

# Garbage Collection using a Finite Liveness Domain

Aman Bansal\*  
IIT Bombay  
India  
amanb@cse.iitb.ac.in

Saksham Goel\*  
IIT Bombay  
India  
saksham.goel@cse.iitb.ac.in

Preey Shah\*  
IIT Bombay  
India  
preey@cse.iitb.ac.in

Amitabha Sanyal  
IIT Bombay  
India  
as@cse.iitb.ac.in

Prasanna Kumar  
IIT Bombay  
India  
prasannak@cse.iitb.ac.in

## Abstract

Functional languages manage heap data through garbage collection. Since static analysis of heap data is difficult, garbage collectors conservatively approximate the *liveness* of heap objects by *reachability* i.e. every object that is reachable from the root set is considered live. Consequently, a large amount of memory that is reachable but not used further during execution is left uncollected by the collector.

Earlier attempts at liveness-based garbage collection for languages supporting structured types were based on analyses that considered arbitrary liveness values, i.e. they assumed that any substructure of the data could be potentially live. This made the analyses complex and unscalable. However, functional programs traverse structured data like lists in identifiable patterns. We identify a set of eight usage patterns that we use as liveness values. The liveness analysis that accompanies our garbage collector is based on this set; liveness arising out of other patterns of traversal are conservatively approximated by this set.

This restriction to a small set of liveness values reaps several benefits – it results in a simple liveness analysis which scales to much larger programs with minimal loss of precision, enables the use of a faster collection technique, and is extendable to higher-order programs.

Our experiments with a purely functional subset of Scheme show a reduction in the analysis time by orders of magnitude. In addition, the minimum heap size required to run programs is comparable with a liveness-based collector with unrestricted liveness values, and in situations where

memory is limited, the garbage collection time is lower than its reachability counterpart.

**CCS Concepts:** • Software and its engineering → Garbage collection; Functional languages; • Theory of computation → Program analysis.

**Keywords:** Garbage Collection, Liveness Analysis, Access Paths, Functional Programming

## ACM Reference Format:

Aman Bansal, Saksham Goel, Preey Shah, Amitabha Sanyal, and Prasanna Kumar. 2020. Garbage Collection using a Finite Liveness Domain. In *Proceedings of the 2020 ACM SIGPLAN International Symposium on Memory Management (ISMM '20)*, June 16, 2020, London, UK. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3381898.3397208>

## 1 Introduction

Modern programming languages have been increasingly adopting automated memory management through garbage collection because of the safety and precision that it offers. Functional languages, where all memory is dynamically allocated on the heap and accessed through references from the stack, have always had garbage collectors as an essential part of their runtime support. Therefore, it is important to investigate possibilities of increasing the effectiveness of garbage collection, especially for use in functional programming languages. Most garbage collectors work by marking heap cells that are reachable from the activation stack and registers (called the *root set*) and collecting unmarked cells as garbage. However, studies across functional programming languages [3, 15, 22] have shown that programs often allocate considerable amount of memory that remains reachable beyond their last use, i.e. even after they cease to be live. Thus a liveness-based garbage collector (LGC) can be expected to collect more garbage in the same program state than a reachability based collector (RGC). Use of a LGC results in several benefits—programs run with lesser memory, there are fewer collections, and, in the case of copying collectors, which copy useful cells instead of garbage, the collections are, on the average faster, since fewer memory cells have

\*Sorted lexicographically by last name

Publication rights licensed to ACM. ACM acknowledges that this contribution was authored or co-authored by an employee, contractor or affiliate of a national government. As such, the Government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only.

ISMM '20, June 16, 2020, London, UK

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-7566-5/20/06...\$15.00

<https://doi.org/10.1145/3381898.3397208>

to copied. Therefore, not only do programs run with lesser heap space, they often run in lesser time.

The primary reason why LGCs have not found acceptance in spite of their claimed and empirically demonstrated benefits [9] is the difficulty of liveness analysis itself<sup>1</sup>. When the analysis involves scalars only, liveness has two values—may-be-live or dead. When the language supports structured data, as it does in the case of functional programming languages with algebraic data types, any reasonably precise analysis has to go beyond this binary classification and identify parts of the structure that are live. Liveness of structured data is commonly represented by sets of paths in the heap [3, 11, 12, 17]. These paths represent parts of the data that can be potentially traversed in the future. However, manipulating arbitrary sets of paths in an analysis that also strives for precision is cumbersome, and the resulting methods do not scale to programs consisting of more than a few functions. The complexity of analysis also inhibits analysis of programs with higher order functions, a staple of most functional programs. Additionally, since the garbage collector has to refer to liveness of root set variables at runtime, the complexity of liveness values and their representation adds overheads to the garbage collection algorithm. This partially nullifies the apparent advantages of a liveness based collection. Finally, the storage of liveness data has its own overheads, as this information has to be stored for every potential garbage collection point.

It is our observation that most functional programs that use datatypes such as lists and trees access the heap in regular identifiable patterns. Therefore, in this paper, we restrict ourselves to a small fixed set of liveness values, each value representing a common pattern of heap accesses. All other values of liveness are safely approximated to a value from this fixed set. This simplifies liveness analysis and leads to dramatic improvements in the scalability of the method. In addition, it also allows us to handle higher-order functions. While the approximation introduced could have led to a reduction in the collected garbage, we largely make up for the loss in precision in two ways, both attributable to our choice of the small fixed set of liveness values.

Firstly, for the body of a function, we maintain separate liveness information for each distinct calling context, where the context includes the liveness value at the call-site. This is different from most analyses<sup>2</sup>. We exploit this added precision by maintaining the calling context at runtime, and using the liveness corresponding to the context in case of a garbage collection in the body of the function.

<sup>1</sup>[9] obtained their liveness information from program traces and not from a static analysis.

<sup>2</sup>While most analyses avoid infeasible interprocedural paths by maintaining separate information regarding call-sites in their calling context, the analysis information (the dataflow values) at the call-sites are generally not part of the calling context and are merged while analysing the body of a function. [20] is one of the few exceptions.

Secondly, our garbage collector has a marking scheme that leverages the fact that the liveness values can be ordered by inclusion. We use this fact in two ways: First we are able to avoid a drawback of most liveness based collectors, namely repeated revisits to the same memory location while marking. And secondly, we are able to do a breadth first traversal during marking, thereby avoiding the stack-space cost of a depth first search.

**A Motivating Example.** As a motivating example, consider the program shown in Figure 1 written in a restricted subset of Scheme. The program defines two functions, **append** (lines 1-8) and **main** (lines 9-13). The function **append** concatenates its list arguments, and **main** creates a pair of lists using **makelist** (definition omitted) and uses **append** to concatenate the pair. The labels denoted by  $\pi$  that annotate some of the expressions refer to *program points* and are not part of the program. Semantically, a program point represents an instant of time just before the execution of the expression that it labels.

For concreteness, assume that the list  $x$  passed to **append** from **main** has four elements indicated by the triangles A, B, C and D in that order (Figure 2). The reference to this list from the activation frame  $AF_1$  is called  $x_1$ ; we differentiate between incarnations of the same variable in different function calls using subscripts. The triangles in the figure indicate structures whose details are not important. Rectangles indicate **cons** cells, the basic building blocks of lists with pointers to the head (**car**) and the tail (**cdr**) of the list. The crossed rectangle is an empty list.

The snapshot in the figure follows a sequence of events in which **main** calls **append**, and **append**, in turn, calls itself recursively four more times, of which the last three calls have returned. As a result, the last two elements of  $x$  are appended to the list  $y$  and the result bound to  $t$  (the  $t_2$  incarnation in  $AF_2$ ). Also assume that the currently active call to **append** has progressed further and invoked the **cons** operation at line 7. If this operation, which requires a fresh **cons** cell to be allocated, fails because the allocator has run out of memory, the garbage collector is invoked. At this time, the topmost activation frame,  $AF_2$ , on the stack contains references represented by the set of variables  $\{x_2, y_2, d_2, t_2, a_2\}$ , and the next frame,  $AF_1$ , contains the set  $\{x_1, y_1, d_1\}$ . In garbage collection parlance, these are called *root set variables*. Note that  $c$  is bound to a boolean value and not a reference.

Our liveness-based garbage collector begins a traversal from, say, the topmost frame  $AF_2$ . It consults the liveness of the root set variables at program point  $\pi_4$ , because this is where the program will resume after garbage collection. Our liveness analysis would have discovered the liveness of the root set variables:  $x_2$ ,  $y_2$  and  $d_2$  as “not live” (denoted  $\perp$ ) as they do not appear in the function body after  $\pi_4$ , and the liveness of  $t_2$  and  $a_2$  as “fully live” (denoted  $\top$ ), which means that *all cons* cells reachable from these variables are likely to

```

1 (define (append x y)
2   (let c ← (null? x) in                ; if x is empty
3     (if c (return y)                   ; return y
4       (let d ← (cdr x) in              ; else bind d to tail of x
5         (let t ← (append d y) in       ; and t to (append d y)
6           π3: (let a ← (car x) in      ; bind a to head of x
7             π4: (let ans ← (cons a t) in ; return a list
8               (return ans))))))))) ; with head a, tail t
9 (define (main)
10  (let p ← (makelist ...) in            ; make a list and bind to p
11    π1: (let q ← (makelist ...) in     ; similarly for q
12      π2: (let final ← (append p q) in ; return (append p q)
13        (return final))))))

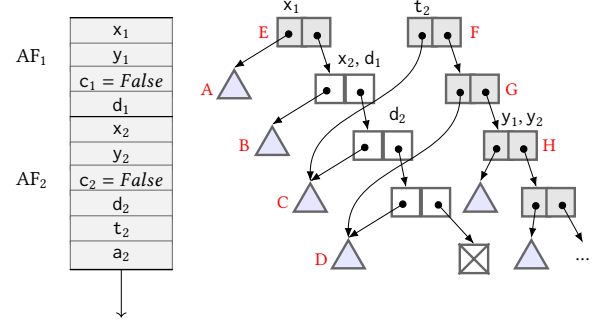
```

**Figure 1.** An example annotated program to illustrate our method.

be used in the future. This is understandable since both these variables and the structures reachable from them are part of the result of the program, and we assume that the agent that invoked this program would use the result in its entirety. Assuming that the underlying garbage collector is a copying collector (though the same idea can also be replayed on other collectors), it would scavenge all the cells in the structures B, C and D, the cons cells F, G and H, and the entire structure under H. Similarly, for garbage collecting the root set in  $AF_1$ , the liveness at program point  $\pi_3$  is consulted. This is because  $\pi_3$  is the program point where the call corresponding to  $AF_1$  would resume after the return of the inner call to **append**. Because of the expression **(car x)** that occurs in the function after  $\pi_3$ , the liveness analysis will determine that the cell pointed by  $x_1$  and the whole substructure under its **car** to be live. We denote this pattern of liveness ("root-and-car" live) as  $\tau_{0_e}$ . The variables  $d_1$  and  $y_1$  on the other hand have  $\perp$  as their liveness values. Traversal from  $x_1$  results in the scavenging of E and all the cells in A. The heap data that will be scavenged in this round of garbage collection are indicated by shading, the remaining cells that are not shaded, while reachable, can be garbage collected and reused for subsequent allocation.

