

Jones-Optimal Partial Evaluation by Specialization-Safe Normalization

MATT BROWN, University of California Los Angeles, USA

JENS PALSBERG, University of California Los Angeles, USA

We present partial evaluation by specialization-safe normalization, a novel partial evaluation technique that is Jones-optimal, that can be self-applied to achieve the Futamura projections and that can be type-checked to ensure it always generates code with the correct type. Jones-optimality is the gold-standard for nontrivial partial evaluation and guarantees that a specializer can remove an entire layer of interpretation. We achieve Jones-optimality by using a novel affine-variable static analysis that directs specialization-safe normalization to always decrease a program's runtime.

We demonstrate the robustness of our approach by showing Jones-optimality in a variety of settings. We have formally proved that our partial evaluator is Jones-optimal for call-by-value reduction, and we have experimentally shown that it is Jones-optimal for call-by-value, normal-order, and memoized normal-order. Each of our experiments tests Jones-optimality with three different self-interpreters.

We implemented our partial evaluator in $F_{\omega}^{\mu i}$, a recent language for typed self-applicable meta-programming. It is the first Jones-optimal and self-applicable partial evaluator whose type guarantees that it always generates type-correct code.

CCS Concepts: • **Software and its engineering** → **Source code generation**; *Translator writing systems and compiler generators*; Interpreters;

Additional Key Words and Phrases: Partial Evaluation, Jones-Optimality, Futamura Projections, Meta-Programming, Lambda Calculus.

ACM Reference Format:

Matt Brown and Jens Palsberg. 2018. Jones-Optimal Partial Evaluation by Specialization-Safe Normalization. *Proc. ACM Program. Lang.* 2, POPL, Article 14 (January 2018), 53 pages. <https://doi.org/10.1145/3158102>

1 INTRODUCTION

A partial evaluator implements Kleene's s-m-n theorem. Given a program and specific values for some of its inputs, a partial evaluator generates a version of the program specialized to those inputs. When given the remaining inputs, the specialized program computes the same result as the original when given all the inputs at once. The goal of partial evaluation is to do some of the computation ahead of time, so that the specialized program runs faster than the original. Researchers have developed partial evaluators for Scheme [Bondorf and Danvy 1991], C [Andersen 1992], and many others. Researchers have also shown how to self-apply partial evaluators and generate the Futamura projections [Jones et al. 1985], which can be used to compile programs and generate compilers and compiler-generators. Finally, researchers have developed Jones-optimal partial evaluators, that is, they can specialize away the overhead of an interpreter.

Authors' addresses: Matt Brown, University of California Los Angeles, USA, msb@cs.ucla.edu; Jens Palsberg, University of California Los Angeles, USA, palsberg@ucla.edu.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2018 Copyright held by the owner/author(s).

2475-1421/2018/1-ART14

<https://doi.org/10.1145/3158102>

What is the type of a partial evaluator? In 1993, Jones, Gomard, and Sestoft [Jones et al. 1993, Section 16.2] suggested the following polymorphic type:

$$\forall A:*. \forall B:*. \text{Exp } (A \rightarrow B) \rightarrow A \rightarrow \text{Exp } B$$

Here, $\text{Exp } (A \rightarrow B)$ is the type of a representation of a program with type $A \rightarrow B$. This technique is called *typed representation*. Implementing a partial evaluator that operates on typed representations guarantees that it always generates well-typed code of the correct type. Jones, Gomard, and Sestoft showed that the above type supports self-application and the Futamura projections. However, the challenge to define a partial evaluator with that type has remained open. Carette, Kiselyov, and Shan [Carette et al. 2009] implemented a partial evaluator with the closely related type that provides the same guarantee:

$$\forall A:*. \forall B:*. \text{Exp } (A \rightarrow B) \rightarrow \text{Exp } A \rightarrow \text{Exp } B$$

Their partial evaluator cannot be self-applied and comes with no claim about Jones optimality. In this paper we present a partial evaluator with the latter type, the first self-applicable partial evaluator that operates on typed representations, supports the Futamura projections, and is Jones optimal.

Our starting point is Mogensen’s partial evaluator [Mogensen 1995] for the untyped lambda-calculus. He used reduction to β -normal form as the basis for his partial evaluator, a technique we call specialization by normalization. A key to supporting the Futamura projections using specialization by normalization is that the partial evaluator itself must have a β -normal form. Mogensen showed this is possible by using Church-encoding to represent λ -terms.

Our first step is to observe that Mogensen’s partial evaluator is not Jones optimal for reduction strategies with *sharing* – where the work of reducing a value bound to a variable is done at most once, either up front or on demand, regardless of how many times that variable is referenced. Such reduction strategies (e.g. call-by-value, call-by-need/lazy) are more commonly used in practice than those without sharing (e.g. call-by-name), because sharing provides significantly better performance. The problem with Mogensen’s partial evaluator is that β -normalization is too aggressive for strategies with sharing; it duplicates work that could have been shared, resulting in specialized programs that are slower than the original. Our solution is to define a subset of β -reduction that is *specialization-safe* – that doesn’t duplicate work and never causes a slowdown. Specialization-safe reduction relies on a novel affine-variable analysis that is used to identify the β -redexes in that safe subset.

Our partial evaluator modifies Mogensen’s specialization by normalization technique to use specialization-safe reduction. Like Mogensen’s partial evaluator, ours has a β -normal form, so it supports the Futamura projections. However, ours never causes a slowdown and is Jones optimal.

Jones optimality is defined in terms of a fixed self-interpreter and a fixed reduction strategy. Given the variety of reduction strategies used in practice, and the many ways to write a self-interpreter, there are many different instances of Jones optimality to consider. Our partial evaluator is the first to be Jones optimal for three different self-interpreters combined with three different reduction strategies. Each self-interpreter operates on a different style of representation, and has a different amount of interpretational overhead. The representation techniques we consider are tagless-final higher-order abstract syntax (HOAS), Mogensen-Scott HOAS, and a tagless-final representation with deBruijn indices.

We apply our techniques to untyped λ -calculus and to a typed λ -calculus F_{ω}^{μ} that supports self-applicable meta-programming. For untyped λ -calculus, we prove that partial evaluation by specialization-safe normalization is Jones optimal for each self-interpreter under call-by-value reduction. For F_{ω}^{μ} , we experimentally verify Jones optimality for each interpreter and each of call-by-value, normal-order, and memoized normal-order reduction, as indicated by the Xs in Table 1.

