

**A peer-reviewed version of this preprint was published in PeerJ on 2 September 2015.**

[View the peer-reviewed version](https://doi.org/10.7717/peerj-cs.22) (peerj.com/articles/cs-22), which is the preferred citable publication unless you specifically need to cite this preprint.

Naish L. 2015. Sharing analysis in the Pawns compiler. PeerJ Computer Science 1:e22 <https://doi.org/10.7717/peerj-cs.22>

# Sharing analysis in the Pawns compiler

Lee Naish

Computing and Information Systems,  
University of Melbourne, Melbourne 3010, Australia  
lee@unimelb.edu.au,  
<http://people.eng.unimelb.edu.au/lee/>

**Abstract.** Pawns is a programming language under development which supports algebraic data types, polymorphism, higher order functions and “pure” declarative programming. It also supports impure imperative features including destructive update of shared data structures via pointers, allowing significantly increased efficiency for some operations. A novelty of Pawns is that all impure “effects” must be made obvious in the source code and they can be safely encapsulated in pure functions in a way that is checked by the compiler. Execution of a pure function can perform destructive updates on data structures which are local to or eventually returned from the function without risking modification of the data structures passed to the function. This paper describes the sharing analysis which allows impurity to be encapsulated. Aspects of the analysis are similar to other published work, but in addition it handles explicit pointers and destructive update, higher order functions including closures and pre- and postconditions concerning sharing for functions.

Keywords: functional programming language, destructive update, mutability, effects, algebraic data type, sharing analysis, aliasing analysis

## 1 Introduction

This paper describes the sharing analysis done by the compiler for Pawns [1], a programming language which is currently under development. Pawns supports both declarative and imperative styles of programming. It supports algebraic data types, polymorphism, higher order programming and “pure” declarative functions, allowing very high level reasoning about code. It also allows imperative code, where programmers can consider the representation of data types, obtain pointers to the arguments of data constructors and destructively update them. Such code requires the programmer to reason at a much lower level and consider aliasing of pointers and sharing of data structures. Low level “impure” code can be encapsulated within a pure interface and the compiler checks the purity. This requires analysis of pointer aliasing and data structure sharing, to distinguish data structures which are only visible to the low level code (and are therefore safe to update) from data structures which are passed in from the high level code (for which update would violate purity). The main aim of Pawns is to get

the benefits of purity for most code but still have the ability to write some key components using an imperative style, which can significantly improve efficiency (for example, a more than twenty-fold increase in the speed of inserting an element into a binary search tree).

There are other functional programming languages, such as ML [2], Haskell [3] and Disciple [4], which allow destructive update of shared data structures but do not allow this impurity to be encapsulated. In these languages the ability to update the data structure is connected to its type<sup>1</sup>. For a data structure to be built using destructive update its type must allow destructive update and any code which uses the data structure can potentially update it as well. This prevents simple declarative analysis of the code and can lead to a proliferation of different versions of a data structure, with different parts being mutable. There is often an efficiency penalty as well, with destructive update requiring an extra level of indirection in the data structure. Pawns avoids this inefficiency and separates mutability from type information, allowing a data structure to be mutable in some contexts and considered “pure” in others. The main cost from the programmer perspective is the need to include extra annotations and information in the source code. This can also be considered a benefit, as they provide useful documentation and error checking. The main implementation cost is additional analysis done by the compiler, which is the focus of this paper.

The rest of this paper assumes some familiarity with Haskell and is structured as follows. Section 2 gives a brief overview of the relevant features of Pawns and Section 3 describes a simple “core” language which source programs are translated into. Section 4 describes the abstract domain used for sharing analysis algorithm, Section 5 defines the algorithm itself and Section 6 gives an extended example. Section 7 briefly discusses precision and efficiency issues. Section 8 discusses related work and Section 9 concludes.

## 2 An overview of Pawns

A more detailed introduction to Pawns is given in [1]. Pawns has many similarities with other functional languages. It supports algebraic data types with parametric polymorphism, higher order programming and curried function definitions. It uses strict evaluation. In addition, it supports destructive update via “references” (pointers) and has a variety of extra annotations to make impure effects more clear from the source code and allow them to be encapsulated in pure code. Pawns also supports a form of global variables (called state variables) which support encapsulated effects, but we do not discuss them further here as they are handled in essentially the same way as other variables in sharing analysis. Pure code can be thought of in a declarative way, where values can be viewed abstractly, without considering how they are represented. Code which uses destructive update must be viewed at a lower level, considering the representation of values, including sharing. We discuss this lower level view first, then briefly

<sup>1</sup> Disciple uses “region” information to augment types, with similar consequences.

81 present how impurity can be encapsulated to support the high level view. We  
 82 use Haskell-like syntax for familiarity.

## 83 2.1 The low level view

84 Values in Pawns are represented as follows. Constants (data constructors with  
 85 no arguments) are represented using a value in a single word. A data constructor  
 86 with  $N > 0$  arguments is represented using a word that contains a tagged pointer  
 87 to a block of  $N$  words in main memory containing the arguments. For simple  
 88 data types such as lists the tag may be empty. In more complex cases some  
 89 bits of the pointer may be used and/or a tag may be stored in a word in main  
 90 memory along with the arguments. Note that constants and tagged pointers  
 91 are not always stored in main memory and Pawns variables may correspond to  
 92 registers that contain the value. Only the arguments of data constructors are  
 93 guaranteed to be in main memory. An array of size  $N$  is represented in the same  
 94 way as a data constructor with  $N$  arguments, with the size given by the tag.  
 95 Functions are represented as either a constant (for functions which are known  
 96 statically) or a closure which is a data constructor with a known function and a  
 97 number of other arguments.

98 Pawns has a `Ref t` type constructor, representing a reference/pointer to a  
 99 value of type `t` (which must be stored in memory). Conceptually we can think of  
 100 a corresponding `Ref` data constructor with a single argument, but this is never  
 101 explicit in Pawns code. Instead, there is an explicit dereference operation: `*vp`  
 102 denotes the value `vp` points to. There are two ways references can be created:  
 103 let bindings and pattern bindings. A let binding `*vp = val` allocates a word  
 104 in main memory, initializes it to `val` and makes `vp` a reference to it (Pawns  
 105 omits Haskell's `let` and `in` keywords; the scope is the following sequence of  
 106 statements/expressions). In a pattern binding, if `*vp` is the argument of a data  
 107 constructor pattern, `vp` is bound to a reference to the corresponding argument  
 108 of the data constructor if pattern matching succeeds (there is also a primitive  
 109 which returns a reference to the  $i^{th}$  element of an array). Note it is not possible  
 110 to obtain a reference to a Pawns variable: variables do not denote memory lo-  
 111 cations. However, a variable `vp` of type `Ref t` denotes a reference to a memory  
 112 location containing a value of type `t` and the memory location can be destruc-  
 113 tively updated by `*vp := val`.

