

# Space-Efficient Manifest Contracts

Michael Greenberg

University of Pennsylvania

mgree@seas.upenn.edu

## Abstract

*Gradual types* mediate the interaction between dynamic and simple types, offering an easy transition from scripts to programs; gradual types allow programmers to evolve prototype scripts into fully fledged, deployable programs. Similarly, *contracts* and *refinement types* mediate the interaction between simple types and more precise types, offering an easy transition from programs to robust, verified programs. A *full-spectrum* language with both gradual and refinement types offers low-level support for the development of programs throughout their lifecycle, from prototype script to verified program.

One attractive formulation of languages with gradual or refinement types uses *casts* to represent the runtime checks necessary for type changes (from dynamic to simple types, and from simple types to refinement types). Briefly, a cast  $\langle T_1 \Rightarrow T_2 \rangle e$  takes a term  $e$  from type  $T_1$  to  $T_2$ , possibly wrapping or tagging  $e$  in the process—or even failing, if  $e$  doesn’t meet the criteria of the type  $T_2$ . Casts are attractive because they offer a unified view of changes in type, have straightforward operational semantics, and enjoy an interesting and fruitful relationship with subtyping.

One longstanding problem with casts is space efficiency: casts in their naïve formulation can consume unbounded amounts of space at runtime both through excessive wrapping as well as through tail-recursion-breaking stack growth. Prior work [20, 21, 33, 35] offers space-efficient solutions exclusively in the domain of gradual types. In this paper, we define a new full-spectrum language that is (a) more expressive than prior languages, and (b) space efficient. We are the first to obtain space-efficient refinement types. Our approach to space efficiency is based on the *coercion* calculi of Herman et al. [20] and Henglein’s work [19], though our explicitly enumerated canonical coercions and our straightforward merge operator are a novel approach to coercions with a simpler theory. We show that space efficiency avoids some checks, failing and diverging less often than naïve calculi—but the two are otherwise observationally equivalent.

**Categories and Subject Descriptors** D.3.1 [Programming Languages]: Formal Definitions and Theory; D.2.4 [Software Engineering]: Software/Program Verification—Programming by contract

**General Terms** Languages, Theory

**Keywords** Contract, refinement type, dynamic checking, precondition, postcondition, space efficiency, coercion, dynamic typing, gradual typing

## 1. Introduction

The promise of a single language admitting the full development cycle—from a small script to a more manageable, statically typed program to a robust, verified system—has held great allure for some time. Prior attempts to fulfill this promise have attacked the problem piecemeal: script to program, and program to verified program. On the one hand, work in the script-to-program category goes back at least to Abadi et al.’s work with type Dynamic, with more recent work falling under the Siek and Taha’s rubric of “gradual typing” [1, 2, 7, 10, 20, 26, 33–35, 37, 38]. On the other hand, program verification is an enormous field in its own right; in this paper, we will focus on dynamic enforcement methods. These methods, often called *contracts*, go back at least as far as Eiffel [3, 11, 14, 16–18, 27, 31]. More recent work takes a type-oriented, or *manifest*, approach to contracts, allowing so-called *refinement types* of the form  $\{x:T \mid e\}$ , inhabited by values  $v$  that satisfy the *predicate*  $e$ , i.e.  $e[v/x] \rightarrow^* \text{true}$ . For example,  $\{x:\text{Int} \mid x \neq 0\}$  denotes the non-zero integers. Dependency increases the expressiveness of refinement type systems. Consider the type  $(x:\text{Real}) \rightarrow \{y:\text{Real} \mid |x - y^2| < \epsilon\}$ . This type specifies the square root function: for any real number  $x$ , functions in this type produce a  $y$  such that  $y^2$  is within  $\epsilon$  of  $x$ . Note that the type of the result depends on the input *value*.

Over the last decade, the state of the art combining these two paradigms—and gradual types and manifest contracts in particular—has steadily progressed [4, 25, 30, 39]. Starting with Sage [25] and continuing with Wadler and Findler [39], many languages have expressed the interactions between dynamic and simple and between simple and refinement types using a single syntactic form: the cast. Written  $\langle T_1 \Rightarrow T_2 \rangle$ , we read the cast form as “cast from type  $T_1$  to type  $T_2$ ”.

Casts are promising. They offer a unified view of changes in type information, have straightforward operational semantics, and enjoy a fruitful relationship with subtyping (see [3, 17, 24, 33, 35, 39]). Our language will derive its semantics from a full-spectrum language with casts. Before we can continue, we must explain how casts work: with type dynamic, with refinement types, and the most interesting part—between function types.

On the dynamic side, casts go into and out of  $?$ , the dynamic type. A cast of the form  $\langle \text{Int} \Rightarrow ? \rangle 5$  asks for the number 5, which has type  $\text{Int}$ , to be treated as a value of type  $?$ , the dynamic type. The operational semantics of such a cast will mark the value with a tag, as in  $5_{\text{Int}}$ . Similarly, a cast of the form  $\langle ? \Rightarrow \text{Int} \rangle 5_{\text{Int}}$  will project the tagged integer out of type dynamic, yielding the original value 5. In the case where the cast’s argument isn’t tagged correctly, the cast must raise an error. Consider the term  $\langle ? \Rightarrow \text{Int} \rangle \text{true}_{\text{Bool}}$ . It tries to project an  $\text{Int}$  out of type dynamic, but the dynamic value

is really a Bool—a type error. All we can reasonably do is abort the program, evaluating to the uncatchable exception fail.<sup>1</sup>

Casting between refinement types works similarly:  $\langle \text{Int} \Rightarrow \{x:\text{Int} \mid x \neq 0\} \rangle 5$  must first check that  $(x \neq 0)[5/x] \rightarrow^* \text{true}$ , i.e., that the refinement’s predicate is satisfied. If so, it will yield a tagged value, just like before:  $5_{\{x:\text{Int} \mid x \neq 0\}!}$ . (The exclamation point on the Int! tag represents *injection* into type dynamic, while the question mark on the  $\{x:\text{Int} \mid x \neq 0\}?$  tag represents the successful *checking* of the predicate.) If the predicate should fail—returning false or failing some nested predicate check—then the check will return fail and the entire program will fail, as well.

Casts on functions are the most interesting: values with functional casts on them, like  $\langle T_{11} \rightarrow T_{12} \Rightarrow T_{21} \rightarrow T_{22} \rangle v_1$ , are themselves values; they are *wrapped* with a *function proxy*. When such wrapped values are applied to a value  $v_2$ , the cast unfolds:

$$\langle \langle T_{11} \rightarrow T_{12} \Rightarrow T_{21} \rightarrow T_{22} \rangle v_1 \rangle v_2 \rightarrow \langle T_{12} \Rightarrow T_{22} \rangle (v_1 \langle \langle T_{21} \Rightarrow T_{11} \rangle v_2 \rangle)$$

Note that this rule is contravariant in the domain.

The casts in this paper (and its most closely related work) have runtime effects and are *not* erasable in general. For example, consider the term  $\langle \text{Int} \Rightarrow \{x:\text{Int} \mid \text{prime } x\} \rangle$ . The cast isn’t erasable: it isn’t in general decidable whether or not its argument will evaluate to a prime number, so a runtime check is necessary. This runtime checking isn’t limited to refinement types: the cast in  $\langle ? \Rightarrow \text{Int} \rangle$  isn’t erasable because we can’t, in general, decide whether or not its argument will evaluate to a dynamic value with an Int! tag as opposed to, say, a Bool! tag. (Some casts *are* erasable, though; e.g., from a subtype to a supertype [3, 24].)

One problem common to calculi with casts is the problem of *space efficiency*. In particular, casts can accumulate in an unbounded way: redundant casts can grow the stack arbitrarily; casting functions can introduce an arbitrary number of function proxies. This unbounded growth of casts can, in the extreme, change the asymptotic complexity of programs. To see why the naïve treatment isn’t space efficient, consider a mutually recursive definition of even and odd, adapted from Herman et al. [20]:

```

even : ?→? = ⟨Int→Bool ⇒ ?→?⟩ λx:Int.
  if x = 0 then true else odd (x - 1)
odd  : Int→Bool = λx:Int.
  if x = 0 then false else (⟨?→? ⇒ Int→Bool⟩ even) (x - 1)

```

In this example, even is written in a more dynamically typed style than odd. (While this example is contrived, it is easy to imagine mutually recursive modules with a mix of typing paradigms.) The cast  $\langle \text{Int} \rightarrow \text{Bool} \Rightarrow ? \rightarrow ? \rangle (\lambda x:\text{Int}. \dots)$  in the definition of even will (a) check that even’s dynamically typed arguments are in fact Ints and (b) tag the resulting booleans into the dynamic type. The symmetric cast on even in the definition of odd serves to make the function even behave as if it were typed. This cast will ultimately cast the integer value  $n - 1$  into the dynamic type,  $?$ , as well as projecting even’s result out of type  $?$  and into type Bool. Now consider the reduction sequence in Figure 1, observing how the number of coercions grows (redexes are highlighted).

While the operational semantics doesn’t have an explicit stack, we can still see the accumulation of pending casts. What’s more, the work is redundant: we must tag and untag true twice. In short, casts have taken an algorithm that should use a constant amount of stack space and turned it into an algorithm that uses  $O(n)$  stack space. Such a large space overhead is impractical: casts aren’t space efficient.

Two of the existing approaches to space efficiency in the world of gradual typing [20, 33, 35] factor casts into their constituent *coercions*, adapting Henglein’s system [19]. We will take a similar

<sup>1</sup>More realistic systems will make such exceptions catchable. To keep things simple, we ignore exception handling.

```

odd 3
→ ⟨⟨?→? ⇒ Int→Bool⟩ even⟩ 2
→ ⟨? ⇒ Bool⟩ (even ⟨⟨Int ⇒ ?⟩ 2⟩)
→ ⟨? ⇒ Bool⟩ (even 2Int!)
→ ⟨? ⇒ Bool⟩ (⟨⟨Int→Bool ⇒ ?→?⟩ (λx:Int. ...)⟩ 2Int!)
→ ⟨? ⇒ Bool⟩ (⟨⟨Bool ⇒ ?⟩ (⟨λx:Int. ...⟩ (⟨? ⇒ Int⟩ 2Int!)))
→ ⟨? ⇒ Bool⟩ (⟨⟨Bool ⇒ ?⟩ (⟨λx:Int. ...⟩ 2)⟩)
→ ⟨? ⇒ Bool⟩ (⟨⟨Bool ⇒ ?⟩ (odd 1)⟩)
→ ⟨? ⇒ Bool⟩ (⟨⟨Bool ⇒ ?⟩ (⟨⟨?→? ⇒ Int→Bool⟩ even) 0)⟩)
→ ⟨? ⇒ Bool⟩ (⟨⟨Bool ⇒ ?⟩ (⟨? ⇒ Bool⟩ (even ⟨⟨Int ⇒ ?⟩ 0)⟩))
→ ⟨? ⇒ Bool⟩ (⟨⟨Bool ⇒ ?⟩ (⟨? ⇒ Bool⟩ (even 0Int!)))
→ ⟨? ⇒ Bool⟩ (⟨⟨Bool ⇒ ?⟩ (⟨? ⇒ Bool⟩ (⟨⟨Int→Bool ⇒ ?→?⟩ (λx:Int. ...)⟩ 0Int!)))
→ ⟨? ⇒ Bool⟩ (⟨⟨Bool ⇒ ?⟩ (⟨? ⇒ Bool⟩ (⟨⟨Bool ⇒ ?⟩ (λx:Int. ...)⟩ (⟨? ⇒ Int⟩ 0Int!)))
→ ⟨? ⇒ Bool⟩ (⟨⟨Bool ⇒ ?⟩ (⟨? ⇒ Bool⟩ (⟨⟨Bool ⇒ ?⟩ (λx:Int. ...)⟩ 0)⟩))
→ ⟨? ⇒ Bool⟩ (⟨⟨Bool ⇒ ?⟩ (⟨? ⇒ Bool⟩ (⟨⟨Bool ⇒ ?⟩ true)⟩))
→ ⟨? ⇒ Bool⟩ (⟨⟨Bool ⇒ ?⟩ (⟨? ⇒ Bool⟩ trueBool!)⟩)
→ ⟨? ⇒ Bool⟩ (⟨⟨Bool ⇒ ?⟩ true)⟩
→ ⟨? ⇒ Bool⟩ trueBool!
→ true

