

Subtleties of the ANSI/ISO C standard

Robbert Krebbers, Freek Wiedijk
Radboud University Nijmegen

Abstract—In our Formalin project to formalize C11 (the ANSI/ISO standard of the C programming language) we discovered many subtleties that make formalization of this standard difficult. We discuss some of these subtleties and indicate how they may be addressed in a formal C semantics.

Furthermore, we argue that the C standard does not allow Turing complete implementations, and that its evaluation semantics does not preserve typing. Finally, we claim that no strictly conforming programs exist. That is, there is no C program for which the standard can guarantee that it will not crash.

Index Terms—C programming language, programming language standardization, formal methods

I. INTRODUCTION

A. Problem

Current programming technology is rather fragile: programs regularly crash, hang, or even allow viruses to have free reign. An important reason is that a lot of programs are developed using low-level programming languages. One of the most extreme instances is the widespread use of the C programming language. In the TIOBE popularity index [23] it is (fall 2012) in the top position.

Whereas most modern programming languages require a compiler to throw an exception when exceptional behavior occurs (e.g. dereferencing a NULL pointer, integer overflow, accessing an array out of its bounds), C [11, 7] does not impose such requirements. Instead, it classifies these behaviors as *undefined* and allows a program to do literally anything in such situations [7: 3.4.3]. On the one hand, this allows a compiler to omit runtime checks and to generate more efficient code, but on the other hand these undefined behaviors often lead to security vulnerabilities [4, 14, 24].

There are two main approaches for improving this situation:

- Switch to a more modern and higher level programming language. This approach reduces the number of programming errors, and if there still is an error, the chance of it being used by an exploit is much lower.

One disadvantage of this approach is that there will be a thicker layer between the program and the hardware of the system. This costs performance, both in execution speed and in memory usage, but it also means a reduction in control over the behavior of the system. Especially for embedded systems and operating system kernels this is an undesired consequence.

- Stick to a low-level programming language like C, but add a formal methods layer on top of it to establish that programs do not exhibit undefined behavior.

Such a layer might allow the developer to annotate their programs with invariants, and to prove that these invari-

ants indeed hold. To be practical most of these invariants should be proven automatically, and the remaining ones by interactive reasoning.

This approach is an extension of *static analysis*. But whereas static analysis tools often yield false-positives, this approach allows the developer to prove that a false-positive is not an actual error.

For functional correctness, this approach has also been successful. There have been various projects to prove the C source code of a microkernel operating system correct [2, 12].

There are many tools for the second approach, like VCC [2], Verifast [10] and Frama-C [19]. However, these tools do not use an explicit formal C semantics and only implicitly ‘know’ about the semantics of C. Therefore the connection between the correctness proof and the behavior of the program when compiled with a real-world compiler is shallow. The soundness of these tools is thus questionable [6].

For this reason, we started in 2011 at the Radboud University a project to provide a formal semantics of the C programming language: the Formalin project [13]. This semantics was to be developed for interactive theorem provers, allowing one to base formal proofs on it. Although there already exist various versions of a formal semantics of significant fragments of C (see Section I-C for an overview), our goal was to formalize the ‘official’ semantics of C, as written down in the C11 standard (back then the target was C99, as C11 was not finished yet). We intended not to skim the more difficult aspects of C and to provide a semantics of the whole language.

Unfortunately, the Formalin project has turned out to be much harder than we anticipated because the C11 standard turned out to be very difficult to formalize. We were aware that C11 includes many features, so that we would need to write a large formalization to include them all. Also, since the standard is written in English, we knew we had to deal with its inherent ambiguity and incompleteness. But we had not realized how difficult things were in this respect.

Already, the very basis of our formalization, the memory model, turned out to be almost impossible to bring into line with the standard text. The reason for this is that C allows both *high-level* (by means of typed expressions) and *low-level* (by means of bit manipulation) access to the memory. The C99 and C11 standards have introduced various restrictions on the interaction between these two levels to allow compilers to make more effective non-aliasing hypotheses based on typing. As also observed in [9, 18] these restrictions have lead to unclarities and ambiguities in the standard text.

B. Approach

The aim of this paper is to discuss the situation. We describe various issues by small example programs, and discuss what the C11 standard says about them, and how a formal semantics may handle these.

During the year that the Formalin project has been running we have developed an (implicit) prototype of a C11 semantics in the form of a large Haskell program. This program can be seen as a *very* critical C interpreter. If the standard says that a program has undefined behavior, our Haskell interpreter will terminate in a state that indicates this.

The intention of our prototype was to develop a clear semantics of the high-level part. To this end, we postponed including low-level details as bytes, object representations, padding and alignment. Due to the absence of low-level details, we were able to support features that are commonly left out, or handled incorrectly, in already existing formal versions of C. In particular, we treat effective types, the common initial segment rule, indeterminate values, pointers to one past the last element, variable length arrays, and `const`-qualified objects. But even without the low-level part, we experienced many other difficulties, that are also described in this paper.

Our prototype is currently being ported to the interactive theorem prover Coq. Nonetheless, the source of the prototype can be inspected at <http://ch2o.cs.ru.nl/>.

While working on a formal version of the C11 standard, we had four rules that guided our thinking:

- 1) If the standard is absolutely clear about something, our semantics should not deviate from that. That means, if the standard clearly states that certain programs should not exhibit undefined behavior, we are not allowed to take the easy way out and let *our* version of the semantics assign undefined behavior to it.
- 2) If it is *not* clear how to read the standard, our semantics should err on the side of caution. Generally this means assigning undefined behavior as we did not want our semantics to allow one to prove that a program has a certain property, when under a different reading of the standard this property might not hold.
- 3) C idiom that is heavily used in practice should not be considered to exhibit undefined behavior, even if the standard is not completely clear about it.
- 4) If real-world C compilers like GCC and clang in AN-SI/ISO C mode exhibit behavior that is in conflict with a straightforward reading of the standard, but that can be explained by a contrived reading of the standard, our semantics should take the side of the compilers and assign undefined behavior.