114 Consider the following code. Two data types are defined. The code creates a  
 115 reference to `Nil` (`Nil` is stored in a newly allocated memory word) and a reference  
 116 to that reference (a pointer to the word containing `Nil` is put in another allocated  
 117 word). It also creates a list containing constants `Blue` and `Red` (requiring the  
 118 allocation of two cons cells in memory; the `Nil` is copied). It deconstructs the  
 119 list to obtain pointers to the head and tail of the list (the two words in the first  
 120 cons cell) then destructively updates the head of the list to be `Red`.


```
121 data Colour = Red | Green | Blue
122 data Colours = Nil | Cons Colour Colours -- like List Colour
```

```

123     ...
124     *np = Nil                -- np = ref to (copy of) Nil
125     *npp = np               -- npp = ref to (copy of) np
126     cols = Cons Blue (Cons Red *np) -- cols = [Blue, Red]
127     case cols of
128     (Cons *headp *tailp) ->   -- get ref to head and tail
129         *headp := Red         -- update head with Red

```

The memory layout after the assignment can be pictured as follows, where boxes represent main memory words and Ref and Cons followed by an arrow represent pointers (no tag is used in either case):



The destructive update above changes the values of both `headp` and `cols` (the representations are shared). One of the novel features of Pawns is that the source code must be annotated with “!” to make it obvious when each “live” variable is updated. If both `headp` and `cols` are used later, the assignment statement above must be written as follows, with `headp` prefixed with “!” and an additional annotation attached to the whole statement indicating `cols` may be updated:

```

141     *!headp := Red !cols      -- update *headp (and cols)

```

We say that the statement *directly* updates `headp` and *indirectly* updates `cols`, due to sharing of representations. Similarly, if `headp` was passed to a function which may update it, additional annotations are required. For example, `(assign !headp Red) !cols` makes the direct update of `headp` and indirect update of `cols` clear. Sharing analysis is used to ensure that source code contains all the necessary annotations. One aim of Pawns is that any effects of code should be made clear by the code. Pawns is an acronym for Pointer Assignment With No Surprises.

Pawns functions have extra annotations in type signatures to document which arguments may be updated. For additional documentation, and help in sharing analysis, there are annotations to declare what sharing may exist between arguments when the function is called (a precondition) and what extra sharing may be added by executing the function (called a postcondition, though it is the union of the pre- and post-condition which must be satisfied after a function is executed). For example, we may have:

```

157 assign :: Ref t -> t -> ()
158     sharing assign !p v = _    -- p may be updated
159     pre nosharing              -- p&v don't share when called
160     post *p = v                -- assign may make *p alias with v

```

As well as checking for annotations on assignments and function calls, sharing analysis is used to check that all arguments which may be updated are annotated in type signatures, and pre- and post-conditions are always satisfied. For example, assuming the previous code which binds `cols`, the call `assign !tailp !cols` annotates all modified variables but violates the precondition of `assign` because there is sharing between `tailp` and `cols` at the time of the call. Violating this precondition allows cyclic structures to be created, which is important for understanding the code. In general, there is an inter-dependence between “!” annotations in the code and pre- and post-conditions. More possible sharing at a call means more “!” annotations are needed, more sharing in (recursive) calls and more sharing when the function returns.

Curried functions and higher order code are supported by attaching sharing and destructive update information to each arrow in a type, though often the information is inferred rather than being given explicitly in the source code. For example, implicit in the declaration for `assign` above is that `assign` called with a single argument of type `Ref t` creates a closure of type `t -> ()` containing that argument (and thus sharing the object of type `t`). The explicit sharing information describes applications of this closure to another argument. There is a single argument in this application, referred to with the formal parameter `v`. The other formal parameter, `p`, refers to the argument of the closure. In general, a type with  $N$  arrows in the “spine” has  $K + N$  formal parameters in the description of sharing, with the first  $K$  parameters being closure arguments.

The following code defines binary search trees of integers and defines a function which takes a pointer to a tree and inserts an integer into the tree. It uses destructive update, as would normally be done in an imperative language. The declarative alternative must reconstruct all nodes in the path from the root down to the new node. Experiments using our prototype implementation of Pawns indicate that for long paths this destructive update version is as fast as hand-written C code whereas the “pure” version is more than twenty times slower, primarily due to the overhead of memory allocation.

```

191 data Tree = Empty | Node Tree Int Tree

192 bst_insert_du :: Int -> Ref Tree -> ()
193     sharing bst_insert_du x !tp = _    -- tree gets updated
194     pre nosharing                      -- integers are atomic so
195     post nosharing                     -- it doesn't share
196 bst_insert_du x !tp =
197     case *tp of
198     Empty ->
199         *!tp := Node Empty x Empty    -- insert new node
200     (Node *lp n *rp) ->
201         if x <= n then
202             (bst_insert_du x !lp) !tp -- update lp (and tp)
203         else
204             (bst_insert_du x !rp) !tp -- update rp (and tp)

```

## 2.2 The high level view

Whenever destructive update is used in Pawns, programmers must be aware of potential sharing of data representations and take a low level view. In other cases it is desirable to have a high level view of values, ignoring how they are represented and any sharing which may be present. Pawns has a mechanism to indicate that such a high level view is taken. Pre- and post-conditions can specify sharing with a special pseudo-variable named **abstract**<sup>2</sup>. No variables which share with **abstract** can be destructively updated. Pawns type signatures which have no annotations concerning destructive update or sharing implicitly indicate no arguments are destructively updated and the arguments and result share with **abstract**. Thus a subset of Pawns code can look like and be considered as pure functional code.