```

Figure 1. Space-inefficient reduction

approach. For example, the cast  $\langle ? \Rightarrow \text{Bool} \rangle$ , which checks that a dynamic value is a boolean and then untags it, is written as the coercion  $\text{Bool}?$ ; the cast  $\langle \text{Bool} \Rightarrow ? \rangle$ , which tags a boolean into type dynamic, is written  $\text{Bool}!$ . Most importantly, coercions can be composed, so  $\langle \text{Bool}! \rangle (\langle \text{Bool}! \rangle e) \rightarrow \langle \text{Bool}!; \text{Bool}! \rangle e$ . Herman et al. normalize the coercion  $\text{Bool}!; \text{Bool}?$  into the no-op coercion  $\text{Id}$ . This normalization process is how they achieve space efficiency. For example:

$$\langle \text{Bool}! \rangle (\langle \text{Bool}! \rangle ((\lambda x:\text{Int}. \dots) (\langle \text{Int}! \rangle 2_{\text{Int}!}))) \rightarrow \langle \text{Id} \rangle ((\lambda x:\text{Int}. \dots) (\langle \text{Int}! \rangle 2_{\text{Int}!}))$$

This composition and normalization of pending coercions allows them to prove a bound on the size of any coercion that occurs during the run of a given program. This bound effectively restores the possibility of tail-call optimization.

However, it isn’t obvious how to extend Herman et al.’s [20] coercion system to refinement types. When do we test that values satisfy predicates? How do refinement type coercions normalize? We show that refinement types should have a checking coercion  $\{x:T \mid e\}?$  and an (un)tagging coercion  $\{x:T \mid e\}!$ ; the key insight for space efficiency is that the composition  $\{x:T \mid e\}?: \{x:T \mid e\}!$  should normalize to  $\text{Id}$ , i.e., checks that are immediately forgotten should be thrown away. Throwing away checks sounds dangerous, but the calculus is still sound—values typed at refinement types must satisfy their refinements. On the one hand, this is great news—space-efficiency is not only more practical, but there are fewer errors! On the other hand, the space-efficient semantics aren’t *exactly* equivalent to the naïve semantics. Whether or not this is good news, these dropped checks are part and parcel of space efficiency. We develop this idea in detail in Section 3; we show that this means that space-efficient programs fail less often than their naïve counterparts in Section 4.

This paper makes several contributions, extending the existing solutions in a number of dimensions.

- We present a simplified approach to coercions that extends the earlier work to refinement types while condensing and clarifying the formulation (Section 3). In particular, earlier systems have given either term rewriting systems modulo equational theories, which lack straightforward implementations [19–21, 33], or they have given complicated algorithms for merging casts [35]. We rework the definition of coercions to admit a straightforward term rewriting system, for which we enumerate

the exact set of canonical coercions. We then define a straightforward merge operator on canonical coercions, offering a simpler and clearer theory.

- We introduce what is, at present, the most expressive full-spectrum language, offering type dynamic as well as refinements of both base types and the dynamic type (Section 2 and Section 3). This new language narrowly edges out Wadler and Findler [39] by including refinements of type dynamic.
- We show that this language is space efficient, i.e., there are a bounded number of coercions in any program, and those coercions are bounded in size (Section 5).
- We also show that our new language is sound with respect to the naïve, inefficient semantics: if the naïve semantics reaches a value, so does the space-efficient one, but occasionally, the naïve semantics will fail when the space-efficient one succeeds (Section 4). Of the prior work, only Siek and Wadler [35] prove a similar soundness theorem, showing that their space-efficient gradual typing calculus is (exactly) equivalent to their original, naïve calculus.

## 2. A naïve coercion calculus

In this section, we define a naïve coercion calculus; to be truly complete, we ought to define a cast calculus, showing that it and the naïve calculus are observationally equivalent. To save space, we omit the definition of a cast calculus and a corresponding proof of observational equivalence with its corresponding naïve coercion calculus, opting to simply give the naïve coercion calculus directly. The relationship between a cast calculus and the naïve coercion calculus is straightforward, and establishing the relationship is not hard. (Siek and Wadler [35] establish a more interesting one, relating their space-efficient threesome casts and a coercion calculus.)

Our language adheres to a design philosophy of “simply types by default, dynamic and refinement types by coercion”. We believe this is a novel philosophy of how to build full-spectrum languages. Our design philosophy has two principles. First, base values have simple types; e.g., all integers are typed at `Int`. Second, we give operations types precise enough to guarantee totality; e.g., division has a type at least as precise as  $\text{Int} \rightarrow \{x:\text{Int} \mid x \neq 0\} \rightarrow \text{Int}$ . We understand refinement types as being designed for protecting partial operations (the original name is due to a method for protecting partial pattern matches [15]); giving operations types that make them total means that our reasoning about runtime errors can entirely revolve around cast (here, coercion) failures.

Before we begin our technical work in earnest, a word about conventions. We are defining two calculi, and they largely share syntax and typing rules (of the form `T_NAME`), relying on context to differentiate terms. The evaluation rules for the naïve calculus in Section 2 are named `F_NAME`, using a subscripted arrow  $\rightarrow_n$  for the reduction relation; the space-efficient evaluation rules in Section 3 are named `E_NAME` using a plain arrow  $\rightarrow$ .

### 2.1 Syntax and typing

Most of the terms here are standard parts of the lambda calculus. The most pertinent extension here is the coercion term,  $\langle c \rangle e$ ; we describe our language of coercions in greater detail below. Evaluation returns *results*: either a value or a failure `fail`. The term `fail` represents coercion failure. Coercion failures can occur when the predicate fails—i.e.,  $e_1[v/x] \rightarrow_n^* \text{false}_{\text{Id}}$  (see `F_CHECKFAIL`)—or when dynamically typed values don’t match their type—e.g.,  $\langle \text{Bool?} \rangle 5_{\text{Int!}}$  (see `F_TAGFAIL`). We treat `fail` as an uncatchable exception. Our values split in two parts: *pre-values*  $u$  are the typical values of other languages: constants and lambdas; *values*  $v$  are pre-values with a stack of primitive coercions. (See below for an expla-

#### Types and base types

$$\begin{aligned} T &::= B \mid T_1 \rightarrow T_2 \mid \{x:B \mid e\} \mid ? \mid \{x:? \mid e\} \\ B &::= \text{Bool} \mid \text{Int} \mid \dots \end{aligned}$$

#### Coercions, primitive coercions, and type tags

$$\begin{aligned} c &::= d_1; \dots; d_n \\ d &::= D! \mid D? \mid c_1 \mapsto c_2 \mid \text{Fail} \\ D &::= B \mid \text{Fun} \mid \{x:B \mid e\} \mid \{x:? \mid e\} \end{aligned}$$

#### Terms, results, values, and pre-values

$$\begin{aligned} e &::= x \mid r \mid \text{op}(e_1, \dots, e_n) \mid e_1 e_2 \mid \langle c \rangle e \mid \\ &\quad \langle \{x:T \mid e_1\}, e_2, v \rangle \\ r &::= v \mid \text{fail} \\ v &::= u_{\text{Id}} \mid v_d \\ u &::= k \mid \lambda x:T. e \end{aligned}$$

#### Typing contexts

$$\Gamma ::= \emptyset \mid \Gamma, x:T$$

Figure 2. Naïve syntax

nation of the different kinds of coercions.) Technically, a value is either pre-values tagged with the identity coercion,  $u_{\text{Id}}$ , or a value *tagged* with an extra coercion  $v_d$ . That is, in this language every value has a list of coercions. Values are introduced in source programs with the identity coercion,  $u_{\text{Id}}$ . Keeping a coercion on *every* value is a slight departure from prior formulations. Doing so is technically expedient—simplifying the structure of the language and clearly differentiating terms with pending coercions and values with tags.

The terms of our calculus are otherwise fairly unremarkable: we have variables, application, and a fixed set of built-in operations. We have two additional runtime terms. The *active check*  $\langle \{x:T \mid e_1\}, e_2, v \rangle$  represents an ongoing check that the value  $v$  satisfies the predicate  $e_1$ ; it is invariant that  $e_1[v/x] \rightarrow_n^* e_2$ .

Our calculus has: simple types, where  $B$  is a base type and  $T_1 \rightarrow T_2$  is the standard function type; the dynamic type  $?$ ; and refinements of both base types and type dynamic. The base refinement  $\{x:B \mid e\}$  includes all constants  $k$  of type  $B$  such that  $e[k_{\text{Id}}/x] \rightarrow_n^* \text{true}_{\text{Id}}$ . Similarly, the dynamic refinement  $\{x:? \mid e\}$  includes all *values*  $v$  such that  $v$  has type  $?$  and  $e[v/x] \rightarrow_n^* \text{true}_{\text{Id}}$ . We sometimes write  $\{x:T \mid e\}$  when it doesn’t matter whether the underlying type is  $?$  or  $B$ . Notice that refinements of dynamic indirectly include refinements of functions. At the cost of having even more canonical coercions in Section 3.1, we could add refinements of functions. We omit them because they would have brought complexity without new insights. Going beyond refinements of functions, however, is challenging future work (see Section 7).

Well formedness of types and contexts is defined straightforwardly in Figure 3. It is worth noting, however, that well formedness of refinements refers back to the term typing judgment.

