

Definite Reference Mutability

Anonymous¹

1 Anonymous Anonymous@Somewhere

Abstract

Reference immutability type systems such as Javari and ReIm ensure that a given reference cannot be used to mutate the referenced object. These systems are conservative in the sense that a mutable references may not be truly mutable, but mutable due to approximation.

In this paper, we present ReM (for definite Re[ference] M[utability]). It separates potentially mutable references into (1) definitely mutable, and (2) maybe mutable, i.e., references whose mutability is due to inherent approximation. In addition, we propose a CFL-reachability system for reference immutability, and prove that it is equivalent to ReIm/ReM, thus building a novel framework for reasoning about correctness of reference immutability type systems. We have implemented ReM and applied it on a large benchmark suite. Our results show that approximately 86.5% of all potentially mutable references are definitely mutable.

1998 ACM Subject Classification D.3.1 Formal Definitions and Theory

Keywords and phrases reference immutability, type inference, CFL-reachability, precision

Digital Object Identifier 10.4230/LIPIcs...

1 Introduction

Reference immutability ensures that a *readonly* reference cannot be used to modify the state of the object, including its transitively reachable state. For example, in the code below

```
1 Date md = new Date(); // mutable by default
2 readonly Date rd = md; // readonly reference
3 rd.setTime(1); // ERROR, rd is readonly
4 md.setTime(1); // OK
```

the `Date` object cannot be modified through the readonly reference `rd`, however, the same object can be modified through the mutable reference `md`.

Reference immutability has a wide variety of applications. It can enrich method specifications. It can help prevent errors due to unwanted aliasing and unwanted object mutation, as well as errors due to concurrency. It can enable compiler and runtime optimizations as well as reasoning about more complex properties such as method purity and object immutability. One application that has not received attention (to the best of our knowledge), is the impact of reference immutability on “flow systems”. Flow systems track and prevent flow from *positive* references to *negative* ones:

```
1 a = b;
2 positive X x = ... ;
3 a.f = x;
4 negative X y = b.f; // ERROR, b.f holds positive x
```



licensed under Creative Commons License CC-BY

Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

XX:2 Definite Reference Mutability

Many interesting analyses fall into this category, most notably approximate computing systems (e.g., EnerJ [25]), which prevent flow of approximate values into precise ones, and taint systems, which prevent flow from sensitive sources to untrusted sinks (e.g., [26, 16]). Unfortunately, the natural subtyping `negative <: positive` is unsound in the presence of mutable references [3]. (In the above example, had we allowed for such subtyping, reference `a` could have been positive, `b` could have been negative, and the program would have type checked.) Therefore, flow systems disallow subtyping for reference types [26, 25, 12]; reference type assignments entail equality constraints instead of subtyping constraints. Reference immutability can alleviate the imprecision arising from equality constraints — clearly, if the left-hand-side of the assignment is readonly, then subtyping is safe — allowing for more correct programs to type check. In summary, because of its many applications, reference immutability has been studied extensively [31, 36, 2, 37, 17, 13, 21], and it remains important to continue research in the area.

Javari [31] is the state-of-the-art in reference immutability. ReIm [17] has similar core semantics but is less expressive and therefore simpler. In this paper we focus on ReIm because of its simplicity and clarity; we believe that our treatment extends to other reference immutability systems. Standard reference immutability systems, like Javari and ReIm capture what we call *definite immutability* — a `readonly` reference is truly immutable. However, a `mutable` reference may be truly mutable, or it may be mutable because of inherent approximation. ReIm (and Javari) approximate in the handling of structure-transmitted dependences [24] (i.e., flow through heap objects). For example, in the code below

```
x.f = y; ... w = z.f; w.g = ...
```

reference `y` is `mutable`. However, it is not necessarily mutable: if `x` and `z` refer to the same object `o`, then `y` is indeed mutable; if they refer to different objects, then it is not mutable (at least not because of the update to `w`). The system does not reason about aliasing, and errs on the safe side marking `y` `mutable`. ReIm (and Javari) handle call-transmitted dependences [24] precisely. In the code below, `id` is the standard “identity” function that returns its argument.

```
x = id(y); // x is readonly
z = id(w); // z is mutable
```

Reference `y` is `readonly`, and `w` is `mutable`. The system properly transmits mutability without mixing the two call sites.

The key contribution of our paper is reasoning about approximation. We propose a new type system ReM (for definite Re[ference] M[utability]). ReM captures definite immutability, and in addition it captures definite mutability — a `mutable` reference is now *definitely mutable*. (We note that, strictly, our use of “definitely mutable” is inaccurate. Of course, whether a given reference is ever mutated is undecidable for various reasons, e.g., it is undecidable whether a given statement is executed, or whether a given path is executed. We use it in the sense of *definitely mutable according to CFL-reachability*, which is a highly precise model of data dependence [24] and analyses are unlikely to improve upon it.) ReM captures approximation explicitly by introducing the `maybe` qualifier. In the earlier example

```
x.f = y; ... w = z.f; w.g = ...
```

`y` is now `maybe mutable`. A key result is that empirically, approximation has limited impact — only about 13% of all potentially mutable references (about 6% of all references) are

maybe mutable, leading to a conclusion that ReIm and ReM are precise, and can be used to power client analyses.

Another contribution of our paper is the interpretation of reference immutability in terms of CFL-reachability. We propose a CFL-reachability system for inference of reference immutability, and prove that it is equivalent to ReIm/ReM’s inference system, thus building a framework for reasoning about correctness, and proving ReIm and ReM correct. To the best of our knowledge, ReIm has not been proven correct, even though it has been used to power client analyses [16, 29]. We envision that our system for reasoning about approximation and correctness will be extended to flow systems [26, 25, 17, 16].

This paper makes the following contributions.

- We present ReM, a novel type system for reference immutability. ReM captures explicitly definite mutability (in the CFL-reachability sense), and approximation.
- We interpret reference immutability in terms of CFL-reachability, and prove ReIm and ReM correct.
- We present an implementation and evaluation. We show that ReIm and ReM are precise — only 13% of potentially mutable references (6% of all references) are maybe mutable. The implementation is publicly available online; we omit the URL to preserve anonymity.

The rest of the paper is organized as follows. Sect. 2 presents the mutability semantics based on CFL-reachability. Sect. 3 interprets ReIm in terms of the mutability semantics, and presents the novel system ReM. Sect. 4 establishes equivalence between the systems in Sects. 2 and 3. Sect. 5 presents the empirical evaluation, Sect. 6 discusses related work, and Sect. 7 concludes.

2 Mutability Semantics

2.1 Flow Graph

The mutability semantics builds a *flow graph* G that represents flow (data) dependences between variables. The nodes in the graph are program variables, e.g., x , y , this , and field access expressions, e.g., $x.f$, $y.f$, $\text{this}.f$. The edges capture flow from one variable/field access expression, to another. The goal is to capture deep reference (im)mutability with data dependence paths in G . For example, in

$$x = y; z = x; z.f = w$$

y is mutable, because there is a path in G from y to z , which is the receiver of the update at field write $z.f = w$. Throughout the paper we refer to receivers at field writes as *updates*.

We restrict our core language to a “named form” in the style of Vaziri et al. [32, 10]. The language models Java with the syntax in Fig. 1, where the results of instantiations, field accesses, and method calls, are immediately stored in a variable. Without loss of generality, we assume that methods have parameter this , and exactly one other formal parameter. Features not strictly necessary are omitted from the formalism, but they are handled correctly in the implementation.

An assignment statement contributes a *direct* (i.e., intraprocedural) edge as follows:

$$x = y \quad \Rightarrow \quad y \xrightarrow{d} x$$

It represents flow from variable y to variable x . Therefore, if x is an update, i.e., there is field write $x.g = z$, the direct edge propagates mutability to reference y .

XX:4 Definite Reference Mutability

$cd ::= \text{class } C \text{ extends } D \{ \overline{fd} \overline{md} \}$	<i>class</i>
$fd ::= t f$	<i>field</i>
$md ::= t m(t \text{ this}, t x) \{ \overline{t y} s; \text{return } y \}$	<i>method</i>
$s ::= s; s \mid x = \text{new } t \mid x = y \mid x.f = y \mid y = x.f \mid x = y.m(z)$	<i>statement</i>
$t ::= q C$	<i>qualified type</i>

■ **Figure 1** Syntax. C and D are class names, f is a field name, m is a method name, and x , y , and z are names of local variables, formal parameters, or parameter **this**. As in the code examples, **this** is explicit. Qualifiers q range over ReIm/ReM qualifiers (defined in Sect. 3).

A field write statement $x.f = y$ contributes a direct edge from y to the field access node $x.f$, and an *approximate* edge from $x.f$ to every $x'.f \in G$, where $x'.f$ is the right-hand-side of a field read $y' = x'.f$. (Without loss of generality we may assume $x' \neq x$.)

$$x.f = y \Rightarrow y \xrightarrow{d} x.f \xrightarrow{a} x'.f$$

We elaborate upon approximate edges shortly. A field read statement $y' = x'.f$ contributes direct edges as follows:

$$y' = x'.f \Rightarrow x' \xrightarrow{d} x'.f \xrightarrow{d} y'$$

Edge $x' \xrightarrow{d} x'.f$ accounts for deep (im)mutability. It links x' to y' , propagating mutability back to x' when y' is update.

Therefore, together a pair of field write $x.f = y$ and field read $y' = x'.f$ contribute a triple

$$x.f = y, y' = x'.f \Rightarrow y \xrightarrow{d} x.f \xrightarrow{a} x'.f \xrightarrow{d} y'$$

creating a path from y to y' . It models flow through heap objects while completely avoiding heap objects. In terms of Reps' terminology [24], our mutability semantics, like ReIm, models structure (i.e., heap)-transmitted dependences approximately.

The approximate edge makes approximation explicit. The approximate path from y to y' propagates mutability from y' back to y , but “with an asterisk”. This is *maybe* mutability — if $x.f$ and $x'.f$ are aliases because x and x' point to the same object, then y is truly mutable, however, if they are not aliases, y is not mutable due to this path. ReIm (and other reference immutability systems) overapproximate, and mark y mutable. A key insight of our work is that the impact of approximate paths is quite muted. Generally, when there is an approximate path from a reference y to an update, there is also a direct path from y to an update, and y would have become mutable regardless of the approximate path.

A method call (method entry) creates the expected *call* edges from actual arguments to formal parameters:

$$i: x = y.m(z) \Rightarrow y \xrightarrow{(i)} \text{this} \quad z \xrightarrow{(i)} p$$

Here **this** and p are the parameters of the compile-time target of the call. The standard CFL-reachability annotation (i) marks call entry at call site i . A method return (method exit) creates a *return* edge from the return value to the left-hand-side of the call assignment:

$$i: x = y.m(z) \Rightarrow \text{ret} \xrightarrow{)i} x$$

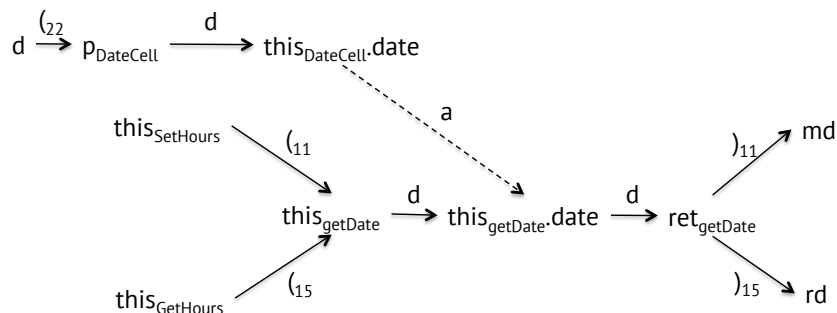
The standard CFL-reachability annotation $)i$ marks a return at site i . In terms of Reps' terminology, the semantics models call-transmitted dependences precisely.

Since the goal is to capture dependences between variables, the semantics eschews objects and object creation. Fig. 2 shows an example including all kinds of statements and their corresponding edges.

```

1 class DateCell {
2   Date date;
3
4   DateCell(DateCell this, Date p) {
5     this.date = p;
6   }
7   Date getDate(DateCell this) {
8     return this.date;
9   }
10  void cellSetHours(DateCell this) {
11    Date md = this.getDate();
12    md.setHours(1); // md is mutated
13  }
14  int cellGetHours(DateCell this) {
15    Date rd = this.getDate();
16    int hour = rd.getHours(); // rd is readonly
17    return hour;
18  }
19
20  public static void main(String[] args) {
21    Date d = new Date();
22    DateCell dc = new DateCell(d);
23    ...
24  }
25 }

```



■ **Figure 2** Running example. Code (adapted from Huang et al. [17]) and corresponding graph.

