

Chapter 7

A Sharing Analysis for SAFE

Ricardo Peña¹, Clara Segura¹, Manuel Montenegro¹

Abstract: We present a sharing analysis for the functional language *Safe*. This is a first-order eager language with facilities for programmer-controlled destruction and copying of data structures. It provides also *regions*, i.e. disjoint parts of the heap where the programmer may allocate data structures. The analysis gives upper approximations to the sets of variables respectively sharing a recursive substructure, or any substructure, of a given variable. Its results will be used to guarantee that destruction facilities and region management are done in a safe way. In order to have a modular and efficient analysis, we provide signatures for functions, which summarize their sharing behaviour. The paper ends up describing the implementation of the analysis and some examples.

7.1 INTRODUCTION

Many imperative languages offer low level mechanisms to allocate and free heap memory, which the programmer may use in order to dynamically create and destroy pointer based data structures. These mechanisms give the programmer complete control over memory usage but are very error prone. Well known problems that may arise when using a programmer-controlled memory management are dangling references, undesired sharing between data structures with complex side effects as a consequence, and polluting memory with garbage.

Functional languages usually consider memory management as a low level issue. Allocation is done implicitly and usually a garbage collector takes care of the memory exhaustion situation.

In a previous paper [11] we proposed a semi-explicit approach to memory control by defining a functional language, called *Safe*, in which the programmer cooperates with the memory management system by providing some information

¹Dpto. Sistemas Informáticos y Computación, Univ. Complutense de Madrid, Spain
ricardo@sip.ucm.es, csegura@sip.ucm.es, manuelmont@gmail.com
Work partially supported by the Spanish project TIN2004-07943-C04.

about the intended use of data structures. For instance, the programmer may indicate that some particular data structure will not be needed in the future and that, as a consequence, it may be safely destroyed by the runtime system and its memory recovered. The language uses regions to locate data structures. It also allows controlling the degree of sharing between different data structures. A garbage collector is not needed. Allocation and destruction of data structures are done as execution proceeds.

More interesting is the definition of a type system guaranteeing that destruction facilities and region management can be done in a safe way. This type system will be the main topic of an ongoing paper (a draft version can be found at [11]). In particular, it guarantees that dangling pointers are never created in the live heap. An ill-constructed program is rejected by the type system. It makes a heavy use of sharing information, given by two functions called *shareall* and *sharerec*. Given a subexpression e of a function body and a free variable x such that e is included in the lexical scope of x ,

- *shareall*(x, e) returns the set of all the variables in scope in e which, at runtime, may share any substructure of the structure pointed to by x .
- *sharerec*(x, e) returns the set of all the variables in scope in e which, at runtime, may share any recursive substructure of the structure pointed to by x .

In this paper an upper approximation to these two functions is computed at compile time by an abstract interpretation-like analysis. Additionally we formally define the operational semantics of the language.

The structure of the paper is as follows: In Section 7.2, we provide a summary of the syntax and the operational semantics of *Safe*. Then, Section 7.3 presents in detail the sharing analysis. In Sections 7.4 and 7.5, the implementation of the analysis is described and it is applied to some illustrative examples. Section 7.6 surveys some related work and concludes.

7.2 SUMMARY OF *SAFE*

7.2.1 Syntax

We start by reproducing some crucial definitions which underlie the language. In Section 7.6 we compare our language design with other approaches using regions and memory management facilities.

Definition 7.1. A **region** is a contiguous memory area in the heap where data structures can be constructed, read, or destroyed. It is allocated and freed as a whole, in constant time.

Definition 7.2. A **cell** is a small memory space, big enough to hold a data constructor. In implementation terms, a cell contains the mark (or code pointer) of the constructor, and a representation of the free variables to which the constructor is applied. These may consist, either of basic values, or of pointers to non-basic values.

Definition 7.3. A **data structure**, in the following a *DS*, is the set of cells obtained by starting at one cell considered as the root, and taking the transitive closure of the relation $C_1 \rightarrow C_2$, where C_1 and C_2 are cells of the same type T , and in C_1 there is a pointer to C_2 .

That means that, for instance in a list of type $[[a]]$, we are considering as a DS the cons-nil spine of the *outermost* list, but not those belonging to the individual innermost lists. Each one of the latter constitute a separate DS.

The following decisions were taken:

1. A DS completely resides in one region.
2. One DS can be part of another DS, or two DSs can share a third one.
3. The basic values —integers, booleans, etc.— do not allocate cells in regions. They live inside the cells of DSs, or in the stack.
4. A function of n parameters can access:
 - Its n parameters, each one residing in a possibly different region.
 - Its **output region**, whenever it builds a DS as a result. There is at most one output region per function. Delivering this region identifier as a parameter is the responsibility of the call. We force functions to leave their result in an output region belonging to the calling context in order to safely delete the intermediate results computed by the function.
 - Its (optional) **working region**, referred to through the reserved identifier *self*, where it may create intermediate DSs. The working region has the same lifetime as the function call: It is allocated at each invocation and freed at function termination.
5. If a parameter of a function is a DS, it can be destroyed by the function. We will say that the parameter is **condemned** because this capability depends on the function definition, not on its use.
6. The capabilities a function has on its accessible DSs and regions are: a function may only read a DS which is a read-only parameter; a function may read (before destroying it), and must destroy, a DS which is a condemned parameter; a function may construct, read, or destroy DSs, in either its output or its working region.

The syntax of *Safe* is shown in Figure 7.1. This is a first-order eager functional language where sharing is expressed using variables in function and constructor applications. We intend *Safe* to be a core language resulting from the desugaring of a higher level language similar to Haskell or ML. The analysis defined in this paper is however done at core level. This is a usual approach in many compilers.

A program *prog* in *Safe* is a sequence of possibly recursive polymorphic function definitions³ followed by a main expression *expr*, calling them, whose value

³The extension to mutual recursion would pose no special problems, but we restrict ourselves to non-mutual recursion in order to ease the presentation.

<i>prog</i>	→	$dec_1; \dots; dec_n; expr$	
<i>dec</i>	→	$f \overline{x_i^n} r = expr$	{recursive, polymorphic function}
		$f \overline{x_i^n} = expr$	
<i>expr</i>	→	a	{atom: literal c or variable x }
		$x@r$	{copy}
		$x!$	{reuse}
		$(f \overline{a_i^n})@r$	{function application}
		$(f \overline{a_i^n})$	{function application}
		$(C \overline{a_i^n})@r$	{constructor application}
		let $x_1 = expr_1$ in $expr$	{non-recursive, monomorphic}
		case x of $\overline{alt_i^n}$	{read-only case}
		case! x of $\overline{alt_i^n}$	{destructive case}
<i>alt</i>	→	$C \overline{x_i^n} \rightarrow expr$	

FIGURE 7.1. First-order functional language *Safe*

is the program result. Function definitions building and returning a new DS will have an additional parameter r , which is the output region, where the resulting DS is to be constructed. In the right hand side expression only r and its own working region *self* may be used. Polymorphic algebraic data types definitions are also allowed. We will assume they are defined separately through **data** declarations.

