

Space-Efficient Gradual Typing in Coercion-Passing Style

YUYA TSUDA, Kyoto University, Japan

ATSUSHI IGARASHI, Kyoto University, Japan

TOMOYA TABUCHI, Kyoto University, Japan

Herman et al. (2007, 2010) pointed out that the insertion of run-time checks into a gradually typed program could hamper tail-call optimization and, as a result, worsen the space complexity of the program. To address the problem, they proposed a space-efficient coercion calculus, which was subsequently improved by Garcia, et al. (2009) and Siek et al. (2015). The semantics of these calculi involves eager composition of run-time checks expressed by coercions to prevent the size of a term from growing. However, it relies also on a nonstandard reduction rule, which does not seem easy to implement. In fact, no compiler implementation of gradually typed languages fully supports the space-efficient semantics faithfully.

In this paper, we study *coercion-passing style*, which Herman et al. have already mentioned, as a technique for straightforward space-efficient implementation of gradually typed languages. A program in coercion-passing style passes “the rest of run-time checks” around—just like continuation-passing style (CPS), in which “the rest of computation” has been passed around—and (unlike CPS) composes coercions eagerly. We give a formal coercion-passing translation from λS by Siek et al. to λS_1 , which is a new calculus of *first-class coercions* tailored for coercion-passing style, and prove correctness of the translation. We also implement our coercion-passing style transformation for the Grift compiler developed by Kuhlenschmidt et al. and give a preliminary experimental result.

Additional Key Words and Phrases: coercion calculus, coercion-passing style, dynamic type checking, gradual typing

ACM Reference Format:

Yuya Tsuda, Atsushi Igarashi, and Tomoya Tabuchi. 2020. Space-Efficient Gradual Typing in Coercion-Passing Style. In *Informal Proceedings of the first ACM SIGPLAN Workshop on Gradual Typing (WGT20)*. ACM, New York, NY, USA, 27 pages.

1 INTRODUCTION

1.1 Space-Efficiency Problem in Gradual Typing

Gradual typing [Siek and Taha 2006; Tobin-Hochstadt and Felleisen 2006] is one of the linguistic approaches to integrating static and dynamic typing. Allowing programmers to mix statically typed program fragments and dynamically typed fragments in a single program, it advocates the “script to program” evolution. Namely, software development starts with simple, often dynamically typed scripts, which evolve to more robust, fully statically typed programs through intermediate stages of partially typed programs. To make this evolution work in practice, it is important that the performance of partially typed programs at intermediate stages is comparable to that of the two ends, that is, dynamically typed scripts and statically typed programs.

However, it has been pointed out that gradual typing suffers from serious efficiency problems from both theoretical and practical view points [Herman et al. 2007, 2010; Takikawa et al. 2016]. In particular, Takikawa et al. [2016] showed that even the state-of-the-art gradual typing implementation could show catastrophic slowdown for partially typed programs due to run-time checking to ensure safety. Worse, such slowdown is not easy to predict because it depends on implicit run-time checks inserted by the language implementation and it requires fairly deep knowledge about the underlying gradual type system to understand how run-time checks are inserted and how they

behave. Since then, several pieces of work have investigated the performance issues [Bauman et al. 2017; Feltey et al. 2018; Kuhlenschmidt et al. 2019; Muehlboeck and Tate 2017; Rastogi et al. 2015; Richards et al. 2017].

Earlier work by Herman et al. [2007, 2010] pointed out a related problem. They showed that, when values are passed between a statically typed part and a dynamically typed part many times, delayed run-time checks may accumulate and make space complexity of a program worse than an unchecked semantics.

To make the discussion more concrete, consider the following mutually recursive functions (written in ML-like syntax):

```
let rec even (x : int) : ★ =
  if x = 0 then true⟨bool!⟩ else (odd (x - 1))⟨bool!⟩
and odd (x : int) : bool =
  if x = 0 then false else (even (x - 1))⟨bool?p⟩
```

Ignoring the gray part, which will be explained shortly, this is a tail-recursive definition of functions to decide a given integer is even or odd, except that the return type of one of the functions is written ★, which is the dynamic type, which can be any tagged value. This definition expresses a situation where a statically typed function and a dynamically typed function calls each other.¹ The gray part represents inserted run-time checks, written by using Henglein’s coercion syntax [Henglein 1994]: $\text{true}\langle\text{bool!}\rangle$ means that (untagged) Boolean value true will be tagged with bool to make a value of the dynamic type and $(\text{even } (x - 1))\langle\text{bool?}^p\rangle$ means that the value returned from recursive call even $(x - 1)$ will be tested whether it is tagged with bool—if so, the run-time check removes the tag and returns the untagged Boolean value; otherwise, it results in *blame*, which is an uncatchable exception (with label p to indicate where the check has failed).

The crux of this example is that the insertion of run-time checks has broken tail recursion: due to the presence of $\langle\text{bool!}\rangle$ and $\langle\text{bool?}^p\rangle$, the recursive calls are not in tail positions any longer. So, according to the original semantics of coercions [Henglein 1994], evaluation of odd 10 as follows:

$$\begin{aligned} \text{odd } 10 &\mapsto^* (\text{even } 9)\langle\text{bool?}^p\rangle \\ &\mapsto^* (\text{odd } 8)\langle\text{bool!}\rangle\langle\text{bool?}^p\rangle \\ &\mapsto^* (\text{even } 7)\langle\text{bool?}^p\rangle\langle\text{bool!}\rangle\langle\text{bool?}^p\rangle \\ &\mapsto \dots \\ &\mapsto^* \text{false}\langle\text{bool!}\rangle\langle\text{bool?}^p\rangle \dots \langle\text{bool!}\rangle\langle\text{bool?}^p\rangle \\ &\mapsto^* \text{false} \end{aligned}$$

Thus, the size of a term being evaluated is proportional to the argument n at its longest, whereas unchecked semantics (without coercions) allows for tail-call optimization and constant-space execution. This is the space-efficiency problem of gradual typing.

1.2 Space-Efficient Gradual Typing

Herman et al. [2007, 2010] also presented a solution to this problem. In the evaluation sequence of odd n above, we could immediately “compress” nested coercion applications $M\langle\text{bool!}\rangle\langle\text{bool?}^p\rangle$ before computation of the target term M ends, because $\langle\text{bool!}\rangle\langle\text{bool?}^p\rangle$ —tagging immediately followed by untagging—does virtually nothing. By doing so, we can maintain that the order of the size of a term in the middle of evaluation is constant. This idea is formalized in terms of a “space-efficient” extension of the coercion calculus [Henglein 1994]. Since then, a few space-efficient coercion/cast calculi have been proposed [Siek et al. 2009, 2015; Siek and Wadler 2010].

¹In this sense, the argument of even should have been ★, too, but it would clutter the code after inserting run-time checks.

<pre> odd 4 ↦* (even 3)⟨bool?P⟩ ↦ (odd (3 - 1))⟨bool!⟩⟨bool?P⟩ ↦ (odd (3 - 1))⟨bool! ; bool?P⟩ = (odd (3 - 1))⟨id_{bool}⟩ ↦ (odd 2)⟨id_{bool}⟩ ↦ (even (2 - 1))⟨bool?P⟩⟨id_{bool}⟩ ↦ (even (2 - 1))⟨bool?P ; id_{bool}⟩ = (even (2 - 1))⟨bool?P⟩ ↦ (even 1)⟨bool?P⟩ ↦ ... </pre>	<pre> oddk (4, id_{bool}) ↦ evenk (4 - 1, bool?P ;; id_{bool}) ↦ evenk (4 - 1, bool?P) ↦ evenk (3, bool?P) ↦ oddk (3 - 1, bool! ;; bool?P) ↦ oddk (3 - 1, id_{bool}) ↦ oddk (2, id_{bool}) ↦ evenk (2 - 1, bool?P ;; id_{bool}) ↦ evenk (2 - 1, bool?P) ↦ evenk (1, bool?P) ↦ ... </pre>
--	--

Fig. 1. Reduction from odd 4 in λS (left) and reduction from odd (4, id_{bool}) in λS_1 (right).

Among them, [Siek et al. \[2015\]](#) have proposed a space-efficient coercion calculus λS . λS is equipped with a composition function that compresses consecutive coercions. The coercion composition is achieved as a simple recursive function thanks to the restriction of coercions to canonical ones. We show evaluation of odd 4 according to the λS semantics in the left of Figure 1.² Here, $s ; t$ is a meta-level operation that composes two coercions s, t (into the canonical form) and yields another coercion that corresponds to their sequential composition. This composition function enables us to prevent the size of a term from growing.

However, in order to ensure that nested coercion applications are always merged, the operational semantics of λS relies on a nonstandard reduction rule and nonstandard evaluation contexts. Although it does not cause any theoretical problems, it does not seem easy to implement—in particular, its compilation method seems nontrivial.

1.3 Our Work: Coercion-Passing Style

In this paper, we study coercion-passing style for space-efficient gradual typing. Just as continuation-passing style, in which “the rest of the computation” is passed around as first-class functions and every function call is at a tail position, a program in coercion-passing style passes “the rest of the run-time checks” around and every function call is at a tail position. Actually, the idea of coercion-passing style has already been listed as one of the possible implementation techniques by [Herman et al. \[2007, 2010\]](#) but it has not been well studied nor formalized.

We use the even/odd example above to describe our approach to the problem. Here are the even/odd functions in coercion-passing style.

```

let rec evenk (x,  $\kappa$ ) =
  if x = 0 then true⟨bool! ;;  $\kappa$ ⟩ else oddk (x - 1, bool! ;;  $\kappa$ )
and oddk (x,  $\kappa$ ) =
  if x = 0 then false⟨ $\kappa$ ⟩ else evenk (x - 1, bool?P ;;  $\kappa$ )

```

Additional parameters named κ are for *first-class coercions*, which are supposed to be applied to return values as in `false(κ)`. We often call these coercions *continuation coercions*. Coercion applications such as `true⟨bool!⟩` and `(oddk (x - 1))⟨bool!⟩` at tail positions in the original program are translated to coercion compositions such as `true⟨bool! ;; κ ⟩` and `oddk (x - 1,`

²Strictly speaking, `bool!` and `bool?P` are abbreviations of `idbool`; `bool!` and `bool?P; idbool`, respectively, in λS .

$\text{bool!} \ ; \ ; \ \kappa$), respectively. When κ is bound to a concrete coercion, it will be composed with bool! before it is applied. Similarly to programs in CPS, function calls pass (composed) coercions.

With these functions in coercion-passing style, the evaluation of $\text{oddk}(4, \text{id}_{\text{bool}})$ (where id_{bool} is an identity coercion, which does nothing) proceeds as in the left of Figure 1. Since tagging followed by untagging (with the same tag) actually does nothing, $\text{bool!} \ ; \ ; \ \text{bool}?^p$ composes to id_{bool} by the (meta-level) coercion composition $\text{bool!} \ ; \ ; \ \text{bool}?^p$.

Similarly to the λS semantics described above, coercion composition in the argument takes place before a recursive call, thus the size of coercions stays at the constant order, overcoming the space efficiency problem. A nice property of our solution is that the evaluation is standard call-by-value.

Contributions. Our contributions in this paper are summarized as follows.

- In the context of the space-efficiency problem of gradual typing, we develop a new space-efficient coercion calculus λS_1 with first-class coercions.
- We formalize a coercion-passing translation from (a slight variant of) space-efficient coercion calculus λS [Siek et al. 2015] to the new calculus λS_1 .
- We prove correctness of the coercion-passing translation.
- We implement the coercion-passing translation on top of the Grift compiler [Kuhlenschmidt et al. 2019], and conduct a preliminary experiment.

Outline. The rest of this paper is organized as follows. We review the space-efficient coercion calculus λS [Siek et al. 2015] in Section 2. We introduce a new space-efficient coercion calculus with first-class coercions λS_1 in Section 3, formalize a translation into coercion-passing style as a translation from λS to λS_1 , and prove correctness of the translation in Section 4. We discuss our implementation of coercion-passing translation on top of the Grift compiler [Kuhlenschmidt et al. 2019] and show a preliminary experimental result in Section 5. Finally, we discuss related work in Section 6 and conclude in Section 7. Proofs of the stated properties can be found in the full version: <https://arxiv.org/abs/1908.02414>.

2 SPACE-EFFICIENT COERCION CALCULUS

In this section, we review the space-efficient coercion calculus λS [Siek et al. 2015], which is the source calculus of our translation. Our definition differs from the original in a few respects, as we will explain later. For simplicity, we do not include (mutually) recursive functions and conditional expressions in the formalization but it is straightforward to add them; in fact, our implementation includes them.

Main novelties of λS over the original coercion calculus λC [Henglein 1994] are (1) space-efficient coercions, which are canonical forms of coercions, whose composition can be defined by a straightforward recursive function and (2) operational semantics in which a sequence of coercion applications are collapsed eagerly—even before they are applied to a value [Herman et al. 2007, 2010; Siek et al. 2009].