The following code defines a function which takes a list of integers and returns a binary search tree containing the same integers. Though it uses destructive update internally, this impurity is encapsulated and it can therefore be viewed as a pure function. The list which is passed in as an argument is never updated. An initially empty tree is created locally. It is destructively updated by inserting each integer of the list into it (using `list_bst_du`, which calls `bst_insert_du`), then the tree is returned. Within the execution of `list_bst` it is important to understand the low level details of how the tree is represented, but this information is not needed outside the call. The sharing analysis of the Pawns compiler allows a distinction between “abstract” variables, which cannot be updated, and “concrete” variables which can be updated. Sharing of concrete variables must be considered and explicitly documented by the programmer.

```

data Ints = Nil | Cons Int Ints

list_bst :: Ints -> Tree -- pure function from Ints to Tree
-- implicit sharing information:
-- sharing list_bst xs = t
-- pre xs = abstract
-- post t = abstract
list_bst xs =
    *tp = Empty           -- create pointer to empty tree
    list_bst_du xs !tp    -- insert integers into tree, updating it
    *tp                  -- return tree

```

<sup>2</sup> There is conceptually a different **abstract** variable for each distinct type.

```

240 list_bst_du :: Ints -> Ref Tree -> ()
241     sharing list_bst_du xs !tp = _ -- tree gets updated
242     pre xs = abstract
243     post nosharing
244 list_bst xs =
245 list_bst_du xs !tp =
246     case xs of
247     (Cons x xs1) ->
248         bst_insert_du x !tp -- insert head of list into tree
249         list_bst_du xs1 !tp -- insert rest of list into tree
250 Nil -> ()

```

### 251 3 Core Pawns

252 An early pass of the Pawns compiler converts all function definitions into a  
 253 core language by flattening nested expressions, introducing extra variables et  
 254 cetera. A variable representing the return value of the function is introduced and  
 255 expressions are converted to bindings for variables. A representation of the core  
 256 language version of code is annotated with type, liveness and other information  
 257 prior to sharing analysis. We just describe the core language here. The right side  
 258 of each function definition is a statement (described using the definition of type  
 259 **Stat** below), which may contain variables, including function names (**Var**), data  
 260 constructors (**DCons**) and pairs containing a pattern (**Pat**) and statement for  
 261 case statements. All variables are distinct except for those in recursive instances  
 262 of **Stat** and variables are renamed to avoid any ambiguity due to scope.

```

263 data Stat =                -- Statement, eg
264     Seq Stat Stat |        -- stat1 ; stat2
265     EqVar Var Var |        -- v = v1
266     EqDeref Var Var |     -- v = *v1
267     DerefEq Var Var |     -- *v = v1
268     DC Var DCons [Var] |  -- v = Cons v1 v2
269     Case Var [(Pat, Stat)] | -- case v of pat1 -> stat1 ...
270     Error |               -- (for uncovered cases)
271     App Var Var [Var] |   -- v = f v1 v2
272     Assign Var Var |     -- *!v := v1
273     Instype Var Var       -- v = v1::instance_of_v1_type
274
275 data Pat =                -- patterns for case, eg
276     Pat DCons [Var]       -- (Cons *v1 *v2)

```

277 Patterns in the core language only bind references to arguments — the argu-  
 278 ments themselves must be obtained by explicit dereference operations. Pawns  
 279 supports “default” patterns but for simplicity of presentation here we assume all  
 280 patterns are covered in core Pawns and we include an error primitive. Similarly,  
 281 we just give the general case for application of a variable to  $N > 0$  arguments;



our implementation distinguishes some special cases. Memory is allocated for **DerefEq**, **DC** (for non-constants) and **App** (for unsaturated applications which result in a closure).

Sharing and type analysis cannot be entirely separated. Destructive update in the presence of polymorphic types can potentially violate type safety or “preservation”. For a variable whose type contains a type variable, we must avoid destructive update with a value with a less general type. For example, in `*x = []` the type of `x` is `Ref [t]`. If `*x` is assigned `[42]`, of type `[Int]`, passing it to a function which expects a `[Bool]` violates type safety. Pawns allows expressions to have their inferred types further instantiated using “`::`”, and the type checking pass of the compiler also inserts some type instantiation. The type checking pass ensures that direct update does not involve type instantiation but to improve flexibility, indirect update is checked during the sharing analysis.

## 4 The abstract domain

The representation of the value of a variable includes some set of main memory words (arguments of data constructors). Two variables share if the intersection of their sets of main memory words is not empty. The abstract domain for sharing analysis must maintain a conservative approximation to all sharing, so we can tell if two variables possibly share (or definitely do not share). The abstract domain we use is a set of pairs (representing possibly intersecting sets of main memory locations) of variable *components*. The different components of a variable partition the set of main memory words for the variable.

The components of a variable depend on its type. For non-recursive types other than arrays, each possible data constructor argument is represented separately. For example, the type `Maybe (Maybe (Either Int Int))` can have an argument of an outer `Just` data constructor, an inner `Just` and `Left` and `Right`. A component can be represented using a list of `x.y` pairs containing a data constructor and an argument number, giving the path from the outermost data constructor to the given argument. For example, the components of the type above can be written as: `[Just.1]`, `[Just.1,Just.1]`, `[Just.1,Just.1,Left.1]` and `[Just.1,Just.1,Right.1]`. If variable `v` has value `Just Nothing`, the expression `v.[Just.1]` represents the single main memory word containing the occurrence of `Nothing`.

For `Ref t` types we proceed as if there was a `Ref` data constructor, so `vp.[Ref.1]` represents the word `vp` points to. For function types, values may be closures. A closure which has had `K` arguments supplied is represented as a data constructor `ClK` with these `K` arguments; these behave in the same way as other data constructor arguments with respect to sharing. Closures also contain a code pointer and an integer which are not relevant to sharing so we ignore them here. We also ignore the subscript on the data constructor for sharing analysis because type and sharing analysis only give a lower bound on the number of closure arguments. Our analysis orders closure arguments so that the most recently supplied argument is first (the reverse of the more natural ordering).

