

Space-Efficient Monotonic References

DEYAAELDEEN ALMAHALLAWI, Indiana University Bloomington, USA

JEREMY G. SIEK, Indiana University Bloomington, USA

Integrating static typing and dynamic typing has received considerable attention in academia and industry. Gradual typing is one approach to this integration that preserves type soundness by casting values at run-time. For higher order values such as functions and mutable references, a cast typically wraps the values in a proxy that performs casts when the values are used. This approach suffers from two problems: (1) chains of proxies can grow and consume unbounded space, and (2) statically typed code regions need to check whether values are proxied. Monotonic references [Siek et al. 2015c] solve both problems for mutable references by directly casting the heap cell instead of wrapping the reference in a proxy. In this paper we show how to integrate monotonic references with the coercion-based solution to the first problem for other values such as functions [Herman et al. 2007]. We present a space-efficient version of the semantics for monotonic references, functions, and pairs and prove an upper bound on space overhead. Furthermore, the prior semantics for monotonic references involved storing cast expressions (not yet values) on the heap and it is not obvious how to implement this behavior efficiently in a compiler and run-time system. In our new semantics, only values are written to the heap, making the semantics straightforward to implement. Finally, we provide a type safety proof that is fully mechanized in Agda.

CCS Concepts: • **Software and its engineering** → **General programming languages; Semantics.**

Additional Key Words and Phrases: gradual-typing, reduction semantics, mutable state

ACM Reference Format:

Deyaaeldeen Almahallawi and Jeremy G. Siek. 2020. Space-Efficient Monotonic References. In *Informal Proceedings of the first ACM SIGPLAN Workshop on Gradual Typing (WGT20)*. ACM, New York, NY, USA, 22 pages.

1 WHAT ARE THE PERFORMANCE CHARACTERISTICS OF GRADUAL TYPING?

Gradual typing combines static and dynamic type checking, giving the programmer control over which typing discipline to be used in each region of code [Anderson and Drossopoulou 2003; Gronski et al. 2006; Siek and Taha 2006; Tobin-Hochstadt and Felleisen 2006]. Siek et al. [2015b] describe five criteria for gradually typed languages, including *type soundness* and the *gradual guarantee*. The type soundness criteria requires that the value of an expression must have the type that was predicted by the type system. The gradual guarantee states that changing type annotations should not change the semantics of the program, except that incorrect type annotations may induce compile-time or run-time errors.

Type soundness is enforced, without compromising the expressive power of the language, by performing runtime type checking at the boundaries between statically typed and dynamically typed regions of code. This runtime overhead could be significant, depending on the nature of the language and what kind of features are supported. For instance, for languages with nominal types, the runtime check for whether a value has a given type is efficient and straightforward to implement. Indeed, Muehlboeck and Tate [2017] show that Nom, a nominally-typed object-oriented language (without generics or function types), exhibits low overheads. On the other hand, with structural types, the runtime check can be much more complex, e.g., for higher-order values it may involve the use of a proxy to mediate between a value and its context. For instance, casting

WGT20, January 25, 2020, New Orleans

2020. ACM ISBN 978-x-xxxx-xxxx-x/YY/MM.

a function value wraps it in a proxy [Findler and Felleisen 2002] that checks, at application sites, whether the types of the input and output values matches that of the cast. Similarly, casting a reference value wraps the reference in a proxy that checks at read sites if the value on the heap could be cast to the type the context expects, and checks at write sites if the written value could be cast to the type of the heap cell.

Using proxies to cast higher-order values causes two problems. First, each time a value of a higher-order type crosses the boundary between a statically typed and dynamically typed code regions, it gets wrapped in a proxy. These chains of proxies can grow without bounds, causing space leaks Herman et al. [2007, 2010]. Furthermore, the chains of proxies can cause runtime slowdowns to the extent of changing the asymptotic complexity of a program. Kuhlenschmidt et al. [2019] demonstrate this problem with a simple quicksort program that is fully statically typed except for one annotation, causing a proxy to be created around the vector being sorted at each recursive call to the sort function. The runtimes collected from running the program on different input sizes suggest that the asymptotic complexity changes from $O(n^2)$ to $O(n^3)$.

The second problem with using proxies is the overhead of runtime dispatch on higher-order values at use sites. At each use site, the implementation must check whether the incoming value is proxied, and if so, perform runtime checks before and/or after processing the underlying value. This runtime dispatch is necessary even in statically typed code regions, because a reference from a dynamically typed region may flow into a static region. Kuhlenschmidt et al. [2019] observe about 20% overhead in a statically-typed matrix multiplication benchmark. This problem prevents the performance of statically-typed code in a gradually-typed programming language to be on par with corresponding code in statically-typed programming languages.

The good news is that these two problems have been solved for the case of functions. Herman et al. [2007, 2010] use the coercions of Henglein [1994] to compress chains by reducing sequences of coercions to normal form and Siek et al. [2015a] define a recursive function for composing coercions in normal form. Furthermore, Siek and Garcia [2012] solve the dynamic dispatch problem for functions, while also compressing coercions, by using a closure representation that can act as either a regular closure or as a proxy but with a uniform calling convention. This technique is used in Grift [Kuhlenschmidt et al. 2019].

For mutable references, Herman et al. [2007, 2010] solve the first problem by compressing chains of proxies using coercions, but their design does not solve the second problem of needing runtime dispatch to handle proxies. Siek et al. [2015c] introduce an alternative design, called *monotonic references*, that does not wrap addresses in proxies and instead casts values in the heap directly, effectively solving the dynamic dispatch issue. However, Siek et al. [2015c] leave as future work the challenge of identifying a normal form for coercions with monotonic references and of developing a function for composing such coercions. Furthermore, their monotonic heap maps addresses to *expressions* that are reduced using a small-step reduction relation. (Normally only values are stored in the heap.) It is not obvious how to implement this reduction-in-the-heap in a compiled implementation of a gradually-typed language.

In this paper, we answer the following questions.

Q.1 How to integrate monotonic references into a space-efficient coercion calculus?

Q.2 How to implement monotonic references in a way that only puts values in the heap?

We answer those questions by presenting a new reduction semantics for monotonic references that is space-efficient for functions and that writes values only to the heap. To achieve the former, we interpret casts as coercions [Henglein 1994; Herman et al. 2007, 2010] and extend the coercion composition operator from prior work [Siek et al. 2015a], to also compose coercions for monotonic references. Regarding the latter, we provide a cast function that returns a value (that can then

be written to the heap) and a list of casts on addresses that are put into a queue for subsequent processing, making the semantics straightforward to implement.

The rest of the paper is organized as follows: Section 2 illustrates the problem of chains of function proxies and the importance of space-efficiency. Section 3 illustrates why prior work on monotonic references wrote cast expressions to the heap. Section 4 gives an informal presentation of our approach. Sections 5 and 6 present our approach formally and prove type safety.

2 THE PROBLEM OF CHAINS OF FUNCTION PROXIES

Proxies ensure that the type of a proxied value matches a certain type. Proxies are needed when it is hard to check immediately if the type of a value matches some type. For instance, it is in general impossible to check whether an arbitrary untyped function will return a Boolean without applying it first. With proxies, the function gets wrapped in a proxy that performs the desired checks at application sites. In this case, it checks if the incoming argument can be cast to the domain type of the underlying function, and also checks if the return value could be cast to Boolean.

An important consequence of that wrapping, however, is that values could accumulate multiple proxies when flowing through the boundaries between statically typed and dynamically typed regions of code. Operations on such proxied values have to go through all the layers of proxies, performing all their checks in order to access the underlying value.

This problem was first observed by Herman et al. [2007, 2010] in two higher-order, mutually recursive functions. One of the functions expects a statically typed function as input and the other expects a dynamically typed function, so implicit function casts are inserted to mediate between the types. A function cast creates a run-time proxy wrapping the function being cast. Every time one of the functions is called, a new proxy is created. These proxies consume space proportional to the number of recursive calls.

With monotonic references there are no proxies on references, however the prior work did not integrate the solutions for space efficiency on other higher-order values such as functions [Siek et al. 2015c]. In this paper we combine monotonic references with space-efficiency for other values.

3 THE PROBLEM OF EXPRESSIONS ON THE HEAP

In this section we review how and why the prior semantics for monotonic references puts expressions on the heap [Siek and Vitousek 2013; Siek et al. 2015c]. Consider the following example that creates a cycle in the heap, written in a variant of Typed Racket syntax that supports the dynamic type `Dyn` to enable fine-grained gradual typing. (Typed Racket uses the term “box” to mean reference. The `Boxof` type corresponds to the `Ref` type.)

```

1 (: f Dyn)
2 (define f (lambda (x) x))
3 (: r Dyn)
4 (define r (box (cons f '())))
5 (set-box! r (cons f r)) ;; establish a cycle
6 (define (g [x : (Boxof (Pair (Dyn -> Integer) (Boxof (Pair (Integer -> Dyn) Dyn)))]))
7   ((car (unbox x)) 42))
8 (g r)

```

On line 2, a reference is allocated at some address, call it a , and placed in variable r ; the heap cell at a is initialized with a pair of values of dynamic type. On line 3, a cycle is created by assigning to r another pair whose second element is r . So at this point the heap is in the following state.

$$\{a \mapsto \langle (\lambda x.x) \langle \star \rightarrow \star \Rightarrow \star \rangle, a \langle \text{Ref} (\star \times \star) \Rightarrow \star \rangle \rangle\}$$

On line 6, function g is applied to r , but g expects a reference to a pair of type

$$(\star \rightarrow \text{Int}) \times \text{Ref} ((\text{Int} \rightarrow \star) \times \star)$$

So a cast will be performed before proceeding with the function application. Casting the pair updates address a with the following expression, where values are colored purple and casts that are about to be applied are colored magenta:

$$a \mapsto \left\langle \begin{array}{l} (\lambda x.x) \langle \star \rightarrow \star \Rightarrow \star \rangle \langle \star \Rightarrow (\star \rightarrow \text{Int}) \rangle, \\ a \langle \text{Ref} (\star \times \star) \Rightarrow \star \rangle \langle \star \Rightarrow \text{Ref} ((\text{Int} \rightarrow \star) \times \star) \rangle \end{array} \right\rangle$$

The cast on the first element of the pair proceeds without further ado.

$$a \mapsto \left\langle \begin{array}{l} (\lambda x.x) \langle \star \rightarrow \star \Rightarrow \star \rightarrow \text{Int} \rangle, \\ a \langle \text{Ref} (\star \times \star) \Rightarrow \star \rangle \langle \star \Rightarrow \text{Ref} ((\text{Int} \rightarrow \star) \times \star) \rangle \end{array} \right\rangle$$

The cast on the second element of the pair is a cast to a reference type, so it will cast address a to the greatest lower bound (with respect to type precision) of its current type and the target type of the cast, which amounts to:

$$(\text{Int} \rightarrow \text{Int}) \times \text{Ref} ((\text{Int} \rightarrow \star) \times \star)$$

This cast is interesting because it causes another cast on the first element of the pair. Proceeding with a reduction step yields the following state.

$$a \mapsto \left\langle \begin{array}{l} (\lambda x.x) \langle \star \rightarrow \star \Rightarrow \star \rightarrow \text{Int} \rangle \langle \star \rightarrow \text{Int} \Rightarrow \text{Int} \rightarrow \text{Int} \rangle, \\ a \end{array} \right\rangle$$

In the above sequence of reductions, it was crucial that the cast on the first element of the pair to $\star \rightarrow \text{Int}$ was written to the heap before the first element was cast again to the type $\text{Int} \rightarrow \star$, enabling a greatest lower bound of $\text{Int} \rightarrow \text{Int}$ that took both types into account.