The program expressions include variables, literals, function and constructor applications, and also **let** and **case** expressions, but there are some additional expressions:

If x is a DS, the expression $x@r$ represents a copy in region r of the DS accessed from x . The DS x must live in a region $r' \neq r$. Both x and $x@r$ have the same recursive structure and they share their non-recursive substructures.

The expression $x!$ means the reusing of the destroyable DS to which x points. This is useful when we do not want to destroy completely a condemned parameter but instead to reuse part of it. In semantic terms, x and $x!$ point to the same physical structure but, in language terms, once $x!$ is used, the name x becomes unaccessible in the subsequent text.

In function application we have a special syntax $@r$ to express the inclusion of the additional output region parameter. Using the same syntax, we express that a constructor application is to be allocated in region r .

The **case!** expression indicates that the outer constructor of x is disposed after the pattern matching so that x is not accessible anymore. The recursive substructures may be explicitly destroyed in the subsequent code via another **case!** or reused via $x!$. A condemned variable may be read but, once its content has been destroyed or reused in another structure, it may not be accessed again. This is what the type system guarantees. It annotates the type of such variable with a $!$.

We show now with several examples how to use the language facilities. In some of them we will write $x!$ or $(C \overline{a_i^n})@r$ as actual parameters of applications in order to abbreviate, when a **let** binding would in fact be needed. In these examples we show also the types that the functions have in the type system previously

$$\begin{aligned}
revD &:: \forall a, \rho_1, \rho_2. [a]!@ \rho_1 \rightarrow \rho_2 \rightarrow [a]@ \rho_2 \\
revD \ x \ r &= (revauxD \ x \ []@r)@r \\
\\
revauxD &:: \forall a, \rho_1, \rho_2. [a]!@ \rho_1 \rightarrow [a]@ \rho_2 \rightarrow \rho_2 \rightarrow [a]@ \rho_2 \\
revauxD \ x \ y \ r &= \mathbf{case!} \ x \ \mathbf{of} \\
&\quad [] \rightarrow y \\
&\quad x : xx \rightarrow (revauxD \ xx \ (x : y)@r)@r
\end{aligned}$$
FIGURE 7.2. Destructive list inversion

$$\begin{aligned}
insertD &:: \forall a, \rho. a \rightarrow Tree \ a!@ \rho \rightarrow \rho \rightarrow Tree \ a@ \rho \\
insertD \ x \ t \ r &= \mathbf{case!} \ t \ \mathbf{of} \\
&\quad Empty \quad \rightarrow (Node \ Empty@r \ x \ Empty@r)@r \\
&\quad Node \ i \ y \ d \rightarrow \mathbf{let} \ c = compare \ x \ y \\
&\quad \quad \mathbf{in} \ \mathbf{case} \ c \ \mathbf{of} \\
&\quad \quad \quad LT \rightarrow (Node \ (insertD \ x \ i)@r \ y \ d!)@r \\
&\quad \quad \quad EQ \rightarrow (Node \ i! \ y \ d!)@r \\
&\quad \quad \quad GT \rightarrow (Node \ i! \ y \ (insertD \ x \ d)@r)@r
\end{aligned}$$
FIGURE 7.3. Destructive insertion with reuse in a binary search tree

mentioned. The first example is the function that reverses a list and, at the same time, destroys it. The code is shown in Figure 7.2. We use the usual auxiliary function with an accumulator parameter. Notice that the differences with the usual functional version are, on the one hand, the use of the region parameter r and, on the other, that a **case!** is used over the original list. The recursive application of the function destroys it completely. Those who call $revD$ should know that the argument is lost in the inversion process, and should not try to use it anymore. This is reflected in the type of the first argument with a **!** annotation.

The next example illustrates the reuse of a condemned structure. It is the function, shown in Figure 7.3, that inserts an element in a binary search tree in such a way that the original tree is partially destroyed. Everything but the path from the root to the inserted element is reused to build the new tree but these parts can no longer be accessed from the original tree.

Notice that when the inserted element is already in the tree (EQ branch) the tree t that has just been destroyed is rebuilt. The purely functional version is obtained by removing the **!** annotations and returning t in the EQ branch.

7.2.2 Big-Step Operational Semantics

We have developed a big-step operational semantics for this language and a small-step operational semantics which have been proved equivalent. In Figure 7.4 we show the big-step operational semantics for *Safe* expressions. A **judgment** of the

$$\begin{array}{c}
\Delta, k : c \Downarrow \Delta, k : c \quad [Lit] \\
\Delta, k : C\bar{a}_i^n @ j \Downarrow \Delta, k : C\bar{a}_i^n @ j \quad [Cons] \\
\Delta[p \mapsto w], k : p \Downarrow \Delta, k : w \quad [Var_1] \\
\frac{j \leq k \quad l \neq j \quad (\Theta, C\bar{a}_i^n) = copy(\Delta, j, C\bar{a}_i^n)}{\Delta[p \mapsto (l, C\bar{a}_i^n)], k : p @ j \Downarrow \Theta, k : C\bar{a}_i^n @ j} \quad [Var_2] \\
\Delta \cup [p \mapsto w], k : p! \Downarrow \Delta, k : w \quad [Var_3] \\
\frac{\Sigma \vdash f \bar{x}_i^n = e \quad \Delta, k+1 : e[\bar{a}_i/\bar{x}_i, k+1/self] \Downarrow \Theta, k'+1 : v}{\Delta, k : f \bar{a}_i^n \Downarrow \Theta |_{k'}, k' : v} \quad [App_1] \\
\frac{\Sigma \vdash f \bar{x}_i^n r = e \quad \Delta, k+1 : e[\bar{a}_i/\bar{x}_i, k+1/self, j/r] \Downarrow \Theta, k'+1 : v}{\Delta, k : f \bar{a}_i^n @ j \Downarrow \Theta |_{k'}, k' : v} \quad [App_2] \\
\frac{\Delta, k : e_1 \Downarrow \Theta, k' : c \quad \Theta, k' : e[c/x_1] \Downarrow \Psi, k'' : v}{\Delta, k : \mathbf{let} \ x_1 = e_1 \ \mathbf{in} \ e \Downarrow \Psi, k'' : v} \quad [Let_1] \\
\frac{\Delta, k : e_1 \Downarrow \Theta, k' : C\bar{a}_i^n @ j \quad j \leq k' \ \mathit{fresh}(p) \quad \Theta \cup [p \mapsto (j, C\bar{a}_i^n)], k' : e[p/x_1] \Downarrow \Psi, k'' : v}{\Delta, k : \mathbf{let} \ x_1 = e_1 \ \mathbf{in} \ e \Downarrow \Psi, k'' : v} \quad [Let_2] \\
\frac{C = C_r \quad \Delta, k : e_r[\bar{a}_j/\bar{x}_j] \Downarrow \Theta, k' : v}{\Delta[p \mapsto (j, C\bar{a}_i^{nr})], k : \mathbf{case} \ p \ \mathbf{of} \ \bar{C}_i \ \bar{x}_i^{nr} \rightarrow e_i^m \Downarrow \Theta, k' : v} \quad [Case] \\
\frac{C = C_r \quad \Delta, k : e_r[\bar{a}_j/\bar{x}_j] \Downarrow \Theta, k' : v}{\Delta \cup [p \mapsto (j, C\bar{a}_i^{nr})], k : \mathbf{case!} \ p \ \mathbf{of} \ \bar{C}_i \ \bar{x}_i^{nr} \rightarrow e_i^m \Downarrow \Theta, k' : v} \quad [Case!]
\end{array}$$

FIGURE 7.4. SAFE big-step operational semantics

form $\Delta, k : e \Downarrow \Theta, k' : v$ means that expression e is successfully reduced to normal form v under heap Δ with $k+1$ regions, ranging from 0 to k , and that a final heap Θ with $k'+1$ regions is produced as a side effect.