Basic forms of coercions are inherited from λC [Henglein 1994], which provides (1) identity coercions id_A (where A is a type), which do nothing; (2) injections $G!$, which add a type tag G to a value to make a value of the dynamic type; (3) projections $G?^p$, which test whether a value of the dynamic type is tagged with G , remove the tag if the test succeeds, or raise blame labeled p if it fails; (4) function coercions $c_1 \rightarrow c_2$, which, when they are applied to a function, coerce an argument to the function by c_1 and a value returned from the function by c_2 ; and (5) sequential compositions $c_1; c_2$, which apply c_1 and c_2 in this order. Space-efficient coercions restrict the way basic coercions are combined by sequential composition; they can be roughly expressed by the

Variables	x, y	Constants	a, b	Operators	op	Blame labels	p
Base types			$\iota ::= \text{int} \mid \text{bool} \mid \dots$				
Types			$A, B, C ::= \star \mid \iota \mid A \rightarrow B$				
Ground types			$G, H ::= \iota \mid \star \rightarrow \star$				
Space-efficient coercions			$s, t ::= \text{id}_\star \mid G^{?P}; i \mid i$				
Intermediate coercions			$i ::= g; G! \mid g \mid \perp^{GpH}$				
Ground coercions			$g, h ::= \text{id}_A \text{ (if } A \neq \star) \mid s \rightarrow t \text{ (if } s \neq \text{id or } t \neq \text{id)}$				
Delayed coercions			$d ::= g; G! \mid s \rightarrow t \text{ (if } s \neq \text{id or } t \neq \text{id)}$				
Terms			$L, M, N ::= V \mid op(M, N) \mid MN \mid M\langle s \rangle \mid \text{blame } p$				
Values			$V, W ::= x \mid U \mid U\langle\langle d \rangle\rangle$				
Uncoerced values			$U ::= a \mid \lambda x. M$				
Type environments			$\Gamma ::= \emptyset \mid \Gamma, x : A$				

Fig. 2. Syntax of λS .

following regular expression:

$$(G^{?P};)^?(id_i + (s_1 \rightarrow s_2)); (G'!)^?$$

(where ι is a base type, s_1 and s_2 stand for space efficient coercions, $(\dots)^?$ stands for an optional element, and $+$ for alternatives). As already mentioned, an advantage of this form is that (meta-level) sequential composition (denoted by $s_1 \circledast s_2$) of two space-efficient coercions results in another space-efficient coercion (if the composition is well typed). For example, the composition

$$((G_1^{?P};)^?(id_i + (s_1 \rightarrow s_2)); G_2!) \circledast (G_3^{?P'}; (id_i + (s_3 \rightarrow s_4)); G_4!)^?$$

will be

$$((G_1^{?P};)^?(id_i + ((s_3 \circledast s_1) \rightarrow (s_2 \circledast s_4)); G_4!)^?)$$

if $G_2 = G_3$ —that is, tagging with G_2 is immediately followed by inspection whether G_2 is present. Notice that the resulting coercion conforms to the regular expression again. (The other case where $G_2 \neq G_3$ means the projection $G_3^{?P'}$ will fail; we will explain such failures later.)

The operational semantics includes the following reduction rule

$$\mathcal{F}[M\langle s \rangle\langle t \rangle] \longrightarrow \mathcal{F}[M\langle s \circledast t \rangle]$$

where \mathcal{F} is an evaluation context that does not include nested coercion applications and whose innermost frame is not a coercion application. This rule intuitively means that two consecutive coercions at the outermost position will be composed *even before* M is evaluated to a value. This eager composition avoids a long chain of coercion applications in an evaluation context.

2.1 Syntax

We show the syntax of λS in Figure 2. The syntax of λS extends that of the simply typed lambda calculus (written in gray) with the dynamic type and (space-efficient) coercions.

Types, ranged over by A, B, C , include the dynamic type \star , base types ι , and function types $A \rightarrow B$. Base types ι include `int` (integer type) and `bool` (Boolean type) and so on. *Ground types*, ranged over by G, H , include base types ι and the function type $\star \rightarrow \star$. They are used for type tags put on values of the dynamic type [Wadler and Findler 2009]. Here, the ground type for functions is always

$\star \rightarrow \star$, reflecting the fact that many dynamically typed languages do not include information on the argument and return types of the function in its type tag.

As we have already discussed, λS restricts coercions to only canonical ones, namely space-efficient coercions s , whose grammar is defined via ground coercions g and intermediate coercions i . Ground coercions correspond to the middle part of space-efficient coercions; unlike the original λS , ground coercions include identity coercions for any function types—such as $\text{id}_{i \rightarrow i}$ —and exclude “virtually identity” coercions such as $\text{id}_i \rightarrow \text{id}_i$. Although these two coercions are extensionally the same, they reduce in slightly different ways: applying $\text{id}_{i \rightarrow i}$ to a function immediately returns the function, whereas applying $\text{id}_i \rightarrow \text{id}_i$ results in a wrapped function whose argument and return values are monitored by id_i , which does nothing. Adopting id_A for any A simplifies our proof that the coercion-passing translation preserves the semantics. An intermediate coercion adds an optional injection to a ground coercion. Coercions of the form \perp^{GpH} trigger blame (labeled p) if applied to a value. They emerge from coercion composition

$$((G_1 ?^p;)^2 (\text{id}_A + (s_1 \rightarrow s_2)); G_2 !) \circ (G_3 ?^{p'}; (\text{id}_A + (s_3 \rightarrow s_4)))(; G_4 !)^2$$

where $A \neq \star$ and $G_2 \neq G_3$, which means the projection $G_3 ?^{p'}$ is bound to a failure. The composition results in $(G_1 ?^p;)^2 \perp^{G_1 p' G_3}$, which means that, unless the optional projection fails—blaming p —it fails with p' . Finally, space-efficient coercions are obtained by adding optional projection to intermediate coercions. id_\star is a special coercion that does not conform to the regular expression above. Strictly speaking, an injection, say $\text{int}!$, has to be written $\text{id}_{\text{int}}; \text{int}!$ and a projection, say $\text{int}?^p$, has to be written $\text{int}?^p; \text{id}_{\text{int}}$. We often omit these identity coercions in examples.

Terms, ranged over by L, M, N , include values V , primitive binary operations $op(M, N)$, function applications MN , coercion applications $M\langle s \rangle$, and coercion failure blame p . The term $M\langle s \rangle$ coerces the value of M with coercion s at run time. The term blame p denotes a run-time type error caused by the failure of a coercion (projection) with blame label p .

Values, ranged over by V, W , include variables x , uncoerced values U , and coerced values $U\langle\langle d \rangle\rangle$. Uncoerced values, ranged over by U , include constants a of base types and lambda abstractions $\lambda x. M$. Unlike λC , where values can involve nested coercion applications, there is at most one coercion in a value—nested coercions will be composed. Coerced values $U\langle\langle d \rangle\rangle$ have two forms: injected values $U\langle\langle g; G! \rangle\rangle$ and wrapped functions $U\langle\langle s \rightarrow t \rangle\rangle$. The check of function coercion is delayed until wrapped functions are applied to a value [Findler and Felleisen 2002; Henglein 1994; Siek and Taha 2006].

Unlike many other studies on coercion and blame calculi, we syntactically distinguish coerced values $U\langle\langle d \rangle\rangle$ from $U\langle d \rangle$ (similarly to Wadler and Findler [2009]). This distinction plays an important role in our correctness proof; roughly speaking, without the distinction, $U\langle d \rangle\langle t \rangle$ would allow two different interpretations: an application of t to a value $U\langle d \rangle$ or two applications of d and t to a value U , which would result in different translation results. We also note that variables x are considered values, not uncoerced values, since they can be bound to coerced values at function calls. In other words, we ensure that values are closed under value substitution.

As usual, applications are left-associative and λ extends as far to the right as possible. We do not commit to a particular choice of precedence between function applications and coercion applications. So, we will always use parentheses to disambiguate terms like $MN\langle t \rangle$. The term $\lambda x. M$ binds x in M as usual. The definitions of free variables and α -equivalence of terms are standard, and thus we omit them. We identify α -equivalent terms.

The metavariable Γ ranges over *type environments*. A type environment is a sequence of pairs of a variable and its type.

Well-formed coercions

$$\begin{array}{c}
\frac{}{G! : G \rightsquigarrow \star} \text{CT-INJ} \quad \frac{}{G?^P : \star \rightsquigarrow G} \text{CT-PROJ} \quad \frac{A \neq \star \quad A \sim G \quad G \neq H}{\perp^{GpH} : A \rightsquigarrow B} \text{CT-FAIL} \\
\frac{}{\text{id}_A : A \rightsquigarrow A} \text{CT-ID} \quad \frac{c_1 : A' \rightsquigarrow A \quad c_2 : B \rightsquigarrow B'}{c_1 \rightarrow c_2 : A \rightarrow B \rightsquigarrow A' \rightarrow B'} \text{CT-FUN} \quad \frac{c_1 : A \rightsquigarrow B \quad c_2 : B \rightsquigarrow C}{(c_1; c_2) : A \rightsquigarrow C} \text{CT-SEQ}
\end{array}$$

Term typing

$$\begin{array}{c}
\frac{}{\Gamma \vdash a : \text{ty}(a)} \text{T-CONST} \quad \frac{\text{ty}(op) = \iota_1 \rightarrow \iota_2 \rightarrow \iota \quad \Gamma \vdash M : \iota_1 \quad \Gamma \vdash N : \iota_2}{\Gamma \vdash op(M, N) : \iota} \text{T-OP} \\
\frac{(x : A) \in \Gamma}{\Gamma \vdash x : A} \text{T-VAR} \quad \frac{\Gamma, x : A \vdash M : B}{\Gamma \vdash \lambda x. M : A \rightarrow B} \text{T-ABS} \quad \frac{\Gamma \vdash M : A \rightarrow B \quad \Gamma \vdash N : A}{\Gamma \vdash MN : B} \text{T-APP} \\
\frac{\Gamma \vdash M : A \quad s : A \rightsquigarrow B}{\Gamma \vdash M\langle s \rangle : B} \text{T-CRC} \quad \frac{\emptyset \vdash U : A \quad d : A \rightsquigarrow B}{\emptyset \vdash U\langle\langle d \rangle\rangle : B} \text{T-CRCV} \quad \frac{}{\emptyset \vdash \text{blame } p : A} \text{T-BLAME}
\end{array}$$

Fig. 3. Typing rules of λS .

2.2 Type System

We give the type system of λS , which consists of three judgments for *type consistency* $A \sim B$, *well-formed coercions* $c : A \rightsquigarrow B$, and *typing* $\Gamma \vdash_S M : A$. The inference rules (except for $A \sim B$) are shown in Figure 3. (We omit the subscript S on \vdash in rules, as some of them are reused for λS_1 .)

The type consistency relation $A \sim B$ is the least reflexive and symmetric and compatible relation that contains $A \sim \star$. As this is standard [Siek and Taha 2006], we omit inference rules here. (We have them in the full version.)

The relation $c : A \rightsquigarrow B$ means that coercion c , which ranges over all kinds of coercions, converts a value from type A to type B . We often call A and B the source and target types of c , respectively. The rule (CT-ID) is for identity coercion id_A . The rule (CT-INJ) is for injection $G!$, which converts type G to type \star . The rule (CT-PROJ) is for projection $G?^P$, which converts type \star to type G . The rule (CT-FUN) is for function coercion $c_1 \rightarrow c_2$. If its argument coercion c_1 converts type A' to type A and its return-value coercion c_2 converts type B to type B' , then function coercion $c_1 \rightarrow c_2$ converts type $A \rightarrow B$ to type $A' \rightarrow B'$. In other words, function coercions are contravariant in their argument coercions and covariant in return-value coercions. The rule (CT-FAIL) is for failure coercion \perp^{GpH} . Here, the source type is not necessarily G but can be any nondynamic type A consistent with G because the source type of a failure coercion may change during coercion composition. For example, the following judgments are derivable:

$$\begin{array}{ccc}
(\text{id}_{\text{int}}; \text{int}!) \rightarrow (\text{int}?^P; \text{id}_{\text{int}}) & : \star \rightarrow \star & \rightsquigarrow \text{int} \rightarrow \text{int} \\
\perp^{\star \rightarrow \star p_{\text{int}}} & : \text{int} \rightarrow \text{bool} & \rightsquigarrow \text{int}
\end{array}$$

Proposition 1 below, which is about the source and target types of intermediate coercions and ground coercions, is useful to understand the syntactic structure of space-efficient coercions. In particular, it states that neither the source nor target type of ground coercions g is the type \star .

PROPOSITION 1 (SOURCE AND TARGET TYPES).

- (1) If $i : A \rightsquigarrow B$ then $A \neq \star$.
- (2) If $g : A \rightsquigarrow B$, then $A \neq \star$ and $B \neq \star$ and there exists a unique G such that $A \sim G$ and $G \sim B$.

The judgment $\Gamma \vdash_S M : A$ means that the λ S-term M is given type A under type environment Γ . When clear from the context, we sometimes write \vdash for \vdash_S with the subscript S omitted. We adopt similar conventions for other relations (such as $\vdash \rightarrow_S$) introduced later.

The rules (T-CONST), (T-OP), (T-VAR), (T-ABS), and (T-APP) are standard. Here, $ty(a)$ maps constant a to a base type ι , and $ty(op)$ maps binary operator op to a (first-order) function type $\iota_1 \rightarrow \iota_2 \rightarrow \iota$. The rule (T-CRC) states that if M is given type A and space-efficient coercion s converts type A to B , then coercion application $M\langle s \rangle$ has type B . The rule (T-CRCV) is similar to (T-CRC), but for coerced values $U\langle\langle d \rangle\rangle$. The rule (T-BLAME) allows blame p to have an arbitrary type A . Here, type environments are always empty \emptyset in (T-CRCV) and (T-BLAME). It is valid because the terms $U\langle\langle d \rangle\rangle$ and blame p arise only during evaluation, which runs a closed term. In other words, these terms are not written by programmers in the surface language, and also they do not appear as the result of coercion insertion.

2.3 Operational Semantics

2.3.1 Coercion Composition. The coercion composition $s \circledast t$ is a recursive function that takes two space-efficient coercions and computes another space-efficient coercion corresponding to their sequential composition. We show the coercion composition rules in Figure 4. The function is defined in such a way that the form of the first coercion decides which rule to apply.