For arrays, `[Array_.1]` is used to represent all words in the array. The expression, `x.[Array_.1,Just.1]` represents the arguments of all `Just` elements in an array `x` of `Maybe` values. For recursive types, paths are “folded” [5] so there are a finite number of components. If a type  $T$  has sub-component(s) of type  $T$  we use the empty path to denote the sub-component(s). In general, we construct a path from the top level and if we come across a sub-component of type  $T$  which is in the list of ancestor types (the top level type followed by the types of elements of the path constructed so far) we just use the path to the ancestor to represent the sub-component. Consider the following mutually recursive types:

```
data RTrees = Nil | Cons RTree RTrees
data RTree = RNode Int RTrees
```

For type `RTrees` we have the components `[]` (this folded path represents both `[Cons.2]` and `[Cons.1,RNode.2]`, since they are of type `RTrees`), `[Cons.1]` and `[Cons.1,RNode.1]`. The expression `t.[Cons.1,RNode.1]` represents the set of memory words which are the first argument of `RNode` in variable `t` of type `RTrees`. For type `RTree` we have the components `[]` (for `[RNode.2,Cons.1]`, of type `RTree`), `[RNode.1]` and `[RNode.2]` (which is also the folded version of `[RNode.2,Cons.2]`, of type `RTrees`). In our sharing analysis algorithm we use a function `fc` (fold component) which takes a  $v.c$  pair, and returns  $v.c'$  where  $c'$  is the correctly folded component for the type of variable  $v$ . For example, `fc (ts.[Cons.2]) = ts.[]`, assuming `ts` has type `RTrees`.

As well as containing pairs of components for distinct variables which may alias, the abstract domain contains “self-sharing” pairs for each possible component of a variable which may exist. Consider the following two bindings:

```
t = RNode 2 Nil
ts = Cons t Nil
```

With our domain, the most precise description of sharing after these two bindings is as follows. We represent a sharing pair as a set of two variable components. The first five are self-sharing pairs and the other two describe the sharing between `t` and `ts`.

```
{t.[RNode.1], t.[RNode.1]},
{t.[RNode.2], t.[RNode.2]},
{ts.[], ts.[]},
{ts.[Cons.1], ts.[Cons.1]},
{ts.[Cons.1,RNode.1], ts.[Cons.1,RNode.1]},
{t.[RNode.1], ts.[Cons.1,RNode.1]},
{t.[RNode.2], ts.[]}
```

Note there is no self-sharing pair for `t.[]` since there is no strict sub-part of `t` which is an `RTree`. Similarly, there is no sharing between `ts.[Cons.1]` and any part of `t`. Although the value `t` is used as the first argument of `Cons` in `ts`, this is not a main memory word which is used to represent the value of `t` (indeed, the value of `t` has no `Cons` cells). The tagged pointer value stored in variable `t` (which may be in a register) is copied into the cons cell.

## 5 The sharing analysis algorithm

We now describe the sharing analysis algorithm. Overall, the compiler attempts to find a proof that for a computation with a depth  $D$  of (possibly recursive) function calls, the following condition  $C$  holds, assuming  $C$  holds for all computations of depth less than  $D$ . This allows a proof by induction that  $C$  holds for all finite computations.

$C$ : For all functions  $f$ , if the precondition of  $f$  is satisfied whenever  $f$  is called, then

1. for all function calls and assignment statements in  $f$ , any live variable that may be updated at that point in an execution of  $f$  is annotated with “!”,
2. there is no update of live “abstract” variables when executing  $f$ ,
3. the union of the pre- and post-conditions of  $f$  is satisfied when  $f$  returns,
4. all parameters of  $f$  which may be updated when executing  $f$  are declared mutable in the type signature of  $f$ ,
5. for all function calls and assignment statements in  $f$ , any live variable that may be directly updated at that point is updated with a value of the same type or a more general type, and
6. for all function calls and assignment statements in  $f$ , any live variable that may be indirectly updated at that point does not share with any variable which has a less general type.

The algorithm is applied to each function definition in core Pawns to compute an approximation to the sharing before and after each statement (we call it the alias set). This can be used to check points 1–3 and 6 above and that preconditions of called functions are satisfied, so the induction hypothesis can be used. Point 4 is established using point 1 and a simple syntactic check that any parameter of  $f$  which is annotated “!” in the definition is declared mutable in the type signature (parameters are considered live throughout the definition). Point 5 relies on 4 and the type checking pass. The core of the algorithm is to compute the alias set after a statement, given the alias set before the statement. This is applied recursively for compound statements.

The alias set used at the start of the definition is the precondition of the function. This implicitly includes self-sharing pairs for all variable components of the arguments of the function and the pseudo-variables **abstract** <sub>$T$</sub>  for each type  $T$  used. Similarly, the postcondition implicitly includes self-sharing pairs for all components of the result (and the **abstract** <sub>$T$</sub>  variable if the result is abstract)<sup>3</sup>. As analysis proceeds, extra variables from the function body are added to the alias set and variables which are no longer live can be removed to improve efficiency. The alias set computed for the end of the definition, with sharing for local variables removed, must be a subset of the union of the pre- and post-condition of the function. We assume type information is given for all variables (a type checking/inference pass is completed before sharing analysis) and sharing

<sup>3</sup> Self-sharing for arguments and results is usually desired. For the rare cases it is not, we may provide a mechanism to override this default in the future.

information is given for all type instances of all (possibly polymorphic) defined functions. All type variables in type assignments in the function definition are replaced by `Ref ()`. This type has a single component which can be shared to represent possible sharing of arbitrary components of an arbitrary type. Finally, we assume there is no type which is an infinite chain of refs, for example, `type Refs = Ref Refs` (for which type folding results in an empty component rather than a `[Ref.1]` component; this is not a practical limitation).

Suppose  $a_0$  is the alias set just before statement  $s$ . The following algorithm computes  $\text{alias}(s, a_0)$ , the alias set just after statement  $s$ . The algorithm structure follows the recursive definition of statements and we describe it using pseudo-Haskell, interspersed with brief discussion. At some points we use high level declarative set comprehensions to describe what is computed and naive implementation may not lead to the best performance.

```
alias (Seq stat1 stat2) a0 =          -- stat1; stat2
    alias stat2 (alias stat1 a0)
alias (EqVar v1 v2) a0 =              -- v1 = v2
    let
        self1 = { {v1.c, v1.c} | {v2.c, v2.c} ∈ a0 }
        share1 = { {v1.c1, v.c2} | {v2.c1, v.c2} ∈ a0 }
    in
        a0 ∪ self1 ∪ share1
alias (DerefEq v1 v2) a0 =            -- *v1 = v2
    let
        self1 = { {v1.[Ref.1], v1.[Ref.1]} } ∪
            { {fc(v1.(Ref.1 :c)), fc(v1.(Ref.1 :c))} | {v2.c, v2.c} ∈ a0 }
        share1 = { {fc(v1.(Ref.1 :c1)), v.c2} | {v2.c1, v.c2} ∈ a0 }
    in
        a0 ∪ self1 ∪ share1
```

