

## Space-efficient gradual typing

David Herman · Aaron Tomb · Cormac Flanagan

Published online: 21 October 2011  
© Springer Science+Business Media, LLC 2011

**Abstract** Gradual type systems offer a smooth continuum between static and dynamic typing by permitting the free mixture of typed and untyped code. The runtime systems for these languages, and other languages with hybrid type checking, typically enforce function types by dynamically generating function proxies. This approach can result in unbounded growth in the number of proxies, however, which drastically impacts space efficiency and destroys tail recursion.

We present a semantics for gradual typing that is based on *coercions* instead of function proxies, and which combines adjacent coercions at runtime to limit their space consumption. We prove bounds on the space consumed by coercions as well as soundness of the type system, demonstrating that programmers can safely mix typing disciplines without incurring unreasonable overheads. Our approach also detects certain errors earlier than prior work.

**Keywords** Gradual typing · Coercions · Casts · Type dynamic

---

David Herman was supported by a grant from the Mozilla Corporation. Cormac Flanagan was supported by a Sloan Fellowship and by NSF grant CCS-0341179. Aaron Tomb was also supported by NSF grant CCS-0341179. For Herman, most of this work was done at Northeastern University, Boston, MA.

D. Herman (✉)  
Mozilla Research, 650 Castro St, Mountain View, CA 94116, USA  
e-mail: [dherman@mozilla.com](mailto:dherman@mozilla.com)

A. Tomb · C. Flanagan  
Computer Science Department, School of Engineering, University of California at Santa Cruz, 1156 High Street, Santa Cruz, CA 95064, USA

A. Tomb  
e-mail: [atomb@soe.ucsc.edu](mailto:atomb@soe.ucsc.edu)

C. Flanagan  
e-mail: [cormac@ucsc.edu](mailto:cormac@ucsc.edu)

## 1 Gradual typing for software evolution

Dynamically typed languages have always excelled at exploratory programming. Languages such as Lisp, Scheme, Smalltalk, and JavaScript support quick early prototyping and incremental development without the overhead of documenting (often-changing) structural invariants as types. For large applications, however, static type systems have proven invaluable. They are crucial for understanding and enforcing key program invariants and abstractions, and they catch many errors early in the development process.

Given these different strengths, it is not uncommon to encounter the following scenario: a programmer builds a prototype in a dynamically-typed scripting language, perhaps even in parallel to a separate, official software development process with an entire team. The team effort gets mired in process issues and over-engineering, and the prototype ends up in production use. Before long, this hastily conceived prototype grows into a full-fledged production system, but without the structure or guarantees provided by static types. The system becomes unwieldy and bugs are cropping up faster than the team can reproduce them. Ultimately, they decide to port the application to a statically typed language, requiring a complete rewrite of the entire system.

The cost of cross-language migration is huge and often insupportable. But the scenario above is avoidable. Several languages combine static and dynamic typing, among them Boo [11], Visual Basic.NET [26], Sage [20], Clean [5, 29], and Racket [34, 35]. This approach of *hybrid typing*, where types are enforced with a combination of static and dynamic checks, has begun to gain attention in the research community [4, 8, 20, 26, 32, 34, 35]. Recently, Siek and Taha [32, 33] coined the slogan *gradual typing* for this important application of hybrid typing: the ability to implement both partially-conceived prototypes and mature, production systems in the same programming language by gradually introducing type discipline. Gradual typing offers the possibility of continuous software evolution from prototype to product, thus avoiding the huge costs of language migration.

Our experience in the technical committee working on the JavaScript [13] language specification provides a more concrete example. JavaScript is a dynamically-typed, higher-order programming language and a key enabling technology for interactive web applications. JavaScript is used both for very simple scripts and for mature, large-scale applications. The enormous popularity of the language is due in no small part to its low barrier to entry; anyone can write a JavaScript program by copying and pasting code from one web page to another. Its dynamic type system and fail-soft runtime semantics allow programmers to produce useful programs with a minimum of effort. The increasing complexity of modern web applications, however, has motivated the desire for more robust mechanisms to support programming in the large. In past years, the technical committee considered the design of a hybrid type system for JavaScript, which would allow the mixing of dynamic and static typing disciplines, to support gradual typing [12]. Gradual typing is still a relatively new research area, however, and so the technical committee would prefer further experimental validation before considering incorporating gradual typing into JavaScript.

This paper extends our earlier conference paper [23] with additional context and perspective on the relationship of our contributions with prior and subsequent work on hybrid and gradual typing.

### 1.1 The cost of gradual typing

The interaction between typed and untyped code imposes certain requirements on the design of the runtime semantics of gradually-typed languages. Gradually-typed languages support

both statically-typed and dynamically-typed code, and include runtime checks (or type casts) at the boundaries between these two typing disciplines, to guarantee that dynamically-typed code cannot violate the invariants of statically-typed code.

To illustrate this idea, consider the following code fragment, which passes a dynamically-typed variable  $x$  into a variable  $y$  of type `Int`:

```
let x : Dyn = true in
let y : Int = x in
...
```

During compilation, the type checker inserts a dynamic type cast `(Int)` to enforce the type invariant on  $y$ , yielding the “compiled” code:

```
let x : Dyn = true in
let y : Int = (Int) x in
...
```

At runtime, this `(Int)` cast detects the attempted type violation, since  $x$  contains `true`:

```
let x : Dyn = true in let y : Int = (Int) x in ...
→ let y : Int = (Int) true in ...
→ Error: "failed cast"
```

Unfortunately, even these simple, first-order type checks can impose unexpected costs, as in the following example, where a programmer has added some type annotations to a previously untyped program:

$$\begin{aligned} even : Dyn \rightarrow Dyn &\stackrel{\text{def}}{=} \lambda n : Dyn. \text{ if } (n = 0) \text{ then true else } odd(n - 1) \\ odd : Int \rightarrow Bool &\stackrel{\text{def}}{=} \lambda n : Int. \text{ if } (n = 0) \text{ then false else } even(n - 1) \end{aligned}$$

This program seems innocuous, but suffers from a subtle space leak. Since  $even$  is dynamically typed, the result of each call to  $even$  must be implicitly cast to `Bool`. That is, after the type checker inserts the necessary casts, the compiled version  $odd_c$  of  $odd$  is:

$$odd_c : Int \rightarrow Bool \stackrel{\text{def}}{=} \lambda n : Int. \text{ if } (n = 0) \text{ then false else } (Bool) (even((Dyn)(n - 1)))$$

These implicit casts result in unbounded growth in the control stack and so destroy tail recursion.

Additional complications arise when first-class functions cross the boundaries between typing disciplines. In general, it is not possible to check if an untyped function satisfies a particular static type. A natural solution is to wrap the function in a proxy that, whenever it is applied, casts its argument and result values appropriately, ensuring that the function is only observed with its expected type. This proxy-based approach was formalized by Findler and Felleisen [15] and widely used in subsequent literature [16, 20, 25, 32], but it has serious consequences for space efficiency.

As a simple example, consider the following program in continuation-passing style, where both mutually recursive functions take a continuation argument  $k$ , but only one of

these arguments is annotated with a precise type:

$$\begin{aligned} evenk : \text{Int} \rightarrow \text{Bool} &\stackrel{\text{def}}{=} \lambda n:\text{Int}. \lambda k:(\text{Dyn} \rightarrow \text{Dyn}). \\ &\quad \text{if } (n = 0) \text{ then } (k \text{ true}) \text{ else } oddk (n - 1) k \\ oddk : \text{Int} \rightarrow \text{Bool} &\stackrel{\text{def}}{=} \lambda n:\text{Int}. \lambda k:(\text{Bool} \rightarrow \text{Bool}). \\ &\quad \text{if } (n = 0) \text{ then } (k \text{ false}) \text{ else } evenk (n - 1) k \end{aligned}$$

During compilation, the type checker inserts a dynamic type cast on the continuation argument  $k$ , yielding the compiled code:

$$\begin{aligned} evenk_c : \text{Int} \rightarrow \text{Bool} &\stackrel{\text{def}}{=} \lambda n:\text{Int}. \lambda k:(\text{Dyn} \rightarrow \text{Dyn}). \\ &\quad \text{if } (n = 0) \\ &\quad \text{then } (\langle \text{Bool} \rangle (k (\langle \text{Dyn} \rangle \text{ true}))) \\ &\quad \text{else } oddk_c (n - 1) (\langle \text{Bool} \rightarrow \text{Bool} \rangle k) \\ oddk_c : \text{Int} \rightarrow \text{Bool} &\stackrel{\text{def}}{=} \lambda n:\text{Int}. \lambda k:(\text{Bool} \rightarrow \text{Bool}). \\ &\quad \text{if } (n = 0) \\ &\quad \text{then } (k \text{ false}) \\ &\quad \text{else } evenk_c (n - 1) (\langle \text{Dyn} \rightarrow \text{Dyn} \rangle k) \end{aligned}$$

Applying these functions to a sufficiently large  $n$  reveals the impact on space consumption of performing repeated function casts to  $k$ :

$$\begin{aligned} &evenk_c \quad n \quad (\langle \text{Dyn} \rightarrow \text{Dyn} \rangle (\lambda x:\text{Bool}. x)) \\ \longrightarrow^* &evenk_c \quad n \quad (\lambda z_1:\text{Dyn}. \langle \text{Dyn} \rangle ((\lambda x:\text{Bool}. x) (\langle \text{Bool} \rangle z_1))) \\ \longrightarrow^* &oddk_c \quad (n - 1) \quad (\langle \text{Bool} \rightarrow \text{Bool} \rangle (\lambda z_1:\text{Dyn}. \langle \text{Dyn} \rangle ((\lambda x:\text{Bool}. x) (\langle \text{Bool} \rangle z_1)))) \\ \longrightarrow^* &oddk_c \quad (n - 1) \quad (\lambda z_2:\text{Bool}. \langle \text{Bool} \rangle \\ &\quad (\lambda z_1:\text{Dyn}. \langle \text{Dyn} \rangle ((\lambda x:\text{Bool}. x) (\langle \text{Bool} \rangle z_1)) \\ &\quad (\langle \text{Dyn} \rangle z_2))) \\ \longrightarrow^* &evenk_c \quad (n - 2) \quad (\langle \text{Dyn} \rightarrow \text{Dyn} \rangle \\ &\quad \lambda z_2:\text{Bool}. \langle \text{Bool} \rangle \\ &\quad (\lambda z_1:\text{Dyn}. \langle \text{Dyn} \rangle ((\lambda x:\text{Bool}. x) (\langle \text{Bool} \rangle z_1)) \\ &\quad (\langle \text{Dyn} \rangle z_2))) \\ \longrightarrow^* &evenk_c \quad (n - 2) \quad (\lambda z_3:\text{Dyn}. \langle \text{Dyn} \rangle \\ &\quad (\lambda z_2:\text{Bool}. \langle \text{Bool} \rangle \\ &\quad (\lambda z_1:\text{Dyn}. \langle \text{Dyn} \rangle ((\lambda x:\text{Bool}. x) (\langle \text{Bool} \rangle z_1)) \\ &\quad (\langle \text{Dyn} \rangle z_2)) \\ &\quad (\langle \text{Bool} \rangle z_3))) \\ \longrightarrow^* &\dots \end{aligned}$$