2.2 Paths in Flow Graph

We classify paths in G into two categories: (1) *exact* paths, which do not contain approximate edges, and (2) *approximate* paths, which contain approximate edges. In our running example in Fig. 2, $\text{this}_{\text{getDate}} \rightsquigarrow \text{rd}$ is an exact path, while $\text{d} \rightsquigarrow \text{md}$ is an approximate path. (We use squiggle arrows \rightsquigarrow to denote multi-edge paths.) Not all paths in G are well-formed, and different well-formed paths have different meaning.

2.2.1 Exact Paths

Fig. 3 defines a context-free grammar that classifies exact paths into 3 categories. This grammar is standard in CFL-reachability theory. There is an M -path from node n to node u if and only if the edge annotations on the path form a string in the language described by M . M -paths are paths with matched parentheses. For example, path $\text{this}_{\text{GetHours}} \rightsquigarrow \text{rd}$ is

XX:6 Definite Reference Mutability

$$\begin{aligned}
E & ::= R' \mid R \mid C \mid M \\
R' & ::= R C \\
R & ::=)_i \mid)_i M \mid)_i R \mid M R \\
C & ::= (_i \mid (_i M \mid (_i C \mid M C \\
M & ::= d \mid (_i M)_i \mid M M
\end{aligned}$$

■ **Figure 3** A context-free grammar for exact paths, i.e., paths that account (solely) for call-transmitted dependences. M captures matched-parentheses strings, e.g., $(_i d)_i$, C captures strings with one or more outstanding calls, e.g., $(_i d (_j d)_j$, and R and R' capture strings with one or more outstanding returns, e.g., $d)_j$.

an M -path. However, $\text{this}_{\text{GetHours}} \rightsquigarrow \text{md}$ is not a well-formed path because call edge $(_{15}$ and return edge $)_{11}$ do not match.

There is a C -path from n to u if and only if the edges from n to u form a string in the language described by C . More intuitively, these are paths with outstanding call edges. For example, $\text{this}_{\text{GetHours}} \rightsquigarrow \text{ret}_{\text{getDate}}$ is a C -path. With respect to reference immutability, if there is an M -path or a C -path from x (or from $x.f$) to an update, then x (or $x.f$) is definitely mutable, in the sense that analysis generally cannot improve from mutable. Because M -paths and C -paths have the same effect, from now on we refer to them as $M|C$ -paths.

The third category is R -paths. There is an R -path from n to u if and only if the edges form a string in R or R' . That is, the path starts with outgoing return edges, and it may or may not descend into a call path before reaching u . For example, $\text{this}_{\text{getDate}} \rightsquigarrow \text{md}$ is an R -path. With respect to reference immutability, if there is an R -path from x (or $x.f$) to an update, then x (or $x.f$) is *polymorphic*. It is mutable in some contexts of invocation of the enclosing method, and readonly in other. For example, $\text{this}_{\text{getDate}}$ and $\text{ret}_{\text{getDate}}$ are polymorphic. They are interpreted as mutable when getDate is called from cellSetHours , and they are interpreted as readonly when getDate is called from cellGetHours .

2.2.2 Approximate Paths

The following grammar rules capture approximate paths:

$$A ::= a \mid a E \mid a A \mid E A$$

There is an A -path from n to u if and only if the edges form a string in A . Pictorially, an A -path consists of exact paths and approximate edges. For example,

$$n \rightsquigarrow \cdot \xrightarrow{E} \cdot \xrightarrow{a} \cdot \rightsquigarrow \cdot \xrightarrow{E} \cdot \xrightarrow{a} \cdot \rightsquigarrow \cdot \xrightarrow{E} u$$

is an A -path, and so is

$$n \xrightarrow{a} \cdot \rightsquigarrow \cdot \xrightarrow{E} \cdot \xrightarrow{a} u$$

The only mandatory component of the A -path is the one approximate edge.

A -paths fall into two categories, $(M|C)A$ -paths, and RA -paths determined by the leading exact path:

- (1) if the leading exact path is an $M|C$ -path, then there is a $(M|C)A$ -path. For example,

$$d \xrightarrow{(20)} \text{pDateCell} \xrightarrow{d} \text{thisDateCell.date} \xrightarrow{a} \text{thisDateCell.date} \xrightarrow{d} \text{ret}_{\text{getDate}} \xrightarrow{9} \text{md}$$

is a $(M|C)A$ -path.

- (2) if the leading exact path is an R -path, then there is an RA -path. For the rest of the paper we use the term R -path to denote both the exact R -path and the RA -path as they have the same effect for our purposes.

Standard reference immutability type systems (e.g., ReIm), conservatively mark mutable every reference x , such that there is a $(M|C)A$ -path from x to an update. As we mentioned earlier, an approximate path introduces uncertainty rather than definite mutability. The key observation of our work is the following. The majority of references x that exhibit an A -path from x to an update, also exhibit a “parallel” M -path or C -path to a (potentially different) update. Therefore, x is indeed definitely mutable and the A -path has no ill impact; an analysis that attempts to handle A -paths, i.e., structure-transmitted dependences, more precisely would not do better regarding x . Roughly speaking, our analysis separates the A -paths that do exhibit a “parallel” path to an update, from the A -paths that do not, thus separating references that are definitely mutable, from ones that are *maybe mutable*. If an analysis that treats structure-transmitted dependences more precisely is to realize precision improvement, the improvement is bounded by the number of maybe mutable references.

3 Type Systems

This section presents two reference immutability type systems, Huang et al.’s [17] ReIm, and our novel proposal ReM. ReIm captures *definite reference immutability*, that is, **readonly** references in ReIm are guaranteed immutable, however, **mutable** references are not necessarily mutable. ReM captures definite immutability and definite mutability — in ReM **readonly** references are still guaranteed immutable, and in addition, **mutable** references are guaranteed mutable (in the CFL-reachability sense).

The reader may wonder why one needs type-based reference immutability like ReIm and Javari, when one has a clear semantics expressed in terms of standard CFL-reachability. First, type-based reference immutability is studied extensively in the literature [31, 36, 37, 17, 13, 21]; its connection to CFL-reachability brings new insights. Second, type-based reference immutability allows programmers to specify immutability requirements with type qualifiers, e.g., **readonly** x , and take advantage of systems such as JSR 308 and the Checker Framework (<https://checkerframework.org/>) to check these immutability requirements; such requirements cannot be easily expressed or checked using CFL-reachability. Third, type systems promote modularity, while CFL-based systems are typically whole-program analyses. Yet another advantage comes when reasoning about complexity. While CFL-reachability is $O(N^3)$, ReIm/ReM inference is $O(N^2)$, where N is the program size.

Sect. 3.1 outlines ReIm, largely following Huang et al. [17]. We add new interpretation in terms of our mutability semantics. Sect. 3.2 builds ReM upon the discussion in Sect. 3.1. Sect. 3.3 discusses type inference for ReIm and ReM.

3.1 ReIm

3.1.1 ReIm Qualifiers

The ReIm type system has three immutability qualifiers: **mutable**, **readonly**, and **poly**. We explain the qualifiers in terms of the mutability semantics defined in Sect. 2.

- **mutable**: A **mutable** reference x can be used to mutate the referenced object. This is the implicit and only option in standard object-oriented languages. In terms of our

XX:8 Definite Reference Mutability

mutability semantics, a **mutable** reference denotes an $M|C$ -path, or a $(M|C)A$ -path from x to an update.

- **readonly**: **readonly** captures “deep” immutability. A **readonly** reference x cannot be used to mutate the referenced object nor anything it references. All of the following are forbidden:
 - $x.f = y$
 - $x.set(z)$ where **set** sets a field of its receiver x
 - $z = id(x); z.f = w$
 - $y = x.f; y.g = z$

In terms of the mutability semantics, a **readonly** reference means that there does not exist either exact or approximate path to update.

- **poly**: This qualifier expresses polymorphism over immutability. **poly** denotes that a reference is interpreted as mutable in some contexts, and it is interpreted as immutable in other contexts. The enclosing method *does not* mutate the reference, however, mutation to the reference or one of its components may happen after return. In terms of the mutability semantics, a **poly** reference denotes that there is an R -path from x to an update — the reference “flows” out of its enclosing method where it is mutated in some caller context.

The subtyping relation between the qualifiers is

mutable <: **poly** <: **readonly**

where $q_1 <: q_2$ denotes q_1 is a subtype of q_2 . For example, it is allowed to assign a **mutable** reference to a **poly** or **readonly** one, but it is not allowed to assign a **readonly** reference to a **poly** or **mutable** one.

3.1.2 Typing Rules

ReIm is independent of the Java type system, which allows us to specify typing rules solely over type qualifiers q . The typing rules, following [17] are presented in Fig. 4. Rule (TASSIGN) is straightforward. It requires that the left-hand-side is a supertype of the right-hand-side. The system does not enforce object immutability and only **mutable** objects are created. The object creation rule becomes redundant and we omit it, just as we did in Sect. 2.

Rules (TREAD), (TWRITE) and (TCALL) make use of *viewpoint adaptation*, a concept from Universe Types [8, 9, 7]. Viewpoint adaptation of a type q' from the point of view of another type q , results in the adapted type q'' . This is written as $q \triangleright q' = q''$.

Below, we explain viewpoint adaptation in terms of the mutability semantics. At field accesses (TREAD) and (TWRITE) \triangleright adapts the field f from the viewpoint of (context of) the receiver. Viewpoint adaptation at field access handles structure-transmitted dependences, *approximately*. At method calls \triangleright adapts formal parameters and the return value from the point of view of the *variable at the left-hand-side* of the call assignment. This variable captures the calling context i . Viewpoint adaptation at calls handles call-transmitted dependences, *precisely*.

Importantly, ReIm restricts fields to **readonly** or **poly**. Javari [31] does allow for **mutable** fields, but this complicates the system. Declaring a field **mutable** in Javari excludes it from the state of the inclosing object, and adaptation of a **mutable** field requires special treatment, as discussed in [31, 17]. One can similarly allow **mutable** fields in ReIm/ReM. We are interested in type inference however, and allowing **mutable** fields would create ambiguity:

$$\begin{array}{c}
\text{(TASSIGN)} \\
\frac{\Gamma(x) = q_x \quad \Gamma(y) = q_y \quad q_y <: q_x}{\Gamma \vdash x = y} \\
\\
\text{(TWRITE)} \\
\frac{\Gamma(x) = q_x \quad q_x = \text{mutable} \quad \Gamma(y) = q_y \quad \text{typeof}(f) = q_f \quad q_y <: q_x \triangleright q_f}{\Gamma \vdash x.f = y} \\
\\
\text{(TREAD)} \\
\frac{\Gamma(x) = q_x \quad \Gamma(y) = q_y \quad \text{typeof}(f) = q_f \quad q_x \triangleright q_f <: q_y}{\Gamma \vdash y = x.f} \\
\\
\text{(TCALL)} \\
\frac{\Gamma(x) = q_x \quad \Gamma(y) = q_y \quad \Gamma(z) = q_z \quad \text{typeof}(m) = q_{\text{this}}, q_p \rightarrow q_{\text{ret}} \quad q_y <: q_x \triangleright q_{\text{this}} \quad q_z <: q_x \triangleright q_p \quad q_x \triangleright q_{\text{ret}} <: q_x}{\Gamma \vdash x = y.m(z)}
\end{array}$$

■ **Figure 4** Typing rules. Function *typeof* retrieves the qualifiers of fields and methods. Γ is a type environment that maps variables to their immutability qualifiers.

if $x.f$ is mutable, does this account for a mutable f , which is excluded from the state of x , or does it account for a deep mutation of x ? Restricting fields to $\{\text{readonly}, \text{poly}\}$ chooses the latter. Javarifier [22], Javari’s inference tool, makes the same choice. Javari is more expressive than ReIm, but its “inferable” semantics appears to be the same as ReIm’s: Huang et al. [17] report essentially identical inference result for Javarifier and ReIm.

