

Monotonic References for Efficient Gradual Typing

Jeremy G. Siek¹, Michael M. Vitousek¹, Matteo Cimini¹, Sam
Tobin-Hochstadt¹, and Ronald Garcia²

¹ Indiana University Bloomington
jsiek@indiana.edu

² University of British Columbia
rxg@cs.ubc.ca

Abstract. Gradual typing enables both static and dynamic typing in the same program and makes it convenient to migrate code regions between the two typing disciplines. One goal of gradual typing is to provide all the benefits of static typing, such as efficiency, in statically-typed regions. However, this goal is elusive: the standard approach to mutable references imposes run-time overhead in statically-typed regions and alternative approaches are too conservative, either statically or at run-time. In this paper we present a new semantics called *monotonic references* which imposes none of the run-time overhead of dynamic typing in statically typed regions. With this design, casting a reference may cause a heap cell to become more statically typed (but not less). Retaining type safety is challenging with strong updates to the heap. Nevertheless, we have a mechanized proof of type safety. Further, we present blame tracking for monotonic references and prove a blame theorem.

1 Introduction

Static and dynamic type systems have well-known strengths and weaknesses. Static type systems provide machine-checked documentation, catch bugs early, and enable efficient code. Dynamic type systems provide the flexibility often needed during prototyping and enable powerful features such as reflection. Over the years, many languages blurred the boundary between static and dynamic typing, such as type hints in Lisp and the addition of a dynamic type to otherwise statically typed languages [Abadi et al., 1989]. But the seamless and sound integration of static and dynamic typing remained problematic until two pieces fell into place: the gradual type system of Siek and Taha [2006] and the blame theorems of Tobin-Hochstadt and Felleisen [2006] and Wadler and Findler [2009].

However, there are challenges regarding the efficiency of gradual typing. One issue concerns mutable references in statically-typed regions of code. Consider the following statically-typed function f that dereferences its parameter x .

```
let f = λx:Ref Int. !x in
f(ref 4);
f(ref (4 as ★))
```

In the first application of f , a normal reference to an integer flows into f . For the second application, we allocate a reference of type $\text{Ref } \star$ (\star is the dynamic type) then implicitly cast it to Ref Int before applying f . According to the semantics of Herman et al. [2007], this cast wraps the reference in a proxy which performs dynamic checks on reads and writes. Thus code generated for the dereference in the body of f must inspect the reference to find out whether it is a normal reference or a proxied reference, and in the proxied case, apply a coercion.

Before discussing solutions to this problem, we recall the *gradual guarantee* of Boyland [2014] and Siek et al. [2015], an important property of the standard semantics for mutable references, and of gradual typing in general. The gradual guarantee promises that removing type annotations, or changing type annotations to be less precise, does not affect the behavior of a program: it should still type check and the result should be the same modulo proxies. (Adding or making type annotations more precise, on the other hand can sometimes induce static type errors and runtime cast errors.) Consider the statically-typed program on the left that allocates a reference to an integer and then dereferences it from within a function. In the code on the right, we change the annotation on h from Ref Int to $\text{Ref } \star$, but the program still type checks and the result remains 42.

$$\begin{array}{ll} \text{let } r = \text{ref } 42 \text{ in} & \text{let } r = \text{ref } 42 \text{ in} \\ \text{let } f = \lambda h:\text{Ref Int}.\!h & \implies \text{let } f = \lambda h:\text{Ref } \star.\!h \\ \text{in } f(r) & \text{in } f(r) \end{array}$$

Wrigstad et al. [2010] address the efficiency problem by introducing a distinction between *like types* and *concrete types*. Concrete types are the usual types of a statically-typed language and incur zero run-time overhead, but dynamically-typed values cannot flow into concrete types. Like types, on the other hand, may refer to dynamically-typed values but incur run-time overhead. The distinction between like types and concrete types achieves the efficiency goals, but the restrictions in their type system mean that removing concrete type annotations, as in the above example, can trigger a static type error.

In this paper we investigate this run-time overhead problem in the context of the gradually-typed lambda calculus with mutable references. We propose a semantics, *monotonic references*, that enables the compilation of statically-typed regions to machine code that is free of any of the indirection or run-time checking associated with dynamic typing, like boxing or bit tags. Monotonic references allow dynamically-typed values to flow into code with (concrete) static types. When a reference flows through a cast, the cast may coerce its underlying heap cell to become more statically typed. In general, this means that values in the heap may evolve monotonically with respect to the precision relation (Section 2). The idea for monotonic references came out of our work on implementing and evaluating gradual typing for Python [Vitousek et al., 2014].

Monotonic references preserve a global invariant that a value in the heap is at least as precise as any reference that points to it. Thus, a static reference always points to a value of the same type, so there is no overhead associated with reading or writing through the reference: the reads and writes may be implemented

as machine loads and stores. By a *static* reference we mean that there are no occurrences of the dynamic type \star in the pointed-to type of the reference, such as `Ref Int` and `Ref (Int × Bool)`. Reads and writes to references that are not static, such as `Ref \star` and `Ref (\star × Bool)`, still require casts: the dynamic regions of code have to pay their own way. The intermediate representation that we compile to contains different instructions for fast, static loads and stores versus non-static loads and stores that require casts.

Swamy et al. [2014] and Rastogi et al. [2014] integrate static and dynamic typing in the context of TypeScript with the TS^\star and Safe TypeScript languages. Both use a notion of monotonicity in the heap, but with respect to subtyping, treating \star as a universal supertype, instead of with respect to the precision relation. Because these languages compile to JavaScript, they inherit the overhead of dynamic typing, whereas with monotonic references, the overhead of dynamic typing occurs only in dynamically-typed code. In the example above, making the type annotation on h less precise causes TS^\star to halt the program with a cast error at the implicit cast from `Ref Int` to `Ref \star` . TS^\star does not allow casts from one mutable reference type to a different one because its references are invariant with respect to subtyping. Thus, TS^\star does not satisfy the gradual guarantee.

In gradually-typed languages with higher-order features such as first-class functions and objects, blame tracking plays an important role in providing meaningful error messages when casts fail. Blame tracking enables fine-grained guarantees, via a blame theorem, regarding which regions of the code are statically type safe. In this paper we present blame tracking for monotonic references and prove a blame theorem. Our design uses the labeled types of Siek and Wadler [2010] as run-time type information (RTTI), together with three new operations on labeled types: a bidirectional cast operator that captures the dual read/write nature of mutable references, a merge operator that models how casts on separate aliases to the same heap cell interact over time, and an operator that casts heap cells between labeled types.

To summarize, this paper presents a new semantics for gradually-typed mutable references that delivers guaranteed efficiency for the statically-typed parts of a program, maintains type safety, and provides blame tracking, while continuing to enable fine-grained migration between static and dynamic code. This paper makes the following technical contributions:

1. We define the semantics of monotonic references (Sections 3 and 5).
2. We discuss our proof of type safety, mechanized in Isabelle (Section 4).
3. We augment monotonic references with blame tracking and prove the blame-subtyping theorem (Section 6).

We review the gradually-typed lambda calculus with references in Section 2 and discuss the run-time overhead associated with mutable references. We address an implementation concern regarding strong updates in Section 7. The paper concludes in Section 9.

2 Background and Problem Statement

Figure 1 reviews the syntax and static semantics of the gradually-typed lambda calculus with references. The primary difference between gradual typing and static typing is that uses of type equality are replaced with *consistency* (aka. compatibility), also defined in Figure 1. The consistency relation enables implicit casts to and from \star . (In contrast, an object-oriented language only allows implicit casts to the top `Object` type.) This consistency relation is a congruence, even for reference types [Herman et al., 2007], which differs from the original treatment of references as invariant [Siek and Taha, 2006]. The more flexible treatment of references enables the passing of references between more and less dynamically typed regions of code, but is also the source of the difficulties that we solve in this paper. The precision relation, which says whether one type is more or less dynamic than another, is also defined in Figure 1, and is closely related to consistency. Two types are consistent when there exists a greatest lower bound with respect to the precision relation. This relation is also known as naïve subtyping [Wadler and Findler, 2009].

All of the types, except for \star , classify unboxed values. So, for example, `Int` is the type for native integers (e.g. 64-bit integers). The auxiliary relations *fun*, *pair*, and *ref*, defined in Figure 1, implement pattern matching on types, enabling a more concise presentation of the typing rules compared to prior presentations of gradual type systems. Labels ℓ represent source code locations that are captured during parsing.

The dynamic semantics of the gradually-typed lambda calculus is defined by a type-directed translation to the coercion calculus [Henglein, 1994], using the standard semantics for mutable references due to Herman et al. [2007].

Each use of consistency between types T_1 and T_2 in the type system, and each use of one of the auxiliary relations, becomes an explicit cast from T_1 to T_2 . The coercion calculus expresses casts in terms of combinators that say how to cast from one type to another. Figure 2 gives the compilation of casts into coercions, written $(T \Rightarrow^\ell T) = c$. The compilation of gradually-typed terms into the coercion-based calculus is otherwise straightforward, so we give just the function application rule as an example:

$$\frac{\begin{array}{cc} \Gamma \vdash e_1 \rightsquigarrow e'_1 : T_1 & \Gamma \vdash e_2 \rightsquigarrow e'_2 : T_2 \\ \text{fun}(T_1, T_{11}, T_{12}) & T_2 \sim T_{11} \\ (T_1 \Rightarrow^\ell T_{11} \rightarrow T_{12}) = c_1 & (T_2 \Rightarrow^\ell T_{11}) = c_2 \end{array}}{\Gamma \vdash (e_1 e_2)^\ell \rightsquigarrow e'_1 \langle c_1 \rangle e'_2 \langle c_2 \rangle : T_{12}}$$

Figures 3 and 4 define the coercion-based calculus. We highlight the parts of the definition related to references, as they are of particular interest here. We review the coercion calculus in the context of discussing the run-time overhead problem in the next subsection. For an introduction to the coercion calculus, we refer to Henglein [1994].