The recursive calls to  $oddk_c$  and  $evenk_c$  quietly cast the continuation argument  $k$  with higher-order casts  $\langle \text{Bool} \rightarrow \text{Bool} \rangle$  and  $\langle \text{Dyn} \rightarrow \text{Dyn} \rangle$ , respectively. This means that the function argument  $k$  is wrapped in an additional function proxy at each recursive call! Despite the fact that the function argument is only used at a total of two different types, its size grows in proportion to the depth of the recursion, due to these dynamically generated proxies.

These examples demonstrate that naïve gradual typing violates space safety of functional programming languages by incurring surprising and unnecessary changes in asymptotic space consumption of programs. The purpose of gradual typing is to support the seamless integration of typing disciplines, especially for the common use case of simple scripts

that mature into large applications. Subtle, significant, and pervasive costs in resource consumption threaten to limit the utility of gradual typing for its intended use.

Of course, if *even* and *odd* are in the same module or maintained by the same development team, a natural solution is to avoid mixing typing disciplines, and to consistently use static typing (or dynamic typing) for both functions. A more problematic situation, however, occurs at interfaces between different modules maintained by different teams, as for example between a GUI library that uses callbacks to a GUI application. In this situation, we might want to migrate the GUI library to be more statically typed, while still maintaining backward compatibility with pre-existing GUI applications.

In short, the flexibility promised by gradual typing can only be achieved if programmers are free to decide how precisely to type various parts of a program. This flexibility is lost if adding type annotations can accidentally trigger drastic changes in its memory usage. The goal of this paper is to combine the flexibility and ease-of-use of gradual type systems with predictable resource consumption.

## 1.2 Space-efficient gradual typing using coercions

We present a semantics for gradual typing that overcomes these problems. Our approach hinges on the simple insight that when proxies accumulate at runtime, they often contain redundant information. In the higher-order example above, the growing chain of function proxies contains only two distinct components,  $\text{Bool} \rightarrow \text{Bool}$  and  $\text{Dyn} \rightarrow \text{Dyn}$ , which could be merged to the simpler but equivalent  $\text{Bool} \rightarrow \text{Bool}$ .

Type casts behave like *error projections* [14], which are closed under composition. However, the syntax of casts does not always compose; for example, there is no cast  $c$  such that  $\langle c \rangle e = \langle \text{Int} \rangle \langle \text{Bool} \rangle e$ . Furthermore, projections are idempotent, which should allow us to eliminate duplicate casts. For example, the following equality holds in a semantic sense:

$$\langle \text{Bool} \rightarrow \text{Bool} \rangle \langle \text{Bool} \rightarrow \text{Bool} \rangle e = \langle \text{Bool} \rightarrow \text{Bool} \rangle e$$

But such transformations are inconvenient, if not impossible, with a representation of higher-order type casts as functions.

Our formalization instead leverages Henglein's *coercion calculus* [22], which provides a syntax for projections, called coercions, which are closed under a composition operator. This allows us to combine adjacent casts whenever they occur at runtime in order to eliminate redundant information and thus guarantee clear bounds on space consumption. The details of combining and simplifying coercions are presented in Sect. 6. By eagerly combining coercions, we can also detect certain errors immediately as soon as a function cast is applied; in contrast, prior approaches would not detect these errors until the casted function is invoked.

Our approach is applicable to many gradually- or hybrid-typed languages [16, 17, 19, 20] that use function proxies and hence are prone to space consumption problems. For clarity, we formalize our approach for the simply-typed  $\lambda$ -calculus with references. Of course, gradual typing is not restricted to simple type systems: the Sage language [20] incorporates gradual typing as part of a very expressive type system with polymorphic functions, type operators, first-class types, dependent types, general refinement types, etc. Concurrently, Siek and Taha [32] developed gradual typing for the simpler language  $\lambda^?$ , which we use as the basis for this presentation.

### 1.3 Early error detection

A second contribution of this work is that it permits early detection of errors, and hence better test coverage. To illustrate this idea, consider the following code fragment, which erroneously attempts to convert an ( $\text{Int} \rightarrow \text{Int}$ ) function to have type ( $\text{Bool} \rightarrow \text{Int}$ ):

```
let g : Dyn = ( $\lambda x : \text{Int}. x$ ) in
let h : ( $\text{Bool} \rightarrow \text{Int}$ ) = g in ...
```

Most prior strategies for gradual typing would not detect this error until  $h$  is applied to a boolean argument, causing the integer cast to fail.

In contrast, our coercion-based implementation allows us to detect this error as soon as  $h$  is defined, which detects gradual typing errors earlier and in more situations. We believe that the eager approach better fits the philosophy of gradual typing, which is to permit maximum flexibility while still detecting errors as early as possible.

### 1.4 Outline

The presentation of our results proceeds as follows. The following section provides some perspective on related work done prior to and following the original publication of this work. Section 3 reviews the syntax and type system of  $\lambda^?$ . Section 4 introduces the coercion algebra underlying our approach. Section 5 describes how we translate source programs into a target language with explicit coercions. Section 6 provides an operational semantics for that target language. Section 7 extends our approach to detect errors earlier and in more situations. Section 8 proves bounds on the space consumption. The last two sections discuss other items of related work and place our results into context.

## 2 Context and perspective

This paper is an extended version of an earlier conference paper [23] that provides additional context and perspective by relating our contributions with both prior and subsequent work in the general area of hybrid and gradual typing.

Figure 1 summarizes some of the central characteristics of a few selected recent papers in this area. These papers vary in regards to whether they support each of the following features:

- precise blame tracking;
- the dynamic type  $\text{Dyn}$ ;
- predicate types, such as  $\{x : \text{Int} \mid x \geq 0\}$ , which denotes non-negative integers;
- early detection of nonsensically-composed function casts at run-time, before the function is applied; and
- guarantees of space efficiency.

The final column in Fig. 1 describes the representation of function casts, either as: functions  $(\lambda x. \dots)$ ; as casts of the form  $\langle S \Leftarrow T \rangle$  converting from type  $T$  to type  $S$ ; as Henglein's coercions; or as *Threesomes* [30]. We discuss these papers and characterizations in more detail below.

All of these papers leverage and extend the general technique of *higher-order casts* formalized by Findler and Felleisen [15]. This technique dynamically enforces contracts on

Paper	Higher Order Casts	Blame	Dyn Type	Predic- icate Types	Early Error Checks	Space Efficient	Cast Represent.
<b>Contract Systems:</b>							
Findler and Felleisen [15]	X	X					Function
<b>Hybrid Type Systems:</b>							
Flanagan [16]	X			X			Function
Knowles and Flanagan [24]	X			X			Function
Gronski et al. [20]	X		X	X			Function
Greenberg et al. [19]	X	X		X			Cast
<b>Gradual Type Systems:</b>							
Siek and Taha [32, 33]	X		X				Cast
<b>Herman et al.</b> [23]	X		X		X	X	Coercion
Siek et al. [31]	X	X	X		X	X	Coercion
Siek and Wadler [30]	X	X	X		X	X	Threesome

**Fig. 1** A comparison of hybrid and gradual typing system

functions by creating wrapper function *proxies*, which check that the appropriate contracts hold on all function arguments and results, in a lazy manner. A key contribution of this work was a mechanism for tracking blame through these function wrappers and for assigning blame for any contract failure back to a particular contract in the source program. This technique of higher-order wrappers (both with and without blame) was subsequently applied in a number of domains.

*Hybrid type checking* uses higher-order casts to overcome decidability and expressiveness limitations of traditional static type systems. The initial presentation [16] developed these ideas for an extension of the simply-typed lambda calculus with *predicate types*, where a base type such as `Int` is refined by an arbitrary boolean expression, to yield a predicate type such as  $\{x : \text{Int} \mid x \geq 0\}$ . In combination with dependent function types, predicate types permit the specification of rather precise function interfaces, which facilitate modular software development. Ou et al. [28] earlier explored similar ideas for interoperation between precisely and coarsely typed code. Subsequent work added a variety of extensions, including a dynamic type `Dyn` [20], clarifications to the metatheory [24], and extensions to support blame tracking and other properties [19].

Concurrently, Siek and Taha used higher-order wrappers to permit convenient interoperation between statically- and dynamically-typed code, targeting both functional [32] and object-oriented [33] languages. The omission of predicate types resulted in a significantly simpler language semantics and metatheory.

This paper (and our earlier conference paper [23]) introduces and satisfies two additional desiderata. Most critical is space efficiency, without which gradually-typed programs can be prone to unexpected blow-ups in stack and heap space requirements. An additional desideratum is early detection of errors, which was lacking in several earlier papers. We leverage Henglein's coercions [22] to represent higher-order casts at run-time, which facilitates cast normalization for space efficiency.

Subsequently, Siek and Taha [31] successfully extended these ideas with blame tracking, while still preserving the eager error detection and space efficiency properties. In addition, their paper also clarified several other aspects of the design space for gradual typing systems.

