

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

[View the peer-reviewed version](https://peerj.com/articles/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 that 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 that 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 post-conditions 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 that 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 that are only visible to the low level code (and are therefore safe to update) from data structures that are passed in from the high level code (for which update would violate purity). The main aim of Pawns is to get the

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

44 There are other functional programming languages, such as ML [2], Haskell
45 [3] and Disciple [4], that allow destructive update of shared data structures but
46 do not allow this impurity to be encapsulated. In these languages the ability
47 to update the data structure is connected to its type¹. For a data structure to
48 be built using destructive update its type must allow destructive update and
49 any code that uses the data structure can potentially update it as well. This
50 prevents simple declarative analysis of the code and can lead to a proliferation
51 of different versions of a data structure, with different parts being mutable. For
52 example, there are four different versions of lists, since both the list elements
53 and the “spine” may (or may not) be mutable, and sixteen different versions
54 of lists of pairs. There is often an efficiency penalty as well, with destructive
55 update requiring an extra level of indirection in the data structure (an explicit
56 “reference” in the type with most versions of ML and Haskell). Pawns avoids
57 this inefficiency and separates mutability from type information, allowing a data
58 structure to be mutable in some contexts and considered “pure” in others. The
59 main cost from the programmer perspective is the need to include extra annota-
60 tions and information in the source code. This can also be considered a benefit,
61 as they provide useful documentation and error checking. The main implemen-
62 tation cost is additional analysis done by the compiler, which is the focus of this
63 paper.

64 The rest of this paper assumes some familiarity with Haskell and is structured
65 as follows. Section 2 gives a brief overview of the relevant features of Pawns.
66 An early pass of the compiler translates Pawns programs into a simpler “core”
67 language; this is described in Section 3. Section 4 describes the abstract domain
68 used for sharing analysis algorithm, Section 5 defines the algorithm itself and
69 Section 6 gives an extended example. Section 7 briefly discusses precision and
70 efficiency issues. Section 8 discusses related work and Section 9 concludes.

71 2 An overview of Pawns

72 A more detailed introduction to Pawns is given in [1]. Pawns has many simi-
73 larities with other functional languages. It supports algebraic data types with
74 parametric polymorphism, higher order programming and curried function defi-
75 nitions. It uses strict evaluation. In addition, it supports destructive update via
76 “references” (pointers) and has a variety of extra annotations to make impure
77 effects more clear from the source code and allow them to be encapsulated in
78 pure code. Pawns also supports a form of global variables (called state variables)
79 which support encapsulated effects, but we do not discuss them further here as
80 they are handled in essentially the same way as other variables in sharing analy-
81 sis. Pure code can be thought of in a declarative way, were values can be viewed

¹ Disciple uses “region” information to augment types, with similar consequences.

82 abstractly, without considering how they are represented. Code that uses de-
 83 structive update must be viewed at a lower level, considering the representation
 84 of values, including sharing. We discuss this lower level view first, then briefly
 85 present how impurity can be encapsulated to support the high level view. We
 86 use Haskell-like syntax for familiarity.

87 2.1 The low level view

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

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

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


```

125 data Colour = Red | Green | Blue
126 data Colours = Nil | Cons Colour Colours -- like List Colour

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

```

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



137
 138 The destructive update above changes the values of both `headp` and `cols`
 139 (the representations are shared). One of the novel features of Pawns is that the
 140 source code must be annotated with “!” to make it obvious when each “live”
 141 variable is updated. If both `headp` and `cols` are used later, the assignment
 142 statement above must be written as follows, with `headp` prefixed with “!” and
 143 an additional annotation attached to the whole statement indicating `cols` may
 144 be updated:

```

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

```

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

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

```

161 assign :: Ref t -> t -> ()
162     sharing assign !p v = _    -- p may be updated
163     pre nosharing              -- p&v don't share when called
164     post *p = v               -- assign may make *p alias with v
165 assign !p v =
166     *!p := v

```

167 The “!” annotation on parameter `p` declares the first argument of `assign`
168 is mutable. The default is that arguments are not mutable. As well as check-
169 ing for annotations on assignments and function calls, sharing analysis is used
170 to check that all parameters which may be updated are declared mutable in
171 type signatures, and pre- and post-conditions are always satisfied. For example,
172 assuming the previous code which binds `cols`, the call `assign !tailp !cols`
173 annotates all modified variables but violates the precondition of `assign` because
174 there is sharing between `tailp` and `cols` at the time of the call. Violating this
175 precondition allows cyclic structures to be created, which is important for un-
176 derstanding the code. If the precondition was dropped, the second argument of
177 `assign` would also need to be declared mutable in the type signature and the
178 assignment to `p` would require `v` to be annotated. In general, there is an inter-
179 dependence between “!” annotations in the code and pre- and post-conditions.
180 More possible sharing at a call means more “!” annotations are needed, more
181 sharing in (recursive) calls and more sharing when the function returns.

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

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


```

201 data Tree = TNil | Node Tree Int Tree
202 bst_insert_du :: Int -> Ref Tree -> ()
203     sharing bst_insert_du x !tp = _ -- tree gets updated
204     pre nosharing -- integers are atomic so
205     post nosharing -- it doesn't share
206 bst_insert_du x !tp =
207     case *tp of
208     TNil ->
209         *!tp := Node TNil x TNil -- insert new node
210     (Node *lp n *rp) ->
211         if x <= n then
212             (bst_insert_du x !lp) !tp -- update lp (and tp)
213         else
214             (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 that may be present. For example, in the two trees `t1` and `t2` depicted below, it is much simpler if we do not have to care or know about the sharing between the trees and within tree `t1`. The high level view is they are both just `Node (Node TNil 123 TNil) 123 (Node TNil 123 TNil)`.



Pawns has a mechanism to indicate that the high level view is taken. Pre- and post-conditions can specify sharing with a special pseudo-variable named `abstract`². The sharing analysis of the Pawns compiler allows a distinction between “abstract” variables, which share with `abstract` and for which the programmer takes a high level view, and “concrete” variables for which the programmer must understand the representation and explicitly declare all sharing in pre- and post-conditions. The analysis checks that no live abstract variables can be destructively updated. Thus if a function has a parameter which is updated, it must be declared mutable and must not be declared to share with `abstract` in the precondition (non-mutable parameters may or may not share with `abstract`). Checking of preconditions ensures that abstract variables are not passed to functions which expect concrete data structures. For example, an abstract tree cannot be passed to `bst_insert_du` because the precondition allows no sharing with `abstract`. It is important that the tree structure is known

² There is conceptually a different `abstract` variable for each distinct type.

238 when `bst_insert_du` is used because the result depends on it. For example,
 239 inserting into the right subtree of `t2` only affects this subtree whereas inserting
 240 into the right subtree of `t1` (which has the same high level value) also changes
 241 the left subtree of both `t1` and `t2`. Note that concrete variables can be passed
 242 to functions which allow abstract arguments. Pawns type signatures that have
 243 no annotations concerning destructive update or sharing implicitly indicate no
 244 arguments are destructively updated and the arguments and result share with
 245 `abstract`. Thus a subset of Pawns code can look like and be considered as pure
 246 functional code.

247 The following code defines a function which takes a list of integers and returns
 248 a binary search tree containing the same integers. Though it uses destructive up-
 249 date internally, this impurity is encapsulated and it can therefore be viewed as
 250 a pure function. The list that is passed in as an argument is never updated and
 251 the tree returned is abstract so it is never subsequently updated (a concrete tree
 252 could be returned if an explicit postcondition with `nosharing` was given). An
 253 initially empty tree is created locally. It is destructively updated by inserting
 254 each integer of the list into it (using `list_bst_du`, which calls `bst_insert_du`),
 255 then the tree is returned. Within the execution of `list_bst` it is important to
 256 understand the low level details of how the tree is represented, but this informa-
 257 tion is not needed outside the call.