Unfortunately, it is not obvious how to implement this process efficiently in a compiler and runtime system. Literally placing expressions in the heap would require a runtime representation of abstract syntax trees and an interpreter, which would be costly to implement and rather slow to execute.

Alternatively, one might try to cast the value in a heap cell in one big step before writing it back to the heap, but the straightforward version of this idea is problematic because of race conditions. In the above example, the second cast on the first element, to $\text{Int} \rightarrow \star$, would see the current type as still $\star \rightarrow \star$, so the greatest lower bound would be $\text{Int} \rightarrow \star$. This would break the monotonicity property (and therefore type safety), because $\text{Int} \rightarrow \star$ is not less precise than $\star \rightarrow \text{Int}$. In this paper we present a refinement of the big-step approach. The key insight is that in the process of casting the value in a heap cell, when we come to an address, instead of immediately casting it, we leave the address as-is and instead put the address and the target type of the cast in a queue. Once the entire value has been processed, we write the result value back to the heap and then cast the addresses in the queue.

4 SPACE-EFFICIENCY AND VALUES ON THE HEAP

This section describes our approach to solve the two problems explained in Sections 2 and 3. Section 4.1 describes the design of normal coercions for monotonic references and Section 4.2 describes our design of the new reduction semantics for monotonic references that only writes values to the heap.

4.1 Coercions in Normal Form for Monotonic References

To answer **Q.1**, we extend the prior work on coercions by [Kuhlen Schmidt et al. 2019], designing normal coercions for monotonic references and extending the coercion composition operator to handle coercions that involve monotonic references.

A monotonic reference coercion holds just a type. To cast an address with some reference coercion $\text{Ref } T$, we cast the value at that address from its type to the greatest lower bound of its type and T . For convenient access, the type of a heap cell is stored next to the cell in the heap, and we refer to it as the run-time type information (RTTI). Casting an address with a sequence of two reference coercions amounts to taking the greatest lower bound of the RTTI and the types those two coercions hold. This suggests that composing two reference coercions should create a new reference coercion that holds a type that is the result of taking the greatest lower bound of the types of the coercions being composed. However this definition causes subtle changes to the coercion grammar and typing. This will be explained in more detail in Section 5.

4.2 Queue of Suspended Casts

To answer **Q.2**, we developed a new dynamic semantics for monotonic references, named λ_S^{ref} , that does not involve storing expressions on the heap. Instead, it uses a simple recursive function to cast values in the heap down to values. To maintain the correctness of reference casts in the presence of cycles, the function does not apply any inner reference casts. Instead it returns the address and puts the inner reference cast in a queue. Reduction proceeds by examining the queue, so a suspended cast is picked from the queue and gets applied. Applying a cast can return more suspended casts, so those get appended to the queue. This process continues until the queue is empty.

The main insight behind this design is that the correct ordering of cast reduction is a top-down breadth first traversal, as discussed in Section 3. With this semantics, the example in Section 3 can be reduced correctly. Recall that the value at address a is allocated at type $\star \times \star$ and is cast to

$$(\star \rightarrow \text{Int}) \times \text{Ref } ((\text{Int} \rightarrow \star) \times \star)$$

Because the top-level cast is a reference cast, it will be pushed to the queue. Subsequently, the cast will be read from the queue and applied to the pair in the heap. The cast on the first projection will be $\star \Rightarrow \star \rightarrow \text{Int}$ and will fully reduce to a value. The cast on the second projection will be suspended by reducing to the address a and the cast itself will be pushed to the queue. So this value will be written to the heap as follows. (The RTTI for the heap cell is written after the colon.)

$$\{a \mapsto \langle \lambda x. x \langle \star \rightarrow \star \Rightarrow \star \rightarrow \text{Int} \rangle, a \rangle : (\star \rightarrow \text{Int}) \times \text{Ref } ((\text{Int} \rightarrow \star) \times \star)\}$$

and the queue will contain the address a with the target type $(\text{Int} \rightarrow \star) \times \star$. Subsequently, the address and its target type will be popped from the queue. The value at the address is cast from its RTTI to the greatest lower bound of the RTTI and the target type. This results in the following final state:

$$\{a \mapsto \langle \lambda x. x \langle \star \rightarrow \star \Rightarrow \text{Int} \rightarrow \text{Int} \rangle, a \rangle : (\text{Int} \rightarrow \text{Int}) \times \text{Ref } ((\text{Int} \rightarrow \star) \times \star)\}$$

Note that casting the function for the second time composed the old cast $\star \rightarrow \star \Rightarrow \star \rightarrow \text{Int}$ and the new one $\star \rightarrow \text{Int} \Rightarrow \text{Int} \rightarrow \text{Int}$ into one cast $\star \rightarrow \star \Rightarrow \text{Int} \rightarrow \text{Int}$. This is a result of our semantics being space-efficient. (The details of this composition is omitted here to simplify the presentation of the reduction.)

It is worth noting that the length of the queue is bounded by the size of the value being cast. This means that the queue cannot grow without bound, so the space overhead caused by the queue is always a constant factor. This new approach is easier to understand and is straightforward to implement in a runtime system for gradually-typed languages.

5 COERCIONS IN NORMAL FORM FOR MONOTONIC REFERENCES

Coercions are combinators on types, originally designed for compile-time optimization for dynamically-typed languages [Henglein 1994]. They were later used by Herman et al. [2007, 2010] to guarantee space-efficiency at run-time for gradual typing. We follow suit and use coercions to guarantee space-efficiency in λ_S^{ref} , answering question **Q.1**.

There are many semantics for coercions [Siek et al. 2009]. In this paper, we use the Lazy-D semantics which allows any type, other than \star itself, to be injected directly to \star . Let I and J range over types other than \star . We adapt the definition of coercions for monotonic references from Siek et al. [2015c], adding a failure coercion of the form $\perp^{I,J}$.

$$c, d ::= \iota \mid I! \mid I? \mid c \rightarrow d \mid c \times d \mid \text{Ref } T \mid \perp^{I,J}$$

To review, the identity coercion ι acts as the identity function. The coercion $I!$ is an injection into \star that tags v with the source type. The projection coercion $I?$ projects the tagged value v from \star to I by checking if the tag is consistent with I . If this is the case, the underlying value is returned. The failure coercion \perp signals an error when the source and target types are not consistent. The function coercion $c \rightarrow d$ creates a proxy around the function v that applies coercion c to the argument and coercion d to the return value. The product coercion $c \times d$ applies coercion c to the first projection and coercion d to the second projection. Finally, and most relevant to this paper, the reference coercion $\text{Ref } T$ casts the value in the heap to be at least as precise as T .

We now turn to identifying a normal form for coercions with monotonic references and a function for composing them. We adapt the coercions in normal form of Kuhlenschmidt et al. [2019] (see the appendix in the auxiliary archive). Coercions in normal form are defined by a grammar consisting of three rules that enable coercion composition by a straightforward recursive function. This restricted grammar makes sure that the longest coercion will be one that starts with a projection, followed by a middle coercion, and ends with an injection. We adapt the normal form of Kuhlenschmidt et al. [2019] by replacing the traditional reference coercion with one for monotonic references, removing the duplicate identity coercion ι , and moving the failure coercion from final coercions to middle coercions. Figure 1 presents our coercions in normal form and how to create and compose them, which we describe next.

Our composition function \circledast takes two coercions in normal form as input and returns a coercion in normal form. It relies on two helper functions: \circledast_i that composes a final coercion with a coercion in normal form and returns a coercion in normal form, and \circledast_g that composes two middle coercions and returns a middle coercion.

Composing reference coercions. Composing two traditional reference coercions is done by composing the coercions for the read and the coercions for the write [Herman et al. 2007, 2010]. Composing two monotonic coercions is different because they hold types instead that are used to cast the heap. We merge the type of the two coercions by taking the greatest lower bound with respect to the precision relation \sqsubseteq (using the function \sqcap). However, the function \sqcap is partial and is undefined when the two input types are not consistent. In that case we return the failure coercion \perp instead, which is why we move failure coercions into the grammar for middle coercions.

Our coercion composition function maps two well-typed coercions to a well-typed coercion.

LEMMA 5.1 (WELL-TYPED COMPOSITION). *if $c : T_1 \Longrightarrow T_2$ and $d : T_2 \Longrightarrow T_3$, then $c \circledast d : T_1 \Longrightarrow T_3$*

PROOF. For the case $\text{Ref } T \circledast_g \text{Ref } T'$: We know that $T' \sqsubseteq T_3$ by $\text{Ref } T' : T_2 \Longrightarrow T_3$. We also know that $(T \sqcap T') \sqsubseteq T'$. By this and the transitivity of \sqsubseteq , we conclude that $(T \sqcap T') \sqsubseteq T_3$. \square

Our proof of that composition terminates relies on the fact that the sum of the sizes of the input coercions gets smaller at each recursive call.

PROPOSITION 5.2 (TERMINATION OF COMPOSITION). $size(c \circ d) \leq size(c) + size(d)$

To establish space-efficiency, we reason about the height of the created and composed coercions. The definition of height we use is different from the one of [Herman et al. \[2007\]](#) because of the difference in semantics between Lazy-UD and Lazy-D. Instead of injection and projection coercions having a height of 1 in the former, they instead have the height of the types they carry.

LEMMA 5.3 (COERCIONS HEIGHT IS BOUNDED). $\|T_1 \Rightarrow T_2\| \leq \max(\|T_1\|, \|T_2\|)$

PROOF. By induction on the coercion creation $T_1 \Rightarrow T_2$. \square

PROPOSITION 5.4 (COMPOSITION HEIGHT IS BOUNDED). $\|c \circ d\| \leq \max(\|c\|, \|d\|)$

PROOF. By Lemma 5.3 and induction on the structure of the compose function $c \circ d$. \square

A well-typed coercion in normal form consists of at most two sequences (three coercions or less), so a space-efficient coercion bounded in height is also bounded in size.

6 λ_S^{ref} DYNAMIC SEMANTICS

In this section, we present the reduction semantics for λ_S^{ref} , and show how it answers both questions [Q.1](#) and [Q.2](#). Most of our answer to question [Q.1](#) has been presented in Section 5. What remains to show is how to restrict the reduction relation to use the composition function \circ to combine adjacent coercions first before reducing expressions underneath. The approach we use is fairly standard [[Herman et al. 2007, 2010](#); [Siek et al. 2015a](#); [Siek and Wadler 2010](#)]. We adjust our notion of values, defined in Figure 2, to only allow at most one layer of casts. Then, a new rule for composing coercions is introduced in the pure coercion reduction relation \longrightarrow_c° , defined in Figure 3. Finally, expression reduction relation \longrightarrow_e , defined in Figure 4, is restricted such that reduction is disallowed under a sequence of two or more casts. Typically, this restriction is implemented by introducing a notion of cast-free evaluation context that does not have a cast application innermost and requiring the reduction of cast expressions to occur in that cast-free context. Instead, we use frames, which is a much simpler structure than evaluation contexts, and there is no need to worry about different ways to decompose them. The frame for cast expressions is removed from the grammar for frames, defined in Figure 2, and congruence rules for cast expressions is added to \longrightarrow_e . Furthermore, the reduction relation is indexed by a flag that indicates whether the cast congruence rule is allowed (cca) in the current context or whether it is disallowed (ccd). The cast congruence rules, `CASTCONG` and `CASTCONGE`, require the flag cca, and they switch the flag of the subexpression reduction to ccd. The rest of the rules are indexed by ccd. To allow rules indexed by ccd to be also available in cca contexts, we include a subsumption rule, `SWITCH`, that flips the flag from cca to ccd. This completes the presentation of our setup for space efficiency.

