan M. Leino and Michał Moskal. VACID-0: Verification of Ample Correctness of Invariants of Data-structures, Edition 0. Technical Report KRML 209, Microsoft Research, Redmond, WA, USA, 2010.
- [20] Richard M. Stallman. The GNU Manifesto. In Joshua Gay, editor, Free Software, Free Society: Selected Essays of Richard M. Stallman, pages 33–41. GNU Press, Boston, 2002.
- [21] Alan Turing. Checking a large routine. In Report of a Conference on High Speed Automatic Calculating Machines, pages 67–69. University Mathematical Laboratory, Cambridge, 1949. http://www.turingarchive.org/browse.php/b/8.
- [22] Freek Wiedijk. A large routine. In Inge Bethke, Bert Bredeweg, and Alban Ponse, editors, Liber Amicorum voor / for Jan A. Bergstra, pages 162–163. Informatics Institute, University of Amsterdam, 2016. http://www.cs.ru.nl/~freek/pilot/janb.pdf.
- [23] Freek Wiedijk. Verifying SHA using VST, 2016. slides for the course Type Theory & Coq 2015-2016, Radboud University Nijmegen, http://www.cs.ru.nl/~freek/talks/ vst-sha.pdf.