Following [17], we define \triangleright as follows:

$$\begin{array}{lcl}
_ \triangleright \text{mutable} & = & \text{mutable} \\
_ \triangleright \text{readonly} & = & \text{readonly} \\
q \triangleright \text{poly} & = & q
\end{array}$$

The underscore denotes a “don’t care” value. Qualifiers **mutable** and **readonly** do not depend on the viewpoint. Qualifier **poly** depends on the viewpoint (context), and is substituted by that viewpoint (context).

Let us take a closer look at rules (TWRITE) and (TREAD). For a pair of field write $x.f = y$ and field read $y' = x'.f$, the rules entail the following constraints:

$$q_y <: q_x \triangleright q_f \quad q_{x'} \triangleright q_f <: q_{y'}$$

Suppose y' is an update, i.e., there is statement $y'.f = z$, and $q_{y'}$ is thus **mutable**. Therefore, q_f must be **poly**. First, recall that $q_f \in \{\text{readonly}, \text{poly}\}$. Since **readonly** adapts to **readonly**, $q_{x'} \triangleright q_f <: \text{mutable}$ does not type check, locking $q_f = \text{poly}$. (TWRITE) sets q_x , the type of the receiver to **mutable**. This serves two purposes in ReIm, (1) to account for the update of x ,

XX:10 Definite Reference Mutability

and (2) to account for the structure-transmitted dependence, i.e, the approximate path from y to the (eventual) update y' . Thus, $q_x \triangleright q_f$ evaluates to **mutable**, forcing q_y to be **mutable** as well. As mentioned earlier, ReIm handles approximate paths conservatively. If there is an $(M|C)A$ -path from a reference y to an update, then ReIm's rules force y to be **mutable**, as is the case above, even though x and x' may refer to different runtime objects.

Now, consider rule (TCALL). Function *typeof* retrieves the type of compile-time target m . q_{this} is the type of parameter **this**, q_p is the type of the formal parameter, and q_{ret} is the type of the return. Rule (TCALL) requires $q_x \triangleright q_{\text{ret}} <: q_x$, which accounts for R -paths. The constraint disallows the return value of m from being **readonly** when there is a call to m , $x = y.m(z)$, where left-hand-side x is **mutable**. Only if the left-hand-sides of all call assignments to m are **readonly**, can the return type of m be **readonly**; otherwise, it is **poly**. A programmer can annotate the return type of m as **mutable**. However, this typing is pointless, as it unnecessarily forces local variables and parameters in m to become **mutable** when they may remain less restrictively **poly**. In Fig. 2, `md = this.getDate();` entails constraint

$$q_{\text{md}} \triangleright q_{\text{ret_getDate}} <: q_{\text{md}} \quad \equiv \quad \text{mutable} \triangleright q_{\text{ret_getDate}} <: \text{mutable}$$

leading to $q_{\text{ret_getDate}} = \text{poly}$. This accounts for the R -path from **ret** to the update through `md`. Continuing with the example, the field read in `DateCell.getDate` (line 8) entails constraint

$$q_{\text{this_getDate}} \triangleright q_{\text{date}} <: q_{\text{ret_getDate}}$$

As a result, $q_{\text{this}} = \text{poly}$, which accounts for the R -path from `this.getDate` to update `md`.

In addition, rule (TCALL) requires $q_y <: q_x \triangleright q_{\text{this}}$. When q_{this} is **readonly** or **mutable**, its adapted value is the same. Thus, when q_{this} is **mutable** (e.g., due to `this.f = 0` in m),

$$q_y <: q_x \triangleright q_{\text{this}} \quad \text{becomes} \quad q_y <: \text{mutable}$$

which disallows q_y from being anything but **mutable**, as expected. This accounts for C - and CA -paths. The interesting case is when q_{this} is **poly**. A **poly** parameter **this** reflects a dependence between **this** and **ret** of m , such as the one in Fig. 2:

```
Date getDate(Date this) { ret = this.date; }
```

It allows the **this** object (or some part of it, in our example the **date** part of it), to be modified in caller context, after m 's return. The type system entails that whenever there is intraprocedural dependence between **this** and **ret**, we have

$$q_{\text{this}} <: q_{\text{ret}}$$

Recall that when there exists a context where the left-hand-side variable x is mutated, q_{ret} must be **poly**. Therefore, constraint $q_{\text{this}} <: q_{\text{ret}}$ forces q_{this} to be **poly** (assuming that **this** is not mutated in the context of its enclosing method). Rule (TCALL) adds the 2 constraints “around” $q_{\text{this}} <: q_{\text{ret}}$ to capture call-transmitted dependences:

$$q_y <: q_x \triangleright q_{\text{this}} \quad q_{\text{this}} <: q_{\text{ret}} \quad q_x \triangleright q_{\text{ret}} <: q_x$$

When m is called in a **mutable** context, i.e., q_x is **mutable**, q_y becomes **mutable**, as expected. Conversely, when m is called in a **readonly** context, i.e., q_x is **readonly**, $q_x \triangleright q_{\text{this}}$ evaluates to **readonly**, leaving q_y unchanged. In terms of our mutability semantics, this behavior captures M - and MA -paths.

3.2 ReM

We now present the ReM type system, which builds upon ReIm.

3.2.1 ReM Qualifiers

The ReM type system adds to the set of ReIm qualifiers, and changes the meaning of some of the ReIm qualifiers. There are 5 qualifiers in ReM: ReIm's **mutable**, **readonly** and **poly**, and two new, **maybe** and **polyormaybe**. Again, we interpret the qualifiers in terms of the mutability semantics defined in Sect. 2.

- **mutable**: A **mutable** reference x is now definitely mutable. It denotes that there is an $(M|C)$ -path from x to update.
- **readonly**: A **readonly** reference x has the same meaning as in ReIm, i.e., there is neither exact nor approximate path to update.
- **maybe**: The **maybe** qualifier is new. A **maybe** reference denotes that there is a $(M|C)A$ -path to update, but there is no R -path to update.
- **poly**: A **poly** reference now denotes that there is an R -path to update, but there is no $(M|C)A$ -path.
- **polyormaybe**: A **polyormaybe** reference denotes that there is an R -path to update, *and* a $(M|C)A$ -path to update.

The subtyping hierarchy is as follows:

$$\begin{array}{ccccccc} & & & & \text{maybe} & & \\ & & & & & & \\ \text{mutable} & <: & \text{polyormaybe} & <: & & <: & \text{readonly} \\ & & & & \text{poly} & & \end{array}$$

3.2.2 Typing Rules

ReM rules extend ReIm. There are two extensions: (1) viewpoint adaptation must account for new qualifiers **maybe** and **polyormaybe**, and (2) rule (TWRITE) must account for approximate paths.

Viewpoint adaptation rules from Sect. 3.1 remain in effect. We add two new rules:

$$\begin{array}{ll} _ \triangleright \text{maybe} & = \text{maybe} \\ q \triangleright \text{polyormaybe} & = (q \triangleright \text{poly}) \wedge \text{maybe} \end{array}$$

Notation \wedge stands for the standard meet operation: the result of $q_1 \wedge q_2$ is the greatest lower bound of q_1 and q_2 in the lattice of ReM types above.

Since field and return types are restricted to $\{\text{readonly}, \text{poly}\}$, adaptation of **maybe** or **polyormaybe** happens *only when adapting parameters at method calls*.

Recall that a **maybe** parameter p denotes a $(M|C)A$ -path from p . Thus, call $x = y.m(z)$ creates an $(M|C)A$ path from z (a CA -path to be precise). Rule (TCALL) requires

$$q_z <: q_x \triangleright q_p \quad \equiv \quad q_z <: \text{maybe}$$

which accounts for the $(M|C)A$ -path from z .

Now recall that a **polyormaybe** parameter p denotes an R -path to update, and an $(M|C)A$ -path to a (possibly different) update. These paths entail paths from z : one through ret , depending on the left-hand-side of the call assignment, and an $(M|C)A$ path. Rule (TCALL)

XX:12 Definite Reference Mutability

applies viewpoint adaptation of q_p , essentially recording the more conservative choice at the caller. Consider (TCALL) constraint

$$q_z <: q_x \triangleright q_p$$

Suppose that x is readonly, and there is no path from the left-hand-side of the call assignment x to update (that is, the R -path from p is due to a different call site). Then there is no new path from z to update through $p \rightsquigarrow \text{ret} \xrightarrow{)_i} x$. However, there is an $(M|C)A$ -path from z to update. Viewpoint adaptation accounts for this:

$$q_z <: (\text{readonly} \triangleright \text{poly}) \wedge \text{maybe} \equiv q_z <: \text{maybe}$$

Conversely, suppose x is mutable, and there is an $M|C$ -path from x to update. Then the R -path leads to an $M|C$ -path from z to update. Viewpoint adaptation accounts for this:


$$q_z <: (\text{mutable} \triangleright \text{poly}) \wedge \text{maybe} \equiv q_z <: \text{mutable}$$

Consider the more detailed example:

```

1  class A {
2    ...
3    C m(A this, B p) {
4      C c = this.f;
5      p.g = c; // Assume some g is mutated
6      return c;
7    }
8
9    public static void main(String[] args) {
10   ...
11   C c1 = a1.m(b);
12   c1.setField(); // c1 is mutated
13   ...
14   C c2 = a2.m(b);
15   int i = c2.getField(); // c2 is readonly
16   ....
17 }
18 }

```



The R -path from this_m to $c1$, entails $q_{\text{this}} <: \text{poly}$, while the $(M|C)A$ -path through $p.g$ entails $q_{\text{this}_m} <: \text{maybe}$. Thus, $q_{\text{this}_m} = \text{polyormaybe}$. At call 11:

$$q_{a1} <: q_{c1} \triangleright q_{\text{this}_m} \equiv q_{a1} <: q_{c1} \triangleright \text{polyormaybe}$$

Since q_{c1} is mutable, polyormaybe adapts to mutable, setting q_{a1} to mutable. On the other hand, at call 14 we have

$$q_{a2} <: q_{c2} \triangleright q_{\text{this}_m} \equiv q_{a2} <: q_{c2} \triangleright \text{polyormaybe} \equiv q_{a2} <: \text{maybe}$$

Thus, since q_{c2} is *readonly*, the *meet* is *maybe*, and q_{a2} is precisely *maybe*.

We now change rule (TWRITE) to account for approximate paths:

$$\frac{\begin{array}{c} \text{(TWRITE)} \\ \Gamma(x) = q_x \quad q_x = \text{mutable} \quad \Gamma(y) = q_y \\ \text{typeof}(f) = q_f \quad q_y <: \text{maybe} \triangleright q_f \end{array}}{\Gamma \vdash x.f = y}$$

q_x remains *mutable* to account for the direct update on x . However, instead of adapting by *mutable* context as in ReIm, we adapt by *maybe*. This reflects the approximate path, which ReIm conservatively made *mutable*. If q_f is *readonly*, then the *maybe*-mutability of x does not affect y . If q_f is *poly*, that reflects an update to some f , and $q_y <: \text{maybe} \triangleright \text{poly}$ propagates the “*maybe*” update to y .

3.3 Type Inference

Type inference for both ReIm and ReM proceeds as outlined in [15, 17] and earlier in [18, 30]. We present novel treatment in terms of dataflow frameworks, and include necessary extensions for ReM.

The inference operates on mappings from keys to values S . The keys in the mapping are (1) local variables and parameters, (2) fields, and (3) method returns. The values in the mapping are *sets* of type qualifiers. For instance, $S(x) = \{\text{poly}, \text{mutable}\}$ in ReIm means the type of reference x can be *poly* or *mutable*.

S is initialized as follows. $S(\text{ret}) = S(f) = \{\text{readonly}, \text{poly}\}$ for all return values ret and fields f . The rest of the variables are initialized to the universal set of qualifiers U , which is $\{\text{readonly}, \text{poly}, \text{mutable}\}$ in ReIm, and $\{\text{readonly}, \text{maybe}, \text{poly}, \text{polyormaybe}, \text{mutable}\}$ in ReM. For the rest of this paper we use U to refer to either ReIm or ReM. We denote the initial mapping by S_0 .