Reduction	Tagless-final HOAS	Mogensen-Scott HOAS	Tagless-final deBruijn
Call-by-value	X \checkmark	X	X
Normal order	X	X	X
Memoized normal order	X	X	X

Table 1. Experimental (X) and formal (\checkmark) verification of Jones optimality of our $F_{\omega}^{\mu_i}$ partial evaluator.

The \checkmark indicates that we have also formally proved Jones optimality for the tagless-final HOAS interpreter and call-by-value reduction.

An appendix containing proofs of the theorems stated in this paper is available from our website [Brown and Palsberg 2018]. We also provide our software artifact that includes implementations of $F_{\omega}^{\mu_i}$ and its self-interpreters and partial evaluator, as well as instructions for reproducing our experimental results.

Rest of the paper. In Section 2, we review Mogensen’s partial evaluator for untyped λ -calculus and demonstrate that it is not Jones optimal. In Section 3, we introduce our key notion of a specialization-safe reduction relation. In Section 4, we present a specialization-safe reduction that operates on untyped λ -terms with affine variable annotations. In Section 5, we discuss partial evaluation by specialization-safe normalization, and prove it Jones optimal for three self-interpreters. In Section 6, we present our typed self-applicable partial evaluator for $F_{\omega}^{\mu_i}$ and the typed Futamura projections. In Section 7, we discuss our experimental results, and in Section 8 we discuss related work.

2 PARTIAL EVALUATION AND JONES OPTIMALITY

A partial evaluator is a meta-program that specializes another program to some of its inputs. When given the remaining inputs, the specialized program will compute the same result as running the original program on all the inputs. Kleene’s s-m-n theorem established that such a specialization process is possible, and since then partial evaluation has been established as a practical technique for program optimization and automatic program generation.

We will define partial evaluation as the specialization of a two-input function to its first input (corresponding to the s-1-1 instance of s-m-n).

Definition 2.1 (Partial Evaluation). mix is a partial evaluator if for every program p and input x , there exists a specialized program p_x such that:

- (1) If $\text{mix } \bar{p} \ \bar{x}$ has a β -normal form, that normal form is \bar{p}_x
- (2) $\forall y. p_x y \equiv_{\beta} p \times y$

The notation \bar{p} denotes the representation of p according to some fixed representation scheme. The first condition indicates that the partial evaluator may not be total: specialization may not terminate. When it does terminate, however, it outputs the representation of a specialized program p_x . The second condition is the correctness condition for the specialized program p_x : it must have the same behavior as $p \times$ on all inputs y . The definition naturally extends to other numbers of inputs. The inputs to which the program is specialized are often called “static” and the remaining inputs “dynamic”. The static inputs are given as representations because they may become part of the output program representation. This is necessary for a λ -calculus without constants [Launchbury 1991].

Futamura projections. A partial evaluator is self-applicable if it can specialize itself. A classical application of a self-applicable partial evaluator is to generate the Futamura projections, which

Description	Definition	Correctness Specification
1. Compile S -program to T	$\text{mix } \overline{\text{int}} \overline{\overline{s}} \equiv_{\beta} \overline{t}$	$\forall x. t \ x \equiv_{\beta} \text{int } \overline{\overline{s}} \ x$
2. Generate an S -to- T compiler	$\text{mix } \overline{\text{mix}} \overline{\overline{\text{int}}} \equiv_{\beta} \overline{\overline{\text{comp}}}$	$\forall s. \text{comp } \overline{\overline{s}} \equiv_{\beta} \text{mix } \overline{\overline{\text{int}}} \overline{\overline{s}}$
3. Generate a compiler-generator	$\text{mix } \overline{\text{mix}} \overline{\overline{\overline{\text{mix}}}} \equiv_{\beta} \overline{\overline{\overline{\text{cogen}}}}$	$\forall \text{int}. \text{cogen } \overline{\overline{\text{int}}} \equiv_{\beta} \text{mix } \overline{\overline{\overline{\text{mix}}}} \overline{\overline{\text{int}}}$
4. Self-generation of cogen		$\overline{\overline{\text{cogen}}} \overline{\overline{\text{mix}}} \equiv_{\beta} \overline{\overline{\overline{\text{cogen}}}}$

Fig. 1. The four Futamura projections for λ -calculus.

$$\text{mix } \overline{f} \overline{d} \equiv_{\beta} \overline{\overline{\text{NF}_{\beta}(f \ d)}}$$

Fig. 2. Mogensen's Specialization by Normalization

use specialization to compile programs and generates compilers and compiler-generators [Futamura 1999]. The four Futamura projections are shown in Figure 1. For each projection we list a correctness specification implied by the correctness of mix . The first Futamura projection compiles programs from a source language S to a target language T by specializing an S -interpreter programmed in T . Given any input to the program, the target program gives the same result as interpreting the source program. The second projection generates a compiler from S -programs to T -programs by specializing the partial evaluator mix to the interpreter. Given any program, the compiler gives the same result as the first projection. The third projection generates a compiler-generator by specializing mix to itself. Given any interpreter, the compiler-generator gives the same result as the second projection. Futamura only discussed these three Futamura projections. Subsequently Glück [Glück 2009] showed that there are more. In particular, the fourth projection demonstrates that the compiler-generator cogen can generate itself when applied to mix . This is sometimes called the “mixpoint”.

The twice-overlined terms in the Futamura projections are representations of representations. These double-representations arise because mix requires its static input to be represented, as we mentioned above.

Jones-optimality. Jones-optimality is the gold-standard for showing a partial evaluator can perform a significant amount of specialization. The definition of Jones-optimality specializes a self-interpreter as a benchmark for partial evaluation.

Definition 2.2. A term u is a self-interpreter for a representation scheme $\overline{\cdot}$ if for any term p , we have $u \ \overline{p} \equiv_{\beta} p$.

A partial evaluator is Jones-optimal if it can specialize away all the computational overhead caused by self-interpretation. More precisely, specializing a self-interpreter to a program should generate a program that is no slower than the original program – on *any* input.

The canonical definition of Jones optimality is from Jones, Gomard and Sestoft [Jones et al. 1993, pg. 139]. We restate it for our context of λ -calculus and using our notation:

Definition 2.3 (Jones-optimality). A partial evaluator mix is Jones-optimal with respect to *time* and a self-interpreter u if, for any terms $p : A \rightarrow B$ and $d : A$, $\text{time}(p' \ d) \leq \text{time}(p \ d)$, where $\text{mix } \overline{u} \overline{\overline{p}} \equiv_{\beta} \overline{\overline{p'}}$.