```

258 data Ints = Nil | Cons Int Ints
259
260 list_bst :: Ints -> Tree -- pure function from Ints to Tree
261 -- implicit sharing information:
262 -- sharing list_bst xs = t
263 -- pre xs = abstract
264 -- post t = abstract
265 list_bst xs =
266     *tp = TNil           -- create pointer to empty tree
267     list_bst_du xs !tp  -- insert integers into tree
268     *tp                 -- return (updated) tree

269 list_bst_du :: Ints -> Ref Tree -> ()
270     sharing list_bst_du xs !tp = _ -- tree gets updated
271     pre xs = abstract
272     post nosharing
273 list_bst_du xs !tp =
274     case xs of
275     (Cons x xs1) ->
276         bst_insert_du x !tp -- insert head of list into tree
277         list_bst_du xs1 !tp -- insert rest of list into tree
278     Nil -> ()

```


279 3 Core Pawns

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

```

291 data Stat =                -- Statement, eg
292     Seq Stat Stat |        -- stat1 ; stat2
293     EqVar Var Var |        -- v = v1
294     EqDeref Var Var |     -- v = *v1
295     DerefEq Var Var |     -- *v = v1
296     DC Var DCons [Var] |  -- v = Cons v1 v2
297     Case Var [(Pat, Stat)] | -- case v of pat1 -> stat1 ...
298     Error |                -- (for uncovered cases)
299     App Var Var [Var] |   -- v = f v1 v2
300     Assign Var Var |     -- *!v := v1
301     Instype Var Var      -- v = v1::instance_of_v1_type
302
303 data Pat =                -- patterns for case, eg
304     Pat DCons [Var]      -- (Cons *v1 *v2)
  
```

305 Patterns in the core language only bind references to arguments — the argu-
 306 ments themselves must be obtained by explicit dereference operations. Pawns
 307 supports “default” patterns but for simplicity of presentation here we assume all
 308 patterns are covered in core Pawns and we include an error primitive. Similarly,
 309 we just give the general case for application of a variable to $N > 0$ arguments;
 310 our implementation distinguishes some special cases. Memory is allocated for
 311 **DerefEq**, **DC** (for non-constants) and **App** (for unsaturated applications which
 312 result in a closure).

313 The runtime behaviour of **Instype** is identical to **EqVar** but it is treated dif-
 314 ferently in type analysis. Sharing and type analysis cannot be entirely separated.
 315 Destructive update in the presence of polymorphic types can potentially violate
 316 type safety or “preservation” — see [5], for example. For a variable whose type
 317 is polymorphic (contains a type variable), we must avoid assigning a value with
 318 a less general type. For example, in $*x = []$ the type of $*x$ is “list of τ ”, where
 319 τ is a type variable. Without destructive update it should be possible to use $*x$
 320 wherever a list of any type is expected. However, if $*x$ is then assigned a list
 321 containing integers (which has a less general type), passing it to a function that


322 expects a list of functions violates type safety (“calling” an arbitrary integer is
 323 not safe). Pawns allows expressions to have their inferred types further instan-
 324 tiated using “::”, and the type checking pass of the compiler also inserts some
 325 type instantiation. The type checking pass ensures that direct update does not
 326 involve type instantiation but to improve flexibility, indirect update is checked
 327 during the sharing analysis.

328 4 The abstract domain

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

337 The components of a variable depend on its type. For non-recursive types
 338 other than arrays, each possible data constructor argument is represented sep-
 339 arately. For example, the type `Maybe (Maybe (Either Int Int))` can have an
 340 argument of an outer `Just` data constructor, an inner `Just` and `Left` and `Right`.
 341 A component can be represented using a list of `x.y` pairs containing a data con-
 342 structor and an argument number, giving the path from the outermost data con-
 343 structor to the given argument. For example, the components of the type above
 344 can be written as: `[Just.1]`, `[Just.1,Just.1]`, `[Just.1,Just.1,Left.1]` and
 345 `[Just.1,Just.1,Right.1]`. If variable `v` has value `Just Nothing`, the expres-
 346 sion `v.[Just.1]` represents the single main memory word containing the occur-
 347 rence of `Nothing`.

348 For `Ref t` types we proceed as if there was a `Ref` data constructor, so
 349 `vp.[Ref.1]` represents the word `vp` points to. For function types, values may
 350 be closures. A closure that has had `K` arguments supplied is represented as a
 351 data constructor `C1K` with these `K` arguments; these behave in the same way
 352 as other data constructor arguments with respect to sharing, except Pawns pro-
 353 vides no way to obtain a pointer to a closure argument. Closures also contain
 354 a code pointer and an integer which are not relevant to sharing so they are ig-
 355 nored in the analysis. We also ignore the subscript on the data constructor for
 356 sharing analysis because type and sharing analysis only give a lower bound on
 357 the number of closure arguments. Our analysis orders closure arguments so that
 358 the most recently supplied argument is first (the reverse of the more natural
 359 ordering). Consider the code below, where `foo` is a function that is defined with
 360 four or more arguments. The sharing analysis proceeds as if the memory layout
 361 was as depicted in the diagram. The pre- and post-conditions of `foo` are part of
 362 the type information associated with `c1`, `c2` and `c3`.



364 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” [6] 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 that 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 that can be used to represent trees which consist of a node containing an integer and a list of sub-trees:

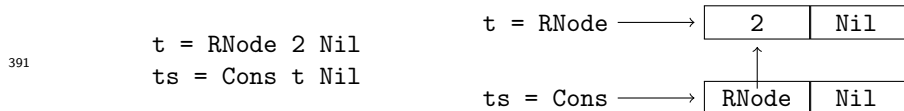
```