On the other hand, answering question [Q.2](#) requires the definition of the heap to be a map from addresses to pairs of values and types (Figure 2). This is achieved by changing how reference reduction works as follows. Instead of writing a delayed cast expression to the heap, reducing a reference cast creates a suspended cast and put it in a queue on the side (referred to by Ψ), without actually changing the heap. This is done using the rule `MONOCAST` of the program reduction relation \longrightarrow_e , defined in Figure 4. The state reduction relation \longrightarrow , defined in Figure 4, pauses expression reduction if the queue is not empty and pops a suspended cast to reduce. For the suspended cast on address a , the value mapped to this address gets cast from $\mu(a)_{\text{rtti}}$ to $\mu(a)_{\text{rtti}} \sqcap T$ using the apply-cast function, defined in Figure 3. The latter function fully reduces the cast to a value while also suspending any inner reference casts creating another queue of suspended casts. The result value gets written to the heap and the new queue gets appended to the old one. This

Coercions (in normal form) $c, d ::= (I?; i) \mid i$
 Final coercions $i ::= (g; I!) \mid g$
 Middle coercions $g ::= \iota \mid c \rightarrow d \mid c \times d \mid \text{Ref } T \mid \perp$

Coercion typing

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

$$\frac{}{\iota : T \Longrightarrow T} \quad \frac{g : T \Longrightarrow I}{g; I! : T \Longrightarrow \star} \quad \frac{i : I \Longrightarrow T}{I?; i : \star \Longrightarrow T} \quad \frac{}{\perp : T \Longrightarrow T'}$$

$$\frac{c : T'_1 \Longrightarrow T_1 \quad d : T_2 \Longrightarrow T'_2}{c \rightarrow d : T_1 \rightarrow T_2 \Longrightarrow T'_1 \rightarrow T'_2} \quad \frac{c : T_1 \Longrightarrow T'_1 \quad d : T_2 \Longrightarrow T'_2}{c \times d : T_1 \times T_2 \Longrightarrow T'_1 \times T'_2} \quad \frac{T \sqsubseteq T_2}{\text{Ref } T : \text{Ref } T_1 \Longrightarrow \text{Ref } T_2}$$

Middle coercion creation

$$\boxed{(I \Rightarrow_g I) = i}$$

$$(I \Rightarrow_g I) = \iota$$

$$(I \Rightarrow_g J) = \perp \quad \text{if } I \not\sim J$$

$$(T_1 \rightarrow T_2 \Rightarrow_g T'_1 \rightarrow T'_2) = (T'_1 \Rightarrow T_1) \rightarrow (T_2 \Rightarrow T'_2)$$

$$(T_1 \times T_2 \Rightarrow_g T'_1 \times T'_2) = (T_1 \Rightarrow T'_1) \times (T_2 \Rightarrow T'_2)$$

$$(\text{Ref } T \Rightarrow_g \text{Ref } T') = \text{Ref } T'$$

Coercion creation

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

$$(I \Rightarrow \star) = \iota; I! \quad (\star \Rightarrow \star) = \iota$$

$$(\star \Rightarrow I) = I?; \iota \quad (I \Rightarrow J) = I \Rightarrow_g J$$

Middle coercion composition

$$\boxed{g \circ_g g = g}$$

Final coercion composition

$$\boxed{i \circ_i c = c}$$

$$\perp \circ_g g = \perp$$

$$g \circ_g \perp = \perp$$

$$\iota \circ_g g = g$$

$$g \circ_g \iota = g$$

$$c \rightarrow d \circ_g c' \rightarrow d' = (c' \circ_g c) \rightarrow (d \circ_g d')$$

$$c \times d \circ_g c' \times d' = (c \circ_g c') \times (d \circ_g d')$$

$$\text{Ref } T \circ_g \text{Ref } T' = \perp \quad \text{if } T \not\sim T'$$

$$\text{Ref } T \circ_g \text{Ref } T' = \text{Ref } (T \sqcap T') \quad \text{if } T \sim T'$$

$$(g; I!) \circ_i (J?; i) = g \circ_i ((I \Rightarrow_g J) \circ_i i)$$

$$g \circ_i (g'; I!) = (g \circ_g g'); I!$$

$$g \circ_i g' = g \circ_g g'$$

$$\perp \circ_i i = \perp$$

$$i \circ_i \perp = \perp$$

$$\iota \circ_i i = i$$

$$i \circ_i \iota = i$$

Coercion composition

$$\boxed{c \circ c = c}$$

$$(I?; i) \circ c = I?; (i \circ_i c) \quad i \circ c = i \circ_i c$$

Type height

$$\boxed{\|T\|}$$

$$\|\text{Int}\| = 1 \quad \|\text{O}\| = 1 \quad \|\star\| = 1 \quad \|T_1 \rightarrow T_2\| = 1 + \max(\|T_1\|, \|T_2\|)$$

$$\|T_1 \times T_2\| = 1 + \max(\|T_1\|, \|T_2\|) \quad \|\text{Ref } T\| = 1 + \|T\|$$

Coercion height

$$\boxed{\|c\|}$$

$$\|\iota\| = 1 \quad \|\perp\| = 1 \quad \|\text{Ref } T\| = 1 + \|T\| \quad \|c \rightarrow d\| = 1 + \max(\|c\|, \|d\|)$$

$$\|c \times d\| = 1 + \max(\|c\|, \|d\|) \quad \|c; I!\| = \max(\|c\|, \|I!\|) \quad \|I?; c\| = \max(\|I!\|, \|c\|)$$

Fig. 1. Normal coercions and their operations

Uncoerced Values	u	$::=$	$k \mid \lambda(x : T). e \mid \langle v, v \rangle \mid a$
Values	v	$::=$	$u \mid u \langle c \rightarrow d \rangle \mid u \langle g ; ! \rangle$
Heap	μ	$::=$	$\emptyset \mid \mu(a \mapsto v : T)$
SuspendedCasts	Ψ	$::=$	$\emptyset \mid (a, T); \Psi$
Frames	F	$::=$	$\square e \mid v \square \mid \langle \square, e \rangle \mid \langle v, \square \rangle \mid \text{fst } \square \mid \text{snd } \square \mid$ $\text{ref } \square @ T \mid ! \square @ T \mid ! \square \mid \square := e @ T \mid v := \square @ T$

Fig. 2. Run-time structures for the space-efficient coercion calculus

$$\begin{aligned}
\text{apply-cast}(\langle u \langle c \rangle, d \rangle) &= \text{apply-cast}(u, (c \ ; \ d)) \\
\text{apply-cast}(a, (\text{Ref } T)) &= \text{just } ((a, T); \emptyset, a) \\
\text{apply-cast}(u, !) &= \text{just } (\emptyset, u) \\
\text{apply-cast}(u, (c \rightarrow d)) &= \text{just } (\emptyset, u \langle c \rightarrow d \rangle) \\
\text{apply-cast}(u, (g ; !)) &= \text{just } (\emptyset, u \langle g ; ! \rangle) \\
\text{apply-cast}(\langle v_1, v_2 \rangle, (c \times d)) &= \text{just } (\Psi \oplus \Psi', \langle v'_1, v'_2 \rangle) \\
&\quad \text{if } \text{apply-cast}(v_1, c) = \text{just } (\Psi, v'_1), \\
&\quad \text{apply-cast}(v_2, d) = \text{just } (\Psi', v'_2) \\
\text{apply-cast}(_, _) &= \text{nothing}
\end{aligned}$$

Fig. 3. apply-cast applies a coercion to a value and possibly returns a value and a queue of suspended casts.

process continues until there is no more suspended casts in the queue. Now the queue is empty, the expression at hand can resume reduction. This way, only values are written to the heap.

State reduction. The state reduction relation \longrightarrow contains 5 rules, 4 of them deal with casting the heap with a suspended cast from the queue. A step can be taken via rule `UPDATEHEAP` if the suspended cast (a, T) is productive, i.e. $T \sqcap \mu(a)_{\text{rtti}} \neq \mu(a)_{\text{rtti}}$, and successful, the result value of the call to `apply-cast` gets written to the heap and the RTTI is updated accordingly. On the other hand, if the cast fails, the `ERROR1` rule reduces to error. Moreover, if the cast is not productive, i.e. $T \sqcap \mu(a)_{\text{rtti}} = \mu(a)_{\text{rtti}}$, the `NOCHANGE` rule drops it from the queue. Finally, if the target type of the suspended cast is not consistent with the RTTI, the `ERROR2` rule reduces to error.

Note that the `apply-cast` function maintains space-efficiency by composing the old coercion on the value being cast and the new coercion.

Example. Now let's consider how the example program in Section 3 reduces with the dynamic semantics of λ_S^{ref} . The cast expression we want to reduce is:

$$r \langle (\text{Ref } ((\star \rightarrow \text{Int}) \times \text{Ref } ((\text{Int} \rightarrow \text{Int}) \times \star))) \rangle$$

and the heap is:

$$a \mapsto \langle f \langle ! ; (\star \rightarrow \star) ! \rangle, a \langle ! ; (\text{Ref } (\star \times \star)) ! \rangle \rangle : \star \times \star$$

Reduction will proceed by applying rule `CAST/SUCCEED` reducing the expression to the address a and putting the cast in the queue. The queue Ψ will be $(a, (\star \rightarrow \text{Int}) \times \text{Ref } ((\text{Int} \rightarrow \text{Int}) \times \star)), \emptyset$. Next, rule `UPDATEHEAP` will be used to pop the cast from the queue, apply it to the value in the heap, write the new value and the new RTTI to the heap, and add the inner reference cast to the queue. The heap will be updated as follows:

$$a \mapsto \langle f \langle (\text{Int} \rightarrow (\text{Int}? ; !)), a \rangle : (\star \rightarrow \text{Int}) \times \text{Ref } ((\text{Int} \rightarrow \text{Int}) \times \star) \rangle$$

and the queue will be $(a, (\text{Int} \rightarrow \text{Int}) \times \star), \emptyset$. Finally, reduction will apply `UPDATEHEAP` to perform the cast in the queue. The final heap will be:

Pure reduction rules

$$e \longrightarrow_e^\circ e$$

$$\begin{array}{ll} \text{fst } \langle v_1, v_2 \rangle \longrightarrow_e^\circ v_1 & (\lambda(x : T). M) v \longrightarrow_e^\circ [x := v]M \\ \text{snd } \langle v_1, v_2 \rangle \longrightarrow_e^\circ v_2 & (u \langle c \rightarrow d \rangle) v \longrightarrow_e^\circ (u(v \langle c \rangle)) \langle d \rangle \end{array}$$