Henglein's coercions provide a convenient run-time representation for higher-order casts, since sequences of coercions can be combined or normalized at runtime, which ensures the desired space efficiency guarantee. However, coercions are a rather low-level and indirect representation of the casts that originated in the source program. To avoid this indirect representation, Siek and Taha [30] introduced the notion of a *threesome* as a replacement for a normalized sequence of coercions. In particular, a threesome  $\langle T_3 \Leftarrow T_2 \Leftarrow T_1 \rangle$  is essentially a downcast from a type  $T_1$  to  $T_2$  followed by an upcast from  $T_2$  to  $T_3$ . Note that this threesome is not equivalent to a conventional type cast  $\langle T_3 \Leftarrow T_1 \rangle$  from  $T_1$  to  $T_3$ , since the intermediate downcast to  $T_2$  may fail on more argument values. This threesome representation corresponds to normalized coercions, and so can achieve similar space efficiency guarantees while still reasoning at the level of types and casts, rather than using lower-level coercions.

Figure 1 provides an overall comparison between the features supported by various hybrid and gradual typed systems. In general, gradual typing avoids the complex metatheory associated with predicate types—this simpler setting has allowed gradual type systems to innovate and explore in other directions, namely space efficiency and early error detection.

We suggest that a fruitful avenue for future research is to merge these two closely related research directions by developing a flexible and expressive type system that supports static, dynamic, and predicate types, with a space-efficient run-time that tracks blame and performs early error checks. We suspect that an appropriate strategy is to appropriately extend the threesome cast representation to support predicate types. Additional directions for future work are discussed in Sect. 10.

### 3 Gradually-typed lambda calculus

This section reviews the gradually-typed  $\lambda$ -calculus  $\lambda_{\rightarrow}^?$  that we use to formalize our ideas. This language is essentially the simply-typed  $\lambda$ -calculus extended with the type `Dyn` to represent dynamic types; it also includes mutable reference cells to demonstrate the gradual typing of assignments. We refer the interested reader to Siek and Taha's work [32] for motivations regarding the design of this language and its type system.

$$\begin{array}{lll} \text{Typing environments:} & E & ::= \emptyset \mid E, x : T \\ \text{Terms:} & e & ::= k \mid x \mid \lambda x : T. e \mid e e \mid \text{ref } e \mid !e \mid e := e \\ \text{Types:} & S, T & ::= B \mid T \rightarrow T \mid \text{Dyn} \mid \text{Ref } T \end{array}$$

Terms include the usual constants, variables, abstractions, and applications, as well as reference allocation, dereference, and assignment. Types include the dynamic type `Dyn`, function types  $T \rightarrow T$ , reference types `Ref T`, and some collection of ground or base types  $B$  (such as `Int` or `Float`).

The  $\lambda_{\rightarrow}^?$  type system is a little unusual in that it is based on an intransitive *consistency* relation  $S \sim T$  instead of the more conventional transitive subtyping relation  $S <: T$ . Any type is consistent with the type `Dyn`, from which it follows that, for example, `Bool`  $\sim$  `Dyn` and `Dyn`  $\sim$  `Int`. However, booleans cannot be used directly as integers, which is why the consistency relation is not transitively closed. We do not assume the consistency relation is symmetric, since our language allows coercions from integers to floats but not vice-versa.

The consistency relation is defined in Fig. 2. Rules [C-DYNL] and [C-DYNR] allow all coercions to and from type `Dyn`. The rule [C-FLOAT] serves as an example of asymmetry by

**Consistency rules**

$$S \sim T$$

$$\begin{array}{cccc} [\text{C-REFL}] & [\text{C-DYNR}] & [\text{C-DYNL}] & [\text{C-FLOAT}] \\ \hline \overline{T \sim T} & \overline{T \sim \text{Dyn}} & \overline{\text{Dyn} \sim T} & \overline{\text{Int} \sim \text{Float}} \\ [\text{C-FUN}] & & [\text{C-REF}] & \\ \frac{T_1 \sim S_1 \quad S_2 \sim T_2}{(S_1 \rightarrow S_2) \sim (T_1 \rightarrow T_2)} & & \frac{T \sim S \quad S \sim T}{\text{Ref } S \sim \text{Ref } T} & \end{array}$$

**Type rules**

$$E \vdash e : T$$

$$\begin{array}{ccc} [\text{T-VAR}] & [\text{T-FUN}] & [\text{T-CONST}] \\ \frac{(x : T) \in E}{E \vdash x : T} & \frac{E, x : S \vdash e : T}{E \vdash (\lambda x : S. e) : (S \rightarrow T)} & \frac{}{E \vdash k : \text{ty}(k)} \\ [\text{T-APP1}] & & [\text{T-APP2}] \\ \frac{E \vdash e_1 : (S \rightarrow T) \quad E \vdash e_2 : S' \quad S' \sim S}{E \vdash (e_1 e_2) : T} & & \frac{E \vdash e_1 : \text{Dyn} \quad E \vdash e_2 : S}{E \vdash (e_1 e_2) : \text{Dyn}} \\ [\text{T-REF}] & [\text{T-DEREFL}] & [\text{T-DEREF2}] \\ \frac{E \vdash e : T}{E \vdash \text{ref } e : \text{Ref } T} & \frac{E \vdash e : \text{Ref } T}{E \vdash !e : T} & \frac{E \vdash e : \text{Dyn}}{E \vdash !e : \text{Dyn}} \\ [\text{T-ASSIGN1}] & & [\text{T-ASSIGN2}] \\ \frac{E \vdash e_1 : \text{Ref } T \quad E \vdash e_2 : S \quad S \sim T}{E \vdash e_1 := e_2 : S} & & \frac{E \vdash e_1 : \text{Dyn} \quad E \vdash e_2 : T}{E \vdash e_1 := e_2 : \text{Dyn}} \end{array}$$

**Fig. 2** Source language type system

allowing coercion from `Int` to `Float` but not the reverse. The rule [C-FUN] is reminiscent of the contravariant/covariant rule for function subtyping. We extend the invariant reference cells of  $\lambda^?$  to allow coercion from `Ref S` to `Ref T` via rule [C-REF], provided `S` and `T` are symmetrically consistent. Unlike functions, reference cells do not distinguish their output (“read”) type from their input (“write”) type, so coercion must be possible in either direction. For example, the two reference types `Ref Int` and `Ref Dyn` are consistent.

Figure 2 also presents the type rules for the source language. Notice the presence of two separate procedure application rules. Rule [T-APP1] checks statically-typed functions; in this case, the argument may have any type consistent with the function’s domain. Rule [T-APP2] handles the case where the operator is dynamically-typed; in this case the argument may be of any type. The assignment rules follow an analogous pattern, accepting a consistent type when the left-hand side is known to have type `Ref T`, and any type when the left-hand side is dynamically-typed. Similarly, dereferences only produce precise types for reference types.

## 4 Coercions

To achieve a space-efficient implementation, we elaborate source programs into a target language with explicit type casts, which allow expressions of one type to be used at any consis-

tent type. Our representation of casts is based on *coercions*, drawn from Henglein's theory of dynamic typing [22]. The key benefit of coercions over prior proxy-based representations is that they are *combinable*; if two coercions are wrapped around a function value, then they can be safely combined into a single coercion, thus reducing the space consumption of the program without changing its semantic behavior.

The coercion language and its typing rules are both defined in Fig. 3. The coercion judgment  $\vdash c : S \rightsquigarrow T$  states that coercion  $c$  serves to coerce values from type  $S$  to type  $T$  (but may fail on some values of type  $S$  that are not coercible to type  $T$ ). The identity coercion  $I$  always succeeds, and has type  $\vdash c : T \rightsquigarrow T$  for any  $T$ . Conversely, the failure coercion  $\text{Fail}$  always fails, but can relate any two types  $S$  and  $T$ , and so  $\vdash \text{Fail} : S \rightsquigarrow T$ . For each base type  $B$ , the coercion  $B?$  “downcasts” a value of type  $\text{Dyn}$  to type  $B$ , and so has type  $\vdash B? : \text{Dyn} \rightsquigarrow B$ . Conversely, the tagging coercion  $B!$  has type  $\vdash B! : B \rightsquigarrow \text{Dyn}$ . Similarly, the function checking coercion  $\text{Fun}?$  converts a value of type  $\text{Dyn}$  to have the dynamic function type  $\text{Dyn} \rightarrow \text{Dyn}$ . If a more precise function type is required, this value can be further coerced via a function coercion  $\text{Fun } c \ d$ , where  $c$  coerces function arguments and  $d$  coerces results. For example, the coercion  $(\text{Fun } \text{Int}?) \ \text{Int}!$  coerces from  $\text{Dyn} \rightarrow \text{Dyn}$  to  $\text{Int} \rightarrow \text{Int}$ , by untagging function arguments (via  $\text{Int}?$ ) and tagging function results (via  $\text{Int}!$ ). Conversely, the tagging coercion  $\text{Fun}!$  coerces a dynamic function of type  $\text{Dyn} \rightarrow \text{Dyn}$  to type  $\text{Dyn}$ . Like function coercions, reference coercions also contain two components: the first for coercing values put into the reference cell; the second for coercing values read from the cell. Finally, the coercion  $c; d$  represents coercion composition, *i.e.*, the coercion  $c$  followed by coercion  $d$ .