The rules (CC-IDYNL) and (CC-PROJL) are applied when the first one is not an intermediate coercion. The rules (CC-INJID), (CC-COLLAPSE), (CC-CONFLICT), and (CC-FAILL) are applied when the first one is a (nonground) intermediate coercion, in which case another intermediate coercion is yielded. Here, (CC-COLLAPSE) and (CC-CONFLICT) perform tag checks if an injection and a projection meet. If type tags do not match, a failure coercion arises.

Failure coercions are necessary for eager coercion composition not to change the behavior of ordinary coercion calculus λ C. The term $M\langle G! \rangle\langle H?^p \rangle$ (if $G \neq H$) in λ C evaluates to blame p —after M evaluates to a value. By contrast, the two coercions $G!$ and $H?^p$ in the term $M\langle id_G; G! \rangle\langle H?^p; id_H \rangle$ are eagerly composed in λ S. Raising blame p immediately would not match the semantics of λ C if M evaluates to another blame. \perp^{GpH} must raise blame p only after M evaluates to a value.

The rules (CC-FAILR) and (CC-INJR) are applied when a ground coercion and an intermediate coercion are composed to another intermediate coercion. The rules (CC-FAILL) and (CC-FAILR) represent the propagation of a failure to the context, somewhat similarly to exceptions. The rule (CC-INJR) represents associativity of sequential compositions but \circledast is propagated to the inside.

The rules (CC-IDL), (CC-IDR), and (CC-FUN) are applied when two ground coercions are composed to another ground coercion. They are straightforward except that $id_A \rightarrow id_B$ has to be normalized to $id_{A \rightarrow B}$ (CC-FUN).

We present a few examples of coercion composition below:

$$\begin{aligned} (id_{\text{bool}}; \text{bool}!) \circledast (\text{bool}?^p; id_{\text{bool}}) &= id_{\text{bool}} \circledast id_{\text{bool}} = id_{\text{bool}} \\ (id_{\star \rightarrow \star}; (\star \rightarrow \star)!) \circledast (\text{int}?^p; id_{\text{int}}) &= \perp^{\star \rightarrow \star \text{int}} \\ ((\iota?^p; id_\iota) \rightarrow (id_{\iota'}; \iota'!)) \circledast ((id_i; \iota!) \rightarrow id_\star) &= ((id_i; \iota!) \circledast (\iota?^p; id_i)) \rightarrow ((id_{\iota'}; \iota'!) \circledast id_\star) \\ &= id_i \rightarrow (id_{\iota'}; \iota'!) \end{aligned}$$

These examples involve situations where an injection meets a projection by (CC-COLLAPSE) or (CC-CONFLICT). The third example is by (CC-FUN).

$$\begin{aligned} (\iota?^p; id_\iota) \circledast (id_i; \iota!) &= \iota?^p; (id_i \circledast (id_i; \iota!)) = \iota?^p; ((id_i \circledast id_i); \iota!) = \iota?^p; (id_i; \iota!) \\ (id_i; \iota!) \circledast (\iota?^p; (id_i; \iota!)) &= id_i \circledast (id_i; \iota!) = (id_i \circledast id_i); \iota! = id_i; \iota! \end{aligned}$$

Evaluation contexts	$\mathcal{E} ::= \mathcal{F} \mid \mathcal{F}[\square \langle s \rangle]$	
	$\mathcal{F} ::= \square \mid \mathcal{E}[op(\square, M)] \mid \mathcal{E}[op(V, \square)] \mid \mathcal{E}[\square M] \mid \mathcal{E}[V \square]$	
Coercion composition		$s \circledast t = s'$
	$id_{\star} \circledast t = t$	CC-IDDYNL
	$(G^{?P}; i) \circledast t = G^{?P}; (i \circledast t)$	CC-PROJL
	$(g; G!) \circledast id_{\star} = g; G!$	CC-INJID
	$(g; G!) \circledast (G^{?P}; i) = g \circledast i$	CC-COLLAPSE
	$(g; G!) \circledast (H^{?P}; i) = \perp^{GpH}$ if $G \neq H$	CC-CONFLICT
	$\perp^{GpH} \circledast s = \perp^{GpH}$	CC-FAILL
	$g \circledast \perp^{GpH} = \perp^{GpH}$	CC-FAILR
	$g \circledast (h; H!) = (g \circledast h); H!$	CC-INJR
	$id_A \circledast g = g$ if $A \neq \star$	CC-IDL
	$g \circledast id_A = g$ if $A \neq \star$ and $g \neq id_A$	CC-IDR
	$(s \rightarrow t) \circledast (s' \rightarrow t') = \begin{cases} id_{A \rightarrow B} & \text{if } s' \circledast s = id_A \text{ and } t \circledast t' = id_B \\ (s' \circledast s) \rightarrow (t \circledast t') & \text{otherwise} \end{cases}$	CC-FUN
Reduction		$M \xrightarrow{e}_S N$ $M \xrightarrow{c}_S N$
	$op(a, b) \xrightarrow{e} \delta(op, a, b)$	R-OP
	$(\lambda x. M) V \xrightarrow{e} M[x := V]$	R-BETA
	$(U \langle \langle s \rightarrow t \rangle \rangle) V \xrightarrow{e} (U(V \langle s \rangle)) \langle t \rangle$	R-WRAP
	$U \langle id_A \rangle \xrightarrow{c} U$	R-ID
	$U \langle \perp^{GpH} \rangle \xrightarrow{c} \text{blame } p$	R-FAIL
	$U \langle d \rangle \xrightarrow{c} U \langle \langle d \rangle \rangle$	R-CRC
	$M \langle s \rangle \langle t \rangle \xrightarrow{c} M \langle s \circledast t \rangle$	R-MERGE C
	$U \langle \langle d \rangle \rangle \langle t \rangle \xrightarrow{c} U \langle d \circledast t \rangle$	R-MERGE V
Evaluation		$M \xrightarrow{e}_{S_1} N$ $M \xrightarrow{c}_{S_1} N$
	$\frac{M \xrightarrow{e} N}{\mathcal{E}[M] \xrightarrow{e} \mathcal{E}[N]}$ E-CTXE	$\frac{M \xrightarrow{c} N}{\mathcal{F}[M] \xrightarrow{c} \mathcal{F}[N]}$ E-CTXC
		$\frac{\mathcal{E} \neq \square}{\mathcal{E}[\text{blame } p] \xrightarrow{e} \text{blame } p}$ E-ABORT

Fig. 4. Reduction/evaluation rules of λS .

As the fourth example shows, a projection followed by an injection does not collapse since the projection might fail. Such a coercion is simplified when it is preceded by another injection (the fifth example).

The following lemma states that composition is defined for two well-formed coercions with matching target and source types.

LEMMA 2. *If $s : A \rightsquigarrow B$ and $t : B \rightsquigarrow C$, then $(s \circ t) : A \rightsquigarrow C$.*

2.3.2 *Evaluation.* We give a small-step operational semantics to λS consisting of two relations on closed terms: the reduction relation $M \longrightarrow_S N$ for basic computation, and the evaluation relation $M \mapsto_S N$ for computing subterms and raising errors.

We show the reduction rules and the evaluation rules of λS in Figure 4. The reduction/evaluation rules are labeled either e or c. The label e is for essential computation, and the label c is for coercion applications. As we see later, this distinction is important in our correctness proof. We write \longrightarrow_S for $\xrightarrow{e}_S \cup \xrightarrow{c}_S$, and \mapsto_S for $\mapsto^e_S \cup \mapsto^c_S$. We sometimes call \mapsto^e_S and \mapsto^c_S e-reduction and c-reduction, respectively.

The rule (R-OP) applies to primitive operations. Here, δ is a (partial) function that takes an operator op and two constants a_1, a_2 , and returns the resulting constant of the primitive operation. We assume that if $ty(op) = \iota_1 \rightarrow \iota_2 \rightarrow \iota$ and $ty(a_1) = \iota_1$ and $ty(a_2) = \iota_2$, then $\delta(op, a_1, a_2) = a$ and $ty(a) = \iota$ for some constant a .

The rule (R-BETA) performs the standard call-by-value β -reduction. We write $M[x := V]$ for capture-avoiding substitution of V for free occurrences of x in M . The definition of substitution is standard, which we omit.

The rule (R-WRAP) applies to applications of wrapped function $U \langle\langle s \rightarrow t \rangle\rangle$ to value V . In this case, we first apply coercion s on the argument to V , and get $V \langle s \rangle$. We next apply function U to $V \langle s \rangle$, and get $U(V \langle s \rangle)$. We then apply coercion t on the returned value, hence $(U(V \langle s \rangle)) \langle t \rangle$.

The rule (R-ID) represents that identity coercion id_A returns the input value U as it is. The rule (R-FAIL) applies to applications of failure coercion \perp^{GpH} to uncoerced value U , which reduces to blame p . The rule (R-CRC) applies to applications $U \langle d \rangle$ of delayed coercion d to uncoerced value U , which reduces to a coerced value $U \langle\langle d \rangle\rangle$.

The rules (R-MERGE) and (R-MERGEV) apply to two consecutive coercion applications, and the two coercions are merged by the composition operation. These rules are key to space efficiency. Thanks to (R-MERGEV), we can assume that there is at most one coercion in a value. The outermost nested coercion applications are merged by (R-MERGE).

Now, we explain *evaluation contexts*, ranged over by \mathcal{E} , shown in the top of Figure 4. Following Siek et al. [2015], we define them in the so-called “inside-out” style [Danvy and Nielsen 2001; Felleisen et al. 1988]. Evaluation contexts represent that function calls in λS are call-by-value and that primitive operations and function applications are evaluated from left to right. The grammar is mutually recursive with \mathcal{F} , which stands for evaluation contexts whose innermost frames are not a coercion application, whereas \mathcal{E} may contain a coercion application as the innermost frame.³ Careful inspection will reveal that both \mathcal{E} and \mathcal{F} contain no consecutive coercion applications. As usual, we write $\mathcal{E}[M]$ for the term obtained by replacing the hole in \mathcal{E} with M . Similarly for $\mathcal{F}[M]$. (We omit their definitions.)

We present a few examples of evaluation contexts below:

$$\begin{aligned} \mathcal{F}_1 &= \square & \mathcal{E}_1 &= \mathcal{F}_1[\square \langle s \rangle] = \square \langle s \rangle \\ \mathcal{F}_2 &= \mathcal{E}_1[V \square] = (V \square) \langle s \rangle & \mathcal{E}_2 &= \mathcal{F}_2[\square \langle t \rangle] = (V(\square \langle t \rangle)) \langle s \rangle \\ \mathcal{F}_3 &= \mathcal{E}_2[\square M] = (V(\square M)) \langle t \rangle \langle s \rangle \end{aligned}$$

We then come back to evaluation rules: The rules (E-CTXE) and (E-CTXC) enable us to evaluate the subterm in an evaluation context. Here, (E-CTXC) requires that computation of coercion applications is only performed under contexts \mathcal{F} —otherwise, the innermost frame may be a coercion application,

³ $\mathcal{F}[\square \langle s \rangle]$ (instead of $\mathcal{F}[\square \langle f \rangle]$) in the definition of \mathcal{E} fixes a problem in Siek et al. [2015] that an identity coercion applied to a nonvalue gets stuck (personal communication).

in which case (R-MERGE) has to be applied first. For example, $U\langle d \rangle \langle t \rangle$ reduces to $U\langle d \circ t \rangle$ rather than $U\langle\langle d \rangle\rangle \langle t \rangle$. The rule (E-ABORT) halts evaluation of a program if it raises blame.

Example 3. Let U be $\lambda x. (x\langle \text{int}^? \rangle + 2)\langle \text{int}! \rangle$. Term $((U\langle \text{int}! \rightarrow \text{int}^? \rangle) 3)\langle \text{int}! \rangle$ evaluates to $5\langle\langle \text{int}! \rangle\rangle$ as follows:

$$\begin{array}{lll}
(U\langle \text{int}! \rightarrow \text{int}^? \rangle) 3 & \mapsto & ((U\langle\langle \text{int}! \rightarrow \text{int}^? \rangle\rangle) 3)\langle \text{int}! \rangle & \text{by (R-CRC)} \\
& \mapsto & (U\langle 3\langle \text{int}! \rangle \rangle)\langle \text{int}^? \rangle \langle \text{int}! \rangle & \text{by (R-WRAP)} \\
& \mapsto & (U\langle 3\langle \text{int}! \rangle \rangle)\langle \text{int}^? ; \text{id} ; \text{int}! \rangle & \text{by (R-MERGE)} \\
& \mapsto & (U\langle 3\langle\langle \text{int}! \rangle\rangle \rangle)\langle \text{int}^? ; \text{id} ; \text{int}! \rangle & \text{by (R-CRC)} \\
& \mapsto & (3\langle\langle \text{int}! \rangle\rangle \langle \text{int}^? \rangle + 2)\langle \text{int}! \rangle \langle \text{int}^? ; \text{id} ; \text{int}! \rangle & \text{by (R-BETA)} \\
& \mapsto & (3\langle\langle \text{int}! \rangle\rangle \langle \text{int}^? \rangle + 2)\langle \text{int}! \rangle & \text{by (R-MERGE)} \\
& \mapsto & (3\langle \text{id} \rangle + 2)\langle \text{int}! \rangle & \text{by (R-MERGEV)} \\
& \mapsto & (3 + 2)\langle \text{int}! \rangle & \text{by (R-ID)} \\
& \mapsto & 5\langle \text{int}! \rangle & \text{by (R-OP)} \\
& \mapsto & 5\langle\langle \text{int}! \rangle\rangle & \text{by (R-CRC)}
\end{array}$$

2.4 Properties

We state a few important properties of λS , including determinacy of the evaluation relation and type safety via preservation and progress [Wright and Felleisen 1994]. We write \mapsto_S^* for the reflexive and transitive closure of \mapsto_S , and \mapsto_S^+ for the transitive closure of \mapsto_S . We say that λS -term M *diverges*, denoted by $M \uparrow_S$, if there exists an infinite evaluation sequence from M .