Jones optimality is defined in terms of a particular fixed *time* function and self-interpreter u . Since different notions of *time* may be most appropriate for different settings, and because there are different ways of implementing a self-interpreter (and different ways of representing programs) for

$$\frac{}{(\lambda x. e)v \rightarrow e[x:=v]} \qquad \frac{e1 \rightarrow e1'}{e1 e2 \rightarrow e1' e2} \qquad \frac{e1 \rightarrow e1'}{(\lambda x. e)e1 \rightarrow (\lambda x. e)e1'}$$

Fig. 3. Call-by-value reduction.

$$\frac{\frac{\frac{x \triangleright x}{e \triangleright q}}{\lambda x. e \triangleright \text{abs}(\lambda x. q)}}{e1 \triangleright q1 \quad e2 \triangleright q2}}{e1 e2 \triangleright \text{app} q1 q2}}{e \triangleright q}}{\bar{e} = \lambda \text{abs}. \lambda \text{app}. q}$$

$$u = \lambda e. e (\lambda x. x) (\lambda x. x)$$

Fig. 4. Mogensen’s Church-encoded representation and self-interpreter for untyped λ -calculus.

a single language, we can consider multiple notions of Jones optimality. We will always explicitly state which *time* function and self-interpreter is under consideration. In this section, we focus on steps of call-by-value (CBV) reduction to a value as our notion of *time*. Call-by-value reduction is defined in Figure 3.

Definition 2.4. $\text{time}_{cbv}(e) = n$ if and only if $e \rightarrow^n v$ for some value v .

Mogensen [Mogensen 1995] implemented a self-applicable partial evaluator for untyped λ -calculus by reducing terms to β -normal form, a technique we call “specialization by normalization”. With this approach, specializing a program f to an input x returns $\text{NF}_\beta(f x)$, the β -normal form of $f x$. One problem with specialization by normalization is that not all terms have a normal form, which leads to non-termination of the partial evaluator: if $f x$ has no normal form, then specializing f to x will not terminate. In particular, if the self-interpreter u and the partial evaluator mix are not strongly normalizing terms, then the Futamura projections will not terminate. Mogensen solved this problem by using a representation based on Church-encoding, for which a strongly normalizing self-interpreter and normalizer can be defined. The partial evaluator is self-applicable and the Futamura projections terminate.

Mogensen’s Church-encoding representation and self-interpreter are shown in Figure 4. A representation is built in two steps: first, the pre-quoter \triangleright applies designated variables `abs` and `app` throughout the term, at each λ -abstraction and application respectively. We assume `abs` and `app` do not occur in the term, which can be ensured by renaming variables. The quoter $\bar{}$ simply abstracts over `abs` and `app` in the pre-quoted term. This formulation of Church encoding produces representations in β -normal form. It is possible to define quotation as a single function (without a pre-quoter), but the representations produced would not be β -normal.

Mogensen’s self-interpreter u is instantiates both `abs` and `app` to identity functions, resulting in a term that is β -equivalent to e . The self-interpreter u is β -normal. While Mogensen’s normalizer is significantly more complex than his self-interpreter, it is also strongly normalizing.

$$\begin{aligned}
n_i &= \lambda s. \lambda z. s^i z \\
\text{succ} &= \lambda n. \lambda s. \lambda z. s (n s z) \\
\\
\text{plus} &= \lambda n. \lambda m. n \text{ succ } m \\
\text{times} &= \lambda n. \lambda m. n (\text{plus } m) n_0 \\
\text{square} &= \lambda n. \text{times } n n \\
\\
f &= \lambda x. \lambda y. \text{times } x (\text{square } y) \\
g &= f n_2
\end{aligned}$$

Fig. 5. An implementation of $g(x) = 2x^2$ on Church numerals

To summarize, Mogensen made three key insights: 1) a self-evaluator that evaluates to normal form can be used as the basis for a partial evaluator (specialization by normalization), 2) specialization by normalization supports the Futamura projections if the self-interpreter and partial evaluator themselves are strongly normalizing, and 3) a representation based on Church-encoding supports a strongly normalizing self-interpreter and specialization by normalization partial evaluator.

Specialization by normalization is not optimal for $time_{cbv}$. For example, consider the λ -term g in Figure 5 that computes the polynomial $g(x) = 2x^2$ on Church numerals. The specialization of u to \bar{g} yields a term u_g that is slower than g on some inputs. For example, $u_g n_3$ evaluates in 89 steps, while $g n_3$ evaluates in 82 steps. The problem is not limited to compilation; specialization in general can cause slowdown. The specialization of the λ -term f in Figure 5 to the church numeral n_2 evaluates in more steps than $f n_2$.

The reason for this slowdown is that β -normalization contracts some β -redexes that cause duplicated work in the residual code – in particular, redexes of the form $(\lambda x. e1)e2$ where $e2$ is not a value. Reducing such a redex can introduce multiple new copies of $e2$ into the term, each of which may need to be evaluated separately. For example, consider the specialization of f to n_2 , which reduces in one step to $\lambda y. \text{times } n_2 (\text{square } y)$, and then in a few more steps to $\lambda y. \text{plus } (\text{square } y) (\text{plus } (\text{square } y) n_0)$. The specializer will continue reducing, but we can see the problem already at this point: we now have two copies of $\text{square } y$ instead of one. We can't compute either copy at specialization-time, because we don't know the value of y until run-time. When we run the specialized program with the input n_2 for y , will have to evaluate $(\text{square } n_2)$ twice.

Returning to our first example, specializing u to g , we see that since $u \bar{g} \equiv_{\beta} g = f n_2$, the β -normal form of $u \bar{g}$ is the same as the β -normal form of $f n_2$. The result is the same as for the specialization of f to n_2 , and we get a slowdown for the same reason.

3 SPECIALIZATION SAFETY

In this section, we consider how Mogensen's specialization-by-normalization approach can be modified so that it can support the Futamura projections and also be Jones optimal. Jones optimality has two requirements: first, the partial evaluator must be strong enough to specialize away the work of a self-interpreter; second, it must not duplicate work or otherwise cause a slowdown.

Our goal is to identify a reduction relation that can be used as the basis for a Jones-optimal specialization-by-normalization partial evaluator. Our main requirements are as follows: first, the reduction relation should be specialization-safe, meaning it never causes a slowdown at run-time;

and second, it should be strong enough to remove all the work of a self-interpreter. We now formally state these two requirements.

