

# A flat reachability-based measure for CakeML’s cost semantics

Alejandro Gomez-Londoño  
alejandrogomez@chalmers.se  
Chalmers University of Technology  
Gothenburg, Sweden

Magnus O. Myreen  
myreen@chalmers.se  
Chalmers University of Technology  
Gothenburg, Sweden

## ABSTRACT

The CakeML project has recently developed a verified cost semantics that allows space cost reasoning of CakeML programs. With this space cost semantics, compiled machine code can be proven to have tight memory bounds ensuring no out-of-memory errors occur during execution. This paper proposes a new cost semantics which is designed to make proofs about space costs significantly simpler than they were with the original version. The work described here has been developed in the HOL4 theorem prover.

## CCS CONCEPTS

• **Software and its engineering** → **Formal software verification; Compilers.**

## KEYWORDS

compiler verification, cost semantics, space usage

### ACM Reference Format:

Alejandro Gomez-Londoño and Magnus O. Myreen. 2018. A flat reachability-based measure for CakeML’s cost semantics. In *Woodstock ’18: ACM Symposium on Neural Gaze Detection, June 03–05, 2018, Woodstock, NY*. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/1122445.1122456>

## 1 INTRODUCTION

Functional languages present programmers with many abstractions (e.g., polymorphism, garbage collection, ADTs, among others) to aid them in the development of complex programs. However, these features often come at the cost of an increase in memory usage and an unclear space usage; making it difficult to judge whether a program has enough memory to run.

The CakeML project has recently developed a verified cost semantics [5] which, at its core, uses a measuring function embedded in the semantics to determine if, at any given point, a program has run out of memory. The cost semantics has been proved sound, which means machine code generated by the CakeML compiler will never produce out-of-memory errors if the CakeML cost semantics has ruled them out. Unfortunately, reasoning using the the original cost semantics for CakeML requires considerable effort.

This paper improves on the original cost semantics. This paper defines an alternative space measuring function, which is defined

in stages: it first computes the set of reachable nodes, and then computes the sum of the size of the data at those nodes. In contrast, the old space measuring function did everything at once: it discovered the reachable part of the heap through recursive descent. The new formulation is expected to lead to tidier proofs than the previous measuring function. Our initial experiments suggest that proofs about space cost are significantly more manageable with the new formulation.

This paper makes the following contributions:

- In this paper, we define a new reachability-based measuring function (Section 3), which is designed to be significantly simpler to work with than the original (Section 2.3).
- We demonstrate (in Section 4) how the new formulation overcomes some of the most significant problems of the original formulation.
- Finally, we discuss (in Section 5) ways in which the new formulation of the space cost semantics can be proved sound so that future space cost proofs can use the new formulation.

All of this work is machined-checked using the HOL4 theorem in the context of the CakeML compiler verification project. It should be noted that this work is currently in an early stage of development, but of sufficient maturity to be presented at IFL. We are confident that the main result of this paper can be fully completed in time for the post-proceedings submission.

## 2 A VERIFIED COST SEMANTICS

The cost semantics for the CakeML compiler [5] is expressed at the level of its `DATALANG` intermediate language.

`DATALANG` is an intermediate language approximately in the middle of the CakeML compiler. It is an imperative intermediate language with nested tuple-like values and reference pointers, but no function values. It appears right before memory becomes finite and the garbage collector is introduced. The semantics of `DATALANG` is expressed in the form of a (functional) big-step semantics.

The semantics for `DATALANG` acts as a cost semantics for CakeML by maintaining a boolean-valued `safe_for_space` field in the semantic state of the operational semantics. This field is set to false whenever a semantic space cost measurement predicts that the current use of space might exceed the configured space limits for heap or stack space.

This paper focuses on the measurement of heap space. At each allocation of new memory, the semantics for `DATALANG` computes the size of the currently live data using a measure called `size_of`. This `size_of` function computes the space consumption of all reachable values from the root values obtained from the stack and global variables. This `size_of` function is defined to carefully track

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

Woodstock ’18, June 03–05, 2018, Woodstock, NY

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-XXXX-X/18/06...\$15.00

<https://doi.org/10.1145/1122445.1122456>

aliasing by keeping track of pointer-equal values, and is unchanged by garbage collection as it only consider live (reachable) data.

To prove the space safety of a CakeML program, one must show that for some limit—concrete or abstract—the semantics of its DATA-LANG representation never sets `safe_for_safe` to false. Once space safety is established, it can be extended all the way to the level of machine code, thanks to the soundness proof for the cost semantics w.r.t. to the CakeML compiler.