This example shows the benefits of liveness-based garbage collection. In fact, a reachability based collector would not have been able to reuse any of the **cons** cells of the original lists and would have required four extra cells to append the two lists. Carried to its limit, if we can tolerate a garbage collection for every **cons** cell allocated by the program, then the two lists, however long, can be appended entirely by reusing the cells of the original lists.

This example uses just three liveness values,  $\perp$ ,  $\top$  and  $\tau_{0_e}$ . Our complete set consists of five more liveness values. As an example of yet another liveness value, consider a list passed to a function that calculates its length. Such a list is



**Figure 2.** A snapshot of the stack and the heap during execution. Stack references and heap data are assumed to be connected through common variable names instead of arrows. Only the activation frames of **append** are shown.

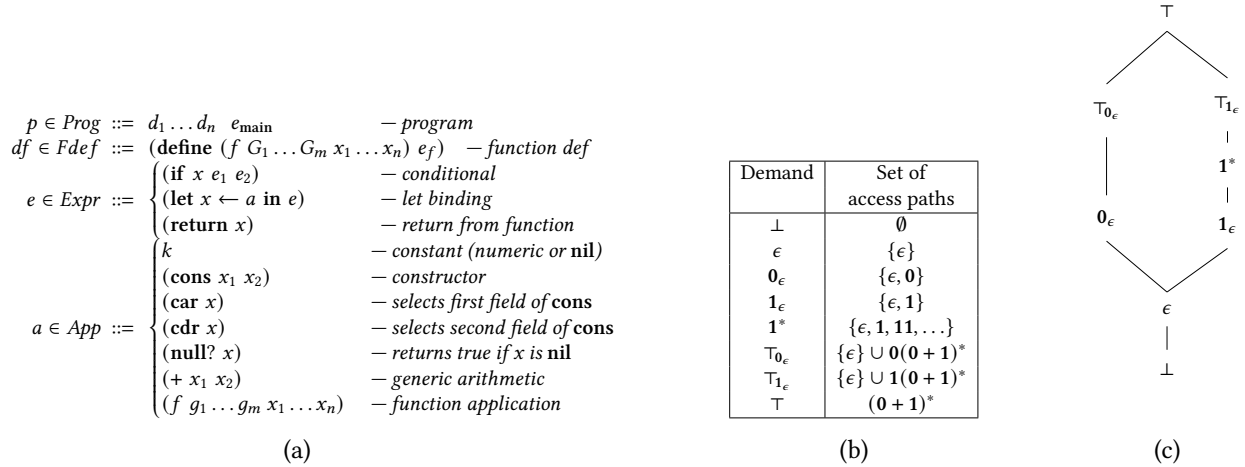
traversed without touching any of the elements. We denote such a liveness value ("spine-only live") as  $1^*$ . The complete set of liveness values that we use is described in Section 2.

The main contributions of the paper are:

1. We propose a liveness-based garbage collection method for a restricted and pure subset of Scheme that allows higher order functions and has lists as the primary data structure.
2. The garbage collection relies on a liveness analysis that is simple and scalable. The scalability derives from our choice of a small fixed set of liveness values that embodies common patterns of heap traversal. Interestingly, experiments on our benchmarks demonstrate that the limited set of liveness values is able to capture heap liveness with a precision that is comparable to a previous method [3] that employs arbitrary liveness values.
3. The small set of liveness values allows us to separate out the liveness of function bodies based on contexts. The context information is carried over to runtime at negligible space costs. However, it enhances the precision of liveness analysis enabling more cells to be identified as garbage.
4. Our choice of liveness patterns also enables the design of a garbage collection mechanism that avoids revisits of memory cells. It also enables breadth-first traversal of the heap, saving the space-cost of a depth-first traversal.
5. We have implemented an interpreter and a copying collector based on our ideas. Our experiments confirm that compared to to an LGC based on arbitrary liveness patterns, our analysis is orders of magnitude faster, comparable in precision and in the number of garbage collections and on the average faster.

However, our study also has its limitations:

1. As mentioned earlier, our method in its current form is limited to lists, and because lists in Scheme are heterogeneous, also to binary trees. Lists and trees are widely used



**Figure 3.** (a) The syntax of the language. (b) Demands and their meaning. (c) The lattice of demand patterns ordered by inclusion.

in functional programs. Other data types in our benchmarks were encoded in terms of lists.

2. While the chosen liveness values have been surprisingly effective for the programs considered, we do not know whether these are general enough to serve a wider range of benchmarks: (i) programs with possibly different heap traversal patterns, (ii) programs using different algebraic data types, and (iii) programs that are larger in size.

## 2 The Target Language

Figure 3(a) shows the target language. It is an eager, higher-order language with a syntax similar to Scheme. For ease of presentation, we restrict the language to Administrative Normal Form (ANF) [23]. In this form, function calls can only have variables as arguments. To avoid dealing with scope-shadowing, we assume that all variables in a program are distinct. Neither of these restrictions affects the expressibility of our language. As mentioned earlier, we sometimes use a label  $\pi$  to refer to the program point before an expression  $e$ , such as  $\pi:e$ . However the label is not part of the language.

A program in our language is a collection of function definitions followed by a call to a distinguished function called **main**. Applications (denoted by the syntactic category *App*) consist of functions or operators applied to their arguments. Constants are regarded as 0-ary functions. Expressions (*Expr*) are either an **if** expression, a **let** expression, or a **return** expression. A **let** expression (**let**  $x \leftarrow a$  **in**  $e$ ) creates a local binding for  $x$  and evaluates the expression  $e$  with this binding. The **return** keyword is used to mark the end of a function.

All functions are named, the language does not support lambda expressions. However, functions can be higher-order, i.e. they can take other functions as arguments. In the function definition (**define**  $(f \ G_1 \dots G_m \ x_1 \dots x_n) \ e_f$ ),  $G_1$  to  $G_m$

are function parameters,  $x_1$  to  $x_n$  are non-function parameters, and  $e_f$  is the function body. Functions have to be fully applied, they cannot be applied to some but not all of their arguments. Partial application is allowed in many functional languages (e.g. Haskell), but not in Scheme. However, while describing liveness analysis, we shall refer to partial applications such as  $(f \ g_1 \dots g_m)$ . This is a first order function that has been derived by applying a higher-order function to all and only its function arguments. Note that partial applications will be used only to describe the analysis. It is not permitted by the language and cannot occur in programs.

## 3 Liveness and its Analysis

A scalar variable is live if there is a possibility of its value being used in future computations and dead if it is definitely not used. In contrast, variables that are bound to heap-allocated structured data need a richer model than the may-be-live or dead model of classical liveness—a model which talks about future accesses of parts of a structure. In this paper, we restrict ourselves to lists. Borrowing notations from [3], we represent an access using the **car** selector by **0** and the **cdr** selector by **1**. An *access path* is any sequence of selectors including the empty sequence  $\epsilon$ . Given a structure, a *demand* is a set of access paths that represent how the structure is possibly accessed. As an example, if  $x$  is a variable with value  $(\text{cons } 1 \ (\text{cons } 2 \ 3))$ . then the accesses of the values in the substructure  $(\text{cons } 2 \ 3)$  with respect to  $x$  is given by the demand  $\{\epsilon, 1, 10, 11\}$ . The access paths in this demand start with the root **cons** cell of  $x$  (represented by  $\epsilon$ ), the path to access the **cons** cell forming the root of  $(\text{cons } 2 \ 3)$  (represented by **1**), and the paths to the leaves **2** and **3** (**10** and **11**). Notice that demands are prefix-closed; a node in a structure cannot be accessed from a root-set variable without going through its ancestors.

$E(\sigma)$	$\sigma$							
	$\perp$	$\epsilon$	$\mathbf{0}_\epsilon$	$\mathbf{1}_\epsilon$	$\mathbf{1}^*$	$\top_{0_\epsilon}$	$\top_{1_\epsilon}$	$\top$
$\epsilon \cup \mathbf{0}\sigma$	$\epsilon$	$\mathbf{0}_\epsilon$	$\top_{0_\epsilon}$	$\top_{0_\epsilon}$	$\top_{0_\epsilon}$	$\top_{0_\epsilon}$	$\top_{0_\epsilon}$	$\top_{0_\epsilon}$
$\epsilon \cup \mathbf{1}\sigma$	$\epsilon$	$\mathbf{1}_\epsilon$	$\top_{1_\epsilon}$	$\top_{1_\epsilon}$	$\mathbf{1}^*$	$\top_{1_\epsilon}$	$\top_{1_\epsilon}$	$\top_{1_\epsilon}$
$\bar{\mathbf{0}}\sigma$	$\perp$	$\perp$	$\epsilon$	$\perp$	$\perp$	$\top$	$\perp$	$\top$
$\bar{\mathbf{1}}\sigma$	$\perp$	$\perp$	$\perp$	$\epsilon$	$\mathbf{1}^*$	$\perp$	$\top$	$\top$

**Figure 4.** The approximation function  $\mathcal{S}$ .

While a demand can in general be an arbitrary set of access paths, recall that the claim of this paper is that restriction to a carefully chosen small set of demands makes liveness analysis more effective with only marginal loss of precision. Towards this we have chosen the following set of demands  $\Sigma = \{\perp, \epsilon, \mathbf{0}_\epsilon, \mathbf{1}_\epsilon, \mathbf{1}^*, \top_{0_\epsilon}, \top_{1_\epsilon}, \top\}$ . These correspond to the sets of access paths shown in Figure 3(b). These demands form a lattice that is ordered by set inclusion, and the join operation on this lattice is denoted by  $\sqcup$ . Any other demand that may arise during liveness analysis is approximated by a demand from this fixed set  $\Sigma$ . Notice that each of the sets is prefix closed. The element  $\mathbf{0}_\epsilon$ , for example, is a demand that denotes the set of access paths  $\{\epsilon, \mathbf{0}\}$  and not simply  $\{\mathbf{0}\}$ .

A *liveness* of a variable at a program point is a demand that approximates all possible future accesses of the data bound to that variable. A variable  $x$  having a liveness of  $\mathbf{0}_\epsilon$  at program point  $\pi$  means that future accesses to  $x$  starting from  $\pi$  are limited to the root of the structure of  $x$  and (the referent of) its **car** field. A *liveness environment* is a mapping from variables to demands, for example  $\{x \mapsto \mathbf{0}_\epsilon, y \mapsto \epsilon, z \mapsto \perp\}$ . The liveness of  $z \mapsto \perp$  indicates that  $z$  is dead. Note that we use the empty sequence of selectors  $\epsilon$  to stand for the root of a structure. Thus base values such as integers have  $\epsilon$  as the only substructure.