The inference iterates over all statements s in the program and removes qualifiers inconsistent with the typing rule for s from S . More precisely, let s consist of variables v_1, v_2, \dots, v_k and let s entail transfer function $c(s)$. Applying $c(s)$ removes each q_1 from $S(v_1)$ when there are no qualifiers $q_2 \in S(v_2), \dots, q_k \in S(v_k)$, such that q_1, q_2, \dots, q_k make s type check; then it removes all q_2 from $S(v_2)$, etc. For example, consider $s: y = x.f$ and corresponding rule (TREAD) triggering constraint $q_x \triangleright q_f <: q_y$. Let $S(x), S(f)$ and $S(y)$ be as follows:

$S(x)$	$S(f)$	$S(y)$
$\{\text{maybe}, \text{polyormaybe}, \text{mutable}\}$	$\{\text{readonly}, \text{poly}\}$	$\{\text{poly}, \text{polyormaybe}, \text{mutable}\}$

$c(s)$ removes *maybe* from $S(x)$ because there does not exist $q_f \in S(f)$ and $q_y \in S(y)$ that satisfy $\text{maybe} \triangleright q_f <: q_y$. Similarly, it removes *readonly* from $S(f)$. After application of $c(s)$:

$S'(x)$	$S'(f)$	$S'(y)$
$\{\text{polyormaybe}, \text{mutable}\}$	$\{\text{poly}\}$	$\{\text{poly}, \text{polyormaybe}, \text{mutable}\}$

One can easily prove that, given S_0 as shown, $c(s)$ only need look at the left-hand-side of the constraint, i.e., the right-hand-side always remains unchanged.

The inference analysis iterates over the statements in the program and removes qualifiers from the sets until it reaches a fixpoint. The problem fits into the standard monotone

XX:14 Definite Reference Mutability

dataflow framework [1, 19]. The lattice L_v for variables v is

$$\begin{array}{ccccccc} & & & & \{\text{maybe, polyormaybe, mutable}\} & & \\ \{\text{mutable}\} & > & \{\text{polyormaybe, mutable}\} & > & & > & U \\ & & & & \{\text{poly, polyormaybe, mutable}\} & & \end{array}$$

and the dataflow lattice L is the crossproduct of L_v , which is standard. Initializing all variables to U corresponds to initializing with the 0 of the lattice. The function space is $L \rightarrow L$ and is monotone. This is a theorem that one can easily show by case-by-case analysis of each $c(s)$. Therefore, the result of fixpoint iteration is the maximal fixpoint solution. (Strictly, it is the least fixpoint solution, however, the term maximal fixpoint solution is widely used in the literature.) Call this solution S_{Fix} . Yet the fixpoint solution is a mapping from references to sets. The actual mapping from references to types is derived as follows: for each reference x we pick the *maximal* element of $S_{Fix}(x)$ according to the following ranking, which mirrors the subtyping lattice:

$$\begin{array}{ccccccc} & & \text{maybe} & & & & \\ \text{readonly} & > & & > & \text{polyormaybe} & > & \text{mutable} \\ & & \text{poly} & & & & \end{array}$$

Importantly, the maximal element exists because each $S_{Fix}(x)$ is an element of L_v . We denote this typing by $max(S_{Fix})$, and call it the *maximal typing*.

The following propositions state that (1) the maximal typing type checks, and (2) the maximal typing is the “best typing”. (Note that setting all references to **mutable** also type checks, but makes up a useless typing.)

► **Proposition 1.** ReIm’s $max(S_{Fix})$ and ReM’s $max(S_{Fix})$ always type check.

Proof. (Sketch) The proof for ReIm is given in [17]. The proof for ReM proceeds by case-by-case analysis. The most difficult case arises at $q_z <: q_x \triangleright q_p$. Let $S_{Fix}(p) = \{\text{poly, polyormaybe, mutable}\}$ and $S_{Fix}(z) = \{\text{maybe, polyormaybe, mutable}\}$. $S_{Fix}(x)$ must be either $\{\text{maybe, polyormaybe, mutable}\}$ or $\{\text{readonly, maybe, poly, polyormaybe, mutable}\}$. If it were any other set, then **maybe** would have been removed from $S(z)$ during fixpoint iteration. Combinations $q_z = \text{maybe}$, $q_x = \text{maybe}$, $q_p = \text{poly}$, and $q_z = \text{maybe}$, $q_x = \text{readonly}$, $q_p = \text{poly}$, maximal typings under the two cases, both typecheck.

A more general statement is true. For every S that satisfies the equations of the dataflow framework, typing $max(S)$ type checks. ◀

Previous work in [15] formalized the notion of “best typing” for ownership type systems, specifically Ownership types [5] and Universe Types [8], by using a heuristic ranking over typings. This formalization applies to ReIm/ReM, as well as other ownership-like type systems, e.g., AJ [32] and EnerJ [25]. Below we extend the treatment of [15] to ReM.

We say that T is a *valid typing* if T type checks. Objective function o ranks valid typings. o takes a valid typing T and returns a tuple of numbers. For ReIm, o is as follows:

$$o_{ReIm}(T) = (|T^{-1}(\text{readonly})|, |T^{-1}(\text{poly})|, |T^{-1}(\text{mutable})|)$$

The tuples are ordered lexicographically. We have $T_1 > T_2$ iff T_1 has more **readonly** references than T_2 , or T_1 and T_2 have the same number of **readonly** references, but T_1 has more **poly** references than T_2 . The preference ranking over typings is based on ranking over qualifiers: naturally, we prefer **readonly** over **poly** and **mutable**, and **poly** over **mutable**.

<pre> 1: procedure CFL 2: $G = \emptyset$ 3: Add $u \xrightarrow{M} u$ to G for all updates u 4: while G changes do 5: for each s in Program do 6: EDGE($e(s)$) 7: end for 8: end while 9: end procedure 1: procedure EDGE($n \xrightarrow{t} n'$) 2: for each $n' \xrightarrow{N} u \in G$ do 3: Add $n \xrightarrow{t \oplus N} u$ to G 4: end for 5: end procedure </pre>	<pre> 1: procedure TYPES 2: $S(n) = U$ 3: $S(n) = \{\text{mutable}\}$ for all updates n 4: while S changes do 5: for each s in Program do 6: CONSTRAINT($c(s)$) 7: end for 8: end while 9: end procedure 1: procedure CONSTRAINT($l <: r$) 2: Remove each q_l from $S(l)$ 3: if $\nexists q_r \in S(r)$ s.t. $q_l <: q_r$ 4: end procedure </pre>
---	---

■ **Figure 5** Algorithm CFL initializes G , then iterates over program statements s adding edges as specified in Sect. 2.1. Algorithm TYPES initializes S , then iterates over program statements s removing qualifiers from S as specified in Sect. 3.3. The algorithms elide details to highlight the “parallel” structure of the two systems.

ReM’s objective function is the following:

$$o_{ReM}(T) = (|T^{-1}(\text{readonly})|, |T^{-1}(\text{poly})| + |T^{-1}(\text{maybe})|, |T^{-1}(\text{polyormaybe})|, |T^{-1}(\text{mutable})|)$$

Again the tuples are ordered based on the natural ranking over qualifiers: `readonly` is the most preferred, followed by `poly` and `maybe`, which are equally preferred, and so on. The following proposition establishes that the maximal typing is the best typing.

► **Proposition 2.** Let o be the objective function over valid typings (either o_{ReIm} over ReIm, or o_{ReM} over ReM). $o(\text{max}(S_{Fix})) > o(T)$ holds for every valid typing $T \neq \text{max}(S_{Fix})$.

Proof. (Sketch) The fact that the maximal typing is the “best typing”, follows from the properties of monotone dataflow frameworks. Let S be another solution of the dataflow framework. (It is easy to see that if T is valid typing, it must be contained into a solution of the dataflow framework.) Since $S > S_{Fix}$ by virtue of S_{Fix} being the maximal fixpoint solution, for every variable x $S(x) \geq S_{Fix}(x)$, and there exist variables y such that $S(y) > S_{Fix}(y)$. Thus, $o_{ReM}(\text{max}(S_{Fix})) > o_{ReM}(\text{max}(S))$. ◀

4 Equivalence

This section formally links the mutability semantics and the maximal typing. Specifically, we establish equivalence between CFL-reachability, as outlined in Sect. 2, and the maximal typing as outlined in Sect. 3.3 and previous work [15, 17, 18, 30].

Fig. 5 states the algorithms for CFL-reachability and type inference explicitly. CFL initializes graph G to \emptyset , then iterates over the program statements adding paths to updates, until no more paths can be added. TYPES initializes S to the 0 of the lattice, then iterates over the statements, removing qualifiers from S until no more qualifiers can be removed. To emphasize the parallel structure, Fig. 5 simplifies the presentation. Most notably, recall that according to Sect. 2 field read $y' = x'.f$ accounts for two edges:

$$y' = x'.f \quad \Rightarrow \quad x' \xrightarrow{d} x'.f \xrightarrow{d} y'$$

XX:16 Definite Reference Mutability

Even though Fig. 5 shows a single invocation of EDGE, in fact CFL processes two edges, first $x'.f \xrightarrow{d} y'$, followed by $x' \xrightarrow{d} x'.f$. Similarly, CFL processes multiple edges at field writes $x.f = y$: for each field read $y' = x'.f$, such that $x'.f \in G$, it processes $x.f \xrightarrow{a} x'.f$, followed by $y \xrightarrow{d} x.f$.

Another detail elided from Fig. 5 is the meaning of \mathfrak{t} , N and concatenation operator \oplus . \mathfrak{t} ranges over the terminals: $(i,)_i, \mathfrak{d}$, and \mathfrak{a} . N ranges over the kinds of paths: $M|C$, $(M|C)A$, and R . Concatenation $\mathfrak{t} \oplus N$ applies the grammar rules, and in all but one case, is straightforward:

$$\begin{aligned} (i \oplus M|C &= M|C \\ (i \oplus (M|C)A &= (M|C)A \\)_i \oplus _ &= R \\ \mathfrak{d} \oplus N &= N \\ \mathfrak{a} \oplus _ &= (M|C)A \end{aligned}$$

$(i \oplus R$ is the difficult case, because it applies rule $M ::= (i M)_i$. EDGE applies concatenation $(i \oplus R$ when processing call edge $z \xrightarrow{(i)} \mathfrak{p}$ and $\mathfrak{p} \xrightarrow{R} u \in G$. (Here \mathfrak{p} is the parameter of the compile-time target method of the call.) $\mathfrak{p} \xrightarrow{R} u \in G$ entails

$$\mathfrak{p} \xrightarrow{M} \text{ret} \xrightarrow{)j} x \xrightarrow{N} u$$

where ret is the return value of the target method, and x is some left-hand-side of a call assignment. Note that $\mathfrak{p} \xrightarrow{M} \text{ret} \xrightarrow{)j} x$ are not explicitly in G , but $x \xrightarrow{N} u$ is in G . There are two cases. If there does not exist j , such that $j = i$, then $(i \oplus R$ does not add new paths; the R -path from \mathfrak{p} is due to a different call site. Otherwise, that is, when $j = i$, concatenation adds $z \xrightarrow{N} u$ to G .

Let us illustrate CFL and EDGE, and TYPES and CONSTRAINT in parallel. Consider

$$y = x; z = y; z.f = w$$

Calling EDGE on $y \xrightarrow{d} z$, which corresponds to statement $z = y$, leads to path $y \xrightarrow{M} z$ in G . Subsequently calling EDGE on edge $x \xrightarrow{d} y$ leads to concatenation of \mathfrak{d} and M and path $x \xrightarrow{M} z$. There are two M -paths, x to z and y to z , as expected. Analogously, calling CONSTRAINT on $q_y <: q_z$, which corresponds to statement $z = y$, removes all qualifiers but `mutable` from $S(y)$. Subsequently calling CONSTRAINT on $q_x <: q_y$ removes all qualifiers but `mutable` from $S(x)$. $S(x) = \{\text{mutable}\}$, and $S(y) = \{\text{mutable}\}$ mirror the two M -paths that CFL finds.

The most interesting case arises (as it has been the case throughout the paper), when adding a call edge. Consider code

$$y = \text{id}(x); y = z; z.f = w$$

and assume $y \xrightarrow{M} z$ and $\mathfrak{p} \xrightarrow{R} z$ are already in G . Concatenation breaks the R -path $\mathfrak{p} \xrightarrow{R} z$ into $\mathfrak{p} \xrightarrow{M} \text{ret} \xrightarrow{)i} y \xrightarrow{M} z$. Since $(i$ and $)_i$ match, it adds path $x \xrightarrow{M} z$. Analogously, for TYPES assume $S(\mathfrak{p}) = \{\text{poly}, \text{polyormaybe}, \text{mutable}\}$ and $S(y) = \{\text{mutable}\}$. Calling CONSTRAINT on $q_x <: q_y \triangleright q_{\mathfrak{p}}$ removes all qualifiers but `mutable` from $S(x)$, which directly corresponds to the M -path from x to z that CFL finds.