The rest of this section introduces: the DATA-LANG semantics, the space semantics, and the definition of the original heap space measure, i.e. `size_of`. More details on the original DATA-LANG’s operational and costs semantics can be found in prior work [5, 9].

## 2.1 DATA-LANG at a glance

DATA-LANG is an imperative language with abstract values, stateful storage of local variables, and a call stack. In the compiler-stack, it sits between the more abstract functional languages and the low-level languages with word-based value representations.

To give a sense of how CakeML programs look when compiled into DATA-LANG consider the following CakeML function expressed in CakeML source syntax (which is very similar to SML syntax).

```
fun app123 x = let a = [1,2,3] in a ++ x end
```

This function prepends the list `[1, 2, 3]` to its input. The result of compiling this function to DATA-LANG is shown in Figure 1.

```
line 0 :   app123 [0] evaluates as
line 1 :     MakeSpace 9
line 2 :     1 := Cons nil_tag []
line 3 :     2 := Const 3
line 4 :     3 := Cons cons_tag [2; 1]
line 5 :     4 := Const 2
line 6 :     5 := Cons cons_tag [4; 3]
line 7 :     6 := Const 1
line 8 :     7 := Cons cons_tag [6; 5]
line 9 :     8 := ListAppend [7; 0]
line 10 :    return 8
```

**Figure 1: DATA-LANG code for a function that prepends [1, 2, 3] to its argument**

At first, the DATA-LANG presentation of the code might seem significantly different. However, on closer inspection, we hope the reader will see the similarity. In DATA-LANG primitive operation are always assigned (`:=`) to a local variables, which are represented as natural numbers. On line 0, argument 0 corresponds to the source code binding `x`. Line 1 allocates 9 slots of space (that’s 3 slots per Cons). Line 2 creates a value representing an empty list using the primitive operation `Cons` and a number tag (`nil_tag`) denoting the `nil` constructor for lists. In line 3 a `Const` operation creates the number literal 3. Line 4 combines local variables 1 (`[]`) and 2 (3) into the singleton list `[3]` using `Cons` and the corresponding list constructor tag (`cons_tag`); using the same process, lines 5 through 8 create the DATA-LANG representation of the list `[1, 2, 3]`. Then, Line 9 applies `ListAppend`—which appends the two lists-shaped values—variables 0 (the argument) and 7 (`[1, 2, 3]`).

```
v = Number int
  | Word64 word64
  | CodePtr num
  | RefPtr num
  | Block timestamp tag (v list)
```

**Figure 2: DATA-LANG’s abstract values**

Primitive values in DATA-LANG are modeled by the data type presented in Figure 2. `Number` is an arbitrarily large integer. `Word64` is a 64-bit machine word. `CodePtr` is a code pointer, and `RefPtr` is a pointer to mutable state (such as arrays). The `Block` constructor represents contiguous values in memory, and encodes datatype constructors, tuples and vectors.

An example of a DATA-LANG value is shown in Figure 3 which shows the `app123_nil`. This is a value that can result from a call to `app123` with the empty list as the argument. `Block` values, with `cons_tag` and `nil_tag` indicate the source-level constructor that each `Block` represents. Furthermore, timestamp values 8, 7, and 6 uniquely identify each block.

```
app123_nil def Block 8 cons_tag [Number 1;
                          Block 7 cons_tag [Number 2;
                          Block 6 cons_tag [Number 3;
                          Block 0 nil_tag [ ]]]]
```

**Figure 3: Block representation of CakeML list [1, 2, 3]**

The semantics state is defined as the record type shown in Figure 4. The fields `locals` and `refs` represent the finite maps of local variables (`v num_map`) and references (`v ref num_map`) respectively. The stack is a list of frames, each frame containing only the relevant variables that should be restored after a call is completed. The `global` field contains an optional reference to an array of global variables. Space limits are kept in a record with fields for heap and stack limits. The boolean flag `safe_for_space` is set to false when space limits have been exceeded. The remaining fields are not of relevance for the presentation here.

The semantics of DATA-LANG is defined as a functional big-step semantics [8]. In this style of semantics, a clocked big-step evaluation function, `evaluate`, takes a (program, state) pair as input, and returns a (result, state) pair as output. As an example, consider the evaluation of `app123` with the empty list as argument, which results in value `app123_nil`. Note that the program is given to `evaluate` as a DATA-LANG AST (`app123_prog`) and arguments are local variables in the state.

```
evaluate (app123_prog,
         s with locals := { 0 ↦ Block 0 nil_tag [ ] })
= (app123_nil, s')
```