This coercion language is sufficient to translate between all consistent types: if types  $S$  and  $T$  are consistent, then the following partial function  $\text{coerce}(S, T)$  is defined and returns the appropriate coercion between these types.

$$\begin{aligned}
 \text{coerce} &: \text{Type} \times \text{Type} \rightarrow \text{Coercion} \\
 \text{coerce}(T, T) &= I \\
 \text{coerce}(B, \text{Dyn}) &= B! \\
 \text{coerce}(\text{Dyn}, B) &= B? \\
 \text{coerce}(\text{Int}, \text{Float}) &= \text{IntFloat} \\
 \text{coerce}(S_1 \rightarrow S_2, T_1 \rightarrow T_2) &= \text{Fun } \text{coerce}(T_1, S_1) \text{ } \text{coerce}(S_2, T_2) \\
 \text{coerce}(\text{Dyn}, T_1 \rightarrow T_2) &= \text{Fun}?; \text{ } \text{coerce}(\text{Dyn} \rightarrow \text{Dyn}, T_1 \rightarrow T_2) \\
 \text{coerce}(T_1 \rightarrow T_2, \text{Dyn}) &= \text{coerce}(T_1 \rightarrow T_2, \text{Dyn} \rightarrow \text{Dyn}); \text{ } \text{Fun}! \\
 \text{coerce}(\text{Ref } S, \text{Ref } T) &= \text{Ref } \text{coerce}(T, S) \text{ } \text{coerce}(S, T) \\
 \text{coerce}(\text{Dyn}, \text{Ref } T) &= \text{Ref}?; \text{ } \text{coerce}(\text{Ref Dyn}, \text{Ref } T) \\
 \text{coerce}(\text{Ref } T, \text{Dyn}) &= \text{coerce}(\text{Ref } T, \text{Ref Dyn}); \text{ } \text{Ref}!
 \end{aligned}$$

Coercing a type  $T$  to itself produces the identity coercion  $I$ . Coercing base types  $B$  to type  $\text{Dyn}$  requires a tagging coercion  $B!$ , and coercing  $\text{Dyn}$  to a base type  $B$  requires a runtime check  $B?$ . Function coercions work by coercing their domain and range types. The type  $\text{Dyn}$  is coerced to a function type via a two-step coercion: first the value is checked to be a function and then coerced from the dynamic function type  $\text{Dyn} \rightarrow \text{Dyn}$  to  $T_1 \rightarrow T_2$ . Dually, typed functions are coerced to type  $\text{Dyn}$  via coercion to a dynamic function type followed by the function tag  $\text{Fun}!$ . Coercing a  $\text{Ref } S$  to a  $\text{Ref } T$  entails coercing all writes from  $T$  to  $S$  and all reads from  $S$  to  $T$ . Coercing reference types to and from  $\text{Dyn}$  is analogous to function coercion.

## Syntax

$$\begin{array}{ll} \text{Coercions: } & c, d ::= I \mid \text{Fail} \mid G! \mid G? \mid \text{IntFloat} \mid \text{Fun } c \ c \mid \text{Ref } c \ c \mid c; c \\ \text{Dynamic tags: } & G \quad ::= \ B \mid \text{Fun} \mid \text{Ref} \end{array}$$

## Coercion rules

$$\boxed{\vdash c : S \rightsquigarrow T}$$

$$\begin{array}{c} \frac{[\text{C-ID}]}{\vdash I : T \rightsquigarrow T} \quad \frac{[\text{C-FAIL}]}{\vdash \text{Fail} : S \rightsquigarrow T} \quad \frac{[\text{C-B!}]}{\vdash B! : B \rightsquigarrow \text{Dyn}} \quad \frac{[\text{C-B?}]}{\vdash B? : \text{Dyn} \rightsquigarrow B} \\ \frac{[\text{C-FUN!}]}{\vdash \text{Fun!} : (\text{Dyn} \rightarrow \text{Dyn}) \rightsquigarrow \text{Dyn}} \quad \frac{[\text{C-FUN?}]}{\vdash \text{Fun?} : \text{Dyn} \rightsquigarrow (\text{Dyn} \rightarrow \text{Dyn})} \\ \frac{[\text{C-FUN}]}{\vdash c_1 : T'_1 \rightsquigarrow T_1 \quad \vdash c_2 : T_2 \rightsquigarrow T'_2} \quad \frac{}{\vdash (\text{Fun } c_1 \ c_2) : (T_1 \rightarrow T_2) \rightsquigarrow (T'_1 \rightarrow T'_2)} \\ \frac{[\text{C-REF!}]}{\vdash \text{Ref!} : (\text{Ref Dyn}) \rightsquigarrow \text{Dyn}} \quad \frac{[\text{C-REF?}]}{\vdash \text{Ref?} : \text{Dyn} \rightsquigarrow (\text{Ref Dyn})} \\ \frac{[\text{C-REF}]}{\vdash c : T \rightsquigarrow S \quad \vdash d : S \rightsquigarrow T} \quad \frac{}{\vdash (\text{Ref } c \ d) : (\text{Ref } S) \rightsquigarrow (\text{Ref } T)} \\ \frac{[\text{C-COMPOSE}]}{\vdash c_1 : T \rightsquigarrow T_1 \quad \vdash c_2 : T_1 \rightsquigarrow T_2} \quad \frac{[\text{C-FLOAT}]}{\vdash \text{IntFloat} : \text{Int} \rightsquigarrow \text{Float}} \end{array}$$

**Fig. 3** Coercion language and type rules

## Lemma 1 (Well-typed coercions)

1.  $S \sim T$  iff  $\text{coerce}(S, T)$  is defined.
2. If  $c = \text{coerce}(S, T)$  then  $\vdash c : S \rightsquigarrow T$ .

*Proof* Inductions on the derivations of  $S \sim T$  and  $\text{coerce}(S, T)$ . □

*Example* Consider again the continuation-passing definition of *evenk* in Sect. 1.1, which passes its continuation argument  $k$  to *oddk*. Because *oddk* expects a continuation argument of type  $\text{Bool} \rightarrow \text{Bool}$  but  $k$  has type  $\text{Dyn} \rightarrow \text{Dyn}$ , the type checker must insert a cast around the use of  $k$ . Note that the types  $\text{Bool} \rightarrow \text{Bool}$  and  $\text{Dyn} \rightarrow \text{Dyn}$  are consistent, from which it follows by Lemma 1 that *coerce* is defined on these arguments:

$$\begin{aligned} & \text{coerce}(\text{Dyn} \rightarrow \text{Dyn}, \text{Bool} \rightarrow \text{Bool}) \\ &= \text{Fun } \text{coerce}(\text{Bool}, \text{Dyn}) \text{ coerce}(\text{Dyn}, \text{Bool}) \\ &= \text{Fun } \text{Bool! Bool?} \end{aligned}$$

## Syntax

<i>Terms:</i>	$s, t ::= x \mid u \mid t \ t \mid \text{ref } t \mid !t \mid t := t \mid \langle c \rangle t$
<i>Values:</i>	$v ::= u \mid \langle vc \rangle u$
<i>Uncoerced values:</i>	$u ::= \lambda x:T. t \mid k \mid a$
<i>Value coercions:</i>	$vc ::= G! \mid \text{Fun } c \ d \mid \text{Ref } c \ d \mid vc; c$
<i>Stores:</i>	$\sigma ::= \emptyset \mid \sigma[a := v]$
<i>Typing environments:</i>	$E ::= \emptyset \mid E, x : T$
<i>Store typings:</i>	$\Sigma ::= \emptyset \mid \Sigma, a : T$

## Type rules for terms

 $E; \Sigma \vdash t : T$ 

$$\begin{array}{c}
 \begin{array}{ccc}
 \text{[T-VAR]} & \text{[T-FUN]} & \text{[T-APP]} \\
 \frac{(x : t) \in E}{E; \Sigma \vdash x : T} & \frac{E, x : S; \Sigma \vdash t : T}{E; \Sigma \vdash (\lambda x : S. t) : (S \rightarrow T)} & \frac{E; \Sigma \vdash t_1 : (S \rightarrow T) \quad E; \Sigma \vdash t_2 : S}{E; \Sigma \vdash (t_1 t_2) : T}
 \end{array} \\
 \begin{array}{ccc}
 \text{[T-REF]} & \text{[T-DEREF]} & \text{[T-ASSIGN]} \\
 \frac{E; \Sigma \vdash t : T}{E; \Sigma \vdash \text{ref } t : \text{Ref } T} & \frac{E; \Sigma \vdash t : \text{Ref } T}{E; \Sigma \vdash !t : T} & \frac{E; \Sigma \vdash t_1 : \text{Ref } T \quad E; \Sigma \vdash t_2 : T}{E; \Sigma \vdash t_1 := t_2 : T}
 \end{array} \\
 \begin{array}{ccc}
 \text{[T-CONST]} & \text{[T-CAST]} & \text{[T-ADDR]} \\
 \frac{}{E; \Sigma \vdash k : \text{ty}(k)} & \frac{\vdash c : S \rightsquigarrow T \quad E; \Sigma \vdash t : S}{E; \Sigma \vdash \langle c \rangle t : T} & \frac{(a : T) \in \Sigma}{E; \Sigma \vdash a : \text{Ref } T}
 \end{array}
 \end{array}$$

## Type rules for stores

 $E; \Sigma \vdash \sigma$ 

$$\frac{\begin{array}{c} \text{[T-STORE]} \\ \text{dom}(\sigma) = \text{dom}(\Sigma) \\ \forall a \in \text{dom}(\sigma). E; \Sigma \vdash \sigma(a) : \Sigma(a) \end{array}}{E; \Sigma \vdash \sigma}$$

**Fig. 4** Target language syntax and type rules

Furthermore, the resulting coercion has the expected type according to the coercion rules of Fig. 3:

$$\frac{\begin{array}{c} \vdash \text{Bool!} : \text{Bool} \rightsquigarrow \text{Dyn} \quad \vdash \text{Bool?} : \text{Dyn} \rightsquigarrow \text{Bool} \\ \vdash (\text{Fun } \text{Bool! } \text{Bool?}) : (\text{Dyn} \rightarrow \text{Dyn}) \rightsquigarrow (\text{Bool} \rightarrow \text{Bool}) \end{array}}{\vdash (\text{Fun } \text{Bool! } \text{Bool?}) : (\text{Dyn} \rightarrow \text{Dyn}) \rightsquigarrow (\text{Bool} \rightarrow \text{Bool})}$$

## 5 Target language and cast insertion

During compilation, we simultaneously type-check the source program and insert explicit type casts to convert between consistent types where necessary. The target language of this cast insertion process is essentially the same as the source, except that it uses explicit casts of the form  $\langle c \rangle t$  as the only mechanism for connecting terms of type Dyn and terms of other types. For example, the term  $\langle \text{Int?} \rangle x$  has type Int, provided that  $x$  has type Dyn. The