A **heap** Δ is a function from fresh variables p (in fact, heap pointers) to closures w of the form $(j, C\bar{a}_i^n)$, meaning that the closure resides in region j . If $[p \mapsto w] \in \Delta$ and $w = (j, C\bar{a}_i^n)$, we will say that $region(w) = j$ and also that $region(p) = j$.

A **normal form** v is either a basic value c or a construction $C\bar{a}_i^n @ j$ to be stored in region j . The actual parameters a_i are either basic values or pointers to other closures. Actual region identifiers j are just natural numbers. Formal regions appearing in a function body are either the formal parameter r or the constant $self$.

By $\Delta[p \mapsto w]$ we denote a heap Δ where the binding $[p \mapsto w]$ is highlighted. In contrast, by $\Delta \cup [p \mapsto w]$ we denote the disjoint union of heap Δ with the binding $[p \mapsto w]$.

The semantics of a complete *Safe* program $d_1; \dots; d_n; e$ (not shown) is the semantics of the main expression e in an environment Σ containing the declarations

d_1, \dots, d_n of all the functions.

Rules *Lit* and *Cons* just say that basic values and constructions are normal forms. Rule *Cons* does not create a closure. Closures are actually created by rule *Let₂* which is the only one allocating fresh memory.

Rule *Var₁* brings a copy of a closure into the main expression. Rule *Var₂* makes a complete copy of the DS pointed to by a variable p into a new region j . Function *copy* follows the pointers in recursive positions of the original structure residing in region l and creates in region j a copy of all recursive closures except for the root closure $C\bar{a}_i^n$. In our runtime system we foresee that some type information is available so that it is possible to implement this function.

Should *copy* find a dangling pointer during the traversal, the whole rule would fail and the derivation would be stuck at this rule. If there is no failure, then the main expression becomes a copy $C\bar{a}'_i^n$ of this root closure where the pointers a_i in recursive positions pointing to closures in region l have been replaced by pointers a'_i to the corresponding closures in region j . The pointers in non recursive positions of all the copied closures are kept identical in the new closures. This implies that both DSs, the old and the new, may share some sub-structures. For instance, if the original DS is a list of lists, the structure created by *copy* is a copy of the outermost list, while the innermost lists become shared between the old and the new list.

Rule *Var₃* is similar to rule *Var₁* except for the fact that the binding $[p \mapsto w]$ is deleted and p does not belong to the domain of the resulting heap. This action may create dangling pointers in the living heap as some closures may have free occurrences of p .

Rules *App₁* and *App₂* show when a new region is created. Notice that the body of the function is executed in a heap with $k+2$ regions. That is, the formal identifier *self* is bound to the new region $k+1$ so that the function body may create DSs in this region or pass this region as a parameter to function calls. By $\Theta|_{k'}$ we denote the heap Θ restricted to closures belonging at most to region k' . In other words, before returning from the function, all closures created in region $k'+1$ are deleted. This action is another source of possible dangling pointers.

Rules *Let₁* and *Let₂* show the eagerness of the language: first, the auxiliary expression e_1 is reduced to normal form and then the main expression is evaluated. The occurrences of the program variable x_1 are replaced either by the normal form if it is a basic value, or by a pointer to it if it is a construction. Notice also that a construction is converted into a closure only if it is bound to a variable in a **let**.

Finally, rule *Case* is the usual one while rule *Case!* expresses what happens in a destructive pattern matching: the binding of the discriminant variable p disappears from the heap. This action is the last source of possible dangling pointers.

Proposition 7.4. *If $\Delta, k : e \Downarrow \Theta, k' : v$ is derivable, then $k = k'$.*

Proof: Straightforward, by induction on the depth of the derivation.

In the following, we will feel free to write the derivable judgments as $\Delta, k : e \Downarrow \Theta, k : v$.

By $fv(e)$ we denote the set of free variables of expression e , excluding function

names and region variables, and by $frv(e)$, the set of free region variables of e . By *Fresh*, we denote the set of names from which the function *fresh* in rule *Let*₂ selects fresh names, and by \mathbb{N} the set of natural numbers. Also, by $dom(\Delta)$ and $range(\Delta)$ we denote the following sets:

$$\begin{aligned} dom(\Delta) &\stackrel{\text{def}}{=} \{p \mid [p \mapsto w] \in \Delta\} \\ range(\Delta) &\stackrel{\text{def}}{=} \bigcup \{fv(w) \mid [p \mapsto w] \in \Delta\} \end{aligned}$$

Proposition 7.5. *If e is an expression satisfying $fv(e) \subseteq \text{Fresh}$ and $frv(e) \subseteq \mathbb{N}$, and $\Delta, k : e \Downarrow \Theta, k : v$ is derivable, and $range(\Delta) \subseteq \text{Fresh}$, then all judgments $\Delta_i, k_i : e_i \Downarrow \Theta_i, k_i : v_i$ of the derivation satisfy:*

1. $fv(e_i) \cup fv(v_i) \subseteq \text{Fresh}$.
2. $frv(e_i) \cup frv(v_i) \subseteq \mathbb{N}$.

Proof: By induction on the depth of the derivation.

For this reason, in the rules of Figure 7.4 we have systematically used letter p —intended to mean a pointer— when referring to free variables, and letter j —intended to mean a natural number— when referring to free region variables.

7.3 SHARING ANALYSIS

In this section we define an analysis that approximates the sharing relations between the variables of a program. At this point, Hindley-Milner types have already been inferred (see implementation details in Section 7.4), so the analysis can ask for the type of a variable through a function called *type*.

7.3.1 Sharing relations

In order to capture sharing, we define four different binary relations between variables:

Definition 7.6. *Given two variables x and y , in scope in an expression,*

1. $x \triangleleft \sim y$ denotes that x is a recursive descendant of y .
2. $x \triangle \sim y$ denotes that x shares a recursive descendant of y .
3. $x \triangleleft y$ denotes that x is any substructure of y .
4. $x \triangle y$ denotes that x shares any substructure of y .