Of course there is a tension between the second and third rule. Furthermore, the fourth rule is a special case of the second, but we included it to stress that compiler behavior can be taken as evidence of where the standard is unclear.

C. Related Work

This related work section consists of three parts: discussion of related work on unclarity in the C standard, discussions

of related work on undefined behavior, and a brief comparison of other versions of a formal semantics of C.

An important related document is a post by Maclaren [18] on the standard committee's mailing list where he expresses his concerns about the standard's notion of an *object* and *effective type*, and discusses their relation to multiprocessing. Like our paper, he presents various issues by considering example programs. Most importantly, he describes three directions to consistency. We will treat those in Section II.

The standard committee's website contains a list of defect reports. These reports describe issues about the standard, and after discussion by the committee, may lead to a revision or clarification of the official standard text. Defect Report #260 [9] raises similar issues as we do and will be discussed thoroughly throughout this paper.

There is also some related work on undefined behavior and its relation to bugs in both programs and compilers. Wang *et al.* [24] classified various kinds of undefined behavior and studied its consequences to real-world systems. They have shown that undefined behavior is a problem in practice and that various popular open-source projects (like the Linux kernel and PostgreSQL) use compiler workarounds for it. However, they do not treat the memory model, and non-aliasing specifically, and also do not consider how to deal with undefined behavior in a formal C semantics.

Yang *et al.* [25] developed a tool to randomly generate C programs to find compiler bugs. This tool has discovered a significant number of previously unknown bugs in state of the art compilers. In order to do this effectively, they had to minimize the number of generated programs that exhibit undefined behavior. However, they do not seem to treat the kinds of undefined behavior that we consider.

Lastly, we will briefly compare the most significant already existing formal versions of a C semantics. There are also many others like [3, 21, 15], but these only cover small fragments of C or are not recent enough to include the troublesome features of C99 and C11 that are the topic of this paper.

Norrish defined a semantics of a large part of C89 in the interactive theorem prover HOL [20]. His main focus was to precisely capture the non-determinism in evaluation of expressions and the standard's notion of *sequence points*. However, the problems described in our paper are due to more recent features of the standard than Norrish's work.

Blazy and Leroy [1] defined a semantics of a large part of C in the interactive theorem prover Coq to prove the correctness of the optimizing compiler CompCert. CompCert treats some of the issues we raise in this paper, but as its main application is to compile code for embedded systems, its developers are more interested in giving a semantics to various undefined behaviors (such as wild pointer casts) and to compile those in a faithful manner, than to support C's non-aliasing features to their full extent (private communication with Leroy).

Ellison and Rosu [5] defined an executable semantics of the C11 standard in the \mathbb{K} -framework. Although their semantics is very complete, has been thoroughly tested, and has some interesting applications, it seems infeasible to be used for

interactive theorem provers. Besides, their current memory model seems not capable of supporting the issues we present. We give more details of the discussed semantics in Section II.

D. Contribution

The contribution of this paper is fourfold:

- We indicate various subtleties of the C11 memory model and type system that we discovered while working on our formal semantics (Section II, III, IV and VII).
- We argue for various properties of C11: lack of Turing completeness (Section V), lack of programs that are guaranteed not to exhibit undefined behavior (Section VI), and lack of preservation of typing (Section VII).
- We present many small example programs that can be used as a ‘benchmark’ for comparing different formal versions of a C semantics.
- We discuss some considerations on how to best proceed with formalizing the C standard, given that the existing standard text is imprecise and maybe even inconsistent.

II. POINTER ALIASING VERSUS BIT REPRESENTATIONS

An important feature of C is to allow both *high-level* (by means of typed expressions) and *low-level* (by means of bit manipulation) access to the memory. For low-level access, the standard requires that each value is represented as a sequence of bytes [7: 3.6, 5.2.4.2.1], called the *object representation* [7: 6.2.6.1p4, 6.5.3.4p2].

In order to allow various compiler optimizations (in particular strong non-aliasing analysis), the standard has introduced various restrictions on the interaction between these two levels of access. Let us consider the following program [9]:

```
int x = 30, y = 31;
int *p = &x + 1, *q = &y;
if (memcmp(&p, &q, sizeof(p)) == 0)
    printf("%d\n", *p);
```

Here we declare two objects `x` and `y` of type `int` and use the `&`-operator to take the address of both (Figure 1a). Increasing the pointer `&x` by one moves it `sizeof(int)` bytes ahead and yields a pointer to the right edge of the `x` block. It may seem strange that such pointers are allowed at all [7: 6.5.6p8] because they cannot be dereferenced, but their use is common programming practice when looping through arrays.

We store these pointers into objects `p` and `q` of type pointer to `int` (Figure 1b). The next step is to check whether these pointers `p` and `q` are equal (note: not whether the memory they point to is equal). We do this by using the `memcmp` function, which checks whether their object representations are equal. It is important to use a bitwise comparison, instead of the ordinary `p == q`, to reveal if additional information is stored. If the object representations of the two are equal, we can conclude that both pointers point to the same memory location and do not contain conflicting bounds information. From this we are allowed to conclude that `x` and `y` are allocated adjacently (Figure 1c).

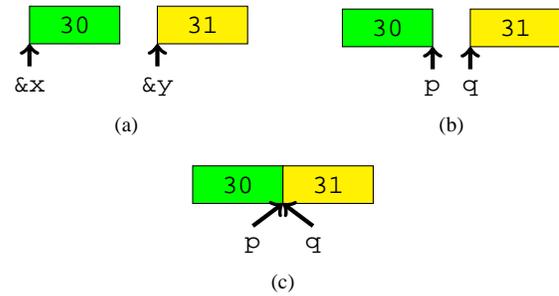


Fig. 1: Adjacent blocks.