Term typing (also defined in Figure 3) is mostly standard. Readers should find the rules for constants (`T_CONST`), variables (`T_VAR`), functions (`T_ABS`), failure (`T_FAIL`), application (`T_APP`), and built-in operations (`T_OP`) familiar. It is worth taking a moment to comment on the type assignment functions  $\text{ty}(k)$  and  $\text{ty}(\text{op})$  used in the `T_CONST` and `T_OP` rules. In line with our philosophy, the rule for constants gives them base types:  $\text{ty}(k) = B$ . We require that no constant have, by default, type dynamic or a refinement type. By the same philosophy, the operator type assignment function  $\text{ty}(\text{op})$  takes operations to first-order types which ensure totality. If  $\text{ty}(\text{op}) = T_1 \rightarrow \dots \rightarrow T_n \rightarrow T$ , then the operation’s denotation  $\llbracket \text{op} \rrbracket$  is a total function from  $(T_1, \dots, T_n)$  to  $T$ . If, for example, division is expressed as the operator `div`, then  $\text{ty}(\text{div}) = \text{Int} \rightarrow \{x:\text{Int} \mid x \neq 0_{\text{Id}}\} \rightarrow \text{Int}$  or some similarly exact

<b>Well formed contexts and types</b>	
$\boxed{\vdash \Gamma}$	$\boxed{\vdash T}$
$\overline{\vdash \emptyset}$ WF_EMPTY	$\frac{\vdash \Gamma \quad \vdash T}{\vdash \Gamma, x:T}$ WF_EXTEND
$\overline{\vdash B}$ WF_BASE	$\frac{\vdash T_1 \quad \vdash T_2}{\vdash T_1 \rightarrow T_2}$ WF_FUN
$\overline{\vdash ?}$ WF_DYN	$\frac{\vdash T \quad x:T \vdash e : \text{Bool}}{\vdash \{x:T \mid e\}}$ WF_REFINE
<b>Well typed terms and values</b>	
$\boxed{\Gamma \vdash u : T}$	$\boxed{\Gamma \vdash e : T}$
$\frac{\vdash \Gamma}{\Gamma \vdash k : \text{ty}(k)}$ T_CONST	$\frac{\vdash T_1 \quad \Gamma, x:T_1 \vdash e_{12} : T_2}{\Gamma \vdash \lambda x:T_1. e_{12} : T_1 \rightarrow T_2}$ T_ABS
$\frac{\Gamma \vdash u : T}{\Gamma \vdash u_{\text{Id}} : T}$ T_PREVAL	$\frac{\Gamma \vdash v : T_1 \quad \vdash d : T_1 \rightsquigarrow T_2 \quad d \neq \{x:T \mid e\} ?}{\Gamma \vdash v_d : T_2}$ T_TAGVAL
$\frac{\Gamma \vdash v : T_1 \quad \vdash d : T_1 \rightsquigarrow T_2 \quad e[v/x] \rightarrow_n^* \text{true}_{\text{Id}}}{\Gamma \vdash v_{\{x:T \mid e\} ?} : T_2}$ T_TAGVALREFINE	
$\frac{\vdash \Gamma \quad x:T \in \Gamma}{\Gamma \vdash x : T}$ T_VAR	$\frac{\vdash T \quad \vdash \Gamma}{\Gamma \vdash \text{fail} : T}$ T_FAIL
$\frac{\vdash c : T_1 \rightsquigarrow T_2 \quad \Gamma \vdash e : T_1}{\Gamma \vdash \langle c \rangle e : T_2}$ T_COERCE	
$\frac{\text{ty}(op) = T_1 \rightarrow \dots \rightarrow T_n \rightarrow T \quad \Gamma \vdash e_i : T_i}{\Gamma \vdash op(e_1, \dots, e_n) : T}$ T_OP	$\frac{\Gamma \vdash e_1 : (T_1 \rightarrow T_2) \quad \Gamma \vdash e_2 : T_1}{\Gamma \vdash e_1 e_2 : T_2}$ T_APP
$\frac{\vdash \Gamma \quad \vdash \{x:T \mid e_1\} \quad \emptyset \vdash v : T \quad \emptyset \vdash e_2 : \text{Bool} \quad e_1[v/x] \rightarrow_n^* e_2}{\Gamma \vdash \langle \{x:T \mid e_1\}, e_2, v \rangle : \{x:T \mid e_1\}}$ T_CHECKNAIVE	
<b>Well typed coercions</b>	
$\boxed{\vdash c : T_1 \rightsquigarrow T_2}$	$\boxed{\vdash d : T_1 \rightsquigarrow T_2}$
$\frac{\vdash T}{\vdash \text{Id} : T \rightsquigarrow T}$ C_ID	$\frac{\vdash d_1 : T_1 \rightsquigarrow T' \quad \vdash \dots \vdash d_n : T' \rightsquigarrow T_2}{\vdash d_1; \dots; d_n : T_1 \rightsquigarrow T_2}$ C_COMPOSE
$\frac{\vdash T_1 \quad \vdash T_2}{\vdash \text{Fail} : T_1 \rightsquigarrow T_2}$ C_FAIL	
$\overline{\vdash B? : ? \rightsquigarrow B}$ C_BUNTAG	$\overline{\vdash B! : B \rightsquigarrow ?}$ C_BTAG
$\frac{\vdash c_1 : T_{21} \rightsquigarrow T_{11} \quad \vdash c_2 : T_{12} \rightsquigarrow T_{22}}{\vdash c_1 \mapsto c_2 : (T_{11} \rightarrow T_{12}) \rightsquigarrow (T_{21} \rightarrow T_{22})}$ C_FUN	
$\overline{\vdash \text{Fun}? : ? \rightsquigarrow (? \rightarrow ?)}$ C_FUNUNTAG	$\overline{\vdash \text{Fun}! : (? \rightarrow ?) \rightsquigarrow ?}$ C_FUNTAG
$\frac{\vdash \{x:T \mid e\}}{\vdash \{x:T \mid e\} ? : T \rightsquigarrow \{x:T \mid e\}}$ C_PREDUNTAG	$\frac{\vdash \{x:T \mid e\}}{\vdash \{x:T \mid e\} ! : \{x:T \mid e\} \rightsquigarrow T}$ C_PREDTAG

Figure 3. Typing

type. This property of operator types is critical: we believe that refinement types are designed to help programmers avoid failures of (fundamentally partial) primitive operations.

Pre-values are typed by T\_CONST and T\_ABS; pre-values tagged with the Id coercion are typed as values by T\_PREVAL. T\_TAGVAL types values that are tagged with anything but a refinement type, for which we use a separate rule. We want all values at a refined type to satisfy their refinement—a key property of refinement types. The T\_TAGVALREFINE rule ensures that values typed at a refinement type actually satisfy their refinement. In our metatheory, the typing rule for active check forms, T\_CHECKNAIVE, holds onto a trace of the evaluation of the predicate. If the check succeeds, the trace can then be put directly into a T\_TAGVALREFINE derivation. Naturally, none of these rule premises about evaluation are necessary for source programs—they are technicalities for our proofs.

The coercions are the essence of this calculus: they represent the step-by-step checks that are done to move values between dynamic, simple, and refinement types. The syntax of coercions in Figure 2

splits coercions into three parts: composite coercions  $c$ , primitive coercions  $d$ , and tags  $D$ . The typing rules for coercions are written in Figure 3. When it is clear from context whether we mean a composite or a primitive coercion, we will simply call either a “coercion”. A composite coercion  $c$  is simply a list of primitive coercions. We write the empty coercion—the composite coercion comprising zero primitive coercions—as **Id**. When we write  $c$ ;  $d$  or  $d$ ;  $c$  in a rule and it matches against a coercion with a single primitive coercion—that is, when we match  $c$ ;  $B!$  against  $B!$ —we let  $c = \text{Id}$ . This is a slight departure from earlier coercion systems; this construction avoids messing around too much with re-association of coercion composition. We compare our coercions to other formulations in related work (Section 6). There are four kinds of primitive coercions: failures **Fail**, tag coercions  $D!$ , checking coercions  $D?$ , and functional coercions  $c_1 \mapsto c_2$ . (Note that  $c_1$  and  $c_2$  are composite.) Finally, the tags  $D$  are a flattening of the type space: each base type  $B$  has a corresponding tag (which we also write  $B$ ); functions have a single tag **Fun**. Intuitively, these are the type tags that are commonly used in dynamically typed languages.

We also have refinement tags for both types of refinement, which we write the same as the corresponding types.

Failure coercions are present only for showing the equivalence with the space-efficient calculus; rule F\_FAIL gives a semantics for **Fail**. (We discuss the operational semantics more fully below, in Section 2.2.) It is worthwhile to contrast our treatment of failure with that of Herman et al. [20]. Whereas Henglein treats mismatched tag/untag operations, such as  $B!$ ; **Fun?**, as stuck, we follow Herman et al. [20] in having an explicit failure coercion, **Fail**, which leads to an uncatchable program failure, fail. The F\_FAIL rule causes the program to fail when a failure coercion appears. (We carefully keep **Fail** out of the tags placed on pre-values and values.) In fact, F\_FAIL will never apply when evaluating sensible source programs—no sane program will start with **Fail** in it, and no naïve evaluation rule generates **Fail**. Instead, failures arise in the naïve calculus when the other rules with FAIL in their name fire. We include F\_FAIL as a technicality for the soundness proof (Theorem 4.6) in Section 4. In Herman et al.’s calculus,  $\langle \mathbf{Fail} \rangle v$  is a value—the program won’t actually fail until this value reaches an elimination form. While systems with lazy error detection have been proposed [23], we say that  $\langle \mathbf{Fail} \rangle e$  raises a program-terminating exception immediately. Eager error detection is more in line with standard error behavior, particularly other calculi where failed casts result in blame [2, 3, 6, 9, 11, 13, 17, 18, 22, 33–35, 38, 39].

The *tagging* coercions  $D!$  and *checking* coercions  $D?$  fall into two groups: those which move values into type dynamic and those which deal with refinements (of either base types or type dynamic).

The tags  $B$  and **Fun** are used to move values to and from the dynamic type  $?$ . The tagging coercions  $B!$  and **Fun!** mark a base value (typed  $B$ ) or functional value (typed  $? \rightarrow ?$ ) as having the dynamic type  $?$ . The checking coercions  $B?$  and **Fun?** are the corresponding *untagging* coercions, taking a dynamic value and checking its tag. If the tags match, the original typed value is returned:  $\langle \text{Bool?} \rangle \text{true}_{B!} \rightarrow_n \text{true}_{\text{Id}}$ . If the tags don’t match, the program fails:  $\langle \text{Bool?} \rangle \text{false}_{\text{Id}} \rightarrow_n^* \text{fail}$ .

The tags  $\{x:B \mid e\}$  and  $\{x:? \mid e\}$  are used for refinements. The checking coercion  $\{x:T \mid e\}?$  checks that a value  $v$  satisfies the predicate  $e$ , i.e., that  $e[v/x] \rightarrow_n \text{true}_{\text{Id}}$ ; see F\_CHECKOK and F\_CHECKFAIL below. The coercion  $\{x:T \mid e\}!$  is correspondingly used to ‘forget’ refinement checks.

Note that in both the dynamic and the refinement cases, the checking coercions are the ones that might fail. Given our philosophy of values starting out simply typed, the two types of coercions differ in that tagging coercions are applied first when moving from simple types to dynamic typing, but checking coercions are applied first when moving to refinements. Any simply typed value is just fine as an appropriately tagged dynamic value, but a simply typed value must be checked to see if it satisfies a refinement.

Finally, while the **Fun!** and **Fun?** coercions injecting and project functions on  $? \rightarrow ?$  into type  $?$ , there is a separate structural coercion that works on typed functions:  $c_1 \mapsto c_2$ , typed by the rule C\_FUN. Note that the C\_FUN rule is contravariant; see the F\_FUN rule below.

## 2.2 Operational semantics

Our rules are adapted from the evaluation contexts used in Herman et al. [20]. F\_BETA and F\_OP are totally standard CBV rules, as are most of the congruence and exception raising rules (F\_APPL, F\_APPR, F\_OPINNER, F\_APPRAISEL, F\_APPRAISER, F\_OPRAISE). Before discussing the coercion evaluation rules, it is worth taking a moment to talk about the denotation of operators. In particular,  $\llbracket op \rrbracket (v_1, \dots, v_n)$  must (a) be total when applied to correctly typed values and (b) ignore the tags on its inputs. This disallows some potentially useful operators—e.g., projecting the

tag from a dynamic value—but greatly simplifies the technicalities relating the two calculi in Section 4. We don’t believe that adding such tag-dependent operators would break anything in a deep way, but omit them for simplicity’s sake.

Most of the rules for coercions take a term of the form  $\langle d; c \rangle v$  and somehow apply the primitive coercion  $d$  to  $v$ . The rest cover more structural uses of coercions. We cover these structural rules first and then explain the ‘tagging’ rules. F\_TAGID applies when we have used up all of the primitive coercions, in which case we simply drop the coercion form.