Definition 3.1. A reduction relation \rightsquigarrow is specialization-safe for *time* if $e1 \rightsquigarrow^* e2$ implies $\text{time}(e2 \ x) \leq \text{time}(e1 \ x)$ for all x .

Definition 3.2. For a given self-interpreter u for a representation function $\widehat{\cdot}$, a reduction relation \rightsquigarrow is u -strong if $u \widehat{p} \rightsquigarrow^* p$ for all p .

We also require that the reduction relation be confluent and have unique normal forms, so that we can implement a specialization-by-normalization partial evaluator for it. Finally, in order to support the Futamura projections, we need the implementation of the partial evaluator itself to have a normal form.

THEOREM 3.3. *If \rightsquigarrow is confluent, specialization-safe for time, and u -strong, then for any terms p and d , $\text{time}(p' \ d) \leq \text{time}(p \ d)$, where $p' = \text{NF}_{\rightsquigarrow}(u \widehat{p})$.*

PROOF. From that \rightsquigarrow is confluent and u -strong, we have $p' = \text{NF}_{\rightsquigarrow}(u \widehat{p}) = \text{NF}(p)$, so $p \rightsquigarrow^* p'$. From that \rightsquigarrow is specialization-safe for *time*, we have $\text{time}(p' \ d) \leq \text{time}(p \ d)$. \square

In the remainder of this section, we examine a variety of β -redexes and determine if they can be included in a specialization-safe reduction relation. In the next section we introduce affine variable annotations that help identify those specialization-safe β -redexes and define a specialization-safe subset of β -reduction that operates on annotated terms.

Intuitively, whether a redex $(\lambda x. e)a$ is specialization-safe depends upon two things: the potential cost of evaluating a at run-time, and the number of times the value of a might be needed.

Consider for example the term $\lambda f. (\lambda x. f \ x \ x) \ n_2$, which contains a single redex $(\lambda x. f \ x \ x) \ n_2$. Since there are two occurrences of x , the value of n_2 will be needed twice. However, n_2 is itself a value, so the cost of evaluating it is 0. Therefore, this redex is specialization-safe – contracting does not duplicate any work.

Consider instead the term $\lambda f. (\lambda x. f \ x \ x)$ (square n_2), which contains the redex $(\lambda x. f \ x \ x)$ (square n_2). This time, the argument (square n_2) is not a value, so contracting this redex would duplicate work: in the reduct f (square n_2) (square n_2), we have to evaluate (square n_2) twice. However, (square n_2) can be reduced to a value at specialization time, so we can reduce the redex $(\lambda x. f \ x \ x)$ (square n_2) to a specialization-safe redex and then contract it.

Next, consider the term $\lambda f. (\lambda x. \text{times } x \ x)$ ($f \ n_2$) and its redex $(\lambda x. \text{times } x \ x)$ ($f \ n_2$). This time, specialization cannot reduce the argument $f \ n_2$ to a value, because f is unknown. If we leave the redex uncontracted, then $f \ n_2$ will be evaluated once at run-time, and the redex will be contracted afterwards. If we contract the redex, the resulting term $\text{times } (f \ n_2)$ ($f \ n_2$) duplicates $f \ n_2$ so it will need to be evaluated twice at run-time, potentially causing a slowdown. Therefore, this redex is not specialization-safe.

Next, consider $\lambda f. (\lambda y. \text{square } y)$ ($f \ n_2$), which is β -equivalent to the previous example, and its redex $(\lambda y. \text{square } y)$ ($f \ n_2$). As before, specialization cannot reduce $f \ n_2$ to a value. But in this case the value of y is needed only once, so this redex is specialization-safe and can be contracted to $\text{square } (f \ n_2)$ without risking a slowdown. The resulting term $\lambda f. \text{square } (f \ n_2) = \lambda f. (\lambda x. \text{times } x \ x)$ ($f \ n_2$) recovers the previous example.

Next, consider the term $(\lambda y. \text{square } y) = \lambda y. (\lambda x. \text{times } x \ x) \ y$ and its redex $(\lambda x. \text{times } x \ x) \ y$. Here, the argument y is a variable and so not a value, and neither can specialization reduce it to a value. However, since call-by-value always reduces arguments to values before β -contracting, we

(values) $v := \lambda x. e \mid \lambda x^\circ. e$
 (terms) $e := x \mid x^\circ \mid e_1 e_2 \mid \lambda x. e \mid \lambda x^\circ. e$

Fig. 6. Untyped λ -calculus with affine-variable annotations

$$\begin{array}{c}
 \text{CBV-1} \frac{}{(\lambda x. e)v \rightarrow e[x:=v]} \\
 \text{CBV-2} \frac{e_1 \rightarrow e_1'}{e_1 e_2 \rightarrow e_1' e_2} \\
 \text{CBV-3} \frac{e_1 \rightarrow e_1'}{(\lambda x. e)e_1 \rightarrow (\lambda x. e)e_1'} \\
 \text{CBV-4} \frac{}{(\lambda x^\circ. e)v \rightarrow e[x^\circ:=v]} \\
 \text{CBV-5} \frac{e_1 \rightarrow e_1'}{(\lambda x^\circ. e)e_1 \rightarrow (\lambda x^\circ. e)e_1'}
 \end{array}$$

Fig. 7. Call-by-value reduction on untyped λ -terms with affine variable annotations.

know that at run-time, y will be substituted with a value. Therefore this redex is also specialization-safe, and $(\lambda y. \text{square } y)$ can be reduced to $(\lambda y. \text{times } y \ y) \equiv_\alpha \text{square}$. This example shows that η -contraction is specialization-safe.

Finally, consider $\lambda f. \lambda g. g \ (\text{times } (f \ n_2))$, and its redex $\text{times } (f \ n_2) = (\lambda n. \lambda m. n \ (\text{plus } m) \ n_0) \ (f \ n_2)$. Since $(f \ n_2)$ cannot be reduced to a value at specialization time, it will have a nonzero cost at run-time. However, since n occurs only once in the definition of times , it is tempting to conclude that this redex is specialization-safe. This conclusion is wrong. Since n occurs under the λ -abstraction $\lambda m. n \ (\text{plus } m) \ n_0$, the value of n will be needed every time that abstraction is applied. Since g is unknown, we cannot predict how many times the λ -abstraction will be applied and n will be needed. If we contract the redex, then we have to recompute $f \ n_2$ every time g is applied. This could cause a lot of duplicated work.