Reference operations reduction rules

$$e, \mu \longrightarrow_e^\circ e, \mu, \Psi$$

$$\begin{array}{ll} \text{ref } v@T, \mu \longrightarrow_e^\circ a, \mu(a \mapsto v : T), \emptyset & \text{if } a \notin \text{dom}(\mu) \quad (\text{ALLOC}) \\ !a, \mu \longrightarrow_e^\circ \mu(a)_{\text{val}}, \mu, \emptyset & (\text{READ}) \\ !a@T, \mu \longrightarrow_e^\circ \mu(a)_{\text{val}}(\mu(a)_{\text{rtti}} \Rightarrow T), \mu, \emptyset & (\text{DYNREAD}) \\ a := v, \mu \longrightarrow_e^\circ \text{Unit}, \mu(a \mapsto v : \mu(a)_{\text{rtti}}), \emptyset & (\text{UPDATE}) \\ a := v@T, \mu \longrightarrow_e^\circ \text{Unit}, \mu(a \mapsto v' : \mu(a)_{\text{rtti}}), \Psi & \text{if } \text{apply-cast}(v, (T \Rightarrow \mu(a)_{\text{rtti}})) = \text{just}(\Psi, v') \quad (\text{DYNUPDATE}) \\ a := v@T, \mu \longrightarrow_e^\circ \text{error}, \mu, \emptyset & \text{if } \text{apply-cast}(v, (T \Rightarrow \mu(a)_{\text{rtti}})) = \text{nothing} \quad (\text{DYNUPDATEFAIL}) \end{array}$$

Program reduction rules

$$e, \mu \longrightarrow_e e, \mu, \Psi$$

$$\begin{array}{lll} \text{SWITCH} \frac{M, \mu \longrightarrow_e^{\text{ccd}} N, \mu, \Psi}{M, \mu \longrightarrow_e^{\text{cca}} N, \mu, \Psi} & \text{PURE} \frac{M \longrightarrow_e^\circ N}{M, \mu \longrightarrow_e^{\text{ccd}} N, \mu, \emptyset} & \text{REF} \frac{M, \mu \longrightarrow_e^\circ N, \mu', \Psi}{M, \mu \longrightarrow_e^{\text{ccd}} N, \mu', \Psi} \\ \text{CAST/SUCCEED} \frac{\text{apply-cast}(v, c) = \text{just}(\Psi, v')}{v \langle c \rangle, \mu \longrightarrow_e^{\text{ccd}} v', \mu, \Psi} & \text{CAST/FAIL} \frac{\text{apply-cast}(v, c) = \text{nothing}}{v \langle c \rangle, \mu \longrightarrow_e^{\text{ccd}} \text{error}, \mu, \emptyset} & \\ \text{CAST/COMPOSE} \frac{}{M \langle c \rangle \langle d \rangle, \mu \longrightarrow_e^{\text{ccd}} M \langle c \ ; \ d \rangle, \mu, \emptyset} & & \\ \text{CONG} \frac{M, \mu \longrightarrow_e^{\text{cca}} N, \mu', \Psi}{F[M], \mu \longrightarrow_e^{\text{ccd}} F[N], \mu', \Psi} & \text{CONGERR} \frac{}{F[\text{error}], \mu \longrightarrow_e^{\text{ccd}} \text{error}, \mu, \emptyset} & \\ \text{CONGCAST} \frac{M, \mu \longrightarrow_e^{\text{ccd}} N, \mu', \Psi}{M \langle c \rangle, \mu \longrightarrow_e^{\text{cca}} N \langle c \rangle, \mu', \Psi} & \text{CONGCASTE} \frac{}{\text{error} \langle c \rangle, \mu \longrightarrow_e^{\text{cca}} \text{error}, \mu, \emptyset} & \end{array}$$

State reduction rules

$$e, \mu, \Psi \longrightarrow_e e, \mu, \Psi$$

$$\begin{array}{ll} \text{PROGREduce} \frac{e, \mu \longrightarrow_e e', \mu, \Psi}{e, \mu, \emptyset \longrightarrow_e e', \mu, \Psi} & \text{UPDATEHEAP} \frac{T' = T \sqcap \mu(a)_{\text{rtti}} \quad T' \neq \mu(a)_{\text{rtti}}}{\text{apply-cast}(\mu(a)_{\text{val}}, (\mu(a)_{\text{rtti}} \Rightarrow T')) = \text{just}(\Psi', v')} \\ & e, \mu, (a, T); \Psi \longrightarrow_e e, \mu(a \mapsto v' : T'), \Psi \oplus \Psi' \\ \text{NOCHANGE} \frac{T \sqcap \mu(a)_{\text{rtti}} = \mu(a)_{\text{rtti}}}{e, \mu, (a, T); \Psi \longrightarrow_e e, \mu, \Psi} & \text{ERROR1} \frac{T' = T \sqcap \mu(a)_{\text{rtti}} \quad T' \neq \mu(a)_{\text{rtti}}}{\text{apply-cast}(\mu(a)_{\text{val}}, (\mu(a)_{\text{rtti}} \Rightarrow T')) = \text{nothing}} \\ & e, \mu, (a, T); \Psi \longrightarrow_e \text{error}, \mu, \Psi \\ \text{ERROR2} \frac{T \approx \mu(a)_{\text{rtti}}}{e, \mu, (a, T); \Psi \longrightarrow_e \text{error}, \mu, \Psi} & \end{array}$$

Fig. 4. Dynamic semantics for the space-efficient coercion calculus

$$a \mapsto \langle f \langle \langle (\iota; \text{Int!}) \rightarrow (\text{Int?}; \iota) \rangle \rangle, a \rangle : (\text{Int} \rightarrow \text{Int}) \times \text{Ref} ((\text{Int} \rightarrow \text{Int}) \times \star)$$

Now the queue is empty, expression reduction will resume and the program will evaluate to 42.

6.1 Type Safety

A lot of the definitions and lemmas used to prove type safety are shared between λ_C^{ref} and λ_S^{ref} , so to highlight the important parts of the proof and to avoid duplication, they will be presented in the appendix only.

Definition 6.1 (Well-defined suspended casts queues). A queue of suspended casts Ψ is well-defined with respect to heap typing Σ , written $\Sigma \vdash \Psi$, iff $\forall a, T. \forall (a, T) \in \Psi$ implies $a \in \Sigma$.

The queue of suspended casts plays a role in typing as well as in reduction. More specifically, the typing of values in the heap depends on both the heap typing and the queue of suspended casts. We combine both using the following function, merge.

Definition 6.2 (Merging Heap Typing and Suspended Casts Queue).

$$\begin{aligned} \text{merge}(\Sigma, \emptyset) &= \Sigma \\ \text{merge}(\Sigma, ((a, T); \Psi)) &= \text{merge}((\Sigma(a \mapsto (T \sqcap \Sigma(a)))), \Psi) \quad \text{if } \Sigma(a) \sim T \\ \text{merge}(\Sigma, ((a, T); \Psi)) &= \text{merge}(\Sigma, \Psi) \quad \text{if } \Sigma(a) \approx T \end{aligned}$$

Definition 6.3 (Well-typed heaps). A heap μ is well-typed with respect to heap typing Σ and queue of suspended casts Ψ , written $\Sigma \mid \Psi \vdash \mu$, iff $\forall a, T. \Sigma(a) = T$ implies $\exists v$ s.t. $\text{merge}(\Sigma, \Psi) \mid \emptyset \vdash v : T$ and $\mu(a) = v : T$.

LEMMA 6.4 (POPULATING THE QUEUE). *If $\Sigma \vdash \Psi$, then, $\forall \Psi'$, s.t. $\Sigma \vdash \Psi'$, $\text{merge}(\Sigma, (\Psi \oplus \Psi')) \sqsubseteq_p \text{merge}(\Sigma, \Psi)$*

From Lemma C.4, we have the following corollary.

COROLLARY 6.5 (HEAP CAST). *If $\Sigma \vdash (a, T'); \Psi$ and $\Sigma \mid (a, T'); \Psi \vdash \mu$ and $(\text{merge}(\Sigma, \Psi)) \mid \emptyset \vdash v : T$ and $\mu(a) = v : T$, then address a can be cast to type $\text{Ref}(T \sqcap T')$ such that $\Sigma(a \mapsto (T \sqcap T')) \mid \Psi \oplus \Psi' \vdash \mu(a \mapsto v' : (T \sqcap T'))$ if such v' and Ψ' exist.*

PROOF. By cases on the result of $(\text{apply-cast}(v, (T \Rightarrow (T \sqcap T'))))$:

nothing : We conclude with a cast error.

just (Ψ', v') : Let $\Sigma' = \Sigma(a \mapsto (T \sqcap T'))$. From $(T \sqcap T') \sqsubseteq T$ and $\Sigma \mid (a, T'); \Psi \vdash \mu$ we have $(T \sqcap T') \sqsubseteq \Sigma(a)$ which implies $\Sigma' \sqsubseteq_p \Sigma$. Also, by the definition of merge, we have $\text{merge}(\Sigma', \Psi) = \text{merge}(\Sigma, ((a, T'); \Psi))$, so from that and Lemma 6.4, we have that $\text{merge}(\Sigma', (\Psi \oplus \Psi')) \sqsubseteq_p \text{merge}(\Sigma', \Psi) \sqsubseteq_p \text{merge}(\Sigma, ((a, T'); \Psi))$. Next, Lemma C.4 is applied to all values in the heap and to v' , so we get $(\text{merge}(\Sigma', (\Psi \oplus \Psi'))) \mid \emptyset \vdash v' : (T \sqcap T')$, thus $\Sigma' \mid \Psi \oplus \Psi' \vdash \mu(a \mapsto v' : (T \sqcap T'))$. □

LEMMA 6.6 (PROGRESS WITH SUSPENDED CASTS). *If $\Sigma \vdash \Psi$ and $\Sigma \mid \Psi \vdash \mu$ and $\Psi \neq \emptyset$ and $\text{merge}(\Sigma, \Psi) \mid \emptyset \vdash M : T$, then $\exists \mu' \text{ s.t. } M, \mu, \Psi \longrightarrow M, \mu', \Psi'$*

PROOF. Because Ψ is not empty, $\exists a, T, \Psi' \text{ s.t. } \Psi = (a, T); \Psi'$. By cases on how T and $\mu(a)_{\text{rtti}}$ are related, a step is taken as follows:

$T \sim \mu(a)_{\text{rtti}}$ and $T \sqcap \mu(a)_{\text{rtti}} \neq \mu(a)_{\text{rtti}}$: By cases on the result of $(\text{apply-cast}(\mu(a)_{\text{val}}, (\mu(a)_{\text{rtti}} \Rightarrow T)))$:

nothing: A step is taken via ERROR1.

just (v, Ψ'') : A step is taken via UPDATEHEAP.

$T \sim \mu(a)_{\text{rtti}}$ **and** $T \sqcap \mu(a)_{\text{rtti}} = \mu(a)_{\text{rtti}}$: A step is taken via NOCHANGE.
 $T \approx \mu(a)_{\text{rtti}}$: A step is taken via ERROR2.

□

If the queue of suspended casts is empty, then the progress lemma is standard.