Now we have ended up in a situation where the low- and high-level world are in conflict. On the one hand, `p` is a pointer to the edge of the `x` block, and thus dereferencing it would be undefined behavior. On the other hand, the object representation of `p` is the same as the object representation of `q`, and so `p` and `q` should behave identically.

Although the standard itself is very unclear about these problems, in Defect Report #260 [9] the committee expressed the following judgment:

Implementations are permitted to track the origins of a bit-pattern and treat those representing an indeterminate value as distinct from those representing a determined value. They may also treat pointers based on different origins as distinct even though they are bitwise identical.

Apparently a value can contain additional information about its origin that is not reflected in its object representation. The reason the committee allows this, is to allow compiler optimizations that would not be correct otherwise.

To show that this is a real issue, we changed the last example slightly and compiled it without warnings with GCC¹:

```
int y = 31, x = 30;
int *p = &x + 1, *q = &y;
if (memcmp(&p, &q, sizeof(p)) == 0) {
    *p = 10;
    printf("%d %d\n", *p, *q);
}
```

This prints two distinct values ‘10 31’. Despite the fact that `p` and `q` are identical on the bit-level (which follows from the fact that the `printf` is executed at all), they still behave differently on the object-level, as indeed Defect Report #260 allows for.

Maclaren describes similar unclarity of the standard with respect to the size of array objects corresponding to pointer values [18]. He presents three directions the standard may take:

- 1) **The original Kernighan & Ritchie approach.** Pointer values are simply addresses, and their object representations are all that matters.
- 2) **Pointers carry their origin.** Each pointer value carries a history on how it is constructed. This history should

¹Using `gcc -O1 -std=c99 -pedantic -Wall`, version 4.1.2.

be preserved by various operations on pointers. Furthermore, if the pointer is refined to a subobject (for example by array indexing, or by taking a field of a structure), it is impossible to get the original pointer back, even if they correspond to the same address in memory.

- 3) **Only visible information is relevant.** This is a weaker variant of the previous one where the history is limited to a certain visibility, for example the current scope.

The first approach is taken by the semantics of Norrish [20] in HOL. Although this approach is clearly understood, and is convenient to formalize, it implies that pointers with equal object representation always have to be treated as equal. Therefore, optimizations as for example performed by GCC and allowed by Defect Report #260 are not sound with respect to Norrish’s semantics.

The second approach allows the most compiler optimizations and is therefore the most appealing one. However, it is unclear how bit-level operations and library functions like `memcpy` should deal with a pointer’s origin. In an earlier version of the CompCert memory model of by Leroy and Blazy [17], bytes were not modeled as simple sequences of bits, but as abstract entities. Although this approach makes many bit-level operations exhibit undefined behavior, it ensures that the abstract information can easily be preserved, even by byte-level operations. In a more recent version of their memory model [16], only bytes that constitute a pointer are abstract entities, whereas those that constitute integers or floating point numbers are real sequences of bits.

The third approach requires a careful definition of ‘visibility’. It is unclear whether such a definition can be given that is both consistent and that allows sufficient compiler optimizations. This approach therefore seems not very attractive.

The memory model of our semantics takes the second approach to its fullest extent. We represent the memory as a finite map from indexes to trees that expose the full structure of values, and pointers are paths through these trees. Subobjects (subarrays, fields of structures or unions) then correspond to subtrees of the memory. The origin of pointers in our semantics is much more detailed than its counterpart in CompCert where only a pointer’s block and offset into that block is stored. In particular, since we do not flatten arrays and structures, we are able to handle subobjects correctly. For example, it allows our semantics to impose undefined behavior on the following example.

```
int a[2][2] = { {13, 21}, {34, 35} };
struct t { int *r, *p, *q; } s;
s.p = &a[0][2]; s.q = &a[1][0];
if (p == q)
    printf("%d\n", *p);
```

We do not use `memcmp` in this example, because our semantics does not support object representations. The comparison `p == q` is the source of the undefined behavior.

Figure 2 displays the representation of the memory after evaluating the first three lines of this code in both a concrete memory and our abstract memory as trees. Our semantics

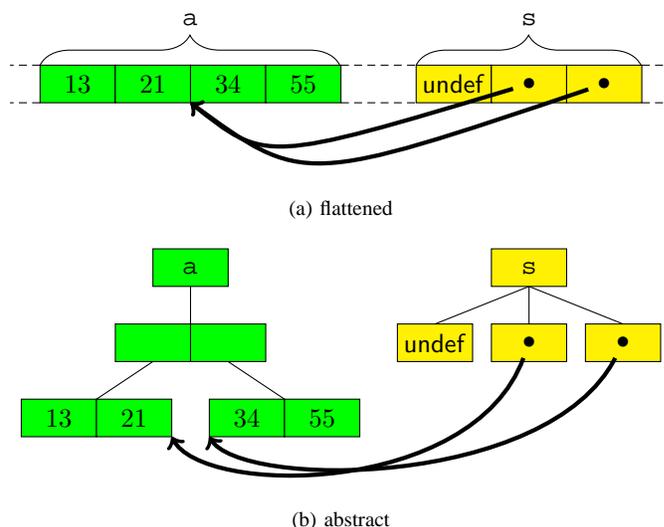


Fig. 2: Example contents of the memory.

imposes special treatment on pointers that point to elements one past the end of an object. We do not allow these pointers to be dereferenced, and they are only allowed in comparisons with other pointers to the same subobject, making the above example undefined.

The semantics by Ellison and Rosu [5] in the \mathbb{K} -framework imposes similar treatment. CompCert [16] does not allow pointers to the end of an object to be used at all. It seems not to be a fundamental limitation of their model to incorporate a similar treatment as ours and Ellison and Rosu’s.