Sequencing is handled by function composition. To bind a fresh variable  $v_1$  to a variable  $v_2$  the self-sharing of  $v_2$  is duplicated for  $v_1$  and the sharing for each component of  $v_2$  is duplicated for  $v_1$ . Binding  $*v_1$  to  $v_2$  is done in a similar way, but the components of  $v_1$  must have `Ref.1` prepended to them and the result folded, and the `[Ref.1]` component of  $v_1$  self-shares.

```
alias (Assign v1 v2) a0 =              -- *v1 := v2
    let
        self1 = { {v1.[Ref.1], v1.[Ref.1]} } ∪
            { {fc(v1.(Ref.1 :c)), fc(v1.(Ref.1 :c))} | {v2.c, v2.c} ∈ a0 }
        share1 = { {fc(v1.(Ref.1 :c1)), v.c2} | {v2.c1, v.c2} ∈ a0 }
        -- al = possible aliases for v1.[Ref.1]
        al = { v_a.c_a | {v1.[Ref.1], v_a.c_a} ∈ a0 }
        -- (live variables in al+v1 must be annotated with !
        -- and must not share with abstract)
```

```

selfal = {{fc(va.(ca++c)),fc(va.(ca++c))} |
          va.ca ∈ a1 ∧ {v2.c,v2.c} ∈ a0}
shareal = {{fc(va.(ca++c1)),v.c2} |
           va.ca ∈ a1 ∧ {v2.c1,v.c2} ∈ a0} ∪
           {{fc(va.(ca++c)),fc(v1.(Ref.1:c))} |
            va.ca ∈ a1 ∧ {v2.c,v2.c} ∈ a0}
-- old1 = old aliases for v1, which can be removed
-- if the assignment doesn't create a cyclic structure
old1 = {{v1.c1,v.c2} | {v1.c1,v.c2} ∈ a0}
in if ∃c {v1.[Ref.1],v2.c} ∈ a0 then
    a0 ∪ self1 ∪ share1 ∪ selfal ∪ shareal
else
    (a0 \ old1) ∪ self1 ∪ share1 ∪ selfal ∪ shareal

```

427 Assignment to an existing variable **\*v1** adds the same sharing as for binding  
 428 a fresh variable, but there are two extra complications. First, **\*v1** may be an alias  
 429 for components of other variables (the live subset of these variables and **v1** must  
 430 be annotated with “!” on the assignment statement; checking such annotations is  
 431 a primary purpose of the sharing analysis). All these variable components must  
 432 have the same sharing added as **\*v1**. The components must be concatenated  
 433 and folded appropriately. Second, if the assignment does not create a cyclic  
 434 structure the existing sharing for **v1** can safely be removed, improving precision.  
 435 It is sufficient to check if any component of **v2** aliases with **v1**. [Ref.1].

```

alias (DC v dc [v1,...vN]) a0 = -- v = Dc v1...vN
let
    self1 = ∪1≤i≤N ({fc(v.[dc.i]),fc(v.[dc.i])} ∪
                    {{fc(v.(dc.i:c)),fc(v.(dc.i:c))} | {vi.c,vi.c} ∈ a0})
    share1 = ∪1≤i≤N {{fc(v.(dc.i:c1)),w.c2} | {vi.c1,w.c2} ∈ a0}
in
    a0 ∪ self1 ∪ share1

```

436 The DerefEq case can be seen as equivalent to **v1 = Ref v2** and binding a  
 437 variable to a data constructor with *N* variable arguments is a generalisation.

```

alias (EqDeref v1 v2) a0 = -- v1 = *v2
let
    self1 = {{v1.c,v1.c} | {v2.(Ref.1:c),v2.(Ref.1:c)} ∈ a0}
    share1 = {{v1.c1,v.c2} | {v2.(Ref.1:c1),v.c2} ∈ a0}
    empty1 = {{v1.[],v.c} | {v1.[],v.c} ∈ (self1 ∪ share1)}
in
    if the type of v1 has a [] component then
        a0 ∪ self1 ∪ share1
    else --- avoid bogus sharing with empty component
        (a0 ∪ self1 ∪ share1) \ empty1

```

438 The `EqDeref` case is similar to the inverse of `DerefEq` in that we are removing  
 439 `Ref.1` rather than prepending it. However, if the empty component results we  
 440 must check that such a component exists for the type of `v1`.

```
alias (App v f [v1,...vN]) a0 =          -- v = f v1...vN
  let
    "f(w1,...wK+N) = r" is used to declare sharing for f
    post = the postcondition of f along with the sharing for
            mutable arguments from the precondition,
            with parameters and result renamed with
            f.[Cl.K],...f.[Cl.1],v1,...vN and v, respectively
    postt = { {x1.c1,x3.c3} | {x1.c1,x2.c2} ∈ post ∧ {x2.c2,x3.c3} ∈ a0 }
    -- (the renamed precondition of f must be a subset of a0,
    -- and mutable arguments of f and live variables they share
    -- with must be annotated with ! and must not share with
    -- abstract)
    -- selfc+sharec not needed for saturated applications
    selfc = { {v.[Cl.i],v.[Cl.i]} | 1 ≤ i ≤ N } ∪
            { {v.((Cl.i):c),v.((Cl.i):c)} |
              1 ≤ i ≤ N ∧ {vi.c,vi.c} ∈ a0 } ∪
            { {v.(Cl.(i+N)):c,v.(Cl.(i+N)):c} |
              {f.(Cl.i):c,f.(Cl.i):c} ∈ a0 }
    sharec = { {v.(Cl.i):c1,x.c2} |
              1 ≤ i ≤ N ∧ {vi.c1,x.c2} ∈ a0 } ∪
            { {v.(Cl.(i+N)):c1,x.c2} |
              {f.(Cl.i):c1,x.c2} ∈ a0 }
  in
    a0 ∪ postt ∪ selfc ∪ sharec
```

441 Function application relies on the sharing information attached to all arrow  
 442 types. Because Pawns uses the syntax of statements to express pre- and post-  
 443 conditions, our the implementation uses the sharing analysis algorithm to derive  
 444 an explicit alias set representation (currently this is done recursively, with the  
 445 level of recursion limited by the fact than pre- and post-conditions must not  
 446 contain function calls). Here we ignore the details of how the alias set represen-  
 447 tation is obtained. The compiler also uses the sharing information to check that  
 448 preconditions are satisfied, all required "!" annotations are present and abstract  
 449 variables are not modified.