Access paths are denoted by  $\alpha$  and demands by  $\sigma$ . Given two access paths  $\alpha_1$  and  $\alpha_2$ , we use  $\alpha_1\alpha_2$  to denote their concatenation. We extend this notation to the concatenation of a symbol with a demand:  $\mathbf{0}\sigma$  is a shorthand for  $\{\alpha \mid \mathbf{0}\alpha \mid \alpha \in \sigma\}$ . However the set of access paths represented by  $\mathbf{0}\sigma$  may not be exactly representable in  $\Sigma$  and may need to be approximated by the least element in the lattice of demands that contains it. This approximation is done by the function  $\mathcal{S}$ . Figure 4 defines  $\mathcal{S}$  for all possible demand expressions  $E(\sigma)$  that can be encountered during liveness analysis.

### 3.1 Liveness Analysis

Liveness analysis determines the liveness environment, i.e. the mapping of variables to demands, at each program point. The analysis is described in terms of the three functions  $\mathcal{A}$ ,  $\mathcal{L}$  and  $\mathcal{FS}$  shown in figure 5. Given a demand  $\sigma$  representing the future uses of an application, the function  $\mathcal{A}$  returns a liveness environment that describes the uses of the variables in the application.  $\mathcal{L}$  plays a similar role for an expression—given a demand, it returns a liveness environment for the *free variables* of the expression. Both  $\mathcal{A}$  and  $\mathcal{L}$  are backward

analyses—they transform a demand on the result of an application or an expression to demands on their arguments.

We start by explaining  $\mathcal{A}$ . Consider the application  $(\mathbf{car} \ x)$  with a demand  $\sigma$ . A program accessing  $\sigma$  part of  $(\mathbf{car} \ x)$  has to start with the root of  $x$ , select the **car**, and then access  $\sigma$  part of the resulting structure. Thus the set of paths that are accessed relative to  $x$  is  $\epsilon \cup \mathbf{0}\sigma$ . However, the demand  $\epsilon \cup \mathbf{0}\sigma$  may not be representable in our domain of patterns and may have to be approximated by  $\mathcal{S}$ . This gives the liveness environment of  $(\mathbf{car} \ x)$  as  $\{x \mapsto \mathcal{S}(\epsilon \cup \mathbf{0}\sigma)\}$ . The application  $(\mathbf{null?} \ x)$  produces the rule  $\{x \mapsto \epsilon\}$  because a list can be tested for nullity by just accessing its root cell. Similarly, a scalar is represented by a single (root) cell which contains its value. Finally, to evaluate  $(+ \ x \ y)$  for any demand, including  $\perp$ , only the roots of  $x$  and  $y$  have to be accessed. Thus the liveness environment of  $(+ \ x \ y)$  is  $\{x \mapsto \epsilon, y \mapsto \epsilon\}$ <sup>3</sup>.

To explain the rule for **cons**, notice that a demand of  $\mathbf{0}_\epsilon$  (or  $\{\epsilon, \mathbf{0}\}$  in expanded form) on  $(\mathbf{cons} \ x \ y)$  means that future uses of  $(\mathbf{cons} \ x \ y)$  will only access the root and the cell obtained by a **car** selection. Since **car** of  $(\mathbf{cons} \ x \ y)$  is  $x$  itself, it follows that the demand on  $x$  is  $\epsilon$ . Generalising this argument, if the demand on  $(\mathbf{cons} \ x \ y)$  is  $\sigma$ , then the demand on  $x$  is  $\{\alpha \mid \mathbf{0}\alpha \in \sigma\}$ . Denoting  $\{\alpha \mid \mathbf{0}\alpha \in \sigma\}$  as  $\bar{\mathbf{0}}\sigma$  and approximating this demand using  $\mathcal{S}$ , we get the liveness of  $x$  as  $\mathcal{S}(\bar{\mathbf{0}}\sigma)$ . A similar reasoning gives the demand on  $y$  as  $\mathcal{S}(\bar{\mathbf{1}}\sigma)$ , where the symbol  $\bar{\mathbf{1}}$  is defined as  $\bar{\mathbf{1}}\sigma = \{\alpha \mid \mathbf{1}\alpha \in \sigma\}$ . Observe that for the same demand  $\mathbf{0}_\epsilon$ , the demand on  $y$  is  $\mathcal{S}(\bar{\mathbf{1}}\mathbf{0}_\epsilon)$  or  $\perp$ , as it should be. The behaviour of  $\mathcal{A}$  for a function call will be explained after we describe the function  $\mathcal{L}$ .

Similar to  $\mathcal{A}$ , the function  $\mathcal{L}$  transforms a demand on an expression to a liveness environment involving the free variables in the expression. The rule for  $(\mathbf{return} \ x)$  is obvious. The liveness environment of  $(\mathbf{if} \ x \ e_1 \ e_2)$  is the pointwise union of the liveness environment that arises out of testing the condition  $x$ , and the liveness environments of  $e_1$  and  $e_2$ . The liveness environments of  $(\mathbf{let} \ x \leftarrow a \ \mathbf{in} \ e)$  is the pointwise union of the liveness environments of the application  $a$  and the expression  $e$ . The liveness environment  $L_1$  of  $e$  is determined with respect to the demand  $\sigma$  on the **let** expression, and the liveness environment of  $a$  is determined with respect to the demand on  $x$  in  $L_1$ . The liveness of  $x$  is erased from the union because the scope of  $x$  is local to  $e$ . This is similar to the backward propagation of liveness across an assignment  $x = e$  in classical dataflow analysis of imperative languages.

The liveness environment arising from a function call  $(f \ g_1 \dots g_m \ x_1 \dots x_n)$  is described using the function  $\mathcal{FS}$ .

<sup>3</sup>Interestingly, the situation would have been different for a lazy language, which does not evaluate an expression unless required. For the  $\perp$  demand, the liveness environment for  $(+ \ x \ y)$  would have been  $\{x \mapsto \perp, y \mapsto \perp\}$ .

$$\boxed{\mathcal{A} :: (\text{App} \times \text{Demand}) \rightarrow \text{LivenessEnvironment}}$$

$$\begin{aligned}
\mathcal{A}(\kappa, \sigma) &= \perp, \text{ for all constants } \kappa \text{ including } \mathbf{nil} \\
\mathcal{A}(\mathbf{null? } x, \sigma) &= \{x \mapsto \epsilon\} \\
\mathcal{A}(\mathbf{+ } x y, \sigma) &= \{x \mapsto \epsilon, y \mapsto \epsilon\} \\
\mathcal{A}(\mathbf{car } x, \sigma) &= \{x \mapsto \mathcal{S}(\epsilon \cup \mathbf{0}\sigma)\} \\
\mathcal{A}(\mathbf{cdr } x, \sigma) &= \{x \mapsto \mathcal{S}(\epsilon \cup \mathbf{1}\sigma)\} \\
\mathcal{A}(\pi: \mathbf{cons } x y, \sigma) &= \{x \mapsto \mathcal{S}(\mathbf{0}\sigma), y \mapsto \mathcal{S}(\mathbf{1}\sigma)\} \\
\mathcal{A}((f g_1 \dots g_m x_1 \dots x_n), \sigma) &= \bigcup_{i=1}^n \{x_i \mapsto \sigma_i\}, \text{ where} \\
&\quad (\sigma_1, \dots, \sigma_n) = \mathcal{FS}((f g_1 \dots g_m), \sigma)
\end{aligned}$$

$$\boxed{\mathcal{L} :: (\text{Exp} \times \text{Demand}) \rightarrow \text{LivenessEnvironment}}$$

$$\begin{aligned}
\mathcal{L}(\mathbf{return } x, \sigma) &= \{x \mapsto \sigma\} \\
\mathcal{L}(\mathbf{if } x e_1 e_2, \sigma) &= \{x \mapsto \epsilon\} \sqcup \mathcal{L}(e_1, \sigma) \sqcup \mathcal{L}(e_2, \sigma) \\
\mathcal{L}(\mathbf{let } x \leftarrow a \mathbf{in } e, \sigma) &= L, \text{ where} \\
&\quad L_1 = \mathcal{L}(e, \sigma), \\
&\quad \sigma' = L_1(x), \\
&\quad L_2 = \mathcal{A}(a, \sigma'), \text{ and} \\
&\quad L = (L_1 \sqcup L_2) - \{x \mapsto L_1(x)\}
\end{aligned}$$

$$\boxed{\mathcal{FS} :: (\text{FirstOrderFunc} \times \text{Demand}) \rightarrow (\text{Demand}_1 \times \dots \times \text{Demand}_n)}$$

$$\begin{aligned}
\mathcal{FS}((f g_1 \dots g_m), \sigma) &= (L(x_1), \dots, L(x_n)), \text{ where} \\
f &\stackrel{\text{def}}{=} (\mathbf{define } (f G_1 \dots G_m x_1 \dots x_n) e_f), \text{ and} \\
L &= \mathcal{L}(e_f[G_1/g_1, \dots, G_m/g_m], \sigma)
\end{aligned}$$

**Figure 5.** Liveness Analysis

$\mathcal{FS}$  takes a first-order function  $(f g_1 \dots g_m)^4$  and a demand  $\sigma$ , and returns a tuple of demands  $(\sigma_1, \dots, \sigma_n)$ . The function  $\mathcal{A}$  maps this tuple to the variables  $x_1 \dots x_n$  to yield the liveness environment of the application  $(f g_1 \dots g_m x_1 \dots x_n)$ .  $(\sigma_1, \dots, \sigma_n)$  is computed by applying  $\mathcal{L}$  to the body  $e_f$  of  $f$  after substituting the formal parameters  $G_1 \dots G_n$  by concrete functions  $g_1 \dots g_m$ .  $\mathcal{FS}$  yields an equation for each function  $f$  and each tuple of functions  $(g_1 \dots g_m)$  that  $f$  may be called with in the program. We have used the control flow analysis from a well-known method for *firstification* [19] to determine all possible concrete functions that a higher order function may be called with.