To better visualize intermediate steps of evaluations, the DATA-LANG semantics can also be expressed as a shallowly embedded state-exception monad. This is the representation used in `app123` which can partially evaluate the first three operations by unfolding `bind` applications:

```

α state = <|
  locals : v num_map;
  refs : v ref num_map;
  stack : stack list;
  global : num option;
  limits : limits;
  safe_for_space : bool;
  clock : num;
  ...
|>

ref = ValueArray (v list) | Bytes bool (word8 list)

limits = <|
  heap_limit : num;
  stack_limit : num;
  ...
|>

```

**Figure 4: The definition of the DATA<sub>LANG</sub> state.**

```

app123 (s with locals := { 0 ↦ Block 0 nil_tag [] })
= (4 := Const 2
   5 := Cons cons_tag [4; 3]
   6 := Const 1
   7 := Cons cons_tag [6; 5]
   8 := ListAppend [7; 0]
   return 8)
s with <| locals := { 0 ↦ Block 0 nil_tag []
                    1 ↦ Block 0 nil_tag []
                    2 ↦ Number 3
                    3 ↦ Block 3 cons_tag
                        [Number 3;
                          Block 0 nil_tag [] ] };
  ...
|>

```

## 2.2 Embedded cost semantics

As previously stated, DATA<sub>LANG</sub>'s costs semantics is embedded into its operational semantics. Therefore, proving space safety of `app123` is a matter of proving the following statement:

$$\vdash s.\text{limits.heap\_limit} = mh \wedge s.\text{limits.stack\_limit} = ms \wedge s.\text{safe\_for\_space} \wedge \text{evaluate}(\text{app123\_prog}, s) = (res, s') \Rightarrow s'.\text{safe\_for\_space}$$

This is, given stack space  $mh$  and heap space  $ms$ ; the evaluation of `app123_prog` preserves `safe_for_space`, thus signalling that the program's memory consumption falls within the given bounds.

Internally, the `safe_for_space` flag is updated at every space-consuming operation, for example, at function calls and whenever new values are created. Auxiliary functions `size_of_heap` and `size_of_stack` are used to update `safe_for_space` in one of two ways. If  $k$  slots of new heap space are to be used (e.g. as part of

`MakeSpace`), then `safe_for_space` is updated as follows:

```

s with
  safe_for_space :=
    (s.safe_for_space ∧
     size_of_heap s + k ≤ s.limits.heap_limit)

```

Similarly, if  $k$  slots of new stack space are to be consumed (e.g. as part of a function call), then `safe_for_space` is updated as follows:

```

s with
  safe_for_space :=
    (s.safe_for_space ∧
     size_of_stack s + k ≤ s.limits.stack_limit)

```

The important work is performed by the `size_of_heap` and `size_of_stack` functions. This paper focuses on improving the formulation of the heap space measure and thus `size_of_heap`.

The original formulation of `size_of_heap` is shown below. Here `stack_to_vs` is an auxiliary function that computes a list of root values from local variables (`s.locals`), the call-stack (`extract_stack`), and global references (`global_to_vs`). The root values are given to the measuring function `size_of`, which computes the size of heap elements reachable from these initial elements.

```

size_of_heap s  $\stackrel{\text{def}}{=}$ 
  let (n, _) =
    size_of (stack_to_vs s) s.refs LN in
  n
stack_to_vs s  $\stackrel{\text{def}}{=}$ 
  toList s.locals ++
  extract_stack s.stack ++
  global_to_vs s.global

```

The main workhorse of this definition is the `size_of` function, which is the topic of the next section.

## 2.3 The original heap measure: size\_of

At the core of DATA<sub>LANG</sub>'s cost semantics is the heap space measuring function `size_of`. This function is responsible for computing a space consumed by all values reachable from the given values. Figure 5 shows its definition with `seen` (a set of timestamps), and `refs` as additional arguments.

The measurement of most values (`CodePtr`, `Word64`, and `Number`) is straightforward, as it is either constant, already accounted for within another structure (e.g. stack frames), or measured by a function without considering other values. Whereas, the handling of `Block` and `RefPtr` values is more involve and where most of the complexity of `size_of` comes from. In the case of `Block` values, a set of already-measured (`seen`) timestamps is kept to avoid counting identical blocks multiple times; this mechanisms assumes a bijection between timestamps and the blocks in memory. For `RefPtr`, pointers are removed from references map (`refs`) once they are counted, this is to only follow a reference once.