Syntax

Base types $B ::= \text{Int} \mid \text{Bool}$
Types $T ::= B \mid T \rightarrow T \mid T \times T \mid \text{Ref } T \mid \star$
Labels ℓ
Operators $op ::= \text{plus} \mid \text{minus} \mid \text{is} \mid \dots$
Expressions $e ::= k \mid op^\ell(\vec{e}) \mid x \mid \lambda x:T. e \mid (e e)^\ell \mid e \text{ as}^\ell T \mid$
 $(e, e) \mid \text{fst}^\ell e \mid \text{snd}^\ell e \mid \text{ref } e \mid !^\ell e \mid e :=^\ell e$
 $\lambda x. e \equiv \lambda x: \star. e$

Consistency

$T \sim T$

$$\frac{}{\star \sim T} \quad \frac{}{T \sim \star} \quad \frac{}{B \sim B} \quad \frac{T_1 \sim T_2}{\text{Ref } T_1 \sim \text{Ref } T_2}$$

$$\frac{T_1 \sim T_3 \quad T_2 \sim T_4}{T_1 \rightarrow T_2 \sim T_3 \rightarrow T_4} \quad \frac{T_1 \sim T_3 \quad T_2 \sim T_4}{T_1 \times T_2 \sim T_3 \times T_4}$$

Precision

$T \sqsubseteq T$

$$\frac{T \sqsubseteq \star \quad B \sqsubseteq B}{\text{Ref } T_1 \sqsubseteq \text{Ref } T_2} \quad \frac{T_1 \sqsubseteq T_2}{\text{Ref } T_1 \sqsubseteq \text{Ref } T_2}$$

$$\frac{T_1 \sqsubseteq T_3 \quad T_2 \sqsubseteq T_4}{T_1 \rightarrow T_2 \sqsubseteq T_3 \rightarrow T_4} \quad \frac{T_1 \sqsubseteq T_3 \quad T_2 \sqsubseteq T_4}{T_1 \times T_2 \sqsubseteq T_3 \times T_4}$$

Expression typing

$\Gamma \vdash e : T$

$$\frac{k : B}{\Gamma \vdash k : B} \quad \frac{\Gamma \vdash \vec{e} : \vec{T} \quad op : \vec{B} \rightarrow B \quad \vec{T} \sim \vec{B}}{\Gamma \vdash op^\ell(\vec{e}) : B} \quad \frac{\Gamma \vdash e : T_1 \quad T_1 \sim T_2}{\Gamma \vdash e \text{ as}^\ell T_2 : T_2}$$

$$\frac{\Gamma(x) = T}{\Gamma \vdash x : T} \quad \frac{\Gamma(x \mapsto T_1) \vdash e : T_2}{\Gamma \vdash \lambda x:T_1. e : T_1 \rightarrow T_2} \quad \frac{\Gamma \vdash e_1 : T_1 \quad \Gamma \vdash e_2 : T_2 \quad \text{fun}(T_1, T_{11}, T_{12}) \quad T_2 \sim T_{11}}{\Gamma \vdash (e_1 e_2)^\ell : T_{12}}$$

$$\frac{\Gamma \vdash e_1 : T_1 \quad \Gamma \vdash e_2 : T_2}{\Gamma \vdash (e_1, e_2) : T_1 \times T_2} \quad \frac{\Gamma \vdash e : T \quad \text{pair}(T, T_1, T_2)}{\Gamma \vdash \text{fst}^\ell e : T_1} \quad \frac{\Gamma \vdash e : T \quad \text{pair}(T, T_1, T_2)}{\Gamma \vdash \text{snd}^\ell e : T_2}$$

$$\frac{\Gamma \vdash e : T}{\Gamma \vdash \text{ref } e : \text{Ref } T} \quad \frac{\Gamma \vdash e : T \quad \text{ref}(T, T')}{\Gamma \vdash !^\ell e : T'} \quad \frac{\Gamma \vdash e_1 : T_1 \quad \Gamma \vdash e_2 : T_2 \quad \text{ref}(T_1, T'_1) \quad T_2 \sim T'_1}{\Gamma \vdash e_1 :=^\ell e_2 : T_1}$$

Type matching

$$\frac{\text{fun}(T_{11} \rightarrow T_{12}, T_{11}, T_{12})}{\text{pair}(T_{11} \times T_{12}, T_{11}, T_{12})} \quad \frac{\text{fun}(\star, \star, \star)}{\text{pair}(\star, \star, \star)}$$

$$\frac{\text{ref}(\text{Ref } T, T)}{\text{ref}(\star, \star)}$$

Fig. 1. Gradually-typed λ calculus with mutable references

$$(T \Rightarrow^\ell T) = c$$

$$\begin{aligned} (B \Rightarrow^\ell B) &= \iota & (I \Rightarrow^\ell \star) &= I! \\ (\star \Rightarrow^\ell \star) &= \iota & (\star \Rightarrow^\ell I) &= I?^\ell \\ (T_1 \rightarrow T_2) \Rightarrow^\ell (T'_1 \rightarrow T'_2) &= (T'_1 \Rightarrow^\ell T_1) \rightarrow (T_2 \Rightarrow^\ell T'_2) \\ (T_1 \times T_2) \Rightarrow^\ell (T'_1 \times T'_2) &= (T_1 \Rightarrow^\ell T'_1) \times (T_2 \Rightarrow^\ell T'_2) \\ \mathbf{Ref} T \Rightarrow^\ell \mathbf{Ref} T' &= \mathbf{Ref} (T \Rightarrow^\ell T') (T' \Rightarrow^\ell T) \end{aligned}$$

Fig. 2. Compile casts to coercions

Expressions	$e ::= k \mid op(\vec{e}) \mid x \mid \lambda x. e \mid e e \mid (e, e) \mid \mathbf{fst} e \mid \mathbf{snd} e \mid \mathbf{ref} e \mid !e \mid e := e \mid e \langle c \rangle \mid \mathbf{blame} \ell$
Injectibles	$I ::= B \mid T \rightarrow T \mid T \times T \mid \mathbf{Ref} T$
Coercions	$c ::= \iota \mid I?^\ell \mid I! \mid c \rightarrow c \mid c \times c \mid c ; c \mid \mathbf{Ref} c c$
Values	$v ::= k \mid \lambda x. e \mid (v, v) \mid v \langle I! \rangle \mid a \mid v \langle \mathbf{Ref} c c \rangle$
Heap	$\mu ::= \emptyset \mid \mu(a \mapsto v)$
Heap Typing	$\Sigma ::= \emptyset \mid \Sigma(a \mapsto T)$
Frames	$F ::= op(\vec{v}, \square, \vec{e}) \mid \square e \mid v \square \mid (\square, e) \mid (v, \square) \mid \mathbf{fst} \square \mid \mathbf{snd} \square \mid \mathbf{ref} \square \mid !\square \mid \square := e \mid v := \square \mid \square \langle c \rangle$

Fig. 3. Syntax for the coercion-based calculus with mutable references

2.1 Run-time overhead in fully-static code

Recall the example in Section 1 in which the dereference of a statically-typed reference must first check whether the reference is proxied or not.

```
let f = λx:Ref Int. !x in
f(ref 4);
let r = ref (4 as ★) in f(r)
```

The overhead can be seen in the dynamic semantics (Figure 4), where there are two reduction rules for dereferencing: (DEREF) and (DEREFCAST), and two reduction rules for updating references: (UPDATE) and (UPDATECAST). Another way to look at this problem is that there are two canonical forms of type $\mathbf{Ref} \mathbf{Int}$, a plain address a and also a value wrapped in a reference coercion, $v \langle \mathbf{Ref} c_1 c_2 \rangle$, so operations on values of this type need to dispatch on which form occurs at runtime. To eliminate this overhead we need a design with only a single canonical form for values of reference type.

The run-time overhead for references affects every read and write to the heap and is particularly detrimental in tight loops over arrays. When adding support for contracts to mutable data structures in Racket, Strickland et al. [2012, Figure 9] measured this overhead at approximately 25% for fully-typed code on a bubble-sort microbenchmark.

Coercion typing

$$\boxed{c : T \Rightarrow T}$$

$$\frac{}{\iota : T \Rightarrow T} \quad \frac{c_1 : T_3 \Rightarrow T_1 \quad c_2 : T_2 \Rightarrow T_4}{c_1 \rightarrow c_2 : (T_1 \rightarrow T_2) \Rightarrow (T_3 \rightarrow T_4)}$$

$$\frac{}{I^{? \ell} : \star \Rightarrow I} \quad \frac{c_1 : T_1 \Rightarrow T_3 \quad c_2 : T_2 \Rightarrow T_4}{c_1 \times c_2 : (T_1 \times T_2) \Rightarrow (T_3 \times T_4)}$$

$$\frac{}{I! : I \Rightarrow \star} \quad \frac{c_1 : T_1 \Rightarrow T_2 \quad c_2 : T_2 \Rightarrow T_3}{c_1 ; c_2 : T_1 \Rightarrow T_3}$$

$$\frac{c_1 : T_1 \Rightarrow T_2 \quad c_2 : T_2 \Rightarrow T_1}{\mathbf{Ref} \ c_1 \ c_2 : \mathbf{Ref} \ T_1 \Rightarrow \mathbf{Ref} \ T_2}$$

Expression typing

$$\boxed{\Gamma; \Sigma \vdash e : T}$$

$$\dots \quad \frac{\Sigma(a) = T}{\Gamma; \Sigma \vdash a : T} \quad \frac{\Gamma; \Sigma \vdash e : T_1 \quad c : T_1 \Rightarrow T_2}{\Gamma; \Sigma \vdash e \langle c \rangle : T_2}$$