To formally establish equivalence we use the bisimulation methodology for proving equivalence between two systems A and B [33, 6]. The methodology requires that we establish

a relation that relates states in A to those in B. In our case, A is constructed by CFL-reachability inference (Algorithm CFL), and B is constructed by type inference (Algorithm TYPES). Our approach defines an explicit equivalence relation between the states in A, captured by G , and those in B, captured by S . Intuitively, assume algorithms CFL and TYPES run in “parallel”. To show equivalence we must show that processing s in each system maintains equivalence.

We state two definitions that form the basis of equivalence. Informally, Def. 3 states that for every path from n to update u in G , n is correspondingly typed in S . In Def. 3 n stands for either a variable node x , or field access node $x.f$.

► **Definition 3.** (Soundness) $G \Rightarrow S$ if and only if

1. $n \xrightarrow{M|C} u \in G \Rightarrow \max(S(n)) <: \text{mutable}$
2. $n \xrightarrow{R} u \in G \Rightarrow \max(S(n)) <: \text{poly}$
3. $n \xrightarrow{(M|C)A} u \in G \Rightarrow \max(S(n)) <: \text{maybe}$

Def. 4 states that n 's maximal type in S implies a corresponding path in G . For example, maximal typing **polyormaybe** must imply that there are both an R -path and a $(M|C)A$ -path in G , but there is no $M|C$ -path.

► **Definition 4.** (Precision) $S \Rightarrow G$ if and only if

1. $\max(S(x)) = \text{mutable} \Rightarrow \exists x \xrightarrow{M|C} u \in G$
2. $\max(S(x)) = \text{polyormaybe} \Rightarrow \exists x \xrightarrow{R} u \in G \wedge \exists x \xrightarrow{(M|C)A} u \in G \wedge \nexists x \xrightarrow{M|C} u \in G$
3. $\max(S(x)) = \text{poly} \Rightarrow \exists x \xrightarrow{R} u \in G \wedge \nexists x \xrightarrow{(M|C)A} u \in G \wedge \nexists x \xrightarrow{M|C} u \in G$
4. $\max(S(x)) = \text{maybe} \Rightarrow \nexists x \xrightarrow{R} u \in G \wedge \exists x \xrightarrow{(M|C)A} u \in G \wedge \nexists x \xrightarrow{M|C} u \in G$
5. $\max(S(x)) = \text{readonly} \Rightarrow \text{no path from } x \text{ in } G$
6. $\max(S(x.f)) = \text{readonly} \Rightarrow \text{no path from } x.f \text{ in } G$

► **Definition 5.** (Equivalence) $G \simeq S$ if and only if $G \Rightarrow S$ and $S \Rightarrow G$.

Let the following Hoare triple denote parallel execution of EDGE and CONSTRAINT on statement s :

$$\{G, S\} \text{ EDGE}(e(s)) \parallel \text{ CONSTRAINT}(c(s)) \{G', S'\}$$

Our key result is the following theorem:

► **Theorem 6.** *If $G \simeq S$ and $\{G, S\} \text{ EDGE}(e(s)) \parallel \text{ CONSTRAINT}(c(s)) \{G', S'\}$ then $G' \simeq S'$.*

Proof. (Sketch) As expected, the proof is by induction on the number of applications of

$$\text{EDGE}(e(s)) \parallel \text{ CONSTRAINT}(c(s))$$

Clearly, the statement holds after initialization, lines 2-3 in CFL and lines 2-3 in TYPES. The inductive step requires case-by-case analysis of each s .

To prove correctness, we must show that given $G \Rightarrow S$, after the execution of $\text{EDGE}(e(s))$ and $\text{CONSTRAINT}(c(s))$, $G' \Rightarrow S'$ still holds. We outline the most difficult case, $\text{EDGE}(z \xrightarrow{(i)} p) \parallel \text{CONSTRAINT}(q_z <: q_x \triangleright q_p)$ (method call naturally breaks into three steps). Consider $z \xrightarrow{(i)} p \oplus p \xrightarrow{R} u$. Let x be the left-hand-side at call assignment i . If there does not exist

XX:18 Definite Reference Mutability

a path $x \xrightarrow{N} u \in G$, then no new paths are added to G' and $G' \Rightarrow S'$ holds. If there exists $x \xrightarrow{N} u \in G$, then a new path $z \xrightarrow{N} u$ is added to G' . We must show that $S'(z)$ reflects N according to Def. 3 (e.g., if N is an $M|C$ -path, then z is mutable). By the inductive hypothesis $p \xrightarrow{R} u \in G \Rightarrow \max(S(p)) <: \text{poly}$. Similarly, the N -path entails appropriate $S(x)$: if N is an $M|C$ path, then $S(x) = \{\text{mutable}\}$, if N is an R path then $\max(S(x)) <: \text{poly}$, and if N is a $(M|C)A$ -path, then $\max(S(x)) <: \text{maybe}$. Constraint $q_z <: q_x \triangleright q_p$ removes qualifiers from $S(z)$. For example, if N is $M|C$, then $S(x) = \{\text{mutable}\}$, and it is easy to see that $\max(S(x)) \triangleright \max(S(p)) <: \text{mutable}$. Thus $\max(S'(z)) <: \text{mutable}$, as needed. We enumerate all cases in Sect. A (Proofs).

In the other direction, we must show that if $S \Rightarrow G$ holds, after the execution of EDGE and CONSTRAINT on s , $S' \Rightarrow G'$ still holds. Consider the analogous case, $\text{EDGE}(z \xrightarrow{(i)} p) \parallel \text{CONSTRAINT}(q_z <: q_x \triangleright q_p)$, and the following values of $S(x), S(p)$ and $S(z)$ (only the maximal element shown):

$S(x)$	$S(p)$	$S(z)$
{ readonly, ... }	{ maybe, ... }	{ readonly, ... }

Constraint $q_z <: q_x \triangleright q_p$ “lowers” $S(z)$ into $S'(z) = \{\text{maybe, ...}\}$. By the inductive hypothesis, $S(x)$ and $S(p)$ entail that there are no paths from x in G , no paths from z in G , and only a $(M|C)A$ -path from p . Therefore, $\text{EDGE}(z \xrightarrow{(i)} p)$ adds an $(M|C)A$ -path from z in G' , and no other kind of path. Thus, $\max(S'(z)) = \text{maybe} \Rightarrow z \xrightarrow{(M|C)A} u$, as expected. We enumerate all cases in Sect. A (Proofs). ◀

For clarity, we omitted method overriding. It is handled in both the mutability semantics and type inference, and equivalence still holds. Concretely, if m' overrides m we add

$$q_{\text{this}_m} \xrightarrow{d} q_{\text{this}_{m'}} \quad q_{p_m} \xrightarrow{d} q_{p_{m'}} \quad q_{\text{ret}_{m'}} \xrightarrow{d} q_{\text{ret}_m}$$

to G . Analogously, we require

$$\text{typeof}(m') <: \text{typeof}(m)$$

which entails

$$(q_{\text{this}_{m'}}, q_{p_{m'}} \rightarrow q_{\text{ret}_{m'}}) <: (q_{\text{this}_m}, q_{p_m} \rightarrow q_{\text{ret}_m})$$

which leads to the standard function subtyping constraints:

$$q_{\text{this}_m} <: q_{\text{this}_{m'}} \quad q_{p_m} <: q_{p_{m'}} \quad q_{\text{ret}_{m'}} <: q_{\text{ret}_m}$$

Our implementation handles function subtyping.

5 Empirical Results

We implemented ReM on top of ReIm. (ReIm is publicly available.) Soot is the underlying platform, and Jimple is the underlying intermediate representation. We evaluate ReM on DaCapo, plus the benchmarks used in Javarifier [22] and ReIm [17]. There are 13 whole programs, and 8 libraries:

- The **Dacapo** suite DaCapo-2006-10MR.

Benchmark	#Readonly	Annotatable References		#Mutable	Time (sec.)
		#Poly	#Maybe/ #PolyOrMaybe		
antlr	15751	1029 (3.4%)	1198/2 (3.9%)	12552 (91.3%)	64.3
bloat	12567	2299 (6.1%)	3880/56 (10.5%)	18844 (82.7%)	36.1
chart	31640	4973 (6.5%)	7570/139 (10.1%)	32063 (80.6%)	79.3
eclipse	16879	4110 (11.3%)	2609/36 (7.3%)	12729 (82.8%)	39.1
fop	47411	5563 (5.9%)	5578/66 (6.0%)	35564 (86.3%)	89.0
hsqldb	31550	4001 (6.2%)	3018/31 (4.7%)	26186 (89.6%)	26.9
luindex	5465	686 (6.0%)	750/11 (4.7%)	4604 (85.8%)	30.1
lusearch	6296	945 (6.9%)	1019/16 (7.6%)	5344 (83.8%)	30.2
pmd	37758	4605 (5.6%)	5248/96 (6.5%)	34096 (86.5%)	78.4
xalan	21254	3337 (7.1%)	2853/63 (6.1%)	19771 (87.1%)	75.4
jolden	1208	281 (10.2%)	206/0 (7.5%)	1069 (83.8%)	29.0
javad	796	45 (3.5%)	15/0 (1.9%)	417 (96.5%)	26.5
ejc	29062	8202 (12.5%)	3612/174 (5.8%)	24768 (86.7%)	49.3
tinySQL	6115	665 (6.0%)	638/7 (5.9%)	3585 (84.8%)	28.1
htmlparser	7787	1215 (8.1%)	1049/12 (7.1%)	4904 (82.2%)	29.2
commons-pool	847	25 (1.4%)	222/0 (12.8%)	640 (74.2%)	27.1
jdbm	1134	317 (9.4%)	322/9 (9.8%)	1584 (82.7%)	28.8
jdbf	4017	312 (4.5%)	289/2 (4.2%)	2379 (89.1%)	58.6
jtds	8628	825 (5.2%)	587/25 (3.9%)	5738 (90.4%)	32.5
java.lang	3336	350 (4.0%)	249/1 (3.3%)	4721 (95.0%)	29.5
java.util	3216	376 (4.3%)	309/0 (3.8%)	4758 (93.9%)	30.9
Average		6.4%	6.4%	86.5%	

■ **Table 1** Inference results for ReM. **Annotatable References** includes all variables of reference type, including locals, parameters, returns, and fields. It does not include variables of primitive type. **#Readonly** shows the number of references inferred as **readonly**. **#Poly** shows the number of variables inferred as **poly**, including, in parentheses, their percentage of the total number of references. **#Maybe/#PolyOrMaybe** shows the number of **maybe** and **polyor-maybe** references, respectively, and their percentage of the total number of references. **#Mutable** shows the number of **mutable** references, which are now definitely mutable. In parentheses is the percentage of definitely mutable references, out of all potentially mutable ones: **#Mutable/(#Maybe+#PolyOrMaybe+#Mutable)**.

- **JOlden** is a classical suite of 10 small whole programs (Javarifier and ReIm).
- **ejc-3.2.0** is the Java Compiler for the Eclipse IDE (Javarifier and ReIm).
- **javad** is a Java disassembler program (ReIm).
- **tinySQL-1.1** is a database engine (Javarifier and ReIm).
- **htmlparser-1.4** is an HTML parser library (Javarifier and ReIm).
- **commons-pool-1.2** is an object pooling library (ReIm).
- **jtds-1.0** is a JDBC driver (ReIm).
- **jdbm-1.0** is a transactional engine (ReIm).
- **jdbf-0.0.1** is an object-oriented mapping system (ReIm).
- **java.lang** and **java.util** are the packages from JDK 1.7.0_75.

All benchmarks are analyzed with JDK 1.7. On whole programs, our analysis performs standard CHA-reachability with all packages included in the analysis. All experiments are done on a MacBook Pro 2.8 GHz Intel Core i7 and 16GB of RAM.