The equations derived from  $\mathcal{FS}$  are, in general, mutually recursive, and are solved by computing the fixed point of a series of approximations. We denote the  $n^{\text{th}}$  approximation of  $\mathcal{FS}$  applied to  $(f g_1 \dots g_m)$  as  $\mathcal{FS}^n(f g_1 \dots g_m)$ .  $\mathcal{FS}^0(f g_1 \dots g_m)$  is the function  $\lambda\sigma.(\perp, \dots, \perp)$ , the function that maps any demand to a tuple of empty demands.  $\mathcal{FS}^n(f g_1 \dots g_m)$  is computed from  $\mathcal{FS}^{n-1}(f' g'_1 \dots g'_m)$ , where  $(f' g'_1 \dots g'_m y_1 \dots y_k)$  is a function call encountered during the computation of  $\mathcal{L}(e_f[G_1/g_1, \dots, G_m/g_m])$ . In practice we use these approximation for computing  $\mathcal{FS}^n(f g_1 \dots g_m)$  for each possible demand  $\sigma$  in  $\Sigma$  and cache the results to be used for the computation of liveness for the body of functions.

<sup>4</sup>As mentioned earlier, the first order function is, in general, obtained by partially applying a higher order function to its function parameters.

The compute intensive part of the analysis is the iterative calculation of  $\mathcal{FS}$ . We use a work-list based algorithm to find the fixed point for  $\mathcal{FS}$ . The work-list is initialised with all function definitions. For any function, say  $f$ , if  $\mathcal{FS}^n \neq \mathcal{FS}^{n-1}$ , all functions which call  $f$  are pushed onto the work-list.

We calculate a bound on the number of worklist iterations with the following argument. For a particular  $\sigma$ ,  $\mathcal{FS}^n \neq \mathcal{FS}^{n-1}$  if there exists an argument  $x_i$  such that  $L^n(x_i) \neq L^{n-1}(x_i)$ , where  $L$  is as defined within the equation for  $\mathcal{FS}$  in Figure 5. Note that for this iterative procedure,  $L$  is a monotonically increasing function and hence the number of updates is bounded by the height of the lattice  $h$ . Extending this argument to all functions, it follows that  $\mathcal{FS}$  can update at most  $kN_f \times h$  times per demand, where  $k$  is the average number of arguments in a function, before reaching the fixed point. The number of functions which call any function is loosely bounded by  $N_f$ , which proves that the worklist can have at most  $kN_f \times h \times N_f$  pushes (and consequently, pops) per demand, or  $kN_f^2 \times h \times l$  in total where  $l$  is  $|\Sigma|$ , the number of lattice points. This is far superior in comparison to the  $O(N_f \cdot 2^{N_f})$  complexity of the liveness-based method with arbitrary precision in [3]. The  $2^{N_f}$  factor arises from a NFA to DFA conversion that is a part of the method. We substantiate this result with experiments in section 5.

### 3.2 Liveness Propagation

While the functions  $\mathcal{A}$  and  $\mathcal{L}$  describe the liveness of an application or an expression with respect to an arbitrary symbolic demand  $\sigma$ , for garbage collection we need to compute the concrete liveness environment at program points. In particular, we have to compute the liveness environment only at certain distinguished program points, called *gc-points*, where the garbage collector can be invoked. The liveness environments at these points are cached for use during garbage collection. We shall, for the moment, assume that a collection of program points called *gc-points* has already been identified, and define these later.

We now describe how to calculate the liveness environments at *gc-points*.

1. We begin by calculating  $\mathcal{L}(e_{\text{main}}, \top)$ , i.e. the liveness environment of  $e_{\text{main}}$ , the body of **main**, for the demand  $\top$ . The demand on  $e_{\text{main}}$  is  $\top$ , because we assume that the entire result returned by the program would be required. While doing this we propagate liveness to the entire program in the following manner:
  - a. Each time we compute  $\mathcal{L}((f\ g_1 \dots g_m\ x_1 \dots x_n), \sigma)$ , we also calculate the liveness environment of  $(e_f[G_1/g_1, \dots, G_m/g_m])$  for the same demand  $\sigma$ . The pair  $((g_1, \dots, g_m), \sigma)$  is called the *calling context* for the call  $((f\ g_1 \dots g_m\ x_1 \dots x_n)$ .
  - b. Consider a *gc-point*  $\pi$  in the body of  $f$ . While computing  $\mathcal{L}(e_f[G_1/g_1, \dots, G_m/g_m], \sigma)$ , we merge the liveness environments at  $\pi$  arising out of different calls only if their calling contexts are the same. Notice that liveness at a program point for different calls to  $f$  are not merged unconditionally as is done in most inter-procedural analyses. One of the key aspects of our method is that we track the calling context at runtime. As a consequence, our garbage collector is guided by a more precise analysis.
  - c. The liveness environment at *gc-points* is recorded in a table called **gctable** that is consulted during garbage collection. **gctable** $[\pi, f, (g_1, \dots, g_m), \sigma]$  contains the merged liveness environment at  $\pi$ , for all calls to  $f$  with the context  $((g_1, \dots, g_m), \sigma)$ .

**An Example.** Consider the same **append** function now expressed as an instance of a higher order function **foldr** (also called **reduce** in some languages). The function **append** is now defined by instantiating **foldr** to appropriate arguments. As before, **main** constructs a pair of lists  $p$  and  $q$ , and appends them. There are six garbage collection points  $\pi_1$ - $\pi_6$ . Of these, we would like to know the liveness environments at  $\pi_2$ ,  $\pi_3$  and  $\pi_4$ . We first compute  $\mathcal{FS}$  for each first order function (either declared or made first order by applying a higher order function to its function parameters). For the *gc-points* of interest, these are **append** and **foldr**. In fact, since **append** calls **foldr** with a swap of its non-function arguments, it

```

1 (define (foldr G id l)
2   (let c← (null? l) in
3     (if c (return id)
4         (let d← (cdr l) in
5           (let t← (foldr G id d) in
6              $\pi_3$ : (let a← (car l) in
7                  $\pi_4$ : (let ans← (G a t) in
8                     (return ans))))))))
9 (define (append x y)
10  (let res ← (foldr cons y x) in
11     $\pi_5$ : (return res)))
12 (define (main)
13  (let p← (makelist ...) in
14     $\pi_1$ : (let q← (makelist ...) in
15         $\pi_2$ : (let final ← (append p q) in
16             $\pi_6$ : (return final))))

```

**Figure 6.** Higher order version of **append**. It uses the higher order function **foldr**

is enough to analyse **foldr** and swap the elements of the resulting tuple.

A control flow analysis reveals that the only function that can be passed for  $G$  in **foldr** is **cons**. Thus one has to solve for  $\mathcal{FS}((\text{foldr cons}), \sigma)$  for all possible values of  $\sigma$ <sup>5</sup>. In the exposition below, if any entity  $E$  has a tuple as a value, we use the notation  $\langle E \rangle_i$  to denote the  $i^{\text{th}}$  element of the tuple. The  $\mathcal{FS}$  rule directly yields the following equation for  $\mathcal{FS}((\text{foldr cons}), \sigma)$ :

$$\mathcal{FS}((\text{foldr cons}), \sigma) = (\text{L}(\text{id}), \text{L}(1)),$$

where  $\text{L} = \mathcal{L}(e_{\text{foldr}}[G/\text{cons}], \sigma)$

This, in turn, expands into the following:

$$\mathcal{FS}((\text{foldr cons}), \sigma) = (\sigma_{\text{id}}, \sigma_1), \text{ where}$$

$$\sigma_{\text{id}} = \sigma \sqcup \langle \mathcal{FS}((\text{foldr cons}), \mathcal{S}(\bar{1}\sigma)) \rangle_1$$

$$\sigma_1 = \epsilon \sqcup \mathcal{S}(\epsilon \cup 1 \langle \mathcal{FS}((\text{foldr cons}), \mathcal{S}(\bar{1}\sigma)) \rangle_2) \sqcup \mathcal{S}(\epsilon \cup 0 \mathcal{S}(\bar{0}\sigma))$$

Let us explain the equation for  $\sigma_1$ . The demands on the three lines joined by  $\sqcup$  correspond to the livenesses of 1 occurring in lines 2, 4 and 6. The liveness  $\epsilon$  is directly from the **null?** rule. The liveness of  $\text{ans}$  is  $\sigma$  because of the **return** rule, and this is transformed to a liveness of  $\mathcal{S}(\bar{0}\sigma)$  for  $a$  and  $\mathcal{S}(\bar{1}\sigma)$  for  $t$  by the **let** and the **cons** rules. The liveness of  $a$  now becomes the demand on  $(\text{car } l)$  at line 6, and this gives the liveness of this occurrence of 1 as  $\mathcal{S}(\epsilon \cup 0 \mathcal{S}(\bar{0}\sigma))$ . The liveness of  $t$ , on the other hand, is propagated to the second (non-function) argument of **foldr**, namely  $d$ , to give a liveness of  $\langle \mathcal{FS}((\text{foldr cons}), \mathcal{S}(\bar{1}\sigma)) \rangle_2$  for  $d$ . This is further propagated through  $(\text{cdr } l)$  to give the liveness of 1 at 3 as  $\mathcal{S}(\epsilon \cup 1 \langle \mathcal{FS}((\text{foldr cons}), \mathcal{S}(\bar{1}\sigma)) \rangle_2)$ .

<sup>5</sup>We shall shortly see that solving for  $\mathcal{FS}((\text{foldr cons}), \top)$  is enough.

Let us calculate the value of  $\mathcal{FS}((\mathbf{foldr\ cons}), \sigma)$ , when  $\sigma = \top$ . We have

$$\begin{aligned} \mathcal{FS}^0((\mathbf{foldr\ cons}), \top) &= (\perp, \perp), \text{ and} \\ \mathcal{FS}^1((\mathbf{foldr\ cons}), \top) &= (\sigma_{\text{id}}^1, \sigma_1^1), \text{ where} \\ \sigma_{\text{id}}^1 &= \top \sqcup \langle \mathcal{FS}^0((\mathbf{foldr\ cons}), \mathcal{S}(\bar{1}\top)) \rangle_1 \\ &= \top \sqcup \langle \mathcal{FS}^0((\mathbf{foldr\ cons}), \top) \rangle_1 \\ &= \top \sqcup \perp = \top, \text{ and} \\ \sigma_1^1 &= \epsilon \sqcup \mathcal{S}(\epsilon \cup \mathbf{1} \langle \mathcal{FS}^0((\mathbf{foldr\ cons}), \mathcal{S}(\bar{1}\top)) \rangle_2) \\ &\quad \sqcup \mathcal{S}(\epsilon \cup \mathbf{0}\mathcal{S}(\bar{0}\top)) \\ &= \epsilon \sqcup \mathcal{S}(\epsilon \cup \mathbf{1} \langle \mathcal{FS}^0((\mathbf{foldr\ cons}), \top) \rangle_2) \\ &\quad \sqcup \mathcal{S}(\epsilon \cup \mathbf{0}\top) \\ &= \epsilon \sqcup \mathcal{S}(\epsilon \cup \mathbf{1}\perp) \sqcup \top_{0\epsilon} \\ &= \epsilon \sqcup \mathcal{S}(\epsilon \cup \perp) \sqcup \top_{0\epsilon} = \epsilon \sqcup \epsilon \sqcup \top_{0\epsilon} = \top_{0\epsilon} \end{aligned}$$