**Cast insertion rules**

$$E \vdash e \hookrightarrow t : T$$

$\begin{array}{c} [\text{C-VAR}] \\ (x : T) \in E \\ \hline E \vdash x \hookrightarrow x : T \end{array}$	$\begin{array}{c} [\text{C-CONST}] \\ E \vdash k \hookrightarrow k : ty(k) \\ \hline E \vdash e_1 \hookrightarrow t_1 : (S \rightarrow T) \quad E \vdash e_2 \hookrightarrow t_2 : S' \quad c = coerce(S', S) \end{array}$	$\begin{array}{c} [\text{C-FUN}] \\ E, x : S \vdash e \hookrightarrow t : T \\ \hline E \vdash (\lambda x : S. e) \hookrightarrow (\lambda x : S. t) : (S \rightarrow T) \end{array}$
		$\begin{array}{c} [\text{C-APP1}] \\ E \vdash e_1 \hookrightarrow t_1 : (S \rightarrow T) \quad E \vdash e_2 \hookrightarrow t_2 : S' \quad c = coerce(S', S) \\ \hline E \vdash e_1 e_2 \hookrightarrow (t_1 ((c) t_2)) : T \end{array}$
		$\begin{array}{c} [\text{C-APP2}] \\ E \vdash e_1 \hookrightarrow t_1 : \text{Dyn} \quad E \vdash e_2 \hookrightarrow t_2 : S' \quad c = coerce(S', \text{Dyn}) \\ \hline E \vdash e_1 e_2 \hookrightarrow (((\text{Fun?}) t_1) ((c) t_2)) : \text{Dyn} \end{array}$
$\begin{array}{c} [\text{C-REF}] \\ E \vdash e \hookrightarrow t : T \\ \hline E \vdash \text{ref } e \hookrightarrow \text{ref } t : \text{Ref } T \end{array}$	$\begin{array}{c} [\text{C-DEREF1}] \\ E \vdash e \hookrightarrow t : \text{Ref } T \\ \hline E \vdash e \hookrightarrow !t : T \end{array}$	$\begin{array}{c} [\text{C-DEREF2}] \\ E \vdash e \hookrightarrow t : \text{Dyn} \\ \hline E \vdash !e \hookrightarrow !(\text{Ref?} t) : \text{Dyn} \end{array}$
		$\begin{array}{c} [\text{C-ASSIGN1}] \\ E \vdash e_1 \hookrightarrow t_1 : \text{Ref } S \quad E \vdash e_2 \hookrightarrow t_2 : T \quad c = coerce(T, S) \\ \hline E \vdash e_1 := e_2 \hookrightarrow (t_1 := ((c) t_2)) : S \end{array}$
		$\begin{array}{c} [\text{C-ASSIGN2}] \\ E \vdash e_1 \hookrightarrow t_1 : \text{Dyn} \quad E \vdash e_2 \hookrightarrow t_2 : T \quad c = coerce(T, \text{Dyn}) \\ \hline E \vdash e_1 := e_2 \hookrightarrow (((\text{Ref?}) t_1) := ((c) t_2)) : \text{Dyn} \end{array}$

**Fig. 5** Cast insertion rules

language syntax and type rules are defined in Fig. 4, and are mostly straightforward. The values include “uncoerced values”  $u$ , consisting of abstractions, constants, and references. The latter are represented as addresses  $a$  in the global store  $\sigma$ . A value may also be wrapped with a *value coercion*  $vc$ . Value coercions are the subset of coercions that begin with either a tag  $G!$  or a higher-order coercion  $\text{Fun } c d$  or  $\text{Ref } c d$ .

The process of type checking and inserting coercions is formalized via the *cast insertion judgment*:

$$E \vdash e \hookrightarrow t : T$$

Here, the type environment  $E$  provides types for free variables,  $e$  is the original source program,  $t$  is a modified version of the original program with additional coercions, and  $T$  is the inferred type for  $t$ . The rules defining the cast insertion judgment are shown in Fig. 5, and they rely on the partial function  $coerce$  to compute coercions between types. For example, rule [C-APP1] compiles an application expression where the operator has a function type  $S \rightarrow T$  by casting the argument expression to type  $S$ . Rule [C-APP2] handles the case where the operator is dynamically-typed by inserting a  $\langle \text{Fun?} \rangle$  check to ensure the operator evaluates to a function, and casting the argument to type Dyn to yield a tagged value. Rules [C-REF] and [C-DEREF1] handle typed reference allocation and dereference. Rule [C-DEREF2] handles dynamically-typed dereference by inserting a runtime  $\langle \text{Ref?} \rangle$  check. Rule [C-ASSIGN1] handles statically-typed assignment, casting the right-hand side to the expected type of the reference cell, and [C-ASSIGN2] handles dynamically-typed assignment, casting the left-hand side with a runtime  $\langle \text{Ref?} \rangle$  check and the right-hand side with a tagging coercion to type Dyn.

Returning to the continuation-passing *evenk* and *oddk* functions, the cast insertion judgment produces the following compiled program (which differs somewhat from the compiled version in the introduction because of our use of coercions to represent casts):

$$\begin{aligned} evenk_c : \text{Int} \rightarrow \text{Bool} &\stackrel{\text{def}}{=} \lambda n : \text{Int}. \lambda k : (\text{Dyn} \rightarrow \text{Dyn}). \\ &\quad \text{if } (n = 0) \\ &\quad \text{then } (\text{Bool}?)(k (\langle \text{Bool!} \rangle \text{ true})) \\ &\quad \text{else } oddk_c(n - 1) (\langle \text{Fun Bool! Bool?} \rangle k) \\ \\ oddk_c : \text{Int} \rightarrow \text{Bool} &\stackrel{\text{def}}{=} \lambda n : \text{Int}. \lambda k : (\text{Bool} \rightarrow \text{Bool}). \\ &\quad \text{if } (n = 0) \\ &\quad \text{then } (k \text{ false}) \\ &\quad \text{else } evenk_c(n - 1) (\langle \text{Fun Bool? Bool!} \rangle k) \end{aligned}$$

Compilation succeeds on all well-typed source programs, and produces only well-typed target programs.

**Theorem 1** (Well-typed programs compile) *For all E, e, and T, the following statements are equivalent:*

1.  $E \vdash e : T$
2.  $\exists t \text{ such that } E \vdash e \hookrightarrow t : T$

**Theorem 2** (Compiled programs are well-typed) *For all E, e, t, and T, if  $E \vdash e \hookrightarrow t : T$  then  $E; \emptyset \vdash t : T$*

*Proof* Inductions on the cast insertion and source language typing derivations. □

## 6 Operational semantics

We now consider how to implement the target language in a manner that limits the space consumed by coercions. The key idea is to combine adjacent casts at runtime to eliminate redundant information while preserving the behavior of programs.

Figure 6 provides the definitions and reduction rules for the target language, using a small-step operational semantics with evaluation contexts. The grammar of evaluation contexts C is defined in terms of contexts D, which cannot begin with a cast. This prevents nesting of adjacent casts and allows the crucial [E-CCAST] reduction rule to ensure that adjacent casts in the program term are always merged.

In order to maintain bounds on their size, coercions are maintained normalized throughout evaluation according to the following rules:

$$\begin{array}{lll} I; c = c & & G!; G? = I \\ c; I = c & & G!; G'? = \text{Fail} \quad \text{if } G \neq G' \\ \text{Fail}; c = \text{Fail} & ( \text{Fun } c_1 c_2 ); ( \text{Fun } d_1 d_2 ) = \text{Fun } (d_1; c_1) (c_2; d_2) \\ c; \text{Fail} = \text{Fail} & ( \text{Ref } c_1 c_2 ); ( \text{Ref } d_1 d_2 ) = \text{Ref } (d_1; c_1) (c_2; d_2) \\ & (c_1; c_2); c_3 = c_1; (c_2; c_3) \end{array}$$

This normalization is applied in a transitive, compatible manner whenever the rule [E-CCAST] is applied, thus bounding the size of coercions during evaluation. Normalized

## Syntax

$$\begin{array}{ll} \text{Evaluation Contexts:} & C ::= D \mid \langle c \rangle D \\ \text{Non-cast Contexts:} & D ::= \bullet \mid (C t) \mid (v C) \mid \text{ref } C \mid !C \mid C := t \mid v := C \end{array}$$

## Evaluation rules

$C[(\lambda x:S. t) v], \sigma \longrightarrow C[t[x := v]], \sigma$	[E-BETA]
$C[\text{ref } v], \sigma \longrightarrow C[a], \sigma[a := v] \quad \text{for } a \notin \text{dom}(\sigma)$	[E-NEW]
$C[!a], \sigma \longrightarrow C[\sigma(a)], \sigma \quad \text{if } a \in \text{dom}(\sigma)$	[E-DEREF]
$C[a := v], \sigma \longrightarrow C[v], \sigma[a := v] \quad \text{if } a \in \text{dom}(\sigma)$	[E-ASSIGN]
$C[k v], \sigma \longrightarrow C[\delta(k, v)], \sigma$	[E-PRIM]
$C[(\langle \text{Fun } c d \rangle u) v], \sigma \longrightarrow C[\langle d \rangle (u (\langle c \rangle v))], \sigma$	[E-CBETA]
$C[!(\langle \text{Ref } c d \rangle a)], \sigma \longrightarrow C[\langle d \rangle !a], \sigma$	[E-CDEREF]
$C[(\langle \text{Ref } c d \rangle a) := v], \sigma \longrightarrow C[\langle d \rangle (a := \langle c \rangle v)], \sigma$	[E-CASSIGN]
$C[(I) u], \sigma \longrightarrow C[u], \sigma \quad \text{if } C \neq C'[\langle c \rangle \bullet]$	[E-ID]
$C[(\text{IntFloat}) n], \sigma \longrightarrow C[\text{nearestFloat}(n)], \sigma \quad \text{if } C \neq C'[\langle c \rangle \bullet]$	[E-FCAST]
$C[\langle c \rangle (\langle d \rangle t)], \sigma \longrightarrow C[\langle nc \rangle t], \sigma \quad \text{if } d;c = nc$	[E-CCAST]

**Fig. 6** Operational semantics

coercions satisfy the following grammar (where  $[-]$  denotes an optional component):

$$\begin{aligned} nc ::= & I \mid \text{Fail} \mid G? \mid G! \mid G?; G! \\ & \mid [\text{Int}?;] \text{IntFloat} [;\text{Float}!] \\ & \mid [\text{Fun}?;] \text{Fun } nc \text{ } nc [;\text{Fun}!] \\ & \mid [\text{Ref}?;] \text{Ref } nc \text{ } nc [;\text{Ref}!] \end{aligned}$$

Most of the remaining reduction rules are straightforward. Rules [E-BETA], [E-NEW], [E-DEREF], and [E-ASSIGN] are standard. The rule [E-PRIM] relies on a function  $\delta : \text{Term} \times \text{Term} \rightarrow \text{Term}$  to define the semantics of constant functions. For simplicity, we assume each  $\delta$  to be defined for all constants  $k$  and arguments  $v$ . Rule [E-CBETA] applies function casts by casting the function argument and result. Rule [E-CDEREF] casts the result of reading a cell and [E-CASSIGN] casts the value written to a cell and casts the value again to the expected output type. Rules [E-ID] and [E-FCAST] respectively perform the identity and float coercions, and are restricted to non-cast contexts to prevent overlap with [E-CCAST].