The F\_MERGE and F\_COERCEINNER rules coordinate coercion merging and congruence. The F\_MERGE rule simply concatenates two adjacent coercions. This concatenation isn’t space efficient—in the space-efficient calculus, we normalize the concatenation to a canonical coercion of bounded size. F\_COERCEINNER steps congruently inside a coerced term—we are careful to ensure that it can only apply after F\_MERGE has fired. Carefully staging F\_COERCEINNER after F\_MERGE helps maintain determinism—if we didn’t force F\_MERGE to apply first, the number of coercions might grow out of control. Note that in F\_MERGE and F\_FAIL, the innermost term need not be a value. If we formulated our semantics as an abstract machine, we could have an explicit stack of coercions; instead, we combine them as they collide.

The remaining coercion rules have TAG in their name and work on a term  $\langle d; c \rangle v$  by combining  $d$  and  $v$ . F\_TAGB and F\_TAGFUN tag base values and functions (of type  $? \rightarrow ?$ ) into type dynamic, using the tagging coercions  $B!$  and **Fun!**, respectively. F\_TAGBB and F\_TAGFUNFUN apply  $B?$  and **Fun?** to values that have matching  $B!$  and **Fun!** tags; the effect is to simply strip the tag off the value. The F\_TAGFUNFAILB, F\_TAGBFAILFUN, and F\_TAGBFAILB rules cause the program to fail when it tries to strip a tag off with a checking coercion that doesn’t match.

F\_CHECK starts the *active check* for a refinement check. An active check  $\langle \{x:T \mid e_1\}, e_2, v \rangle$  is a special kind of condition: if  $e_2 \rightarrow_n^* \text{true}_{\text{Id}}$ , then it returns  $v_{\{x:T \mid e_1\}!}$  (rule F\_CHECKOK); if  $e_2 \rightarrow_n^* \text{false}_{\text{Id}}$ , then the active check returns fail (rule F\_CHECKFAIL). Note that the typing rules for active checks make sure that  $e_1[v/x] \rightarrow_n^* e_2$ , i.e., that the active check is actually checking whether or not the value satisfies the predicate. The F\_TAGPREDPRED rule is similar the untagging rules F\_TAGBB and F\_TAGFUNFUN, though there is no chance of failure here. Having  $\{x:T \mid e\}!$  eliminate the tag  $\{x:T \mid e\}?$  is reminiscent of the coercion normalization rule given in the introduction—which we said occasionally skips checks. But here in the naïve calculus, F\_TAGPREDPRED only applies when removing a tag from a value, i.e., when the check has already been done. In our space-efficient semantics in Section 3, our coercion normalization will actually skip checks.

The F\_TAGFUNWRAP rule wraps a value in a functional coercion. The F\_FUN rule unwinds applications of wrapped values, coercing the wrapped function’s argument and result.

In Figure 5, we translate the cast example from the introduction (Figure 1). It is easy to see that this calculus isn’t space efficient, either: coercions can consume an unbounded amount of space. As the function evaluates, a stack of coercions builds up—here, proportional to the size of the input. Again, we highlight redexes; note that the whole term is highlighted when the outermost coercions merge. The casts of the earlier example match the coercions here, e.g. the cast  $\langle ? \Rightarrow \text{Bool} \rangle e$  is just like the coercion  $\langle \text{Bool?} \rangle e$ .

Finally, the naïve calculus is type sound.

**2.1 Theorem [Type soundness]:** If  $\emptyset \vdash e : T$  then either  $e \rightarrow_n^* r$  or  $e$  diverges.

$$\begin{array}{c}
\frac{}{(\lambda x:T. e_{12})_{\mathbf{Id}} v_2 \rightarrow_n e_{12}[v_2/x]} \text{ F\_BETA} \qquad \frac{}{v_1 (c_1 \mapsto c_2) v_2 \rightarrow_n \langle c_2 \rangle (v_1 (\langle c_1 \rangle v_2))} \text{ F\_FUN} \\
\frac{}{op(v_1, \dots, v_n) \rightarrow_n [op](v_1, \dots, v_n)} \text{ F\_OP} \qquad \frac{}{\langle \{x:T \mid e\}?\rangle; c \rangle v \rightarrow_n \langle c \rangle \langle \{x:T \mid e\}, e[v/x], v \rangle} \text{ F\_CHECK} \\
\frac{}{\langle \{x:T \mid e\}, \mathbf{true}_{\mathbf{Id}}, v \rangle \rightarrow_n v_{\{x:T \mid e\} ?}} \text{ F\_CHECKOK} \qquad \frac{}{\langle \{x:T \mid e\}, \mathbf{false}_{\mathbf{Id}}, v \rangle \rightarrow_n \mathbf{fail}} \text{ F\_CHECKFAIL} \\
\frac{}{(\mathbf{Id}) v \rightarrow_n v} \text{ F\_TAGID} \qquad \frac{}{\langle \mathbf{Fun}?\rangle; c \rangle v_{B!} \rightarrow_n \mathbf{fail}} \text{ F\_TAGFUNFAILB} \qquad \frac{}{\langle \mathbf{Fun}?\rangle; c \rangle v_{\mathbf{Fun}!} \rightarrow_n \langle c \rangle v} \text{ F\_TAGFUNFUN} \\
\frac{}{\langle B!; c \rangle v \rightarrow_n \langle c \rangle v_{B!}} \text{ F\_TAGB} \qquad \frac{}{\langle \mathbf{Fun}!\rangle; c \rangle v \rightarrow_n \langle c \rangle v_{\mathbf{Fun}!}} \text{ F\_TAGFUN} \\
\frac{}{\langle B?; c \rangle v_{B!} \rightarrow_n \langle c \rangle v} \text{ F\_TAGBB} \qquad \frac{B \neq B'}{\langle B?; c \rangle v_{B'!} \rightarrow_n \mathbf{fail}} \text{ F\_TAGBFAILB} \qquad \frac{}{\langle B?; c \rangle v_{\mathbf{Fun}!} \rightarrow_n \mathbf{fail}} \text{ F\_TAGBFAILFUN} \\
\frac{}{\langle (c_1 \mapsto c_2); c \rangle v \rightarrow_n \langle c \rangle v_{c_1 \mapsto c_2}} \text{ F\_TAGFUNWRAP} \qquad \frac{}{\langle \{x:T \mid e\}!\rangle; c \rangle v_{\{x:T \mid e\} ?} \rightarrow_n \langle c \rangle v} \text{ F\_TAGPREDPRED} \\
\frac{e \neq \langle c' \rangle e'}{(\mathbf{Fail}; c) e \rightarrow_n \mathbf{fail}} \text{ F\_FAIL} \qquad \frac{e_1 \rightarrow_n e'_1}{e_1 e_2 \rightarrow_n e'_1 e_2} \text{ F\_APPL} \qquad \frac{e_2 \rightarrow_n e'_2}{v_1 e_2 \rightarrow_n v_1 e'_2} \text{ F\_APPR} \\
\frac{e_i \rightarrow_n e'_i}{op(v_1, \dots, v_{i-1}, e_i, \dots, e_n) \rightarrow_n op(v_1, \dots, v_{i-1}, e'_i, \dots, e_n)} \text{ F\_OPINNER} \qquad \frac{e \neq \langle c' \rangle e'' \quad e \rightarrow_n e'}{\langle c \rangle e \rightarrow_n \langle c \rangle e'} \text{ F\_COERCEINNER} \\
\frac{}{\langle c_1 \rangle (\langle c_2 \rangle e) \rightarrow_n \langle c_2; c_1 \rangle e} \text{ F\_MERGE} \qquad \frac{e_2 \rightarrow_n e'_2}{\langle \{x:T \mid e_1\}, e_2, v \rangle \rightarrow_n \langle \{x:T \mid e_1\}, e'_2, v \rangle} \text{ F\_CHECKINNER} \\
\frac{}{\langle c \rangle \mathbf{fail} \rightarrow_n \mathbf{fail}} \text{ F\_COERCERAISE} \qquad \frac{}{\mathbf{fail} e_2 \rightarrow_n \mathbf{fail}} \text{ F\_APPRAISEL} \qquad \frac{}{v_1 \mathbf{fail} \rightarrow_n \mathbf{fail}} \text{ F\_APPRAISER} \\
\frac{}{op(v_1, \dots, v_{i-1}, \mathbf{fail}, \dots, e_n) \rightarrow_n \mathbf{fail}} \text{ F\_OPRAISE} \qquad \frac{}{\langle \{x:T \mid e\}, \mathbf{fail}, v \rangle \rightarrow_n \mathbf{fail}} \text{ F\_CHECKRAISE}
\end{array}$$

Figure 4. Naïve operational semantics

$$\begin{array}{c}
\frac{\Gamma \vdash u : T_1 \quad \vdash c : T_1 \rightsquigarrow T_2 \quad c \neq \mathbf{Fail} \quad c \neq c'; \{x:T \mid e\} ?}{\Gamma \vdash u_c : T_2} \text{ T\_VAL} \\
\frac{\Gamma \vdash u_c : T \quad \vdash \{x:T \mid e\} ? : T \rightsquigarrow \{x:T \mid e\} \quad e[u_c/x] \rightarrow^* \mathbf{true}_{\mathbf{Id}}}{\Gamma \vdash u_c; \{x:T \mid e\} ? : \{x:T \mid e\}} \text{ T\_VALREFINE} \\
\frac{\vdash \Gamma \quad \vdash \{x:T \mid e_1\} \quad \emptyset \vdash v : T \quad \emptyset \vdash e_2 : \mathbf{Bool} \quad e_1[v/x] \rightarrow^* e_2}{\Gamma \vdash \langle \{x:T \mid e_1\}, e_2, v \rangle : \{x:T \mid e_1\}} \text{ T\_CHECK}
\end{array}$$

Figure 7. Updated typing rules for the space-efficient calculus

### 3. A space-efficient coercion calculus

Having developed the naïve semantics for our language, we now turn to space efficiency. There are two loci of inefficiency: coercion merges and function proxies (functional coercions). When F\_MERGE applies, it merely concatenates two coercions:  $\langle c_1 \rangle (\langle c_2 \rangle e) \rightarrow_n \langle c_2; c_1 \rangle e$ . Our space-efficient semantics combines  $c_2$  and  $c_1$  to eliminate redundant checks. To bound the number of function proxies, we'll make sure that coercion merging combines adjacent functional coercions and change the tagging scheme on values—while the naïve semantics allows an arbitrary stack of tags values, the space-efficient semantics will keep the size of value tags bounded. The solution to both of these problems lies in *canonical coercions* and our *merge* algorithm. Before we describe them below in Section 3.1, we discuss changes to the syntax of values and to the typing rules.

We make changes to both the syntax (Figure 6) and typing rules (Figure 7) of the naïve calculus of Section 2; we define an entirely new operational semantics (Figure 10). Throughout the new typing rules, we assume that coercions are canonical (see below).

In the naïve calculus, tagging is stacked: a value is either a pre-value tagged with **Id** or a value tagged with a single primitive coercion. In the space-efficient calculus, we collapse this stack: values are pre-values tagged with a composite coercion,  $u_c$ . Naïve, stacked values were typed using T\_PREVAL, T\_TAGVAL, and T\_TAGVALREFINE; we now use rules T\_VAL and T\_VALREFINE to type values.