Proofs of the stated properties are in the full version.

LEMMA 4 (DETERMINACY). *If $M \mapsto_S N$ and $M \mapsto_S N'$, then $N = N'$.*

THEOREM 5 (PROGRESS). *If $\emptyset \vdash_S M : A$, then one of the following holds.*

- (1) $M \mapsto_S M'$ for some M' .
- (2) $M = V$ for some V .
- (3) $M = \text{blame } p$ for some p .

THEOREM 6 (PRESERVATION). *If $\emptyset \vdash_S M : A$ and $M \mapsto_S N$, then $\emptyset \vdash_S N : A$.*

COROLLARY 7 (TYPE SAFETY). *If $\emptyset \vdash_S M : A$, then one of the following holds.*

- (1) $M \mapsto_S^* V$ and $\emptyset \vdash_S V : A$ for some V .
- (2) $M \mapsto_S^* \text{blame } p$ for some p .
- (3) $M \uparrow_S$.

3 SPACE-EFFICIENT FIRST-CLASS COERCION CALCULUS

In this section, we introduce λS_1 , a new space-efficient coercion calculus with first-class coercions; λS_1 serves as the target calculus of the translation into coercion-passing style. The design of λS_1 is tailored to coercion-passing style and, as a result, first-class coercions are not as general as one might expect: for example, coercions for coercions are restricted to identity coercions (e.g., $\text{id}_{\langle \text{int}! \rangle}$).

Since coercions are first-class in λS_1 , the use of (space-efficient) coercions s is not limited to coercion applications $M\langle s \rangle$; they can be passed to a function as an argument, for example. We equip λS with the infix (object-level) operator $M ;; N$ to compute the composition of two coercions: if M and N evaluate to coercions s and t , respectively, then $M ;; N$ reduces to their composition $s \circ t$, which is another space-efficient coercion. The type of (first-class) coercions from A to B is written $A \rightsquigarrow B$.⁴

⁴In λS , \rightsquigarrow is the symbol used in the three-place judgment form $c : A \rightsquigarrow B$, whereas \rightsquigarrow is also a type constructor in λS_1 .

Variables	x, y, κ	Type variables	X, Y
Types	$A, B, C ::= \star \mid \iota \mid A \rightsquigarrow B \mid A \Rightarrow B \mid X$		
Ground types	$G, H ::= \iota \mid \star \Rightarrow \star$		
Space-efficient coercions	$s, t ::= \text{id}_\star \mid G^{?P}; i \mid i$		
Intermediate coercions	$i ::= g; G! \mid g \mid \perp^{GpH}$		
Ground coercions	$g, h ::= \text{id}_A \text{ (if } A \neq \star) \mid s \Rightarrow t \text{ (if } s \neq \text{id or } t \neq \text{id)}$		
Delayed coercions	$d ::= g; G! \mid s \Rightarrow t \text{ (if } s \neq \text{id or } t \neq \text{id)}$		
Terms	$L, M, N ::= V \mid \text{op}(M, N) \mid L(M, N) \mid \text{let } x = M \text{ in } N$ $\mid M ;; N \mid M\langle N \rangle \mid \text{blame } p$		
Values	$V, W, K ::= x \mid U \mid U\langle\langle d \rangle\rangle$		
Uncoerced values	$U ::= a \mid \lambda(x, \kappa). M \mid s$		
Type environments	$\Gamma ::= \emptyset \mid \Gamma, x : A$		

Fig. 5. Syntax of λS_1 .

In λS_1 , every function abstraction takes two arguments, one of which is a parameter for a continuation coercion to be applied to the value returned from this abstraction. For example, $\lambda x. 1$ in λS corresponds to $\lambda(x, \kappa). 1\langle\kappa\rangle$ in λS_1 —here, κ is a coercion parameter. Correspondingly, a function application takes the form $M(N, L)$, which calls function M with an argument pair (N, L) , in which L is a coercion argument, which is applied to the value returned from M . For example, $(f\ 3)\langle s \rangle$ in λS corresponds to $f(3, s)$ in λS_1 ; $(f\ 3)$ (without a coercion application) corresponds to $f(3, \text{id})$.

The type of a function abstraction in λS_1 is written $A \Rightarrow B$, which means that the type of the first argument is the type A and the source type of the second coercion argument is B . An abstraction is polymorphic over the target type of the coercion argument; so, if a function of type $A \Rightarrow B$ is applied to a pair of A and $B \rightsquigarrow C$, then the type of the application will be C . Polymorphism is useful—and in fact required—for coercion-passing translation to work because coercions with different target types may be passed to calls to the same function in λS . Intuitively, $A \Rightarrow B$ means $\forall X. A \times (B \rightsquigarrow X) \rightarrow X$ but we do not introduce \forall -types explicitly because our use of \forall is limited to the target-type polymorphism. However, we do have to introduce type variables for typing function abstractions.

Following the change to function types, function coercions in λS_1 take the form $s \Rightarrow t$. Roughly speaking, its meaning is the same: it coerces an input to a function by s and coerces an output by t . However, due to the coercion passing semantics, there is slight change in how t is used at a function call. Consider $f\langle\langle s \Rightarrow t \rangle\rangle$, i.e., coercion-passing function f wrapped by coercion $s \Rightarrow t$. If the wrapped function is applied to (V, t') , V is coerced by s before passing to f as in λS ; instead of coercing the return value by t , however, t is prepended to t' and passed to f (together with the coerced V) so that the return value is coerced by t and then t' . In the reduction rule, prepending t to t' is represented by composition $t ;; t'$.

3.1 Syntax

We show the syntax of λS_1 in Figure 5. We reuse the same metavariables from λS . We also use κ for variables, and K for values.

Well-formed coercions (replace)

$$\boxed{c : A \rightsquigarrow B}$$

$$\frac{c_1 : A' \rightsquigarrow A \quad c_2 : B \rightsquigarrow B'}{c_1 \Rightarrow c_2 : A \Rightarrow B \rightsquigarrow A' \Rightarrow B'} \text{CT-FUN}$$

Term typing (excerpt)

$$\boxed{\Gamma \vdash_{S_1} M : A}$$

$$\frac{s : A \rightsquigarrow B}{\Gamma \vdash s : A \rightsquigarrow B} \text{T-CRCN} \quad \frac{\Gamma \vdash M : A \rightsquigarrow B \quad \Gamma \vdash N : B \rightsquigarrow C}{\Gamma \vdash M ;; N : A \rightsquigarrow C} \text{T-CMP}$$

$$\frac{\Gamma \vdash M : A \quad \Gamma \vdash N : A \rightsquigarrow B}{\Gamma \vdash M\langle N \rangle : B} \text{T-CRC} \quad \frac{\emptyset \vdash U : A \quad \emptyset \vdash d : A \rightsquigarrow B}{\emptyset \vdash U\langle\langle d \rangle\rangle : B} \text{T-CRCV}$$

$$\frac{\Gamma, x : A, \kappa : B \rightsquigarrow X \vdash M : X \quad (X \text{ does not appear in } \Gamma, A, B)}{\Gamma \vdash \lambda(x, \kappa). M : A \Rightarrow B} \text{T-ABS}$$

$$\frac{\Gamma \vdash M : A \quad \Gamma, x : A \vdash N : B}{\Gamma \vdash \text{let } x = M \text{ in } N : B} \text{T-LET} \quad \frac{\Gamma \vdash L : A \Rightarrow B \quad \Gamma \vdash M : A \quad \Gamma \vdash N : B \rightsquigarrow C}{\Gamma \vdash L(M, N) : C} \text{T-APP}$$

Fig. 6. Typing rules of λS_1 .

We replace $A \rightarrow B$ with $A \Rightarrow B$ and add $A \rightsquigarrow B$ and type variables to types. The syntax for ground types and space-efficient, intermediate, ground, and delayed coercions is the same except that \rightarrow is replaced with \Rightarrow , similarly to types. As we have mentioned, we replace abstractions and applications with two-argument versions. We also add let-expressions (although they could be introduced as derived forms) and coercion composition $M ;; N$. The syntax for coercion applications is now $M\langle N \rangle$, where N is a general term (of type $A \rightsquigarrow B$). Uncoerced values now include space-efficient coercions.

The term $\lambda(x, \kappa). M$ binds x and κ in M , and the term $\text{let } x = M \text{ in } N$ binds x in N . The definitions of free variables and α -equivalence of terms are standard, and thus we omit them. We identify α -equivalent terms.

The definition of type environments, ranged over by Γ , is the same as λS .

3.2 Type System

The type system of λS_1 is straightforward adaption of that of λS . Main rules are shown in Figure 6.

The relation $c : A \rightsquigarrow B$ is mostly the same as that of λS . We replace the rule (CT-FUN) as shown. As in λS , function coercions are contravariant in their argument coercions and covariant in their return-value coercions.

The judgment $\Gamma \vdash_{S_1} M : A$ means that term M of λS_1 has type A under type environment Γ . The rules (T-CONST), (T-OP), (T-VAR), and (T-BLAME) are the same as λS , and so we omit them. The rule (T-LET) is standard.

The rules (T-ABS) and (T-APP) look involved but the intuition that $A \Rightarrow B$ corresponds to $\forall X. A \times (B \rightsquigarrow X) \rightarrow X$ should help to understand them. The rule (T-ABS) assigns type $A \Rightarrow B$ to an abstraction $\lambda(x, \kappa). M$ if the body is well typed under the assumption that x is of type A and κ is of type $B \rightsquigarrow X$ for fresh X . The type variable X should not appear in Γ, A, B so that the target type can be polymorphic at call sites. The rule (T-APP) for applications is already explained.

The rule (T-CRCN) assigns type $A \rightsquigarrow B$ to space-efficient coercion s if it converts a value from type A to type B . The rules (T-CRC) and (T-CRCV) are similar to the corresponding rules of λS , but adjusted to first-class coercions.

Evaluation contexts	$\mathcal{E} ::= \square \mid \mathcal{E}[\square(M, N)] \mid \mathcal{E}[V(\square, N)] \mid \mathcal{E}[V(W, \square)]$ $\mid \mathcal{E}[op(\square, M)] \mid \mathcal{E}[op(V, \square)] \mid \mathcal{E}[\text{let } x = \square \text{ in } M]$ $\mid \mathcal{E}[\square ;; M] \mid \mathcal{E}[V ;; \square] \mid \mathcal{E}[\square \langle M \rangle] \mid \mathcal{E}[V \langle \square \rangle]$	
Coercion composition (replace)	$(s \Rightarrow t) \circledast (s' \Rightarrow t') = \begin{cases} \text{id}_{A \Rightarrow B} & \text{if } s' \circledast s = \text{id}_A \text{ and } t \circledast t' = \text{id}_B \\ (s' \circledast s) \Rightarrow (t \circledast t') & \text{otherwise} \end{cases}$	<div style="border: 1px solid black; padding: 2px; display: inline-block;">$s \circledast t = s'$</div> CC-FUN
Reduction	$op(a, b) \xrightarrow{e} \delta(op, a, b) \quad \text{R-OP}$ $(\lambda(x, \kappa). M)(V, W) \xrightarrow{e} M[x := V, \kappa := W] \quad \text{R-BETA}$ $(U\langle\langle s \Rightarrow t \rangle\rangle)(V, W) \xrightarrow{e} \text{let } \kappa = t ;; W \text{ in } U(V\langle s \rangle, \kappa) \quad \text{R-WRAP}$ $\text{let } x = V \text{ in } M \xrightarrow{c} M[x := V] \quad \text{R-LET}$ $s ;; t \xrightarrow{c} s \circledast t \quad \text{R-CMP}$ $U\langle \text{id}_A \rangle \xrightarrow{c} U \quad \text{R-ID}$ $U\langle \perp^{GpH} \rangle \xrightarrow{c} \text{blame } p \quad \text{R-FAIL}$ $U\langle d \rangle \xrightarrow{c} U\langle\langle d \rangle\rangle \quad \text{R-CRC}$ $U\langle\langle d \rangle\rangle \langle t \rangle \xrightarrow{c} U\langle d ;; t \rangle \quad \text{R-MERGEV}$	<div style="border: 1px solid black; padding: 2px; display: inline-block;">$M \xrightarrow{e}_{S_1} N$</div> <div style="border: 1px solid black; padding: 2px; display: inline-block;">$M \xrightarrow{c}_{S_1} N$</div>
Evaluation	$\frac{M \xrightarrow{x} N \quad \mathcal{X} \in \{e, c\}}{\mathcal{E}[M] \mapsto^{\mathcal{X}} \mathcal{E}[N]} \quad \text{E-CTX} \quad \frac{\mathcal{E} \neq \square}{\mathcal{E}[\text{blame } p] \mapsto^e \text{blame } p} \quad \text{E-ABORT}$	<div style="border: 1px solid black; padding: 2px; display: inline-block;">$M \mapsto^e_{S_1} N$</div> <div style="border: 1px solid black; padding: 2px; display: inline-block;">$M \mapsto^c_{S_1} N$</div>

Fig. 7. Reduction/evaluation rules of λS_1 .

3.3 Operational Semantics

The composition function $s \circledast t$ is mostly the same as that of λS . We only replace (CC-FUN) as shown in Figure 7. As in λS , function coercions are contravariant in their argument coercions and covariant in their return-value coercions.