In Figure 7.5 we illustrate these relations using trees to represent data structures in the heap. A black subtree represents a recursive substructure while a white subtree represents any substructure (recursive or not).

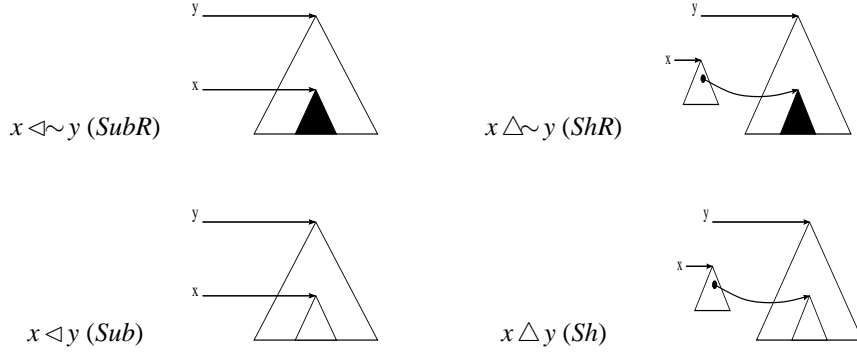


FIGURE 7.5. Sharing relations

We note that all the four relations are reflexive, Δ is also symmetric, and $\triangleleft\sim$ and \triangleleft are transitive. Moreover, the following implications hold:

$$\begin{aligned} x \triangleleft\sim y &\Rightarrow x \triangleleft y \Rightarrow x \Delta y \\ x \triangleleft\sim y &\Rightarrow x \Delta\sim y \Rightarrow x \Delta y \end{aligned}$$

but \triangleleft and $\Delta\sim$ do not necessarily imply each other.

The interpretation defined in Figure 7.6 does a top-down traversal of a program, accumulating these relations as soon as bound variables become free variables.

Whenever convenient, non-symmetric relations can be read as functions $Var \rightarrow \{Var\}$, giving $R(x)$ the set of all y such that yRx (i.e. $(y,x) \in R$). Also we will write $R = [x \rightarrow S]$ to indicate that $S = R(x)$.

The symmetric relation Δ is kept in a set of sets of variables. If $S \in \Delta$ then $x \Delta y$ for all $x, y \in S$.

Based on the above considerations, we will define an abstract interpretation S (meaning *sharing*) which, given an expression e delivers the following seven sets:

$$(SubRP, ShRP, SubP, SubR, ShR, Sub, Sh)$$

which contain respectively all the variables z such that $e \triangleleft\sim z$, $e \Delta\sim z$, $e \triangleleft z$, $z \triangleleft\sim e$, $z \Delta\sim e$, $z \triangleleft e$ and $z \Delta e$, where eRx and xRe means that the normal form of e , when evaluated at runtime, is related to x through R .

7.3.2 Function signatures

In order to achieve a modular analysis, we decide to reflect the result of the analysis of a function f in a *function signature*. We keep these signatures in a function environment ρ . A function signature $\rho(f)$ has the following type: $(\{Int\}, \{Int\}, \{Int\}, \{Int\}, \{Int\}, \{Int\}, \{Int\})$.

The meaning of the seven sets is as above, except for the fact that these contain only parameter indexes instead of all the (free and bound) variables of the body

expression. This is reasonable as the effect of a function should be completely reflected in the relationship between the parameters and the result.

In Figure 7.6 the interpretation S for expressions is defined. We explain it in detail later. When applied to a function definition $f\ x_1 \dots x_n = e$, it is straightforward to extract the signature of the function while computing the least fixpoint, in case it is recursive. The interpretation of a definition adds the signature of the new definition to the signatures environment:

$$\begin{aligned} S_d[[f\ x_1 \dots x_n = e]]\ \rho &= \text{fix } (\lambda\rho.\rho\ [f \rightarrow \text{extract}([x_1, \dots, x_n], S[[e]]\ R_0\ R_0\ R_0\ R_0\ \rho)])\ \rho_0 \\ \text{where } \rho_0 &= \rho\ [f \rightarrow (\emptyset, \emptyset, \emptyset, \emptyset, \emptyset, \emptyset)] \\ R_0 &= \{(x_i, x_i) \mid i \in \{1..n\}\} \\ \text{extract}(xs, (S_1, \dots, S_7)) &= (\{i \mid x_i \in xs \cap S_1\}, \dots, \{i \mid x_i \in xs \cap S_7\}) \end{aligned}$$

where $\rho\ [f \rightarrow s]$ either adds signature s for f or replaces it in case there was already one for it. As function S and function extract are monotone over a finite lattice, the least fixpoint exists and can be computed using Kleene's ascending chain.

Given a whole program $P = \text{dec}_1; \dots \text{dec}_k$; e the analysis first builds an increasing function environment and then analyses the main expression given initially empty relations (there are no free variables but function names):

$$S_p[[P]] = S[[e]]\ \emptyset\ \emptyset\ \emptyset\ \emptyset\ (S_d[[\text{dec}_k]]\ (\dots (S_d[[\text{dec}_1]]\ [])\ \dots))$$

Notice that the right hand sides of the definitions are analysed given relations where each parameter is only related to itself. This means that the signatures are computed assuming that all the parameters are disjoint. When they are not, the function application computes the additional sharing.

7.3.3 Interpretation of expressions

We explain now the details of the interpretation S for expressions. By abuse of notation, we will write $Sh(x)$ even though Sh is not a function, with the following convention:

$$Sh(x) \stackrel{\text{def}}{=} \bigcup \{S \mid x \in S \wedge S \in Sh\}$$

A basic value c neither has substructures nor is part of any structure, so its interpretation is just seven empty sets.

If x is returned as the result of a function, we use the information in the accumulator parameters of S to extract all the relevant information about its sharing. Notice that, from the operational semantics point of view, $x!$ is just the same structure as x , hence its interpretation. The semantics of $x@r$ is the creation of a copy of the recursive part of x in a new region r . As a consequence, the first, third, fourth and fifth sets of the interpretation are empty, and the third set excludes those variables with the same (recursive) type as x . The non-recursive part of $x@r$ is shared with x and potentially with any variable sharing substructures with x , hence the seventh set. However only the non-recursive children of x may be children of $x@r$, hence the sixth set.