375 data RTrees = Nil | Cons RTree RTrees
376 data RTree = RNode Int RTrees

```

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

387 As well as containing pairs of components for distinct variables which may
388 alias, the abstract domain contains “self-alias” pairs for each possible component
389 of a variable which may exist. Consider the following two bindings and the
390 corresponding diagram (as with `Cons`, no tag is used for `RNode`):



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

```

396   {{t.[RNode.1], t.[RNode.1]},
397    {t.[RNode.2], t.[RNode.2]},
398    {ts.[], ts.[]},
399    {ts.[Cons.1], ts.[Cons.1]},
400    {ts.[Cons.1,RNode.1], ts.[Cons.1,RNode.1]},
401    {t.[RNode.1], ts.[Cons.1,RNode.1]},
402    {t.[RNode.2], ts.[]}}

```

403 Note there is no self-alias pair for $t.[]$ since there is no strict sub-part of t
404 that is an `RTree`. Similarly, there is no alias between $ts.[Cons.1]$ and any part
405 of t . Although the value t is used as the first argument of `Cons` in ts , this is not
406 a main memory word that is used to represent the value of t (indeed, the value
407 of t has no `Cons` cells). The tagged pointer value stored in variable t (which
408 may be in a register) is copied into the cons cell. Such descriptions of sharing are
409 an abstraction of computation states. The set above abstracts all computation
410 states in which t is a tree with a single node, ts is a list of trees, elements of
411 ts may be t or have t as a subtree, and there are no other live variables with
412 non-atomic values.

413 5 The sharing analysis algorithm

414 We now describe the sharing analysis algorithm. Overall, the compiler attempts
415 to find a proof that for a computation with a depth D of (possibly recursive)
416 function calls, the following condition C holds, assuming C holds for all compu-
417 tations of depth less than D . This allows a proof by induction that C holds for
418 all computations that terminate normally.

419 C : For all functions f , if the precondition of f is satisfied (abstracts the compu-
420 tation state) whenever f is called, then

- 421 1. for all function calls and assignment statements in f , any live variable that
422 may be updated at that point in an execution of f is annotated with “!”,
- 423 2. there is no update of live “abstract” variables when executing f ,
- 424 3. all parameters of f which may be updated when executing f are declared
425 mutable in the type signature of f ,
- 426 4. the union of the pre- and post-conditions of f abstracts the state when
427 f returns plus the values of mutable parameters in all states during the
428 execution of f ,
- 429 5. for all function calls in f , the sharing information among the actual pa-
430 rameters is a subset of the sharing information among formal parameters as
431 declared in the precondition, modulo variable renaming,
- 432 6. for all function calls and assignment statements in f , any live variable that
433 may be directly updated at that point is updated with a value of the same
434 type or a more general type, and
- 435 7. for all function calls and assignment statements in f , any live variable that
436 may be indirectly updated at that point only shares with variables of the
437 same type or a more general type.

438 The algorithm is applied to each function definition in core Pawns to compute
 439 an approximation to the sharing before and after each statement (we call it the
 440 alias set). This can be used to check points 1, 2, 4, 5 and 7 above; 5 allows
 441 the induction hypothesis to be used. Point 3 is established using point 1 and
 442 a simple syntactic check that any parameter of f that is annotated “!” in the
 443 definition is declared mutable in the type signature (parameters are considered
 444 live throughout the definition). Point 6 relies on 3 and the type checking pass.
 445 The core of the algorithm is to compute the alias set after a statement, given
 446 the alias set before the statement. This is applied recursively for compound
 447 statements in a form of abstract execution.

448 We do not prove correctness of the algorithm but hope our presentation is
 449 sufficiently detailed to have uncovered any bugs. A proof would have a separate
 450 case for each kind of statement in the core language, showing that if the initial
 451 alias set abstracts the execution state before the statement the resulting alias set
 452 abstracts the execution state after the statement. This would require a more formal
 453 description of execution states and their relationship with the core language
 454 and the abstract domain. The abstract domain relies on type information so the
 455 sharing analysis relies on type preservation in the execution. Type preservation
 456 also relies on sharing analysis. Thus a completely formal approach must tackle
 457 both problems together. Although our approach is not formal, we do state the
 458 key condition C , which has points relating to both sharing and types, and we
 459 include `Instype` in the core language.

460 The alias set used at the start of a definition is the precondition of the func-
 461 tion. This implicitly includes self-alias pairs for all variable components of the
 462 arguments of the function and the pseudo-variables `abstractT` for each type T
 463 used. Similarly, the postcondition implicitly includes self-alias pairs for all com-
 464 ponents of the result (and the `abstractT` variable if the result is abstract)³. As
 465 abstract execution proceeds, extra variables from the function body are added
 466 to the alias set and variables that are no longer live can be removed to improve
 467 efficiency. For each program point, the computed alias set abstracts the compu-
 468 tation state at that point in all concrete executions of the function that satisfy
 469 the precondition. For mutable parameters of the function, the sharing computed
 470 also includes the sharing from previous program points. The reason for this spe-
 471 cial treatment is explained when we discuss the analysis of function application.
 472 The alias set computed for the end of the definition, with sharing for local vari-
 473 ables removed, must be a subset of the union of the pre- and post-condition of
 474 the function.

475 Before sharing analysis, a type checking/inference pass is completed which
 476 assigns a type to each variable and function application. This determines the
 477 components for each variable. Polymorphism is also eliminated as follows. Sup-
 478 pose we have a function `take :: Int -> [a] -> [a] sharing take n xs =`
 479 `ys pre nosharing post ys = xs` which returns the list containing the first n
 480 elements of `xs`. For each call to `take`, the pre- and post-conditions are deter-

³ Self-aliasing 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.

481 mined based on the type of the application. An application to lists of Booleans
 482 will have two components for each variable whereas an application to lists of lists
 483 of Booleans will have four. When analysing the definition of `take` we instantiate
 484 type variables such as `a` above to `Ref ()`. This type has a single component
 485 which can be shared to represent possible sharing of arbitrary components of
 486 an arbitrary type. Finally, we assume there is no type which is an infinite chain
 487 of refs, for example, `type Refs = Ref Refs` (for which type folding results in
 488 an empty component rather than a `[Ref.1]` component; this is not a practical
 489 limitation).

490 Suppose a_0 is the alias set just before statement s . The following algo-
 491 rithm computes `alias(s, a0)`, the alias set just after statement s . The algorithm
 492 structure follows the recursive definition of statements and we describe it using
 493 pseudo-Haskell, interspersed with discussion. The empty list is written `[]`, non-
 494 empty lists are written `[a, b, c]` or `a:b:c:[]` and `++` denotes list concatenation.
 495 At some points we use high level declarative set comprehensions to describe what
 496 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