The definition of `size_of` succeeds at providing tight bounds, mitigating the effects of aliasing, and traversing all live data; however, perhaps due to its precise and concrete nature, it can be challenging to reason about. The main hurdle with `size_of` is the linearity of its traversal, where initial measurements at the front

```

size_of [] refs seen  $\stackrel{\text{def}}{=} (0, \text{refs}, \text{seen})$ 
size_of (x :: xs) refs seen  $\stackrel{\text{def}}{=} \text{let } (n_1, \text{refs}_1, \text{seen}_1) = \text{size\_of } x \text{ refs seen ;}$ 
  (n2, refs2, seen2) = size_of [x] refs1 seen1 in
  (n1 + n2, refs2, seen2)
size_of [Word64 v0] refs seen  $\stackrel{\text{def}}{=} (3, \text{refs}, \text{seen})$ 
size_of [Number i] refs seen  $\stackrel{\text{def}}{=} \text{(if is\_smallnum } i \text{ then } 0 \text{ else bignum\_size } i, \text{refs}, \text{seen})$ 
size_of [CodePtr v1] refs seen  $\stackrel{\text{def}}{=} (0, \text{refs}, \text{seen})$ 
size_of [RefPtr r] refs seen  $\stackrel{\text{def}}{=} \text{case lookup } r \text{ refs of}$ 
  None  $\Rightarrow (0, \text{refs}, \text{seen})$ 
  | Some (ValueArray vs)  $\Rightarrow$ 
    (let (n, refs', seen') = size_of vs (delete r refs) seen
     in
     (n + |vs| + 1, refs', seen'))
  | Some (ByteArray v2 bs)  $\Rightarrow$ 
    (|bs| div (arch_size lims div 8) + 2, delete r refs, seen)
size_of [Block ts tag vs] refs seen  $\stackrel{\text{def}}{=} \text{if } vs = [] \vee \text{isSome (lookup ts seen) then } (0, \text{refs}, \text{seen})$ 
  else
    let (n, refs', seen') =
      size_of vs refs (insert ts () seen) in
    (n + |vs| + 1, refs', seen')

```

Figure 5: Definition of the old size\_of.

of the argument list directly affect subsequent ones through pointers or timestamps. Thus, conceptually simple properties (e.g. the reordering of values) are hard to prove and apply.

### 3 A FLAT REACHABILITY-BASED MEASUREMENT

This section shows the definition of our new heap cost measuring function, flat\_size\_of, which is to replace the original size\_of. In a nutshell, flat\_size\_of takes a set of root addresses, computes the set of all addresses reachable from that initial set, and then sums the size of the heap element that is at one of the reachable addresses. We will go through the detail below.

First of all, DATA<sub>LANG</sub> has no immediate notion of heap address. For the purposes of the definition of flat\_size\_of, we define a type for DATA<sub>LANG</sub> addresses. An address is either a timestamp (TStamp) of a Block (remember each block has a unique timestamp) or a pointer to a reference (RStamp).

$$\text{addr} = \text{TStamp num} \mid \text{RStamp num}$$

Given a list of root values, we can compute, using to\_addr, the root addresses. Note that this function does not recurse into the values inside Block, because it only wants to collect the immediate addresses of these values.

```

to_addr []  $\stackrel{\text{def}}{=} \emptyset$ 
to_addr (Block ts v0 v1 :: xs)  $\stackrel{\text{def}}{=} \{ \text{TStamp } ts \} \cup \text{to\_addr } xs$ 
to_addr (RefPtr ref :: xs)  $\stackrel{\text{def}}{=} \{ \text{RStamp } ref \} \cup \text{to\_addr } xs$ 

```

The following value kinds do not have addresses in this representation. We will deal with these in a different way below.

```

to_addr (Number v4 :: xs)  $\stackrel{\text{def}}{=} \text{to\_addr } xs$ 
to_addr (Word64 v5 :: xs)  $\stackrel{\text{def}}{=} \text{to\_addr } xs$ 

```

Once the root addresses have been collected, we can neatly compute the set of all reachable addresses using the reflexive transitive closure (\*) of a next-relation, which is defined further down.

```

reachable_v refs blocks roots  $\stackrel{\text{def}}{=} \{ y \mid \exists x. x \in \text{roots} \wedge (\text{next refs blocks})^* x y \}$ 

```

The next-relation relates an address *a* to the addresses that are one-step reachable from *a*. Here blocks is a mapping from timestamps to Block values.

```

next refs blocks (TStamp ts) r  $\stackrel{\text{def}}{=} r \in \text{block\_to\_addrs blocks } ts$ 
next refs blocks (RStamp ref) r  $\stackrel{\text{def}}{=} r \in \text{ptr\_to\_roots refs } ref$ 
block_to_addrs blocks ts  $\stackrel{\text{def}}{=} \text{case lookup } ts \text{ blocks of}$ 
  | Some (Block _ _ vs)  $\Rightarrow \text{to\_addr } vs$ 
  | _  $\Rightarrow \emptyset$ 
ptr_to_roots refs p  $\stackrel{\text{def}}{=} \text{case lookup } p \text{ refs of}$ 
  | Some (ValueArray vs)  $\Rightarrow \text{to\_addr } vs$ 
  | _  $\Rightarrow \emptyset$ 