4 AFFINE VARIABLES AND SPECIALIZATION-SAFE REDUCTION

The examples in the previous section demonstrate that full β -reduction is not specialization-safe: some β -contractions at specialization-time can cause a slowdown at run-time. In this section, we define *specialization-safe reduction*, a subset of full β -reduction that is indeed specialization-safe: Theorem 4.5 states that no sequence of specialization-safe reductions can result in a term that evaluates in more call-by-value steps than the original.

We achieve specialization safety by only contracting β -redexes that cannot cause a slowdown. We use affine-variable annotations to indicate which variables are guaranteed to be referenced at most once at run-time. Affine variables are similar to linear variables, except that linear variables are guaranteed to be referenced *exactly* once.

Figure 6 defines the untyped λ -calculus (ULC) with *affine* variables. Affine variables are annotated like x° , while unlimited variables are unannotated. The values are λ -abstractions as usual, and they may abstract over affine or unlimited variables. Throughout this section, all λ -terms can have annotations.

Figure 7 defines call-by-value reduction on annotated λ -terms. Call-by-value reduction ignores annotations – the rules CBV-1 and CBV-3 are identical except for the annotations, as are CBV-4 and CBV-5. We overload the arrow \rightarrow that denotes call-by-value reduction for unannotated terms. We will make clear which kind of term is under consideration at all times. We define time_{cbv} for annotated terms similarly to the definition for unannotated terms.

Definition 4.1. For an annotated term e , $\text{time}_{cbv}(e) = n$ if and only if $e \rightarrow^n v$ for some value v .

$$\begin{array}{c}
\text{(unlimited contexts)} \Gamma = \langle \rangle \mid \Gamma, x \\
\text{(affine contexts)} \Sigma = \langle \rangle \mid x^\circ \\
\\
\frac{x \in \Gamma}{\Gamma; \Sigma \vdash x} \\
\frac{}{\Gamma; x^\circ \vdash x^\circ} \\
\frac{\Gamma; \Sigma_1 \vdash e_1 \quad \Gamma; \Sigma_2 \vdash e_2}{\Gamma; \Sigma \vdash e_1 e_2} \quad \Sigma = \Sigma_1 \uplus \Sigma_2 \\
\frac{}{\langle \rangle \uplus x = x} \\
\frac{}{x \uplus \langle \rangle = x} \\
\frac{}{\langle \rangle \uplus \langle \rangle = \langle \rangle} \\
\frac{}{(\Gamma, x); \langle \rangle \vdash e} \\
\frac{}{\Gamma; \Sigma \vdash \lambda x. e} \\
\frac{}{\Gamma; x^\circ \vdash e} \\
\frac{}{\Gamma; \Sigma \vdash \lambda x^\circ. e}
\end{array}$$

Fig. 8. Well-formedness rules for affine-variable annotations.

$$\begin{array}{c}
\text{SSR-1} \frac{}{(\lambda x. e)v \rightarrow_s e[x:=v]} \\
\text{SSR-2} \frac{e_1 \rightarrow_s e_1'}{e_1 e_2 \rightarrow_s e_1' e_2} \\
\text{SSR-3} \frac{e_1 \rightarrow_s e_1'}{e e_1 \rightarrow_s e e_1'} \\
\text{SSR-4} \frac{}{(\lambda x^\circ. e)e' \rightarrow_s e[x^\circ:=e']} \\
\text{SSR-5} \frac{}{(\lambda x. e)y \rightarrow_s e[x:=y]} \\
\text{SSR-6} \frac{e \rightarrow_s e'}{(\lambda x. e) \rightarrow_s (\lambda x. e')} \\
\text{SSR-7} \frac{e \rightarrow_s e'}{(\lambda x^\circ. e) \rightarrow_s (\lambda x^\circ. e')}
\end{array}$$

Fig. 9. Specialization-Safe Reduction

Figure 8 defines a simple set of rules for checking affine-variable annotations. We track variables in two contexts: a context Γ that contains unlimited variables, and a context Σ that contains affine variables. Σ is allowed to contain at most one variable at a time. The rule for unlimited variables checks for the variable in the unlimited context, and the one for affine variables checks the affine context. For an application $e_1 e_2$, we check e_1 and e_2 using the same unlimited context, but disjoint affine contexts. This ensures that only e_1 or e_2 may reference the current affine variable. We always check the body of an unlimited λ -abstraction using an empty affine context. We check the body of an affine λ -abstraction $(\lambda x^\circ. e)$ using the affine context containing only x° . This means that affine variables may never occur under a nested λ abstraction.

Figure 9 defines specialization-safe reduction on annotated terms. The rule SSR-1 is ordinary call-by-value β -reduction. The rules SSR-2 and SSR-3 allow a step to occur in either sub-term of an application, like in full β -reduction. Call-by-value reduction would restrict SSR-3 to when e is a value. The rule SSR-4 is the key rule enabled by the affine-variable annotations. It allows a β -reduction when the argument is not a value, as long as the λ -abstraction binds an affine variable. The rule SSR-5 allows a β -reduction when the argument is an unlimited variable, as long as the λ -abstraction is also unlimited. Note that we do not allow β -reduction when the argument is an affine variable and the λ -abstraction is unlimited, because that could duplicate the affine variable. If the λ -abstraction is affine, SSR-4 applies for any argument. The last two rules, SSR-6 and SSR-7, allow reduction under λ -abstractions.

We now apply specialization-safe reduction to the examples from the previous section. First was the term $\lambda f. (\lambda x. f x x) n_2$. The variable x cannot be annotated as affine, because it occurs twice. Neither can f be annotated, because it occurs under the abstraction over x . However, since n_2 is a value, SSR-1 applies, so this is a redex for specialization-safe reduction.

The next example was $\lambda f. (\lambda x. f \times x)$ (square n_2). Again, x and f are not affine, but square n_2 can be reduced to a value, after which SSR-1 will apply.

The next example was $\lambda f. (\lambda x. \text{times } x \times x)$ ($f \ n_2$), which has the β -redex $(\lambda x. \text{times } x \times x)$ ($f \ n_2$). Again x is not affine and $(f \ n_2)$ is not a value, so this is not a specialization-safe redex. Also, since $(f \ n_2)$ is not reducible to a value, the β -redex can't be reduced to a specialization-safe redex.

The next example was $\lambda f. (\lambda y. \text{square } y)$ ($f \ n_2$), which can be annotated $\lambda f^\circ. (\lambda y^\circ. \text{square } y^\circ)$ ($f^\circ \ n_2$). Now consider the β -redex $(\lambda y^\circ. \text{square } y^\circ)$ ($f^\circ \ n_2$). The argument $(f^\circ \ n_2)$ is not value and can't be reduced to one. However, since y° is affine, this is still a specialization-safe redex, reducible using rule SSR-4.