The changes to the typing rules aren't major: T\_VAL and T\_VALREFINE account for flattened values: T\_VAL will apply to  $u_c$  unless the coercion  $c$  ends in  $\{x:T \mid e\} ?$ , in which case the typing derivation for the value will be T\_VALREFINE around T\_VAL. We separate the two rules to make sure we have value inversion.

```

odd 3Id
→n evenInt!→Bool? 2Id
→n ⟨Bool?⟩ (even ((Int!) 2Id))
→n ⟨Bool?⟩ (((λx:Int. ...) Id)Int?→Bool! (2Id)Int!)
→n ⟨Bool?⟩ ((Bool!) ((λx:Int. ...) Id ((Int?) (2Id)Int!)))
→n ⟨Bool!; Bool?⟩ ((λx:Int. ...) Id ((Int?) (2Id)Int!))
→n ⟨Bool!; Bool?⟩ ((λx:Int. ...) Id 2Id)
→n ⟨Bool!; Bool?⟩ (odd 1Id)
→n ⟨Bool!; Bool?⟩ (evenInt!→Bool? 0Id)
→n ⟨Bool!; Bool?⟩ ((Bool?) (even ((Int!) 0Id)))
→n ⟨Bool!; Bool!; Bool?⟩ (even ((Int!) 0Id))
→n ⟨Bool!; Bool!; Bool?⟩ (((λx:Int. ...) Id)Int?→Bool! (0Id)Int!)
→n ⟨Bool!; Bool!; Bool?⟩ ((Bool!) ((λx:Int. ...) Id ((Int?) (0Id)Int!)))
→n ⟨Bool!; Bool!; Bool!; Bool?⟩
  ((λx:Int. ...) Id ((Int?) (0Id)Int!))
→n ⟨Bool!; Bool!; Bool!; Bool?⟩ ((λx:Int. ...) Id 0Id)
→n ⟨Bool!; Bool!; Bool!; Bool?⟩ trueId
→n ⟨Bool!; Bool!; Bool!⟩ (trueId)Bool!
→n ⟨Bool!; Bool?⟩ trueId
→n ⟨Bool?⟩ (trueId)Bool!
→n ⟨Id⟩ trueId
→n trueId

```

Figure 5. Naïve reduction

```

r ::= v | fail
v ::= uc
u ::= k | λx:T. e

```

Figure 6. Updated syntax for the space-efficient calculus

That is, we want to ensure that if a value has a refinement check tag on it, it satisfies that refinement. The T\_CHECK rule changes to use the space-efficient semantics, but it remains a technical rule for supporting the evaluation of programs. Finally, we change all of the typing rules to require that coercions appearing in the program source are *canonical*. Before discussing the new evaluation rules, we discuss our space-efficient coercions and what we mean by a canonical coercion.

### 3.1 Space-efficient coercions

We define a set of *canonical coercions*, further subdivided into value coercions for constants and for functions. We list these coercions in Table 1 below; we prove that they are in fact *the* normal coercions for a standard set of rewrite rules given in Figure 8. Next, we define a set of rules for merging coercions, proving that merging two well typed canonical coercions yields a well typed canonical coercion—no bigger than the previous two combined. This is how we will show space efficiency: where a naïve implementation would accumulate and discharge all checks, our calculus will keep its coercions in canonical form for which we have a bounded size.

Henglein and Herman et al.’s systems work by taking a single rewrite rule, the so-called  $\phi$  rule:

$$B!; B? \longrightarrow \text{Id} \quad (\phi)$$

They then define a term rewriting system modulo an equational theory obtained by completing the following rules with reflexivity, symmetry, transitivity, and compatibility:

$$\begin{aligned} c_{11} \mapsto c_{12}; c_{21} \mapsto c_{22} &= (c_{21}; c_{11}) \mapsto (c_{12}; c_{22}) \\ (c_1; c_2); c_3 &= c_1; (c_2; c_3) \end{aligned}$$

We, however, directly define a rewrite system in Figure 8, where we lift  $c \longrightarrow c'$  over all possible redexes in a composite coercion  $d_1; \dots; d_n$ . Note that many of these rules correspond to reduction rules in the naïve operational semantics; we’ve written the reduc-

```

Fail; c → Fail (F_FAIL)
c; Fail → Fail
(c11 ↦ c12); (c21 ↦ c22) → (c21; c11) ↦ (c12; c22)
B!; B? → Id (F_TAGBB)
B!; B'? → Fail when B ≠ B' (F_TAGBFAILB)
Fun!; Fun? → Id (F_TAGFUNFUN)
B!; Fun? → Id (F_TAGBFAILFUN)
Fun!; B? → Id (F_TAGFUNFAILB)
{x:T | e}?; {x:T | e'}! → Id (F_TAGPREDPRED)
{x:T | e}?; {x:T | e}? → Id

```

Figure 8. Rewriting rules

tion rule names next to such rewrite rules. The rules for predicates over base types and type dynamic are new. Just as  $B?$  takes a less specific type,  $?$ , to a more specific type  $B$  while performing a check (C\_BUNTAG), we have  $\{x:B \mid e\}?$  that take a less specific type,  $B$  to a more specific type  $\{x:B \mid e\}$  while performing a check (C\_PREDUNTAG). By the same analogy,  $\{x:B \mid e\}!$  takes a more specific type to a less specific one. The rules for refinements don’t have just a  $\phi$  rule—they must have what Henglein calls a  $\psi$  rule:

$$\{x:T \mid e\}?; \{x:T \mid e\}! \longrightarrow \text{Id}$$

Consider the rule F\_TAGPREDPRED from the naïve operational semantics (Figure 4): we must have a  $\psi$  rule if  $\{x:T \mid e\}!$  is going to untag  $v_{\{x:T \mid e\}?$ . But if  $\{x:T \mid e\}?$ ;  $\{x:T \mid e\}!$  on tags, it must also hold for coercions on the stack if we want space efficiency. That is, space-efficient refinement checking *must* drop some checks on the floor. The  $\phi$  rule for refinements is an optimization, unnecessary for soundness and space efficiency.

**3.1 Lemma:** The rewrite system is strongly normalizing on well typed terms.

**3.2 Lemma:** The canonical coercions are normal.

**3.3 Lemma:** If  $\vdash c : T_1 \rightsquigarrow T_2$  and  $c$  is normal, then  $c$  is canonical.

**3.4 Corollary:** All well typed coercions rewrite to a canonical coercion.

Having developed the rewrite rules for our (somewhat relaxed) coercions, we define a merge algorithm in Figure 9 that takes two coercions and merges them from the edges. We will only ever merge *canonical* coercions, greatly simplifying the algorithm. Our merging relation implements the  $\phi$  rule in N\_BB and N\_FUNFUN. We implement failure rules in N\_BFAILB, N\_BFAILFUN, and N\_FUNFAILB. We relate our coercion systems to others in Section 6.

Looking at Table 1, we can see why T\_VALREFINE only needs to check the last coercion on a tagged pre-value: if  $\{x:T \mid e\}?$  appears in a canonical coercion, it appears at the end. Similarly, the observation that certain coercions are *value coercions*, i.e., are the only coercions that can be applied to values, can be made based on typing: if constants are typed at simple types and lambdas are assigned functional types, then all value coercions must come from  $B$  or  $T_1 \rightarrow T_2$ .

We write  $c_1 \Downarrow c_2$  (read “merge  $c_1$  and  $c_2$ ”) for canonical coercions  $c_1$  and  $c_2$  to mean the coercion  $c$  such that  $c_1 * c_2 \Rightarrow c$ . This notation is justified by Lemma 3.6, which shows that merging is an operator on canonical coercions.

**3.5 Lemma [Preservation for merge]:** If  $\vdash c_1 : T_1 \rightsquigarrow T_2$  and  $\vdash c_2 : T_2 \rightsquigarrow T_3$  and  $c_1 * c_2 \Rightarrow c_3$  then  $\vdash c_3 : T_1 \rightsquigarrow T_3$ .

$$\begin{array}{c}
\frac{c_1; c_2 \text{ is canonical}}{c_1 * c_2 \Rightarrow (c_1; c_2)} \quad \text{N\_CANONICAL} \qquad \frac{c_{21} * c_{11} \Rightarrow c_{31} \quad c_{12} * c_{22} \Rightarrow c_{32} \quad c_1 * (c_{31} \mapsto c_{32}); c_2 \Rightarrow c}{c_1; (c_{11} \mapsto c_{12}) * (c_{21} \mapsto c_{22}); c_2 \Rightarrow c} \quad \text{N\_FUN} \\
\\
\frac{}{\text{Id} * c \Rightarrow c} \quad \text{N\_IDL} \qquad \frac{c \neq \text{Id}}{c * \text{Id} \Rightarrow c} \quad \text{N\_IDR} \qquad \frac{}{\text{Fail} * c \Rightarrow \text{Fail}} \quad \text{N\_FAILL} \qquad \frac{c \neq \text{Fail}}{c * \text{Fail} \Rightarrow \text{Fail}} \quad \text{N\_FAILR} \\
\\
\frac{c_1 * c_2 \Rightarrow c}{c_1; B! * B?; c_2 \Rightarrow c} \quad \text{N\_BB} \qquad \frac{c_1 * c_2 \Rightarrow c}{c_1; \text{Fun!} * \text{Fun?}; c_2 \Rightarrow c} \quad \text{N\_FUNFUN} \\
\\
\frac{B \neq B'}{c_1; B! * B?; c_2 \Rightarrow \text{Fail}} \quad \text{N\_BFAILB} \qquad \frac{}{c_1; B! * \text{Fun?}; c_2 \Rightarrow c} \quad \text{N\_BFAILFUN} \qquad \frac{}{c_1; \text{Fun!} * B?; c_2 \Rightarrow \text{Fail}} \quad \text{N\_FUNFAILB} \\
\\
\frac{c_1 * c_2 \Rightarrow c}{c_1; \{x: T \mid e\} * \{x: T \mid e\}!; c_2 \Rightarrow c} \quad \text{N\_PREDPRED} \qquad \frac{c_1 * c_2 \Rightarrow c}{c_1; \{x: T \mid e\}! * \{x: T \mid e\}?; c_2 \Rightarrow c} \quad \text{N\_PREDSAME}
\end{array}$$

Figure 9. Merging coercions

Coercion	Type
<b>Id</b> :	$T \rightsquigarrow T$
<b>Fail</b> :	$T \rightsquigarrow T'$
$\{x:? \mid e\}?$ :	$? \rightsquigarrow \{x:? \mid e\}$
$B?; c$ :	$? \rightsquigarrow T$
<b>Fun?</b> ; $c$ :	$? \rightsquigarrow ? \rightarrow ?$
$\vdash c : B \rightsquigarrow T$ is a value coercion	
$B!$ :	$B \rightsquigarrow ?$
$B!; \{x:? \mid e\}?$ :	$B \rightsquigarrow \{x:? \mid e\}$
$\{x:B \mid e\}?$ :	$B \rightsquigarrow \{x:B \mid e\}$
<b>Fun!</b> :	$(? \rightarrow ?) \rightsquigarrow ?$
<b>Fun!</b> ; $\{x:? \mid e\}?$ :	$(? \rightarrow ?) \rightsquigarrow \{x:? \mid e\}$
$c_1 \mapsto c_2$ :	$(T_{11} \rightarrow T_{12}) \rightsquigarrow (T_{21} \rightarrow T_{22})$
$c_1 \mapsto c_2; \text{Fun!}$ :	$(T_{11} \rightarrow T_{12}) \rightsquigarrow ?$
$c_1 \mapsto c_2; \text{Fun!}; \{x:? \mid e\}?$ :	$(T_{11} \rightarrow T_{12}) \rightsquigarrow \{x:? \mid e\}$
$\{x:? \mid e\}!$ :	$\{x:? \mid e\} \rightsquigarrow ?$
$\{x:? \mid e\}!; \{x:? \mid e'\}?$ :	$\{x:? \mid e\} \rightsquigarrow ?$ where $e \neq e'$
$\{x:? \mid e\}!; B?; c$ :	$\{x:? \mid e\} \rightsquigarrow T$
$\vdash c : B \rightsquigarrow T$ is a value coercion	
$\{x:? \mid e\}!; \text{Fun?}; c$ :	$\{x:? \mid e\} \rightsquigarrow ? \rightarrow ?$
$\vdash c : (? \rightarrow ?) \rightsquigarrow T$ is a value coercion	
$\{x:B \mid e\}!; c$ :	$\{x:B \mid e\} \rightsquigarrow T$
$\vdash c : B \rightsquigarrow T$ is a value coercion and $c \neq \{x:B \mid e\}?$	

Rows with a blue background are *value coercions*, and are the only coercions that can appear as tags on pre-values. Horizontal rules mark a change of source type.