$$\begin{aligned}
S \llbracket c \rrbracket \text{SubR ShR Sub Sh } \rho &= (\emptyset, \emptyset, \emptyset, \emptyset, \emptyset, \emptyset) \\
S \llbracket x \rrbracket \text{SubR ShR Sub Sh } \rho &= (\{z \mid x \in \text{SubR}(z)\}, \\
&\quad \{z \mid x \in \text{ShR}(z)\}, \\
&\quad \{z \mid x \in \text{Sub}(z)\}, \\
&\quad \text{SubR}(x), \text{ShR}(x), \text{Sub}(x), \text{Sh}(x)) \\
S \llbracket x! \rrbracket \text{SubR ShR Sub Sh } \rho &= S \llbracket x \rrbracket \text{SubR ShR Sub Sh } \rho \\
S \llbracket x@r \rrbracket \text{SubR ShR Sub Sh } \rho &= (\emptyset, \\
&\quad \{z \mid x \in \text{ShR}(z) \wedge \text{type}(z) \neq \text{type}(x)\}, \\
&\quad \emptyset, \emptyset, \emptyset, \\
&\quad \text{Sub}(x) - \text{SubR}(x), \text{Sh}(x)) \\
S \llbracket g \overline{a_i^m} @r \rrbracket \text{SubR ShR Sub Sh } \rho &= (\{z \mid \exists j \in \text{SubR}P_g.a_j \in \text{SubR}(z)\}, \\
&\quad \{z \mid \exists j \in \text{ShR}P_g.a_j \in \text{ShR}(z)\}, \\
&\quad \{z \mid \exists j \in \text{Sub}P_g.a_j \in \text{Sub}(z)\}, \\
&\quad \bigcup_j \{\text{SubR}(a_j) \mid j \in \text{SubR}g\}, \\
&\quad \bigcup_j \{\text{ShR}(a_j) \mid j \in \text{ShR}g\}, \\
&\quad \bigcup_j \{\text{Sub}(a_j) \mid j \in \text{Sub}g\}, \\
&\quad \bigcup_j \{\text{Sh}(a_j) \mid j \in \text{Sh}g\}) \\
\text{where } (\text{SubR}P_g, \text{ShR}P_g, \text{Sub}P_g, \text{SubR}g, \text{ShR}g, \text{Sub}g, \text{Sh}g) = \rho(g) \\
S \llbracket C \overline{a_i^m} @r \rrbracket \text{SubR ShR Sub Sh } \rho &= (\emptyset, \\
&\quad \{z \mid \exists a_j \in \text{ShR}(z)\}, \\
&\quad \emptyset, \\
&\quad \bigcup_j \{\text{SubR}(a_j) \mid j \in \text{RecPos}(C)\}, \\
&\quad \bigcup_j \{\text{ShR}(a_j) \mid j \in \text{RecPos}(C)\}, \\
&\quad \bigcup_j \{\text{Sub}(a_j) \mid j \in \{1..m\}\}, \\
&\quad \bigcup_j \{\text{Sh}(a_j) \mid j \in \{1..m\}\}) \\
S \llbracket \text{let } x_1 = e_1 \text{ in } e \rrbracket \text{SubR ShR Sub Sh } \rho &= (S \llbracket e \rrbracket \text{SubR}_2 \text{ShR}_2 \text{Sub}_2 \text{Sh}_2 \rho) \setminus \{x_1\} \\
\text{where } (\text{SubR}P_1, \text{ShR}P_1, \text{Sub}P_1, \text{SubR}_1, \text{ShR}_1, \text{Sub}_1, \text{Sh}_1) &= S \llbracket e_1 \rrbracket \text{SubR ShR Sub Sh } \rho \\
\text{SubR}_2 &= (\text{SubR} \cup [x_1 \mapsto \text{SubR}_1] \cup \{[z \mapsto \{x_1\}] \mid z \in \text{SubR}P_1\})^* \\
\text{ShR}_2 &= \text{ShR} \cup [x_1 \mapsto \text{ShR}_1] \cup \{[z \mapsto \{x_1\}] \mid z \in \text{ShR}P_1\} \cup \text{SubR}_2 \\
\text{Sub}_2 &= (\text{Sub} \cup [x_1 \mapsto \text{Sub}_1] \cup \{[z \mapsto \{x_1\}] \mid z \in \text{Sub}P_1\})^* \\
\text{Sh}_2 &= \text{Sh} \cup \{x_1\} \cup \text{Sh}_1 \uplus (\text{Sub}_2 \cup \text{ShR}_2) \\
S \llbracket \text{case } x \text{ of } \overline{C_i \overline{x_{ij}^{n_i}} \rightarrow e_i} \rrbracket \text{SubR ShR Sub Sh } \rho &= \bigcup_i (S \llbracket e_i \rrbracket \text{SubR}_i \text{ShR}_i \text{Sub}_i \text{Sh}_i \rho) \setminus \{\overline{x_{ij}^{n_i}}\} \\
\text{where } \text{SubR}_i &= (\text{SubR} \cup [x \mapsto \{x_{ij} \mid j \in \text{RecPos}(C_i)\}] \\
&\quad \cup \{[x_{ij} \mapsto \text{SubR}(x) \setminus \{x\}] \mid j \in \text{RecPos}(C_i)\})^* \\
\text{ShR}_i &= \text{ShR} \cup \{[x_{ij} \mapsto \text{ShR}(x)] \mid j \in \text{RecPos}(C_i)\} \cup \text{SubR}_i \\
\text{Sub}_i &= (\text{Sub} \cup [x \mapsto \{x_{ij} \mid j \in \{1..n_i\}\}] \cup \{[x_{ij} \mapsto \text{Sub}(x) \setminus \{x\}] \mid j \in \{1..n_i\}\})^* \\
\text{Sh}_i &= (\text{Sh} \cup \{y, x_{ij}\} \mid y \in \text{Sh}(x) \wedge j \in \{1..n_i\}) \uplus (\text{Sub}_i \cup \text{ShR}_i)
\end{aligned}$$

FIGURE 7.6. Definition of the abstract interpretation S

The interpretation of a function application $g \overline{a}_i^m @ r$ returning a DS is rather involved. Regarding the first set, the recursive descendant relation is transitive. So, the result of g is a recursive descendant of a variable z if and only if an actual parameter a_j of g is a recursive descendant of z and the result of g is a recursive descendant of a_j . The same transitivity applies to the third set. Regarding the second set, the result of g shares a recursive descendant of a variable z if an actual parameter a_j of g shares a recursive descendant of z , and a_j is in sharing relation with the result of g . This probably will give us more variables than the ones actually sharing a recursive descendant of z , but it is a safe approximation. This is a place where signatures may lose information. The fourth and sixth sets are defined taking respectively into account the transitivity of the relations $\triangleleft \sim$ and \triangleleft . The fifth and seventh sets are safe, but may be imprecise, approximations to respectively the set of variables sharing a recursive substructure and sharing any substructure with the result of g . The interpretation of a function application $g \overline{a}_i^m$ of a function g not having an output region as a parameter is identical to the previous one.

In the interpretation of a data construction $C \overline{a}_i^m @ r$, the first and third sets are empty because a newly created DS cannot be a substructure of any other. However, it will share a recursive descendant of a variable z if any of its substructures a_j already shared it. Any variable being a recursive descendant of a recursive parameter a_j of C will also be a recursive descendant of the construction. The set $RecPos(C)$ contains the recursive positions of the constructor C . Similar reasoning can be applied to the fifth set containing the variables which share a recursive descendant of the construction. The next set definition exploits the transitivity of the \triangleleft relation. The last set consists also of a union over all the parameters of C , because the construction inherits the sharing of all its substructures.