### 6.1 Space-efficient coercions

Using these evaluation rules, the continuation-passing  $\text{evenk}_c$  and  $\text{oddk}_c$  functions of Sect. 1.1 now maintain normalized coercions during evaluation:

$$\begin{aligned} & \text{evenk}_c \quad n \quad (\langle \text{Fun Bool? Bool!} \rangle \lambda x:\text{Bool}. x) \\ \longrightarrow^* & \text{oddk}_c \quad (n - 1) \quad (\langle \text{Fun Bool! Bool?} \rangle \langle \text{Fun Bool? Bool!} \rangle \lambda x:\text{Bool}. x) \\ \longrightarrow & \text{oddk}_c \quad (n - 1) \quad (\langle \langle \text{Fun Bool? Bool!} \rangle \text{Fun Bool! Bool?} \rangle \lambda x:\text{Bool}. x) \\ & \quad (\text{Fun Bool! Bool?}) \lambda x:\text{Bool}. x \quad [\text{E-CCAST}] \\ = & \text{oddk}_c \quad (n - 1) \quad (\langle \text{Fun (Bool!; Bool?) (Bool!; Bool?)} \rangle \lambda x:\text{Bool}. x) \\ = & \text{oddk}_c \quad (n - 1) \quad (\langle \text{Fun } I \text{ } I \rangle \lambda x:\text{Bool}. x) \\ \longrightarrow^* & \text{evenk}_c \quad (n - 2) \quad (\langle \text{Fun Bool? Bool!} \rangle \langle \text{Fun } I \text{ } I \rangle \lambda x:\text{Bool}. x) \\ = & \text{evenk}_c \quad (n - 2) \quad (\langle \text{Fun Bool? Bool!} \rangle \lambda x:\text{Bool}. x) \\ \longrightarrow^* & \dots \end{aligned}$$

Note the crucial role played by rule [E-CCAST] in this reduction sequence to combine the adjacent function casts. Once these casts are combined, the coercion normalization rules simplify them down to ( $\text{Fun } I \ I$ ) at each subsequent invocation of  $odd_c$ .

$$\begin{aligned} & (\text{Fun Bool? Bool!}); (\text{Fun Bool! Bool?}) \\ & = \text{Fun} (\text{Bool!}; \text{Bool?}) (\text{Bool!}; \text{Bool?}) \\ & = \text{Fun } I \ I \end{aligned}$$

## 6.2 Tail recursion

Consider the compiled version of the direct-style *even* and *odd* functions:

$$\begin{aligned} even_c : \text{Dyn} \rightarrow \text{Dyn} & \stackrel{\text{def}}{=} \lambda n: \text{Dyn}. \\ & \quad \text{if } (\langle \text{Int?} \rangle n) = 0 \\ & \quad \text{then } \langle \text{Bool!} \rangle \text{ true} \\ & \quad \text{else } \langle \text{Bool!} \rangle (odd_c ((\langle \text{Int?} \rangle n) - 1))) \\ odd_c : \text{Int} \rightarrow \text{Bool} & \stackrel{\text{def}}{=} \lambda n: \text{Int}. \\ & \quad \text{if } (n = 0) \\ & \quad \text{then } \text{false} \\ & \quad \text{else } \langle \text{Bool?} \rangle (even_c ((\langle \text{Int!} \rangle (n - 1)))) \end{aligned}$$

In the following evaluation sequence, the definition of evaluation contexts forces the evaluation of rule [E-CCAST] before reducing the application of *odd*. In other words, tail calls are evaluated only *after* normalizing the cast at the top of the program stack. Thus, as shown below, there is no unbounded growth in the control stack.

$$\begin{aligned} & \langle \text{Bool?} \rangle (even_c ((\langle \text{Int!} \rangle 4))) \\ & \longrightarrow^* \langle \text{Bool?} \rangle \langle \text{Bool!} \rangle (odd_c 3) \\ & \longrightarrow \langle \text{Bool!}; \text{Bool?} \rangle (odd_c 3) \\ & = \langle I \rangle (odd_c 3) \\ & \longrightarrow^* \langle \text{Bool?} \rangle (even_c ((\langle \text{Int!} \rangle 2))) \\ & \longrightarrow^* \dots \end{aligned}$$

## 6.3 Type soundness

Evaluation satisfies the usual preservation and progress lemmas.

**Theorem 3** (Soundness of evaluation) *If  $\emptyset; \emptyset \vdash t : T$  then either*

1.  $t, \emptyset$  diverges,
2.  $t, \emptyset \longrightarrow^* C[\langle \text{Fail} \rangle u], \sigma$  or
3.  $t, \emptyset \longrightarrow^* v, \sigma$  and  $\exists \Sigma$  such that  $\emptyset; \Sigma \vdash v : T$  and  $\emptyset; \Sigma \vdash \sigma$ .

*Proof* Via standard subject reduction and progress lemmas in the style of Wright and Felleisen [37].  $\square$

## 7 Early error detection

Our coercion-based implementation of casts allows the runtime to detect errors earlier in the evaluation sequence and in more situations than earlier systems for gradual typing. To illustrate this idea, consider the following code fragment, which erroneously attempts to convert an  $(\text{Int} \rightarrow \text{Int})$  function to have type  $(\text{Bool} \rightarrow \text{Int})$ :

```
let f : (\text{Int} \rightarrow \text{Int}) = (\lambda x : \text{Int}. x) in
let h : (\text{Bool} \rightarrow \text{Int}) = f in ...
```

Most gradual type systems (including ours) will detect this error at compile time, since these types are not consistent:  $(\text{Int} \rightarrow \text{Int}) \not\sim (\text{Bool} \rightarrow \text{Int})$ . However, the introduction of an additional variable of type  $\text{Dyn}$  into the data-flow path hides this error from the type system, since  $(\text{Int} \rightarrow \text{Int}) \sim \text{Dyn} \sim (\text{Bool} \rightarrow \text{Int})$ :

```
let f : (\text{Int} \rightarrow \text{Int}) = (\lambda x : \text{Int}. x) in
let g : \text{Dyn} = f in
let h : (\text{Bool} \rightarrow \text{Int}) = g in ...
```

For this revised example, the function representation of higher-order casts would not detect this error until  $h$  is applied to a boolean argument, causing the integer cast to fail.

In contrast, our coercion-based implementation allows us to detect this error as soon as  $h$  is initialized. After cast insertion and evaluation, the value of  $h$  is

$$\begin{aligned} h &= (\langle \text{Fun } \text{Bool! Int?}; (\text{Fun Int? Int!}) \rangle (\lambda x : \text{Int}. x)) \\ &= \langle \text{Fun } (\text{Int?}; \text{Bool!}) (\text{Int?}; \text{Int!}) \rangle (\lambda x : \text{Int}. x) \\ &= \langle \text{Fun Fail I} \rangle (\lambda x : \text{Int}. x) \end{aligned}$$

where the domain coercion  $\text{Fail}$  explicates the inconsistency of the two domain types  $\text{Int}$  and  $\text{Bool}$ . We can modify our semantics to halt as soon as such inconsistencies are detected, by adding the following coercion normalization rules to lift this failure to the top level:

$$\begin{array}{ll} \text{Fun } c \text{ Fail} = \text{Fail} & \text{Ref } c \text{ Fail} = \text{Fail} \\ \text{Fun Fail } c = \text{Fail} & \text{Ref Fail } c = \text{Fail} \end{array}$$

Under these additional normalization rules, we have:

$$\begin{aligned} h &= \langle \text{Fun Fail I} \rangle (\lambda x : \text{Int}. x) \\ &= \langle \text{Fail} \rangle (\lambda x : \text{Int}. x) \end{aligned}$$

which is not a value, and so evaluation gets stuck.

Using these rules, our implementation strategy halts as soon as  $h$  is initialized, resulting in earlier runtime error detection and the detection of errors in more situations than under the function representation of higher-order casts. We believe that the eager approach better fits the philosophy of gradual typing, which is to permit maximum flexibility while still detecting errors as early as possible.

## 8 Space efficiency

We now consider how much space coercions consume at runtime, beginning with an analysis of how much space each individual coercion can consume.

### 8.1 Space consumption

The size of a coercion  $\text{size}(c)$  is defined as the size of its abstract syntax tree representation. When two coercions are sequentially composed and normalized during evaluation, the size of the normalized, composed coercion may of course be larger than either of the original coercions. In order to reason about the space required by such composed coercions, we introduce a notion of the *height* of a coercion:

$$\begin{aligned} \text{height}(I) &= \text{height}(\text{Fail}) = \text{height}(G!) = \text{height}(G?) = 1 \\ \text{height}(\text{Ref } c\ d) &= \text{height}(\text{Fun } c\ d) = 1 + \max(\text{height}(c), \text{height}(d)) \\ \text{height}(c; d) &= \max(\text{height}(c), \text{height}(d)) \end{aligned}$$

Notably, the height of a composed coercion is bounded by the maximum height of its constituents. In addition, normalization never increases the height of a coercion. Thus, the height of any coercion created during program evaluation is never larger than the height of some coercion in the original elaborated program.