Table 1. Canonical coercions

**3.6 Lemma [Merge is an operator]:** Given canonical coercions  $\vdash c_1 : T_1 \rightsquigarrow T_2$  and  $\vdash c_2 : T_2 \rightsquigarrow T_3$ , then there exists a unique canonical coercion  $c$  such that  $c_1 * c_2 \Rightarrow c$ .

### 3.2 Operational semantics

We give the relevant rules of the changed operational semantics in Figure 10. The biggest change to our operational semantics is that E\_MERGE explicitly merges the two coercions. Herman et al. simply say that they keep their coercions in normal form—that is, we should interpret normalization happening automatically when E\_MERGE applies, even though they write E\_MERGE as simply concatenating the two coercions into  $c_2; c_1$ . Our semantics explicitly normalizes the coercions (rule E\_MERGE), possibly stopping the program on the next step (the no-longer-useless rule E\_FAIL).

Otherwise, the rules are largely the same as the naïve semantics, though we’re now able to use merges to distill the tag rules into a few possibilities: E\_TAG replaces all of the success-

```

odd 3Id
→ evenInt!→Bool? 2Id
→ (Bool?) (even ((Int!) 2Id))
→ (Bool?) (((λx:Int. ...) Int?→Bool! 2Int!))
→ (Bool?) ((Bool!) ((λx:Int. ...) Id ((Int?) 2Int!)))
→ (Id) ((λx:Int. ...) Id ((Int?) 2Int!))
→ (Id) (λx:Int. ...) Id 2Id
→ (Id) (odd 1Id)
→ (Id) (evenInt!→Bool? 0Id)
→ (Id) ((Bool?) (even ((Int!) 0Id)))
→ (Bool?) (even ((Int!) 0Id))
→ (Bool?) ((λx:Int. ...) Int?→Bool! 0Int!)
→ (Bool?) ((Bool!) (λx:Int. ...) Id ((Int?) 0Int!))
→ (Id) ((λx:Int. ...) Id ((Int?) 0Int!))
→ (Id) (λx:Int. ...) Id 0Id
→ (Id) trueId
→ trueId

```

Figure 11. Space-efficient reduction

ful F\_TAG\* rules; E\_TAGFAIL replaces all of the failing F\_TAG\* rules. E\_CHECKOK is essentially F\_CHECKOK, though it uses a merge instead of concatenation (though the typing rules mean that the merge will apply N\_CANONICAL every time).

We can finally observe that our reduction is space efficient: the coercions in Figure 11 don’t grow with the size of the input like the coercions in Figure 5 or the casts in Figure 1. We discuss this claim in more detail in Section 5.

### 3.3 Proofs

Our space-efficient calculus enjoys type soundness; we show as much using standard syntactic methods.

**3.7 Lemma [Progress]:** If  $\emptyset \vdash e : T$  then either  $e$  is a result, or there exists an  $e'$  such that  $e \rightarrow e'$ .

**3.8 Lemma [Preservation]:** If  $\emptyset \vdash e : T$  and  $e \rightarrow e'$ , then  $\emptyset \vdash e' : T$ .

## 4. Soundness of the space-efficient calculus

We will never get an exact semantic match between the naïve and space-efficient semantics: the  $\psi$  rule for refinements in the space-efficient semantics mean that some checks will happen in the naïve semantics that won’t happen in the space-efficient semantics. All we can hope for is that *if* the naïve semantics produces a value, the space-efficient calculus will produce a similar one.



$$\begin{array}{c}
\frac{}{u_1(c_1 \mapsto c_2) v_2 \longrightarrow \langle c_2 \rangle (u_1 \mathbf{Id} (\langle c_1 \rangle v_2))} \text{E.FUN} \\
\frac{d_1 \neq \{x:T \mid e\}?\quad c * d_1 \Rightarrow c' \quad c' \neq \mathbf{Fail}}{\langle d_1; c_2 \rangle u_c \longrightarrow \langle c_2 \rangle u_{c'}} \text{E.TAG} \\
\frac{}{\langle \mathbf{Id} \rangle v \longrightarrow v} \text{E.TAGID} \\
\frac{}{\langle c_1 \rangle (\langle c_2 \rangle e) \longrightarrow \langle c_2 \downarrow c_1 \rangle e} \text{E.MERGE} \\
\frac{d_1 \neq \{x:T \mid e\}?\quad c * d_1 \Rightarrow \mathbf{Fail}}{\langle d_1; c_2 \rangle u_c \longrightarrow \mathbf{fail}} \text{E.TAGFAIL} \\
\frac{}{\langle \{x:T \mid e\}, \mathbf{true}_{\mathbf{Id}}, u_c \rangle \longrightarrow u_{c \downarrow \{x:T \mid e\}}?} \text{E.CHECKOK}
\end{array}$$

Figure 10. Operational semantics

### Pre-value and value rules

$$\begin{array}{c}
k_{\mathbf{Id}} \sim k_{\mathbf{Id}} : B \iff \text{ty}(k) = B \\
v_{11} \sim v_{21} : T_1 \rightarrow T_2 \iff \\
\quad \forall v_{12} \sim v_{22} : T_1. v_{11} v_{12} \simeq v_{21} v_{22} : T_2 \\
v_{1B!} \sim u_{2c \downarrow B!} : ? \iff v_1 \sim u_{2c} : B \\
v_{1\mathbf{Fun}!} \sim u_{2c \downarrow \mathbf{Fun}!} : ? \iff v_1 \sim u_{2c} : ? \rightarrow ? \\
v_{\{x:T \mid e_1\}?\} \sim u_{c \downarrow \{x:T \mid e_2\}?\} : \{x:T \mid e_1\} \\
\iff \\
v_1 \sim u_{2c} : T \wedge \{x:T \mid e_1\} \sim \{x:T \mid e_2\}
\end{array}$$

### Term rules

$$\begin{array}{c}
e_1 \simeq e_2 : T \\
\iff \\
e_1 \text{ diverges} \vee e_1 \xrightarrow{*}_n \mathbf{fail} \vee \\
(e_1 \xrightarrow{*}_n v_1 \wedge e_2 \xrightarrow{*}_n v_2 \wedge v_1 \sim v_2 : T)
\end{array}$$

### Type rules

$$\begin{array}{c}
B \sim B \quad ? \sim ? \\
T_{11} \rightarrow T_{12} \sim T_{21} \rightarrow T_{22} \iff T_{11} \sim T_{21} \wedge T_{12} \sim T_{22} \\
\{x:T \mid e_1\} \sim \{x:T \mid e_2\} \iff \\
\quad \forall v_1 \sim v_2 : T. e_1[v_1/x] \simeq e_2[v_2/x] : \mathbf{Bool}
\end{array}$$

### Closing substitutions

$$\begin{array}{c}
\Gamma \models \delta \iff \forall x \in \text{dom}(\Gamma). \delta_1(x) \sim \delta_2(x) : T \\
\Gamma \vdash e_1 \simeq e_2 : T \iff \forall \Gamma \models \delta. \delta_1(e_1) \simeq \delta_2(e_2) : T
\end{array}$$

Figure 12. Relating the naïve and space-efficient semantics

We adapt the asymmetric logical relations from Greenberg et al. [17] to show that the two calculi behave mostly the same, with the naïve calculus diverging and failing more often. We define the logical relation in Figure 12. The definitions begin by defining a relation  $v_1 \sim v_2 : T$  for closed values and a relation  $e_1 \simeq e_2 : T$  for closed terms as a fixpoint on types (it is a known bug that this isn't quite right—rather, we ought to use step-indexed logical relations). We lift the definitions to open terms by defining dual closing value substitutions  $\delta$ ; if  $\Gamma \models \delta$  and  $x:T \in \Gamma$ , then  $\delta_1(x) \sim \delta_2(x) : T$ .

Naïve terms are on the left of the relation, while space-efficient terms are on the right. We require that both sides be well typed. We obtain the asymmetry we seek by saying that when the naïve semantics yields a value, then the space-efficient yields a similar one—but otherwise, the naïve semantics will fail or diverge. This definition still allows the space-efficient calculus to diverge or to fail, but then the naïve semantics must also diverge or fail—but note that it's possible for the left-hand side of the relation to diverge and the right-hand side to fail, and vice versa. This is possible because the naïve semantics could run a check diverges, while the space-efficient semantics skips that check and instead runs a failing one.

### Calculating canonical coercions

$$\begin{array}{c}
\text{canonical}(\mathbf{Id}) = \mathbf{Id} \\
\text{canonical}(\mathbf{Fail}) = \mathbf{Fail} \\
\text{canonical}(d_1; \dots; d_n) = \\
\quad \text{canonical}(d_1) \downarrow \text{canonical}(d_2; \dots; d_n) \\
\text{canonical}(D!) = \text{canonical}(D)! \\
\text{canonical}(D?) = \text{canonical}(D)? \\
\text{canonical}(c_1 \mapsto c_2) = \text{canonical}(c_1) \mapsto \text{canonical}(c_2) \\
\text{canonical}(B) = B \\
\text{canonical}(\mathbf{Fun}) = \mathbf{Fun} \\
\text{canonical}(\{x:T \mid e\}) = \{x:T \mid \text{canonical}(e)\} \\
\text{canonical}(u_{\mathbf{Id}}) = \text{canonical}(u)_{\mathbf{Id}} \\
\text{canonical}(v_d) = u_{c \downarrow d} \\
\text{where } \text{canonical}(v) = u_c
\end{array}$$

Figure 13. Canonicalizing naïve terms

The value relation  $v_1 \sim v_2 : T$  is subtler than usual for logical relation: the definitions at  $?$  and  $\{x:T \mid e\}$  must shuffle some tags around. In particular, the rule for type  $?$  is split into cases by the underlying tag of values. The case for refinements  $\{x:T \mid e\}$  requires that the values be related at the underlying type  $T$  (recalling that  $T = B$  or  $T = ?$ ) and also that the values be tagged as satisfying the predicate (or a related predicate, in the case of the space-efficient calculus).

Our proof works by showing that a well-typed naïve term  $e$  is related to its translation into the space-efficient term  $\text{canonical}(e)$ . We define  $\text{canonical}$  in Figure 13, omitting most of the cases since they are homomorphic. We give the cases for coercions and for values because they are the most interesting. In particular,  $\text{canonical}(v)$  must unfold the stacked tags on a naïve-calculus value and merge them into a single coercion.

Our ultimate goal is soundness: if  $\Gamma \vdash e : T$  then  $\Gamma \vdash e \simeq \text{canonical}(e) : T$ . Our proof works in a few stages: first we define relations  $\vdash c$  **ignorable** (coercions which are equivalent to  $\mathbf{Id}$  or  $\mathbf{Fail}$ ) and  $\vdash c$  **failable** (coercions which are equivalent to  $\mathbf{Fail}$ ). We omit the details of these relations for space, though they are available in the supplementary materials. We then prove lemmas that allow us to easily work with ignorable and failable coercions (Lemma 4.1 and Lemma 4.2, respectively). Then we relate non-canonical coercions to canonical ones (using a separate inductive relation  $\vdash c_1 \sim c_2$ , defined in Figure 14). We show that such related coercions are logically related on logically related values. We then prove a separate lemma showing that related coercions are logically related on related terms—this not a trivial extension of the similar lemma for values, due to coercion merges. With those lemmas to hand, we show soundness. Don't worry—we explain the proof less tersely as we go.