The **let** expression introduces a new bound variable x_1 which may appear free in the main expression e . First, the interpretation of the auxiliary expression e_1 is launched and the sharing created by it is accumulated in the parameters. Then, the main expression e is interpreted taking into account the new sharing. If R represents a reflexive, non-symmetric, transitive relation, by R^* we mean its reflexive, transitive closure. Operator \uplus computes the union of a reflexive, symmetric and non-transitive relation and a reflexive, non-symmetric transitive one. Notice that the addition of $SubR_2$ to ShR_2 , and the addition of this latter set and that of Sub_2 to Sh_2 just implements the inclusion of the underlying relations, as explained above. Finally, the information related to x_1 is deleted as the variable will not be in scope in the context.

As usual, the interpretation of a **case** is the least upper bound of the interpretation of its alternatives, and this involves a loss of information. Before each alternative is interpreted, we accumulate the sharing of the bound variables x_{ij} introduced by it. Part of this sharing is straightforward: all these variables are descendants of the parent structure x and some of them are recursive descendants of it. Additionally, if we have $y \in SubR(x) \wedge y \neq x$, that means $y \triangleleft \sim x$. As there is no more information available, it may be the case that $y \triangleleft \sim x_{ij}$ for some recursive child of x . The only safe way to cope with this possibility is to include in $SubR_i$

the pairs $y \triangleleft \sim x_{ij}$ for all the recursive children of x . Similar reasoning applies to the rest of the sets.

The interpretation of **case!** is the same as the previous one. Although the discriminant variable is being condemned we cannot eliminate its sharing information as we do not know whether the rest of variables are safely used. For example, we could write $z = \mathbf{case!} \ x \ \mathbf{of} \ C \ y \rightarrow x$. The analysis says that variables x and z share a substructure, although such sharing is unsafe because x has been destroyed.

7.4 IMPLEMENTATION AND EXAMPLES

In this section we present the implementation of the analysis and give some examples of functions to which it has been applied. We have defined a concrete sugared syntax for *Safe* in which programs look very much like Haskell programs, i.e. functions are defined by means of equations and pattern matching, guards and **where** clauses are allowed, as well as data type declarations and infix operators and constructors.

A complete front-end has been developed from scratch by using standard tools such as lexical analyzer and parser generators. In Figure 7.7 we show its (already implemented) phases. The renamer phase ensures that every identifier is well defined and that every bound variable is given a different name. A Hindley-Milner type inference is done at this level in order to reject ill-typed programs, and to provide report messages related to the sugared syntax. Also, the sharing analysis needs the underlying type of a variable and the recursive positions of data constructors (cf. Figure 7.6). This phase decorates each expression in the abstract syntax tree with its Hindley-Milner type.

The desugarer transforms the high-level syntax into the *Safe* core syntax presented in Section 7.2. During this transformation new bound variables may be introduced. They are given appropriate types and fresh names.

After these steps, the sharing analysis described in this paper is done. Its main function has the following type:

```
analyzePrg :: Prog TypeExp -> Prog (TypeExp, Maybe ShareInfo)
```

That is, given a program decorated with Hindley-Milner types, it returns a program additionally decorated with sharing information. This sharing information has different shapes depending on the entity being decorated:

- If it is a function definition, it consists of its signature.
- If it is a **let** or a **case!** expression, it consists of the sharing information accumulated from the beginning of the function body this expression belongs to, up to the root of the expression. These are the only expressions where we need to keep the sharing information, which consists of the seven sets corresponding to the variable either defined by the **let** expression, or inspected by the **case!** expression.
- Binding occurrences of variables are not decorated.




FIGURE 7.7. Phases of the Safe compiler implementation

In this way, it is easy to extract the *shareall* and *sharerec* sets for any given variable x in any given context. Then, $ShR(x)$ and $Sh(x)$ give us the desired information. Such information is used in the following phase, whose result consists of Hindley-Milner types decorated with destruction (!) annotations, i.e. *Safe* types. This phase is also implemented but it is not part of this paper.

The front-end and the analysis have been implemented in Haskell using the GHC to compile it. In total, about 3,000 Haskell lines have been written. In order to improve efficiency, the analyzer stores the four relations in a single balanced tree, using the modules `Map` and `Set` of the GHC library [1]. Also, the inverses of the three first relations are kept in the tree. In this way, the symmetric and/or transitive closures, the union, and some other operations on relations needed by the analysis, are done in a more concise and efficient way. Let n be the number of bound variables of a function body, m the size of its abstract syntax tree, and p the number of function arguments. Then, the analysis cost is in $O(nmp)$ in the worst case. The analysis of a function can be done independently of each other.

When applied to the functions defined in Section 7.2, the analysis computes the following signatures:

$$\begin{aligned}
 \rho(\text{revauxD}) &= (\{2\}, \{2\}, \{2\}, \{2\}, \{2\}, \{2\}, \{1, 2\}) \\
 \rho(\text{revD}) &= (\emptyset, \emptyset, \emptyset, \emptyset, \emptyset, \emptyset, \{1\}) \\
 \rho(\text{insertD}) &= (\emptyset, \{1, 2\}, \emptyset, \emptyset, \{2\}, \{1\}, \{1, 2\})
 \end{aligned}$$

which are accurate descriptions of the input-output sharing relations of these functions.

Function *revauxD* appends the reverse of its first parameter to its second one. Since it does not reuse the recursive cells of its first parameter, the only remaining recursive sharing is related to its second parameter. Nevertheless the sharing with the non recursive elements of the first list is reflected in the last set of the signature. Function *revD* consists of a simple call to *revauxD* passing it an empty list as the second actual parameter, so the only remaining sharing is that between the non-

$$\begin{aligned}
\mathit{splitD} &:: \forall a, \rho. \mathit{Int} \rightarrow [a]!@ \rho \rightarrow \rho \rightarrow ([a]@ \rho, [a]@ \rho)@ \rho \\
\mathit{splitD} \ 0 \ xs! \ r &= ([]@ r, xs!)@ r \\
\mathit{splitD} \ n \ []! \ r &= ([]@ r, []@ r)@ r \\
\mathit{splitD} \ n \ (x : xs)! \ r &= ((x : xs_1)@ r, xs_2)@ r \\
&\quad \mathbf{where} \ (xs_1, xs_2) = \mathit{splitD} \ (n - 1) \ xs \ r \\
\\
\mathit{mergeD} &:: \forall a, \rho. [a]!@ \rho \rightarrow [a]!@ \rho \rightarrow \rho \rightarrow [a]@ \rho \\
\mathit{mergeD} \ []! \ ys! \ r &= ys! \\
\mathit{mergeD} \ (x : xs)! \ []! \ r &= (x : xs!)@ r \\
\mathit{mergeD} \ (x : xs)! \ (y : ys)! \ r &= \\
&\quad \begin{cases} x \leq y &= (x : \mathit{mergeD} \ xs \ (y : ys!)@ r \ @r)@ r \\ \textit{otherwise} &= (y : \mathit{mergeD} \ (x : xs!)@ r \ ys \ @r)@ r \end{cases} \\
\\
\mathit{msortD} &:: \forall a, \rho. [a]!@ \rho \rightarrow \rho \rightarrow [a]@ \rho \\
\mathit{msortD} \ xs \ r &= \\
&\quad \begin{cases} n \leq 1 &= xs! \\ \textit{otherwise} &= \mathit{mergeD} \ (\mathit{msortD} \ xs_1 \ @r) \ (\mathit{msortD} \ xs_2 \ @r) \ @r \end{cases} \\
&\quad \mathbf{where} \ (xs_1, xs_2) = \mathit{splitD} \ (n \ \textit{div} \ 2) \ xs \ @r \\
&\quad \quad n = \mathit{length} \ xs
\end{aligned}$$

FIGURE 7.8. Destructive mergesort

recursive structures of the input and output lists.