Reduction rules for functions, primitives, and pairs

$$\boxed{e \longrightarrow e}$$

$$\begin{array}{ll} (\lambda x. e) v \longrightarrow [x := v]e & \mathbf{fst} \ (v_1, v_2) \longrightarrow v_1 \\ \mathit{op}(\vec{k}) \longrightarrow \delta(\mathit{op}, \vec{k}) & \mathbf{snd} \ (v_1, v_2) \longrightarrow v_2 \end{array}$$

Cast reduction rules

$$\boxed{e \longrightarrow_c e}$$

$$\begin{array}{l} v \langle \iota \rangle \longrightarrow_c v \\ v \langle I_1! \rangle \langle I_2^{? \ell} \rangle \longrightarrow_c v \langle I_1 \Rightarrow^\ell I_2 \rangle \quad \text{if } I_1 \sim I_2 \\ v \langle I_1! \rangle \langle I_2^{? \ell} \rangle \longrightarrow_c \mathbf{blame} \ \ell \quad \text{if } I_1 \not\sim I_2 \\ v \langle c_1 \rightarrow c_2 \rangle \longrightarrow_c \lambda x. v \ (x \langle c_1 \rangle) \langle c_2 \rangle \\ (v_1, v_2) \langle c_1 \times c_2 \rangle \longrightarrow_c (v_1 \langle c_1 \rangle, v_2 \langle c_2 \rangle) \\ v \langle c_1 ; c_2 \rangle \longrightarrow_c v \langle c_1 \rangle \langle c_2 \rangle \end{array}$$

Reference reduction rules

$$\boxed{e, \mu \longrightarrow_r e, \mu}$$

$$\begin{array}{ll} \mathbf{ref} \ v, \mu \longrightarrow_r a, \mu(a \mapsto v) & \text{if } a \notin \mathit{dom}(\mu) \quad (\mathbf{ALLOCREF}) \\ !a, \mu \longrightarrow_r \mu(a), \mu & (\mathbf{DEREF}) \\ !(v \langle \mathbf{Ref} \ c_1 \ c_2 \rangle), \mu \longrightarrow_r (!v) \langle c_1 \rangle, \mu & (\mathbf{DEREFCAST}) \\ a := v, \mu \longrightarrow_r a, \mu(a \mapsto v) & (\mathbf{UPDATE}) \\ v_1 \langle \mathbf{Ref} \ c_1 \ c_2 \rangle := v_2, \mu \longrightarrow_r v_1 := v_2 \langle c_2 \rangle, \mu & (\mathbf{UPDATECAST}) \end{array}$$

State reduction rules

$$\frac{e \longrightarrow e'}{e, \mu \longrightarrow e', \mu} \quad \frac{e \longrightarrow_c e'}{e, \mu \longrightarrow e', \mu} \quad \frac{e, \mu \longrightarrow_r e', \mu'}{e, \mu \longrightarrow e', \mu'}$$

$$\frac{e, \mu \longrightarrow e', \mu'}{F[e], \mu \longrightarrow F[e'], \mu'} \quad \frac{}{F[\mathbf{blame} \ \ell], \mu \longrightarrow \mathbf{blame} \ \ell, \mu}$$

Fig. 4. Coercion-based calculus with mutable references

2.2 Non-determinism in multi-threaded code

This standard semantics for mutable references produces an error only if type inconsistency is witnessed by some read or write to a particular reference, so in a non-deterministic multi-threaded program, whether a check will fail at run-time is difficult to predict.

The contract system in Racket implements the standard semantics [Flatt and PLT, 2014]. For example, the following program sometimes fails and blames `b1`, sometimes fails and blames `b2`, and sometimes succeeds, as explained below.

```
#lang racket
(define b (box #f))
(define/contract b1 (box/c integer?) b)
(define/contract b2 (box/c string?) b)

(thread (lambda ()
  (for ([i 2])
    (set-box! b1 5)
    (sleep 0.000000001)
    (add1 (unbox b1))))))
(thread (lambda ()
  (for ([i 2])
    (set-box! b2 "hello")
    (sleep 0.000000001)
    (string-append "world" (unbox b2))))))
```

The program creates a single heap cell `b`, and accesses it through two distinct proxies, `b1` and `b2`, each with its own dynamic check. When the two threads do not interleave, the program succeeds, but if the second thread changes `b2` to contain a string between the `set-box!` and `unbox` calls for `b1`, the system halts, blaming one of the parties.

In contrast, if `box/c` implemented monotonic references, then an error would *deterministically* occur when `define/contract` is used for the second time.

3 Monotonic References Without Blame

Figures 5 and 6 define the syntax and semantics of our new coercion calculus with monotonic references, but without blame. Figure 8 defines the compilation of casts to monotonic coercions, also without blame. The addition of blame adds considerable complexity, so we postpone its treatment to Section 5. Typical of gradually-typed languages, there is a value form for values that have been boxed and injected to \star , which is $v\langle I! \rangle$. The I plays the role of a tag that records the type of v . The values at all other types are unboxed, as they would be in a statically-typed language.

With monotonic references, only one kind of value has reference type: normal addresses. When a cast is applied to a reference, instead of wrapping the reference with a cast, we cast the underlying value on the heap. To make sure that the

new type of the value is consistent with all the outstanding references, we require that a cast only make the type of the value more precise (Figure 1). Otherwise the cast results in a run-time error. Thus, we maintain the heap invariant that the type of each reference in the program is less or equally precise as the type of the value on the heap that it points to, as captured in the typing rule (WTREF).

One might wonder why our heap invariant uses the precision relation instead of subtyping. Could we obtain the same efficiency goals using subtyping instead? Consider the following program in which a function of type $\star \rightarrow \text{Int}$ is referenced from the static type $\text{Int} \rightarrow \text{Int}$. (We have $\star \rightarrow \text{Int} <: \text{Int} \rightarrow \text{Int}$.)

```
let r1 = ref (λx : ⋆. x as Int) in
let r2 = (r1 as Ref (Int → Int)) in
!r2 42
```

The dereference of r_2 should not require overhead, but we have a function of type $\star \rightarrow \text{Int}$ that is to be applied to an integer, and the conversion from Int to \star requires boxing in our setting. Thus, the dereference of r_2 is not simply a load instruction, but it must handle the casting from $\star \rightarrow \text{Int}$ to $\text{Int} \rightarrow \text{Int}$. (Other systems, such as Reticulated Python and TS^{*}, box all values. In these systems, upcasts on dereferences are unnecessary, but instead overhead is incurred in nearly every operation.) In general, given a reference of type $\text{Ref } T_2$, even when T_2 is a static type, there are many types T_1 such that $T_1 <: T_2$ and $T_1 \neq T_2$.

The syntax of the monotonic calculus differs from the standard calculus in that there are two kinds of dereference and update expressions. Programmers need not worry about choosing which of the two dereference or update expressions to use because this choice is type-directed and therefore is handled during compilation from the source language to the coercion calculus. We reserve the forms $!e$ and $e_1 := e_2$ for situations in which the reference type is fully static (See the typing rules in Figure 6). In these situations we know that the value in the heap has the same type as the reference. Thus, if a reference has a fully static type, such as Ref Int , the corresponding value on the heap must be an actual integer (and not an injection to \star), so we need only one reduction rule for dereferencing a fully-static reference (DEREFM), and one rule for updating a fully-static reference (UPDM).

For expressions of reference type that are not fully-static, we introduce the syntactic forms $!e@T$ and $e_1 := e_2@T$ for dereference and update, respectively. The type annotation T records the compile-time type of e , that is, e has type $\text{Ref } T$. For example, T could be \star , $\star \times \star$, or $\star \times \text{Int}$. Because the value on the heap might be more precise than T , a cast is needed to mediate between T and the run-time type of the heap cell.

The reduction rule (DYNDEREFM) casts from the addresses' run-time type, which we store next to the heap cell, to the compile-time type T . We write $\mu(a)_{\text{rtti}}$ for the run-time type information for reference a and we write $\mu(a)_{\text{val}}$ for the value in the heap cell. The reduction rule (DYNUPDM) casts the incoming value v from T to the address's run-time type, so the new content of the cell is $cv = v\langle T \Rightarrow \mu(a)_{\text{rtti}} \rangle$. This cv is not a value yet, so storing it in the heap is

Expressions	$e ::= .. \mid \mathbf{ref}_T e \mid !e@T \mid e := e@T \mid \mathbf{error}$
Coercions	$c ::= \iota \mid I? \mid I! \mid c \rightarrow c \mid c \times c \mid c ; c \mid \mathbf{Ref } T$
Values	$v ::= k \mid \lambda x. e \mid (v, v) \mid v\langle I! \rangle \mid a$
Casted Values	$cv ::= v \mid cv\langle c \rangle \mid (cv, cv)$
Heap	$\mu ::= \emptyset \mid \mu(a \mapsto v : T)$
Evolving Heap	$\nu ::= \emptyset \mid \nu(a \mapsto cv : T)$
Frames	$F ::= .. \mid !\square@T \mid \square := e@T \mid v := \square@T$

Fig. 5. Syntax for monotonic references without blame

unusual. In earlier versions of the semantics we tried to reduce cv to a value before storing it in the heap, but there are complications that force this design, which we discuss later in this section. To summarize our treatment of dereference and update, we present efficient semantics for the fully-static dereference and update but have slightly increased the overhead for dynamic dereferences and updates. This is a price we are willing to pay to have dynamic typing “pay its own way”.