LEMMA 6.7 (PROGRESS WITH NO SUSPENDED CASTS). *If $\Sigma \mid \emptyset \vdash \mu$ and $\Sigma \mid \emptyset \vdash M : T$, then either*

- (1) M is a value, or
- (2) $M = \text{error}$, or
- (3) $\exists \Psi, N, \mu'$ s.t. $M, \mu, \emptyset \longrightarrow N, \mu', \Psi$

From Lemmas 6.6 and 6.7, we assemble the full proof of progress by cases on whether the queue of suspended casts is empty:

COROLLARY 6.8 (PROGRESS). *If $\Sigma \vdash \Psi$ and $\Sigma \mid \Psi \vdash \mu$ and $\text{merge}(\Sigma, \Psi) \mid \emptyset \vdash M : T$, then either*

- (1) M is a value, or
- (2) $M = \text{error}$, or
- (3) $\exists N, \mu'$ s.t. $M, \mu, \Psi \longrightarrow N, \mu', \Psi'$

LEMMA 6.9 (TYPE PRESERVATION). *If $\Sigma \vdash \Psi$ and $\Sigma \mid \Psi \vdash \mu$ and $\text{merge}(\Sigma, \Psi) \mid \emptyset \vdash M : T$ and $M, \mu, \Psi \longrightarrow N, \mu', \Psi'$, then $\exists \Sigma'$ s.t. $\Sigma' \vdash \Psi'$ and $\text{merge}(\Sigma', \Psi') \mid \emptyset \vdash N : T$ and $\Sigma' \mid \Psi' \vdash \mu'$ and $\Sigma' \sqsubseteq_{p/e} \Sigma$ and $\text{merge}(\Sigma', \Psi') \sqsubseteq_{p/e} \text{merge}(\Sigma, \Psi)$*

From Corollary 6.8 and Lemma 6.9, we prove type safety.

THEOREM 6.10 (TYPE SAFETY). *If $\Sigma \vdash \Psi$ and $\Sigma \mid \Psi \vdash \mu$ and $\text{merge}(\Sigma, \Psi) \mid \emptyset \vdash e : T$, then either*

- (1) e diverges, or
- (2) $e = \text{error}$, or
- (3) $\exists v, \Sigma', \mu'$ s.t. $e, \mu, \Psi \longrightarrow^* v, \mu', \emptyset$ and $\Sigma' \mid \emptyset \vdash v : T$ and $\Sigma' \mid \emptyset \vdash \mu'$

The semantics and proof of type safety is fully mechanized in Agda proof assistant and it is available in the following URL:

<https://github.com/deyaaeldeen/monotonic>.

7 RELATED WORK

We conjecture that the space-efficient semantics for λ_S^{ref} that we define in this paper yields the same observable behavior as that of the non-space-efficient semantics λ_C^{ref} , which is the semantics of [Siek et al. 2015c] with some corrections. We define λ_C^{ref} in Appendix B.

In the following we discuss some additional related work regarding both monotonicity and space efficiency.

Monotonicity. Fährdrich and Leino [2003] present a statically checkable typestate system where objects may flow from less restrictive to more restrictive typestates, but not vice versa. Monotonic references is different in that they require run-time checks in partially typed code regions to maintain the monotonicity invariant.

Bierman et al. [2010] provide formal operational semantics for C^\sharp with the dynamic type added, named FC_4^\sharp . It uses an RTTI-based approach and subtype checks to implement casts. However, their approach is for a class-based object-oriented language as opposed to ML-like structural reference type that our work is based on. In particular, their work does not use the consistency relation and instead extends the implicit type conversion relation such that a type can be cast to dynamic if it can be cast to object.

Rastogi et al. [2015] and Swamy et al. [2014] integrate static and dynamic typing in the context of TypeScript with the TS[★] and Safe TypeScript languages. Both use a notion of monotonicity in the heap where each object is instrumented with RTTI, but with respect to subtyping, treating [★] as a universal supertype, instead of with respect to the precision relation. Furthermore, they are different from monotonic references in that they inherit the overhead of dynamic typing because they compile to JavaScript. Moreover, in Safe TypeScript, the object's RTTI gets updated when a read field operation is performed. In monotonic references, reading from the heap does not change the RTTI.

Space efficiency. Siek and Wadler [2010] represent casts by threesomes and provide a composition operator of threesomes that uses the greatest lower bound function on types. Furthermore, they provide a compilation from threesomes to coercions and show how to compose them based on threesomes composition. Siek and Garcia [2012] define a space-efficient abstract machine for the gradually-typed lambda calculus based on coercions. Greenberg [2015] Provides a composition operator for manifest contracts, contracts that check stronger properties than simple types. However, reference types are not addressed in all aforementioned work.

8 CONCLUSION

Monotonic references is a design for gradually-typed mutable references that improves over the traditional design as they incur zero-overhead for reference operations in statically-typed code regions while maintaining the full expressiveness of gradual typing. However, prior work on monotonic references presented reduction semantics that writes cast expressions to the heap that reduce using a small-step reduction relation. It is not straightforward to implement such a process efficiently in a compiler and runtime system for gradually-typed languages. Furthermore, earlier work did not guarantee space-efficiency, a key property to ensure runtime efficiency in implementations of gradually-typed languages.

In this paper, we present novel dynamic semantics for monotonic references that is space-efficient and that only writes values to the heap. The former is accomplished by identifying a normal form for coercions with monotonic references and by extending the composition function accordingly. Moreover, in our dynamic semantics, we apply a cast to a value using a simple recursive function that suspends inner casts of references by putting them in a queue. This process is straightforward to implement in a runtime system for a gradually-typed language. Finally, we provide a type safety proof that is fully mechanized in Agda.

ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under Grant No. 1763922.

REFERENCES

- Christopher Anderson and Sophia Drossopoulou. 2003. BabyJ - From Object Based to Class Based Programming via Types. In *WOOD '03*, Vol. 82. Elsevier.
- Gavin Bierman, Erik Meijer, and Mads Torgersen. 2010. Adding dynamic types to C# (*ECOOP'10*). Springer-Verlag, 25.
- Manuel Fähndrich and K Rustan M Leino. 2003. Heap monotonic typestates. In *International Workshop on Aliasing, Confinement and Ownership in object-oriented programming (IWACO)*.
- Robert Bruce Findler and Matthias Felleisen. 2002. *Contracts for Higher-Order Functions*. Technical Report NU-CCS-02-05. Northeastern University.
- Michael Greenberg. 2015. Space-Efficient Manifest Contracts (*POPL '15*). ACM, New York, NY, USA, 181–194. <https://doi.org/10.1145/2676726.2676967>
- Jessica Gronski, Kenneth Knowles, Aaron Tomb, Stephen N. Freund, and Cormac Flanagan. 2006. Sage: Hybrid Checking for Flexible Specifications. In *Scheme and Functional Programming Workshop*. 93–104.

- Fritz Henglein. 1994. Dynamic typing: syntax and proof theory. *Science of Computer Programming* 22, 3 (June 1994), 197–230.
- David Herman, Aaron Tomb, and Cormac Flanagan. 2007. Space-Efficient Gradual Typing. In *Trends in Functional Prog. (TFP)*. XXVIII.
- David Herman, Aaron Tomb, and Cormac Flanagan. 2010. Space-efficient gradual typing. *Higher-Order and Symbolic Computation* 23, 2 (2010), 167–189.
- Andre Kuhlenschmidt, Deyaaeldeen Almahallawi, and Jeremy G. Siek. 2019. Toward Efficient Gradual Typing for Structural Types via Coercions. In *Conference on Programming Language Design and Implementation (PLDI)*. ACM.
- Fabian Muehlboeck and Ross Tate. 2017. Sound Gradual Typing is Nominally Alive and Well. *Proc. ACM Program. Lang.* 1, OOPSLA, Article 56 (Oct. 2017), 30 pages. <https://doi.org/10.1145/3133880>
- Aseem Rastogi, Nikhil Swamy, Cédric Fournet, Gavin Bierman, and Panagiotis Vekris. 2015. Safe & Efficient Gradual Typing for TypeScript (POPL '15). ACM, New York, NY, USA, 167–180. <https://doi.org/10.1145/2676726.2676971>
- Jeremy G. Siek and Ronald Garcia. 2012. Interpretations of the Gradually-Typed Lambda Calculus. In *Scheme and Functional Programming Workshop*.
- Jeremy G. Siek, Ronald Garcia, and Walid Taha. 2009. Exploring the Design Space of Higher-Order Casts. In *European Symposium on Programming (ESOP)*. 17–31.
- 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. 2015a. Blame and coercion: Together again for the first time. In *Conference on Programming Language Design and Implementation (PLDI)*.
- Jeremy G. Siek and Michael M. Vitousek. 2013. Monotonic References for Gradual Typing. *CoRR* abs/1312.0694 (2013).
- Jeremy G. Siek, Michael M. Vitousek, Matteo Cimini, and John Tang Boyland. 2015b. Refined Criteria for Gradual Typing. In *SNAPL: Summit on Advances in Programming Languages (LIPLCS: Leibniz International Proceedings in Informatics)*.
- Jeremy G. Siek, Michael M. Vitousek, Matteo Cimini, Sam Tobin-Hochstadt, and Ronald Garcia. 2015c. Monotonic References for Efficient Gradual Typing. In *European Symposium on Programming (ESOP)*.
- Jeremy G. Siek and Philip Wadler. 2010. Threesomes, with and without blame (POPL). 365–376.
- Nikhil Swamy, Cedric Fournet, Aseem Rastogi, Karthikeyan Bhargavan, Juan Chen, Pierre-Yves Strub, and Gavin Bierman. 2014. Gradual Typing Embedded Securely in JavaScript. In *ACM Conference on Principles of Programming Languages (POPL)*.
- Sam Tobin-Hochstadt and Matthias Felleisen. 2006. Interlanguage Migration: From Scripts to Programs. In *Dynamic Languages Symposium*.
- Philip Wadler and Robert Bruce Findler. 2009. Well-typed programs can't be blamed. In *European Symposium on Programming (ESOP)*. 1–16.