```

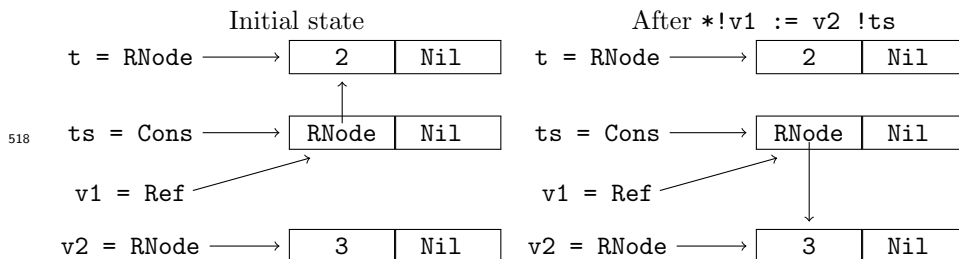
497 Sequencing is handled by function composition. To bind a fresh variable `v1` to
 498 a variable `v2` the self-aliasing of `v2` is duplicated for `v1` and the aliasing for each
 499 component of `v2` is duplicated for `v1`. Binding `*v1` to `v2` is done in a similar way,
 500 but the components of `v1` must have `Ref.1` prepended to them and the result
 501 folded, and the `[Ref.1]` component of `v1` self-aliases. Folding is only needed for
 502 the rare case of types with recursion through `Ref`.

```


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 = {va.ca | {v1.[Ref.1], va.ca} ∈ 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 ∈ al ∧ {v2.c, v2.c} ∈ a0}
    shareal = {{fc(va.(ca++c1)), v.c2} |
              va.ca ∈ al ∧ {v2.c1, v.c2} ∈ a0} ∪
              {{fc(va.(ca++c)), fc(v1.(Ref.1 :c))} |
              va.ca ∈ al ∧ {v2.c, v2.c} ∈ a0}
  in if v1 is a mutable parameter then
      a0 ∪ self1 ∪ share1 ∪ selfal ∪ shareal
    else let
      -- old1 = old aliases for v1, which can be removed
      old1 = {{v1.(Ref.1 :d : c1), v.c2} | {v1.(Ref.1 :d : c1), v.c2} ∈ a0}
    in (a0 \ old1) ∪ self1 ∪ share1 ∪ selfal ∪ shareal

```

503 Assignment to an existing variable `*v1` adds the same sharing as for binding
 504 a fresh variable, but there are two extra complications. First, `*v1` may be an
 505 alias for components of other variables (the live subset of these variables and
 506 `v1` must be annotated with “!” on the assignment statement; checking such
 507 annotations is a primary purpose of the sharing analysis). All these variable
 508 components must have the same sharing added as `*v1`. The components must be
 509 concatenated and folded appropriately. Second, if `v1` is not a mutable parameter
 510 the existing sharing with a path strictly longer than `[Ref.1]` can safely be
 511 removed, improving precision. The component `v1.[Ref.1]` represents the single
 512 memory word which is overwritten and whatever the old contents shared with
 513 is no longer needed to describe the sharing for `v1`. For mutable parameters the
 514 old value may share with variables from the calling context and we retain this
 515 information, as explained later. Consider the example below, where `t` and `ts` are
 516 as before, local variable `v1` is a reference to the element of `ts` and it is assigned
 517 `v2`, which is `RNode 3 Nil`.



519 Because $\{ts.[Cons.1], v1.[Ref.1]\}$ is in any correct approximation to the
 520 initial state, ts will be in $a1$ and will have sharing with $v2$ added. The old sharing
 521 of $v1$ with t will be discarded. Note that we cannot discard the old sharing of
 522 ts with t for two reasons. First, the assignment updates only one memory word
 523 whereas there may be other words also represented by $ts.[Cons.1]$. Second,
 524 we only know $ts.[Cons.1]$ possibly aliases $v1.[Ref.1]$ — no definite aliasing
 525 information is maintained. In some cases the old sharing of $v1$ is discarded and
 526 immediately added again. Consider the following example, which creates a cyclic
 527 list.



529 The old sharing between $v1$ and $v3$ is discarded but added again (via `share1`)
 530 because $v2$ also shares with $v3$. Correctness of the algorithm when cyclic terms
 531 are created depends on the abstract domain we use. A more expressive domain
 532 could distinguish between different cons cells in a list. For example, if types are
 533 “folded” at the third level of recursion rather than the first, the domain can
 534 distinguish three classes of cons cells, where the distance from the first cons cell,
 535 modula three, is zero, one or two. For a cyclic list with a single cons cell, that
 536 cons cell must be in all three classes and our algorithm would need modification
 537 to achieve this. However, in our domain types are folded at the first level of
 538 recursion so we have a unique folded path for each memory cell in cyclic data
 539 structure (cyclic terms can only be created with recursive types). There is no
 540 distinction between the first and second cons cell in a list, for example.

```

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

541 The `DerefEq` case can be seen as equivalent to $v1 = Ref\ v2$ and binding a
 542 variable to a data constructor with N variable arguments is a generalisation. If
 543 there are multiple v_i that share, the corresponding components of v must also
 544 share; these pairs are included in `self1`.