It is clear that the value of  $\sigma_{\text{id}}$  converges to  $\top$ . However, to find the convergence point of  $\sigma_1$ , we iterate once more.

$$\begin{aligned} \mathcal{FS}^2((\mathbf{foldr\ cons}), \top) &= (\top, \sigma_1^2), \text{ where} \\ \sigma_1^2 &= \epsilon \sqcup \mathcal{S}(\epsilon \cup \mathbf{1} \langle \mathcal{FS}^1((\mathbf{foldr\ cons}), \mathcal{S}(\bar{1}\top)) \rangle_2) \\ &\quad \sqcup \mathcal{S}(\epsilon \cup \mathbf{0}\mathcal{S}(\bar{0}\top)) \\ \sigma_2^1 &= \epsilon \sqcup \mathcal{S}(\epsilon \cup \mathbf{1} \langle \mathcal{FS}^1((\mathbf{foldr\ cons}), \top) \rangle_2) \sqcup \top_{0\epsilon} \\ \sigma_2^1 &= \epsilon \sqcup \mathcal{S}(\epsilon \cup \mathbf{1}\top_{0\epsilon}) \sqcup \top_{0\epsilon} = \epsilon \sqcup \top_{1\epsilon} \sqcup \top_{0\epsilon} = \top \end{aligned}$$

Since  $\sigma_1$  converges to  $\top$ ,  $\mathcal{FS}^2((\mathbf{foldr\ cons}), \top) = (\top, \top)$ . Armed with this information, we calculate the liveness environments at the three gc-points  $\pi_2$ ,  $\pi_3$  and  $\pi_4$  as follows:

1. The liveness at  $\pi_2$  is  $\{p \mapsto \mathcal{FS}^1((\mathbf{append}), \top), q \mapsto \mathcal{FS}^2((\mathbf{append}), \top)\}$ . This is the same as  $\{p \mapsto \mathcal{FS}^2((\mathbf{foldr\ cons}), \top), q \mapsto \mathcal{FS}^1((\mathbf{foldr\ cons}), \top)\}$  and further simplifies to  $\{p \mapsto \top, q \mapsto \top\}$ .
2. The liveness at  $\pi_3$  is  $\{a \mapsto \mathcal{S}(\bar{0}\top), t \mapsto \mathcal{S}(\bar{1}\top), l \mapsto \mathcal{S}(\epsilon \cup \mathcal{S}(\bar{0}\top))\}$ . This simplifies to  $\{a \mapsto \top, t \mapsto \top, l \mapsto \top_{1\epsilon}\}$ .
3. The liveness at  $\pi_4$  is  $\{a \mapsto \perp, t \mapsto \perp, l \mapsto \perp, \text{ans} \mapsto \top\}$ .

The liveness of  $l$  at  $\pi_3$  and  $\pi_4$  are the same as in our informal description in Section 1.

**GC-Points.** Before proceeding further we explain what gc-points are. As mentioned earlier, these are program points where a program can potentially resume after pausing for garbage collection.

1. First consider the template  $\pi_1$ : **(let**  $a \leftarrow (\mathbf{cons\ } x\ y)$  **in**  $e$ ). Clearly the garbage collection happens because of the failed allocation of a **cons** cell. After garbage collection is over, the execution resumes at  $\pi_1$  to complete the evaluation of the **let** expression. Thus a program point before a **let** expression that binds a variable to a **cons** cell is a gc-point.
2. Next consider the template **(let**  $a \leftarrow (f\ x_1 \dots x_n)$  **in**  $\pi_2$ :  $e$ ). A garbage collection could potentially take place during the call to  $f$ . After garbage collection, the execution will resume in the current scope at program point  $\pi_2$ . Thus a program point just after a function call is a gc-point.

A tricky case is the function call  $\pi_1$ : **(let**  $a \leftarrow (G\ x_1\ x_2)$  **in**  $\pi_2$ :  $e$ ), where  $G$  is a parameter that may be bound to **cons**. This is a situation where the idea of a separate liveness analysis

for each calling context comes to use. If the calling context binds  $G$  to **cons**, then the gc-point is identified as  $\pi_1$ , else it is  $\pi_2$ . Thus gc-point is not a purely syntactic notion but could also depend on the calling context.

## 4 The Garbage Collector

We can think of a liveness based traversal as a generalization of reachability. Given a node  $n$  in the heap and a demand  $\sigma$  associated with it, we define a node  $n'$  to be  $\sigma$ -reachable from  $n$ , if  $n'$  can be reached from  $n$  by following an access path in  $\sigma$ . Thus an RGC traverses all cells that are  $\top$ -reachable from the referents of the root set.

Our experimental setup includes an interpreter for the language shown in Figure 3 along with a single-generational copying collector. Assume that the program has been paused for garbage collection. Consider an activation frame in the stack and assume that it corresponds to the function call  $(f\ g_1 \dots g_m\ x_1 \dots x_n)$ . The activation frame will contain, along with other information, references corresponding to the root set variables  $x_1 \dots x_n$  and the function names  $g_1 \dots g_m$ . As additional information, we also maintain the name of the function  $f$  itself, the gc-point  $\pi$  inside the function body where the program is paused and also the demand  $\sigma$  on the function call during analysis. Thus the activation frame has all information for the garbage collector to consult  $\mathbf{gctable}[\pi, f, (g_1, \dots, g_m), \sigma]$  and read the liveness of the root set variables at  $\pi$ . This information is used by the garbage collector for traversal during copying.

LGCs [3, 21] with arbitrary demand patterns suffer from a drawback over RGCs during runtime. On visiting a memory cell  $n$  that has already been marked (i.e. visited before), an RGC does not initiate further traversals from  $n$ . This is safe because all cells reachable from  $n$  would have already been visited earlier. This property does not hold for LGC. An earlier visit to  $n$  with a demand, say  $\sigma$ , would have resulted in visits to only  $\sigma$ -reachable cells from  $n$ . Since there is no guarantee that the next visit to  $n$  will be accompanied by a demand that is a subset of  $\sigma$ , LGC methods with arbitrary patterns may need to initiate further traversals from an already visited cell.

With a set of finite demand patterns and an ordering on them, we are able to address this issue. The garbage collector creates a sequence of bins for the demand patterns  $\top$ ,  $1^*$ ,  $1_\epsilon$ ,  $0_\epsilon$  and  $\epsilon$ , and populates them by variables references in the root set. The order of the patterns in the sequence is important. The traversal for garbage collection now takes place by picking references from the bins in the sequence above. For example, while traversing a reference from the bin  $1^*$ , if we visit a marked cell  $n$ , we know that this cell had been visited earlier with either a  $1^*$  pattern itself or a  $\top$  pattern. Since the  $1^*$ -reachable cells from  $n$  would have already been visited earlier, further traversals from  $n$  are not necessary. The situation is slightly different however, if we



reach a marked cell with a  $0_\epsilon$  pattern. The marking could have been due to an earlier arrival with a  $1_\epsilon$  or  $1^*$  pattern, and neither of these two patterns subsume  $0_\epsilon$  (i.e.  $0_\epsilon \not\sqsubseteq 1_\epsilon$  and  $0_\epsilon \not\sqsubseteq 1^*$ ). Therefore, we have to traverse a single step from  $n$ , but this step may be redundant if  $n$  had been visited earlier with  $0_\epsilon$  or  $\top$ .

For obvious reasons, the demand  $\perp$  does not have a bin. Also notice the absence of bins for  $\top_{0_\epsilon}$  and  $\top_{1_\epsilon}$ . If a root set variable, say  $z$ , has the liveness  $\top_{0_\epsilon}$ , then, while initializing bins, the reference corresponding to  $z$  is inserted in the bin for  $0_\epsilon$  and the reference corresponding to  $(\text{car } z)$  is inserted in the bin for  $\top$ . Note that the traversal with very few revisits was made possible because of a total order amongst the bins corresponding to the infinite demands  $1^*$  and  $\top$ .

A second drawback of LGC addressed in our method is that we do a breadth-first traversal of the heap, thus avoiding the space-costs that come from using a stack for depth-first traversal. Earlier methods that had arbitrary sets of access paths as demands [3, 21] used an automaton for representing liveness, and collection involved a joint depth-first traversal of the heap and the automaton. Our collector uses the property that if the liveness  $\sigma$  of a root set variable is one of the infinite demands  $\top$  or  $1^*$ , then every  $\sigma$ -reachable cell from this root set variable also has the same demand. So we can copy these in breadth-first order without having to expend space to explicitly keep track of their liveness.

**The Garbage Collection Algorithm.** Algorithm 1 shows the garbage collection algorithm. As in any other single generation copying collector, the available memory is divided into two semispaces, their starting locations denoted by the variables `fromSpace` and `toSpace`. A round of garbage collection starts with a swap of roles of these variables, so that the occupied memory starts with `fromSpace`, from where live cells are copied into `toSpace`.

The garbage collector works as follows. The helper function `copy(r)` copies the content of  $r$  to `toSpace` if it is not already copied, and, in any case, returns a pointer to the copy. The collector function `gc` picks up root sets in order of bins. If the current bin corresponds to a demand  $\sigma$ , then all cells that are  $\sigma$ -reachable from  $r$  are copied into `toSpace`, if they are not already copied. The referent of each root in the bins is copied first, and the rest of the cells that are  $\sigma$ -reachable from  $r$  are either copied within the `if` ( $\sigma == 0_\epsilon$ ) or in the copy function (for other values of  $\sigma$ ). `startp` is a logical variable unrelated to garbage collection. Its only role is in describing an invariant that would be used in arguing the correctness of the collector.