Next consider $\lambda y. (\lambda x. \text{times } x \times x) \ y$. The variable y is affine, so we can annotate the term $\lambda y^\circ. (\lambda x. \text{times } x \times x) \ y^\circ$. In this case, specialization-safe reduction does not allow the β -redex $(\lambda x. \text{times } x \times x) \ y^\circ$ to be reduced, because the reduct $\text{times } y^\circ \ y^\circ$ would not be well-formed. We could instead leave y unannotated, and the term is still well-formed. Then $(\lambda x. \text{times } x \times x) \ y$ can be reduced by rule SSR-5 because y is unlimited, and the reduct $(\lambda y. \text{times } y \ y)$ is well-formed. This demonstrates two important points: a term can have multiple annotations, and different annotations can lead to different specialization-safe reduction behavior. Our partial evaluator always inserts the maximum number of annotations allowed by the rules, as described in the next section.

The last example was $\lambda f. \lambda g. g \ (\text{times } (f \ n_2))$, and its redex $\text{times } (f \ n_2) = (\lambda n. \lambda m. n \ (\text{plus } m)) \ n_0$ ($f \ n_2$). The variable n cannot be annotated as affine because it occurs on the λ -abstraction for m . Also, the argument $(f \ n_2)$ is not a value, so this is not reducible by specialization-safe reduction.

The following theorem states that well-formedness is preserved by specialization-safe reduction.

THEOREM 4.2. *If $\Gamma; \Sigma \vdash e$ and $e \rightarrow_s e'$, then $\Gamma; \Sigma \vdash e'$.*

A term e is in specialization-safe normal form (or is specialization-safe-normal) if there is no e' such that $e \rightarrow_s e'$. All β -normal forms are also specialization-safe-normal, but a specialization-safe normal form may contain β -redexes (specifically, specialization-unsafe redexes). Specialization-safe reduction is confluent, and so has unique normal forms. Our proof of confluence is based on Nipkow's lecture notes on λ -calculus [Nipkow 2012]. We define $\text{NF}_s(e) = e'$ if and only if $e \rightarrow_s^* e'$ and e' is specialization-safe-normal.

THEOREM 4.3. *Specialization-safe reduction is confluent.*

COROLLARY 4.4. *Specialization-safe normal forms are unique.*

The following theorem states that specialization-safe reduction cannot cause a slowdown: no sequence of specialization-safe reductions can result in a term that evaluates in more more call-by-value steps than the original.

THEOREM 4.5. *If $\langle \rangle; \langle \rangle \vdash e$ and $e \rightarrow_s^* e'$ and $e \rightarrow^i v$, then there exists an $i' \leq i$ and a value v' such that $e' \rightarrow^{i'} v'$, and $v \rightarrow_s^* v'$.*

Theorem 4.5 is illustrated in Figure 10. The solid arrows correspond to the premises of the theorem, and the dashed arrows correspond the conclusion. After specialization-safe reduction, we may end up with a value v' that is different than what we'd get with call-by-value reduction alone. This is possible because specialization-safe reduction allows reduction under λ -abstractions. The two values are equivalent and converge with some additional specialization-safe reduction.

Specialization-safety follows from Theorem 4.5:

COROLLARY 4.6 (SPECIALIZATION-SAFETY). *If $\langle \rangle; \langle \rangle \vdash p$ and $e \rightarrow_s^* e'$, then for all x , $\text{time}_{cbv}(e \ x) = i$ implies $\text{time}_{cbv}(e' \ x) \leq i$.*




Fig. 10. Illustration of Theorem 4.5.




Fig. 11. Specialization-safe reduction for unannotated terms.

PROOF. $time_{cbv}(e\ x) = i$ implies that $e\ x \rightarrow^i v$ for some v . Since $e \rightarrow_s^* e'$, we have $e\ x \rightarrow_s^* e'\ x$. By Theorem 4.5, there exists an $i' \leq i$ and a value v' such that $e'\ x \rightarrow^{i'} v'$ and $v \rightarrow_s^* v'$. Therefore, $time_{cbv}(e'\ x) = i' \leq i = time_{cbv}(e\ x)$. \square

5 PARTIAL EVALUATION BY SPECIALIZATION-SAFE NORMALIZATION

The specification for a partial evaluator based on specialization-safe normalization is as follows:

$$\text{mix } \bar{p} \ \bar{d} \equiv_{\beta} \overline{\text{erase}(\text{NF}_s(\text{annotate}(\bar{p} \ \bar{d})))}$$

It operates on unannotated terms, but first annotates them with the maximum number of affine annotations allowed by the rules. There is a unique maximally annotated term for each unannotated term, and the annotation process is straightforward. A bound variable can be annotated as affine if 1) the number of occurrences is at most one, and 2) the number of occurrences inside inner lambda abstractions is zero. After reducing to specialization-safe normal form it erases the annotations. Since *annotate*, NF_s , and *erase* are all well-defined functions, there is a unique specialized program for each program and static input.

For unannotated terms $e1$ and $e2$, we write $e1 \Longrightarrow_s e2$ to mean the maximally annotated version of $e1$ reduces in some number of steps to an annotated version of $e2$. This is depicted in Figure 11.

Since call-by-value reduction on annotated terms ignores affine-variable annotations, an annotated term evaluates in the same amount of time as the corresponding unannotated term:

LEMMA 5.1. *For any annotated term e , $time_{cbv}(e) = time_{cbv}(\text{erase}(e))$.*

PROOF. Straightforward. \square

The following Lemma combines Corollary 4.6 and Lemma 5.1 to bridge the gap between unannotated and annotated λ -terms.

LEMMA 5.2. *For any self-interpreter u for a representation function $\widehat{\cdot}$, and for any term p , if $u \widehat{p} \Rightarrow_s p$, then mix is Jones-optimal for u and $time_{cbv}$.*

$$\begin{array}{c}
\widehat{x} = \lambda \text{var} . \lambda \text{abs} . \lambda \text{app} . \text{var } x \\
\widehat{e} = q \\
\hline
\widehat{\lambda x . e} = \lambda \text{var} . \lambda \text{abs} . \lambda \text{app} . \text{abs } (\lambda x . q) \\
\widehat{e1} = q1 \quad \widehat{e2} = q2 \\
\hline
\widehat{e1 e2} = \lambda \text{var} . \lambda \text{abs} . \lambda \text{app} . \text{app } q1 q2 \\
\text{fix}_f = (\lambda x . f (\lambda y . x x y)) (\lambda x . f (\lambda y . x x y)) \\
\text{fix} = \lambda f . \text{fix}_f \\
\text{msint} = \text{fix } (\lambda u . \lambda e . e (\lambda x . x) (\lambda f . \lambda x . u (f ((\lambda y . x)x)))) (\lambda f . \lambda x . u f (u x))
\end{array}$$

Fig. 12. Mogensen-Scott encoding and self-interpreter.