```

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

```

545 The EqDeref case is similar to the inverse of DerefEq in that we are removing
 546 Ref.1 rather than prepending it. However, if the empty component results we
 547 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
    mut = the arguments that are declared mutable
    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
    -- (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 needed for possible closure creation
    selfc = {{v.[Cl.i], v.[Cl.i]} | 1 ≤ i ≤ N} ∪
            {{v.((Cl.(N + 1 - i)) :c1), v.((Cl.(N + 1 - j)) :c2)} |
             {vi.c1, vj.c2} ∈ a0} ∪
            {{v.((Cl.(i + N)) :c1), v.((Cl.(j + N)) :c2)} |
             {f.((Cl.i) :c1), f.((Cl.j) :c2)} ∈ a0}
    sharec = {{v.((Cl.(N + 1 - i)) :c1), x.c2} | {vi.c1, x.c2} ∈ a0} ∪
             {{v.((Cl.(i + N)) :c1), x.c2} | {f.((Cl.i) :c1), x.c2} ∈ a0}
    -- postt+postm needed for possible function call
    postt = {{x1.c1, x3.c3} | {x1.c1, x2.c2} ∈ post ∧ {x2.c2, x3.c3} ∈ a0}
    postm = {{x1.c1, x2.c2} | {x1.c1, vi.c3} ∈ a0 ∧ {x2.c2, vj.c4} ∈ a0 ∧
             {vi.c3, vj.c4} ∈ post ∧ vi ∈ mut ∧ vj ∈ mut}
  in
    a0 ∪ selfc ∪ sharec ∪ postt ∪ postm


```

548 For many App occurrences the function is known statically and we can deter-
 549 mine if the function is actually called or a closure is created instead. However,

550 in general we must assume either could happen and add sharing for both. If a
 551 closure is created, the first N closure arguments share with the N arguments of
 552 the function call and any closure arguments of \mathbf{f} share with additional closure
 553 arguments of the result (this requires renumbering of these arguments). Anal-
 554 ysis of function calls relies on the sharing and mutability information attached
 555 to all arrow types. Because Pawns uses the syntax of statements to express pre-
 556 and post-conditions, our implementation uses the sharing analysis algorithm to
 557 derive an explicit alias set representation (currently this is done recursively, with
 558 the level of recursion limited by the fact than pre- and post-conditions must not
 559 contain function calls). Here we ignore the details of how the alias set represen-
 560 tation is obtained. The compiler also uses the sharing information immediately
 561 before an application to check that the precondition is satisfied, all required “!”
 562 annotations are present and abstract variables are not modified.

563 Given that the precondition is satisfied, the execution of a function results in
 564 sharing of parameters that is a subset of the union of the declared pre- and post-
 565 conditions (we assume the induction hypothesis holds for the sub-computation,
 566 which has a smaller depth of recursion). However, any sharing between non-
 567 mutable arguments that exists immediately after the call must exist before the
 568 call. The analysis algorithm does not add sharing between non-mutable argu-
 569 ments in the precondition as doing so would unnecessarily restrict how “high
 570 level” and “low level” code can be mixed. It is important we can pass a variable
 571 to a function that allows an abstract argument without the analysis conclud-
 572 ing the variable subsequently shares with **abstract**, and therefore cannot be
 573 updated. Thus **post** is just the declared postcondition plus the subset of the
 574 precondition which involves mutable parameters of the function, renamed ap-
 575 propriately. The last N formal parameters, $w_{K+1} \dots w_{K+N}$ are renamed as the
 576 arguments of the call, $v_1 \dots v_N$ and the formal result r is renamed \mathbf{v} . The formal
 577 parameters $w_1 \dots w_K$ represent closure arguments $K \dots 1$ of \mathbf{f} . Thus a variable
 578 component such as w_1 . [**Cons**. 1] is renamed \mathbf{f} . [**C1**. K , **Cons**. 1].

579 It is also necessary to include one step of transitivity in the sharing informa-
 580 tion: if variable components $x_1.c_1$ and $x_2.c_2$ alias in **post** and $x_2.c_2$ and $x_3.c_3$
 581 (may) alias before the function call, we add an alias of $x_1.c_1$ and $x_3.c_3$ (in **postt**).
 582 Function parameters are proxies for the argument variables as well as any vari-
 583 able components they may alias and when functions are analysed these aliases
 584 are not known. This is why the transitivity step is needed, and why mutable
 585 parameters also require special treatment. If before the call, $x_1.c_1$ and $x_2.c_2$ may
 586 alias with mutable parameter components $v_i.c_3$ and $v_j.c_4$, respectively, and the
 587 two mutable parameter components alias in **post** then $x_1.c_1$ and $x_2.c_2$ may alias
 588 after the call; this is added in **postm**. Consider the example below, where we
 589 have a pair $\mathbf{v1}$ (of references to references to integers) and variables \mathbf{x} and \mathbf{y}
 590 share with the two elements of $\mathbf{v1}$, respectively. When $\mathbf{v1}$ is passed to function
 591 $\mathbf{f1}$ as a mutable parameter, sharing between \mathbf{x} and \mathbf{y} is introduced. The sharing
 592 of the mutable parameter in the postcondition, $\{\mathbf{v1}$. [**Pair**. 1, **Ref**. 1, **Ref**. 1],
 593 $\mathbf{v1}$. [**Pair**. 2, **Ref**. 1, **Ref**. 1] $\}$, results in sharing between \mathbf{x} and \mathbf{y} being added in
 594 the analysis.

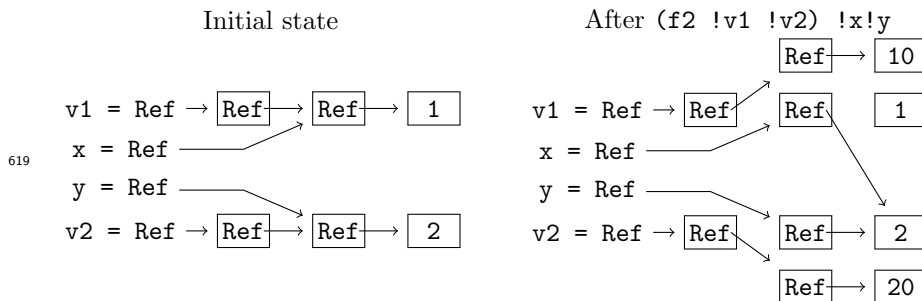


```

596 f1 :: Pair (Ref (Ref Int)) -> ()
597     sharing f1 !v1 = _
598     pre nosharing
599     post *a = *b; v1 = Pair a b
600 f1 !v1 =
601     case v1 of (Pair rr1 rr2) -> *rr1 := *rr2 !v1