Function *insertD* builds a new tree which shares with the original tree everything but the path from the root to the inserted element. This means that the resulting tree and the original one share both recursive and non-recursive parts. This is the reason why 2 appears in the second, fifth and seventh sets of the signature. Also the resulting tree has *x* as a non-recursive descendant, so the 1 in the second, sixth and seventh sets.

7.5 A MORE INVOLVED EXAMPLE

In this section we show a more involved example, a *mergesort* algorithm. In order to give compact code, the functions shown in this section are sugared although the analysis is executed over their desugared versions.

First, we define auxiliary functions to split the input list and merge two ordered lists in a single ordered list. In Figure 7.8 (top) we show a destructive version of the splitting function. As in the previous examples, there are small differences with a purely functional version. In the base case ($n = 0$) we reuse the list in the output; in the recursive case we use a **case!** (written as a destructive pattern) over the argument list. We also have to add *@r* where necessary.

The sharing analysis produces the following signature for this function:

$$\rho(\mathit{splitD}) = (\emptyset, \{2\}, \emptyset, \emptyset, \emptyset, \{2\}, \{2\})$$

meaning that:

- The result of the function may share a recursive substructure of the argument list, which is obvious.
- The argument list may be a child of the result, which is true when n is 0.
- The argument list and the result share some substructure, which again is obvious.

Figure 7.8 (middle) shows the destructive version of the merging function. In the recursive calls to *mergeD* one of the parameters is one of the original lists. But the original list may not be referenced as its top cell has been destroyed by a **case!**, so the original list is rebuilt by reusing its components. This is the only detail to care about.

The sharing analysis produces the following signature for this function:

$$\rho(\text{mergeD}) = (\{2\}, \{1, 2\}, \{2\}, \{2\}, \{1, 2\}, \{2\}, \{1, 2\})$$

meaning that the argument lists and the result may share recursive and non-recursive substructures one of the other. Notice that only the second argument list may be a recursive child of the result (and viceversa) because we build a new cell for each cell of the first argument while we reuse the second argument list when the first one is empty.

Finally, in Figure 7.8 (bottom) we show the destructive mergesort *msortD*, that uses the previously defined functions. Both the input list xs and the intermediate results are either destroyed or reused into the result. This allows us to conclude that this function consumes a constant additional heap space. In [11] we proved this by induction on the length of the argument list. The sharing analysis produces the following signature for this function:

$$\rho(\text{msortD}) = (\{1\}, \{1\}, \{1\}, \{1\}, \{1\}, \{1\}, \{1\})$$

meaning that the argument list and the result may share recursive and non-recursive substructures one of the other.

Recall that this sharing analysis does not take into account the fact that some substructures are destroyed because it is not known yet whether the program is type-safe. In this sense the analysis is an upper approximation of the sharing.

7.6 RELATED WORK AND CONCLUSIONS

Several approaches have been taken to memory management, some of which have inspired our work. In [5] a comparison of some of them are presented by using a Game of Life example:

```

nextgen g = {create and return new generation}
life n g = if n = 0 then g
           else life (n - 1) (nextgen g)

```


Assuming that a generation g is a big data structure allocated in the heap, a functional program like this would allocate n generations in the heap until a garbage collector would decide to dispose the intermediate ones. However, if the intended use of one intermediate generation is only the creation of the next one, it seems reasonable to dispose the intermediate data structure as soon as possible. In *Safe* we would modify the program as follows in order to get such behaviour:

$$\begin{aligned} \text{nextgen } g \ r &= \mathbf{case!} \ g \ \mathbf{of} \ \rightarrow \dots \\ &\quad \{\text{create new generation in region } r\} \\ \text{life } n \ g \ r &= \mathbf{if} \ n = 0 \ \mathbf{then} \ g! \quad \{\text{reuse argument } g\} \\ &\quad \mathbf{else} \ \text{life } (n - 1) \ (\text{nextgen } g@r)@r \end{aligned}$$

Tofte and Talpin [13] introduced the use of nested regions with a **letregion** ρ construct as an extension to Core ML. Like ours, regions are memory areas where DSs can be constructed, and they are allocated and deallocated as a whole. A difference is that, in our system, region allocation/deallocation are synchronized with function calls. Also, we have an additional mechanism that allows us to selectively destroy DSs in the working or in the output region. In their framework, in the previous example a single region is forced to contain all the intermediate data structures and no memory advantages are obtained.

An extension to their work [3, 12] allows to *reset* all the data structures in a region without deallocating the whole region. In the previous example the old generation region is resetted once the new generation is created. So, a new temporary region is created to allocate the new generation which must be copied into the output region after resetting it. The user is responsible for introducing the copy functions but not for annotating the program with resetting annotations.

$$\begin{aligned} \text{nextgen } g &= \{\text{create and return new generation}\} \\ \text{life } n \ g &= \mathbf{if} \ n = 0 \ \mathbf{then} \ g \\ &\quad \mathbf{else} \ \text{life } (n - 1) \ (\mathbf{copy} \ (\text{nextgen } g)) \end{aligned}$$

The *copy* function allows to build the new generation in a separated region and makes possible to run *life* in constant heap space. However this version may waste a lot of time just in copying, once for each recursive call. Additionally, inserting the *copy* function requires a deep knowledge of the resetting mechanism as this is not explicit in the program. In our opinion, the **case!** annotation is more intuitive: the user just says that the old generation may be liberated as it will not be used anymore. And it is only said for a particular data structure, not for the whole region.

The AFL system [2] inserts (as a result of an analysis) allocation and deallocation commands separated from the **letregion** construct which now only brings new regions into scope. In the example, this allows to free the old region as soon as the new generation is computed, without needing a copy in each recursive call. This is only required in the base case:

$$\begin{aligned} \text{nextgen } g &= \{\text{create and return new generation}\} \\ \text{life } n \ g &= \mathbf{if} \ n = 0 \ \mathbf{then} \ \mathbf{copy} \ g \\ &\quad \mathbf{else} \ \text{life } (n - 1) \ (\text{nextgen } g) \end{aligned}$$

Again, inserting the *copy* function in the appropriate place requires a deep knowledge of the annotations that will be inserted after the analysis.