Expressions $e, M, N ::= k \mid \lambda x. e \mid e e \mid \langle e, e \rangle \mid \text{fst } e \mid \text{snd } e \mid \text{ref } e \mid !e \mid e := e$

Fig. 5. GTLC⁺ syntax

Types

Base Types $B ::= () \mid \text{Bool} \mid \text{Int}$
 Injectables $I, J ::= B \mid T \rightarrow T \mid T \times T \mid \text{Ref } T$
 Types $T, S ::= \star \mid I$

Consistency

$$\star \sim T \quad T \sim \star \quad B \sim B \quad \frac{T_1 \sim T_2}{\text{Ref } T_1 \sim \text{Ref } T_2} \quad \frac{T_1 \sim T_3 \quad T_2 \sim T_4}{T_1 \rightarrow T_2 \sim T_3 \rightarrow T_4} \quad \frac{\boxed{T \sim T} \quad T_1 \sim T_3 \quad T_2 \sim T_4}{T_1 \times T_2 \sim T_3 \times T_4}$$

Shallow consistency

$$\star \smile T \quad T \smile \star \quad B \smile B \quad \frac{}{\text{Ref } T_1 \smile \text{Ref } T_2} \quad \frac{}{T_1 \rightarrow T_2 \smile T_3 \rightarrow T_4} \quad \frac{\boxed{T \smile T}}{T_1 \times T_2 \smile T_3 \times T_4}$$

Type preciseness

$$T \sqsubseteq \star \quad B \sqsubseteq B \quad \frac{T_1 \sqsubseteq T_2}{\text{Ref } T_1 \sqsubseteq \text{Ref } T_2} \quad \frac{T_1 \sqsubseteq T_3 \quad T_2 \sqsubseteq T_4}{T_1 \rightarrow T_2 \sqsubseteq T_3 \rightarrow T_4} \quad \frac{T_1 \sqsubseteq T_3 \quad T_2 \sqsubseteq T_4}{T_1 \times T_2 \sqsubseteq T_3 \times T_4} \quad \boxed{T \sqsubseteq T}$$

Static types

$$\int B \quad \frac{\int T}{\int \text{Ref } T} \quad \frac{\int T_1 \quad \int T_2}{\int T_1 \rightarrow T_2} \quad \frac{\int T_1 \quad \int T_2}{\int T_1 \times T_2} \quad \boxed{\int T}$$

Meet operation (greatest lower bound)

$$\star \sqcap T = T \sqcap \star = T \quad B \sqcap B = B \quad (T_1 \times T_2) \sqcap (T_3 \times T_4) = (T_1 \sqcap T_3) \times (T_2 \sqcap T_4)$$

$$\text{Ref } T_1 \sqcap \text{Ref } T_2 = \text{Ref } (T_1 \sqcap T_2) \quad (T_1 \rightarrow T_2) \sqcap (T_3 \rightarrow T_4) = (T_1 \sqcap T_3) \rightarrow (T_2 \sqcap T_4)$$

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

Fig. 6. Types and their operations

A TYPE SYSTEM OF GTLC⁺

The syntax of the surface language, GTLC⁺, is given in Figure 5. GTLC⁺ has functions and pairs in addition to references. Figure 6 describes supported types and relations and functions on them. We use the symbol \star to represent the dynamic type. Consistency and type preciseness relations are typical [Siek and Taha 2006; Siek et al. 2015b; Wadler and Findler 2009]. Shallow consistency [Siek and Wadler 2010] is needed to implement lazy error detection. The unary relation \int holds for types that do not contain \star . Finally, the \sqcap function computes the greatest lower bound of the input types.

As with any other gradually-typed language, GTLC⁺ is compiled to a statically-typed cast calculus, by inserting casts. The cast calculus uses coercions [Henglein 1994] for cast specification. Figure 7 presents the syntax and the type system. Addresses a are required by rule WT-ADDR to have a type $\text{Ref } T$ where $\Sigma(a) \sqsubseteq T$. This restriction makes sure for each address a , $\forall r = a$ and $r : \text{Ref } T$, that $\Sigma(a) \sim T$, i.e. the types of all references are always consistent with the type of the value in the heap. Expression forms for operations on references are unusual in that there are two forms for each operation, one to be used in fully static code regions and the other to be used in other regions [Siek et al. 2015c]. In particular, $!e$ and $e := e$ correspond to load and store machine

Expressions	e, M, N	$::=$	$k \mid \lambda(x : T). e \mid e e \mid \langle e, e \rangle \mid \text{fst } e \mid \text{snd } e \mid e \langle c \rangle \mid a \mid \text{ref } e@T \mid !e \mid !e@T \mid e := e \mid e := e@T \mid \text{error}$
Typing environment	Γ	$::=$	$\emptyset \mid \Gamma, x : T$
Heap typing	Σ	$::=$	$\emptyset \mid \Sigma, a : T$

Expression typing

 $\Sigma \mid \Gamma \vdash M : T$

$\frac{}{\Sigma \mid \Gamma \vdash k : \text{TYPE}(k)}$	$\frac{(x : T) \in \Gamma}{\Sigma \mid \Gamma \vdash x : T}$	$\frac{\Sigma \mid \Gamma \vdash M : T_1 \rightarrow T_2 \quad \Sigma \mid \Gamma \vdash N : T_1}{\Sigma \mid \Gamma \vdash MN : T_2}$
$\frac{\Sigma \mid \Gamma, x : T_1 \vdash M : T_2}{\Sigma \mid \Gamma \vdash \lambda(x : T_1). M : T_1 \rightarrow T_2}$	$\frac{\Sigma \mid \Gamma \vdash M : T_1 \quad \Sigma \mid \Gamma \vdash N : T_2}{\Sigma \mid \Gamma \vdash \langle M, N \rangle : T_1 \times T_2}$	
$\frac{\Sigma \mid \Gamma \vdash \langle M, N \rangle : T_1 \times T_2}{\Sigma \mid \Gamma \vdash \text{fst } M : T_1}$	$\frac{\Sigma \mid \Gamma \vdash \langle M, N \rangle : T_1 \times T_2}{\Sigma \mid \Gamma \vdash \text{snd } M : T_2}$	
$\frac{\Sigma \mid \Gamma \vdash M : T}{\Sigma \mid \Gamma \vdash \text{ref } M@T : \text{Ref } T}$	$\frac{\Sigma \mid \Gamma \vdash M : \text{Ref } T \quad \wp T}{\Sigma \mid \Gamma \vdash !M : T}$	
$\frac{\Sigma \mid \Gamma \vdash M : \text{Ref } T \quad \wp T}{\Sigma \mid \Gamma \vdash !M@T : T}$	$\frac{\Sigma \mid \Gamma \vdash M : \text{Ref } T \quad \Sigma \mid \Gamma \vdash N : T \quad \wp T}{\Sigma \mid \Gamma \vdash M := N : ()}$	
$\frac{\Sigma \mid \Gamma \vdash M : \text{Ref } T \quad \Sigma \mid \Gamma \vdash N : T \quad \wp T}{\Sigma \mid \Gamma \vdash M := N@T : ()}$		
$\text{WT-ADDR} \frac{\Sigma(a) \sqsubseteq T}{\Sigma \mid \Gamma \vdash a : \text{Ref } T}$	$\frac{}{\Sigma \mid \Gamma \vdash \text{error} : T}$	$\frac{\Sigma \mid \Gamma \vdash M : T \quad c : T \Longrightarrow T'}{\Sigma \mid \Gamma \vdash M \langle c \rangle : T'}$

Fig. 7. Syntax and type system of the coercion calculus

instructions and should be used in static code regions as they do not perform any checks. This is true because monotonic references have a unique canonical form which is just addresses, so unlike proxied references [Herman et al. 2007, 2010], there is no dynamic dispatch on whether the address is proxied. Also, by the typing rule for addresses WT-ADDR, the type of a heap cell must be as precise as or more precise than all references to it, so it must be the case that the type of the value in the heap matches the type of the statically typed reference. On the other hand, $!e@T$ and $e := e@T$ are reserved for other code regions where the reference type is not static. The type annotation T records the type of the reference. The type of the value in the heap might be more precise than T , so a cast is needed between them. Asking the programmers to use both forms would be inconvenient. Fortunately, the compiler can automatically select which form to emit. Cast insertion, defined in Figure 8, is a type-directed compilation from GTLC⁺ to the coercion calculus including compiling to the right form of a reference operation based on the type of the reference. Rules CI-SREAD and CI-SWRITE check if the reference type is static to compile the operation form to the corresponding efficient one. Rules CI-READ and CI-WRITE compile to the other forms. The CI-ALLOCATION rule compiles the allocation form to one that records the reference type. The recorded type annotation is used to initialize the run-time type information (RTTI) of the heap cell. The other rules for typing and cast insertion are standard.

Cast insertion

 $\Sigma \mid \Gamma \vdash M \hookrightarrow N : T$

$$\begin{array}{c}
\frac{}{\Sigma \mid \Gamma \vdash k \hookrightarrow k : \text{TYPE}(k)} \quad \frac{(x : T) \in \Gamma}{\Sigma \mid \Gamma \vdash x \hookrightarrow x : T} \\
\frac{\Sigma \mid x : T_1, \Gamma \vdash M \hookrightarrow N : T_2}{\Sigma \mid \Gamma \vdash \lambda(x : T_1). M \hookrightarrow \lambda(x : T_1). N : T_1 \rightarrow T_2} \\
\frac{\Sigma \mid \Gamma \vdash M \hookrightarrow M' : T_1 \rightarrow T_2 \quad \Sigma \mid \Gamma \vdash N \hookrightarrow N' : T'_1}{\Sigma \mid \Gamma \vdash MN \hookrightarrow M'(N'\langle T'_1 \Rightarrow T_1 \rangle) : T_2} \\
\frac{\Sigma \mid \Gamma \vdash M \hookrightarrow M' : \star \quad \Sigma \mid \Gamma \vdash N \hookrightarrow N' : T}{\Sigma \mid \Gamma \vdash MN \hookrightarrow (M'\langle \star \Rightarrow (\star \Rightarrow \star) \rangle)(N'\langle T \Rightarrow \star \rangle) : \star} \\
\frac{\Sigma \mid \Gamma \vdash M \hookrightarrow M' : T_1 \quad \Sigma \mid \Gamma \vdash N \hookrightarrow N' : T_2}{\Sigma \mid \Gamma \vdash \langle M, N \rangle \hookrightarrow \langle M', N' \rangle : T_1 \times T_2} \\
\frac{\Sigma \mid \Gamma \vdash M \hookrightarrow N : T_1 \times T_2}{\Sigma \mid \Gamma \vdash \text{fst } M \hookrightarrow \text{fst } N : T_1} \quad \frac{\Sigma \mid \Gamma \vdash M \hookrightarrow N : T_1 \times T_2}{\Sigma \mid \Gamma \vdash \text{snd } M \hookrightarrow \text{snd } N : T_2} \\
\frac{\Sigma \mid \Gamma \vdash M \hookrightarrow N : \star}{\Sigma \mid \Gamma \vdash \text{fst } M \hookrightarrow \text{fst } (N\langle \star \Rightarrow \star \times \star \rangle) : \star} \quad \frac{\Sigma \mid \Gamma \vdash M \hookrightarrow N : \star}{\Sigma \mid \Gamma \vdash \text{snd } M \hookrightarrow \text{snd } (N\langle \star \Rightarrow \star \times \star \rangle) : \star} \\
\text{CI-ALLOCATION} \frac{\Sigma \mid \Gamma \vdash M \hookrightarrow N : T}{\Sigma \mid \Gamma \vdash \text{ref } M \hookrightarrow \text{ref } N@T : \text{Ref } T} \quad \text{CI-SREAD} \frac{\Sigma \mid \Gamma \vdash M \hookrightarrow N : \text{Ref } T \quad \}T}{\Sigma \mid \Gamma \vdash !M \hookrightarrow !N : T} \\
\text{CI-READ} \frac{\Sigma \mid \Gamma \vdash M \hookrightarrow N : \text{Ref } T}{\Sigma \mid \Gamma \vdash !M \hookrightarrow !N@T : T} \quad \frac{\Sigma \mid \Gamma \vdash M \hookrightarrow N : \star}{\Sigma \mid \Gamma \vdash !M \hookrightarrow !(N\langle \star \Rightarrow \text{Ref } \star \rangle)@* : \star} \\
\text{CI-SWRITE} \frac{\Sigma \mid \Gamma \vdash M \hookrightarrow M' : \text{Ref } T \quad \Sigma \mid \Gamma \vdash N \hookrightarrow N' : T \quad \}T}{\Sigma \mid \Gamma \vdash M := N \hookrightarrow M' := N' : ()} \\
\text{CI-Write} \frac{\Sigma \mid \Gamma \vdash M \hookrightarrow M' : \text{Ref } T \quad \Sigma \mid \Gamma \vdash N \hookrightarrow N' : T' \quad T \sim T' \quad T \neq T'}{\Sigma \mid \Gamma \vdash M := N \hookrightarrow M' := (N'\langle T' \Rightarrow T \rangle)@T' : ()} \\
\frac{\Sigma \mid \Gamma \vdash M \hookrightarrow M' : \star \quad \Sigma \mid \Gamma \vdash N \hookrightarrow N' : T}{\Sigma \mid \Gamma \vdash M := N \hookrightarrow (M'\langle \star \Rightarrow \text{Ref } \star \rangle) := (N'\langle T' \Rightarrow \star \rangle)@* : ()}
\end{array}$$

Fig. 8. Compiling GTLC⁺ to the coercion calculus

B DYNAMIC SEMANTICS OF λ_C^{ref}

Monotonic references were first introduced in terms of an abstract machine mechanized in Isabelle/HOL [Siek and Vitousek 2013]. Later, Siek et al. [2015c] presented a reduction semantics for monotonic references based on coercions along with the addition of blame tracking. However, The latter had a few minor problems that λ_C^{ref} fix.