450 The main thing done for function application is to add the declared post-  
 451 condition of the function, renamed appropriately. The  $N$  arguments of the call  
 452 replace the last  $N$  formal parameters and `v` replaces the formal result. The first  
 453  $K$  formal parameters represent closure arguments of `f`, so those variables are  
 454 replaced with `f` and the components are prefixed with the representation for a

closure argument. As well as the declared postcondition, sharing for the mutable arguments of the precondition must be included. The analysis of a function definition guarantees that the union of pre- and post-conditions are satisfied when the function returns (assuming the precondition is satisfied initially), but execution cannot add sharing between non-mutable arguments so it is not added here. Thus by not including preconditions in the declared postconditions, precision is improved. It is also necessary to include one step of transitivity in the sharing information: if the renamed postcondition introduces sharing between variable components  $x_1.c_1$  and  $x_2.c_2$  and before the function call  $x_2.c_2$  shared with  $x_3.c_3$  we add sharing between  $x_1.c_1$  and  $x_3.c_3$ .

For some calls we can know statically that a closure cannot result, but in general we must assume that a closure is created and the first  $N$  closure arguments share with the  $N$  arguments of the function call and any closure arguments of  $\mathbf{f}$  share with additional closure arguments of the result (this requires renumbering of these arguments).

```

alias Error a0 =  $\emptyset$  -- error
alias (Case v [(p1, s1), ... (pN, sN)]) a0 = -- case v of ...
  let
    old = { {v.c1, v2.c2} | {v.c1, v2.c2} ∈ a0 }
  in
     $\bigcup_{1 \leq i \leq N}$  aliasCase a0 old v pi si

aliasCase a0 av v (Pat dc [v1, ... vN]) s = -- (Dc *v1...*vN) -> s
  let
    avdc = { {fc(v.(dc.i : c1)), w.c2} | {fc(v.(dc.i : c1)), w.c2} ∈ av }
    rself = { {vi.Ref.1, vi.Ref.1} | 1 ≤ i ≤ N }
    vishare = { {fc(vi.Ref.1 : c1), fc(vj.Ref.1 : c2))} |
               {fc(v.(dc.i : c1), fc(v.(dc.j : c2))} ∈ av }
    share = { {fc(vi.Ref.1 : c1), w.c2} | {fc(v.(dc.i : c1), w.c2))} ∈ av }
  in
    alias s (rself ∪ vishare ∪ share ∪ (a0 \ av) ∪ avdc)

```

For a case expression we return the union of the alias sets obtained for each of the different branches. For each branch we only keep sharing information for the variable we are switching on which is compatible with the data constructor in that branch (we remove all the old sharing,  $\mathbf{av}$ , and add the compatible sharing,  $\mathbf{avdc}$ ). Note we use a high level declarative definition for  $\mathbf{avdc}$  (and other variables) which implicitly uses the inverse of  $\mathbf{fc}$ . To deal with individual data constructors we consider pairs of components of arguments  $i$  and  $j$  which may alias in order to compute possible sharing between  $v_i$  and  $v_j$ , including self-sharing when  $i = j$ . The corresponding component of  $v_i$  (prepended with **Ref** and folded) may alias the component of  $v_j$ . For example, if  $\mathbf{v}$  of type **RTrees** is matched with  $\mathbf{Cons} \ *v_1 \ *v_2$  and  $\mathbf{v} . []$  self-shares, we need to find the components which fold

481 to  $v.[]$  ( $v.[\text{Cons}.2]$  and  $v.[\text{Cons}.1, \text{RNode}.2]$ ) in order to compute the sharing  
 482 for  $v2$  and  $v1$ . Thus we compute that  $fc(v2.[\text{Ref}.1, \text{Cons}.2]) = v2.[\text{Ref}.1]$   
 483 may alias  $v1.[\text{Ref}.1, \text{Cons}.1, \text{RNode}.2]$ , which can occur if the data structure is  
 484 cyclic. The DC case cannot introduce cycles as the variable on the left is distinct  
 485 from the variables in the right but **Assign** can introduce cycles.

```
alias (Instype v1 v2) a0 =          -- v1 = v2::t
    alias (EqVar v1 v2) a0
    -- (if any sharing is introduced between v1 and v2,
    -- v2 must not be indirectly updated later while live)
```

486 Type instantiation is dealt with in the same way as variable equality, with  
 487 the additional check that if any sharing is introduced, the variable with the more  
 488 general type is not implicitly updated later while still live (it is sufficient to check  
 489 there is no “!v2” annotation attached to a later statement).

## 490 6 Example

491 We now show how this sharing analysis algorithm is applied to the binary search  
 492 tree code given earlier. We give a core Pawns version of each function and the  
 493 alias set before and after each statement, plus an additional set at the end which  
 494 is the union of the pre- and post-conditions of the function. To save space,  
 495 we write the alias set as a set of sets where each inner set represents all sets  
 496 containing exactly two of its members. Thus  $\{\{a, b, c\}\}$  represents a set of six  
 497 sharing pairs: sharing between all pairs of elements, including self-sharing. The  
 498 return value is given by variable **ret** and variables **absL** and **absT** are the versions  
 499 of **abstract** for type **Ints** and **Tree**, respectively.

```
500 list_bst xs =          -- 0
    v1 = Empty          -- 1
    *tp = v1            -- 2
    list_bst_du xs !tp   -- 3
    ret = *tp            -- 4
```

505 We start with the precondition:  $a_0 = \{\{xs.[\text{Cons}.1], \text{absL}.[\text{Cons}.1]\},$   
 506  $\{xs.[], \text{absL}.[]\}\}$ . Binding to a constant introduces no sharing so  $a_1 = a_0$ .  
 507  $a_2 = a_1 \cup \{\{tp.[\text{Ref}.1]\}\}$ . The function call has precondition  $a_0 \cup \{\{tp.[\text{Ref}.1]\},$   
 508  $\{tp.[\text{Ref}.1, \text{Node}.2]\}\}$ , which is a superset of  $a_2$ . Since **tp** is a mutable argu-  
 509 ment the precondition sharing for **tp** is added:  $a_3 = a_2 \cup \{\{tp.[\text{Ref}.1,$   
 510  $\text{Node}.2]\}\}$ . The final sharing includes the return variable, **ret**:  $a_4 = a_3 \cup$   
 511  $\{\{ret.[], tp.[\text{Ref}.1]\}, \{ret.[\text{Node}.2], tp.[\text{Ref}.1, \text{Node}.2]\}\}$ . After remov-  
 512 ing sharing for the dead (local) variable **tp** we obtain a subset of the union of  
 513 the pre- and post-conditions, which is  $a_0 \cup \{\{ret.[], \text{absT}.[]\}, \{ret.[\text{Node}.2],$   
 514  $\text{absT}.[\text{Node}.2]\}\}$ .



```

515 list_bst_du xs !tp =                -- 0
516     case xs of
517         (Cons *v1 *v2) ->          -- 1
518             x = *v1                -- 2
519             xs1 = *v2               -- 3
520             v3 = bst_insert_du x !tp -- 4
521             v4 = list_bst_du xs1 !tp -- 5
522             ret = v4                -- 6
523     Nil ->                          -- 7
524         ret = ()                   -- 8
525     -- after case                  -- 9