```

602 The need to be conservative with the sharing of mutable parameters in the
603 analysis of function definitions (the special treatment in `Assign`) is illustrated
604 by the example below. Consider the initial state, with variables `v1` and `v2` which
605 share with `x` and `y`, respectively. After `f2` is called `x` and `y` share, even though
606 the parameters `v1` and `v2` do not share at any point in the execution of `f2`. If
607 mutable parameters were not treated specially in the `Assign` case, `nosharing`
608 would be accepted as the postcondition of `f2` and the analysis of the call to
609 `f2` would then be incorrect. The sharing is introduced between memory cells
610 that were once shared with `v1` and others that were once shared with `v2`. Thus
611 in our algorithm, the sharing of mutable parameters reflects all memory cells
612 that are reachable from the parameters during the execution of the function.
613 Where the mutable parameters are assigned in `f2`, the sharing of the parameters
614 previous values (`rr1` and `rr2`) is retained. Thus when the final assignment is
615 processed, sharing between the parameters is added and this must be included
616 in the postcondition. Although this assignment does not modify `v1` or `v2`, the
617 “!” annotations are necessary and alert the reader to potential modification of
618 variables that shared with the parameters when the function was called.



```

620 f2 :: Ref (Ref (Ref Int)) -> Ref (Ref (Ref Int)) -> ()
621     sharing f2 !v1 !v2 = _
622     pre nosharing
623     post **v1 = **v2
624 f2 !v1 !v2 =
625     *r10 = 10           -- ref to new cell containing 10
626     *rr10 = r10        -- ref to above ref
627     *r20 = 20          -- ref to new cell containing 20
628     *rr20 = r20        -- ref to above ref
629     rr1 = *v1          -- save *v1
630     rr2 = *v2          -- save *v2
631     *!v1 := rr10       -- update *v1 with Ref (Ref 10)
632     *!v2 := rr20       -- update *v2 with Ref (Ref 20)
633     *rr1 := *rr2 !v1!v2 -- can create sharing at call


alias Error a0 = ∅ -- error
alias (Case v [(p1, s1), ..., (pN, sN)]) a0 = -- case v of ...
  let
    old = {{v.c1, v2.c2} | {v.c1, v2.c2} ∈ a0}
  in
    ⋃1 ≤ i ≤ 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)

```

634 For a case expression we return the union of the alias sets obtained for each of
635 the different branches. For each branch we only keep sharing information for the
636 variable we are switching on that is compatible with the data constructor in that
637 branch (we remove all the old sharing, `av`, and add the compatible sharing, `avdc`).
638 Note we use a high level declarative definition for `avdc` (and other variables)
639 which implicitly uses the inverse of `fc`. To deal with individual data constructors
640 we consider pairs of components of arguments i and j which may alias in order
641 to compute possible sharing between v_i and v_j , including self-aliases when $i = j$.
642 The corresponding component of v_i (prepended with `Ref` and folded) may alias
643 the component of v_j . For example, if v of type `RTrees` is matched with `Cons *v1`
644 `*v2` and v . `[]` self-aliases, we need to find the components which fold to v . `[]`
645 (v . `[Cons.2]` and v . `[Cons.1, RNode.2]`) in order to compute the sharing for $v2$

646 and `v1`. Thus we compute that `v2.[Ref.1]`, may alias `v1.[Ref.1,RNode.2]`.
 647 This can occur if the data structure is cyclic, such as the example below where `v`
 648 is a list containing a single tree with 2 in the node and `v` as the children (hence it
 649 represents a single infinite branch). Note that `v1.[Ref.1,RNode.2]` represents
 650 both the memory cell containing the `Cons` pointer and the cell containing `Nil`.



```

651
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)
  
```

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

656 6 Example

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

```

666 list_bst xs =                -- 0
667     v1 = TNil                -- 1
668     *tp = v1                 -- 2
669     list_bst_du xs !tp       -- 3
670     ret = *tp                -- 4
  
```

671 We start with the precondition: $a_0 = \{\{xs.[Cons.1], absL.[Cons.1]\},$
 672 $\{xs.[], absL.[]\}$. Binding to a constant introduces no sharing so $a_1 = a_0.$
 673 $a_2 = a_1 \cup \{tp.[Ref.1]\}$. The function call has precondition $a_0 \cup \{tp.[Ref.1]\},$
 674 $\{tp.[Ref.1,Node.2]\}$, which is a superset of a_2 . Since `tp` is a mutable argu-
 675 ment the precondition sharing for `tp` is added: $a_3 = a_2 \cup \{tp.[Ref.1,$

676 Node.2]}}. The final sharing includes the return variable, $\text{ret}: a_4 = a_3 \cup$
 677 $\{\{\text{ret}. [], \text{tp}. [\text{Ref}. 1]\}, \{\text{ret}. [\text{Node}. 2], \text{tp}. [\text{Ref}. 1, \text{Node}. 2]\}\}$. After remov-
 678 ing sharing for the dead (local) variable tp we obtain a subset of the union of
 679 the pre- and post-conditions, which is $a_0 \cup \{\{\text{ret}. [], \text{absT}. []\}, \{\text{ret}. [\text{Node}. 2],$
 680 $\text{absT}. [\text{Node}. 2]\}\}$.

```

681 list_bst_du xs !tp =                -- 0
682     case xs of
683     (Cons *v1 *v2) ->                -- 1
684         x = *v1                       -- 2
685         xs1 = *v2                      -- 3
686         v3 = bst_insert_du x !tp       -- 4
687         v4 = list_bst_du xs1 !tp      -- 5
688         ret = v4                       -- 6
689     Nil ->                             -- 7
690         ret = ()                       -- 8
691     -- after case                     -- 9

```

692 We start with the precondition, $a_0 = \{\{\text{tp}. [\text{Ref}. 1]\}, \{\text{tp}. [\text{Ref}. 1, \text{Node}. 2]\},$
 693 $\{\text{xs}. [\text{Cons}. 1], \text{absL}. [\text{Cons}. 1]\}, \{\text{xs}. [], \text{absL}. []\}\}$. The Cons branch of the
 694 case introduces sharing for $v1$ and $v2$: $a_1 = a_0 \cup \{\{\text{xs}. [\text{Cons}. 1], \text{absL}. [\text{Cons}. 1],$
 695 $v1. [\text{Ref}. 1], v2. [\text{Ref}. 1, \text{Cons}. 1]\}, \{v2. [\text{Ref}. 1], \text{xs}. [], \text{absL}. []\}\}$. The list
 696 elements are atomic so $a_2 = a_1$. The next binding makes the sharing of xs1 and
 697 xs the same: $a_3 = a_2 \cup \{\{v2. [\text{Ref}. 1], \text{xs}. [], \text{xs1}. [], \text{absL}. []\}, \{v1. [\text{Ref}. 1],$
 698 $\text{xs}. [\text{Cons}. 1], \text{xs1}. [\text{Cons}. 1], \text{absL}. [\text{Cons}. 1], v2. [\text{Ref}. 1, \text{Cons}. 1]\}\}$. This can
 699 be simplified by removing the dead variables $v1$ and $v2$. The precondition of the
 700 calls are satisfied and $a_6 = a_5 = a_4 = a_3$. For the Nil branch we remove the in-
 701 compatible sharing for xs from a_0 : $a_7 = \{\{\text{tp}. [\text{Ref}. 1]\}, \{\text{tp}. [\text{Ref}. 1, \text{Node}. 2]\},$
 702 $\{\text{absL}. [\text{Cons}. 1]\}, \{\text{absL}. []\}\}$ and $a_8 = a_7$. Finally, $a_9 = a_6 \cup a_8$. This contains
 703 all the sharing for mutable parameter tp and, ignoring local variables, is a subset
 704 of the union of the pre- and post-conditions, a_0 .

```

705 bst_insert_du x !tp =                -- 0
706     v1 = *tp                           -- 1
707     case v1 of
708     TNil ->                             -- 2
709         v2 = TNil                       -- 3
710         v3 = TNil                       -- 4
711         v4 = Node v2 x v3               -- 5
712         *!tp := v4                     -- 6
713         ret = ()                       -- 7
714     (Node *lp *v5 *rp) ->              -- 8
715         n = *v5                         -- 9
716         v6 = (x <= n)                  -- 10
717         case v6 of
718         True ->                         -- 11