## Relating coercions

$$\begin{array}{c}
\frac{\vdash c'_i \text{ ignorable} \quad \vdash c_i \sim d_i}{\vdash c'_0; c_1; c'_1; c_2; \dots; c'_{n-1}; c_n; c'_n \sim d_1; \dots; d_n} \text{ R\_COMPOSITE} \qquad \frac{\vdash c \text{ ignorable}}{\vdash c \sim \text{Id}} \text{ R\_ID} \\
\\
\frac{\vdash c'_i \text{ ignorable}}{\vdash (c_{1n1}; \dots; c_{111}) \mapsto (c_{112}; \dots; c_{1n2}) \sim c_{21} \mapsto c_{22}} \text{ R\_FUN} \qquad \frac{\vdash c_1 \text{ failable}}{\vdash c'_1; c_1; c'_2 \sim c_2} \text{ R\_FAIL} \\
\\
\frac{\vdash D \sim D}{\vdash D! \sim D!} \text{ R\_TAG} \qquad \frac{\vdash D \sim D}{\vdash D? \sim D?} \text{ R\_CHECK} \qquad \frac{}{\vdash B \sim B} \text{ R\_DB} \qquad \frac{}{\vdash \text{Fun} \sim \text{Fun}} \text{ R\_DFUN} \\
\\
\frac{}{\vdash \{x:T \mid e_1\} \sim \{x:T \mid \text{canonical}(e_1)\}} \text{ R\_DPREDCANONICAL} \qquad \frac{\forall v_1 \sim v_2 : T. e_1[v_1/x] \simeq e_2[v_2/x] : \text{Bool}}{\vdash \{x:T \mid e_1\} \sim \{x:T \mid e_2\}} \text{ R\_DPREDLR}
\end{array}$$

Figure 14. Relating coercions

Ignorable coercions can be freely added or removed to naïve terms while preserving logical relation to space-efficient terms.

**4.1 Lemma:** If  $\langle c_1; c_2 \rangle v_1 \simeq e_2 : T$  and  $\vdash c_1$  **ignorable** then  $\langle c_2 \rangle v_1 \simeq e_2 : T$ .

We prove a similar lemma that *failable* coercions always fail.

**4.2 Lemma:** If  $\vdash c_1$  **failable**, then  $\langle c'_1; c_1; c'_2 \rangle v_1 \simeq e_2 : T$ .

With ignorable and failable coercions, we can characterize *all* non-canonical coercions, relating them to canonical coercions. The relation  $\vdash c_1 \sim c_2$  relates a non-canonical coercion  $c_1$  to a canonical coercion  $c_2$ . Note that this inductively defined relation isn't the same thing as the logical relation. First we show that well typed coercions  $c$  in the naïve calculus are related to  $\text{canonical}(c)$ . Then we'll use this general relation to relate *in the logical relation* how coercion forms work on logically related values. We define two rules for relating the tag  $D = \{x:T \mid e\}$ , one which uses canonical (rule R\_DPREDCANONICAL) and one which uses the logical relation (rule R\_DPREDLR). We'll use the former rule in the first lemma, and the latter rule in the second lemma; the IH on term sizes in the soundness theorem will let us convert the derivations using the first rule to use the second.

**4.3 Lemma [Characterizing non-canonical coercions]:** If  $\vdash c : T_1 \rightsquigarrow T_2$  then  $\vdash c \sim \text{canonical}(c)$  (using R\_DPREDCANONICAL).

We now show that *any* coercions  $\vdash c_1 \sim c_2$  yield related results when applied to related values. We defined the relation  $\vdash c_1 \sim c_2$  because this lemma is easier to prove on the relation than on the canonical function itself.

**4.4 Lemma [Relating canonical coercions]:**

If  $v_1 \sim v_2 : T_1$  and  $\vdash c_1 : T_1 \rightsquigarrow T_2$  and  $\vdash c_1 \sim c_2$  (using R\_DPREDLR), then  $\langle c_1 \rangle v_1 \simeq \langle c_2 \rangle v_2 : T_2$ .

In Greenberg et al. [17], a similar characterization of casts is sufficient. In a standard lambda calculus, we would be able to use the logical relation to reduce  $e_1$  and  $e_2$  to values and then directly apply Lemma 4.4. But that strategy won't work here: reducing those terms may put new coercions on the outside; these extra coercions will merge into  $c_1$  and  $c_2$  (by F\_MERGE or E\_MERGE), possibly disrupting  $\vdash c_1 \sim c_2$ .

**4.5 Lemma [Relating coercions with merges]:** If  $e_1 \simeq e_2 : T_1$  and  $\vdash c_1 : T_1 \rightsquigarrow T_2$  and  $\vdash c_1 \sim c_2$  (using R\_DPREDLR), then  $\langle c_1 \rangle e_1 \simeq \langle c_2 \rangle e_2 : T_2$ .

**Proof:** First, we can ignore the cases where  $e_1 \rightarrow_n^* \text{fail}$  or  $e_1$  diverges—those are immediately related. So  $e_1 \rightarrow_n^* v_1$  and

$e_2 \rightarrow_n^* v_2$ . By induction on the length of the evaluation derivations, with a long case analysis on the space-efficient side.  $\square$

**4.6 Theorem [Soundness]:** If  $\Gamma \vdash e : T$  then  $\Gamma \vdash e \simeq \text{canonical}(e) : T$ .

**Proof:** We generalize the proof to work also on values and terms. By lexicographic induction on the typing derivation and the size of the term ( $v$  or  $e$ , respectively), using Lemma 4.5 in the T\_COERCE case. We use Lemma 4.3 and Lemma 4.4 in the T\_TAGVAL and T\_TAGVALREFINE cases.  $\square$

The definition of  $\emptyset \vdash e_1 \simeq e_2 : T$  gives us our approximate observational equivalence: either  $e \rightarrow_n^* \text{fail}$ ,  $e$  diverges, or  $e \rightarrow_n^* v_1$  and  $\text{canonical}(e) \rightarrow_n^* v_2$  such that  $v_1 \sim v_2 : T$ . Note that for base values, we have exactly the same result on both sides.

## 5. Space efficiency

The structure of our space-efficiency proof is largely the same as in prior work. Coercion size is broken down by the order of the types involved; the maximum size of any coercion is  $|\text{largest coercion}| \cdot 2^{\text{tallest type}}$ . Inspecting the canonical coercions, the largest is  $\{x:? \mid e\}!; \text{Fun?}; c_1 \mapsto c_2; \text{Fun!}; \{x:? \mid e'\}?$ , with a size of 5. The largest possible canonical coercion therefore has size  $M = 5 \cdot 2^h$ .

Formally, observe that merging canonical coercions  $c_1$  and  $c_2$  either produces a smaller coercion or  $c_1; c_2$  is canonical.

**5.1 Lemma [Merge reduces size]:** If  $c_1 * c_2 \Rightarrow c_3$ , then either  $\text{size}(c_1) + \text{size}(c_2) > \text{size}(c_3)$ , or  $c_3 = c_1; c_2$  is canonical.

**Proof:** By induction on the derivation of  $c_1 * c_2 \Rightarrow c_3$ .  $\square$

Rules with merges (and E\_MERGE in particular) don't increase the size of the largest coercion in the program. Applying this lemma across an evaluation  $e \rightarrow_n^* e'$ , we can see that no coercion ever exceeds the size of the largest coercion in  $e$ . If  $M$  is the size of the largest coercion, then there is at most an  $M$ -fold space overhead of coercions. But this size bound is galactic; we find it hard to believe that this overhead is observable in practice. A much more interesting notion of space efficiency—not studied here—is to determine implementation schemes for space-efficient layout of coercions in memory and time-efficient merges of coercions. We believe that explicitly enumerating the canonical coercions is a step towards this goal: the canonical coercions in Table 1 are exactly those which must be represented.

## 6. Related work

There are two related threads of work: a more recent line of work on gradual types, refinement types, and full-spectrum programming

languages; and an older, more general line of work on coercions, which may or may not have runtime semantics. Space efficiency and representation have been studied in both settings.

### Space efficiency, gradual typing, and refinement types

In Siek and Taha’s seminal work on gradual typing [34], space efficiency is already a concern—they point out that the canonical forms lemma has implications for which values can be unboxed (the typed ones). Herman, Tomb, and Flanagan [20] compiled a language like Siek and Taha’s into a calculus with Henglein’s coercions [19], proving a space-efficiency result with a galactic bound similar to ours. Herman et al. stop at proving that their compilation is type preserving without proving soundness of their compilation. (We compare our system to Herman et al.’s in greater detail below.) Siek, Garcia, and Taha [33] explore the design space around Herman et al.’s result, this time with an observational equivalence theorem exactly relating two coercion semantics.

Siek and Wadler [35] study an alternative, cast-based formulation of space efficiency, proving tighter bounds than Herman et al. [20] and an exact observational equivalence. Their insight is that casts can be factored not merely as a “twosome”  $\langle S \Rightarrow T \rangle$ , but rather as a threesome:  $\langle S \xrightarrow{R} T \rangle$ . They maintain the invariant that  $S$  downcasts to  $R$ , and  $R$  upcasts to  $T$ ; merging casts amounts to calculating a greatest lower bound. They come up with an elegant theory of merging casts, with a detailed accounting for blame. While the mathematics is beautiful, we believe that their algorithm is overkill: Herman et al.’s journal article [21] cleanly enumerates the recursive structure of the canonical coercions for dynamic and simple types, with only 17 possible structures at the top level. Siek and Wadler’s theory is the theory of these 17 structures. Many of the solutions can be simply pre-computed and looked up in a table at runtime. We have 37 canonical coercions. We don’t study the question here, but we believe that a careful analysis would allow for very compact representations with very fast merges—by pointer comparison and table lookup when functional coercions aren’t involved. We discuss this issue further in future work (Section 7).

Before considering other full-spectrum languages, we compare our work to the most closely related work: Herman, Tomb, Flanagan [20, 21] and Henglein [19]. Henglein is trying to reason carefully about programs written in a dynamic style, rather than thinking about multi-paradigm programming (though it is clear that he knows that his work applies to “dynamic typing in a static language”). His theory of coercions has no `Fail` coercion and treats `Id` slightly differently at function types. Herman et al. adapt his calculus to match the setting of gradual types, though they never rebuild his theory. Henglein develops a general theory *characterizing* canonical coercions, but we *enumerate* them.

Perhaps the biggest difference is that Henglein and Herman et al. formulate coercions as having arbitrary composition:  $c_1; c_2$  is a coercion that can be used freely. As a consequence, it is somewhat difficult to reason directly about coercions in the calculus: what should  $\langle (c_1 \mapsto c_2; \text{Fun!}); \{x: ? \mid e\} \rangle v$  do? Their solution is to work with coercions up to an equivalence relation that includes associativity of coercion composition; coercions normalize in a term rewriting system modulo this equivalence relation. Henglein studies some algorithmic rewriting systems. But Herman et al. don’t develop the rewriting system at all, never showing that their rewriting system is strongly normalizing, and even when they enumerate canonical coercions in their journal version [21], they do so without proof. We feel that term rewriting modulo equational theories is somewhat insufficient for guiding an implementation of a coercion calculus: the compiler needs a concrete representation for coercions and a concrete algorithm for merging them. We accordingly adopt a constrained form of coercion composition out of a desire to aid implementation, but also out of expedience: we don’t need to

worry about associativity at all. We don’t believe that free composition buys anything, anyway: we don’t expect programmers to be writing coercions by hand, so ease of expression in the coercion language isn’t particularly important.