```

526 We start with the precondition,  $a_0 = \{\{tp.[Ref.1]\}, \{tp.[Ref.1, Node.2]\},$   
 527  $\{xs.[Cons.1], absL.[Cons.1]\}, \{xs.[], absL.[]\}\}$ . The Cons branch of the  
 528 case introduces sharing for v1 and v2:  $a_1 = a_0 \cup \{\{xs.[Cons.1], absL.[Cons.1],$   
 529  $v1.[Ref.1], v2.[Ref.1, Cons.1]\}, \{v2.[Ref.1], xs.[], absL.[]\}\}$ . The list  
 530 elements are atomic so  $a_2 = a_1$ . The next binding makes the sharing of xs1 and  
 531 xs the same:  $a_3 = a_2 \cup \{\{v2.[Ref.1], xs.[], xs1.[], absL.[]\}, \{v1.[Ref.1],$   
 532  $xs.[Cons.1], xs1.[Cons.1], absL.[Cons.1], v2.[Ref.1, Cons.1]\}\}$ . This can  
 533 be simplified by removing the dead variables v1 and v2. The precondition of the  
 534 calls are satisfied and  $a_6 = a_5 = a_4 = a_3$ . For the Nil branch we remove the in-  
 535 compatible sharing for xs from  $a_0$ :  $a_7 = \{\{tp.[Ref.1]\}, \{tp.[Ref.1, Node.2]\},$   
 536  $\{absL.[Cons.1]\}, \{absL.[]\}\}$  and  $a_8 = a_7$ . Finally,  $a_9 = a_6 \cup a_8$ . Ignoring local  
 537 variables, this is a subset of the union of the pre- and post-conditions,  $a_0$ .

```

538 bst_insert_du x !tp =                -- 0
539     v1 = *tp                         -- 1
540     case v1 of
541         Empty ->                     -- 2
542             v2 = Empty               -- 3
543             v3 = Empty               -- 4
544             v4 = Node v2 x v3        -- 5
545             *!tp := v4               -- 6
546             ret = ()                 -- 7
547         (Node *lp *v5 *rp) ->        -- 8
548             n = *v5                  -- 9
549             v6 = (x <= n)             -- 10
550             case v6 of
551                 True ->              -- 11
552                     v7 = (bst_insert_du x !lp) !tp -- 12
553                     ret = v7         -- 13
554                 False ->            -- 14
555                     v8 = (bst_insert_du x !rp) !tp -- 15
556                     ret = v8         -- 16
557             -- end case              -- 17
558     -- end case                      -- 18

```

Here  $a_0 = \{\{\text{tp. [Ref. 1]}\}, \{\text{tp. [Ref. 1, Node. 2]}\}\}$  and  $a_1 = a_0 \cup \{\{\text{v1. []}, \text{tp. [Ref. 1]}\}, \{\text{tp. [Ref. 1, Node. 2]}, \text{v1. [Node. 2]}\}\}$ . For the **Empty** branch we remove the **v1** sharing so  $a_4 = a_3 = a_2 = a_0$  and  $a_5 = a_4 \cup \{\{\text{v4. []}\}, \{\text{v4. [Node. 2]}\}\}$ . After the destructive update,  $a_6 = a_5 \cup \{\{\text{v4. []}, \text{tp. [Ref. 1]}\}, \{\text{v4. [Node. 2]}, \text{tp. [Ref. 1, Node. 2]}\}\}$  (**v4** is dead and can be removed) and  $a_7 = a_6$ . For the **Node** branch we have  $a_8 = a_1 \cup \{\{\text{v1. []}, \text{tp. [Ref. 1]}, \text{lp. [Ref. 1]}, \text{rp. [Ref. 1]}\}, \{\text{tp. [Ref. 1, Node. 2]}, \text{lp. [Ref. 1, Node. 2]}, \text{rp. [Ref. 1, Node. 2]}, \text{v5. [Ref. 1]}, \text{v1. [Node. 2]}\}\}$ . The same set is retained for  $a_9 \dots a_{17}$  (assuming the dead variable **v5** is retained), the preconditions of the function calls are satisfied and the required annotations are present. Finally,  $a_{18} = a_{17} \cup a_7$  and after eliminating local variables we get the postcondition, which is the same as the precondition.

## 7 Discussion

It is inevitable we lose some precision with recursion in types. However, it seems that some loss of precision could be avoided relatively easily. The use of the empty path to represent sub-components of recursive types results in imprecision when references are created. For example, the analysis of `*vp = Nil; v = *vp` concludes that the empty component of **v** may share with itself and the **Ref** component of **vp** (in reality, **v** has no sharing). Instead of the empty path, a dummy path of length one could be used. A more aggressive approach would be to unfold the recursion an extra level, at least for some types. This could allow us to express (non-)sharing of separate subtrees and whether data structures are cyclic, at the cost of more variable components and more complex pre- and postconditions.

Increasing the number of variable components also affects efficiency. The algorithmic complexity is affected by the representation of alias sets. Currently we use a naive implementation, using just ordered pairs of variable components as the set elements and a set library which uses an ordered binary tree. The size of the set can be  $O(N^2)$ , where  $N$  is the maximum number of live variable components of the same type at any program point (each such variable component can share with all the others). In typical code the number of live variables at any point is not particularly large. If the size of alias sets does become problematic, a more refined set representation could be used, such as the set of sets of pairs representation we used in Section 6, where sets of components which all share with each other are optimised. We have not stress tested our implementation as it is intended to be a prototype, but performance has not been concerning at this stage.