**Correctness of the Algorithm.** We use a variation of the coloring abstraction that is often used to argue about the soundness of garbage collectors. We colour memory cells as follows: A cell is black if it and its  $\sigma$ -reachable children have been copied to `toSpace`. A cell is grey, if it has been copied into `toSpace`, but at least one of its  $\sigma$ -reachable children has

```

Function gc():
  swap(fromSpace, toSpace)
  allocp ← beginning of toSpace
  for each bin  $\sigma$  in  $\{\top, 1^*, 1_\epsilon, 0_\epsilon, \epsilon\}$  do
    for each root  $r$  in  $\sigma$  do
      scanp ← allocp; startp ← allocp
      r ← copy(r)
      if  $\sigma == 0_\epsilon$  then
        car(r) ← copy(car(r))
        scanp ← allocp
      while scanp < allocp do
        if scanp points to a cons cell then
          switch  $\sigma$  do
            case  $\top$  do
              car(scanp) ←
                copy(car(scanp))
              cdr(scanp) ←
                copy(cdr(scanp))
              scanp ←
                scanp + size(scanp)
            case  $1^*$  do
              cdr(scanp) ←
                copy(cdr(scanp))
              scanp ←
                scanp + size(scanp)
            case  $1_\epsilon$  do
              cdr(scanp) ←
                copy(cdr(scanp))
              scanp ← allocp
            case  $\epsilon$  do
              scanp ← allocp
          else
            /*scanp points to nil or integer*/
            scanp ← allocp

Function copy(r):
  if  $r$  has no forwarding address then
    r' ← allocp
    copy *r to *r'
    allocp ← allocp + size(r)
    set forwarding address of  $r$  as  $r'$ 
  return (forwarding address of  $r$ )

```

**Algorithm 1:** The garbage collection algorithm

not been copied. Finally, a cell is white if the cell itself has not been copied to `toSpace`. The correctness of the algorithm follows from the following invariant at the beginning of the while loop:

1.  $\text{startp} \leq \text{scanp} \leq \text{allocp}$ , and
2. all cells with starting addresses within  $[\text{startp}, \text{scanp})$ , are black and are  $\sigma$ -reachable from  $r$ .

Consider a bin  $\sigma$  and a root variable  $r$  from the  $\sigma$  selected by the headers of the two `for` loops. While we do

not prove the invariant stated above, we show as its consequence that when the while loop terminates, the algorithm would have correctly copied all cells that are  $\sigma$ -reachable from  $r$ . First notice that just after the termination of the while loop,  $\text{scanp} = \text{allocp}$ , and thus there are no grey cells. Now assume to the contrary that there is a cell  $n$  that is  $\sigma$ -reachable from the root and is not colored black. Since the root has been copied before the while loop, we can find a cell in  $[\text{startp}, \text{scanp})$  with at least one white  $\sigma$ -reachable child. Thus the color of this cell is grey, which is contradictory to the loop invariant which says that all cells such that within  $[\text{startp}, \text{scanp})$  are black.

## 5 Experimental Results

To demonstrate the effectiveness of our method, we have built an interpreter for the target language and integrated it separately with the following three collectors: Our collector based on the finite liveness domain  $\Sigma$  (referred to as  $\text{LGC}_F$ ), an implementation of the collector based on arbitrary liveness values [3] (referred to as  $\text{LGC}_A$ ) and a reachability based garbage collector (referred to as RGC). The hardware platform for our experiments was a third generation *i7* processor with 8GB RAM. Any difference in the reported timings in [3] and what we report for the same benchmark is attributable to two reasons: the input sizes of the benchmark and the hardware platform (both not reported in [3]). The collectors were benchmarked with programs from various sources including the standard no-fib suite. In addition,  $\text{LGC}_F$  was tested with four medium-sized higher-order programs: Barnes-Hut simulation, a list based floating point arithmetic program (IEEE 754), a functional parser, and curves—a program to generate fractals.

**Results.** The graphs in figure 7 indicate memory usage patterns of some of the benchmarks. The x-axis measures progress in execution in terms of the cumulative number of cells allocated, which we call *GC Ticks*. The y-axis is the number of cells in *toSpace*. The black line denotes the actual number of live cells at every tick. This metric is computed by recording the GC-ticks at creation and at last use of each allocated cell at runtime. Note that the graphs represent memory usage of programs when run with the heap size set to the minimum required for execution under RGC. The saw-tooth nature of the curves is explained by the number of cells in *toSpace* growing one-on-one with each GC-Tick and then instantaneously falling during the release of cells during garbage collection. A line connecting the lowest points of the curve for a GC scheme represents the *analysed liveness* of the analysis accompanying the scheme

**Precision of Liveness and Minimum Heap Requirement.** Figure 7 shows the precision of liveness analysis. This shows up as the extent to which the analysed liveness comes

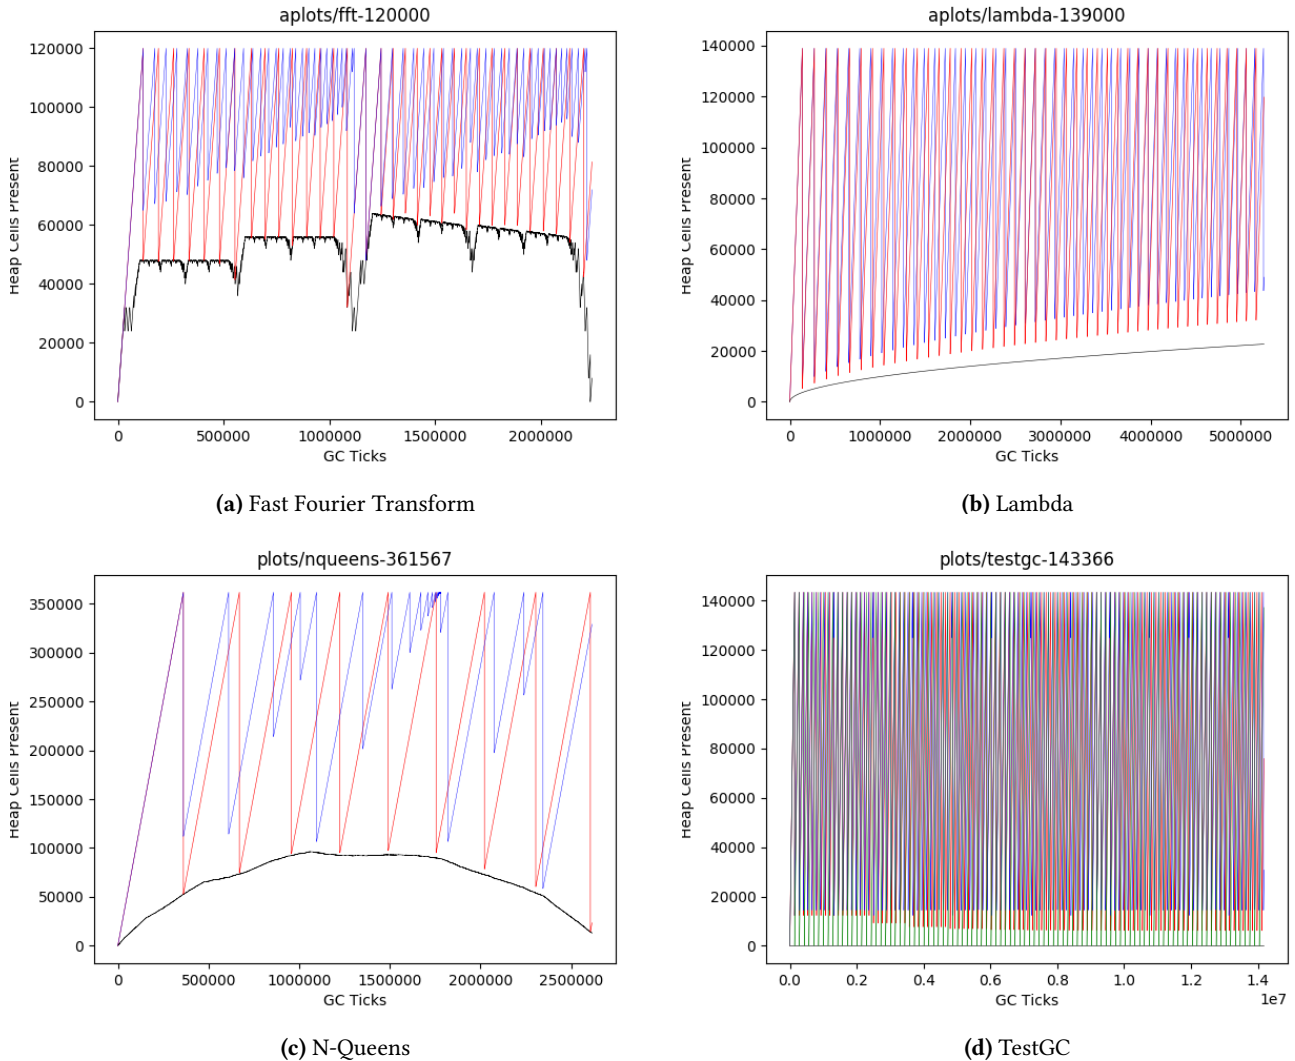
**Table 1.** Minimum Heap Cells (in bytes) required for execution under various GC schemes.

	RGC	$\text{LGC}_F$	$\text{LGC}_A$
lambda	138340	102500	102500
fft	113500	64502	64502
testgc*	143366	12291	41
treejoin	525488	7150	7150
fibheap	254520	13558	13558
nperm	930264	8920	8920
huffman	150041	50029	50023
gc_bench*	131071	102401	131071
nqueens	361567	98531	98531
knightstour	508225	307091	307090
lcss	5M+	< 50000	< 50000
paraffins	892090	891194	891194
sudoku	160277	85363	85347
deriv	491462	491446	491446
curves	204000	96000	-
parser	262216	131096	-
barneshut	6830	1530	-
IEEE-754	39360	122	-

close to real liveness during garbage collection. The effectiveness of  $\text{LGC}_F$  over RGC in identifying live memory is evident— $\text{LGC}_F$  comes much closer to the real liveness curve than RGC. As expected, the gaps in the curves between analysed and real liveness for lambda and testgc also show that liveness analysis in  $\text{LGC}_F$  is an approximation. The memory usage plots also show that saw-toothed pattern for  $\text{LGC}_F$  is more widely spaced because of fewer GC invocations.

We avoid plotting memory usage for  $\text{LGC}_A$ , since it is indistinguishable from  $\text{LGC}_F$  for all benchmarks except testgc. Thus our restricted set of liveness values captures liveness almost as precisely as  $\text{LGC}_A$ . The program testgc (Figure 7 (d)), is a notable exception. It is a crafted program with an irregular memory access pattern that cannot be captured by our set of liveness values.

Since  $\text{LGC}_F$ , in general, releases as much memory (and often more) than RGC during each collection, it is natural to expect programs executing under  $\text{LGC}_F$  to run with smaller heap sizes. This is confirmed by Table 1. The minimum heap requirements of  $\text{LGC}_F$  and  $\text{LGC}_A$  are comparable on most programs. For the program testgc  $\text{LGC}_F$  falls short because of imprecise liveness analysis. Surprisingly,  $\text{LGC}_F$  performs better than  $\text{LGC}_A$  for gcbench, a synthetic program from the no-fib benchmark suite which creates a tree but does not use any of its parts. Given the comparable precision of liveness analysis for both  $\text{LGC}_F$  and  $\text{LGC}_A$  for this program, the superior minheap performance of  $\text{LGC}_F$  can only be attributed to the additional precision from the context-based separation of analysis of function bodies in our method.