We prove that a partial evaluator based on specialization-safe normalization is strong enough to specialize away the work of three different self-interpreters. Therefore it is also Jones-optimal for each interpreter and time_{cbv} . The first self-interpreter is the one by Mogensen shown in Figure 4. It operates on the same Church-encoded representation $\bar{\cdot}$ that mix does. The second uses Mogensen-Scott encoding, another encoding of λ -terms by Mogensen [Mogensen 1992] that is based on Scott's encoding of natural numbers. The third is a tagless-final encoding that uses deBruijn indices to represent variables, based on Kiselyov's representation of STLC in Haskell [Kiselyov 2012]. Whenever mix specializes a self-interpreter other than Mogensen's, we have two representations involved: mix itself always operates on representations produced by $\bar{\cdot}$.

The following two theorems prove that a partial evaluator based on specialization-safe normalization is Jones-optimal for Mogensen's Church-encoding self-interpreter (Figure 4) and time_{cbv} .

LEMMA 5.3. *For any closed term e , $u \bar{e} \implies_s e$.*

THEOREM 5.4. *mix is Jones optimal for time_{cbv} and Mogensen's self-interpreter.*

PROOF. Follows from Lemma 5.2 and Lemma 5.3. □

Mogensen-Scott encoding. Quotation for a Mogensen-Scott-encoded representation is defined in Figure 12. It's a higher-order abstract syntax (HOAS) representation like Mogensen's Church-encoded representation. The key difference between the two representations is that while Church-encoding defines a fold over the syntax of the term, Mogensen-Scott-encoding defines pattern matching on the term. The two representations are isomorphic, but each encoding technique is well-suited to particular applications. Church-encoding is good for folds, while Mogensen-Scott is preferable for operations that aren't easily programmed as folds. Programming a fold like our self-interpreter takes a bit more work on a Mogensen-Scott representation. In particular, we have to use a fixpoint combinator to traverse the term. Still, this extra work can be done using only specialization-safe reduction, so our partial evaluator is Jones-optimal for the Mogensen-Scott interpreter as well.

LEMMA 5.5. *For any term e , $\text{msint } \widehat{e} \implies_s e$.*

THEOREM 5.6. *mix is Jones-optimal for time_{cbv} and msint .*

PROOF. Follows from Lemma 5.2 and Lemma 5.5. □

$$\begin{array}{c}
\frac{}{\Gamma, x \vdash x \blacktriangleright \text{fst}} \\
\frac{\Gamma \vdash x \blacktriangleright e}{\Gamma, y \vdash x \blacktriangleright (\lambda p. e (\text{snd } p))} \\
\\
\frac{\Gamma \vdash x \blacktriangleright e}{\Gamma \vdash x \triangleright \text{var } e} \\
\frac{\Gamma, x \vdash e \triangleright q}{\Gamma \vdash (\lambda x. e) \triangleright \text{abs } q} \\
\frac{\Gamma \vdash e1 \triangleright q1 \quad \Gamma \vdash e2 \triangleright q2}{\Gamma \vdash e1 e2 \triangleright \text{app } q1 q2} \\
\frac{\langle \rangle \vdash e \triangleright q}{\widehat{e} = \lambda \text{var}. \lambda \text{abs}. \lambda \text{app}. q} \\
\\
\text{dbVar} = (\lambda f. \lambda e. f e) \\
\text{dbAbs} = (\lambda b. \lambda e. \lambda x. b (\lambda f. f x e)) \\
\text{dbApp} = (\lambda a. \lambda b. \lambda e. a e (b e)) \\
\text{dbInt} = (\lambda q. q \text{dbVar } \text{dbAbs } \text{dbApp } (\lambda x. x))
\end{array}$$

Fig. 13. Tagless-final encoding with deBruijn indices.

deBruijn indices. Figure 13 shows a tagless-final representation that uses deBruijn indices to represent variables, unlike the two previous HOAS representations. It's based on Kiselyov's representation of simply-typed λ -calculus in Haskell [Kiselyov 2012], which uses nested pairs to represent environments, and composed projection functions to represent deBruijn indices. For example, the deBruijn index 0 is represented as the `fst`, index 1 is represented as `$\lambda p. \text{fst } (\text{snd } p)$` , and so on. The quoter uses a function $\Gamma \vdash x \blacktriangleright e$ to construct the projection function for variables. The self-interpreter for this representation is quite different than the ones for either the tagless-final or the Mogensen-Scott HOAS, but our partial evaluator is Jones optimal for it as well.

LEMMA 5.7. *For any closed term e , $\text{dbInt } \widehat{e} \implies_s e$.*

THEOREM 5.8. *`mix` is Jones-optimal for time_{cbv} and `dbInt`.*

PROOF. Follows from Lemma 5.2 and Lemma 5.7. \square

Discussion. We now demonstrate why each of the three specialization-safe β -reductions — rules SSR-1, SSR-4, and SSR-5 — are needed to achieve Jones-optimality for our three self-interpreters.

First, consider how the Church-encoding interpreter (Figure 4) operates on the representation of $(\lambda a. a a)$. We show the affine variable annotations for clarity.

$$\begin{aligned}
& \overline{u (\lambda a. a a)} \\
&= (\lambda e^\circ. e^\circ (\lambda x^\circ. x^\circ) (\lambda x^\circ. x^\circ)) (\lambda \text{abs}. \lambda \text{app}. \text{abs } (\lambda a. \text{app } a a)) \\
&\rightarrow_s (\lambda \text{abs}. \lambda \text{app}. \text{abs } (\lambda a. \text{app } a a)) (\lambda x^\circ. x^\circ) (\lambda x^\circ. x^\circ) \\
&\rightarrow_s^2 (\lambda x^\circ. x^\circ) (\lambda a. (\lambda x^\circ. x^\circ) a a) \\
&\rightarrow_s^2 (\lambda a. a a)
\end{aligned}$$

The first step is derived by SSR-4, since e is affine. The next two steps are derived by by SSR-1 because `abs` and `app` are not affine and their parameters are values. The last two steps are derived by by SSR-4 because the two x s are affine.