Furthermore, this bound on the height of each coercion in turn guarantees a bound on the coercion's size, according to the following lemma. In particular, even though the length of a coercion sequence  $c_1; \dots; c_n$  does not contribute to its height, the restricted structure of well-typed, normalized coercions constrains the length (and hence size) of such sequences.

**Lemma 2** *For all well-typed normalized coercions  $c$ ,  $\text{size}(c) \leq 5(2^{\text{height}(c)} - 1)$ .*

*Proof* Induction on  $c$ . Assume  $c = c_1; \dots; c_n$ , where each  $c_i$  is not a sequential composition.

Suppose some  $c_i = \text{Ref } d_1\ d_2$ . So  $\vdash c_i : \text{Ref } S \rightsquigarrow \text{Ref } T$ . Hence  $c_i$  can be preceded only by  $\text{Ref}?$ , which must be the first coercion in the sequence, and similarly can be followed only by  $\text{Ref}!$ , which must be the last coercion in the sequence. Thus in the worst case  $c = \text{Ref}?; \text{Ref } d_1\ d_2; \text{Ref}!$  and  $\text{size}(c) = 5 + \text{size}(d_1) + \text{size}(d_2)$ . Applying the induction hypothesis to the sizes of  $d_1$  and  $d_2$  yields:

$$\text{size}(c) \leq 5 + 2(5(2^{\text{height}(c)-1} - 1)) = 5(2^{\text{height}(c)} - 1)$$

The case for  $\text{Fun } d_1\ d_2$  is similar. The coercions  $I$  and  $\text{Fail}$  can only appear alone. Finally, coercions of the form  $G?; G!$  are valid. However, composition of a coercion  $c$  matching this pattern with one of the other valid coercions is either ill-typed or triggers a normalization that yields a coercion identical to  $c$ .  $\square$

In addition, the height of any coercion created during cast insertion is bounded by the height of some type in the typing derivation (where the height of a type is the height of its abstract syntax tree representation).

**Lemma 3** *If  $c = \text{coerce}(S, T)$ , then  $\text{height}(c) \leq \max(\text{height}(S), \text{height}(T))$ .*

*Proof* Induction on the structure of  $\text{coerce}(S, T)$ .  $\square$

**Theorem 4** (Size of coercions) *For any  $e_0, c$  such that*

1.  $\emptyset \vdash e_0 \hookrightarrow t_0 : T$  and
2.  $t_0, \emptyset \longrightarrow^* t, \sigma$  and
3.  $c$  occurs in  $(t, \sigma)$ ,

there exists some type  $S$  in the derivation of  $\emptyset \vdash e_0 \leftrightarrow t_0 : T$  such that  $\text{height}(c) \leq \text{height}(S)$  and  $\text{size}(c) \leq 5(2^{\text{height}(S)} - 1)$ .

*Proof* Induction on the length of the reduction sequence. The base case proceeds by induction on the derivation of  $\emptyset \vdash e_0 \leftrightarrow t_0 : T$ . In cases [C-APP1], [C-APP2], [C-DEREF2], [C-ASSIGN1], and [C-ASSIGN2], the new coercions are introduced by applications of  $\text{coerce}(S, T)$  for types  $S$  and  $T$  in the derivation. In each such case, Lemma 3 guarantees that  $\text{height}(c) \leq \max(\text{height}(S), \text{height}(T))$ , and by Lemma 2,  $\text{size}(c) \leq 5(2^{\text{height}(\max(\text{height}(S), \text{height}(T)))} - 1)$ . The remaining cases are trivial inductions.

If the reduction sequence is non-empty, we have  $t_0, \emptyset \longrightarrow^k t', \sigma' \longrightarrow t, \sigma$ . If  $c$  occurs in  $t'$ ,  $\sigma'$  then by induction, there exists an  $S$  in the derivation of  $\emptyset \vdash e_0 \leftrightarrow t_0 : T$  such that  $\text{height}(c) \leq \text{height}(S)$  and  $\text{size}(c) \leq 5(2^{\text{height}(S)} - 1)$ . Otherwise, the only reduction rules that produce fresh coercions are [E-CBETA], [E-CDEREF], [E-CASSIGN], all of which produce strictly smaller coercions, and [E-CCAST]. In the latter case, the fresh coercion  $c = c_1; c_2$  is synthesized from coercions  $c_1$  and  $c_2$  in the redex. By induction, there are types  $S_1$  and  $S_2$  in the derivation of  $\emptyset \vdash e_0 \leftrightarrow t_0 : T$  such that  $\text{height}(c_i) \leq \text{height}(S_i)$  and  $\text{size}(c_i) \leq 5(2^{\text{height}(S_i)} - 1)$  for  $i \in 1, 2$ . Since  $\text{height}(c) = \max(\text{height}(c_1), \text{height}(c_2))$ , let  $S'$  be the taller of the two types. Then  $\text{height}(c) \leq \text{height}(S')$  and by Lemma 2,  $\text{size}(c) \leq 5(2^{\text{height}(c)} - 1)$ , so  $\text{size}(c) \leq 5(2^{\text{height}(S')} - 1)$ .  $\square$

We now bound the total cost of maintaining coercions in the space-efficient semantics. We define the size of a program state inductively as the sum of the sizes of its components. In order to construct a realistic measure of the store, we count only those cells that an idealized garbage collector would consider live by restricting the  $\text{size}$  function to the domain of reachable addresses.

$$\begin{aligned}\text{size}(t, \sigma) &= \text{size}(t) + \text{size}(\sigma|_{\text{reachable}(t)}) \\ \text{size}(\sigma) &= \sum_{a \in \text{dom}(\sigma)} (1 + \text{size}(\sigma(a))) \\ \text{size}(k) &= \text{size}(a) = \text{size}(x) = 1 \\ \text{size}(\lambda x : T. t) &= 1 + \text{size}(T) + \text{size}(t) \\ \text{size}(\text{ref } t) &= \text{size}(!t) = 1 + \text{size}(t) \\ \text{size}(t_1 := t_2) &= \text{size}(t_1 t_2) = 1 + \text{size}(t_1) + \text{size}(t_2) \\ \text{size}(\langle c \rangle t) &= 1 + \text{size}(c) + \text{size}(t)\end{aligned}$$

To show that coercions occupy bounded space in the model, we compare the size of program states in reduction sequences to program states in an “oracle” semantics where coercions require no space. The oracular measure  $\text{size}_{\text{OR}}$  is defined similarly to  $\text{size}$ , but without a cost for maintaining coercions; that is,  $\text{size}_{\text{OR}}(c) = 0$ . The following theorem then bounds the fraction of the program state occupied by coercions in the space-efficient semantics.

**Theorem 5** (Space consumption) *If  $\emptyset \vdash e \leftrightarrow t : T$  and  $t, \emptyset \longrightarrow^* t', \sigma$ , then there exists some type  $S$  in the derivation of  $\emptyset \vdash e \leftrightarrow t : T$  such that  $\text{size}(t', \sigma) \in O(2^{\text{height}(S)} \cdot \text{size}_{\text{OR}}(t', \sigma))$ .*

*Proof* During evaluation, the [E-CCAST] rule prevents nesting of adjacent coercions in any term in the evaluation context, redex, or store. Thus the number of coercions in the program state is proportional to the size of the program state. By Theorem 4 the size of each coercion is in  $O(2^{\text{height}(S)})$  for the largest  $S$  in the typing of  $e$ .  $\square$

## 8.2 Implementing tail recursion

Theorem 5 has the important consequence that coercions do not affect the control space consumption of tail-recursive programs. This important property is achieved by always combining adjacent coercions on the stack via the [E-CCAST] rule. This section sketches three implementation techniques for this rule.

*Coercion-passing style* This approach is similar to security-passing style [36], which adds an extra argument to every procedure representing its security context; in this case, the argument represents the result coercion. Thus the representation of a function type is

$$\lceil S \rightarrow T \rceil = \lceil S \rceil \times \text{Coercion} \rightarrow \lceil T \rceil$$

Tail calls coalesce but do not perform coercions, instead passing them along to the next function. Coercions are instead discharged when a value is returned to a context with a non-trivial coercion.

*Trampoline* A trampoline [18] is a well-known technique for implementing tail-recursive languages where tail calls are implemented by returning a thunk to a top-level loop. Tail-recursive functions with coercions return both a thunk and a coercion to the driver loop, which accumulates and coalesces returned coercions. The representation of a function type is then

$$\lceil S \rightarrow T \rceil = \lceil S \rceil \rightarrow (\lceil T \rceil + ((\mathbf{1} \rightarrow \lceil T \rceil) \times \text{Coercion}))$$

*Continuation marks* Continuation marks [9, 10] allow programs to annotate continuation frames with arbitrary data. When a marked frame performs a tail call, the subsequent frame can inherit and modify the destroyed frame's marks. Coercions on the stack are stored as marks and coalesced on tail calls. In this approach, the representation of a function type is direct:

$$\lceil S \rightarrow T \rceil = \lceil S \rceil \rightarrow \lceil T \rceil$$

since the presence of continuation marks makes the target language expressive enough to store coercions on the stack. However, the implementation of tail calls requires a protocol for coalescing coercions, so compilation is still required.

## 9 Related work

There is a large body of work combining static and dynamic typing. The simplest approach is to use reflection with the type `Dyn`, as in Amber [6]. Since case dispatch cannot be precisely type-checked with reflection alone, many languages provide statically-typed `typecase` on dynamically-typed values, including Simula-67 [2] and Modula-3 [7].

For dynamically-typed languages, *soft typing* systems provide type-like static analyses for optimization and early error reporting [37]. These systems may provide static type information but do not allow explicit type annotations, whereas enforcing documented program invariants (i.e., types) is a central feature of gradual typing.

Similarly, Henglein's theory of dynamic typing [22] provides a framework for static type optimizations but only in a purely dynamically-typed setting. We use Henglein's coercions instead for structuring the algebra of our cast representation. Our application is essentially

different: in the gradually-typed setting, coercions serve to enforce explicit type annotations, whereas in the dynamically-typed setting, coercions represent checks required by primitive operations.