```

```

719         v7 = (bst_insert_du x !lp) !tp      -- 12
720         ret = v7                            -- 13
721     False ->                                -- 14
722         v8 = (bst_insert_du x !rp) !tp      -- 15
723         ret = v8                            -- 16
724     -- end case                             -- 17
725 -- end case                                -- 18

```

726 Here $a_0 = \{\{tp. [Ref. 1]\}, \{tp. [Ref. 1, Node. 2]\}\}$ and $a_1 = a_0 \cup \{\{v1. [],$
727 $tp. [Ref. 1]\}, \{tp. [Ref. 1, Node. 2], v1. [Node. 2]\}\}$. For the TNil branch we
728 remove the v1 sharing so $a_4 = a_3 = a_2 = a_0$ and $a_5 = a_4 \cup \{\{v4. []\},$
729 $\{v4. [Node. 2]\}\}$. After the destructive update, $a_6 = a_5 \cup \{\{v4. [], tp. [Ref. 1]\},$
730 $\{v4. [Node. 2], tp. [Ref. 1, Node. 2]\}\}$ (v4 is dead and can be removed) and $a_7 =$
731 a_6 . For the Node branch we have $a_8 = a_1 \cup \{\{v1. [], tp. [Ref. 1], lp. [Ref. 1],$
732 $rp. [Ref. 1]\}, \{tp. [Ref. 1, Node. 2], lp. [Ref. 1, Node. 2], rp. [Ref. 1, Node. 2],$
733 $v5. [Ref. 1], v1. [Node. 2]\}\}$. The same set is retained for $a_9 \dots a_{17}$ (assuming
734 the dead variable v5 is retained), the preconditions of the function calls are sat-
735 isfied and the required annotations are present. Finally, $a_{18} = a_{17} \cup a_7$, which
736 contains all the sharing for **tp**, and after eliminating local variables we get the
737 postcondition, which is the same as the precondition.

738 7 Discussion

739 Imprecision in the analysis of mutable parameters could potentially be reduced
740 by allowing the user to declare that only certain parts of a data structure are
741 mutable, as suggested in [1]. It is inevitable we lose some precision with recursion
742 in types, but it seems that some loss of precision could be avoided relatively
743 easily. The use of the empty path to represent sub-components of recursive types
744 results in imprecision when references are created. For example, the analysis of
745 `*vp = Nil; v = *vp` concludes that the empty component of **v** may alias with
746 itself and the **Ref** component of **vp** (in reality, **v** has no sharing). Instead of the
747 empty path, a dummy path of length one could be used. Flagging data structures
748 which are known to be acyclic could also improve precision for **Case**. A more
749 aggressive approach would be to unfold the recursion an extra level, at least for
750 some types. This could allow us to express (non-)sharing of separate subtrees
751 and whether data structures are cyclic, at the cost of more variable components,
752 more complex pre- and post-conditions and more complex analysis for **Assign**
753 and **Case**.

754 Increasing the number of variable components also decreases efficiency. The
755 algorithmic complexity is affected by the representation of alias sets. Currently
756 we use a naive implementation, using just ordered pairs of variable components
757 as the set elements and a set library which uses an ordered binary tree. The size
758 of the set can be $O(N^2)$, where N is the maximum number of live variable com-
759 ponents of the same type at any program point (each such variable component
760 can alias with all the others). In typical code the number of live variables at

761 any point is not particularly large. If the size of alias sets does become problem-
762 atic, a more refined set representation could be used, such as the set of sets of
763 pairs representation we used in Section 6, where sets of components that all alias
764 with each other are optimised. There are also simpler opportunities for efficiency
765 gains, such as avoiding sharing analysis for entirely pure code. We have not stress
766 tested our implementation or run substantial benchmarks as it is intended to be
767 a prototype, but performance has been encouraging. Translating the tree inser-
768 tion code plus a test harness to C, which includes the sharing analysis, takes
769 around half the time of compiling the resulting C code using GCC. Total com-
770 pilation time is less than half that of GHC for equivalent Haskell code and less
771 than one tenth that of MLton for equivalent ML code. The Pawns executable is
772 around 3–4 times as fast as the others.

773 8 Related work

774 Related programming languages are discussed in [1]; here we restrict attention
775 to work related to the sharing analysis algorithm. The most closely related work
776 is that done in the compiler for Mars [7], which extends similar work done for
777 Mercury [8] and earlier for Prolog [9]. All use a similar abstract domain based on
778 the type folding method first proposed in [6]. Our abstract domain is somewhat
779 more precise due to inclusion of self-aliasing, and we have no sharing for con-
780 stants. In Mars it is assumed that constants other than numbers can share. Thus
781 for code such as `xs = []; ys = xs` our analysis concludes there is no sharing
782 between `xs` and `ys` whereas the Mars analysis concludes there may be sharing.

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

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

804 which constructs paths with edges alternating from s_1 and s_2 , for example **a-b**
 805 (from s_1), **b-c** (from s_2), **c-d** (from s_1) and **d-e** (from s_2).