```

With these functions we can state the set of all reachable addresses as follows.

$$\text{reachable\_v refs blocks (to\_addr roots)}$$

Next we need to compute the heap space consumed by a heap element at a specific address using size\_of\_addr. The flat\_measure function will be explained below.

```

size_of_addr lims refs blocks (TStamp ts)  $\stackrel{\text{def}}{=} \text{case lookup } ts \text{ blocks of}$ 
  Some (Block _ _ vs)  $\Rightarrow 1 + |vs| + \text{flat\_measure lims } vs$ 
  | _  $\Rightarrow 0$ 
size_of_addr lims refs blocks (RStamp p)  $\stackrel{\text{def}}{=} \text{case lookup } p \text{ refs of}$ 
  None  $\Rightarrow 0$ 
  | Some (ValueArray vs)  $\Rightarrow 1 + |vs| + \text{flat\_measure lims } vs$ 
  | Some (ByteArray _ bs)  $\Rightarrow |bs| \text{ div } (\text{arch\_size lims div } 8) + 2$ 

```

In the above definition, we can see that an address of a Block *t n vs* has size  $1 + |vs| + \text{flat\_measure lims } vs$ . Here 1 is the for the header of the heap element; |vs| is for the length of the payload of the heap element; and flat\_measure lims vs is to account for the heap elements that are immediately reachable from this block, but have no address. The definition of flat\_measure, shown in Figure 6, counts Block and RefPtr values as having zero size, because they are already counted elsewhere.

Now we have a way to compute the set of reachable addresses and to compute size of a heap element at each address. Our final

```

flat_measure lims []  $\stackrel{\text{def}}{=} 0$ 
flat_measure lims (x :: y :: ys)  $\stackrel{\text{def}}{=} \text{flat\_measure lims } [x] + \text{flat\_measure lims } (y :: ys)$ 
flat_measure lims [Word64 v0]  $\stackrel{\text{def}}{=} 3$ 
flat_measure lims [Number i]  $\stackrel{\text{def}}{=} \text{if small\_num lims.arch\_64\_bit } i \text{ then } 0 \text{ else bignum\_size lims.arch\_64\_bit } i$ 
flat_measure lims [Block v13 v14 v15]  $\stackrel{\text{def}}{=} 0$ 
flat_measure lims [CodePtr v16]  $\stackrel{\text{def}}{=} 0$ 
flat_measure lims [RefPtr v17]  $\stackrel{\text{def}}{=} 0$ 

```

**Figure 6: The definition of flat\_measure**

definition makes use of `sum_img` which sums the application of a given function  $f$  to all elements of a finite set  $s$ .

$$\text{sum\_img } f \ s \stackrel{\text{def}}{=} \text{fold\_finite\_set } (\lambda a \ b. f \ a + b) \ s \ 0$$

The top-level definition of the new heap measure is the following. This definition sums the size of all `Word64` and large `Number` values in the roots using `flat_measure`. This is added to `sum_img` of `size_of_addr` applied to every reachable address in the heap.

$$\text{flat\_size\_of } \text{lims } \text{refs } \text{blocks } \text{roots} \stackrel{\text{def}}{=} \text{flat\_measure lims roots} + \text{sum\_img (size\_of\_addr lims refs blocks) (reachable\_v refs blocks (to\_addr roots))}$$

Even though this definition is very different in formulation from the original `size_of`, shown in Figure 5, it computes the same number.

## 4 IMPROVING ON SIZE\_OF

To illustrate the challenges of reasoning about `size_of`, consider the following reordering property:

$$\text{size\_of } [x, y] \ \text{refs LN} = \text{size\_of } [y, x] \ \text{refs LN}$$

Intuitively, this property must hold for a measuring function as the values considered are the same. However, with `size_of` of both sides of the equality might perform completely different traversals:

$$\begin{aligned} \text{size\_of } [y] \ \text{refs LN} &= (n_{y1}, \text{refs}_{y1}, \text{seen}_{y1}) \wedge \\ \text{size\_of } [x] \ \text{refs LN} &= (n_{x1}, \text{refs}_{x1}, \text{seen}_{x1}) \wedge \\ \text{size\_of } [y] \ \text{refs}_{x1} \ \text{seen}_{x1} &= (n_{y2}, \text{refs}_{y2}, \text{seen}_{y2}) \wedge \\ \text{size\_of } [x] \ \text{refs}_{y1} \ \text{seen}_{y1} &= (n_{x2}, \text{refs}_{x2}, \text{seen}_{x2}) \Rightarrow \\ (n_{y1} + n_{x2}, \text{refs}_{x2}, \text{seen}_{x2}) &= (n_{x1} + n_{y2}, \text{refs}_{y2}, \text{seen}_{y2}) \end{aligned}$$