None of these approaches facilitates *migration* between dynamically and statically-typed code, at best requiring hand-coded interfaces between them. The gradually-typed approach, exemplified by Sage [20] and  $\lambda^?$  [32], lowers the barrier for code migration by allowing mixture of expressions of type Dyn with more precisely-typed expressions. Our work improves gradual typing by eliminating the drastic effects on space efficiency subtly incurred by crossing the boundary between typing disciplines. Siek and Taha's subsequent work on gradual typing for objects [33] merges some casts at runtime, but they do not address tail recursion and so the remaining casts may still occupy unbounded space. In particular, our Space Consumption Theorem does not hold in their setting.

In earlier work in this area, [28] explored similar ideas in the context of dependent type systems. Their approach uses function proxies to support clean interoperation and hence migration between simply-typed code and dependently-typed code. Several other systems employ dynamic function proxies, including hybrid type checking [16], software contracts [15], and recent work on software migration by Tobin-Hochstadt and Felleisen [34, 35]. We believe our approach to coalescing redundant proxies could improve the efficiency of all of these systems.

Minamide and Garrigue [27] address a similar issue with unbounded proxies but in the context of specialization of polymorphic functions rather than general type conversions. Their approach does not handle tail calls correctly and only addresses runtime complexity. Our work presents a representation of casts as coercions in order to improve the memory guarantees of gradual typing. Nevertheless, our representation appears to improve the runtime efficiency similarly to Minamide and Garrigue's work. We intend to explore the runtime improvements of our approach as well.

## 10 Conclusion

We have presented a space-efficient implementation strategy for the gradually-typed  $\lambda$ -calculus. By ensuring that the presence of coercions does not asymptotically affect the space consumption of programs, we guarantee that programmers can gradually annotate programs with types without inadvertently introducing significant changes in resource consumption. Since the publication of our original conference results [23], other researchers have extended these ideas to support blame tracking [31] and to express normalized coercions as more elegant threesomes [30].

Most work on hybrid and gradual type systems to date has been developed in the context of rather idealized languages. The motivation for these idealized languages has been to enable researchers to focus on some of the essential and theoretical complexities of these type systems. However, some systems do not support side effects, and others, such as [19], assume the language is strongly-normalizing.

There remain some difficulties in scaling up these ideas to realistic languages with full-featured type systems that include features such as classes, subtyping, method overriding, recursive types, or polymorphism, particularly while preserving the desirable property of space efficiency.

Polymorphic types present some challenges for maintaining constant bounds on the size of coercions, in part because polymorphic types require representation of data structures with binders. The integration of polymorphism into various contract systems and hybrid/gradual type systems has been a topic of active research over the last several years.

Polymorphism was incorporated into dynamic contracts by Guha et al. [21]. More recently, gradual typing was extended to support polymorphism, using dynamic sealing to ensure that values of polymorphic type satisfy relational parametricity [1]. Neither of these results investigate space overheads.

Predicate types pose additional challenges for bounding the size of coercions, and may require additional assumptions, such as a bound on the number of predicates used in types. Combining polymorphism with predicate types may be tricky, and most research restricts predicates to base types. Some systems support predicates on non-base types (see for example, [3, 20, 38]), and this generalization seems necessary to support polymorphism, since predicates could be used to refine polymorphic arguments. In particular, the programming language SAGE [20] was an experiment at applying hybrid typing to a rich type system with first-class types. First-class types can then be used to express explicit polymorphism, where type parameters are passed explicitly, just like other parameters. Unfortunately, the lack of type inference in SAGE made type annotations and explicit type parameters somewhat burdensome. There was no evaluation of the space overheads in SAGE.

The integration of recursive types into hybrid/gradual type systems also deserves further study. Contract systems have long supported recursive contracts (for example, a list contract). SAGE's first-class types can express recursive types, but there are no results on the space overheads of the resulting coercions. It appears fairly plausible that the recursive types can be supported in a gradual type system in a space-efficient manner, although the details remain to be worked out.

In addition to the technical difficulties of scaling gradual and hybrid types to rich languages in a space-efficient manner, another key challenge for future work is to demonstrate the software engineering benefits of these techniques for code development and evolution.

## References

1. Ahmed, A., Findler, R.B., Matthews, J., Wadler, P.: Blame for all. In: Workshop on Script to Program Evolution (2009)
2. Birtwhistle, G.M., et al.: Simula Begin. Chartwell-Brett Ltd., Bromley (1979)
3. Blume, M., McAllester, D.: Sound and complete models of contracts. *J. Funct. Program.* **16**(4–5), 375–414 (2006). ISSN 0956-7968. doi:[10.1007/s10990-011-9066-z](https://doi.org/10.1007/s10990-011-9066-z)
4. Bracha, G.: Pluggable type systems. In: Workshop on Revival of Dynamic Languages, October 2004
5. Brus, T., van Eekelen, M., van Leer, M., Plasmeijer, M.: Clean nnn a language for functional graph rewriting. In: Kahn, G. (ed.) Functional Programming Languages and Computer Architecture. Lecture Notes in Computer Science, vol. 274, pp. 364–384. Springer, Berlin (1987)
6. Cardelli, L.: Amber. In: Spring School of the LITP on Combinators and Functional Programming Languages, pp. 21–47 (1986)
7. Cardelli, L., Donahue, J., Glassman, L., Jordan, M., Kalsow, B., Nelson, G.: Modula-3 report (revised). Technical Report 52, DEC SRC (1989)
8. Chambers, C.: The Cecil Language Specification and Rationale: Version 3.0. University of Washington (1998)
9. Clements, J.: Portable and high-level access to the stack with Continuation Marks. PhD thesis, Northeastern University (2005)
10. Clements, J., Felleisen, M.: A tail-recursive machine with stack inspection. *ACM Trans. Program. Lang. Syst.* **26**(6), 1029–1052 (2004)
11. de Oliveira, R.B.: The Boo programming language (2005)
12. ECMAScript Edition 4 group wiki. Ecma International (2007)
13. ECMAScript Language Specification. Ecma International, 5th edn. (2009). URL <http://www.ecma-international.org/publications/files/ECMA-ST/ECMA-262.pdf>
14. Findler, R.B., Blume, M.: Contracts as pairs of projections. In: Proceedings of the International Symposium on Functional and Logic Programming (FLOPS), pp. 226–241 (2006)
15. Findler, R.B., Felleisen, M.: Contracts for higher-order functions. In: Proceedings of the ACM SIGPLAN International Conference on Functional Programming (ICFP), pp. 48–59, October 2002

16. Flanagan, C.: Hybrid type checking. In: Proceedings of the ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL), pp. 245–256 (2006)
17. Flanagan, C., Freund, S.N., Tomb, A.: Hybrid types, invariants, and refinements for imperative objects. In: The International Workshop on Foundations and Developments of Object-Oriented Languages (FOOL/WOOD) (2006)
18. Ganz, S.E., Friedman, D.P., Wand, M.: Trampolined style. In: Proceedings of the ACM SIGPLAN International Conference on Functional Programming (ICFP), pp. 18–27 (1999)
19. Greenberg, M., Pierce, B., Weirich, S.: Contracts made manifest. In: Proceedings of the ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL) (2010)
20. Gronski, J., Knowles, K., Tomb, A., Freund, S.N., Flanagan, C.: Sage: Hybrid checking for flexible specifications. In: Scheme and Functional Programming Workshop, September 2006
21. Guha, A., Matthews, J., Findler, R.B., Krishnamurthi, S.: Relationally-parametric polymorphic contracts. In: DLS, pp. 29–40 (2007)
22. Henglein, F.: Dynamic typing: Syntax and proof theory. *Sci. Comput. Program.* **22**(3), 197–230 (1994)
23. Herman, D., Tomb, A., Flanagan, C.: Space-efficient gradual typing. In: Trends in Functional Programming (TFP) (2007)
24. Knowles, K.W., Flanagan, C.: Hybrid type checking. *ACM Trans. Program. Lang. Syst.* **32**(2) (2010)
25. Matthews, J., Findler, R.B.: Operational semantics for multi-language programs. In: Proceedings of the ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL) (2007)
26. Meijer, E., Drayton, P.: Static typing where possible, dynamic typing when needed. In: Workshop on Revival of Dynamic Languages (2005)
27. Minamide, Y., Garrigue, J.: On the runtime complexity of type-directed unboxing. In: Proceedings of the ACM SIGPLAN International Conference on Functional Programming (ICFP), pp. 1–12 (1998)
28. Ou, X., Tan, G., Mandelbaum, Y., Walker, D.: Dynamic typing with dependent types. In: Proceedings of the IFIP International Conference on Theoretical Computer Science, pp. 437–450 (2004)
29. Pil, M.: First class file I/O. In: IFL '96: Selected Papers from the 8th International Workshop on Implementation of Functional Languages, pp. 233–246, London, UK. Springer, Berlin (1997). ISBN 3-540-63237-9
30. Siek, J., Wadler, P.: Exploring the design space of higher-order casts. In: Proceedings of the ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL) (2010)
31. Siek, J., Garcia, R., Taha, W.: Exploring the design space of higher-order casts. In: Proceedings of the European Symposium on Programming Languages and Systems (ESOP), pp. 17–31. Springer, Berlin (2009)
32. Siek, J.G., Taha, W.: Gradual typing for functional languages. In: Scheme and Functional Programming Workshop, September 2006
33. Siek, J.G., Taha, W.: Gradual typing for objects. In: Proceedings of the European Conference on Object-Oriented Programming (ECOOP), Berlin, Germany, July 2007
34. Tobin-Hochstadt, S., Felleisen, M.: Interlanguage migration: From scripts to programs. In: Dynamic Languages Symposium, October 2006
35. Tobin-Hochstadt, S., Felleisen, M.: The design and implementation of typed scheme. In: Proceedings of the ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL), pp. 395–406 (2008)
36. Wallach, D.S.: A New Approach to Mobile Code Security. PhD thesis, Princeton University, Department of Computer Science (1999)
37. Wright, A.K.: Practical Soft Typing. PhD thesis, Rice University (August 1998)
38. Xu, D.N., Peyton Jones, S.L., Claessen, K.: Static contract checking for Haskell. In: POPL, pp. 41–52 (2009)