Similarly to λS , we give a small-step operational semantics to λS_1 consisting of two relations on closed terms: the reduction relation $M \longrightarrow_{S_1} N$ and the evaluation relation $M \mapsto_{S_1} N$. We show the reduction/evaluation rules of λS_1 in Figure 7. As in λS , they are labeled either e or c. We write \longrightarrow_{S_1} for $\xrightarrow{e}_{S_1} \cup \xrightarrow{c}_{S_1}$, and \mapsto_{S_1} for $\mapsto^e_{S_1} \cup \mapsto^c_{S_1}$.

The rules (R-OP) and (R-BETA) are standard. Note that (R-BETA) is adjusted for pair arguments. We write $M[x := V, \kappa := K]$ for capture-avoiding simultaneous substitution of V and K for x and κ , respectively, in M .

The rule (R-WRAP) applies to applications of wrapped function $U\langle\langle s \Rightarrow t \rangle\rangle$ to value V . Since coercion s is for function arguments, it is applied to V , as in λS . Additionally, we compose coercion t on the return value with continuation coercion W . Thus, $V\langle s \rangle$ and $t ;; W$ are passed to function U . Note that we use a let expression to evaluate the second argument $t ;; W$ before $V\langle s \rangle$. It is a necessary adjustment for the semantics of λS and λS_1 to match.

The rule (R-LET) is standard; it is labeled as c because we use let-expressions only for coercion compositions. The rule (R-CMP) applies to coercion compositions $s ;; t$, which is evaluated by meta-level coercion composition function $s \circledast t$. The rules (R-ID), (R-FAIL), (R-CRC), and (R-MERGEV) are the same as λS .

Now, we explain evaluation contexts, ranged over by \mathcal{E} , shown in the top of Figure 7. In contrast to λS , evaluation contexts are standard in λS_1 . The definition of evaluation contexts \mathcal{E} represents that function calls in λS_1 are call-by-value, and primitive operations, function applications, coercion compositions, and coercion applications are all evaluated from left to right.

We then come back to evaluation rules: The evaluation rules (E-CTX) and (E-ABORT) are the same as λS . (However, evaluation contexts in (E-CTX) are more straightforward in λS_1 .)

Finally, we should emphasize that we no longer need (R-MERGEV) in λS_1 . So, λS_1 is an ordinary call-by-value language and its semantics should be easy to implement.

Example 8. Let U be $\lambda(x, \kappa). \text{let } \kappa' = \text{int!} ;; \kappa \text{ in } (x \langle \text{int}^{?p} \rangle + 2) \langle \kappa' \rangle$, which corresponds to the λS -term $\lambda x. (x \langle \text{int}^{?p} \rangle + 2) \langle \text{int!} \rangle$ in Example 3. In fact, we will obtain this term as a result of our coercion-passing translation defined in the next section. The term $(U \langle \text{int!} \Rightarrow \text{int}^{?p} \rangle) (3, \text{int!})$ evaluates to $5 \langle \langle \text{int!} \rangle \rangle$ as follows:

$$\begin{array}{ll}
(U \langle \text{int!} \Rightarrow \text{int}^{?p} \rangle) (3, \text{int!}) & \\
\mapsto (U \langle \langle \text{int!} \Rightarrow \text{int}^{?p} \rangle \rangle) (3, \text{int!}) & \text{by (R-CRC)} \\
\mapsto \text{let } \kappa'' = \text{int}^{?p} ;; \text{int! in } U (3 \langle \text{int!} \rangle, \kappa'') & \text{by (R-WRAP)} \\
\mapsto \text{let } \kappa'' = \text{int}^{?p}; \text{id; int! in } U (3 \langle \text{int!} \rangle, \kappa'') & \text{by (R-CMP)} \\
\mapsto U (3 \langle \text{int!} \rangle, (\text{int}^{?p}; \text{id; int!})) & \text{by (R-LET)} \\
\mapsto U (3 \langle \langle \text{int!} \rangle \rangle, (\text{int}^{?p}; \text{id; int!})) & \text{by (R-CRC)} \\
\mapsto \text{let } \kappa' = \text{int!} ;; (\text{int}^{?p}; \text{id; int!}) \text{ in } (3 \langle \langle \text{int!} \rangle \rangle \langle \text{int}^{?p} \rangle + 2) \langle \kappa' \rangle & \text{by (R-BETA)} \\
\mapsto \text{let } \kappa' = \text{int! in } (3 \langle \langle \text{int!} \rangle \rangle \langle \text{int}^{?p} \rangle + 2) \langle \kappa' \rangle & \text{by (R-CMP)} \\
\mapsto (3 \langle \langle \text{int!} \rangle \rangle \langle \text{int}^{?p} \rangle + 2) \langle \text{int!} \rangle & \text{by (R-LET)} \\
\mapsto (3 \langle \text{id} \rangle + 2) \langle \text{int!} \rangle & \text{by (R-MERGEV)} \\
\mapsto (3 + 2) \langle \text{int!} \rangle & \text{by (R-ID)} \\
\mapsto 5 \langle \text{int!} \rangle & \text{by (R-OP)} \\
\mapsto 5 \langle \langle \text{int!} \rangle \rangle & \text{by (R-CRC)}
\end{array}$$

It is easy to see that the steps by (R-MERGEV) in Example 3 are simulated by (R-CMP) followed by (R-LET).

3.4 Properties

We state a few properties of λS_1 below. Their proofs are in the full version.

LEMMA 9 (DETERMINACY). *If $M \mapsto_{S_1} N$ and $M \mapsto_{S_1} N'$, then $N = N'$.*

THEOREM 10 (PROGRESS). *If $\emptyset \vdash_{S_1} M : A$, then one of the following holds.*

- (1) $M \mapsto_{S_1}^* M'$ for some M' .
- (2) $M = V$ for some V .
- (3) $M = \text{blame } p$ for some p .

THEOREM 11 (PRESERVATION). *If $\emptyset \vdash_{S_1} M : A$ and $M \mapsto_{S_1} N$, then $\emptyset \vdash_{S_1} N : A$.*

COROLLARY 12 (TYPE SAFETY). *If $\emptyset \vdash_{S_1} M : A$, then one of the following holds.*

- (1) $M \mapsto_{S_1}^* V$ and $\emptyset \vdash_{S_1} V : A$ for some V .
- (2) $M \mapsto_{S_1}^* \text{blame } p$ for some p .
- (3) $M \uparrow_{S_1}$.

Type translation	$\Psi(\star) = \star$	$\Psi(\iota) = \iota$	$\Psi(A \rightarrow B) = \Psi(A) \Rightarrow \Psi(B)$	$\Psi(A) = A'$
Coercion translation	$\Psi(s) = s'$		Value translation	$\Psi(V) = V'$
	$\Psi(\text{id}_A) = \text{id}_{\Psi(A)}$		$\Psi(x) = x$	
	$\Psi(g; G!) = \Psi(g); \Psi(G)!$		$\Psi(a) = a$	
	$\Psi(G?^p; i) = \Psi(G)?^p; \Psi(i)$		$\Psi(\lambda x. M) = \lambda(x, \kappa). (\mathcal{K} \llbracket M \rrbracket \kappa)$	
	$\Psi(s \rightarrow t) = \Psi(s) \Rightarrow \Psi(t)$		$\Psi(U \langle\langle d \rangle\rangle) = \Psi(U) \langle\langle \Psi(d) \rangle\rangle$	
	$\Psi(\perp^{GpH}) = \perp^{GpH}$			
Term translation			$\mathcal{C} \llbracket M \rrbracket = M'$	$\mathcal{K} \llbracket M \rrbracket K = M'$
	$\mathcal{C} \llbracket V \rrbracket = \Psi(V)$			
	$\mathcal{C} \llbracket M^A \rrbracket = \mathcal{K} \llbracket M \rrbracket \text{id}_{\Psi(A)}$	if M is not a value		
	$\mathcal{K} \llbracket V \rrbracket K = \Psi(V) \langle K \rangle$			TR-VAL
	$\mathcal{K} \llbracket \text{op}(M, N) \rrbracket K = \text{op}(\mathcal{C} \llbracket M \rrbracket, \mathcal{C} \llbracket N \rrbracket) \langle K \rangle$			TR-OP
	$\mathcal{K} \llbracket MN \rrbracket K = (\mathcal{C} \llbracket M \rrbracket) (\mathcal{C} \llbracket N \rrbracket, K)$			TR-APP
	$\mathcal{K} \llbracket M \langle s \rangle \rrbracket \text{id} = \mathcal{K} \llbracket M \rrbracket \Psi(s)$			TR-CRCID
	$\mathcal{K} \llbracket M \langle s \rangle \rrbracket K = \text{let } \kappa = \Psi(s) ;; K \text{ in } (\mathcal{K} \llbracket M \rrbracket \kappa)$	if $K \neq \text{id}$		TR-CRC
	$\mathcal{K} \llbracket \text{blame } p \rrbracket K = \text{blame } p$			TR-BLAME

Fig. 8. Translation into coercion-passing style (from λS to λS_1).

4 TRANSLATION INTO COERCION-PASSING STYLE

In this section, we formalize a translation into coercion-passing style as a translation from λS to λS_1 and state its correctness. As its name suggests, this translation is similar to transformations into continuation-passing style (CPS transformations) for the call-by-value λ -calculus [Plotkin 1975].

4.1 Definition of Translation

We give the translation into coercion-passing style by the translation rules presented in Figure 8. In order to distinguish metavariables of λS and λS_1 , we often use blue for the source calculus λS . When we need static type information in translation rules, we write M^A to indicate that term M has type A . Thus, strictly speaking, the translation is defined for type derivations in λS .

Translations for types $\Psi(A)$ and coercions $\Psi(s)$ are very straightforward, thanks to the special type constructor \Rightarrow : they just recursively replace type/coercion constructor \rightarrow with \Rightarrow .

Value translation $\Psi(V)$ and term translation $\mathcal{K} \llbracket M \rrbracket K$ are defined in a mutually recursive manner. In $\mathcal{K} \llbracket M \rrbracket K$, M is a λS -term whereas K is a λS_1 -term, which is either a variable or a λS_1 -coercion. $\mathcal{K} \llbracket M \rrbracket K$ returns a λS_1 -term—in coercion-passing style—that applies K to the value of M .

Value translation $\Psi(V)$ is rather straightforward: every function $\lambda x. M$ is translated to a λS_1 -abstraction that takes as the second argument κ a coercion which is to be applied to the return value. So, the body is translated by term translation $\mathcal{K} \llbracket M \rrbracket \kappa$.

We now describe the translation for terms. We write $\mathcal{K} \llbracket M \rrbracket K$ for the translation of λS -term M with continuation coercion K . We first explain the basic transformation scheme given by the

following simpler rules:

$$\begin{array}{ll}
\mathcal{K}' \llbracket V \rrbracket K = \Psi(V) \langle K \rangle & \text{TR}'\text{-VAL} \\
\mathcal{K}' \llbracket \text{op}(M^1, N^2) \rrbracket K = \text{op}(\mathcal{K}' \llbracket M \rrbracket \text{id}_{i_1}, \mathcal{K}' \llbracket N \rrbracket \text{id}_{i_2}) \langle K \rangle & \text{TR}'\text{-OP} \\
\mathcal{K}' \llbracket M^{A \rightarrow B} N^A \rrbracket K = (\mathcal{K}' \llbracket M \rrbracket \text{id}_{\Psi(A \rightarrow B)}) (\mathcal{K}' \llbracket N \rrbracket \text{id}_{\Psi(A)}, K) & \text{TR}'\text{-APP} \\
\mathcal{K}' \llbracket M \langle s \rangle \rrbracket K = \text{let } \kappa = \Psi(s) ;; K \text{ in } (\mathcal{K}' \llbracket M \rrbracket \kappa) & \text{TR}'\text{-CRC} \\
\mathcal{K}' \llbracket \text{blame } p \rrbracket K = \text{blame } p & \text{TR}'\text{-BLAME}
\end{array}$$

(We put a prime on \mathcal{K} to avoid confusion.)

The rule (TR'-VAL) applies to values V , where we apply coercion K to the result of value translation $\Psi(V)$.

The rule (TR'-OP) applies to primitive operations $\text{op}(M, N)$. We translate the arguments M and N with identity continuation coercions by $\mathcal{K}' \llbracket M \rrbracket \text{id}$ and $\mathcal{K}' \llbracket N \rrbracket \text{id}$ and pass them to the primitive operation. The given continuation coercion K is applied to the result. Translating subexpressions with id is one of the main differences from CPS transformation. While continuations in continuation-passing style capture the whole rest of computation, continuation coercions in coercion-passing style capture only the coercion applied right after the current computation. Since neither M nor N is surrounded by a coercion, they are translated with identity coercions of appropriate types. (Cases where a subexpression itself is a coercion application will be discussed shortly.) Careful readers may notice at this point that left-to-right evaluation of arguments is enforced by the semantics (or the definition of evaluation contexts) of λS , not by the translation. In other words, the correctness of the translation relies on the fact that λS evaluation is left-to-right and call-by-value. This is another point that is different from CPS transformation, which dismisses the distinction of call-by-name and call-by-value.

The rule (TR'-APP) applies to function applications $M N$. We translate function M and argument N with identity continuation coercions just like the case for primitive operations. We then pass the continuation coercion K as the second argument to function $\mathcal{K}' \llbracket M \rrbracket \text{id}$.

The rule (TR'-CRC) applies to coercion applications $M \langle s \rangle$. We can think of the sequential composition of $\Psi(s)$ and K as the continuation coercion for M . Thus, we first compute the composition $\Psi(s) ;; K$, bind its result to κ , and translate M with continuation κ . The let -expression is necessary to compose $\Psi(s)$ and K before evaluating $\mathcal{K}' \llbracket M \rrbracket \kappa$. In general, it is not necessarily the case that $\mathcal{K}' \llbracket M \rrbracket K$ evaluates K first, so if we set $\mathcal{K}' \llbracket M \langle s \rangle \rrbracket K = (\mathcal{K}' \llbracket M \rrbracket (\Psi(s) ;; K))$, then the order of computation would change by the translation and correctness of translation would be harder.