Of course, the drawback of our current memory model is that bytes are not present at all, whereas the CompCert memory model at least allows byte-level operations on objects solely consisting of integers and floating point numbers. But since the intention of our semantics was to obtain a better understanding of the high-level part of the C memory, we do not consider this to be a problem for the moment.

III. THE COMMON INITIAL SEQUENCE

The C programming language supports various data types to be used to build more complex types. In particular, structure, union, and array types. Structures are like product types and unions are like sum types. Due to the low-level nature of C, unions are *untagged* rather than *tagged*. This means that the current variant of the union is not stored. Consider:

```
union int_or_float { int x; float y; };
```

Given an object of type `int_or_float`, it is not possible to (portably) test whether it contains the `int` or the `float` variant. This may seem unnatural, but it is in the spirit of C to let the programmer decide whether or not to store the tag.

An interesting consequence of untagged unions is the standard’s *common initial sequence* rule [7: 6.5.2.3]. Consider:

```
struct t1 { int m; float f; };
struct t2 { int m; short s; };
union { struct t1 s1; struct t2 s2; } u;
```



Fig. 3: A union containing two structures with a common initial sequence.

The representation of the object `u` might look as pictured in Figure 3. Although, due to alignment there may be additional space between the members, the standard guarantees that the integer parts always coincide. Even stronger, in this case it guarantees the following [7: 6.5.2.3p6]:

... it is permitted to inspect the common initial part of any of them anywhere that a declaration of the completed type of the union is visible.

For example, that means we are allowed to do things like:

```
int main(void) {
    u.s2.m = 20;
    printf("%d\n", u.s1.m);
}
```

So, we set the integer part via the one variant of the union, and read it out via the other. However, the following program exhibits undefined behavior as in the function body the relation to the union type is not visible.

```
int f(struct t1 *p1, struct t2 *p2) {
    p2->m = 20;
    return p1->m;
}
int main(void) {
    printf("%d\n", f(&u.s1, &u.s2));
}
```

This is to allow compilers to make additional aliasing assumptions about `p1` and `p2` because their types are different. Real compilers, like GCC, happily use this, and indeed this example can be adapted for GCC such that something different from the naively expected ‘20’ is printed.

The standard’s definition of ‘visible’ is rather unclear, especially when a pointer to a common initial segment is passed through another function. For example, in

```
int *f(int *p) { return p; }
int main(void) {
    u.s2.m = 20;
    printf("%d\n", f(u.s1.m));
}
```

does passing the pointer through `f` remove the visibility of the common initial segment? Our semantics takes the cautious way, and annotates each structure fragment in the path corresponding to a pointer with a flag whether the common initial segment rule may be used. When storing objects in the memory, this flag is set to false. Hence, it imposes undefined behavior on this last example.

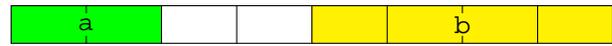


Fig. 4: `struct T { short a ; int b }` on a standard 32 bits architecture.

IV. INDETERMINATE VALUES

Uninitialized variables and padding bytes of objects of structure type take an *indeterminate value* [7: 6.2.6.1p6]. An indeterminate value is an object that either describes an unspecified value or is a *trap representation* [7: 3.17.2]. A trap representation is an object representation that does not represent a value of the object type and reading it leads to undefined behavior [7: 6.2.6.1p5].

A. Uninitialized variables

Since an object of type `unsigned char` cannot have a trap value, reading it does not exhibit undefined behavior. Instead it just gives an unspecified value. This property is important to allow simple minded bit-wise copying of structures, without having to worry about padding bytes between members. For example:

```
struct T { short a ; int b; } x = {10, 11}, y;
for (size_t i = 0; i < sizeof(x); i++)
    ((unsigned char*)&y)[i] =
    ((unsigned char*)&x)[i];
```

Figure 4 displays the representation of `x` on a standard 32 bits architecture with a 4 bytes alignment requirement for integers. This means that integers should be aligned at addresses that are multiples of 4, and therefore we have 2 padding bytes between the members. In case reading indeterminate values of type `unsigned char` (and in particular these padding bytes) would exhibit undefined behavior, this copy would also exhibit undefined behavior.

An interesting property of indeterminate values is that Defect Report #260 [9] allows them to change arbitrarily, so reading an indeterminate value twice might yield different results. This is useful for an optimizing compiler because it may figure out the actual lifetime of two values is disjoint and thereby share the storage location of both. As an example (the mysterious `&x` will be explained later)

```
unsigned char x; &x;
printf("%d\n", x);
printf("%d\n", x);
```

does *not* exhibit undefined behavior (an object of type `unsigned char` cannot contain a trap value), but Defect Report #260 allows the two printed values to be different. This is not so strange: the compiler might do liveness analysis and decide that `x` does not need to be saved on the stack when calling `printf`. And then of course `printf` might clobber the register that contains `x`.

Unfortunately, the standard is very unclear about the imposed behavior of various operations on indeterminate values, *e.g.*, what happens when they are copied or used in expressions. For example, should `y` be indeterminate after evaluating

```

unsigned char x, y; &x;
y = x/2;

```

Surely the most significant bit of `y` will be 0 after this? Or is this not something that the standard guarantees? But if `y` is *not* indeterminate after this, what about:

```
y = x/1;
```

We just changed the constant, and therefore after this statement `y` also should be determinate? But after:

```
y = x;
```

will it still not be indeterminate? This seems almost indistinguishable from the `x/1` case, but the liveness analysis argument surely also will apply to this `y`? A formal C semantics will have to take a stance on this.