- The grammar for casted values¹ does not allow casts on other casted values. However, the rule CASTREF1 in the cast reduction relation reads a casted value from the heap and then

¹ λ_C^{ref} uses the term *delayed casts* instead to avoid confusion with values that have casts on them

Coercions	$c, d ::= \iota \mid I! \mid I? \mid c \rightarrow d \mid c \times d \mid \text{Ref } T \mid \perp^{I,J}$
Inert coercions	$c \uparrow ::= I! \mid c \rightarrow d$
Active coercions	$c \downarrow ::= \iota \mid I? \mid c \times d \mid \text{Ref } T \mid \perp^{I,J}$

Coercion typing

$c : T \Longrightarrow T$

$$\begin{array}{c}
\frac{}{\iota : T \Longrightarrow T} \quad \frac{}{I! : I \Longrightarrow \star} \quad \frac{}{I? : \star \Longrightarrow I} \quad \frac{I \approx J}{\perp^{I,J} : I \Longrightarrow J} \\
\frac{c : T'_1 \Longrightarrow T_1 \quad d : T_2 \Longrightarrow T'_2}{c \rightarrow d : T_1 \rightarrow T_2 \Longrightarrow T'_1 \rightarrow T'_2} \quad \frac{c : T_1 \Longrightarrow T'_1 \quad d : T_2 \Longrightarrow T'_2}{c \times d : T_1 \times T_2 \Longrightarrow T'_1 \times T'_2} \quad \frac{}{\text{Ref } T' : \text{Ref } T \Longrightarrow \text{Ref } T'}
\end{array}$$

Coercion creation

$(T \Longrightarrow T) = c$

$$\begin{array}{ll}
(T \Longrightarrow T) = \iota & (T_1 \rightarrow T_2 \Longrightarrow T'_1 \rightarrow T'_2) = (T_1 \Longrightarrow T_1) \rightarrow (T_2 \Longrightarrow T'_2) \\
(I \Longrightarrow \star) = I! \quad (\star \Longrightarrow I) = I? & (T_1 \times T_2 \Longrightarrow T'_1 \times T'_2) = (T_1 \Longrightarrow T'_1) \times (T_2 \Longrightarrow T'_2) \\
(I \Longrightarrow J) = \perp^{I,J} \quad \text{if } I \not\sim J & (\text{Ref } T \Longrightarrow \text{Ref } T') = \text{Ref } T'
\end{array}$$

Fig. 9. Simple coercions and their operations

wraps it in another cast expression that is written to the heap. λ_C^{ref} fixes this issue by changing the grammar of delayed casts to allow casts on delayed casts (Figure 10).

- A casted value could reduce to an expression that has an inner error using a congruence rule. Such expression is not a casted value and can not be written to the heap. The state reduction relation has a rule to handle reduction to an error expression but does not have one to handle reduction to expressions that have an inner error. λ_C^{ref} fixes this issue by updating the ERROR rule to make sure the result delayed cast could be decomposed into a context and an error expression.
- The type preservation lemma, the second item in Lemma 2, relates the old heap typing and the new heap typing with the precision relation on heap typing (Definition 1). However, this does not account for the fact that the heap could grow by creating new references. λ_C^{ref} fixes this issue by recognizing that the heap typing could move along two orthogonal dimensions, it could become either more precise (Definition C.1) or larger (Definition C.2). The type preservation lemma is updated accordingly (Lemma C.18).

The variant of the semantics of $\lambda_C^{\text{ref}2}$ and associated type safety proof is fully mechanized in Agda proof assistant. Figure 9 presents simple coercions and their operations. Inert coercions, denoted $c \uparrow$, form values without checking at the time of the cast. Injection and function coercions are inert coercions. On the other hand, all other coercions do not form values and are called active coercions, denoted $c \downarrow$.

Figure 10 presents the definition of values, delayed casts, heaps, evolving heaps, and frames. Figures 11 and 12 presents the reduction rules for the cast calculus with simple coercions. The symbol \odot indicates a reduction relation that does not relate heaps. The symbol \circ indicates a reference reduction relation.

C TYPE SAFETY FOR λ_C^{ref}

We prove type safety for the corrected version of λ_C^{ref} (defined in Appendix B).

²The difference is state reduction relation has more verbose rules in the mechanized version

Run-Time Structures

Values	$v ::= k \mid \lambda(x : T). e \mid \langle v, v \rangle \mid a \mid v \langle c \uparrow \rangle$
Delayed Casts	$dc ::= v \mid dc \langle c \rangle \mid \langle dc, dc \rangle$
Heap	$\mu ::= \emptyset \mid \mu(a \mapsto v : T)$
Evolving Heap	$\nu ::= \mu \mid \nu(a \mapsto dc : T)$
Frames	$F ::= \square e \mid \nu \square \mid \langle \square, e \rangle \mid \langle \nu, \square \rangle \mid \text{fst } \square \mid \text{snd } \square \mid \square \langle c \rangle \mid$ $\text{ref } \square @ T \mid !\square @ T \mid !\square \mid \square := e @ T \mid \nu := \square @ T$
Evaluation Contexts	$\mathcal{E} ::= \emptyset \mid F, \mathcal{E}$

Fig. 10. Run-Time structures for the simple coercion calculus

Pure coercion reduction rules

$$e \longrightarrow_c^\circ e$$

$$\begin{array}{l} v \langle i \rangle \longrightarrow_c^\circ v \\ v \langle ! \rangle \langle J ? \rangle \longrightarrow_c^\circ v \langle I \Rightarrow J \rangle \end{array} \quad \begin{array}{l} v \langle \perp^{I,J} \rangle \longrightarrow_c^\circ \text{error} \\ \langle v_1, v_2 \rangle \langle c \times d \rangle \longrightarrow_c^\circ \langle v_1 \langle c \rangle, v_2 \langle d \rangle \rangle \end{array}$$

Cast reduction rules

$$e, v \longrightarrow_c e, v$$

$$\begin{array}{c} \text{PURECAST} \frac{M \longrightarrow_c^\circ N}{M, v \longrightarrow_c N, v} \\ \text{CASTREF1} \frac{v(a) = dc : T_1 \quad T_1 \sim T_2 \quad T_1 \sqcap T_2 \neq T_1}{a \langle \text{Ref } T_2 \rangle, v \longrightarrow_c a, v(a \mapsto dc \langle T_1 \Rightarrow (T_1 \sqcap T_2) \rangle) : T_1 \sqcap T_2} \\ \text{CASTREF2} \frac{v(a)_{\text{rtti}} \sim T \quad v(a)_{\text{rtti}} \sqcap T = v(a)_{\text{rtti}}}{a \langle \text{Ref } T \rangle, v \longrightarrow_c a, v} \quad \text{CASTREF3} \frac{v(a)_{\text{rtti}} \approx T}{a \langle \text{Ref } T \rangle, v \longrightarrow_c \text{error}, v} \\ \text{CONGC} \frac{M, v \longrightarrow_c N, v'}{F[M], v \longrightarrow_c F[N], v'} \quad \text{CONGERRC} \frac{}{\mathcal{E}[\text{error}], v \longrightarrow_c \text{error}, v} \end{array}$$

Fig. 11. Dynamic semantics for casts in the simple coercion calculus

We first recognize that the heap typing can move along two orthogonal dimensions, it could become either more precise (Definition C.1) or its domain can grow (Definition C.2)³. Combining both relations results in one, $\sqsubseteq_{p/e}$, that fully captures how the heap typing changes after each step, called Heap Typing Progress (Definition C.3).

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

Definition C.2 (Extension relation on heap typings). $\Sigma \sqsubseteq_e \Sigma'$ iff $\text{dom}(\Sigma) \subseteq \text{dom}(\Sigma')$ and $\Sigma(a) = \Sigma'(a)$.

Definition C.3 (Progress relation for heap typings). $\Sigma' \sqsubseteq_{p/e} \Sigma$ if $\Sigma' \sqsubseteq_p \Sigma$ or $\Sigma \sqsubseteq_e \Sigma'$.

Lemma C.4 asserts that expression typing is preserved when moving to a more precise heap typing. Similarly, Lemma C.5 asserts that expression typing is also preserved when moving to a larger heap typing. Corollary C.6 combines both lemmas and asserts that expression typing is preserved when the heap typing changes with respect to the $\sqsubseteq_{p/e}$ relation.