The crux of the monotonic semantics is in the reduction rules that apply a reference coercion to an address: (CASTREF1), (CASTREF2), and (CASTREF3). In (CASTREF1) we have an address that maps to cv of type T_1 and we cast cv so that it is no more dynamic than (i.e. at least as static as) both the target type T_2 and all of the existing references to the cell. To accomplish this, we take the greatest lower bound $T_3 = T_1 \sqcap T_2$ (Figure 7) to be the new type of the cell, so the new contents is $cv' = cv\langle T_1 \Rightarrow T_3 \rangle$. There are two side conditions on (CASTREF1): $T_1 \sqcap T_2$ must be defined and $T_3 \neq T_1$. If $T_1 \sqcap T_2$ is undefined, or equivalently, if $T_1 \not\sim T_2$, we instead signal an error, as handled by (CASTREF3). If $T_3 = T_1$, then there is no need to cast cv , which is handled by (CASTREF2).

The last coercion reduction rule (PURECAST) imports the reduction rules from the standard semantics (Figure 4) though here we ignore blame, i.e., replace $\mathbf{blame } \ell$ with \mathbf{error} , $I_2?^\ell$ with $I_2?$, and $I_1 \Rightarrow^\ell I_2$ with $I_1 \Rightarrow I_2$.

The meet function defined in Figure 7 computes the greatest lower bound with respect to the precision relation.

To motivate our organization of the heap, we present two examples that demonstrate why we store run-time type information and casted values, not just values, on the heap.

Cycles and termination The first complication is that there can be cycles in the heap and we need to make sure that when we apply a cast to an address in a cycle, the cast terminates. Consider the following example in which we create a pair whose second element is a reference back to itself.

```

let r1 = ref (42, 0 as *) in
r1 := (42, r1 as *);
let r2 = r1 as Ref (Int × Ref *) in
fst !r2

```

Once the cycle is established, we cast r_1 from type $\mathbf{Ref } (\mathbf{Int } \times \star)$ to $\mathbf{Ref } (\mathbf{Int } \times \mathbf{Ref } \star)$. The presence of the nested $\mathbf{Ref } \star$ in the target type means that the cast

Expression typing

$$\boxed{\Gamma; \Sigma \vdash e : T}$$

$$\frac{\Gamma; \Sigma \vdash e : \mathbf{Ref} T}{\Gamma; \Sigma \vdash !e : T} \quad \frac{\Gamma; \Sigma \vdash e_1 : \mathbf{Ref} T \quad \Gamma; \Sigma \vdash e_2 : T \quad \mathit{static} T}{\Gamma; \Sigma \vdash e_1 := e_2 : \mathbf{Ref} T} \quad \frac{\Gamma; \Sigma \vdash e : \mathbf{Ref} T}{\Gamma; \Sigma \vdash !e@T : T}$$

$$\frac{\Gamma; \Sigma \vdash e_1 : \mathbf{Ref} T \quad \Gamma; \Sigma \vdash e_2 : T}{\Gamma; \Sigma \vdash e_1 := e_2@T : \mathbf{Ref} T} \quad \dots \quad \frac{\Sigma(a) \sqsubseteq T_2}{\Gamma; \Sigma \vdash a : T_2} \quad (\text{WTREF})$$

Cast reduction rules

$$\boxed{e, \nu \longrightarrow_{cr} e, \nu}$$

$$\frac{e \longrightarrow_c e'}{e, \nu \longrightarrow_{cr} e', \nu} \quad (\text{PURECAST})$$

$$\frac{\nu(a) = cv : T_1 \quad T_3 = T_1 \sqcap T_2 \quad T_3 \neq T_1 \quad cv' = cv \langle T_1 \Rightarrow T_3 \rangle}{a(\mathbf{Ref} T_2), \nu \longrightarrow_{cr} a, \nu(a \mapsto cv' : T_3)} \quad (\text{CASTREF1})$$

$$\frac{\nu(a) = cv : T_1 \quad T_1 = T_1 \sqcap T_2}{a(\mathbf{Ref} T_2), \nu \longrightarrow_{cr} a, \nu} \quad (\text{CASTREF2})$$

$$\frac{\nu(a) = cv : T_1 \quad T_1 \not\sqsubseteq T_2}{a(\mathbf{Ref} T_2), \nu \longrightarrow_{cr} \mathbf{error}, \nu} \quad (\text{CASTREF3})$$

Program reduction rules

$$\boxed{e, \mu \longrightarrow_e e, \nu}$$

$$e, \mu \longrightarrow_e e', \mu \quad \text{if } e \longrightarrow e'$$

$$\mathbf{ref}_T v, \mu \longrightarrow_e a, \mu(a \mapsto v : T) \quad \text{if } a \notin \text{dom}(\mu)$$

$$!a, \mu \longrightarrow_e \mu(a)_{\text{val}}, \mu \quad (\text{DEREFM})$$

$$!a@T, \mu \longrightarrow_e \mu(a)_{\text{val}} \langle \mu(a)_{\text{rtti}} \Rightarrow T \rangle, \mu \quad (\text{DYNDEREFM})$$

$$a := v, \mu \longrightarrow_e a, \mu(a \mapsto v : \mu(a)_{\text{rtti}}) \quad (\text{UPDM})$$

$$a := v@T, \mu \longrightarrow_e a, \mu(a \mapsto cv : \mu(a)_{\text{rtti}}) \quad (\text{DYNUPDM})$$

where $cv = v \langle T \Rightarrow \mu(a)_{\text{rtti}} \rangle$

For $X \in \{cr, e\}$:

$$\frac{e, \nu \longrightarrow_X e', \nu'}{F[e], \nu \longrightarrow_X F[e'], \nu'} \quad \frac{}{F[\mathbf{error}], \nu \longrightarrow_X \mathbf{error}, \nu}$$

State reduction rules

$$\boxed{e, \nu \longrightarrow e, \nu}$$

$$\frac{e, \mu \longrightarrow_X e', \nu \quad X \in \{cr, e\}}{e, \mu \longrightarrow e', \nu} \quad \frac{\nu(a) = cv : T \quad cv, \nu \longrightarrow_{cr} cv', \nu' \quad \nu'(a)_{\text{rtti}} = T}{e, \nu \longrightarrow e, \nu'(a \mapsto cv' : T)} \quad (\text{HCAST})$$

$$\frac{\nu(a) = cv : T \quad cv, \nu \longrightarrow_{cr} \mathbf{error}, \nu'}{e, \nu \longrightarrow \mathbf{error}, \nu'} \quad \frac{\nu(a) = cv : T \quad cv, \nu \longrightarrow_{cr} cv', \nu' \quad \nu'(a)_{\text{rtti}} \neq T}{e, \nu \longrightarrow e, \nu'} \quad (\text{HDROP})$$

Fig. 6. Monotonic references without blame

$$\boxed{T \sqcap T = T}$$

$$\begin{array}{ll} \star \sqcap T = T & (T_1 \times T_2) \sqcap (T_3 \times T_4) = (T_1 \sqcap T_3) \times (T_2 \sqcap T_4) \\ T \sqcap \star = T & (T_1 \rightarrow T_2) \sqcap (T_3 \rightarrow T_4) = (T_1 \sqcap T_3) \rightarrow (T_2 \sqcap T_4) \\ B \sqcap B = B & \end{array}$$

Fig. 7. The meet function (greatest lower bound)

on r_1 will trigger another cast on r_1 . The correct result of this program is 42 but a naïve dynamic semantics would diverge. Our semantics avoids divergence by checking whether the new run-time type is equal to the old run-time type; in such cases the heap cell is left unchanged (see rule (CASTREF2)).

Casted values in the heap Consider the following example in which we create a triple of type $\star \times \star \times \star$ whose third element is a reference back to itself.

```
let r0 = ref (42 as  $\star$ , 7 as  $\star$ , 0 as  $\star$ ) in
r0 := (42 as  $\star$ , 7 as  $\star$ , r0 as  $\star$ );
let r1 = r0 as Ref (Int  $\times$   $\star$   $\times$  Ref (Int  $\times$  Int  $\times$   $\star$ )) in
fst (fst !r1)
```

Suppose a_0 is the address created in the allocation on the first line. On line three we cast a_0 in such a way that we trigger two casts on a_0 . Consider the action of these casts on just the first two elements of the triple, we have:

$$\star \times \star \Rightarrow \text{Int} \times \star \Rightarrow \text{Int} \times \text{Int}$$

The second cast occurs while the first is still in progress. Now, suppose we delayed updating the heap cell until we finished reducing to a value. At the moment when we apply the second cast, we would still have the original value, of type $\star \times \star$, in the heap. This is problematic because our next step would be to apply a cast from $\text{Int} \times \star \Rightarrow \text{Int} \times \text{Int}$ to this value, but the value's type and the source type of the cast don't match! In fact, in this example the result would be incorrect; we would get $42\langle \text{Int}! \rangle$ instead of 42.

There are several solutions to this problem, and they all require storing more information on the heap or as a separate map. Here we take the most straightforward approach of immediately updating the heap with casted values, that is, with values that are in the process of being cast.