Also, should the following print '0', or may it print a different value as well, because the `x` changed during the evaluation of the subtraction expression?

```

unsigned char x; &x;
printf("%d\n", x - x);

```

Defect Report #338 [9] remarks that on some architectures (e.g. IA-64) registers may hold trap values that do not exist in memory. Thus, for such architectures, programs cannot safely copy uninitialized variables of type **unsigned char** because these might reside in registers. In the C11 standard this problem has been fixed by including the following workaround [7: 6.3.2.1p2]:

If the lvalue designates an object of automatic storage duration that could have been declared with the register storage class (never had its address taken), and that object is uninitialized (not declared with an initializer and no assignment to it has been performed prior to use), the behavior is undefined.

This is the reason we had to take the address of the uninitialized variables in the preceding examples. Of course even with the `&x` present, a compiler can decide to ignore it and still use a register for `x`, as this address is never used.

This workaround again shows that treating uninitialized objects as indeterminate has its difficulties. In the memory model of our semantics we keep track of uninitialized (or freed) memory by special undef nodes (see Figure 2). Since objects of structure or union type cannot be indeterminate, we only allow these undef nodes on the leaves. Again, we take the most cautious way, and let operations on these special undef nodes, like reading them, exhibit undefined behavior.

B. Freeing memory

The standard states that the value of a pointer becomes indeterminate when the object it points to reaches the end of its lifetime [7: 6.2.4]. In particular this means that whenever some memory is freed, all pointers to it become indeterminate. For example, assuming the `mallocs` do not fail, the following program can still exhibit undefined behavior

```

int *p = malloc(sizeof(int));
free(p);
int *q = malloc(sizeof(int));
printf("%d\n", p == q);

```

because the value of the pointer `p` has become indeterminate and now can be a trap value. Of course, we can still compare the bit patterns of both pointers, and if they are equal, try to use `p` instead.

```

int *p = malloc(sizeof(int));
free(p);
int *q = malloc(sizeof(int));
if(memcmp(&p, &q, sizeof(p)) == 0)
    *p = 10;

```

Again, Defect Report #260 [9] states that this program exhibits undefined behavior.

The fact that a pointer object becomes indeterminate after the block it points to has been freed means that if we copy pointers to various places in memory, then all copies should become indeterminate and not just the argument of `free` (which does not even need to be an lvalue). This means that a `free` operation will affect the formal memory state globally. And what about individual bytes of a pointer that have been copied? Will these also become indeterminate after freeing that pointer?

C. Padding bytes

The standard states that when a value is stored in a member of an object of structure or union type, padding bytes take an unspecified value [7: 6.2.6.1p6], and that when a value is stored in a member of a union type, bytes that do not correspond to that member take an unspecified value [7: 6.2.6.1p7]. Consider:²

```

union { int a[1] ; int b[2]; } x;
x.b[0] = 10; x.b[1] = 11;
x.a[0] = 12;
printf("%d\n", x.b[1]);

```

This example can lead to undefined behavior, as assigning to `x.a` makes the bytes of `x.b` that do not belong to `x.a` unspecified, and thus `x.b[1]` possibly indeterminate. Taking Defect Report #260 into account, it is unclear whether the bytes that belong to `x.b[1]` after this may change arbitrarily.

This program also suggests that the representation of pointers in a formal semantics should contain information describing which parts of the memory should become indeterminate upon assignment to them. If instead of assigning to `x.a[0]` directly we do:

```

int *p = &x.a[0];
*p = 12;

```

`x.b[1]` will *still* become indeterminate. But the assignment to `p` might happen anywhere in the program, even in a context where the union type is not visible at all.

²Adapted from Shao Miller, Bounds Checking as Undefined Behaviour, `comp.std.c` newsgroup, July 29, 2010.

V. C IS NOT TURING COMPLETE

In this section we will argue that *no* Turing complete implementation of the C standard exists. Obviously there do exist *some* implementations that are not Turing complete, but that is weaker than what we claim. We argue that for *any* implementation that behaves as described by the standard, all programs will correspond to a finite state machine, and hence that their termination is decidable. This implies that C is not Turing complete, as for a Turing complete language termination is undecidable by the halting problem. Our argument is similar to the well-known pumping lemma [22].

Unfortunately, we cannot give a mathematical proof that C is not Turing complete, as that would require a mathematically precise version of the standard, which does not exist yet.

Let us put a restriction on what we will establish. With sufficient I/O, an external file may be used as the tape to create a Turing machine emulator. Even then the existence of the function

```
long int ftell(FILE *stream);
```

that gives the position in a file [7: 7.21.9.4], seems to imply that files as described in the standard need to have a bounded length as the return type of `ftell` is bounded. However, an implementation could have its own infinite-file access interface, allowing it to become Turing complete. Therefore we will here only consider C programs without I/O. (Our argument *can* be adapted to show that the programs that restrict I/O to reading `stdin` and writing `stdout` are Turing incomplete as well, but we will not do so here.)

The first step in our argument is the observation that each type that is not an array has finitely many distinct values, as a value of a type T needs to have a representation that consists of `sizeof(T) * CHAR_BIT` many bits. Now for array types the elements of the array need to have an address that corresponds to a pointer, and those pointers *also* only have finitely many distinct values. This means that there is an upper bound on the length of arrays, and thus array types also only have a bounded number of distinct values.

There are three ways to get storage in C: (1) through global or static variables, (2) through allocated storage (*e.g.* using `malloc`), (3) through automatic variables in function calls. The size of the storage obtained through (1) is clearly bounded. All allocated storage needs to be addressable and there are only finitely many pointers, therefore storage obtained through (2) is also bounded. As for (3), since the standard does not restrict recursion depth, there could be unboundedly many automatic variables. This means that the size of this part of the state of a program does not have to be bounded. However, the size of the *visible* part of that state, the automatic variables that are in scope, together with those that can be reached by following pointers, *is* bounded.