³We ignore garbage collection and assume heaps can not become smaller

Pure reduction rules

$$e \xrightarrow{\circledast_e} e$$

$$\begin{array}{ll} \text{fst } \langle v_1, v_2 \rangle \xrightarrow{\circledast_e} v_1 & (\lambda(x : T). M) v \xrightarrow{\circledast_e} [x := v]M \\ \text{snd } \langle v_1, v_2 \rangle \xrightarrow{\circledast_e} v_2 & (v_1 \langle c \rightarrow d \rangle) v_2 \xrightarrow{\circledast_e} (v_1 (v_2 \langle c \rangle)) \langle d \rangle \end{array}$$

Reference operations reduction rules

$$e, \mu \xrightarrow{\circ_e} e, v$$

$$\begin{array}{ll} \text{ref } v@T, \mu \xrightarrow{\circ_e} a, \mu(a \mapsto v : T) & \text{if } a \notin \text{dom}(\mu) & (\text{NSEALLOC}) \\ !a, \mu \xrightarrow{\circ_e} \mu(a)_{\text{val}}, \mu & & (\text{READ}) \\ !a@T, \mu \xrightarrow{\circ_e} \mu(a)_{\text{val}} \langle \mu(a)_{\text{rtti}} \Rightarrow T \rangle, \mu & & (\text{DYNREAD}) \\ a := v, \mu \xrightarrow{\circ_e} \text{Unit}, \mu(a \mapsto v : \mu(a)_{\text{rtti}}) & & (\text{UPDATE}) \\ a := v@T, \mu \xrightarrow{\circ_e} \text{Unit}, \mu(a \mapsto v \langle T \Rightarrow \mu(a)_{\text{rtti}} \rangle : \mu(a)_{\text{rtti}}) & & (\text{DYNUPDATE}) \end{array}$$

State reduction rules

$$e, v \longrightarrow e, v$$

$$\begin{array}{ll} \text{PURE} \frac{M \xrightarrow{\circledast_e} N}{M, \mu \longrightarrow N, \mu} & \text{MONO} \frac{M, \mu \xrightarrow{\circ_e} N, v}{M, \mu \longrightarrow N, v} \\ \text{CONG} \frac{M, v \longrightarrow N, v'}{F[M], \mu \longrightarrow F[N], v} & \text{CONGERR} \frac{}{\mathcal{E}[\text{error}], \mu \longrightarrow \text{error}, v} \\ \text{CAST} \frac{e, \mu \longrightarrow_c e', v}{e, \mu \longrightarrow e', v} & \text{ERROR} \frac{v(a) = dc : T \quad dc, v \longrightarrow_c \mathcal{E}[\text{error}], v'}{e, v \longrightarrow \text{error}, v'} \\ & \text{NORTTICHANGE} \frac{v(a) = dc : T \quad dc, v \longrightarrow_c dc', v' \quad v'(a)_{\text{rtti}} = T}{e, v \longrightarrow e, v'(a \mapsto dc' : T)} \\ & \text{RTTICHANGED} \frac{v(a) = dc : T \quad dc, v \longrightarrow_c dc', v' \quad v'(a)_{\text{rtti}} \neq T}{e, v \longrightarrow e, v'} \end{array}$$

Fig. 12. Dynamic semantics for the simple coercion calculus

LEMMA C.4 (STRENGTHENING WRT. THE HEAP TYPING PRECISION). *If $\Sigma \mid \Gamma \vdash e : T$ and $\Sigma' \sqsubseteq_p \Sigma$, then $\Sigma' \mid \Gamma \vdash e : T$.*

PROOF. For the case of addresses: From $\Sigma' \sqsubseteq_p \Sigma$ and the WT-ADDR rule and transitivity of \sqsubseteq , we have $\Sigma'(a) \sqsubseteq T$. Therefore $\Sigma' \mid \Gamma \vdash a : T$. \square

LEMMA C.5 (WEAKENING WRT. THE HEAP TYPING EXTENSION). *If $\Sigma \mid \Gamma \vdash e : T$ and $\Sigma \sqsubseteq_e \Sigma'$, then $\Sigma' \mid \Gamma \vdash e : T$.*

PROOF. For the case of addresses: we have $\Sigma'(a) \sqsubseteq T$ from $\Sigma \sqsubseteq_e \Sigma'$ and the WT-ADDR rule. Therefore $\Sigma' \mid \Gamma \vdash a : T$. \square

COROLLARY C.6 (WEAKENING WRT. THE HEAP TYPING PROGRESS). *If $\Sigma \mid \Gamma \vdash e : T$ and $\Sigma' \sqsubseteq_{p/e} \Sigma$, then $\Sigma' \mid \Gamma \vdash e : T$.*

PROOF. By applying Lemmas C.4 and C.5 in the cases of \sqsubseteq_p and \sqsubseteq_e respectively. \square

Heaps are well-typed when the delayed casts inside respect heap typing.

Definition C.7 (Well-typed heaps). A heap v is well-typed with respect to heap typing Σ , written $\Sigma \vdash v$, iff $\forall a, T. \Sigma(a) = T$ implies $\exists dc$ s.t. $\Sigma \mid \emptyset \vdash dc : T$ and $v(a) = dc : T$.

From Lemma C.4, we have the following corollary.

COROLLARY C.8 (HEAP CAST). *If $\Sigma \vdash v$ and $v(a) = dc : T$ and $T' \sqsubseteq T$, then a can be cast to $\text{Ref } T'$ such that $\Sigma(a \mapsto T') \vdash v(a \mapsto (dc \langle T \Rightarrow T' \rangle)) : T'$.*

PROOF. Let $\Sigma' = \Sigma(a \mapsto T')$ and $dc' = dc \langle T \Rightarrow T' \rangle$. From $T' \sqsubseteq T$ and $\Sigma \vdash v$ we have $T' \sqsubseteq \Sigma(a)$ which implies $\Sigma' \sqsubseteq_p \Sigma$. Next, Lemma C.4 is applied to all delayed casts in the heap and to dc' , so we get $\Sigma' \mid \emptyset \vdash dc' : T'$, thus $\Sigma' \vdash v(a \mapsto dc' : T')$. \square

Similarly, from Lemma C.5, we have the following corollary.

COROLLARY C.9 (HEAP EXTENSION). *If $\Sigma \vdash \mu$ and $\Sigma \mid \emptyset \vdash v : T$ and $a \notin \text{dom}(\Sigma)$, then v can be added at a such that $\Sigma(a \mapsto T) \vdash \mu(a \mapsto v : T)$.*

PROOF. Let $\Sigma' = \Sigma(a \mapsto T)$. We have $\Sigma \sqsubseteq_e \Sigma'$ from $a \notin \text{dom}(\Sigma)$. Next, Lemma C.5 is applied to all values in the heap and to v to get $\Sigma' \mid \emptyset \vdash v : T$. We add v to the heap at address a , thus $\Sigma' \vdash \mu(a \mapsto v : T)$. \square

Lemma C.10 shows that coercions are either inert or active.

LEMMA C.10 (INERT AND ACTIVE COERCIONS). $\forall c : T \Rightarrow T', c \uparrow$ or $c \downarrow$

Active coercions do not create values, so it better be the case that they can take a step. Lemma C.11 asserts this fact:

LEMMA C.11 (ACTIVE CAST PROGRESS). *If $\Sigma \vdash v$ and $\Sigma \mid \emptyset \vdash v : T$ and $c : T \Rightarrow T'$ and $c \downarrow$, then $\exists e, v'$ s.t. $v \langle c \rangle, v \longrightarrow_c e, v'$*

PROOF. By case analysis on $c \downarrow$. In the case of monotonic reference coercion, v must be some address a , and $v \langle \text{Ref } T' \rangle$ takes a step via \longrightarrow_c . In particular:

CASTREF1 if $T' \sim v(a)_{\text{rtti}}$ and $T' \sqcap v(a)_{\text{rtti}} \neq v(a)_{\text{rtti}}$.

CASTREF2 if $T' \sim v(a)_{\text{rtti}}$ and $T' \sqcap v(a)_{\text{rtti}} = v(a)_{\text{rtti}}$.

CASTREF3 if $T' \approx v(a)_{\text{rtti}}$.

\square

The following lemma proves that a delayed cast can take step if it is not a value.

LEMMA C.12 (REDUCIBLE DELAYED CAST PROGRESS). *If $\Sigma \vdash v$ and $\Sigma \mid \emptyset \vdash dc : T$ and dc is not a value, then $\exists e, v'$ s.t. $dc, v \longrightarrow_c e, v'$*

PROOF. By case analysis on dc .

Case v Contradiction by the premise that dc is not a value.

Case $dc' \langle c \rangle$ By cases on whether dc' is a value

Case yes We know that dc is not a value, so it must be the case that c is neither an injection or a function coercion, so Lemma C.11 is applied.

Case no The induction hypothesis is applied to dc and a step is taken via CONG_c .

Case $\langle dc_1, dc_2 \rangle$ It must be the case that either or both projections are not values, so the induction hypothesis is applied to first projection that is not a value and a step is taken via CONG_c .

\square

However this is not true for arbitrary cast expressions because inert coercions form values.

COROLLARY C.13 (CAST PROGRESS). *If $\Sigma \vdash v$ and $\Sigma \mid \emptyset \vdash v : T$ and $c : T \Longrightarrow T'$, then either*

- (1) $v\langle c \rangle$ is a value, or
- (2) $\exists e, v'$ s.t. $v\langle c \rangle, v \longrightarrow_c e, v'$

PROOF. By Lemma C.10, c is either $c \uparrow$ or $c \downarrow$. In the case of the former, we conclude that $v\langle c \rangle$ is a value. Otherwise, Lemma C.11 is applied. \square

If a delayed cast took a reduction step, the result expression should be another delayed cast or an erroneous expression.

LEMMA C.14 (REDUCIBLE DELAYED CAST PRESERVATION). *If $\Sigma \vdash v$ and $\Sigma \mid \emptyset \vdash dc : T$ and $\Sigma' \mid \emptyset \vdash e : T$ and $\Sigma' \vdash v'$ and $dc, v \longrightarrow_c e, v'$, then either*

- (1) e is a delayed cast, or
- (2) $e = \mathcal{E}[\text{error}]$

we prove progress when the heap contains at least one delayed cast that is not a value.

LEMMA C.15 (PROGRESS WITH EVOLVING HEAP). *If $\Sigma \vdash v$ and v is not a normal heap and $\Sigma \mid \emptyset \vdash M : T$, then $\exists v'$ s.t. $M, v \longrightarrow M, v'$*

PROOF. Because v is not normal, then $\exists a, dc$ s.t. $v(a)_{\text{val}} = dc$ and dc is not a value. By Lemma C.12, dc takes a step to an expression e . By Lemma C.14, there are two cases:

Case $e = \mathcal{E}[\text{error}]$: A step is taken via ERROR.

Case e is a delayed cast : By cases on whether $v(a)_{\text{rtti}} = v'(a)_{\text{rtti}}$

Case $v(a)_{\text{rtti}} = v'(a)_{\text{rtti}}$: A step is taken via NoRTTICheck

Case $v(a)_{\text{rtti}} \neq v'(a)_{\text{rtti}}$: A step is taken via RTTIChecked

\square

Progress in the case of a normal heap is standard.

LEMMA C.16 (PROGRESS WITH NORMAL HEAP). *If $\Sigma \vdash \mu$ and $\Sigma \mid \emptyset \vdash M : T$, then either*

- (1) M is a value, or
- (2) $M = \text{error}$, or
- (3) $\exists N, v$ s.t. $M, \mu \longrightarrow N, v$

PROOF. By Lemma C.13 and standard induction on the typing derivation. \square

From Lemmas C.15 and C.16, we assemble the full proof of progress by cases on whether the heap is in a normal state:

COROLLARY C.17 (PROGRESS). *If $\Sigma \vdash v$ and $\Sigma \mid \emptyset \vdash M : T$, then either*

- (1) M is a value, or
- (2) $M = \text{error}$, or
- (3) $\exists N, v'$ s.t. $M, v \longrightarrow N, v'$

LEMMA C.18 (TYPE PRESERVATION). *If $\Sigma \vdash v$ and $\Sigma \mid \emptyset \vdash M : T$ and $M, v \longrightarrow N, v'$, then $\exists \Sigma'$ s.t. $\Sigma' \mid \emptyset \vdash N : T$ and $\Sigma' \vdash v'$ and $\Sigma' \sqsubseteq_{p/e} \Sigma$*

From Corollary C.16 and Lemma C.18, we prove type safety.

THEOREM C.19 (TYPE SAFETY). *If $\Sigma \vdash v$ and $\Sigma \mid \emptyset \vdash e : T$, then either*

- (1) e diverges, or
- (2) $e = \text{error}$, or
- (3) $\exists v, \Sigma', v'$ s.t. $e, v \longrightarrow^* v, v'$ and $\Sigma' \mid \emptyset \vdash v : T$ and $\Sigma' \vdash v'$