Benchmark	#Readonly	Parameters, Returns and Fields			Time (sec.)
		#Poly	#Maybe/ #PolyOrMaybe	#Mutable	
antlr	2821	647 (11.2%)	334/0 (5.8%)	1972 (85.5%)	60.9
bloat	2803	1746 (19.5%)	783/5 (8.8%)	3637 (82.2%)	34.7
chart	9464	3529 (15.8%)	2013/34 (9.2%)	7229 (77.9%)	79.6
eclipse	4722	3235 (28.5%)	685/6 (6.1%)	2697 (79.6%)	35.7
fop	13695	4052 (15.1%)	1617/14 (6.1%)	7528 (82.2%)	84.6
hsqldb	7663	3164 (20.2%)	662/5 (4.3%)	4147 (86.1%)	25.5
luindex	1726	555 (16.0%)	199/1 (5.8%)	987 (83.2%)	31.6
lusearch	1949	749 (17.7%)	281/2 (6.7%)	1260 (81.7%)	31.7
pmd	11085	3300 (13.7%)	1480/15 (6.2%)	8256 (84.7%)	74.1
xalan	6849	2375 (16.4%)	788/15 (5.5%)	4448 (84.7%)	65.5
jolden	375	215 (27.2%)	42/0 (5.3%)	159 (79.1%)	27.0
javad	272	44 (11.2%)	0/0 (0.0%)	78 (100%)	26.9
ejc	6208	7033 (38.2%)	603/15 (3.4%)	4544 (88.0%)	47.3
tinySQL	2267	470 (13.2%)	152/0 (4.3%)	683 (81.8%)	57.9
htmlparser	2074	1067 (24.8%)	191/1 (4.5%)	975 (83.5%)	59.3
commons-pool	229	25 (4.5%)	101/0 (18.0%)	205 (67.0%)	25.6
jdbm	331	247 (26.9%)	94/0 (10.2%)	247 (77.8%)	27.6
jdbf	1392	239 (11.4%)	33/0 (1.6%)	433 (92.9%)	57.5
jtids	2280	636 (16%)	138/6 (3.6%)	922 (86.5%)	31.9
java.lang	1228	247 (10.6%)	57/0 (2.4%)	806 (93.4%)	26.7
java.util	1143	272 (11.6%)	74/0 (3.2%)	851 (92.0%)	30.0
Average		17.6%	5.8%	84.1%	

■ **Table 2 Parameters, returns and fields** includes all parameters, returns and fields of reference type. It does not include local variables.

Tab. 1 presents the results of running ReM inference on the benchmarks. On average, only 6.4% of all references are inferred as *maybe* or *polyormaybe*. They make up only about 13.5% of all potentially mutable references; nearly 87% of all potentially mutable references are *definitely mutable*. Tab. 2 presents results of another run, reporting on parameters and returns only. We include these results to assess the impact of the intermediate representation (Soot’s Jimple). The results show that 5.8% of all references are *maybe* or *polyormaybe*, and about 84% of all potentially mutable parameters are in fact *definitely mutable*. These are similar results, suggesting that the intermediate representation does not lead to overestimate of the number of *definitely mutable* references. (In fact, our investigation suggests that it leads to underestimate, as we explain shortly.) Running times do not exceed 90 seconds, with most benchmarks completing in under 60 seconds on a commodity laptop.

The results demonstrate that ReM, ReIm and Javari, are *precise* and *scalable*. They can be used to power inference for approximate computing (e.g. EnerJ [25] and Rely [4]), taint analysis (e.g., DroidInfer [16]), and method purity [17], as well as other client analyses. Even if one designed a more complex system that handled structure-transmitted dependences more precisely, improvement would lead to at most 5-6% of all references (12-13% of ReIm’s *mutable* references), being promoted from *mutable*. ReM/ReIm’s complexity is $O(N^2)$ which leads to fast running times.

Finally, to better understand the results, we examined all 15 *maybe* references from javad, and 15 randomly selected *maybe* references from ejc. We looked to identify definite paths to mutation, or more precisely, we examined $y \xrightarrow{d} x.f \xrightarrow{a} x'.f \xrightarrow{d} y'$ and attempted to prove

that there exist a runtime object o such that x points to o , x' points to o , and the value of $x.f$ indeed flows to $x'.f$. We immediately identified such definite paths in 16 out of 30 cases. The remaining 14 cases exhibited difficult data and control flow, and we could not identify definite paths. A typical case of obvious definite paths was the following. Consider this typical code for initializing an array field f :

```

1  f = new X[10];
2  for (int i=0; i<cnt; i++)
3    f[i] = new X();
4  ...

```

This code snippet is translated into the following Jimple:

```

1  r1 = newarray (X)[10];
2  this.f = r1;
3  ...
4  r2 = this.f;
5  x = new X();
6  r2[i] = x;

```

Mutation of the array, which is definite, is captured by the approximate path, and the **maybe** typing of $r1$. We conjecture that this case leads to overestimate of the number of **maybe** variables. A simple optimization (which we have chosen not to implement to preserve the simplicity and clarity of ReM) would reduce the number of **maybe/polyormaybe** references. Another source of **maybe** mutability is containers. E.g., in

```

1  class Container {
2    Data data;
3    void set(Container this, Data p) {
4      this.data = p;
5    }
6    Data get(Container this) {
7      return this.data;
8    }
9  }

```

parameter p of `set` is rightfully **maybe** mutable. p and the `data` object will be mutable in some clients of `Container` and **readonly** in others.

6 Related Work

The most closely related work is Huang et al.'s ReIm and ReImInfer [17]. Our work builds upon ReIm and ReImInfer but extends them in two directions. First, we build a theoretical framework that interprets ReIm and ReM in terms of CFL-reachability, and we prove them correct within this simple framework. To the best of our knowledge, ReIm has not been proven correct even though it has been used to power client analyses [16, 29]. Second, we propose ReM and definite mutability, which extends the expressive power of ReIm, and also, establishes a bound on (im)precision. Our empirical results show that ReIm/ReM are highly precise and highly scalable.

Reference immutability has been an active area of research for many years. Tschantz et al. propose Javari [31] and Javari's inference tool Javarifier [22]. Javari is more expressive

and more complex than ReIm, but the inferable features are essentially the same. (Huang et al. report that Javarifier and ReImInfer produce essentially the same result.) Zibin et al.’s IGJ [36] and OIGJ [37] are type systems that support reference immutability and object immutability. Haack and Poll [14] propose a type system for object immutability as well. Gordon et al. [13] propose a reference immutability system for safe parallelism. Potanin et al. [21] survey work on reference and object immutability, and method purity. As it is standard, these systems support definite immutability (like ReIm). They do not attempt to estimate precision (imprecision), or connect reference immutability and CFL-reachability.

Artzi et al. [2] propose a hybrid static and dynamic analysis for inference of parameter reference immutability. In contrast, our work focuses on static analysis.

Salcianu and Rinard’s JPPA [28] and Pierce’s JPure [20] infer method purity for Java. ReIm/ReM is more general, in the sense that it enables reasoning about method purity, as well as other client analyses (e.g., EnerJ [25] and DroidInfer [16]). The fact that ReIm/ReM is precise, suggests that client analyses would be precise as well.

CFL-reachability is a standard program analysis framework [24]. Rehof and Fahndrich [23] connect type-based flow analysis and CFL-reachability. This is similar to our interpretation of type-based reference immutability in terms of CFL-reachability. However, Rehof and Fahndrich do not discuss mutable references and it is unclear how they handle such references or structure-transmitted dependences. Fahndrich et al. [11] apply the theory of [23] to build a context-sensitive Steensgard-style points-to analysis for C, thus using equality constraints instead of subtyping constraints. (Equality constraints is the standard approach to the handling of mutable references [26, 25, 12], as we mention earlier.) Our work focuses specifically on reference immutability and reasoning about its precision. The result that ReIm/ReM is precise, indicates that they can be incorporated into flow analyses [26, 25, 12, 16].

Sridharan and Bodik [27] present refinement-based points-to analysis for Java using CFL-reachability. Xu et al. [34] improve the scalability of CFL-reachability-based points-to analysis. These works focus on points-to analysis and require heap abstraction. Therefore, they inherit known issues with reflection. Type-based reference immutability and the parallel CFL-reachability analysis avoid heap abstraction and thus, they completely avoid issues due to reflective object creation (`x = Class.forName("className").newInstance()`), for free. We still face issues with reflective method invocation (`getMethod`). However, reflective object creation is by far most common, and has been studied extensively in the points-to analysis community. Recent work by Zhang and Su [35] propose new approximation algorithms based on CFL-reachability that can handle both structure-transmitted and call transmitted dependences precisely. Our work focuses on type-based reference immutability, for which handling of structure-transmitted dependences approximately appears sufficient.

7 Conclusion

We presented ReM, a novel reference immutability type system. ReM separated potentially mutable references into definitely mutable, and maybe mutable, i.e., references that may be mutable due to inherent approximation. In addition, we proposed a CFL-reachability system for reference immutability, thus building a novel framework for reasoning about correctness of reference immutability type systems. We implemented ReM and showed that about 86.5% of all potentially mutable references were definitely mutable.

References

- 1 Alfred V. Aho, Monica S. Lam, Ravi Sethi, and Jeffrey D. Ullman. *Compilers: Principles, Techniques, and Tools (2nd Edition)*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 2006.
- 2 Shay Artzi, Adam Kiezun, David Glasser, and Michael D. Ernst. Combined static and dynamic mutability analysis. In *Proceedings of the Twenty-second IEEE/ACM International Conference on Automated Software Engineering, ASE '07*, pages 104–113, New York, NY, USA, 2007. ACM. URL: <http://doi.acm.org/10.1145/1321631.1321649>, doi:10.1145/1321631.1321649.
- 3 Joseph A. Bank, Andrew C. Myers, and Barbara Liskov. Parameterized types for java. In *Proceedings of the 24th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '97*, pages 132–145, New York, NY, USA, 1997. ACM. URL: <http://doi.acm.org/10.1145/263699.263714>, doi:10.1145/263699.263714.
- 4 Michael Carbin, Sasa Misailovic, and Martin C. Rinard. Verifying quantitative reliability for programs that execute on unreliable hardware. In *Proceedings of the 2013 ACM SIGPLAN International Conference on Object Oriented Programming Systems Languages & Applications, OOPSLA 2013, part of SPLASH 2013, Indianapolis, IN, USA, October 26-31, 2013*, pages 33–52, 2013.
- 5 David G. Clarke, John M. Potter, and James Noble. Ownership types for flexible alias protection. In *Proceedings of the 13th ACM SIGPLAN Conference on Object-oriented Programming, Systems, Languages, and Applications, OOPSLA '98*, pages 48–64, New York, NY, USA, 1998. ACM. URL: <http://doi.acm.org/10.1145/286936.286947>, doi:10.1145/286936.286947.
- 6 Rance Cleaveland and Matthew Hennessy. Testing equivalence as a bisimulation equivalence. In *Automatic Verification Methods for Finite State Systems, International Workshop, Grenoble, France, June 12-14, 1989, Proceedings*, pages 11–23, 1989. URL: https://doi.org/10.1007/3-540-52148-8_2, doi:10.1007/3-540-52148-8_2.
- 7 Dave Cunningham, Werner Dietl, Sophia Drossopoulou, Adrian Francalanza, Peter Müller, and Alexander J. Summers. Formal methods for components and objects. chapter Universe Types for Topology and Encapsulation, pages 72–112. Springer-Verlag, Berlin, Heidelberg, 2008. URL: http://dx.doi.org/10.1007/978-3-540-92188-2_4, doi:10.1007/978-3-540-92188-2_4.
- 8 Werner Dietl and Peter Müller. Universes: Lightweight ownership for JML. *Journal of Object Technology*, 4(8):5–32, 2005.
- 9 Werner Dietl and Peter Müller. Runtime universe type inference. In *IWACO*, pages 72–80, 2007. URL: http://sct.ethz.ch/projects/student_docs/Frank_Lyner/Frank_Lyner_MA_paper.pdf.
- 10 Julian Dolby, Christian Hammer, Daniel Marino, Frank Tip, Mandana Vaziri, and Jan Vitek. A data-centric approach to synchronization. *ACM Transactions on Programming Languages and Systems*, 34(1):1–48, April 2012.
- 11 Manuel Fähndrich, Jakob Rehof, and Manuvir Das. Scalable context-sensitive flow analysis using instantiation constraints. In *Proceedings of the ACM SIGPLAN 2000 Conference on Programming Language Design and Implementation, PLDI '00*, pages 253–263, New York, NY, USA, 2000. ACM. URL: <http://doi.acm.org/10.1145/349299.349332>, doi:10.1145/349299.349332.
- 12 Robert Fuhrer, Frank Tip, Adam Kiezun, Julian Dolby, and Markus Keller. Efficiently refactoring java applications to use generic libraries. In *Proceedings of the 19th European Conference on Object-Oriented Programming, ECOOP'05*, pages 71–96, Berlin, Heidelberg, 2005. Springer-Verlag. URL: http://dx.doi.org/10.1007/11531142_4, doi:10.1007/11531142_4.