We walk through the execution of the above example, explaining our rules for reducing casted values in the heap and showing snapshots of the heap. We use the following abbreviations.

$$\begin{array}{l} T_0 = \star \times \star \times \star \\ T_1 = \text{Int} \times \star \times \text{Ref } T_2 \\ T_2 = \text{Int} \times \text{Int} \times \star \\ c = \text{Int}^? \times \iota \times (\text{Ref } T_2)^? \end{array}$$

The first line of the program allocates a triple.

$$a_0 \mapsto (42\langle \text{Int!} \rangle, 7\langle \text{Int!} \rangle, 0\langle \text{Int!} \rangle) : T_0$$

The second line sets the third element to be a reference to itself.

$$a_0 \mapsto (42\langle \text{Int!} \rangle, 7\langle \text{Int!} \rangle, a_0\langle \text{Ref } T_0! \rangle) : T_0$$

The third line casts the reference to $\text{Ref } T_1$ via (CASTREF1).

$$a_0 \mapsto (42\langle \text{Int!} \rangle, 7\langle \text{Int!} \rangle, a_0\langle \text{Ref } T_0! \rangle)\langle c \rangle : T_1$$

We have a casted value in the heap that needs to be reduced. We apply (HCAST) and (PURECAST) to get

$$a_0 \mapsto (42, 7\langle \text{Int!} \rangle, a_0\langle \text{Ref } T_2 \rangle) : T_1$$

We cast address a_0 again, this time to $T_1 \sqcap T_2$, via rule (HDROP) and (CASTREF1).

$$a_0 \mapsto (42, 7\langle \text{Int!} \rangle, a_0\langle \text{Ref } T_2 \rangle)\langle \iota \times \text{Int?} \times \text{Ref } T_2 \rangle : \text{Int} \times \text{Int} \times \text{Ref } T_2$$

A few reductions via (HCAST) and (PURECAST) give us

$$a_0 \mapsto (42, 7, a_0\langle \text{Ref } T_2 \rangle) : \text{Int} \times \text{Int} \times \text{Ref } T_2$$

The final cast applied to a_0 is a no-op because the run-time type is already more precise than T_2 . So we reduce via (HCAST) and (CASTREF2) to:

$$a_0 \mapsto (42, 7, a_0) : \text{Int} \times \text{Int} \times \text{Ref } T_2$$

Even though we allow casted values on the heap, we require the normalization of all such casts before returning to the execution of the program. We distinguish between normal heaps of values, μ , and evolving heaps, ν , that may contain both values and casted values. Normal heaps are a subset of the evolving heaps.

Encoding permissive references The monotonic discipline and its run-time invariant-enforcement seems to restrict how developers can formulate their programs. It is natural to ask whether monotonic references are compatible with the flexibility that is expected in dynamic languages. In this section we show that the monotonic discipline admits permissive references through a syntactic discipline that can be conveniently provided to programmers.

Consider the following program that uses an allocated reference cell at two incompatible types, Int and Bool .

```

let x = ref (4 as ★) in
let y = (x as Ref Int) in
let z = (x as Ref Bool) in
!y;
z := true;
!z

```

Under the standard reference semantics, this program runs without incident, but under monotonic references it fails in the cast to `Ref Bool`. We can regain this flexibility under monotonic references via a disciplined use of `★` typed reference cells. Consider the following rewrite of this program:

```

let x = ref (4 as ★) in
let y = x in           // treat y like Ref Int
let z = x in           // treat z like Ref Bool
(!y) as Int;
(z := (true) as ★) as Bool;
(!z) as Bool

```

In this encoding, all references have type `Ref ★`, and typing is enforced only at dereferences and updates, using ascriptions. This program runs successfully under the monotonic semantics, but it would be tedious and error prone to insert these ascriptions by hand.

Luckily there is no need: we codify this permissive reference discipline by introducing a surface language that makes this convenient. We extend the expressions with *permissive references* `ref e`, and the types with a corresponding type `Ref T`. Consistency is extended so that permissive references have the same consistency properties as monotonic references, but permissive references are not consistent with monotonic references.

Finally we introduce a type-directed transformation $\Gamma \vdash e : T \rightsquigarrow e'$ that translates permissive references to monotonic references. The interesting cases are presented below.

$$\begin{array}{c}
\frac{x : \widetilde{\text{Ref}} T \in \Gamma}{\Gamma \vdash x : \widetilde{\text{Ref}} T \rightsquigarrow x} \quad \frac{\Gamma \vdash e : T \rightsquigarrow e'}{\Gamma \vdash \widetilde{\text{ref}} e : \widetilde{\text{Ref}} T \rightsquigarrow \text{ref}(e' \text{ as } \star)} \\
\frac{\Gamma \vdash e : \widetilde{\text{Ref}} T \rightsquigarrow e'}{\Gamma \vdash !e : T \rightsquigarrow (!e') \text{ as } T} \quad \frac{\Gamma \vdash e_1 : \widetilde{\text{Ref}} T_1 \rightsquigarrow e'_1 \quad \Gamma \vdash e_2 : T_2 \rightsquigarrow e'_2 \quad T_1 \sim T_2}{\Gamma \vdash e_1 := e_2 : T_1 \rightsquigarrow (e'_1 := (e'_2 \text{ as } \star)) \text{ as } T_1}
\end{array}$$

Note that the static semantics for permissive references enforces type consistency at assignments, even though the assigned value is ultimately cast to `★`. Furthermore, reference values translate to themselves, so object identity is preserved. However cast overhead is introduced at each dereference and update, so permissive references pay their own way with respect to performance.

If we revisit the initial example in this section and replace `ref` with `widetilde{ref}` and `Ref` with `widetilde{Ref}`, then this judgment translates the first program above into the second.

Proposition 1 (Translation). *If $\Gamma \vdash e : T \rightsquigarrow e'$ then $|\Gamma| \vdash e' : |T|$, Where $|\cdot|$ is the compatible extension of the equation $|\widetilde{\text{Ref}} T| = \text{Ref } \star$.*

This syntactic extension gives programmers access to both permissive references and monotonic references as desired.

$$\boxed{(T \Rightarrow T) = c}$$

$$\begin{array}{ll} (B \Rightarrow B) = \iota & (I \Rightarrow \star) = I! \\ (\star \Rightarrow \star) = \iota & (\star \Rightarrow I) = I? \end{array}$$

$$\begin{array}{l} (T_1 \rightarrow T_2) \Rightarrow (T'_1 \rightarrow T'_2) = (T'_1 \Rightarrow T_1) \rightarrow (T_2 \Rightarrow T'_2) \\ (T_1 \times T_2) \Rightarrow (T'_1 \times T'_2) = (T_1 \Rightarrow T'_1) \times (T_2 \Rightarrow T'_2) \\ \mathbf{Ref} T \Rightarrow \mathbf{Ref} T' = \mathbf{Ref} T' \end{array}$$

Fig. 8. Compile casts to monotonic coercions (without blame)

Permissive references are a useful abstraction for the programmer and provide strong guarantees. However, such guarantees are provided only as long as permissive references do not flow into monotonic references. Consider the program above (with permissive references) where the following code comes after the `let` statements.

```
let w1 = (x as  $\star$ ) in
let w2 = (w1 as Ref Bool) in
w2 := true;
```

The program places us in a same situation as the original program that the monotonic semantics could not run without error. This example shows an important syntactic discipline for programmers that want to employ the monotonic paradigm for gradual references: *permissive references should not flow into monotonic references.*

4 Type Safety for Monotonic References

We present the high-points of the type safety proof here. The full proof is mechanized in Isabelle 2013 and available on arxiv [Siek and Vitousek, 2013]. The semantics in the mechanized version differs from the semantics presented here in that it uses an abstract machine instead of a reduction semantics, as we found the mechanized proof easier to carry out on an abstract machine while the reduction semantics is more approachable.

We begin by lifting the precision relation to heap typings.

Definition 1 (Precision relation on heap typings). $\Sigma' \sqsubseteq \Sigma$ iff $\text{dom}(\Sigma') = \text{dom}(\Sigma)$ and $\Sigma(a) = T$ implies $\Sigma'(a) = T'$ where $T' \sqsubseteq T$.

Our first lemma below is important: expression typing is preserved when moving to a more precise heap typing.

Lemma 1 (Strengthening wrt. the heap typing). If $\Gamma; \Sigma \vdash e : T$ and $\Sigma' \sqsubseteq \Sigma$, then $\Gamma; \Sigma' \vdash e : T$.

Proof (Proof sketch). The interesting case is for addresses. We have

$$\frac{\Sigma(a) \sqsubseteq T}{\Gamma; \Sigma \vdash a : T}$$

From $\Sigma' \sqsubseteq \Sigma$ and transitivity of \sqsubseteq , we have $\Sigma'(a) \sqsubseteq T$. Therefore $\Gamma; \Sigma' \vdash a : T$.

The definition of well-typed heaps is standard.

Definition 2 (Well-typed heaps). A heap ν is well-typed with respect to heap typing Σ , written $\Sigma \vdash \nu$, iff $\forall a \bar{T}. \Sigma(a) = T$ implies $\nu(a) = cv : T$ and $\emptyset; \Sigma \vdash cv : T$ for some cv .

From the strengthening lemma, we have the following corollary.

Corollary 1 (Monotonic heap update). If $\Sigma \vdash \nu$ and $\Sigma(a) = T$ and $T' \sqsubseteq T$ and $\emptyset; \Sigma \vdash cv : T'$, then $\Sigma(a \mapsto T') \vdash \nu(a \mapsto cv : T')$.

Proof (sketch). Let $\Sigma' = \Sigma(a \mapsto T')$. From $T' \sqsubseteq T$ we have $\Sigma' \sqsubseteq \Sigma$, so by Lemma 1 we have $\emptyset; \Sigma' \vdash cv : T'$ and $\Sigma' \vdash \nu$. Thus, $\Sigma(a \mapsto T') \vdash \nu(a \mapsto cv : T')$.

Lemma 2 (Progress and Preservation). Suppose $\emptyset; \Sigma \vdash e : T$ and $\Sigma \vdash \nu$. Exactly one of the following holds:

1. (a) e is a value, or
 (b) $e = \mathbf{error}$, or
 (c) $e, \nu \longrightarrow e', \nu'$ for some e' and ν' .
2. for all e', ν' , if $e, \nu \longrightarrow e', \nu'$ then $\emptyset; \Sigma' \vdash e' : T$ and $\Sigma' \vdash \nu'$ and $\Sigma' \sqsubseteq \Sigma$ for some Σ' .