Lastly, the rule (TR'-BLAME) defines the translation of $\text{blame } p$ with continuation K as $\text{blame } p$.

The translation \mathcal{K}' seems acceptable but, just as naive CPS transformation leaves administrative redexes, it leaves many applications of id , which we call *administrative coercions*. We expect M and $\mathcal{K}' \llbracket M \rrbracket K$ behave “similarly” but administrative redexes make it hard to show such semantic correspondence. So, we will optimize the translation so that administrative coercions are eliminated, similarly to CPS transformations that eliminates administrative redexes [Appel 1992; Danvy and Filinski 1992; Danvy and Nielsen 2003; Plotkin 1975; Sabry and Felleisen 1993; Sabry and Wadler 1997; Wand 1991].

The bottom of Figure 8 shows the optimized translation rules. The idea to eliminate administrative coercions is close to the colon translation by Plotkin [1975]: we avoid translating values with administrative coercions. So, we introduce an auxiliary translation function $\mathcal{C} \llbracket M \rrbracket$, which returns $\Psi(V)$ —without a coercion application—if M is a value V and returns $\mathcal{K}' \llbracket M \rrbracket \text{id}$ otherwise. Translation rules for primitive operations and function applications are adapted so that they use $\mathcal{C} \llbracket M \rrbracket$ to translate subexpressions. We also split the rule for coercion applications according to whether K is

id or not. The rule (TR-CRCID), which applies if K is id, optimizes the trivial composition $\Psi(s) \circledast \text{id}$ away.

We present a few examples of the translation below:

$$\begin{aligned}\Psi(5) &= 5 \\ \Psi(\lambda x. x + 1) &= \lambda(x, \kappa). (x + 1)\langle \kappa \rangle \\ \mathcal{K} \llbracket (\lambda x. x) 5 \rrbracket \text{int!} &= (\lambda(x, \kappa). x\langle \kappa \rangle) (5, \text{int!}) \\ \mathcal{K} \llbracket ((\lambda x. x) 5)\langle \text{int!} \rangle \rrbracket \text{int?}^p &= \text{let } \kappa = \text{int!} \;; \text{int?}^p \text{ in } (\lambda(x, \kappa). x\langle \kappa \rangle) (5, \kappa)\end{aligned}$$

The following example shows the translation of the λS -term in Example 3 will be the λS_1 -term in Example 8.

Example 13. Let U be a λS -term $\lambda x. (x\langle \text{int?}^p \rangle + 2)\langle \text{int!} \rangle$.

$$\begin{aligned}\Psi(U) &= \lambda(x, \kappa). (\mathcal{K} \llbracket (x\langle \text{int?}^p \rangle + 2)\langle \text{int!} \rangle \rrbracket \kappa) \\ &= \lambda(x, \kappa). \text{let } \kappa' = \text{int!} \;; \kappa \text{ in } (\mathcal{K} \llbracket (x\langle \text{int?}^p \rangle + 2) \rrbracket \kappa') \\ &= \lambda(x, \kappa). \text{let } \kappa' = \text{int!} \;; \kappa \text{ in } (\mathcal{C} \llbracket x\langle \text{int?}^p \rangle \rrbracket + \mathcal{C} \llbracket 2 \rrbracket)\langle \kappa' \rangle \\ &= \lambda(x, \kappa). \text{let } \kappa' = \text{int!} \;; \kappa \text{ in } ((\mathcal{K} \llbracket x\langle \text{int?}^p \rangle \rrbracket \text{id}) + 2)\langle \kappa' \rangle \\ &= \lambda(x, \kappa). \text{let } \kappa' = \text{int!} \;; \kappa \text{ in } ((\mathcal{K} \llbracket x \rrbracket \text{int?}^p) + 2)\langle \kappa' \rangle \\ &= \lambda(x, \kappa). \text{let } \kappa' = \text{int!} \;; \kappa \text{ in } (x\langle \text{int?}^p \rangle + 2)\langle \kappa' \rangle \\ \mathcal{K} \llbracket ((U\langle \text{int!} \rightarrow \text{int?}^p \rangle) 3) \rrbracket \text{id} &= (\mathcal{K} \llbracket (U\langle \text{int!} \rightarrow \text{int?}^p \rangle) \rrbracket \text{id}) (\mathcal{K} \llbracket 3 \rrbracket \text{id}, \text{id}) \\ &= (\mathcal{K} \llbracket U \rrbracket (\text{int!} \rightarrow \text{int?}^p)) (3, \text{id}) \\ &= (\Psi(U)\langle \text{int!} \Rightarrow \text{int?}^p \rangle) (3, \text{id})\end{aligned}$$

4.2 Correctness of Translation

Having defined the translation, we now state its correctness properties with auxiliary lemmas. (Their proofs are in the full version.)

To begin with, the translation preserves typing. Here, we write $\Psi(\Gamma)$ for the type environment satisfying the following:

$$(x : A) \in \Gamma \text{ if and only if } (x : \Psi(A)) \in \Psi(\Gamma).$$

THEOREM 14 (TRANSLATION PRESERVES TYPING). *If $\Gamma \vdash_S M : A$, then $\Psi(\Gamma) \vdash_{S_1} (\mathcal{K} \llbracket M \rrbracket \text{id}_{\Psi(A)}) : \Psi(A)$.*

As for the preservation of semantics, we will prove the following theorem that states the semantics is preserved by the translation:

THEOREM 15 (TRANSLATION PRESERVES SEMANTICS). *If $\emptyset \vdash_S M : \iota$, then*

- (1) $M \mapsto_S^* a$ iff $\mathcal{K} \llbracket M \rrbracket \text{id}_i \mapsto_{S_1}^* a$;
- (2) $M \mapsto_S^* \text{blame } p$ iff $\mathcal{K} \llbracket M \rrbracket \text{id}_i \mapsto_{S_1}^* \text{blame } p$; and
- (3) $M \uparrow_S$ iff $\mathcal{K} \llbracket M \rrbracket \text{id}_i \uparrow_{S_1}$.

To prove this theorem, it suffices to show the left-to-right direction (Theorem 16 below) for each item because the other direction follows from Theorem 16 together with other properties: for example, if $\emptyset \vdash_S M : \iota$ and $\mathcal{K} \llbracket M \rrbracket \text{id}_i \uparrow_{S_1}$, then M can neither get stuck (by type soundness of λS) nor terminate (as it contradicts the left-to-right direction and the fact that \mapsto_{S_1} is deterministic).

THEOREM 16 (TRANSLATION SOUNDNESS). *Suppose $\Gamma \vdash_S M : A$.*

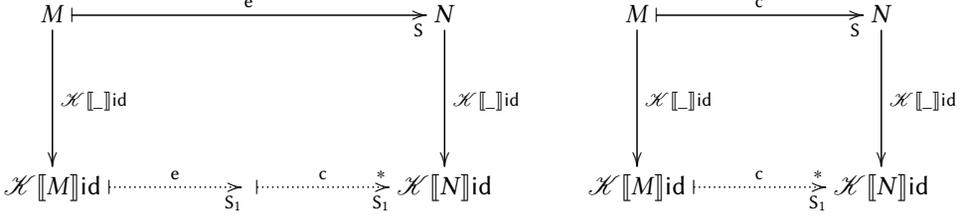
- (1) *If $M \mapsto_S^* V$, then $\mathcal{K} \llbracket M \rrbracket \text{id} \mapsto_{S_1}^* \Psi(V)$.*
- (2) *If $M \mapsto_S^* \text{blame } p$, then $\mathcal{K} \llbracket M \rrbracket \text{id} \mapsto_{S_1}^* \text{blame } p$.*

(3) If $M \uparrow_S$, then $\mathcal{K} \llbracket M \rrbracket \text{id} \uparrow_{S_1}$.

A standard proof strategy would be to show that single-step reduction in the source language is simulated by multi-step reduction in the target language. In fact, we prove the following lemma:

LEMMA 17 (SIMULATION).

- (1) If $M \xrightarrow{e}_S N$, then $\mathcal{K} \llbracket M \rrbracket \text{id} \xrightarrow{e}_{S_1} \xrightarrow{c}_{S_1}^* \mathcal{K} \llbracket N \rrbracket \text{id}$.
- (2) If $M \xrightarrow{c}_S N$, then $\mathcal{K} \llbracket M \rrbracket \text{id} \xrightarrow{c}_{S_1}^* \mathcal{K} \llbracket N \rrbracket \text{id}$.



Whereas a single-step e-reduction in λ_S is translated to one or more steps in λ_{S_1} starting from an e-reduction step, a single-step c-reduction in λ_S can be translated to *zero* steps in λ_{S_1} . An example is $0\langle \text{int!} \rangle \langle \text{id} \rangle \xrightarrow{c}_S 0\langle \text{int!} \rangle$; the two terms both translate to $0\langle \text{int!} \rangle$ by removing `id`. Still an infinite reduction sequence in λ_S is preserved by translation because c-reduction is terminating (proved in the full version) and there is an infinite number of e-reductions.

As is the case for simulation proofs for CPS translation [Appel 1992; Danvy and Filinski 1992; Danvy and Nielsen 2003; Plotkin 1975; Sabry and Felleisen 1993; Sabry and Wadler 1997; Wand 1991], this simulation property⁵ is quite subtle. We discuss this subtlety below.

First, it is important that the translation removes administrative identity coercions by distinguishing values and nonvalues in $\mathcal{C} \llbracket M \rrbracket$. For example, $(\lambda x. x) 5 \xrightarrow{e}_S 5$ holds in λ_S but the translation $\mathcal{K}' \llbracket (\lambda x. x) 5 \rrbracket K$ without removing administrative redexes would yield $(\lambda(x, \kappa). x \langle \kappa \rangle) \langle \text{id} \rangle (5 \langle \text{id} \rangle, K)$, which performs c-reduction before calling the function. More formally, we prove the following lemma, which means the redex in the source is also the redex in the target.

LEMMA 18.

- (1) For any \mathcal{F} , there exist \mathcal{E}' such that for any M , $\mathcal{K} \llbracket \mathcal{F}[M] \rrbracket \text{id} = \mathcal{E}'[\mathcal{C} \llbracket M \rrbracket]$.
- (2) For any \mathcal{F} and s , there exists \mathcal{E}' such that for any M , $\mathcal{K} \llbracket \mathcal{F}[M\langle s \rangle] \rrbracket \text{id} = \mathcal{E}'[\mathcal{K} \llbracket M \rrbracket \Psi(s)]$.

To prove this lemma, the rule (TR-CRCID) also plays an important role: for example, if we removed (TR-CRCID) and used (TR-CRC) for all coercion applications, $\mathcal{K} \llbracket (1 + 1) \langle \text{int!} \rangle \rrbracket \text{id}$ would translate to $\text{let } \kappa = \text{int!} \;; \text{id in } (1 + 1) \langle \kappa \rangle$, which performs c-reduction before adding 1 and 1, which is the first thing the original term $(1 + 1) \langle \text{int!} \rangle$ will do.

Second, optimizing too many (identity) coercions can break simulation. Consider $M \stackrel{\text{def}}{=} (((\lambda x. M_1) \langle \text{id}_i \rightarrow \iota! \rangle)) a \langle \iota?^p \rangle$ and $N \stackrel{\text{def}}{=} ((\lambda x. M_1) (a \langle \text{id}_i \rangle)) \langle \iota! \rangle \langle \iota?^p \rangle$, for which $M \mapsto_S N$ holds by (R-WRAP). Then,

$$\begin{aligned} \mathcal{K} \llbracket M \rrbracket \text{id} &= ((\mathcal{K} \llbracket \lambda(x, \kappa). M_1 \rrbracket \kappa) \langle \langle \text{id}_i \Rightarrow \iota! \rangle \rangle) (a, \iota?^p) \\ &\xrightarrow{S_1} \text{let } \kappa' = \iota! \;; \iota?^p \text{ in } (\mathcal{K} \llbracket \lambda(x, \kappa). M_1 \rrbracket \kappa) (a \langle \text{id}_i \rangle, \kappa') \\ &= \mathcal{K} \llbracket N \rrbracket \text{id}. \end{aligned}$$

⁵If we had been interested only in that translation preserves term equivalence, we could have simplified the technical development by, say, removing the distinction between $U\langle s \rangle$ and $U\langle\langle s \rangle\rangle$. However, simulation is crucial to show that divergence is preserved by translation.

At one point, we defined the translation (let's call it \mathcal{K}'') so that applications of identity coercions would be removed as much as possible, namely,

$$\mathcal{K}''\llbracket N \rrbracket \text{id} = \text{let } \kappa' = \iota! \;; \iota?^p \text{ in } (\mathcal{K}''\llbracket \lambda(x, \kappa). M_1 \rrbracket \kappa) (a, \kappa')$$

(notice that $\langle \text{id}_t \rangle$ on a is removed). Although $\mathcal{K}''\llbracket M \rrbracket \text{id}$ and $\mathcal{K}''\llbracket N \rrbracket \text{id}$ reduced to the same term, we did not quite have $\mathcal{K}''\llbracket M \rrbracket \text{id} \mapsto^+ \mathcal{K}''\llbracket N \rrbracket \text{id}$ as we had expected.