## 8 Related work

Related programming languages are discussed in [1]; here we restrict attention to work related to the sharing analysis algorithm. The most closely related work is that done in the compiler for Mars [6], which extends similar work done for

Mercury [7] and earlier for Prolog [8]. All use a similar abstract domain based on the type folding method first proposed in [5]. Our abstract domain is somewhat more precise due to inclusion of self-aliasing, and we have no sharing for constants. In Mars it is assumed that constants other than numbers can share. Thus for code such as `xs = []; ys = xs` our analysis concludes there is no sharing between `xs` and `ys` whereas the Mars analysis concludes there may be sharing.

One important distinction is that in Pawns sharing (and mutability) is declared in type signatures of functions so the Pawns compiler just has to check the declarations are consistent, rather than infer all sharing from the code. However, it does have the added complication of destructive update. As well as having to deal with the assignment primitive, it complicates handling of function calls and case statements (the latter due to the potential for cyclic structures). Mars, Mercury and Prolog are essentially declarative languages. Although Mars has assignment statements the semantics is that values are copied rather than destructively updated — the variable being assigned is modified but other variables remain unchanged. Sharing analysis is used in these languages to make the implementation more efficient. For example, the Mars compiler can often emit code to destructively update rather than copy a data structure because sharing analysis reveals no other live variables share it. In Mercury and Prolog the analysis can reveal when heap-allocated data is no longer used, so the code can reuse or reclaim it directly instead of invoking a garbage collector.

These sharing inference systems use an explicit graph representation of the sharing behaviour of each segment of code. For example, code  $s_1$  may cause sharing between (a component of) variables **a** and **b** (which is represented as an edge between nodes **a** and **b**) and between **c** and **d** and code  $s_2$  may cause sharing between **b** and **c** and between **d** and **e**. To compute the sharing for the sequence  $s_1; s_2$  they use the “alternating closure” of the sharing for  $s_1$  and  $s_2$ , which constructs paths with edges alternating from  $s_1$  and  $s_2$ , for example **a-b** (from  $s_1$ ), **b-c** (from  $s_2$ ), **c-d** (from  $s_1$ ) and **d-e** (from  $s_2$ ).

The sharing behaviour of functions in Pawns is represented explicitly, by a pre- and postcondition and set of mutable arguments but there is no explicit representation for sharing of statements. The (curried) function `alias s` represents the sharing behaviour of `s` and the sharing behaviour of a sequence of statements is represented by the composition of functions. This representation has the advantage that the function can easily use information about the current sharing, including self-sharing, and remove some if appropriate. For example, in the `[]` branch of the case in the code below the sharing for `xs` is removed and we can conclude the returned value does not share with the argument.

```
map_const_1 :: [t] -> [Int]
  sharing map_const_1 xs = ys pre nosharing post nosharing
map_const_1 xs =
  case xs of
    [] -> xs      -- can look like result shares with xs
    (_:xs1) -> 1:(map_const_1 xs1)
```

Given the extra precision that can be achieved, it may be worth attempting to adapt our approach to inferring alias information in languages without destructive update. There are other approaches to and uses of alias analysis for imperative languages, such as [9] and [10], but these are not aimed at precisely capturing information about dynamically allocated data. A more detailed discussion of such approaches is given in [6].

## 9 Conclusion

Purely declarative languages have the advantage of avoiding side effects, such as destructive update of function arguments. This makes it easier to combine program components, but some algorithms are hard to code efficiently without flexible use of destructive update. A function can behave in a purely declarative way if destructive update is allowed, but restricted to data structures which are created inside the function. The Pawns language uses this idea to support flexible destructive update encapsulated in a declarative interface. It is designed to make all side effects “obvious” from the source code. Because there can be sharing between the representations of different arguments of a function, local variables and the value returned, sharing analysis is an essential component of the compiler. It is also used to ensure “preservation” of types in computations. Sharing analysis has been used in other languages to improve efficiency and to give some feedback to programmers but we use it to support important features of the programming language.

The algorithm operates on (heap allocated) algebraic data types, including arrays and closures. In common with other sharing analysis used in declarative languages it supports binding of variables, construction and deconstruction (combined with selection or “case”) and function/procedure calls. In addition, it supports explicit pointers, destructive update via pointers, creation and application of closures and pre- and post-conditions concerning sharing attached to type signatures of functions. It also uses an abstract domain with additional features to improve precision.

## References

1. Naish, L.: An informal introduction to Pawns: a declarative/imperative language. <http://people.eng.unimelb.edu.au/lee/papers/pawns> (2015) Accessed 2015-03-16.
2. Milner, R., Tofte, M., Macqueen, D.: The Definition of Standard ML. MIT Press, Cambridge, MA, USA (1997)
3. Jones, S.P., Hughes, J., Augustsson, L., Barton, D., Boutel, B., Burton, W., Fasel, J., Hammond, K., Hinze12, R., Hudak, P., et al.: Report on the programming language Haskell 98, a non-strict purely functional language, February 1999. (1999)
4. Lippmeier, B.: Type Inference and Optimisation for an Impure World. PhD thesis, Australian National University (June 2009)
5. Bruynooghe, M.: Compile time garbage collection or how to transform programs in an assignment free languages into code with assignments. IFIP TC2/WG2.1 Working Conference on Program Specification and Transformation, Bad-Tölz, Germany, April 15-17, 1986 (1986)

- 687 6. Giuca, M.: Mars: An imperative/declarative higher-order programming language  
688 with automatic destructive update. PhD thesis, University of Melbourne (2014)
- 689 7. Mazur, N., Ross, P., Janssens, G., Bruynooghe, M.: Practical aspects for a working  
690 compile time garbage collection system for Mercury. In Codognet, P., ed.: Proceed-  
691 ings of ICLP 2001, Lecture Notes in Computer Science. Volume 2237., Springer  
692 (2001) 105–119
- 693 8. Mulkers, A.: Live data structures in logic programs, derivation by means of abstract  
694 interpretation. Springer-Verlag (1993)
- 695 9. Landi, W., Ryder, B.G.: A safe approximate algorithm for interprocedural aliasing.  
696 SIGPLAN Not. **27**(7) (July 1992) 235–248
- 697 10. Emami, M., Ghiya, R., Hendren, L.J.: Context-sensitive interprocedural points-to  
698 analysis in the presence of function pointers. In: Proceedings of the ACM SIGPLAN  
699 1994 Conference on Programming Language Design and Implementation. PLDI  
700 '94, New York, NY, USA, ACM (1994) 242–256