Theorem 1 (Type Safety). Suppose $\emptyset; \Sigma \vdash e : T$ and $\Sigma \vdash \nu$. Exactly one of the following holds:

1. $e, \nu \longrightarrow^* v, \nu'$ and $\emptyset; \Sigma' \vdash v : T$ for some Σ' , or
2. $e, \nu \longrightarrow^* \mathbf{error}, \nu'$, or
3. e diverges.

Proof. If e diverges we immediately conclude the proof. Otherwise, suppose e does not diverge. Then $e, \nu \longrightarrow^* e', \nu'$ and e' cannot reduce. We proceed by induction on the length $e, \nu \longrightarrow^* e', \nu'$, and use Lemma 2 to conclude.

5 Monotonic References with Blame

We turn to the challenge of designing blame tracking for monotonic references, presenting several examples that motivate and provide intuitions for the design. The later part of this section presents the dynamic semantics of monotonic references with blame tracking.

Consider the following example in which we allocate a reference of dynamic type and then, separately, cast from `Ref \star` to `Ref Int` and to `Ref Bool`.

```
let r0 = ref (42 asℓ1  $\star$ ) in
let r1 = r0 asℓ2 Ref Int in
let r2 = r0 asℓ3 Ref Bool in
!r2
```

$$\begin{array}{c}
\frac{}{B <: B} \quad \frac{}{T <: \star} \quad \frac{}{\mathbf{Ref} T <: \mathbf{Ref} T} \\
\frac{T'_1 <: T_1 \quad T_2 <: T'_2}{T_1 \rightarrow T_2 <: T'_1 \rightarrow T'_2} \quad \frac{T_1 <: T'_1 \quad T_2 <: T'_2}{T_1 \times T_2 <: T'_1 \times T'_2}
\end{array}$$

Fig. 9. Subtyping relation

With monotonic references, the cast at ℓ_3 triggers an error, because `Int` and `Bool` are inconsistent. But what blame labels should the error message include? Is it only the fault of ℓ_3 ? Not really; because ℓ_3 would not cause an error if it were not for the cast at ℓ_2 . The casts at ℓ_2 and ℓ_3 disagree with each other regarding the type of the heap cell, so we blame both. The result of this program is `blame { ℓ_2, ℓ_3 }`.

Next consider an example in which we allocate a reference at type `Ref Int`, cast it to `Ref *`, and then attempt to write a Boolean.

```

let r0 = ref 42 in
let r1 = r0 asℓ1 Ref * in
r1 :=ℓ3 (true asℓ2 *)

```

The update on the third line triggers an error, and we have three possible locations to blame: ℓ_1 , ℓ_2 , and ℓ_3 . The cast at ℓ_2 is from `Bool` to `*`, which is harmless. There is no cast at ℓ_3 , we are just writing a value of type `*` to a reference of type `Ref *`. The real culprit here is ℓ_1 , which casts from `Ref Int` to `Ref *`, thereby opening up the potential for the later cast error. Naïvely, this looks like an upcast, but a proper treatment of subtyping for references makes references invariant. So we have `Ref Int` $\not<$: `Ref *` and the result of this program is `blame { ℓ_1 }`. Figure 9 presents the subtyping relation.

We consider a pair of examples below that differ only on the fourth line. We allocate a reference to a pair at type `Ref (* × *)` then cast it to `Ref (Int × *)` and to `Ref (* × Int)`. In the first example, we update through the original reference, writing a Boolean and integer, whereas in the second example we write an integer and a Boolean. Here is the first example:

```

let r0 = ref (1 asℓ1 *, 2 asℓ2 *) in
let r1 = r0 asℓ3 Ref (Int × *) in
let r2 = r0 asℓ4 Ref (* × Int) in
r0 := (true asℓ5 *, 2 asℓ6 *);
fst !r0

```

and here is the second example, just showing the fourth line:

```

...
r0 := (1 asℓ7 *, true asℓ8 *);
...

```

The first example should produce `blame {ℓ3}` while the second example should produce `blame {ℓ4}`, but the challenge is how can we associate multiple blame labels with the same heap cell?

We take inspiration from Siek and Wadler [2010] and use *labeled types* for our run-time type information. With a labeled type, each type constructor within the type can be labeled with a blame label. Figure 10 gives the syntax of labeled types and operations on them, which we shall explain later in this section. In the above examples, the run-time type information for the heap cell evolves as follows:

$$(\star \times^\emptyset \star) \Rightarrow (\mathbf{Int}^{\ell_3} \times^\emptyset \star) \Rightarrow (\mathbf{Int}^{\ell_3} \times^\emptyset \mathbf{Int}^{\ell_4})$$

In the first example, when we write `true` into the first element of the pair, the cast to `Int` fails and blames ℓ_3 , as desired. In the second example, when we write `true` into the second element, the cast to `Int` fails and blames ℓ_4 , as desired.

Our next example brings up a somewhat ambiguous situation. We allocate a reference at type `Ref \star` , cast it to `Ref Int` twice, then write a Boolean.

```
let r0 = ref (42 asℓ1  $\star$ ) in
let r1 = r0 asℓ2 Ref Int in
let r2 = r0 asℓ3 Ref Int in
r0 := (true asℓ4  $\star$ )
```

Should we blame ℓ_2 or ℓ_3 ? In some sense, they are both just as guilty and the ideal would be to blame them both. On the other hand, maintaining potentially large sets of blame labels would induce some space overhead. Our design instead blames the first cast with respect to execution order, in this case ℓ_2 .

For our final example, we adapt the above example to have a function in the heap cell so that we can consider the behavior to the left of the arrow.

```
let r0 = ref (λx:  $\star$ . true) in
let r1 = r0 asℓ1 Ref (Int → Bool) in
let r2 = r0 asℓ2 Ref (Int → Bool) in
r0 := λx: Int. zero?(x);
!r0 (true asℓ3  $\star$ )
```

The run-time type information for the heap cell evolves in the following way:

$$(\star \rightarrow^\emptyset \mathbf{Bool}^\emptyset) \Rightarrow (\mathbf{Int}^{\ell_1} \rightarrow^\emptyset \mathbf{Bool}^\emptyset) \Rightarrow (\mathbf{Int}^{\ell_1} \rightarrow^\emptyset \mathbf{Bool}^\emptyset)$$

The function application on the last line of the example triggers a cast error, with the blame going to ℓ_1 , again because we wish to blame the first cast with respect to execution order. However, to obtain this semantics some care must be taken. On the second cast, we merge the labeled type for the second cast with the current run-time type information:

$$(\mathbf{Int}^{\ell_1} \rightarrow^\emptyset \mathbf{Bool}^\emptyset) \Delta (\mathbf{Int}^{\ell_2} \rightarrow^\emptyset \mathbf{Bool}^\emptyset)$$

If we were to use the composition function from Siek and Wadler [2010], the result would be $\mathbf{Int}^{\ell_2} \rightarrow^\emptyset \mathbf{Bool}^\emptyset$ because that composition function is contravariant for

function parameters. Here we instead want to be covariant on function parameters, so the result is $\text{Int}^{\ell_1} \rightarrow^{\emptyset} \text{Bool}^{\emptyset}$. We define a new function for merging labeled types, Δ , in Figure 10.

5.1 Semantics of monotonic references with blame

Armed with the intuitions from the above examples, we discuss the semantics of monotonic references with blame, defined in Figures 12 and 13. The semantics is largely similar to the semantics without blame except that the run-time type information is represented as labeled types and we replace the functions, such as meet (\sqcap) that operate on types, with functions such as merge (Δ) that operate on labeled types.

Proposition 2 (Meet is the erasure of merge).

If $|P_1| \sim |P_2|$, then $|P_1 \Delta P_2| = |P_1| \sqcap |P_2|$.

If $|P_1| \not\sim |P_2|$, then $P_1 \Delta P_2 = \perp^L$ for some L .

As discussed with the example above, the definition of $P_1 \Delta P_2$ takes into account that P_1 is temporally prior to P_2 and should therefore take precedence with respect to blame responsibility. We use the auxiliary function $p \Delta q$ to choose between two optional labels, returning the first if it is present and the second otherwise.

When we cast a reference via rule (CASTR1B), we need to update the heap cell from labeled type P_1 to P_3 . We accomplish this with a new operator $P_1 \Rightarrow P_3$ that produces a coercion. The most interesting line of its definition is for reference types. There we use a different operator, $P \Leftrightarrow Q$, that produces a labeled type and captures the bidirectional read/write nature of mutable references.

The definitions of Δ , \Rightarrow , and \Leftrightarrow need to percolate errors, which we write as \perp^L where L is a set of blame labels. We use “smart” constructors $\hat{\hookrightarrow}$, $\hat{\times}$, and $\hat{\text{Ref}}$ that return \perp^L if either argument is \perp^L (with precedent to the left if both arguments are errors), but otherwise act like the underlying constructor.

In the rule for allocation, we initialize the RTTI to T^{\emptyset} . (Figure 11 defines converting a type to a labeled type.) In the rule for a dynamic dereference, (DYNDRFMB), we cast from the reference’s run-time labeled type to T by promoting T to the labeled type T^{\emptyset} and then applying the \Rightarrow function to cast between labeled types, so we have $\mu(a)_{\text{rtti}} \Rightarrow T^{\emptyset}$. Suppose that $\mu(a)_{\text{rtti}}$ is Ref Int^{ℓ} and T is $\text{Ref } \star$. Then the coercion we apply during the dereference is $\text{Int}^{\ell!}$; so our injection coercions contain labeled types. The rule for dynamic update, (DYNUPDMB), is dual: we perform the cast $T^{\emptyset} \Rightarrow \mu(a)_{\text{rtti}}$.

Because our injection and projection coercions contain labeled types, the (COLLAPSE) rule becomes

$$v\langle P_1! \rangle \langle P_2? \rangle \longrightarrow_c v\langle P_1 \Rightarrow P_2 \rangle \quad \text{if } |P_1| \sim |P_2|$$

We make similar changes to the (CONFLICT) rule.