The work discussed so far consisted of calculi devised expressly for space-efficient gradual typing. Findler et al. [12] discuss space efficiency from the perspective of an implementation in PLT Racket (then PLT Scheme). Their setting—latent contracts, no type system—is rather different from the foregoing systems; they address datatypes, while the foundational calculi omit datatypes.

Considering the world of full-spectrum programming languages more broadly, we summarize existing solutions. None of the following are space efficient; we are the first to combine space efficiency, gradual types, and refinement types. Ou et al. [30] cover the spectrum and include dependent types, but allow only a constrained set of refinement predicates; Sage [25] covers the entire spectrum and also includes dependent types, but lacks a soundness proof; Wadler and Findler’s [39] development covers dynamic types through refinements of base types; Bierman et al. [4] cover the whole spectrum but (also with dependency) only for first-order types.

### Coercions

There are many other systems that use coercions to other ends. Henglein gives an excellent summary of work up to 1994 in the related work section of his article [19]. One of the classic uses of coercions is subtyping [5, 28]; more recent work relates subtyping and polymorphism [8]. Work on unboxing [29, 32] confronts similar issues of space efficiency. Many of these examples carefully ensure that coercions are erasable, while our coercions are definitely not.

Swamy, Hicks, and Bierman [36] study coercion insertion in general, showing that their framework can encode gradual types. We haven’t studied coercion insertion at all, though Swamy et al.’s framework would be a natural one to use. We are not aware of work on how coercion insertion algorithms affect space consumption.

## 7. Future work

The obvious next step is adding blame [11]. Siek and Wadler [35] are the only space-efficient calculus to have blame, which they obtain with some effort—their threesome merging takes place outside in, making it hard to compute which label to blame. We believe that our inside-out coercion merge algorithm offers a straightforward way to compute blame: we conjecture that blame comes from left to right.

Extending the calculus to general refinements, where any type  $T$  can be refined to  $\{x:T \mid e\}$ , would be a challenging but important step towards adding polymorphism. (You can’t allow refinement of type variables unless *any* type can be refined, since there’s no way to know what type will be substituted in for the variable.) It wouldn’t be too difficult to add function refinements  $\{x:(T_1 \rightarrow T_2) \mid e\}$  to this calculus, but refinements of refinements seem to break space efficiency: if  $\{x:B \mid e\}$  is canonical, so is  $\{x:B \mid e\}; \{x:\{x:B \mid e\} \mid e'\}$ —there are an infinite number of canonical coercions. In a monomorphic calculus, the number of canonical coercions can be bounded by the types in the original program, but that doesn’t hold for a polymorphic calculus. Prior work relating dynamic types and polymorphism will apply here, as well [2, 26].

Adding dependent functions to the coercion calculus above would complicate matters significantly, but would also add a great deal of expressiveness. Programs with dependent types have a potentially infinite set of types (and so coercions) as the program evaluates, but the *shape* of canonical coercions remain unchanged.

Set semantics for refinement types extend refinement types to have a set of predicates, rather than a single one. The `N.PREDSAME`

merge rule skips a check when we would have projected out of and then back into a refinement type. If refinements were sets, we could broaden this optimization to allow the space-efficient calculus to avoid even more redundant checks.

## Acknowledgments

We thank Stephanie Weirich and Benjamin Pierce for looking over drafts, and also for their excellent advice and support. Brent Yorgey contributed to some discussions. Our work has been supported by the National Science Foundation under grant 0915671 *Contracts for Precise Types*. This material is based upon work supported by the DARPA CRASH program through the United States Air Force Research Laboratory (AFRL) under Contract No. FA8650-10-C-7090. The views expressed are those of the authors and do not reflect the official policy or position of the Department of Defense or the U.S. Government. Finally, the support of Hannah de Keijzer and the patience of the staff of B2 are greatly appreciated.

## References

- [1] M. Abadi, L. Cardelli, B. Pierce, and G. Plotkin. Dynamic typing in a statically-typed language. In *Principles of Programming Languages (POPL)*, 1989.
- [2] A. Ahmed, R. B. Findler, J. Siek, and P. Wadler. Blame for all. In *Principles of Programming Languages (POPL)*, 2011.
- [3] J. F. Belo, M. Greenberg, A. Igarashi, and B. C. Pierce. Polymorphic contracts. In *European Symposium on Programming (ESOP)*, 2011.
- [4] G. M. Bierman, A. D. Gordon, C. Hrițcu, and D. Langworthy. Semantic subtyping with an SMT solver. In *International Conference on Functional Programming (ICFP)*, 2010.
- [5] V. Breazu-Tannen, T. Coquand, C. A. Gunter, and A. Scedrov. Inheritance as implicit coercion. *Information and Computation*, 93(1):172–221, 1991. ISSN 0890-5401. doi:http://dx.doi.org/10.1016/0890-5401(91)90055-7. Selections from 1989 {IEEE} Symposium on Logic in Computer Science.
- [6] O. Chitil and F. Huch. Monadic, prompt lazy assertions in haskell. In *Asian Symposium on Programming Languages and Systems (APLAS)*, 2007.
- [7] R. Chugh, P. M. Rondon, and R. Jhala. Nested refinements: a logic for duck typing. In *Principles of Programming Languages (POPL)*, POPL '12, pages 231–244, New York, NY, USA, 2012. ACM. ISBN 978-1-4503-1083-3. doi:10.1145/2103656.2103686.
- [8] J. Cretin and D. Rémy. On the power of coercion abstraction. In *Principles of Programming Languages (POPL)*, POPL '12, pages 361–372, New York, NY, USA, 2012. ACM. ISBN 978-1-4503-1083-3. doi:10.1145/2103656.2103699.
- [9] C. Dimoulas, R. B. Findler, C. Flanagan, and M. Felleisen. Correct blame for contracts: no more scapegoating. In *Principles of Programming Languages (POPL)*, 2011. doi:10.1145/1926385.1926410.
- [10] T. Disney and C. Flanagan. Gradual information flow typing. In *Workshop on Script-to-Program Evolution (STOP)*, 2011.
- [11] R. B. Findler and M. Felleisen. Contracts for higher-order functions. In *International Conference on Functional Programming (ICFP)*, 2002.
- [12] R. B. Findler, S.-Y. Guo, and A. Rogers. Implementation and application of functional languages. chapter Lazy Contract Checking for Immutable Data Structures, pages 111–128. Springer-Verlag, Berlin, Heidelberg, 2008. ISBN 978-3-540-85372-5. doi:10.1007/978-3-540-85373-2\_7.
- [13] C. Flanagan. Hybrid type checking. In *Principles of Programming Languages (POPL)*, 2006.
- [14] C. Flanagan, K. R. M. Leino, M. Lillibridge, G. Nelson, J. B. Saxe, and R. Stata. Extended static checking for java. In *Programming Language Design and Implementation (PLDI)*, 2002.
- [15] T. Freeman and F. Pfenning. Refinement types for ML. In *Programming Language Design and Implementation (PLDI)*, June 1991.
- [16] A. D. Gordon and C. Fournet. Principles and applications of refinement types. Technical Report MSR-TR-2009-147, 2008.
- [17] M. Greenberg, B. C. Pierce, and S. Weirich. Contracts made manifest. In *Principles of Programming Languages (POPL)*, 2010.
- [18] J. Gronski and C. Flanagan. Unifying hybrid types and contracts. In *Trends in Functional Programming (TFP)*, 2007.
- [19] F. Henglein. Dynamic typing: Syntax and proof theory. *Sci. Comput. Program.*, 22(3):197–230, 1994.
- [20] D. Herman, A. Tomb, and C. Flanagan. Space-efficient gradual typing. In *Trends in Functional Programming (TFP)*, pages 404–419, Apr. 2007.
- [21] D. Herman, A. Tomb, and C. Flanagan. Space-efficient gradual typing. *Higher Order Symbol. Comput.*, 23(2):167–189, June 2010. ISSN 1388-3690. doi:10.1007/s10990-011-9066-z.
- [22] R. Hinze, J. Jeuring, and A. Löh. Typed contracts for functional programming. In *Functional and Logic Programming (FLOPS)*, 2006.
- [23] C. Hritcu, M. Greenberg, B. Karel, B. C. Pierce, and G. Morrisett. All your ifcexception are belong to us. In *Security and Privacy (SP)*, 2013 *IEEE Symposium on*, pages 3–17, 2013. doi:10.1109/SP.2013.10.
- [24] K. Knowles and C. Flanagan. Hybrid type checking. To appear in *TOPLAS.*, 2010.
- [25] K. Knowles, A. Tomb, J. Gronski, S. N. Freund, and C. Flanagan. Sage: Hybrid checking for flexible specifications. In *Scheme and Functional Programming Workshop*, 2006.
- [26] J. Matthews and A. Ahmed. Parametric polymorphism through runtime sealing or, theorems for low, low prices! In *European Symposium on Programming (ESOP)*, 2008.
- [27] B. Meyer. *Eiffel: the language*. Prentice-Hall, Inc., 1992. ISBN 0-13-247925-7.
- [28] Y. Minamide. Runtime behavior of conversion interpretation of subtyping. In *Selected Papers from the 13th International Workshop on Implementation of Functional Languages*, IFL '02, pages 155–167, London, UK, UK, 2002. Springer-Verlag. ISBN 3-540-43537-9.
- [29] Y. Minamide and J. Garrigue. On the runtime complexity of type-directed unboxing. In *International Conference on Functional Programming (ICFP)*, ICFP '98, pages 1–12, New York, NY, USA, 1998. ACM. ISBN 1-58113-024-4. doi:10.1145/289423.289424.
- [30] X. Ou, G. Tan, Y. Mandelbaum, and D. Walker. Dynamic typing with dependent types. In *IFIP Conference on Theoretical Computer Science (TCS)*, 2004.
- [31] P. M. Rondon, M. Kawaguchi, and R. Jhala. Liquid types. In *Programming Language Design and Implementation (PLDI)*, 2008.
- [32] Z. Shao. Flexible representation analysis. In *International Conference on Functional Programming (ICFP)*, pages 85–98, Amsterdam, The Netherlands, June 1997.
- [33] J. Siek, R. Garcia, and W. Taha. Exploring the design space of higher-order casts. In G. Castagna, editor, *Programming Languages and Systems*, volume 5502 of *Lecture Notes in Computer Science*, pages 17–31. Springer Berlin Heidelberg, 2009. ISBN 978-3-642-00589-3. doi:10.1007/978-3-642-00590-9\_2.
- [34] J. G. Siek and W. Taha. Gradual typing for functional languages. In *Scheme and Functional Programming Workshop*, September 2006.
- [35] J. G. Siek and P. Wadler. Threesomes, with and without blame. In *Principles of Programming Languages (POPL)*, pages 365–376, New York, NY, USA, 2010. ACM. ISBN 978-1-60558-479-9. doi:10.1145/1706299.1706342.
- [36] N. Swamy, M. Hicks, and G. M. Bierman. A theory of typed coercions and its applications. In *International Conference on Functional Programming (ICFP)*, pages 329–340, New York, NY, USA, 2009. ACM. ISBN 978-1-60558-332-7. doi:10.1145/1596550.1596598.
- [37] S. Thatte. Quasi-static typing. In *Principles of Programming Languages (POPL)*, 1990.
- [38] S. Tobin-Hochstadt and M. Felleisen. The design and implementation of typed scheme. In *Principles of Programming Languages (POPL)*, 2008.

- [39] P. Wadler and R. B. Findler. Well-typed programs can't be blamed. In *European Symposium on Programming (ESOP)*, 2009.