This mismatch between applications exposes the following problems:

- There is no straightforward relation between the two measurements of  $[x]$  (or those of  $[y]$ ) as `size_of` is applied to different arguments.
- All blocks in  $[x]$  and  $[y]$  with the same timestamps must have the same contents; otherwise, the order in which blocks are counted will affect the result due to aliasing mitigation.

These issues can be overcome by introducing well-formedness conditions on  $[x]$  and  $[y]$ , and by generalizing the property statement to one more suited for induction (e.g. list permutations). However, these kinds of hurdles appear more often than one might want for such a crucial function.

In stark contrast, reordering can be trivially proved for `flat_size_of`. First, a call to `flat_measure` traverses a list to add non-root values, and is thus unaffected by permutations. Similarly, the initial root set computed by `to_addr` is the union of all addresses in the list of values and is again unaffected by reordering. Therefore, the remaining application of `sum_img` is being applied to the same arguments.

This ease of reasoning is what makes `flat_size_of` better suited for proofs of space safety as shown in the rest of this section.

## 4.1 A layout for space safety proofs

As mentioned before, to prove the space safety of a `DATALANG` program one must show the preservation of `safe_for_space` through its evaluation (Section 2.2). As most `DATALANG` programs are composed of multiple recursive functions, it is often necessary to separately prove space safety for some of them. To prove a function is space safe, one generally needs three kinds of assumptions:

- (A1) The current space consumption is below the limits or roughly  $\text{size\_of} + M \leq \text{heap\_limit}$ , where  $M$  is any extra space the function body needs.
- (A2) A description of the arguments to the function, e.g., a list-shaped block, a number within 0 and 255, among others.
- (A3) That the function is defined in  $s$ .code and its body corresponds with the code being evaluated

Resulting in the following layout:

$$\begin{aligned} \vdash A1 \wedge A2 \wedge A3 \wedge \\ \text{s.safe\_for\_space} \wedge \\ \text{evaluate (fun\_body, s)} = (\text{res}, s') \Rightarrow \\ s'.\text{safe\_for\_space} \end{aligned}$$

Proofs are by complete induction on the semantic clock and symbolic evaluation of the function body. Assumption (A2) should allow the evaluation of most of the function body. Moreover, intermediate updates to `safe_for_space` can be resolved using (A1). Once the recursive call is reached, assumption (A3) replaces the function call with the function's body such that the inductive hypothesis can be applied. At this point in the proof, assumptions must be established again for the state at the function call. (A3) is trivial as  $s$ .code does not change. (A2) might require work, but well-formed function code correctly operates on its values and thus provides good arguments. The proof of (A1) shown below is where things are more likely to become tricky:

$$\begin{aligned} \vdash \dots \\ \text{size\_of\_heap } s + M \ s \leq \text{s.limits.heap\_limit} \Rightarrow \\ \text{size\_of\_heap } s' + M \ s' \leq \text{s'.limits.heap\_limit} \end{aligned}$$

Here, we must show that the space required at the recursive call ( $\text{size\_of\_heap } s' + M \ s'$ ) is still less than `heap_limit`, assuming the space was enough in the original call. This amounts to proving

that the required space decreases as the function recurses:

$$\vdash \dots \Rightarrow \text{size\_of\_heap } s' + M s' \leq \text{size\_of\_heap } s + M s$$

This follows the intuition that function calls should take either progressively less space, or require an extra amount of memory bounded by  $M$ .

## 4.2 A tail recursive example

Consider a hypothetical tail-recursive function `ftail` with the following features:

- Takes a list of numbers as argument.
- Operates over the head of the list consuming constant space.
- Makes a tail-recursive call with the tail of the list.

Now assume we wanted to prove `ftail` space safe for concrete argument  $[1, 2, 3]$ . Instantiating the proof layout from the previous section, we arrive at the proof goal shown below:

$$\begin{aligned} \vdash & \text{size\_of\_heap } s + C \leq s.\text{limits.heap\_limit} \wedge \\ & \text{lookup "ftail"s.code} = \text{Some ftail\_body} \wedge \\ & s.\text{locals} = \\ & \quad \{ 0 \mapsto \text{Block 8 cons\_tag [Number 1,} \\ & \quad \quad \text{Block 7 cons\_tag [Number 2, \dots]]} \} \wedge \\ & s.\text{safe\_for\_space} \wedge \\ & \text{evaluate (ftail\_body, } s) = (res, s') \Rightarrow \\ & s'.\text{safe\_for\_space} \end{aligned}$$

Where  $C$  is the (constant) space the function uses to operate.