Now there will be a (very large) bound to the recursion depth, such that if that bound is exceeded, a function will end up in a visible state that it has been in before. Clearly, in that case, the program will be in an infinite loop. Therefore, for an arbitrary program there are two possibilities. Either the

bound to the recursion depth will not be exceeded, in which case the size of the state is bounded, and it behaves as a finite state machine. Or, the recursion depth *will* be exceeded, in which case the program will infinitely go deeper and deeper in recursion, repeating the pattern of function calls, in which case the program also behaves as a finite state machine. This finishes our argument.

The result that C is not Turing complete is surprising: it seems possible to implement a Turing machine emulator in C, which seems to imply Turing completeness. However, the problem is that the C data types are too small to represent an arbitrarily long tape. C integers only have a specific number of bits, and even when one wants to use arrays or linked lists to hold bigger data structures, the *pointers* that are involved in this also only have a finite number of bits, and therefore one cannot have arbitrarily large data structures as one would run out of different pointer values.

VI. STRICTLY CONFORMING C PROGRAMS DO NOT EXIST

An implementation of a programming language typically organizes its storage into two parts: the *stack* and the *heap*. On a function call, the stack is extended with a frame containing the function's arguments, local variables and return address (the address of the instruction to be executed when the function is finished). The heap contains dynamically allocated storage.

The standard abstracts from implementation details like these, and thus also allows implementations that do not organize their memory in this way. Although we agree that this is the right approach, we do believe it should at the very least account for (an abstract version of) the problem of *stack overflow*. Unfortunately, the notion of stack overflow is not mentioned by the standard [7] or the standard's rationale [8] at all. This is very troublesome, as for most actual implementations stack overflow is a real problem. Let us consider the following function.

```
/*@ decreases n; ensures \result == n; */
unsigned long f(unsigned long n) {
    if (n != 0) return f(n - 1) + 1;
    else return 0;
}
```

With most tools for C verification one can prove by induction that the function `f` behaves as the identity function (for example, the Jessie plug-in for Frama-C [19] turns these four lines into 92 lines of Why3ML code, leading to four verification conditions that are all trivially provable). However, in a real C implementation, a call like `f(10000000)` will *not* return 10000000, but instead will crash with a message like 'segmentation fault'.

Furthermore, stack overflow does not necessarily have to result in a crash with a nice error message, but might also overwrite non-stack parts of the memory (possibly putting the address of virus code there).

Stack overflow even can occur without function calls. For example, the program

```
int main(void) { int a[10000000]; }
```

might also crash. In order to make this really happen with an optimizing compiler, one might need to use the array `a` to prevent it from being optimized away. For example when compiled with GCC or clang (with `-O2`) the following crashes when running it with the standard stack size:

```
int main(void) {
    int a[10000000];
    for (int i = 0; i < 10000000; i++) a[i] = i;
    return a[10];
}
```

This all means that a proof of correctness of a program with respect to the standard only guarantees correctness relative to the assumption that the stack does not overflow. As we have seen, this assumption does not hold in general. But even worse, it is not clear for which n the program

```
int main(void) { int a[n]; }
```

is guaranteed *not* to overflow the stack. On a microcontroller this might already happen for rather small n . Therefore, as there is no a clear division between the n that are safe and the ones that are not, it seems reasonable that even the extreme program, which is a variation on the case for $n = 0$,

```
int main(void) { }
```

potentially could overflow a (very) small stack. For one thing, there need to be space on the stack for the `int` return value, and then the runtime environment that calls `main` also might need a bigger stack than is there.

The obvious way to change the text of the C standard to address this issue would be to add to [7: 6.5.2.2] something like:

A function call (even the initial call to `main`) might overflow the stack, in which case the behavior is undefined. It is implementation defined under what circumstances this is guaranteed not to happen.

It is important that it is not just *unspecified* when these overflows happen, for then it still would be impossible to reason about any program formally from the text of the standard.

A consequence of this addition is that *strictly conforming programs* no longer exist as one of their defining properties is that they ‘shall not produce output dependent on ... implementation defined behavior’ [7: 4.5]. Of course, once one takes stack overflow into account no program has that property anymore.

Another consequence is that *all* C implementations become *conforming*, as that notion is defined by quantification over all strictly conforming programs [7: 4.6]. Therefore, the definition of conforming C implementations should also be adapted.

The fact that the standard does not allow to portably test for stack overflow is to us one of the main omissions of the language. A call to `malloc` will return a `NULL` pointer when there is no storage available, which means that a program can test for that situation. But there is no counterpart to this for function calls, or for entering a block.

VII. FAILURE OF THE SUBJECT REDUCTION PROPERTY

A desired property for a typed programming language is *subject reduction*, which means that evaluation preserves typing. As proven by Norrish, this property holds for (his small-step expression semantics of) C89 [20]. We will argue that due to the introduction of variable length arrays in C99, this property no longer holds. Before pursuing the problem of subject reduction, we briefly introduce variable length arrays and some other problems of these.

Before C99, arrays were required to have a size that could be determined at compile-time. To avoid this restriction programmers had to use dynamically allocated memory (e.g. through `malloc` and `free`) for arrays of dynamic size. To loosen this restriction, C99 introduced support for *variable length arrays* (VLAs) to the language.

The first shortcoming of VLAs is related to the fact that there is no portable way to test if there is sufficient storage available (on the stack) when entering a block (see Section VI). Since most implementations use the stack to store VLAs, not being able to perform such a test, makes VLAs dangerous in practice. Consider the following program.

```
int main(void) {
    int n;
    scanf("%d", &n);
    int a[n];
}
```

Since there is no way to test if there is enough space on the stack, it is impossible to portably ensure that the program does not crash.