Figure 11 defines the compilation of casts to monotonic coercions. Compared to the compilation without blame (Figure 8), there are three differences. The

Optional labels $p, q ::= \emptyset \mid \{\ell\}$
 Label sets $L ::= \emptyset \mid \{\ell\} \mid \{\ell_1, \ell_2\}$
 Labeled types $P, Q ::= B^p \mid P \rightarrow^p P \mid P \times^p P \mid \mathbf{Ref}^p P \mid \star$

Erase labels $|P| = T$

$|B^p| = B \quad |P \rightarrow^p Q| = |P| \rightarrow |Q| \quad |P \times^p Q| = |P| \times |Q| \quad |\mathbf{Ref}^p P| = \mathbf{Ref} |P| \quad |\star| = \star$

Top label $lab(P) = L$

$lab(B^p) = p \quad lab(P \rightarrow^p Q) = p \quad lab(P \times^p Q) = p \quad lab(\mathbf{Ref}^p P) = p \quad lab(\star) = \emptyset$

Merge optional labels $p \Delta p = p$

$\{\ell\} \Delta q = \{\ell\} \quad \emptyset \Delta q = q$

Merge labeled types $P \Delta P = P \text{ or } \perp^L$

$$\begin{aligned} B^p \Delta B^q &= B^{p \Delta q} \\ P \Delta \star &= P \quad \star \Delta Q = Q \\ (P \rightarrow^p P') \Delta (Q \rightarrow^q Q') &= (P \Delta Q) \rightarrow^{p \Delta q} (P' \Delta Q') \\ (P \times^p P') \Delta (Q \times^q Q') &= (P \Delta Q) \hat{\times}^{p \Delta q} (P' \Delta Q') \\ \mathbf{Ref}^p P \Delta \mathbf{Ref}^q Q &= \hat{\mathbf{Ref}}^{p \Delta q} (P \Delta Q) \\ P \Delta Q &= \perp^{lab(P) \cup lab(Q)} \quad \text{otherwise} \end{aligned}$$

Bidirectional cast between labeled types $P \Leftrightarrow P = P \text{ or } \perp^L$

$$\begin{aligned} B^p \Leftrightarrow B^q &= B^\emptyset \\ P \Leftrightarrow \star &= P \quad \star \Leftrightarrow Q = Q \\ (P \rightarrow^p P') \Leftrightarrow (Q \rightarrow^q Q') &= (P \Leftrightarrow Q) \rightarrow^\emptyset (P' \Leftrightarrow Q') \\ (P \times^p P') \Leftrightarrow (Q \times^q Q') &= (P \Leftrightarrow Q) \hat{\times}^\emptyset (P' \Leftrightarrow Q') \\ \mathbf{Ref}^p P \Leftrightarrow \mathbf{Ref}^q Q &= \hat{\mathbf{Ref}}^\emptyset (P \Leftrightarrow Q) \\ P \Leftrightarrow Q &= \perp^{lab(P) \cup lab(Q)} \quad \text{otherwise} \end{aligned}$$

Cast between labeled types $P \Rightarrow P = c \text{ or } \perp^L$

$$\begin{aligned} B^p \Rightarrow B^q &= \iota \quad \star \Rightarrow \star = \iota \\ P \Rightarrow \star &= P! \quad \star \Rightarrow Q = Q? \\ (P \rightarrow^p P') \Rightarrow (Q \rightarrow^q Q') &= (Q \Rightarrow P) \rightarrow^\emptyset (P' \Rightarrow Q') \\ (P \times^p P') \Rightarrow (Q \times^q Q') &= (P \Rightarrow Q) \hat{\times}^\emptyset (P' \Rightarrow Q') \\ \mathbf{Ref}^p P \Rightarrow \mathbf{Ref}^q Q &= \hat{\mathbf{Ref}}^\emptyset (P \Rightarrow Q) \\ P \Rightarrow Q &= \perp^{lab(P) \cup lab(Q)} \quad \text{otherwise} \end{aligned}$$

Fig. 10. Labeled types and their operations

$$(T \Rightarrow^\ell T) = c$$

$$\begin{aligned} (B \Rightarrow^\ell B) &= \iota & (T \Rightarrow^\ell \star) &= T^{\emptyset!} \\ (\star \Rightarrow^\ell \star) &= \iota & (\star \Rightarrow^\ell T) &= T^{\ell?} \end{aligned}$$

$$\begin{aligned} (T_1 \rightarrow T_2) \Rightarrow^\ell (T'_1 \rightarrow T'_2) &= (T'_1 \Rightarrow^\ell T_1) \rightarrow (T_2 \Rightarrow^\ell T'_2) \\ (T_1 \times T_2) \Rightarrow^\ell (T'_1 \times T'_2) &= (T_1 \Rightarrow^\ell T'_1) \times (T_2 \Rightarrow^\ell T'_2) \\ \mathbf{Ref} T_1 \Rightarrow^\ell \mathbf{Ref} T_2 &= \mathbf{Ref} (T_1^\ell \Leftrightarrow T_2^\ell) \end{aligned}$$

Add labels to a type

$$T^\ell = P$$

$$\begin{aligned} B^\ell &= B^\ell & (T_1 \rightarrow T_2)^\ell &= T_1^\ell \rightarrow^\ell T_2^\ell & (T_1 \times T_2)^\ell &= T_1^\ell \times^\ell T_2^\ell \\ (\mathbf{Ref} T)^\ell &= \mathbf{Ref}^\ell T^\ell & \star^\ell &= \star \end{aligned}$$

Fig. 11. Compile casts to monotonic coercions (with blame)

Expressions	$e ::= \dots \mid \mathbf{blame} L$
Coercions	$c ::= \iota \mid P? \mid P! \mid c \rightarrow c \mid c \times c \mid c ; c \mid \mathbf{Ref} P$
Values	$v ::= k \mid \lambda x. e \mid (v, v) \mid v \langle P! \rangle \mid a$
Casted Values	$cv ::= v \mid cv \langle c \rangle \mid (cv, cv)$
Heap	$\mu ::= \emptyset \mid \mu(a \mapsto v : P)$
Evolving Heap	$\nu ::= \emptyset \mid \nu(a \mapsto cv : P)$

Fig. 12. Syntax for monotonic references with blame

first two concern injection and projection coercions: instead of only having a blame label on projections we have labeled types inside both injections and projections, as noted above. In the compilation of a cast labeled ℓ , we generate a labeled type for the injection from T by adding the empty label to T , and for the projection to T by adding ℓ to T . The third difference is in the formation of the reference coercion. Instead of simply taking the target type, we use the bidirectional operator \Leftrightarrow . Recall the second example of this section in which we blamed the cast from $\mathbf{Ref} \text{Int}$ to $\mathbf{Ref} \star$. By using \Leftrightarrow , the resulting coercion is $\mathbf{Ref} \text{Int}^{\ell_1}$ instead of $\mathbf{Ref} \star$.

6 The Blame-Subtyping Theorem

The blame-subtyping theorem pin-points the source of cast errors in gradually-typed programs. The blame-subtyping theorem states that if a program results in a cast error, $\mathbf{blame} L$, then the blame labels in L identify the location of implicit casts that did not respect subtyping. That is, the blame labels that occur in a safe implicit cast, $T_1 \Rightarrow T_2$ where $T_1 <: T_2$, can never be blamed.

We prove the blame-subtyping theorem via a preservation-style proof in which we preserve the $e\text{safe} \ell$ predicate [Wadler and Findler, 2009]. This proof is

Coercion typing

$$\boxed{c : T \Rightarrow T}$$

$$\frac{}{P? : \star \Rightarrow |P|} \quad \frac{}{P! : |P| \Rightarrow \star} \quad \dots$$

Pure cast reduction rules

$$\boxed{e \longrightarrow_c e}$$

$$\dots \quad v\langle P_1! \rangle \langle P_2? \rangle \longrightarrow_c v\langle P_1 \Rightarrow P_2 \rangle \quad \text{if } |P_1| \sim |P_2| \quad (\text{COLLAPSE})$$

$$v\langle P_1! \rangle \langle P_2? \rangle \longrightarrow_c \mathbf{blame} L \quad \text{if } P_1 \Rightarrow P_2 = \perp^L \quad (\text{CONFLICT})$$

Cast reduction rules

$$\boxed{e, \nu \longrightarrow_{cr} e, \nu}$$

$$\frac{e \longrightarrow_c e'}{e, \nu \longrightarrow_{cr} e', \nu} \quad (\text{PCASTB})$$

$$\frac{\nu(a) = cv : P_1 \quad P_3 = P_1 \Delta P_2 \quad |P_3| \neq |P_1| \quad cv' = cv\langle P_1 \Rightarrow P_3 \rangle}{a\langle \mathbf{Ref} P_2 \rangle, \nu \longrightarrow_{cr} a, \nu(a \mapsto cv' : P_3)} \quad (\text{CASTR1B})$$

$$\frac{\nu(a) = cv : P_1 \quad P_1 = P_1 \Delta P_2}{a\langle \mathbf{Ref} P_2 \rangle, \nu \longrightarrow_{cr} a, \nu} \quad (\text{CASTR2B})$$

$$\frac{\nu(a) = cv : P_1 \quad P_1 \Delta P_2 = \perp^L}{a\langle \mathbf{Ref} P_2 \rangle, \nu \longrightarrow_{cr} \mathbf{blame} L, \nu} \quad (\text{CASTR3B})$$