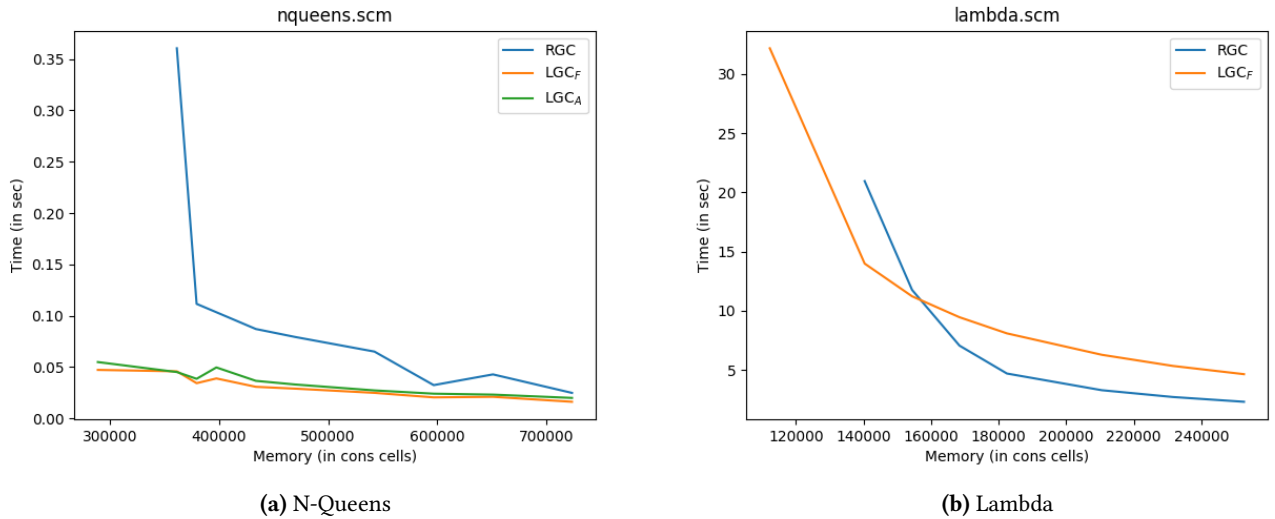


**Figure 7.** Memory plots of programs: x-axis is the cumulative number of `cons` cells allocated and y-axis is the number of `cons` cells in `toSpace`. The black curve represents real liveness, the blue and the orange curves are the number of cells in `toSpace` for RGC and LGC<sub>F</sub> respectively. The LGC<sub>A</sub> curve is indistinguishable from LGC<sub>F</sub> for programs (a) to (c). The exception amongst all benchmarks is `testgc` shown in (d). The green curve representing LGC<sub>A</sub> in this plot reaches real liveness at each GC, whereas LGC<sub>F</sub> does not.

**Analysis Time.** A comparison of the time taken for liveness analysis for the two liveness-based collectors is in Table 2. LGC<sub>F</sub> has much lower analysis time with some examples showing a difference of several orders of magnitude. This is expected since, as explained at the end of Section 3.1, the work-list algorithm of LGC<sub>F</sub> with its short lattice height is computationally much less expensive than the analysis of LGC<sub>A</sub>.

Going beyond the table that shows the analysis time of our benchmark programs, we believe that the analysis will scale for larger programs for the following reasons. The  $N_f^2$  term in the theoretical worst case bound of the expensive  $\mathcal{FS}$  computation arises out of the assumption that each function

can potentially call every other function in the program. However, if we make the more reasonable assumption that the number of functions called by a function is bounded by a constant, the bound is linear in  $N_f$ . We have experimentally observed that, during the computation of  $\mathcal{FS}$ , the number of iterations per function over all programs in our benchmarks is less than 10. Similarly, the worst time per iteration over our benchmark programs is 0.25 ms. Considering these figures to be representative over arbitrary programs, a realistic yet conservative estimate of the analysis time for a large program with 1000 functions is  $0.25 * 10 * 1000 = 2.5\text{secs}$ .



**Figure 8.** Time taken by garbage collector vs the amount of heap memory available. In Figure (b) the LGC<sub>A</sub> plot is avoided due to difference in scale.

**Table 2.** Time taken for compile-time analysis.

	LGC <sub>F</sub> (in ms)	LGC <sub>A</sub> (in ms)
lambda	6.38	6067.00
fft	17.24	546.17
testgc	4.32	29.19
treejoin	5.71	3278.70
fibheap	30.53	2049.89
nperm	48.58	174.67
huffman	11.18	1459.72
gc_bench	18.75	7.61
nqueens	3.86	296.11
knightstour	4.90	432.86
lcss	14.59	980.44
paraffins	33.56	3784.40
sudoku	23.52	187162.00
deriv	6.80	327.75
curves	37.40	-
parser	11.44	-
barneshut	152.76	-
IEEE-754	98.52	-

**Time for Garbage Collection.** The graphs in Figure 8 plot the average garbage collection time (gc-time) for different collectors over a range of heap sizes. The time measurements are accurate modulo the vagaries of capturing CPU Time. We reiterate the factors that affect gc-time overheads in LGC<sub>F</sub>. First, as mentioned in Section 4, LGC<sub>F</sub> does a binning operation that incurs the costs of a bucket sort during garbage collection time. On the other hand, given the precision of our liveness analysis arrived by several means, the number of copied cells is generally lower than RGC. In addition, the number of collector invocations is also lower

for LGC<sub>F</sub>. All of these add up to significant gc-time gains for LGC<sub>F</sub> over RGC in memory constrained regions.

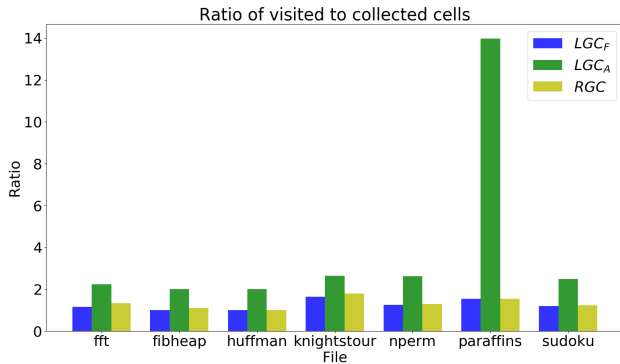
In contrast, the regions where memory is abundantly available, the difference between number of collections is not significant (in many cases, non-existent) and we find the benefits diminishing. For systems where memory is readily available, performing such an analysis is not computationally beneficial and RGC performs better, as seen in Figure 8 lambda.

An issue related to gc-time is the number of redundant traversals of the heap during collection. This value will be high if there are repeated heap traversal during a garbage collection. LGC<sub>F</sub> has an advantage over LGC<sub>A</sub> in that the average number of visits per copied cell is much lower (Figure 9). For LGC<sub>F</sub>, this ratio is close to 1, and almost always matches RGC. Some programs are not represented in this plot because their LGC<sub>A</sub> ratios go beyond scale.

Since our method carries over large amounts of liveness data from analysis to execution, a natural concern would be the size of this data. The amount of space needed to store the analysis data is  $N_p * (N_l * (l + 1))$  bytes, where  $N_p$  is the number of the garbage collection points,  $N_l$  is the average number of live variables at each garbage collection point and  $l$  is the number of points in the lattice. Empirically, the size of this data is in the range of a few (around 5) KBs for all programs, with the only outlier being sudoku, which requires around 20KB.

## 6 Related Work

There is a substantial body of work related to the the use of static and dynamic analyses for improving usage of heap memory. The empirical study in [9], done in the context of



**Figure 9.** Ratio of total heap cell visited to total heap cells collected, both figures cumulative over all GC invocations. A higher ratio points towards repeated traversal of heap during collection.

C, C++ and Eiffel, was the first to recommend the use of liveness to enhance the effectiveness of garbage collection. The garbage collector was based on a liveness analysis obtained from program traces. The authors discovered that heaps often contained a large number of cells that are reachable but not live (upto 32%). Subsequent studies by [3, 15, 22, 24] have confirmed this feature across several languages: Java, Haskell and Scheme. The conclusion in [9] is that for liveness-based garbage collection to be effective, one must analyze global variables with interprocedural analysis. This message is relevant for us since our analysis is interprocedural, and instead of globals we have heap-allocated data that, like globals, is shared across different functions.

We first discuss the use of liveness in garbage collection for imperative languages. The heap safety analysis in [25] suggests a method for answering queries regarding the safety of inserting a `free` statement for a variable or nullifying a reference at a given location, so that the referent object can be collected subsequently. The approach in [16] is similar in its end goals. However, while heap safety analysis merely solves a decision problem—it answers queries regarding whether a reference can be nullified at a program point, the approach in [16] also identifies candidate references which can be nullified at each program point. Both these studies are in the context of Java programs. Similar approaches that either verify the safety of explicitly freed memory or insert memory freeing statements are [6, 8].

In the space of functional languages, many of the approaches are aimed at reallocation through analysis followed by code insertion, also called *compile time garbage collection*. The most well known of these is, of course, deforestation [4, 5, 27]. Deforestation methods essentially identify commonly occurring patterns of code that allocate a chunk of memory and then consume it. Such patterns are fused so that the allocation-consumption takes places at a finer level of granularity, and a small constant amount of memory is repeatedly allocated and consumed immediately. Other

approaches to compile time garbage collection are [12], [18] and [7]. In addition, there are sharing analysis based reallocation [14] and region based analysis [26]. Hoffman [10] uses linear types to compile functional programs to C programs that perform in-place updates. This is similar to our motivating example in its end goal. An approach more relevant to garbage collection is [11], which uses linear type system to infer an abstract form of reference counting. However, the method is built on limited datatypes that does not even include lists in their full generality. In addition, the method is not supported by an implementation.

The approaches that are closest to ours are [3] and [17]. The first paper suggests a liveness based garbage collection for a first-order eager language. The second paper reworks the method for a lazy language that also manages to garbage collects parts of closures. However to represent liveness values, they use an unrestricted set of demands (arbitrary subsets of  $(0 + 1)^*$ ). A demand is represented as an automaton, and garbage collection is a simultaneous depth first traversal of the heap and the automaton. Construction of the automaton is costly. It involves approximating a context free grammar to a nondeterministic automaton, which is then converted to a deterministic finite automaton—both costly processes. Since the set of demands is infinite with an absence of a total order, they cannot use a binning technique and avoid revisits as we do in our method. The idea of using a finite domain of demands to simplify the analysis comes from strictness analysis [13, 28]. Strictness analysis is similar to liveness, and is used for space optimization of lazy languages. Recollect that a lazy language does not evaluate an expression until it is required. Strictness analysis determines the parts of an argument of a function that are guaranteed to be used inside it. The identified parts can then be evaluated before being passed to the function instead of being passed as space consuming closures. Just as in our case, limiting the strictness values to a small domain enables the analysis to address higher order functions. The idea of separating the liveness environment by the calling context in a function body also appears in [20], albeit with a different notion of context.