Another problem of VLAs is that C allows casts to variable length types. Since size expressions in such casts may impose side-effects, this makes the situation rather complex. For example, consider:

```
1 ? (int(*)[f(5)])0 : (int(*)[f(3)])0;
```

where `f` is the identity function on `int`. We were unable to find anything in the standard on which of the function calls `f(5)` and `f(3)` may (or should) be evaluated. It is reasonable to allow implementations to evaluate neither of them, as programs can generally be executed correctly without needing to know the size of array bounds in pointer casts. But what about:

```
printf("%d\n",
    sizeof*(1 ? 0 : (int(*)[f(3)])0));
```

Here an implementation clearly needs to evaluate the function call, to obtain the value of the `sizeof`, even though it is in the branch of the conditional that is not taken. The standard includes the following related clause [7: 6.7.6.2p5]:

Where a size expression is part of the operand of a `sizeof` operator and changing the value of the size expression would not affect the result of the operator, it is unspecified whether or not the size expression is evaluated.

Since in this example the size expression in the ‘not to be taken’ branch has to be evaluated, one may wonder whether they are also allowed to be evaluated in the ‘not to be taken’ branch of the earlier example.

It also is unclear how these function calls in casts are evaluated with respect to sequence points. As they are evaluated in branches that are not taken, it seems they are exempt from the normal execution flow of an expression. But if they already can be executed before the sequence point that starts the subexpression that contains them, is there a reason they cannot be executed before the evaluation of the statement that contains them starts?

The standard’s (implicit) distinction between *static* and *run-time* typing is the reason that subject reduction breaks. This means that an expression can get a more restricted type during its evaluation. For example, statically the expression `(int(*)[f(5)])0` has type ‘pointer to integer array of variable length’, whereas at run-time it will become of the more restricted type ‘pointer to integer array of length n ’ where n is the result of evaluating `f(5)`.

The previous example already indicates that one has to sacrifice either subject reduction or uniqueness of types (that is, each expression has a unique type). However, we will argue that the situation is worse, and that for a reduction semantics whose rules are applied locally, subject reduction will fail even if expressions are allowed to have multiple types. Consider:

```
1 ? (int(*)[f(5)])0 : (int(*)[3])0
```

In this example the two subexpressions `(int(*)[f(5)])0` and `(int(*)[3])0` have (static) types `int(*)[*]` and `int(*)[3]`, respectively, where $T[*]$ denotes the variable length array type over T . By typing of the conditional and composite types [7: 6.5.15p6, 6.2.7] the full expression has type `int(*)[3]`.

If the function call `f(5)` gets evaluated, and returns a value different from 3, typing breaks, *i.e.*, a well typed expression is evaluated to a non-well typed expression. Luckily, the standard imposes undefined behavior on this example by means of the following clause [7: 6.7.6.2p5]:

If the two array types are used in a context which requires them to be compatible, it is undefined behavior if the two size specifiers evaluate to unequal values.

Currently, our C semantics deals with this problem by allowing evaluation of size expressions at any place in a bigger expression. If the conditional gets reduced, we check if both arguments have compatible types, and if not, assign undefined behavior. This approach has two obvious drawbacks. First of all, it breaks subject reduction, and secondly, undefined behavior gets caught at a late moment, or might not get caught at all. For example, in

```
g() ? (int(*)[f(5)])0 : (int(*)[3])0
```

it may happen that `f(5)` is evaluated first, and returns a value unequal to 3, in which case typing has already failed. After that, it may happen that the call to `g()` does not return (or

example because it invokes an `exit` or loops indefinitely), resulting in the undefined behavior not getting described by the semantics.

But does this mean that in a formal C semantics each reduction step has to establish that typing does not break in order to catch undefined behavior? This would obviously fix subject reduction, but destroys the locality of a small-step reduction semantics.

VIII. CONCLUSION

A. Discussion

We will finish our discussion of the subtleties of the C standard by asking the following questions:

- 1) Is the interpretation of the C standard that we presented in this paper a reasonable reading of the standard?
- 2) Is the C standard itself (under this interpretation) reasonable?

First of all, we claim that the standard is not fully clear. The standard committee’s response in Defect Report #260 [9] is not obvious from the text of the standard. Also, Maclaren [18] presents various issues about which the standard does not have an unambiguous reading, even when taking Defect Report #260 into account. One might even claim that the exact text of the standard (without giving it a liberal reading) is inconsistent, *i.e.* that from a logical point of view the standard implies anything one likes. But of course that position is not very productive.

Even the relation between Defect Report #260 and the official standard text is not completely clear. The report *does* include a response by the standard committee and as such should not be taken lightly. It was a clarification of the C99 version of the standard, and hence it seems obvious that there was some defect in that text. However, the parts of the text of the standard in the C11 version relevant for the issues discussed in the report have not changed from their counterpart in the C99 standard at all.

The standard makes a very clear distinction between ‘normative’ and ‘informative’ text in the standard, and the explanations of Defect Report #260 certainly are not part of the ‘normative’ text of the C11 standard. Therefore, it seems an option to decide to ignore the committee’s response in this report, especially the notion of the *origin* of an object, which does not occur in the standard at all. In that case, one could read the standard in a ‘Kernighan & Ritchie’ manner. But of course in that case the optimizations of for example GCC *will* be incorrect, and one will get standard compliant behavior only when compiling using flags like `-fno-strict-aliasing`. Many real world programs, like for example the Linux kernel, are compiled with such flags anyway [24], but on the other hand this attitude would mean that strictly conforming programs would have less than optimal performance, because they would need to generate more code to test for changes due to aliasing.

The second question, whether the standard can and should be improved is even more interesting. We think that the fact

that it is very difficult (or maybe even impossible) to give a mathematically precise version of the memory model from the standard, is a sign that one should aim for a clearer explanation of the issues involved. However, it seems difficult to give a version that is both fully clear and still is reasonably compatible with the current standard.

For example, in the memory model of CompCert, bytes (objects of type `unsigned char`) that constitute a pointer are *not* simple sequences of bits, but instead are abstract entities [16]. This means that treating them like numbers – calculating with them, printing them – will exhibit undefined behavior. This seems not in the spirit of the language. On the other hand, for actual practical programs, this restriction on the programs that are considered meaningful probably is not important at all.