Program reduction rules

$$\boxed{e, \mu \longrightarrow_e e, \mu}$$

$$\mathbf{ref}_T v, \mu \longrightarrow_e a, \mu(a \mapsto v : T^0) \quad \text{if } a \notin \text{dom}(\mu)$$

$$!a, \mu \longrightarrow_e \mu(a)_{\text{val}}, \mu \quad (\text{DEREFMB})$$

$$!a@T, \mu \longrightarrow_e \mu(a)_{\text{val}} \langle \mu(a)_{\text{rtti}} \Rightarrow T^0 \rangle, \mu \quad (\text{DYNDRFMB})$$

$$a := v, \mu \longrightarrow_e a, \mu(a \mapsto v : \mu(a)_{\text{rtti}}) \quad (\text{UPDMB})$$

$$a := v@T, \mu \longrightarrow_e a, \mu(a \mapsto cv : \mu(a)_{\text{rtti}}) \quad (\text{DYNUPDMB})$$

where $cv = v\langle T^0 \Rightarrow \mu(a)_{\text{rtti}} \rangle$

For $X \in \{cr, e\}$:

$$\frac{e, \nu \longrightarrow_X e', \nu'}{F[e], \nu \longrightarrow_X F[e'], \nu'} \quad \frac{}{F[\mathbf{blame} L], \nu \longrightarrow_X \mathbf{blame} L, \nu}$$

State reduction rules

$$\boxed{e, \nu \longrightarrow_e e, \nu}$$

$$\frac{e, \mu \longrightarrow_X e', \nu \quad X \in \{cr, e\}}{e, \mu \longrightarrow e', \nu} \quad \frac{\nu(a) = cv : P \quad cv, \nu \longrightarrow_{cr} \mathbf{blame} L, \nu'}{e, \nu \longrightarrow \mathbf{blame} L, \nu'}$$

$$\frac{\nu(a) = cv : P \quad cv, \nu \longrightarrow_{cr} cv', \nu' \quad |\nu'(a)_{\text{rtti}}| = |P|}{e, \nu \longrightarrow e, \nu'(a \mapsto cv' : P)}$$

$$\frac{\nu(a) = cv : P \quad cv, \nu \longrightarrow_{cr} cv', \nu' \quad |\nu'(a)_{\text{rtti}}| \neq |P|}{e, \nu \longrightarrow e, \nu'}$$

Fig. 13. Monotonic references with blame

conducted on the coercion calculus, so to relate the result back to the gradually-typed λ -calculus, we need a theorem concerning the relationship between subtyping and coercion blame safety, Theorem 2. Recall that subtyping is defined in Figure 9 and compilation to coercions is defined in Figure 11.

Theorem 2 (Blame-Subtyping Theorem for coercion calculus). *For all T_1, T_2 , and ℓ , it holds that $T_1 <: T_2$ iff $(T_1 \Rightarrow^\ell T_2)$ safe ℓ .*

Lemma 3 (Preservation of blame safety).

For all e, e', ν, ν' , and ℓ , if e, ν safe ℓ and $e, \nu \longrightarrow e', \nu'$ then e', ν' safe ℓ .

We now move away from the coercion calculus and prove these important results on the gradually typed λ calculus with references. This latter language is indeed the one that programmers are expected to use. The following definitions will help to recast the results into the setting of the gradually typed language.

Definition 3 (Casts for a label in an expression). *Let e be an expression and ℓ a label, we say that e contains the cast $T_1 \Rightarrow T_2$ for ℓ whenever, in the derivation of $\Gamma \vdash e \rightsquigarrow e' : T$, there is the creation of a coercion via $T_1 \Rightarrow^\ell T_2$.*

Definition 4 (Blame safety for gradually-typed expressions). *A gradually-typed expression e is safe for ℓ if all the casts contained in e labeled ℓ respect subtyping.*

We now have all the ingredients to state and prove one of the main contributions of the paper, i.e. the Blame-Subtyping Theorem for the gradually-typed λ calculus with references.

Lemma 4 (Translation preserves blame safety). *If e safe ℓ and $\Gamma \vdash e \rightsquigarrow e' : T$, then e' safe ℓ .*

Proof. The proof is a straightforward induction on $\Gamma \vdash e \rightsquigarrow e' : T$.

Theorem 3 (Blame-Subtyping Theorem). *For all e, e', T_1, T_2, ℓ , if $\emptyset \vdash e \rightsquigarrow e' : T$, e safe ℓ , and $e', \emptyset \longrightarrow \text{blame } L, \nu$, then $\ell \notin L$.*

Proof. From the assumptions we have e' safe ℓ by Lemma 4. Then we conclude by applying the Blame-Subtyping Theorem for the coercion calculus.

7 Implementation concerns w.r.t. strong updates

The monotonic semantics for references performs in-place updates to the heap with values of different type. In languages where values have uniform size, like many functional and object-oriented languages, this does not pose a problem. However, for languages where values may have different sizes, in-place updates pose a problem. This issue can be addressed using an approach inspired by garbage collection techniques. When the semantics is to update a cell with a larger value than the current one, the implementation allocates a new piece

of memory and places a forwarding pointer in the old location. When reading and writing through dynamic references, the implementation must check for and follow the forwarding pointers. However, when reading and writing through fully-static references, the implementation does not need to consider forwarding pointers because fully-static heap cells never move. Then during a garbage collection, the implementation can collapse sequences of forwarding pointers to reduce overhead in subsequent execution.

8 Related Work

Here we mention related work that is not discussed in the introduction or elsewhere in the paper.

The casts and coercions studied in this paper bear many similarities with contracts [Findler and Felleisen, 2002]. Racket [Flatt and PLT, 2014] provides contracts for mutable values in the form of impersonators [Strickland et al., 2012], which, for our purposes, can be viewed as implementing the standard semantics of Herman et al. [2007], as we saw in Section 2.

Fähndrich and Leino [2003] introduce a technique similar to monotonic references with their monotonic typestate. In this design, objects may flow from less restrictive to more restrictive typestates, but not vice versa. Unlike monotonic references, which require runtime checks due to the existence of dynamically-typed regions of code, their system enforces monotonicity statically.

Gradual typing was added to C^\sharp with the addition of the `dynamic` type. Bierman et al. [2010] define a formal model of C^\sharp , named FC_4^\sharp , and present an operational semantics. The semantics is similar to that of Swamy et al. [2014] in that they use an RTTI-based approach and subtype checks to implement casts.

9 Conclusion

We have presented a new design for gradually-typed mutable references, called monotonic references, the first to incur zero-overhead for reference accesses in statically typed code while maintaining the full expressiveness of a gradual type system. We defined a dynamic semantics for monotonic references and presented a mechanized proof of type safety. Further, we defined blame tracking based on using labeled types in the run-time type information and proved the blame-subtyping theorem.

Bibliography

- M. Abadi, L. Cardelli, B. Pierce, and G. Plotkin. Dynamic typing in a statically-typed language. In *Symposium on Principles of programming languages*, 1989.
- G. Bierman, E. Meijer, and M. Torgersen. Adding dynamic types to C#. In *European Conference on Object-Oriented Programming*, 2010.
- J. T. Boyland. The problem of structural type tests in a gradual-typed language. In *Foundations of Object Oriented Languages*, FOOL '14, pages 675–681. ACM, 2014.
- M. Fähndrich and K. R. M. Leino. Heap monotonic typestate. In *International Workshop on Alias Confinement and Ownership*, 2003.
- R. B. Findler and M. Felleisen. Contracts for higher-order functions. In *International Conference on Functional Programming*, ICFP, pages 48–59, 2002.
- M. Flatt and PLT. The Racket reference 6.0. Technical report, PLT Inc., 2014. <http://docs.racket-lang.org/reference/index.html>.
- F. Henglein. Dynamic typing: syntax and proof theory. *Science of Computer Programming*, 22(3):197–230, June 1994.
- D. Herman, A. Tomb, and C. Flanagan. Space-efficient gradual typing. In *Trends in Functional Prog. (TFP)*, page XXVIII, April 2007.
- A. Rastogi, N. Swamy, C. Fournet, G. Bierman, and P. Vekris. Safe & efficient gradual typing for TypeScript. Technical Report MSR-TR-2014-99, 2014.
- J. G. Siek and W. Taha. Gradual typing for functional languages. In *Scheme and Functional Programming Workshop*, pages 81–92, September 2006.
- J. G. Siek and M. M. Vitousek. Monotonic references for gradual typing. *Computing Research Repository*, 2013. URL <http://arxiv.org/abs/1312.0694>.
- J. G. Siek and P. Wadler. Threesomes, with and without blame. In *Symposium on Principles of Programming Languages*, POPL, pages 365–376, January 2010.
- J. G. Siek, M. M. Vitousek, M. Cimini, and J. T. Boyland. Refined criteria for gradual typing. Under review for publication at SNAPL 2015, 2015.
- T. S. Strickland, S. Tobin-Hochstadt, R. B. Findler, and M. Flatt. Chaperones and impersonators: run-time support for reasonable interposition. OOPSLA, 2012.
- N. Swamy, C. Fournet, A. Rastogi, K. Bhargavan, J. Chen, P.-Y. Strub, and G. Bierman. Gradual typing embedded securely in JavaScript. In *Symposium on Principles of Programming Languages (POPL)*, January 2014.
- S. Tobin-Hochstadt and M. Felleisen. Interlanguage migration: From scripts to programs. In *Dynamic Languages Symposium*, 2006.
- M. M. Vitousek, J. G. Siek, A. Kent, and J. Baker. Design and evaluation of gradual typing for Python. In *Dynamic Languages Symposium*, 2014.
- P. Wadler and R. B. Findler. Well-typed programs can't be blamed. In *European Symposium on Programming*, ESOP, pages 1–16, March 2009.
- T. Wrigstad, F. Z. Nardelli, S. Lebesne, J. Östlund, and J. Vitek. Integrating typed and untyped code in a scripting language. In *Symposium on Principles of Programming Languages*, POPL, pages 377–388, 2010.