Third, the distinction between $U\langle s \rangle$ and $U\langle\langle s \rangle\rangle$ is crucial to ensure that substitution commutes with the translation:

LEMMA 19 (SUBSTITUTION). *If $\kappa \notin FV(M) \cup FV(V)$, then $(\mathcal{K}\llbracket M \rrbracket \kappa)[x := \Psi(V), \kappa := K] \xrightarrow{c}_{S_1}^* \mathcal{K}\llbracket M[x := V] \rrbracket K$.*

Roughly speaking, if we identified a value $U\langle\langle s \rangle\rangle$ and an application $U\langle s \rangle$ of s to an uncoerced value U , the term $U\langle s \rangle\langle t \rangle$ would allow two interpretations: an application of t to a value $U\langle s \rangle$ and applications of s and t to U and committing to either interpretation would break this property.

5 IMPLEMENTATION AND PRELIMINARY EVALUATION

We have implemented the coercion-passing translation described in Section 4 and the semantics of λS_1 for Grift [Kuhlenschmidt et al. 2019], an experimental compiler for gradually typed languages. GTLC+, the language that the Grift compiler implements, supports integers, floating-point numbers, Booleans, higher-order functions, local binding by `let`, (mutually) recursive definitions by `letrec`, conditional expressions, iterations, sequencing, and mutable references. The compiler supports different run-time check schemes, those based on type-based casts [Siek and Taha 2006] and space-efficient coercions [Siek et al. 2015]. Note that, although space-efficient coercions are supported, only nested coercions on *values* are composed; in other words, (R-MERGEC) is not implemented. Thus, implicit run-time checks may break tail calls and seemingly tail-recursive functions may cause stack overflow.

We modify the compiler phases for run-time checking based on the space-efficient coercions. After typechecking a user program, the compiler inserts type-based casts to the program and converts type-based casts to space-efficient coercions, following the translation from blame calculus λB to λS [Siek et al. 2015]. Our implementation performs the coercion-passing translation after the translation into λS . It is straightforward to extend the translation scheme to language features that are not present in λS . For example, here is translation for conditional expressions:

$$\mathcal{K}\llbracket \text{if } M \text{ then } N_1 \text{ else } N_2 \rrbracket K = \text{if } (\mathcal{K}\llbracket M \rrbracket \text{id}) \text{ then } (\mathcal{K}\llbracket N_1 \rrbracket K) \text{ else } (\mathcal{K}\llbracket N_2 \rrbracket K).$$

Intermediate languages used in Grift already supports first-class coercions, making it straightforward to implement our approach. We modify another compiler phase that generates operations on coercions such as $M \;; N$ and (R-WRAP). The current implementation, which generates C code and uses `clang`⁶ for compilation to machine code, relies on the C compiler to perform tail-call optimizations. We have found the original compiler's handling of recursive types⁷ hampers tail-call optimizations, so our implementation does not deal with recursive types. We leave their implementation for future work.

We have conducted a preliminary experiment to measure the overhead of the coercion-passing translation. The benchmark programs we have used are the tail-recursive even-odd functions in direct style:

⁶<https://clang.llvm.org/>

⁷The C function to compose coercions takes a pointer to a stack-allocated object as an argument and writes into the object when recursive coercions are composed.

```
(letrec ([even (lambda ([n : A1]) : A3
  (if (= 0 n) #f
      (odd (- n 1)))))]
 [odd (lambda ([n : A2]) : A4
  (if (= 0 n) #t
      (even (- n 1)))))]
  (odd n))
```

and its CPS-transformed version:

```
(letrec ([evenk (lambda ([n : A1] [k : (A3 -> A3)])) : A3
  (if (= n 0) (k #t)
      (oddk (- n 1) k)))]
 [oddk (lambda ([n : A2] [k : (A4 -> A4)])) : A4
  (if (= n 0) (k #f)
      (evenk (- n 1) k)))]
  (oddk n (lambda ([v : Bool]) : Bool v)))
```

We run the former with the original and modified compilers and the latter with the original for all combinations of A_1 and A_2 , which are either `Int` or `Dyn`, and A_3 and A_4 , which are either `Bool` or `Dyn`. They result in 48 different configurations. We call the direct-style program compiled by the original compiler `Base`, the direct-style program compiled by the modified compiler `CrcPS`, the CPS program compiled by the original compiler `CPS`.

First, we have confirmed that, as n increases, 12 of 16 configurations of `Base` cause stack overflow.⁸ In the four configurations that survived, both A_3 and A_4 are set to `Bool`. `CrcPS` never causes stack overflow for any configuration. For `CPS`, 8 of 16 configurations cause stack overflow. In all the crashed configurations A_3 and A_4 are different.

To our surprise, `Base` causes stack overflow even when $A_3 = A_4 = \text{Dyn}$ and `CPS` crashes for some configurations. We have found that it is due to the typing rule of `Grift` for conditional expressions. In `Grift`, if one of the branches is given a static type, say `Bool`, and the other is `Dyn`, the whole `if`-expression is given the static type and the compiler inserts a cast from `Dyn` to the branch of type `Dyn`. In the direct-style program where both A_3 and A_4 are `Dyn`, the two then-branches are always Boolean constants and the recursive calls in the two else-branches involve casts from `Dyn` to `Bool`, hence the insertion of projections `bool?P`. However, since the return types are declared to be `Dyn`, the whole `if`-expressions are cast back to `Dyn`, inserting injections `bool!`. Thus, every recursive call involves a projection immediately followed by an injection, as shown below, eventually causing stack overflow.

```
(letrec ([even (lambda ([n : Dyn]) : Dyn
  (if (= 0 n⟨int?P1⟩) #f
      (odd (- n⟨int?P2⟩ 1)⟨bool?P3⟩)⟨bool!⟩))]
 [odd (lambda ([n : Dyn] : Dyn) : Dyn
  (if (= 0 n⟨int?P4⟩) #t
      (even (- n⟨int?P5⟩ 1)⟨bool?P6⟩)⟨bool!⟩))]
  (odd n))
```

Similarly, a projection and an injection are inserted implicitly to the CPS program with $A_3 \neq A_4$, causing stack overflow.

⁸The size of the run-time stack is 32MB.

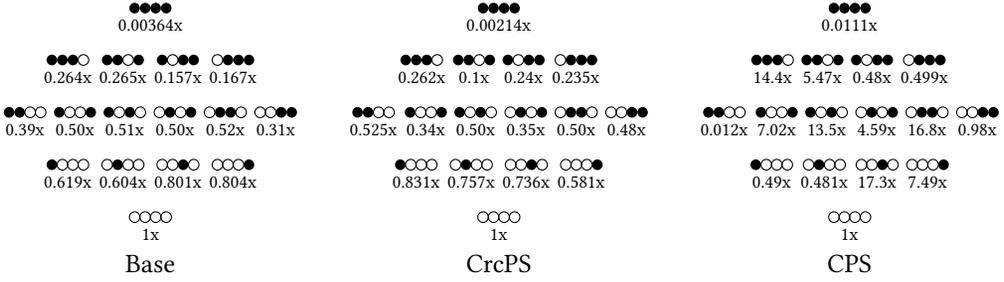


Fig. 10. Performance lattices. The four black/white circles represent whether each A_i is static (represented by a black circle) or dynamic (by a white circle).

6.1 Space-Efficient Coercion/Cast Calculi

As we have already mentioned, it is fairly well known that coercions [Henglein 1994] and casts [Wadler and Findler 2009] hamper the tail-call optimization and make the space complexity of the execution of a program worse than the execution under an unchecked semantics. We discuss below a few pieces of work [Garcia 2013; Herman et al. 2007, 2010; Siek et al. 2009, 2015; Siek and Wadler 2010] addressing the problem.

To the best of our knowledge, Herman et al. [2007, 2010] were the first to observe the space-efficiency problem of inserted dynamic checks. They developed a variant of Henglein’s coercion calculus with semantics such that a sequence of coercion applications is eagerly composed to reduce the size of coercions. However, their coercion composition operator is defined to be associative, equating $(c_1; c_2); c_3$ and $c_1; (c_2; c_3)$; thus, an algorithm for computing coercion composition was not very clear. They did not take blame tracking [Findler and Felleisen 2002] into account, either.

Later, Siek et al. [2009] extended Herman et al. [2007, 2010] with a few different blame tracking strategies. The issue of identifying $(c_1; c_2); c_3$ and $c_1; (c_2; c_3)$ remained. According to their terminology, our work, which follows previous work [Siek et al. 2015], adopts the UD semantics, which allows only $\star \rightarrow \star$ as a tag to functional values, as opposed to the D semantics, which allows any function types to be used as a tag.

Siek and Wadler [2010] introduced threesomes to a blame calculus. Threesome casts have a third type (called a mediating type) in addition to the source and target types; a threesome cast is considered a downcast from the source to the mediating type, followed by an upcast from the mediating to the target. Threesome casts allow a simple recursive algorithm to compose two threesome casts but blame tracking is rather complicated.

Garcia [2013] gave a translation from the threesome calculus to a coercion calculus and the two solutions are equivalent. They introduced supercoercions and a recursive algorithm to compute composition of supercoercions but they were complex, too.

Siek et al. [2015] proposed yet another space-efficient coercion calculus λS , in which they succeeded in developing a simple recursive algorithm for coercion composition by restricting coercions to be in certain canonical forms—what they call space-efficient coercions. They also gave a translation from blame calculus λB to λS (via Henglein’s coercion calculus λC) and showed that the translation is fully abstract. As we have discussed already, our λS has introduced syntax that distinguishes an application $U\langle s \rangle$ of a coercion to (uncoerced) values from $U\langle\langle d \rangle\rangle$ for a value wrapped by a delayed coercion. Such distinction, which can be seen some blame calculi [Wadler and Findler 2009], is not just an aesthetic choice but crucial for proving correctness of the translation.

All the above-mentioned calculi adopt a nonstandard reduction rule to compose coercions or casts even before the subject evaluates to a value, together with a nonstandard form of evaluation contexts and as a result it has not been clear how to implement them efficiently. [Herman et al. \[2007, 2010\]](#) sketched a few possible implementation strategies, including coercion passing but details were not discussed. [Siek and Garcia \[2012\]](#) showed an interpreter which performs coercion composition at tail calls. Although not showing correctness of the interpreter, their interpreter would give a hint to direct low-level implementation of space-efficient coercions. Our work addresses the problem of the nonstandard semantics in a different way—by translating a program into coercion-passing style. The difference, however, may not be so large as it may appear at first: in [Siek and Garcia \[2012\]](#), a state of the abstract machine includes an evaluation context, which contains the information on a coercion to be applied to a return value and such a coercion roughly corresponds to our continuation coercions. More detailed analysis of the relationship between the two implementation schemes is left for future work.

[Kuhlenschmidt et al. \[2019\]](#) built an experimental compiler Grift for gradual typing with structural types. It supports run-time checking with the space-efficient coercions of λS but does not support composition of coercions at tail positions. We have implemented our coercion-passing translation for the Grift compiler.

[Greenberg \[2015\]](#) has studied the same space-efficiency problem in the context of manifest contract calculi [[Greenberg et al. 2010, 2012](#); [Knowles and Flanagan 2010](#)] and proposed a few semantics for composing casts that involve contract checking. [Feltey et al. \[2018\]](#) have recently implemented Greenberg’s eidetic contracts on top of Typed Racket [[Tobin-Hochstadt and Felleisen 2008](#)] but, similarly to [Kuhlenschmidt et al. \[2019\]](#), composition is limited on a sequence of contracts applied to values.

There are other recent work for making gradual typing efficient [[Bauman et al. 2017](#); [Muehlboeck and Tate 2017](#); [Rastogi et al. 2015](#); [Richards et al. 2017](#)] but as far as we know, none of them addresses the problem caused by run-time checking applied to tail positions. Additionally, [Castagna et al. \[2019\]](#) implemented a virtual machine for space-efficient gradual typing in presence of set-theoretic types, but without blame tracking. They address the problem caused by casts applied to tail positions by an approach similar to the one in the interpreter by [Siek and Garcia \[2012\]](#). They implemented their virtual machine and evaluated their implementation by benchmarks such as the even/odd program.

6.2 Continuation-Passing Style

Obviously, our coercion-passing style translation is inspired by continuation-passing style translation, first formalized by [Plotkin \[1975\]](#). However, coercions represent only a part of the rest of computation and are, in this sense, closer to delimited continuations [[Danvy and Filinski 1990](#)]. Roughly speaking, translating a subexpression with `id` corresponds to the reset operation [[Danvy and Filinski 1990](#)] to delimit continuations. Unlike (delimited) continuations, which are usually expressed by first-class functions, coercions have compact representations and compactness can be preserved by composition.

[Wallach and Felten \[1998\]](#) proposed security-passing style to implement Java stack inspection [[Lindholm and Yellin 1999](#)]. The idea is indeed similar to ours: each function is augmented by an additional argument to pass information on run-time security checking.

In CPS, it is crucial to eliminate administrative redexes to achieve a simulation property [[Appel 1992](#); [Danvy and Filinski 1992](#); [Danvy and Nielsen 2003](#); [Plotkin 1975](#); [Sabry and Felleisen 1993](#); [Sabry and Wadler 1997](#); [Wand 1991](#)], which says that a reduction in the source is simulated by a sequence of (one-directional) reductions in the translation. Simulation is usually achieved by applying different translations to an application MN , depending whether M and N are values or

not. In addition to such value/nonvalue distinction, our coercion-passing translation also relies on whether a given continuation coercion is id or not when a coercion application is translated.

Continuation-passing style eliminates the difference between call-by-name and call-by-value but our coercion-passing style translation works only under the call-by-value semantics of the target language because coercions have to be eagerly composed. It would be interesting to investigate call-by-name for either the source and/or the target language.