The study in [1] observes that liveness may aid generational garbage collection. Because of the early collection of objects, fewer objects would be moved into the tenured generation and thus result in fewer major collections. We believe that our implementation can be easily modified to work in generational setting; the only additional issue in our case being that inter-generational pointers have to be recorded along with their liveness values.

## 7 Conclusions

The disciplined use of recursive functions on algebraic data types is a common paradigm that is followed while programming in a functional language. This has an interesting

consequence on heap usage. The traversals on the heap are regular and follow identifiable patterns. In this paper we have identified a few such patterns that are used to approximate liveness values of structured data. The liveness values form a lattice and the liveness analysis based on this lattice is much more efficient and as precise as previous methods that use an unrestricted set of liveness values. The use of a restricted set of liveness values also leads to an improvement in the design of the collector which does a breadth-first traversal of the heap and also minimizes redundant traversals. An implementation of a garbage collector based on the restricted liveness set ( $LGC_F$ ) confirms several benefits over collectors that use an arbitrary set of liveness values ( $LGC_A$ ), and also over reachability based collectors (RGC): The analysis time of  $LGC_F$  is a significant improvement over  $LGC_A$ , the minimum memory required for programs to run is comparable with  $LGC_A$ , the number of redundant traversals over memory cells reduce to the level of RGC, and garbage collection times are better than both  $LGC_A$  and RGC.

Since the choice of liveness values in our case is not backed by sufficient evidence of its applicability to a wider range of programs, an avenue for improvement is a data-driven study in the spirit of [9] to identify a better choice of a lattice of liveness values. The choice could even be specialized to a set of target programs. While it is not very difficult to change the lattice of liveness values in our current implementation, a more careful implementation could fully parameterize the collector implementation with respect to the choice of lattice.

One could also extend the scope of the garbage collector to support a more realistic functional language, and even to a lazy language like Haskell by extending the techniques in [17]. However, whether our technique could be extended to programs in mainstream imperative languages like Java with their state changing features and unrestricted usage of heap, requires investigation.

It has been already been demonstrated in [24] that Java programs too carry considerable data that are dead but reachable. Given this, one could ask whether the techniques used here could carry over to programs written in Java. A possible approach is based on [2]. This paper shows a close relation between SSA, an intermediate form of programs used by most compilers and functional programs of the kind that we have considered in the paper. It does seem possible to take the SSA form of programs written in Java by tapping into the workflow of a compiler and following the technique suggested in [2] convert these into the variant of Scheme considered in this paper. However, given the different styles of programming, whether it would be possible to identify a small set of liveness patterns that can be used to yield an analysis of reasonable precision is an open question.

## References

- [1] Ole Agesen, David Detlefs, and J. Eliot Moss. 1998. Garbage Collection and Local Variable Type-Precision and Liveness in Java Virtual Machines. In *Proceedings of the ACM SIGPLAN 1998 Conference on Programming Language Design and Implementation (Montreal, Quebec, Canada) (PLDI '98)*. ACM, New York, NY, USA, 269–279. <https://doi.org/10.1145/277650.277738>

- [2] Andrew W. Appel. 1998. SSA is Functional Programming. *ACM SIGPLAN NOTICES* 33, 4 (1998), 17–20.
- [3] Rahul Asati, Amitabha Sanyal, Amey Karkare, and Alan Mycroft. 2014. Liveness-Based Garbage Collection. In *Compiler Construction*, Albert Cohen (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 85–106.
- [4] Olaf Chitil. 1999. Typer Inference Builds a Short Cut to Deforestation. In *Fourth ACM SIGPLAN International Conference on Functional Programming (Paris, France) (ICFP '99)*. ACM, New York, NY, USA, 249–260. <https://doi.org/10.1145/317636.317907>
- [5] Andrew Gill, John Launchbury, and Simon L. Peyton Jones. 1993. A Short Cut to Deforestation. In *Conference on Functional Programming Languages and Computer Architecture (Copenhagen, Denmark) (FPCA '93)*. ACM, New York, NY, USA, 223–232. <https://doi.org/10.1145/165180.165214>
- [6] Samuel Z. Guyer, Kathryn S. McKinley, and Daniel Frampton. 2006. Free-Me: A Static Analysis for Automatic Individual Object Reclamation. In *27th ACM SIGPLAN Conference on Programming Language Design and Implementation (Ottawa, Ontario, Canada) (PLDI '06)*. ACM, New York, NY, USA, 364–375. <https://doi.org/10.1145/1133981.1134024>
- [7] G. W. Hamilton. 1995. Compile-Time Garbage Collection for Lazy Functional Languages. In *International Workshop on Memory Management (IWMM '95)*. Springer-Verlag, Berlin, Heidelberg, 119–144.
- [8] Matthew Hertz and Emery D. Berger. 2005. Quantifying the Performance of Garbage Collection vs. Explicit Memory Management. *SIGPLAN Not.* 40, 10 (Oct. 2005), 313–326.
- [9] Martin Hirzel, Amer Diwan, and Johannes Henkel. 2002. On the Usefulness of Type and Liveness Accuracy for Garbage Collection and Leak Detection. *ACM Trans. Program. Lang. Syst.* 24, 6 (Nov. 2002), 593–624. <https://doi.org/10.1145/586088.586089>
- [10] Martin Hofmann. 2000. A Type System for Bounded Space and Functional In-Place Update—Extended Abstract. In *Programming Languages and Systems, 9th European Symposium on Programming, ESOP 2000 (Lecture Notes in Computer Science)*, Gert Smolka (Ed.), Vol. 1782. Springer, New York, NY, USA, 165–179. [https://doi.org/10.1007/3-540-46425-5\\_11](https://doi.org/10.1007/3-540-46425-5_11)
- [11] Atsushi Igarashi and Naoki Kobayashi. 2000. Garbage Collection Based on a Linear Type System.
- [12] Katsuro Inoue, Hiroyuki Seki, and Hikaru Yagi. 1988. Analysis of Functional Programs to Detect Run-Time Garbage Cells. *ACM Trans. Program. Lang. Syst.* 10, 4 (Oct. 1988), 555–578. <https://doi.org/10.1145/48022.48025>
- [13] Kristian Damm Jensen, Peter Hjøresen, and Mads Rosendahl. 1994. Efficient strictness analysis of Haskell. In *Static Analysis*, Baudouin Le Charlier (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 346–362.
- [14] Simon B. Jones and Daniel Le Métayer. 1989. Compile-Time Garbage Collection by Sharing Analysis. In *Fourth International Conference on Functional Programming Languages and Computer Architecture (Imperial College, London, United Kingdom) (FPCA '89)*. ACM, New York, NY, USA, 54–74. <https://doi.org/10.1145/99370.99375>
- [15] Amey Karkare, Amitabha Sanyal, and Uday Khedker. 2006. Effectiveness of Garbage Collection in MIT/GNU Scheme. <http://arxiv.org/abs/cs/0611093>.
- [16] Uday P. Khedker, Amitabha Sanyal, and Amey Karkare. 2007. Heap Reference Analysis Using Access Graphs. *ACM Trans. Program. Lang. Syst.* 30, 1 (Nov. 2007), 1–es. <https://doi.org/10.1145/1290520.1290521>
- [17] Prasanna Kumar K., Amitabha Sanyal, and Amey Karkare. 2016. Liveness-Based Garbage Collection for Lazy Languages. In *2016 ACM SIGPLAN International Symposium on Memory Management (Santa Barbara, CA, USA) (ISMM 2016)*. ACM, New York, NY, USA, 122–133. <https://doi.org/10.1145/2926697.2926698>

- [18] Oukse Lee, Hongseok Yang, and Kwangkeun Yi. 2003. Inserting Safe Memory Reuse Commands into ML-like Programs. In *10th International Conference on Static Analysis* (San Diego, CA, USA) (SAS'03). Springer-Verlag, Berlin, Heidelberg, 171–188.
- [19] Neil Mitchell and Colin Runciman. 2009. Losing Functions without Gaining Data: Another Look at Defunctionalisation. In *2nd ACM SIGPLAN Symposium on Haskell* (Edinburgh, Scotland) (Haskell '09). ACM, New York, NY, USA, 13–24. <https://doi.org/10.1145/1596638.1596641>
- [20] Rohan Padhye and Uday P. Khedker. 2013. Interprocedural Data Flow Analysis in Soot Using Value Contexts. In *2nd ACM SIGPLAN International Workshop on State Of the Art in Java Program Analysis* (Seattle, Washington) (SOAP 13). ACM, New York, NY, USA, 31–36.
- [21] K Prasanna Kumar, Amitabha Sanyal, Amey Karkare, and Saswat Padhi. 2019. A Static Slicing Method for Functional Programs and Its Incremental Version. In *28th International Conference on Compiler Construction* (Washington, DC, USA) (CC 2019). ACM, New York, NY, USA, 53–64.
- [22] Niklas Røjemo and Colin Runciman. 1996. Lag, Drag, Void and Use—Heap Profiling and Space-Efficient Compilation Revisited. In *First ACM SIGPLAN International Conference on Functional Programming* (Philadelphia, Pennsylvania, USA) (ICFP '96). ACM, New York, NY, USA, 34–41. <https://doi.org/10.1145/232627.232633>
- [23] Amr Sabry and Matthias Felleisen. 1992. Reasoning about Programs in Continuation-Passing Style.. In *1992 ACM Conference on LISP and Functional Programming* (San Francisco, California, USA) (LFP '92). ACM, New York, NY, USA, 288–298. <https://doi.org/10.1145/141471.141563>
- [24] Ran Shaham, Elliot K. Kolodner, and Mooly Sagiv. 2000. On Effectiveness of GC in Java. In *2nd International Symposium on Memory Management* (Minneapolis, Minnesota, USA) (ISMM '00). ACM, New York, NY, USA, 12–17. <https://doi.org/10.1145/362422.362430>
- [25] Ran Shaham, Eran Yahav, Elliot K. Kolodner, and Mooly Sagiv. 2003. Establishing Local Temporal Heap Safety Properties with Applications to Compile-Time Memory Management. In *10th International Conference on Static Analysis* (San Diego, CA, USA) (SAS'03). Springer-Verlag, Berlin, Heidelberg, 483–503.
- [26] Mads Tofte and Lars Birkedal. 1998. A Region Inference Algorithm. *ACM Trans. Program. Lang. Syst.* 20, 4 (July 1998), 724–767. <https://doi.org/10.1145/291891.291894>
- [27] Philip Wadler. 1988. Deforestation: Transforming programs to eliminate trees. In *ESOP '88*, H. Ganzinger (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 344–358.
- [28] Philip Wadler and R. J. M. Hughes. 1987. Projections for strictness analysis. In *Functional Programming Languages and Computer Architecture*, Gilles Kahn (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 385–407.