Using assumptions (A1), (A2), and (A3), most of the proof can proceed by evaluation; until the tail recursive call to `ftail` is reached and we must establish assumption (A1) again:

$$\text{size\_of\_heap } s' \leq \text{size\_of\_heap } s$$

Which by definition of `size_of_heap` and the abbreviation of `extract_stack s.stack ++ global_to_vs s.global` as `rest` simplifies to:

$$\begin{aligned} & \text{size\_of } ([\text{Block 7 cons\_tag [Number 2, \dots]]] ++ \text{rest}) \\ & \quad \text{refs LN} \\ & \leq \\ & \text{size\_of } ([\text{Block 8 cons\_tag [Number 1,} \\ & \quad \text{Block 7 cons\_tag [Number 2, \dots]]} ++ \text{rest}) \\ & \quad \text{refs LN} \end{aligned}$$

And given the equality `size_of rest refs LN = (n, refs', seen)` can be rewritten further to:

$$\text{size\_of [Block 7 cons\_tag \dots] refs' seen} \leq \text{size\_of [Block 8 cons\_tag \dots] refs' seen}$$

At this point, it would appear the proof is almost done, as we are essentially testing if the space occupied by a list  $([1, 2, 3])$  is greater than that of its tail  $([2, 3])$ , which it must be. However, due to `size_of`'s handling of timestamps and the fact that `seen` is symbolic, one can not show this inequality without additional assumptions. Concretely, one can think of a scenario where from the timestamps in the block only 8 is in `seen`; this will result in the

measurement being 0 at the right of the inequality and 4 on the left, a clear falsehood.

$$\begin{aligned} & 8 \in \text{seen} \wedge 7 \notin \text{seen} \wedge \dots \wedge \\ & \text{size\_of [Block 7 \dots] refs' seen} = (4, \text{refs}'', \text{seen}'') \wedge \\ & \text{size\_of [Block 8 \dots] refs' seen} = (0, \text{refs}'', \text{seen}'') \Rightarrow \\ & 4 \leq 0 \end{aligned}$$

Therefore, the proof goal must be extended with a predicate ensuring that if timestamp 8 is in `seen` it must be the case that 7 and all other subsequent timestamps in the list are also in `seen`.

Proving such predicate and all its associated lemmas takes considerable work, to the point that, similar mechanisms in existing space safety proofs take around 25% of the Theory file. The issue is further aggravated by the fact that this kind of predicates can not be easily generalized for all types of values and must be re-written every time a new type is used.

If we were to switch our reasoning to `flat_size_of` our proof goal could be greatly simplified:

$$\begin{aligned} & \text{flat\_size\_of refs blocks ([Block 7 \dots] ++ \text{rest})} \\ & \leq \\ & \text{flat\_size\_of refs blocks ([Block 8 \dots] ++ \text{rest})} \end{aligned}$$

While we can no longer “drop” `rest` from the roots, `flat_size_of` more than makes up for this with its use of sets and relations to represent the reachable memory. To showcase this, consider the following lemma:

$$\begin{aligned} & \text{flat\_measure } \text{lims } x = \text{flat\_measure } \text{lims } x \wedge \\ & \text{reachable\_v refs blocks (to\_addrs } x) \subseteq \\ & \quad \text{reachable\_v refs blocks (to\_addrs } y) \Rightarrow \\ & \text{flat\_size\_of } \text{lims refs blocks } x \leq \\ & \quad \text{flat\_size\_of } \text{lims refs blocks } y \end{aligned}$$

Which states that if the reachable set of addresses from two roots  $x$  and  $y$  are subsets, then the space measurement of  $x$  done by `flat_size_of` must be less than that of  $y$ . Using this lemma the proof goal becomes trivial:

$$\begin{aligned} & \{\text{TStamp 7}\} \cup \text{reachable\_v } \dots (\text{to\_addrs } \text{rest}) \\ & \subseteq \{\text{TStamp 8, TStamp 7}\} \cup \text{reachable\_v } \dots (\text{to\_addrs } \text{rest}) \end{aligned}$$

This would conclude the proof with little more than basic set reasoning.

Is this ease of reasoning in the presence of (possibly) aliased values what makes `flat_size_of` a suitable measuring function for a cost semantics. In particular, the reachability-based approach to gathering live data aids the function, and its reasoning, to not be concerned with “where” a value is or how it is structured, and focus solely in its effect on the space measurement. In contrast, reasoning about `size_of` constantly requires additional safeguards structural guarantees that should only concern the memory model.