Our region system is simpler than those of the above approaches and it does not require such complex inference algorithms. Although the version of the language presented here has explicit regions, we have designed a region inference algorithm which hides them from the programmer. It is a simple extension of the Hindley-Milner type inference one.

Hughes and Pareto [8] also incorporate in Embedded-ML the concept of region to their sized-types system so that heap and stack consumption can be type-checked. In this approach, region sizes are bounded. Our main differences to them are again the region-function association and the explicit disposal of structures. Their sized types system could be a good starting point for our future work, as we also intend to compute region sizes at compile time.

More recently, in a proof carrying code framework, Hofmann and Jost [6] have developed a type system to infer heap consumption. Theirs is also a first-order eager functional language with a construct *match'* which has inspired our **case!**. They associate to each function the space required by its execution and the remaining space after it. They also use a linear type system but they do not achieve a complete safety in using destructive facilities. Unlike us they do not use the concept of nested regions where DSs are allocated, so that sharing is not controlled in the same way.

There are many works devoted to sharing analysis in functional and logic languages, some of them rather old. In the functional field, the aim of most analyses has been performing part of the garbage collection at compile time, or detecting when destructive updating of data structures could be done safely.

In Hudak's approach [7] a reference count of shared data is done at compile time by using abstract interpretation on a first order, eager functional language with updatable arrays. The abstract domains consist of just natural numbers. In order to have a terminating analysis the domains are restricted to finite intervals $\{1 \dots n\}$, for an arbitrary n , and topped with ∞ meaning 'too much sharing'. An array based quicksort algorithm using in place updating is shown correct by the analysis.

Jones and Le Métayer [10] use also abstract interpretation on a first order, eager functional language with non-homogeneous lists in order to avoid allocation of fresh cells and to reuse instead cells not needed by the rest of the computation. Their analysis is a combination of sharing and absence analyses and the abstract domains are nested tuples of booleans. Again, domains are forced to be finite by bounding the nesting depth of the tuples by an arbitrary number n . The analysis looks rather complex and not very efficient as it does several traversals of the same code. Also the authors do not show evidence of having implemented it.

Inoue et al [9] use non-standard techniques, such as context free languages and intersection between such languages, in order to perform garbage collection at compile time. The language analyzed is a first order subset of LISP. The idea is to detect cells created by a function and not belonging to the result. Such those cells are disposed at the end of the function body. They show good results for

some LISP test programs. In the logic programming field, Gudjonsson et al [4] provide a comprehensive survey of sharing analyses. Sharing is important here for much the same reasons than in the functional field but also to detect opportunities for parallel evaluation.

The main novelty of our approach is, on the one hand the context—a functional language with explicit destruction—and on the other its modularity. In the previously described works, the analyses are done at the whole program level while ours is done function by function. Reflecting the result of a function analysis in a signature provides the connection between the different functions of the program. The subsequent safety analysis, based on a special type system, will also be done function by function, so the sharing signature can be seen as an annotation associated to the function type.

We find our sharing analysis to be precise enough for successfully analysing the examples we have tried so far, but its quality will be evaluated when both phases, the sharing and the safety analysis, work together. The safety type system, not described here, has some characteristics of linear types (see [14] as a basic reference and [6] as a nearer one) and, as it has already been said, it heavily uses the result of the sharing analysis.

As future work, we will prove the correctness of the analysis with respect to the small-step operational semantics (not shown in this paper) of the language. Also, as we have already said, the region annotations $@r$ will be inferred so that the programmer will forget about regions: each data structure not sharing any substructure with the function result will be considered local, and consequently built in the working (self) region of the function. The rest of them will be built in the output region.

Our final aim is to develop a type based analysis that automatically infers memory consumption. A sized-types system could automate induction reasoning like the one mentioned in Section 7.5.

REFERENCES

- [1] S. Adams. Efficient sets –a balancing act. *Journal of Functional Programming*, 3(4):553–561, 1993.
- [2] A. Aiken, M. Fähndrich, and R. Levien. Better static memory management: improving region-based analysis of higher-order languages. In *Proceedings of the ACM SIGPLAN 1995 conference on Programming language design and implementation, PLDI'95*, pages 174–185. ACM Press, 1995.
- [3] L. Birkedal, M. Tofte, and M. Vejlstrup. From region inference to von neumann machines via region representation inference. In *Conference Record of POPL '96: The 23rd ACM SIGPLAN-SIGACT*, pages 171–183, 1996.
- [4] G. Gudjónsson and W. H. Winsborough. Compile-time memory reuse in logic programming languages through update in place. *ACM TOPLAS*, 21(3):430–501, 1999.
- [5] F. Henglein, H. Makholm, and H. Niss. A direct approach to control-flow sensitive region-based memory management. In *Proceedings of the 3rd ACM SIGPLAN international conference on Principles and Practice of Declarative Programming, PPDP'01*, pages 175–186. ACM Press, 2001.
- [6] M. Hofmann and S. Jost. Static prediction of heap space usage for first-order functional programs. In *Proceedings of the 30th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 185–197. ACM Press, 2003.
- [7] P. Hudak. A Semantic Model of Reference Counting and its Abstraction. In *Lisp and Functional Programming Conference*, pages 351–363. ACM Press, 1986.
- [8] R. J. M. Hughes and L. Pareto. Recursion and Dynamic Data-Structures in Bounded Space; Towards Embedded ML Programming. In *Proceedings of the Fourth ACM SIGPLAN International Conference on Functional Programming, ICFP'99*, ACM Sigplan Notices, pages 70–81, Paris, France, September 1999. ACM Press.
- [9] K. Inoue, H. Seki, and H. Yagi. Analysis of Functional Programs to Detect Run-Time Garbage Cells. *ACM TOPLAS*, 10(4):555–578, 1988.
- [10] S. B. Jones and D. Le Metayer. Compile Time Garbage Collection by Sharing Analysis. In *Int. Conf. on Functional Programming and Computer Architecture*, pages 54–74. ACM Press, 1989.
- [11] R. Peña and C. Segura. A First-Order Functional Language for Reasoning about Heap Consumption. In *16th International Workshop on Implementation and Application of Functional Languages, IFL'04. Technical Report 0408, Christian-Albrechts University of Kiel*, pages 64–80, 2004.
- [12] M. Tofte, L. Birkedal, M. Elsmann, N. Hallenberg, T. H. Olesen, and P. Sestoft. Programming with regions in the MLKit (revised for version 4.3.0). Technical report, IT University of Copenhagen, Denmark, 2006.
- [13] M. Tofte and J.-P. Talpin. Region-based memory management. *Information and Computation*, 132(2):109–176, 1997.
- [14] P. Wadler. Linear types can change the world! In *IFIP TC 2 Working Conference on Programming Concepts and Methods*. North Holland, 1990.