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

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

821 There is also substantial work on sharing analysis for logic programming
 822 languages using other abstract domains, notably the set-sharing domain of [10]
 823 (a set of sets of variables), generally with various enhancements — see [11] for a
 824 good summary and evaluation. Applications include avoiding the “occurs check”
 825 in unification [12] and exploiting parallelism of independent sub-computations
 826 [13]. These approaches are aimed at identifying sharing of logic variables rather
 827 than sharing of data structures. For example, although the two Prolog goals `p(X)`
 828 and `q(X)` share `X`, they are considered independent if `X` is instantiated to a data
 829 structure that is ground (contains no logic variables). Ground data structures in
 830 Prolog are read-only and cause no problem for parallelism or the occurs check,
 831 whether they are shared or not. For this reason, the set-sharing domain is often
 832 augmented with extra information related to groundness [11]. In Pawns there
 833 are no logic variables but data structures are mutable, hence their sharing is
 834 important.

835 However, the set-sharing domain (with enhancements) has been adapted to
 836 analysis of sharing of data structures in object oriented languages such as Java
 837 [14]. One important distinction is that Pawns supports algebraic data types
 838 which allow a “sum of products”: there can be a choice of several data con-
 839 structors (a sum), where each one consists of several values as arguments (a
 840 product). Java and most other imperative and object oriented languages do not.
 841 Products are supported by objects containing several values but the only choice
 842 (sum) is whether the object is null or not. Java objects and pointers in most
 843 imperative languages are similar to a `Maybe` algebraic data type, with `Nothing`
 844 corresponding to null. A `Ref` cannot be null. The abstract domain of [14] uses
 845 set-sharing plus additional information about what objects are definitely not
 846 null. For Pawns code that uses `Refs` this information is given by the data type
 847 — the more expressive types allow us to trivially infer some information that is

848 obscured in other languages. For code that uses `Maybe`, our domain can express
849 the fact that a variable is definitely `Nothing` by not having a self-alias of the
850 `Just` component. The rich structural information in our domain fits particularly
851 well with algebraic data types. There are also other approaches to and uses of
852 alias analysis for imperative languages, such as [15] and [16], but these are not
853 aimed at precisely capturing information about dynamically allocated data. A
854 more detailed discussion of such approaches is given in [7].

855 9 Conclusion

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

870 The algorithm operates on (heap allocated) algebraic data types, including
871 arrays and closures. In common with other sharing analysis used in declarative
872 languages it supports binding of variables, construction and deconstruction
873 (combined with selection or “case”) and function/procedure calls. In addition, it
874 supports explicit pointers, destructive update via pointers, creation and applica-
875 tion of closures and pre- and post-conditions concerning sharing attached to type
876 signatures of functions. It also uses an abstract domain with additional features
877 to improve precision. Early indications are that the performance is acceptable:
878 compared with other compilers for declarative languages, the prototype Pawns
879 compiler supports encapsulated destructive update, is fast and produces fast
880 executables.

881 Acknowledgements

882 Reviewers’ comments on an earlier version of this paper were very helpful in
883 ironing out some important bugs in the algorithm and improving the presenta-
884 tion.

885 References

- 886 1. Naish, L.: An informal introduction to Pawns: a declarative/imperative language.
887 <http://people.eng.unimelb.edu.au/lee/papers/pawns> (2015) Accessed 2015-03-16.

- 888 2. Milner, R., Tofte, M., Macqueen, D.: *The Definition of Standard ML*. MIT Press,
889 Cambridge, MA, USA (1997)
- 890 3. Jones, S.P., Hughes, J., Augustsson, L., Barton, D., Boutel, B., Burton, W., Fasel,
891 J., Hammond, K., Hinze12, R., Hudak, P., et al.: Report on the programming
892 language Haskell 98, a non-strict purely functional language, February 1999. (1999)
- 893 4. Lippmeier, B.: *Type Inference and Optimisation for an Impure World*. PhD thesis,
894 Australian National University (June 2009)
- 895 5. Wright, A.: Simple imperative polymorphism. In: *LISP and Symbolic Computa-*
896 *tion*. (1995) 343–356
- 897 6. Bruynooghe, M.: *Compile time garbage collection or how to transform programs in*
898 *an assignment free languages into code with assignments*. IFIP TC2/WG2.1 Work-
899 *ing Conference on Program Specification and Transformation, Bad-Tölz, Germany,*
900 *April 15-17, 1986* (1986)
- 901 7. Giuca, M.: *Mars: An imperative/declarative higher-order programming language*
902 *with automatic destructive update*. PhD thesis, University of Melbourne (2014)
- 903 8. Mazur, N., Ross, P., Janssens, G., Bruynooghe, M.: Practical aspects for a working
904 compile time garbage collection system for Mercury. In Codognet, P., ed.: *Proceed-*
905 *ings of ICLP 2001, Lecture Notes in Computer Science. Volume 2237.*, Springer
906 (2001) 105–119
- 907 9. Mulkers, A.: *Live data structures in logic programs, derivation by means of abstract*
908 *interpretation*. Springer-Verlag (1993)
- 909 10. Jacobs, D., Langen, A.: Accurate and efficient approximation of variable aliasing
910 in logic programs. In Lusk, E.L., Overbeek, R.A., eds.: *NACLP*, MIT Press (1989)
911 154–165
- 912 11. Bagnara, R., Zaffanella, E., Hill, P.M.: Enhanced sharing analysis techniques: a
913 comprehensive evaluation. *Theory and Practice of Logic Programming* **5** (1 2005)
914 1–43
- 915 12. Søndergaard, H.: An application of abstract interpretation of logic programs: Occur
916 check reduction. In: *Proc. Of the European Symposium on Programming on ESOP*
917 *86, New York, NY, USA, Springer-Verlag New York, Inc.* (1986) 327–338
- 918 13. Bueno, F., García de la Banda, M., Hermenegildo, M.: Effectivness of abstract
919 interpretation in automatic parallelization: A case study in logic programming.
920 *ACM Trans. Program. Lang. Syst.* **21**(2) (March 1999) 189–239
- 921 14. Méndez-Lojo, M., Hermenegildo, M.: Precise set sharing analysis for Java-style
922 programs. In Logozzo, F., Peled, D., Zuck, L., eds.: *Verification, Model Checking,*
923 *and Abstract Interpretation. Volume 4905 of Lecture Notes in Computer Science.*
924 Springer Berlin Heidelberg (2008) 172–187
- 925 15. Landi, W., Ryder, B.G.: A safe approximate algorithm for interprocedural aliasing.
926 *SIGPLAN Not.* **27**(7) (July 1992) 235–248
- 927 16. Emami, M., Ghiya, R., Hendren, L.J.: Context-sensitive interprocedural points-to
928 analysis in the presence of function pointers. In: *Proceedings of the ACM SIGPLAN*
929 *1994 Conference on Programming Language Design and Implementation. PLDI*
930 *'94, New York, NY, USA, ACM* (1994) 242–256