## 5 SOUNDNESS

At the time of submission, we have not yet started proving soundness of the new version of the cost semantics. However, our intention is to prove soundness of the new cost semantics so that all future uses of CakeML's space cost semantics can make use of this improved formulation that we have presented in this paper.

There are two options for proving soundness: (1) soundness of the new formulation can be proved w.r.t. the original formulation;

or (2) the old formulation could be replaced by the new formulation in the definition of the `DATA LANG` semantics. Option (1) is neatly self-contained to mostly a proof about the relationship between the old and the new space measures. However, if the new formulation is to truly replace the old one, then option (2) is the right way to go, even though it requires redoing some fiddly proofs in the middle of the correctness proofs of the CakeML compiler. At the time of writing, we are leaning towards option (2), since we do not want the old one to become a burden for proof maintenance.

We expect to have soundness of the new version of the cost semantics proved by IFL's post-proceedings deadline.

## 6 RELATED WORK

Work on verified cost semantics of verified compiler is available for the CompCert [7] and CakeML [6] compilers. Carbonneaux et al. [4] develop a source level logic for stack space reasoning that translates to the CompCert compiler output. Besson *et al.* extends CompCert's memory model with finite memory and integer pointers in CompCertS [1–3]; which allows for memory usage estimates of C functions that are proven to be bounds of the compiled code. The cost semantics of the CakeML compiler [5] is to our knowledge the only verified costs semantics for a high-level garbage-collected language and as such is a good candidate for further research in the topic.

## 7 CONCLUSION

This paper has proposed a new flat reachability based heap space cost measure for CakeML's verified space cost semantics. Early experiments suggests that the new formulation of the heap measure is significantly more pleasant to use in proofs of space safety. We plan to replace the original formulation of the heap measure with this new one. The hope is that the new formulation will make reasoning of space cost scale beyond simple examples.

## REFERENCES

- [1] Frédéric Besson, Sandrine Blazy, and Pierre Wilke. 2014. A Precise and Abstract Memory Model for C Using Symbolic Values. In *Programming Languages and Systems*, Jacques Garrigue (Ed.). Springer International Publishing, Cham, 449–468.
- [2] Frédéric Besson, Sandrine Blazy, and Pierre Wilke. 2015. A Concrete Memory Model for CompCert. In *Interactive Theorem Proving*. Springer International Publishing, Cham, 67–83.
- [3] Frédéric Besson, Sandrine Blazy, and Pierre Wilke. 2019. CompCertS: A Memory-Aware Verified C Compiler Using a Pointer as Integer Semantics. *Journal of Automated Reasoning* 63, 2 (01 Aug 2019), 369–392.
- [4] Quentin Carbonneaux, Jan Hoffmann, Tahina Ramananandro, and Zhong Shao. 2014. End-to-end Verification of Stack-space Bounds for C Programs. *SIGPLAN Not.* 49, 6 (June 2014), 270–281.
- [5] Alejandro Gómez-Londoño, Johannes Åman Pohjola, Hira Taqdees Syeda, Magnus O. Myreen, and Yong Kiam Tan. 2020. Do you have space for dessert? a verified space cost semantics for CakeML programs. *Proc. ACM Program. Lang.* 4, OOPSLA (2020), 204:1–204:29. <https://doi.org/10.1145/3428272>
- [6] Ramana Kumar, Magnus O. Myreen, Michael Norrish, and Scott Owens. 2014. CakeML: a verified implementation of ML. In *The 41st Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '14, San Diego, CA, USA, January 20–21, 2014*, Suresh Jagannathan and Peter Sewell (Eds.). ACM, 179–192. <https://doi.org/10.1145/2535838.2535841>
- [7] Xavier Leroy. 2009. Formal Verification of a Realistic Compiler. *Communications of the ACM* 52, 7 (2009). <https://doi.org/10.1145/1538788.1538814>
- [8] Scott Owens, Magnus O. Myreen, Ramana Kumar, and Yong Kiam Tan. 2016. Functional Big-Step Semantics. In *European Symposium on Programming (ESOP) (Lecture Notes in Computer Science)*, Peter Thiemann (Ed.). Springer, 589–615.
- [9] Yong Kiam Tan, Magnus O. Myreen, Ramana Kumar, Anthony C. J. Fox, Scott Owens, and Michael Norrish. 2019. The Verified CakeML Compiler Backend. *J. Funct. Program.* 29 (2019), e2. <https://doi.org/10.1017/S0956796818000229>