B. Future work

The future work in our Formalin project has two directions. On the one hand we are porting the prototype semantics that we have developed in Haskell to the theorem prover Coq. For the moment, this development mainly ignores the issues described in this paper and aims at an elegant semantics for non-local control flow, block scoped variables, expressions with side effects, and sequence points. In this formalization we both give an operational semantics and an axiomatic semantics, and we prove that the axiomatic semantics is correct with respect to the operational semantics.

On the other hand we are experimenting with a formal memory model in which the issues from this paper *are* taken into account. The state of the memory in this model consists of a pair of an *abstract* version of the memory (like the one in our Haskell prototype), and a *concrete* version of the memory (which consists of sequences of bytes encoding the data). The puzzle is how to deal with information flow between these two components. For the moment we have not been able to resolve these issues in a satisfactory way.

If we succeed in creating such a ‘Defect Report #260-compliant’ formal memory model, we will incorporate it into the operational semantics that we already are working on. We then also will be able to take a guaranteed consistent position on the issues discussed in this paper.

Acknowledgments. We thank Rob Arthan, Sandrine Blazy, Herman Geuvers, Andreas Lochbihler, Michael Norrish and Erik Poll for discussions on this topic, Xavier Leroy for details on the CompCert memory model, and the people on the `comp.std.c` newsgroup for helpful discussions on the C standard. This work is financed by the Netherlands Organisation for Scientific Research (NWO).

REFERENCES

- [1] Sandrine Blazy and Xavier Leroy. Mechanized semantics for the Clight subset of the C language. *Journal of Automated Reasoning*, 43(3):263–288, 2009.
- [2] Ernie Cohen, Markus Dahlweid, Mark A. Hillebrand, Dirk Leinenbach, Michal Moskal, Thomas Santen, Wolfram Schulte, and Stephan Tobies. VCC: A Practical System for Verifying Concurrent C. In *TPHOLS*, volume 5674 of *LNCS*, pages 23–42, 2009.
- [3] Jeffrey Cook and Sakthi Subramanian. A Formal Semantics for C in Nqthm. Technical Report 517D, Trusted Information Systems, 1994.
- [4] Crispin Cowan, Perry Wagle, Calton Pu, Steve Beattie, and Jonathan Walpole. Buffer Overflows: Attacks and Defenses for the Vulnerability of the Decade. In *DISCEX*, 2000.
- [5] Chucky Ellison and Grigore Rosu. An executable formal semantics of C with applications. In *POPL*, pages 533–544, 2012.
- [6] Chucky Ellison and Grigore Rosu. Slides for the POPL’12 talk on [5], 2012. <http://fsl.cs.uiuc.edu/pubs/ellison-rosu-2012-popl-slides.pdf>.
- [7] International Organization for Standardization. *INCITS/ISO/IEC 9899-2011: Programming languages – C*. ISO Working Group 14, 2012.
- [8] International Organization for Standardization. Rationale for International Standard – Programming Languages – C, 2003. Revision 5.10.
- [9] International Organization for Standardization. WG14 Defect Report Summary for ISO/IEC 9899:1999, 2008. <http://www.open-std.org/jtc1/sc22/wg14/www/docs/>.
- [10] Bart Jacobs and Frank Piessens. The VeriFast program verifier. Technical Report CW-520, Department of Computer Science, Katholieke Universiteit Leuven, Belgium, 2008.
- [11] Brian W. Kernighan and Dennis M. Ritchie. *The C Programming Language*. Prentice Hall, 2nd edition, 1988.
- [12] Gerwin Klein et al. seL4: formal verification of an OS kernel. In *SOSP*, pages 207–220, 2009.
- [13] Robbert Krebbers and Freek Wiedijk. A Formalization of the C99 Standard in HOL, Isabelle and Coq. In *CICM*, volume 6824 of *LNAI*, pages 297–299, 2011.
- [14] Patrice Lacroix and Jules Desharnais. Buffer Overflow Vulnerabilities in C and C++. Rapport de Recherche DIUL-RR-0803, Université Laval, Québec, Canada, 2008.
- [15] Dirk Leinenbach. *Compiler Verification in the Context of Pervasive System Verification*. PhD thesis, Saarland University, Saarbrücken, 2008.
- [16] Xavier Leroy, Andrew W. Appel, Sandrine Blazy, and Gordon Stewart. The CompCert Memory Model, Version 2. Research report RR-7987, INRIA, 2012.
- [17] Xavier Leroy and Sandrine Blazy. Formal verification of a C-like memory model and its uses for verifying program transformations. *Journal of Automated Reasoning*, 41(1):1–31, 2008.
- [18] Nick Maclaren. What is an Object in C Terms?, 2001. Mailing list message, <http://www.open-std.org/jtc1/sc22/wg14/9350>.
- [19] Yannick Moy and Claude Marché. *Jessie Plugin Tutorial, Beryllium version*. INRIA, 2009.
- [20] Michael Norrish. *C formalised in HOL*. PhD thesis, University of Cambridge, 1998.
- [21] Nikolaos Papaspyrou. *A Formal Semantics for the C Programming Language*. PhD thesis, National Technical University of Athens, 1998.
- [22] Thomas Sudkamp. *Languages and machines - an introduction to the theory of computer science*. Addison-Wesley, 1988.
- [23] TIOBE Software. TIOBE index. <http://www.tiobe.com/content/paperinfo/tpci/>, 2012.
- [24] Xi Wang, Haogang Chen, Alvin Cheung, Zhihao Jia, Nickolai Zeldovich, and M. Frans Kaashoek. Undefined Behavior: What Happened to My Code? In *APSys*, 2012.
- [25] Xuejun Yang, Yang Chen, Eric Eide, and John Regehr. Finding and understanding bugs in C compilers. In *PLDI*, pages 283–294, 2011.