- 13 Colin S. Gordon, Matthew J. Parkinson, Jared Parsons, Aleks Bromfield, and Joe Duffy. Uniqueness and reference immutability for safe parallelism. In *Proceedings of the ACM International Conference on Object Oriented Programming Systems Languages and Applications*, OOPSLA '12, pages 21–40, New York, NY, USA, 2012. ACM. URL: <http://doi.acm.org/10.1145/2384616.2384619>, doi:10.1145/2384616.2384619.
- 14 Christian Haack and Erik Poll. Type-based object immutability with flexible initialization. In *Proceedings of the 23rd European Conference on ECOOP 2009 — Object-Oriented Programming*, Genoa, pages 520–545, Berlin, Heidelberg, 2009. Springer-Verlag. URL: http://dx.doi.org/10.1007/978-3-642-03013-0_24, doi:10.1007/978-3-642-03013-0_24.
- 15 Wei Huang, Werner Dietl, Ana Milanova, and Michael D. Ernst. Inference and checking of object ownership. In *Proceedings of the 26th European Conference on Object-Oriented Programming*, ECOOP'12, pages 181–206, Berlin, Heidelberg, 2012. Springer-Verlag. URL: http://dx.doi.org/10.1007/978-3-642-31057-7_9, doi:10.1007/978-3-642-31057-7_9.
- 16 Wei Huang, Yao Dong, Ana Milanova, and Julian Dolby. Scalable and precise taint analysis for android. In *Proceedings of the 2015 International Symposium on Software Testing and Analysis*, ISSTA 2015, pages 106–117, New York, NY, USA, 2015. ACM. URL: <http://doi.acm.org/10.1145/2771783.2771803>, doi:10.1145/2771783.2771803.
- 17 Wei Huang, Ana Milanova, Werner Dietl, and Michael D. Ernst. Reim & reiminfer: Checking and inference of reference immutability and method purity. In *Proceedings of the ACM International Conference on Object Oriented Programming Systems Languages and Applications*, OOPSLA '12, pages 879–896, New York, NY, USA, 2012. ACM. URL: <http://doi.acm.org/10.1145/2384616.2384680>, doi:10.1145/2384616.2384680.
- 18 Adam Kiezun, Michael D. Ernst, Frank Tip, and Robert M. Fuhrer. Refactoring for parameterizing java classes. In *Proceedings of the 29th International Conference on Software Engineering*, ICSE '07, pages 437–446, Washington, DC, USA, 2007. IEEE Computer Society. URL: <http://dx.doi.org/10.1109/ICSE.2007.70>, doi:10.1109/ICSE.2007.70.
- 19 Flemming Nielson, Hanne Riis Nielson, and Chris Hankin. *Principles of program analysis*. Springer-Verlag New York, Inc., 1999.
- 20 David J. Pearce. Jpure:: A modular purity system for java. In *Proceedings of the 20th International Conference on Compiler Construction: Part of the Joint European Conferences on Theory and Practice of Software*, CC'11/ETAPS'11, pages 104–123, Berlin, Heidelberg, 2011. Springer-Verlag. URL: <http://dl.acm.org/citation.cfm?id=1987237.1987247>.
- 21 Alex Potanin, Johan Östlund, Yoav Zibin, and Michael D. Ernst. Immutability. In *Aliasing in Object-Oriented Programming*, volume 7850 of *LNCS*, pages 233–269. Springer-Verlag, April 2013.
- 22 Jaime Quinonez, Matthew S. Tschantz, and Michael D. Ernst. Inference of reference immutability. In *Proceedings of the 22Nd European Conference on Object-Oriented Programming*, ECOOP '08, pages 616–641, Berlin, Heidelberg, 2008. Springer-Verlag. URL: http://dx.doi.org/10.1007/978-3-540-70592-5_26, doi:10.1007/978-3-540-70592-5_26.
- 23 Jakob Rehof and Manuel Fähndrich. Type-base flow analysis: From polymorphic subtyping to cfl-reachability. In *Proceedings of the 28th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '01, pages 54–66, New York, NY, USA, 2001. ACM. URL: <http://doi.acm.org/10.1145/360204.360208>, doi:10.1145/360204.360208.
- 24 Thomas Reps. Undecidability of context-sensitive data-independence analysis. *ACM Transactions on Programming Languages and Systems*, 22(1):162—186, 2000.
- 25 Adrian Sampson, Werner Dietl, Emily Fortuna, Danushen Gnanapragasam, Luis Ceze, and Dan Grossman. Enerj: approximate data types for safe and general low-power computation.

- In *Proceedings of the 32nd ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2011, San Jose, CA, USA, June 4-8, 2011*, pages 164–174, 2011.
- 26 Umesh Shankar, Kunal Talwar, Jeffrey S. Foster, and David Wagner. Detecting format string vulnerabilities with type qualifiers. In *Proceedings of the 10th Conference on USENIX Security Symposium - Volume 10, SSYM'01*, pages 16–16, Berkeley, CA, USA, 2001. USENIX Association. URL: <http://dl.acm.org/citation.cfm?id=1267612.1267628>.
 - 27 Manu Sridharan and Rastislav Bodík. Refinement-based context-sensitive points-to analysis for java. In *Proceedings of the 27th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '06*, pages 387–400, New York, NY, USA, 2006. ACM. URL: <http://doi.acm.org/10.1145/1133981.1134027>, doi:10.1145/1133981.1134027.
 - 28 Alexandru Sălciuanu and Martin Rinard. Purity and side effect analysis for java programs. In *Proceedings of the 6th International Conference on Verification, Model Checking, and Abstract Interpretation, VMCAI'05*, pages 199–215, Berlin, Heidelberg, 2005. Springer-Verlag. URL: http://dx.doi.org/10.1007/978-3-540-30579-8_14, doi:10.1007/978-3-540-30579-8_14.
 - 29 Rishi Surendran and Vivek Sarkar. Automatic parallelization of pure method calls via conditional future synthesis. In *Proceedings of the 2016 ACM SIGPLAN International Conference on Object-Oriented Programming, Systems, Languages, and Applications, OOPSLA 2016*, pages 20–38, New York, NY, USA, 2016. ACM. URL: <http://doi.acm.org/10.1145/2983990.2984035>, doi:10.1145/2983990.2984035.
 - 30 Frank Tip, Robert M. Fuhrer, Adam Kiezun, Michael D. Ernst, Ittai Balaban, and Bjorn De Sutter. Refactoring using type constraints. *ACM Transactions on Programming Languages and Systems*, 33(3):1–47, April 2011. doi:10.1145/1961204.1961205.
 - 31 Matthew S. Tschantz and Michael D. Ernst. Javari: Adding reference immutability to java. In *Proceedings of the 20th Annual ACM SIGPLAN Conference on Object-oriented Programming, Systems, Languages, and Applications, OOPSLA '05*, pages 211–230, New York, NY, USA, 2005. ACM. URL: <http://doi.acm.org/10.1145/1094811.1094828>, doi:10.1145/1094811.1094828.
 - 32 Mandana Vaziri, Frank Tip, Julian Dolby, Christian Hammer, and Jan Vitek. A type system for data-centric synchronization. In *Proceedings of the 24th European Conference on Object-oriented Programming, ECOOP'10*, pages 304–328, Berlin, Heidelberg, 2010. Springer-Verlag. URL: <http://dl.acm.org/citation.cfm?id=1883978.1884000>.
 - 33 Yuepeng Wang, Isil Dillig, Shuvendu K. Lahiri, and William R. Cook. Verifying equivalence of database-driven applications. *Proc. ACM Program. Lang.*, 2(POPL):56:1–56:29, December 2017. URL: <http://doi.acm.org/10.1145/3158144>, doi:10.1145/3158144.
 - 34 Guoqing Xu, Atanas Rountev, and Manu Sridharan. Scaling cfl-reachability-based points-to analysis using context-sensitive must-not-alias analysis. In *Proceedings of the 23rd European Conference on ECOOP 2009 — Object-Oriented Programming*, Genoa, pages 98–122, Berlin, Heidelberg, 2009. Springer-Verlag. URL: http://dx.doi.org/10.1007/978-3-642-03013-0_6, doi:10.1007/978-3-642-03013-0_6.
 - 35 Qirun Zhang and Zhendong Su. Context-sensitive data-dependence analysis via linear conjunctive language reachability. In *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2017*, pages 344–358, New York, NY, USA, 2017. ACM. URL: <http://doi.acm.org/10.1145/3009837.3009848>, doi:10.1145/3009837.3009848.
 - 36 Yoav Zibin, Alex Potanin, Mahmood Ali, Shay Artzi, Adam Kie, un, and Michael D. Ernst. Object and reference immutability using java generics. In *Proceedings of the the 6th Joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT*

XX:26 Definite Reference Mutability

Symposium on The Foundations of Software Engineering, ESEC-FSE '07, pages 75–84, New York, NY, USA, 2007. ACM. URL: <http://doi.acm.org/10.1145/1287624.1287637>, doi:10.1145/1287624.1287637.

- 37 Yoav Zibin, Alex Potanin, Paley Li, Mahmood Ali, and Michael D. Ernst. Ownership and immutability in generic java. In *Proceedings of the ACM International Conference on Object Oriented Programming Systems Languages and Applications*, OOPSLA '10, pages 598–617, New York, NY, USA, 2010. ACM. URL: <http://doi.acm.org/10.1145/1869459.1869509>, doi:10.1145/1869459.1869509.

A Poofs

Our main theorem follows from the following two lemmas.

► **Lemma 7.** *If $G \Rightarrow S$ and $\{G, S\} \text{ EDGE}(e(s)) \parallel \text{CONSTRAINT}(c(s)) \{G', S'\}$ then $G' \Rightarrow S'$.*

Proof. The proof relies on the fact that viewpoint adaptation preserves subtyping. That is, for each x, x' and $p, x <: x' \Rightarrow x \triangleright p <: x' \triangleright p$. Also, for each x, p and $p', p <: p' \Rightarrow x \triangleright p <: x \triangleright p'$. Therefore, for each $x, x', p,$ and $p', x <: x' \wedge p <: p' \Rightarrow x \triangleright p <: x' \triangleright p'$.

We proceed by induction on the number of applications of

$\text{EDGE}(e(s)) \parallel \text{CONSTRAINT}(c(s))$

and case by case analysis.

Consider the most difficult case, case 1: s is $x = y.m(z)$. Naturally, it breaks into 3 smaller cases, (1) $z \xrightarrow{(i)} p$, (2) $y \xrightarrow{(i)}$ this and (3) $\text{ret} \xrightarrow{(i)} x$. (2) is analogous to (1).

For (1), suppose EDGE adds $z \xrightarrow{(i \oplus M|C)} u$ to G' . By the inductive hypothesis, $p \xrightarrow{M|C} u$ implies $\text{max}(S(p)) <: \text{mutable}$. The corresponding constraint $q_z <: q_x \triangleright q_p$ sets $\text{max}(S'(z)) = \text{mutable}$. Similarly, $p \xrightarrow{(M|C)A} u \Rightarrow \text{max}(S(p)) <: \text{maybe}$ and constraint $q_z <: q_x \triangleright q_p$ leads to $q_z <: q_x \triangleright \text{maybe} <: \text{maybe}$ by the above theorem. Thus, $\text{max}(S'(z)) <: \text{maybe}$, as needed. Now, suppose that (1) adds $z \xrightarrow{(i \oplus R)} u$ to G' . This entails $\text{max}(S(\text{ret})) <: \text{poly}$. If $x \xrightarrow{M|C} u \in G$ then $\text{max}(S(x)) <: \text{mutable}$, leading to $\text{max}(S'(z)) <: \text{mutable}$. If $x \xrightarrow{(M|C)A} u \in G$ then $\text{max}(S(x)) <: \text{maybe}$, leading to $\text{max}(S'(z)) = \text{maybe}$, as needed. If $x \xrightarrow{R} u \in G$ then $\text{max}(S(x)) <: \text{poly}$; constraint $q_z <: q_x \triangleright q_p \equiv q_z <: \text{poly} \triangleright \text{poly}$ leads to $\text{max}(S'(z)) <: \text{poly}$, as needed.