7 CONCLUSION

We have developed a new coercion calculus λS_1 with first-class coercions as a target language of coercion-passing style translation from λS , an existing space-efficient coercion calculus. We have proved the translation preserves both typing and semantics. To achieve a simulation property, it is important to reduce administrative coercions, just as in CPS transformations. Our coercion-passing style translation solves the difficulty in implementing the semantics of λS in a faithful manner and, with the help of first-class coercions, makes it possible to implement in a compiler for a call-by-value language. We have modified an existing compiler for a gradually typed language and conducted a preliminary experiment. We have confirmed that our implementation successfully overcomes memory overflow caused by coercions at tail positions, which Kuhlenschmidt et al. [2019] did not support. Although the overhead is not small, causing slowdown of 2.61 times at maximum, the overall performance is stable under different type annotations and tends to be better than manually CPS translated programs.

Aside from completing the implementation by adding recursive types, which the original Grift compiler supports, more efficient implementation is an obvious direction of future work. Our coercion-passing style translation introduces a lot of identity coercions and optimizing operations on coercions will be necessary.

From a theoretical point of view, it would be interesting to extend the technique to gradual typing in the presence of parametric polymorphism [Ahmed et al. 2011, 2017; Igarashi et al. 2017; Toro et al. 2019; Xie et al. 2018], for which a polymorphic coercion calculus has to be studied first—Kießling and Luo [2003]; Luo [2008], who study coercive subtyping in polymorphic settings, may be relevant. The present design of λS_1 is geared towards coercion-passing style. For example, in λS_1 , trivial (namely identity) coercions for coercion types $A \rightsquigarrow B$ are allowed; passing coercions to dynamically typed code is prohibited; variables cannot appear as an argument to coercion constructors, like $x \Rightarrow s$. It may be interesting to study more general first-class coercions without such restrictions.

ACKNOWLEDGMENTS

This work was partially supported by JSPS KAKENHI Grant Number JP17H01723. We would like to thank John Toman for proofreading.

REFERENCES

- Amal Ahmed, Robert Bruce Findler, Jeremy G. Siek, and Philip Wadler. 2011. Blame for all. In *Proceedings of the 38th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2011, Austin, TX, USA, January 26-28, 2011*. 201–214. <https://doi.org/10.1145/1926385.1926409>
- Amal Ahmed, Dustin Jamner, Jeremy G. Siek, and Philip Wadler. 2017. Theorems for free for free: parametricity, with and without types. *PACMPL* 1, ICFP (2017), 39:1–39:28. <https://doi.org/10.1145/3110283>
- Andrew W. Appel. 1992. *Compiling with Continuations*. Cambridge University Press.
- Spenser Bauman, Carl Friedrich Bolz-Tereick, Jeremy G. Siek, and Sam Tobin-Hochstadt. 2017. Sound gradual typing: only mostly dead. *PACMPL* 1, OOPSLA (2017), 54:1–54:24. <https://doi.org/10.1145/3133878>
- Giuseppe Castagna, Guillaume Duboc, Victor Lanvin, and Jeremy G. Siek. 2019. A space-efficient call-by-value virtual machine for gradual set-theoretic types. In *Proceedings of the 31st Symposium on Implementation and Application of Functional Languages, IFL 2019, Singapore, September 25-27, 2019*.

- Julien Cretin and Didier Rémy. 2012. On the power of coercion abstraction. In *Proceedings of the 39th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2012, Philadelphia, Pennsylvania, USA, January 22-28, 2012*. 361–372. <https://doi.org/10.1145/2103656.2103699>
- Olivier Danvy and Andrzej Filinski. 1990. Abstracting Control. In *LISP and Functional Programming*. 151–160. <https://doi.org/10.1145/91556.91622>
- Olivier Danvy and Andrzej Filinski. 1992. Representing Control: A Study of the CPS Transformation. *Mathematical Structures in Computer Science* 2, 4 (1992), 361–391. <https://doi.org/10.1017/S0960129500001535>
- Olivier Danvy and Lasse R. Nielsen. 2001. Syntactic Theories in Practice. *Electr. Notes Theor. Comput. Sci.* 59, 4 (2001), 358–374. [https://doi.org/10.1016/S1571-0661\(04\)00297-X](https://doi.org/10.1016/S1571-0661(04)00297-X)
- Olivier Danvy and Lasse R. Nielsen. 2003. A first-order one-pass CPS transformation. *Theor. Comput. Sci.* 308, 1-3 (2003), 239–257. [https://doi.org/10.1016/S0304-3975\(02\)00733-8](https://doi.org/10.1016/S0304-3975(02)00733-8)
- Matthias Felleisen, Mitchell Wand, Daniel P. Friedman, and Bruce F. Duba. 1988. Abstract Continuations: A Mathematical Semantics for Handling Full Jumps. In *LISP and Functional Programming*. 52–62. <https://doi.org/10.1145/62678.62684>
- Daniel Feltey, Ben Greenman, Christophe Scholliers, Robert Bruce Findler, and Vincent St-Amour. 2018. Collapsible contracts: fixing a pathology of gradual typing. *PACMPL* 2, OOPSLA (2018), 133:1–133:27. <https://doi.org/10.1145/3276503>
- Robert Bruce Findler and Matthias Felleisen. 2002. Contracts for higher-order functions. In *Proceedings of the Seventh ACM SIGPLAN International Conference on Functional Programming (ICFP '02), Pittsburgh, Pennsylvania, USA, October 4-6, 2002*. 48–59. <https://doi.org/10.1145/581478.581484>
- Ronald Garcia. 2013. Calculating threesomes, with blame. In *ACM SIGPLAN International Conference on Functional Programming, ICFP'13, Boston, MA, USA - September 25 - 27, 2013*. 417–428. <https://doi.org/10.1145/2500365.2500603>
- Michael Greenberg. 2015. Space-Efficient Manifest Contracts. In *Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2015, Mumbai, India, January 15-17, 2015*. 181–194. <https://doi.org/10.1145/2676726.2676967>
- Michael Greenberg, Benjamin C. Pierce, and Stephanie Weirich. 2010. Contracts made manifest. In *Proceedings of the 37th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2010, Madrid, Spain, January 17-23, 2010*. 353–364. <https://doi.org/10.1145/1706299.1706341>
- Michael Greenberg, Benjamin C. Pierce, and Stephanie Weirich. 2012. Contracts made manifest. *J. Funct. Program.* 22, 3 (2012), 225–274. <https://doi.org/10.1017/S0956796812000135>
- Fritz Henglein. 1994. Dynamic Typing: Syntax and Proof Theory. *Sci. Comput. Program.* 22, 3 (1994), 197–230. [https://doi.org/10.1016/0167-6423\(94\)00004-2](https://doi.org/10.1016/0167-6423(94)00004-2)
- David Herman, Aaron Tomb, and Cormac Flanagan. 2007. Space-Efficient Gradual Typing. In *Proceedings of the Eighth Symposium on Trends in Functional Programming, TFP 2007, New York City, New York, USA, April 2-4, 2007*. 1–18.
- David Herman, Aaron Tomb, and Cormac Flanagan. 2010. Space-efficient gradual typing. *Higher-Order and Symbolic Computation* 23, 2 (2010), 167–189. <https://doi.org/10.1007/s10990-011-9066-z>
- Yuu Igarashi, Taro Sekiyama, and Atsushi Igarashi. 2017. On polymorphic gradual typing. *PACMPL* 1, ICFP (2017), 40:1–40:29. <https://doi.org/10.1145/3110284>
- Robert Kießling and Zhaohui Luo. 2003. Coercions in Hindley-Milner Systems. In *Types for Proofs and Programs, International Workshop, TYPES 2003, Torino, Italy, April 30 - May 4, 2003, Revised Selected Papers*. 259–275. https://doi.org/10.1007/978-3-540-24849-1_17
- Kenneth Knowles and Cormac Flanagan. 2010. Hybrid type checking. *ACM Trans. Program. Lang. Syst.* 32, 2 (2010), 6:1–6:34. <https://doi.org/10.1145/1667048.1667051>
- Andre Kuhlenschmidt, Deyaaeldeen Almahallawi, and Jeremy G. Siek. 2019. Toward Efficient Gradual Typing for Structural Types via Coercions. In *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI 2019)*. ACM, New York, NY, USA, 517–532. <https://doi.org/10.1145/3314221.3314627>
- Tim Lindholm and Frank Yellin. 1999. *The Java Virtual Machine Specification* (2nd ed.). Addison-Wesley.
- Zhaohui Luo. 2008. Coercions in a polymorphic type system. *Mathematical Structures in Computer Science* 18, 4 (2008), 729–751. <https://doi.org/10.1017/S0960129508006804>
- Fabian Muehlboeck and Ross Tate. 2017. Sound gradual typing is nominally alive and well. *PACMPL* 1, OOPSLA (2017), 56:1–56:30. <https://doi.org/10.1145/3133880>
- Gordon D. Plotkin. 1975. Call-by-Name, Call-by-Value and the λ -Calculus. *Theor. Comput. Sci.* 1, 2 (1975), 125–159. [https://doi.org/10.1016/0304-3975\(75\)90017-1](https://doi.org/10.1016/0304-3975(75)90017-1)
- Aseem Rastogi, Nikhil Swamy, Cédric Fournet, Gavin M. Bierman, and Panagiotis Vekris. 2015. Safe & Efficient Gradual Typing for TypeScript. In *Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2015, Mumbai, India, January 15-17, 2015*. 167–180. <https://doi.org/10.1145/2676726.2676971>
- Gregor Richards, Ellen Arteca, and Alexi Turcotte. 2017. The VM already knew that: leveraging compile-time knowledge to optimize gradual typing. *PACMPL* 1, OOPSLA (2017), 55:1–55:27. <https://doi.org/10.1145/3133879>

- Amr Sabry and Matthias Felleisen. 1993. Reasoning about Programs in Continuation-Passing Style. *Lisp and Symbolic Computation* 6, 3-4 (1993), 289–360.
- Amr Sabry and Philip Wadler. 1997. A Reflection on Call-by-Value. *ACM Trans. Program. Lang. Syst.* 19, 6 (1997), 916–941. <https://doi.org/10.1145/267959.269968>
- Jeremy G. Siek and Ronald Garcia. 2012. Interpretations of the gradually-typed lambda calculus. In *Proceedings of the 2012 Annual Workshop on Scheme and Functional Programming, Scheme 2012, Copenhagen, Denmark, September 9-15, 2012*. 68–80. <https://doi.org/10.1145/2661103.2661112>
- Jeremy G. Siek, Ronald Garcia, and Walid Taha. 2009. Exploring the Design Space of Higher-Order Casts. In *Programming Languages and Systems, 18th European Symposium on Programming, ESOP 2009, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2009, York, UK, March 22-29, 2009. Proceedings*. 17–31. https://doi.org/10.1007/978-3-642-00590-9_2
- Jeremy G. Siek and Walid Taha. 2006. Gradual typing for functional languages. In *Scheme and Functional Programming Workshop*. 81–92.
- Jeremy G. Siek, Peter Thiemann, and Philip Wadler. 2015. Blame and coercion: together again for the first time. In *Proceedings of the 36th ACM SIGPLAN Conference on Programming Language Design and Implementation, Portland, OR, USA, June 15-17, 2015*. 425–435. <https://doi.org/10.1145/2737924.2737968>
- Jeremy G. Siek and Philip Wadler. 2010. Threesomes, with and without blame. In *Proceedings of the 37th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2010, Madrid, Spain, January 17-23, 2010*. 365–376. <https://doi.org/10.1145/1706299.1706342>
- Asumu Takikawa, Daniel Feltey, Ben Greenman, Max S. New, Jan Vitek, and Matthias Felleisen. 2016. Is sound gradual typing dead?. In *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2016, St. Petersburg, FL, USA, January 20 - 22, 2016*. 456–468. <https://doi.org/10.1145/2837614.2837630>
- Sam Tobin-Hochstadt and Matthias Felleisen. 2006. Interlanguage migration: from scripts to programs. In *Proc. of Dynamic Languages Symposium*. 964–974. <https://doi.org/10.1145/1176617.1176755>
- Sam Tobin-Hochstadt and Matthias Felleisen. 2008. The design and implementation of typed scheme. In *Proceedings of the 35th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2008, San Francisco, California, USA, January 7-12, 2008*. 395–406. <https://doi.org/10.1145/1328438.1328486>
- Matías Toro, Elizabeth Labrada, and Éric Tanter. 2019. Gradual parametricity, revisited. *PACMPL* 3, POPL (2019), 17:1–17:30. <https://doi.org/10.1145/3290330>
- Philip Wadler and Robert Bruce Findler. 2009. Well-Typed Programs Can’t Be Blamed. In *Programming Languages and Systems, 18th European Symposium on Programming, ESOP 2009, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2009, York, UK, March 22-29, 2009. Proceedings*. 1–16. https://doi.org/10.1007/978-3-642-00590-9_1
- Dan S. Wallach and Edward W. Felten. 1998. Understanding Java Stack Inspection. In *Security and Privacy - 1998 IEEE Symposium on Security and Privacy, Oakland, CA, USA, May 3-6, 1998, Proceedings*. 52–63. <https://doi.org/10.1109/SECPRI.1998.674823>
- Mitchell Wand. 1991. Correctness of Procedure Representations in Higher-Order Assembly Language. In *Mathematical Foundations of Programming Semantics, 7th International Conference, Pittsburgh, PA, USA, March 25-28, 1991, Proceedings*. 294–311. https://doi.org/10.1007/3-540-55511-0_15
- Andrew K. Wright and Matthias Felleisen. 1994. A Syntactic Approach to Type Soundness. *Information and Computation* 115, 1 (Nov. 1994), 38–94.
- Ningning Xie, Xuan Bi, and Bruno C. d. S. Oliveira. 2018. Consistent Subtyping for All. In *Programming Languages and Systems - 27th European Symposium on Programming, ESOP 2018, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2018, Thessaloniki, Greece, April 14-20, 2018, Proceedings*. 3–30. https://doi.org/10.1007/978-3-319-89884-1_1