For (3), suppose EDGE adds $\text{ret} \xrightarrow{(i \oplus N)} u$ to G' . This happens only if there is $x \xrightarrow{N} u \in G$. Therefore by the inductive hypothesis $\text{max}(S(x)) \neq \text{readonly}$, and therefore, constraint $q_x \triangleright \text{ret} <: q_x$ entails that $\text{max}(S'(\text{ret})) <: \text{poly}$, as needed.

Case 2: s is $x = y$ is straightforward.

Case 3: s is $x.f = y$. Suppose EDGE processes approximate edge $x.f \xrightarrow{a} x'.f$ followed by direct edge $y \xrightarrow{d} x.f$. EDGE adds to G' only if $x'.f$ in G , which by the inductive hypothesis entails $S(f) = \{\text{poly}\}$. Thus, $\text{max}(S(x.f))$ evaluates to mutable , and the desired subtyping is preserved (even though this typing is not precise). Furthermore, field write constraint $q_y <: q_x \triangleright q_f$ evaluates to $q_y <: \text{maybe} \triangleright \text{poly}$, meaning that $\text{max}(S(y)) <: \text{maybe}$ as needed to account for the $(M|C)A$ -path from y in G' .

Finally, consider case 4: s is $y = x.f$. There is new path in G' only if there is a path in G from y . If there is a path from y , then $\text{max}(S(y)) \neq \text{readonly}$. The constraint for field write $q_x \triangleright q_f <: q_y$ entails $\text{max}(S(f)) = \text{poly}$ and $\text{max}(S(x)) <: \text{max}(S(y))$. Therefore, $\text{max}(S(x.f))$ and $\text{max}(S(x))$ reflect the new paths through y . ◀

► **Lemma 8.** *If $S \Rightarrow G$ and $\{G, S\} \text{ EDGE}(e(s)) \parallel \text{CONSTRAINT}(c(s)) \{G', S'\}$ then $S' \Rightarrow G'$.*

Proof. The proof is by induction on the number of applications of

$\text{EDGE}(e(s)) \parallel \text{CONSTRAINT}(c(s))$

We begin with the most difficult case, case 1: s is $x = y.m(z)$.

XX:28 Definite Reference Mutability

The tables below examine all cases for parameter constraint $q_z <: q_x \triangleright q_p$. Constraint $q_y <: q_x \triangleright q_{\text{this}}$ is completely analogous. For brevity, sets S show only the maximal element. For example, when

$$S(x) = \{\text{readonly}, \dots \quad S(p) = \{\text{maybe}, \dots \quad S(z) = \{\text{readonly}, \dots$$

the constraint $q_z <: q_x \triangleright q_p$ removes `readonly` and `poly` from $S(z)$ resulting in:

$$S'(z) = \{\text{maybe}, \dots$$

To preserve precision, we must show that after execution of the parallel EDGE operation, there exist only an $(M|C)A$ -path from z to an update. The last column of the table enumerates the paths from z : added by $\text{EDGE}(z \xrightarrow{(i)} p)$ given $S(x)$ and $S(p)$, and existing ones in G . Continuing with the example, $(M|C)A[p] + \text{none}[\text{ret}] + \text{none}[z]$ reads as follows: (1) an $(M|C)A$ -path is added through p (the inductive hypothesis $S \Rightarrow G$ entails that the `maybe` typing of p implies an $(M|C)A$ -path from p), and no paths are added through `ret` and x (again, the `maybe` typing implies that there is no path through `ret` and x), and (2) there are no existing paths from z (due to the `readonly` typing of z in S). Therefore, there is only an $(M|C)A$ -path from z , and this case preserves precision.

The cases of $S(z) = \{\text{mutable}\}$ and $S(p) = \{\text{readonly}, \dots\}$ are not shown because neither of these cases triggers change to S , and it is trivial to argue $S' \Rightarrow G'$.

The following table enumerates the cases for $S(x) = \{\text{readonly}, \dots\}$:

$S(x)$	$S(p)$	$S(z)$	$S'(z)$	$G, \text{EDGE}(z \xrightarrow{(i)} p)$
{ readonly,	{ maybe,			$(M C)A[p] + \text{none}[\text{ret}]$
		{ readonly, { maybe, { poly, { polyormaybe,	{ maybe, NO CHANGE { polyormaybe, NO CHANGE	+none[z] + $(M C)A[z]$ + $R[z]$ + $(M C)A[z] + R[z]$
{ readonly,	{ poly,			$\text{none}[p] + \text{none}[x]$
		{ readonly, { maybe, { poly, { polyormaybe,	NO CHANGE NO CHANGE NO CHANGE NO CHANGE	+none[z] + $(M C)A[z]$ + $R[z]$ + $(M C)A[z] + R[z]$
{ readonly,	{ polyormaybe,			$(M C)A[p] + \text{none}[x]$
		{ readonly, { maybe, { poly, { polyormaybe,	{ maybe, NO CHANGE { polyormaybe, NO CHANGE	+none[z] + $(M C)A[z]$ + $R[z]$ + $(M C)A[z] + R[z]$

The table below enumerates the cases for $S(x) = \{\text{maybe}, \dots\}$:

$S(x)$	$S(p)$	$S(z)$	$S'(z)$	$G, \text{EDGE}(z \xrightarrow{i} p)$
{ maybe,	{ maybe,			$(M C)A[p] + \text{none}[\text{ret}]$
		{ readonly, { maybe, { poly, { polyormaybe,	{ maybe, NO CHANGE { polyormaybe, NO CHANGE	$+\text{none}[z]$ $+(M C)A[z]$ $+R[z]$ $+(M C)A[z] + R[z]$
{ maybe,	{ poly,			$\text{none}[p] + (M C)A[x]$
		{ readonly, { maybe, { poly, { polyormaybe,	{ maybe, NO CHANGE { polyormaybe NO CHANGE	$+\text{none}[z]$ $+(M C)A[z]$ $+R[z]$ $+(M C)A[z] + R[z]$
{ maybe,	{ polyormaybe,			$(M C)A[p] + (M C)A[x]$
		{ readonly, { maybe, { poly, { polyormaybe,	{ maybe, NO CHANGE { polyormaybe NO CHANGE	$+\text{none}[z]$ $+(M C)A[z]$ $+R[z]$ $+(M C)A[z] + R[z]$

The table below enumerates the cases for $S(x) = \{\text{poly}, \dots\}$:

$S(x)$	$S(p)$	$S(z)$	$S'(z)$	$G, \text{EDGE}(z \xrightarrow{i} p)$
{ poly,	{ maybe,			$(M C)A[p] + \text{none}[\text{ret}]$
		{ readonly, { maybe, { poly, { polyormaybe,	{ maybe, NO CHANGE { polyormaybe, NO CHANGE	$+\text{none}[z]$ $+(M C)A[z]$ $+R[z]$ $+(M C)A[z] + R[z]$
{ poly,	{ poly,			$\text{none}[p] + R[x]$
		{ readonly, { maybe, { poly, { polyormaybe,	{ poly, { polyormaybe NO CHANGE NO CHANGE	$+\text{none}[z]$ $+(M C)A[z]$ $+R[z]$ $+(M C)A[z] + R[z]$
{ poly,	{ polyormaybe,			$(M C)A[p] + R[x]$
		{ readonly, { maybe, { poly, { polyormaybe,	{ polyormaybe, { polyormaybe, { polyormaybe, NO CHANGE	$+\text{none}[z]$ $+(M C)A[z]$ $+R[z]$ $+(M C)A[z] + R[z]$

XX:30 Definite Reference Mutability

The table below enumerates the cases for $S(x) = \{\text{polyormaybe}, \dots\}$:

$S(x)$	$S(p)$	$S(z)$	$S'(z)$	$G, \text{EDGE}(z \xrightarrow{i} p)$
{ polyormaybe,	{ maybe,			$(M C)A[p] + \text{none}[\text{ret}]$
		{ readonly, { maybe, { poly, { polyormaybe,	{ maybe, NO CHANGE { polyormaybe, NO CHANGE	$+\text{none}[z]$ $+(M C)A[z]$ $+R[z]$ $+(M C)A[z] + R[z]$
{ polyormaybe,	{ poly,			$\text{none}[p] + R[x] + (M C)A[x]$
		{ readonly, { maybe, { poly, { polyormaybe,	{ polyormaybe, { polyormaybe, { polyormaybe, NO CHANGE	$+\text{none}[z]$ $+(M C)A[z]$ $+R[z]$ $+(M C)A[z] + R[z]$
{ polyormaybe,	{ polyormaybe,			$(M C)A[p] + (M C)A[x] + R[x]$
		{ readonly, { maybe, { poly, { polyormaybe,	{ polyormaybe, { polyormaybe, { polyormaybe, NO CHANGE	$+\text{none}[z]$ $+(M C)A[z]$ $+R[z]$ $+(M C)A[z] + R[z]$

Consider constraint $q_x \triangleright q_{\text{ret}} <: q_x$. If $S(x)$ is $\{\text{readonly}, \dots\}$ then there is no change to S and no change to G , and the statement holds. If $S(\text{ret})$ is $\{\text{poly}\}$, then, there is no change to S and EDGE “adds” a path already in G , resulting in no change in G as well. Let $S(x)$ be any other value but $\{\text{readonly}, \dots\}$ and let $S(\text{ret})$ be $\{\text{readonly}, \dots\}$. $S'(\text{ret})$ becomes $\{\text{poly}\}$ implying an R -path. Since $S(x)$ is of any other value but $\{\text{readonly}, \dots\}$, this means that a path from x to update, $x \xrightarrow{N} u$ does exist in G , and $\text{EDGE}(\text{ret} \xrightarrow{i} x)$ results in R path from ret in G' . Therefore, the theorem holds.

Consider case 2, s is $x = y$. We enumerate all possibilities analogously. Again we omit the cases when $S(x) = \{\text{mutable}, \dots\}$ as well as the case when $S(y) = \{\text{readonly}, \dots\}$, as they are trivial.

$S(x)$	$S(y)$	$S'(y)$	$G, \text{EDGE}(y \xrightarrow{d} x)$
{ maybe,	{ readonly, { maybe, { poly, { polyormaybe,	{ maybe, NO CHANGE { polyormaybe, NO CHANGE	$(M C)A[x] + \text{none}[y]$ $(M C)A[x] + (M C)A[y]$ $(M C)A[x] + R[y]$ $(M C)A[x] + (M C)A[y] + R[y]$
{ poly,	{ readonly, { maybe, { poly, { polyormaybe,	{ poly, { polyormaybe, NO CHANGE NO CHANGE	$R[x] + \text{none}[y]$ $R[x] + (M C)A[y]$ $R[x] + R[y]$ $R[x] + (M C)A[y] + R[y]$
{ polyormaybe,	{ readonly, { maybe, { poly, { polyormaybe,	{ polyormaybe, { polyormaybe, { polyormaybe, NO CHANGE	$(M C)A[x] + R[x] + \text{none}[y]$ $(M C)A[x] + R[x] + (M C)A[y]$ $(M C)A[x] + R[x] + R[y]$ $(M C)A[x] + R[x] + (M C)A[y] + R[y]$

Now consider case 3, s is $x.f = y$, and corresponding constraints $q_y <: \text{maybe} \triangleright q_f$. If f is readonly , then $\text{maybe} \triangleright q_f$ is readonly , and there is no change in S . By the inductive hypothesis, $x'.f$ being readonly implies that there does not exist a read $x'.f$ such that there is a path from $x'.f$ to update in G . Thus, no path is added to G' through $x.f$ thus preserving

the paths from y and the theorem. If f is **poly**, then y becomes **maybe**, or lower in S' , thus properly accounting for the $(M|C)A$ -path from y through $x.f$ that appears in G' .

Finally, consider case 4. s is $y' = x'.f$. If f is **readonly**, then there does not exist a path from $x'.f$ in G . If y' is not **readonly**, then there exists a path from y' . Thus, f becomes **poly** in S' , and $x' <: y'$ in S' . In G' , there are new paths from $x'.f$ and x reflecting S' .