To demonstrate the need for SSR-5, consider how the Mogensen-Scott interpreter (Figure 12) operates on the representation of $(\lambda a. a a)$:

$$\begin{aligned}
& \text{msint } (\overline{\lambda a. a a}) \\
\rightarrow_s^* & \lambda x. \text{msint } ((\lambda a. \overline{a a}) ((\lambda y^\circ. x)x)) \\
\rightarrow_s & \lambda x. \text{msint } ((\lambda a. \overline{a a}) x) \\
\rightarrow_s & \lambda x. \text{msint } (\overline{x x}) \\
\rightarrow_s^* & \lambda x. x x
\end{aligned}$$

The first sequence of steps does not need SSR-5. The term $((\lambda y^\circ. x)x)$ ensures that x is unlimited, and reduces x in the next step, which is derived by SSR-4 because y is affine. The next step is not reducible by SSR-1 because x is a variable (and variables are not values), and is also not reducible by SSR-4 because a is not affine. However, it is reducible by SSR-5 because we forced x to be unlimited. This step allows the interpreter to proceed and eventually return the original term (up to renaming).

6 TYPED SELF-APPLICABLE PARTIAL EVALUATION FOR $F_\omega^{\mu i}$

We implement a typed version of our partial evaluator for $F_\omega^{\mu i}$, a Turing-complete typed λ -calculus with support for typed self-representation. The language is defined in Figure 14. It extends System F_ω with iso-recursive types and intensional type functions, which together are useful for typed meta-programming. It is type-safe and type checking is decidable. For more details, we refer the interested reader to Brown and Palsberg [Brown and Palsberg 2017].

Our self-representation type Exp is defined in Figure 15. It is a tagless-final style representation – intuitively, a typed version of Mogensen’s Church-encoded representation. It is also Parametric Higher-Order Abstract Syntax (PHOAS) [Chlipala 2008; Washburn and Weirich 2003]. In particular, the type PExp is parametric in V , which determines the type of free variables in a representation. Intuitively, PExp can be understood as the type of representations that may contain free variables. The type Exp quantifies over V , which ensures that the representation is closed. The self-representation of $F_\omega^{\mu i}$ presented by Brown and Palsberg was a typed Mogensen-Scott representation – intuitively, a Generalized Algebraic Data Type (GADT) encoded using λ -terms. We instead use the tagless-final style because it supports a normalizing self-interpreter and a normalizing self-applicable partial evaluator, allowing us to generate the Futamura projections.

Quotation for our $F_\omega^{\mu i}$ representation is defined in Figure 16. The quotation function $\bar{\cdot}$ is defined only on closed terms, and depends on a pre-quotation function \triangleright from type derivations to terms. In the judgment $\Gamma \vdash e : T \triangleright q$, the input is the type derivation $\Gamma \vdash e : T$, and the output is a term q . We call q the pre-representation of e . In addition to the cases for variables, λ -abstractions and applications, we have cases for type abstraction and application and the folding and unfolding of iso-recursive types. The case of type abstractions and type applications use the $\text{IsAll } T$ proof terms, which prove that the type T is a quantified type. They also use utility functions stripAll_K , $\text{underAll}_{X,K,T}$, and $\text{inst}_{X,K,T,S}$, which are useful to meta-programs for operating on type abstractions and applications. These IsAll proofs and utility functions are part of Brown and Palsberg’s extensional representation of polymorphism.

We prove two theorems about our representation: first, every closed and well-typed term has a representation that is closed and well-typed, and the type of a term determines the type of its representation. Ill-typed terms cannot be represented. Second, representations are β -normal forms.

THEOREM 6.1. *If $\langle \rangle \vdash e : T$, then $\langle \rangle \vdash \bar{e} : \text{Exp } T$.*

THEOREM 6.2. *If $\langle \rangle \vdash e : T$, then \bar{e} is β -normal.*

(kinds) $K ::= *$	$ K_1 \rightarrow K_2$	$ \forall X:K. T$	$ \lambda X:K. T \mid T_1 T_2$	$ \text{Typecase}$
(types) $T ::= X$	$ T_1 \rightarrow T_2$	$ \lambda x:T. e$	$ e_1 e_2$	$ \text{fold } T_1 T_2 e \mid \text{unfold } T_1 T_2 e$
(terms) $e ::= x$	$ \lambda x:T. e$	$ \Gamma, (X:K)$		
(environments) $\Gamma ::= \langle \rangle$				
(normal form terms) $v ::= n \mid (\lambda x:T. v) \mid (\Lambda X:K. v) \mid \text{fold } T_1 T_2 v$				
(neutral terms) $n ::= x \mid n v \mid n T \mid \text{unfold } T_1 T_2 n$				
Grammar				
$\frac{(X:T) \in \Gamma}{\Gamma \vdash x : T}$ $\frac{(X:K) \in \Gamma}{\Gamma \vdash X : K}$ $\frac{\Gamma \vdash T_1 : * \quad \Gamma \vdash T_2 : *}{\Gamma \vdash T_1 \rightarrow T_2 : *}$ $\frac{\Gamma, (X:K) \vdash T : *}{\Gamma \vdash (\forall X:K. T) : *}$ $\frac{\Gamma, (X:K_1) \vdash T : K_2}{\Gamma \vdash (\lambda X:K_1. T) : K_1 \rightarrow K_2} \quad \frac{\Gamma \vdash T_1 : K_2 \rightarrow K \quad \Gamma \vdash T_2 : K_2}{\Gamma \vdash T_1 T_2 : K}$ $\Gamma \vdash \mu : ((* \rightarrow *) \rightarrow * \rightarrow *) \rightarrow * \rightarrow *$ $\Gamma \vdash \text{Typecase} : (* \rightarrow * \rightarrow *) \rightarrow$ $(* \rightarrow *) \rightarrow (* \rightarrow *) \rightarrow$ $(((* \rightarrow *) \rightarrow * \rightarrow *) \rightarrow * \rightarrow *) \rightarrow$ $*$				
Type Formation				
$T \equiv T$ $\frac{T_1 \equiv T_2 \quad T_2 \equiv T_1}{T_1 \equiv T_1}$ $\frac{T_1 \equiv T_2 \quad T_2 \equiv T_3}{T_1 \equiv T_3}$ $\frac{T_1 \equiv T_1' \quad T_2 \equiv T_2'}{T_1 \rightarrow T_2 \equiv T_1' \rightarrow T_2'}$ $\frac{T \equiv T'}{(\forall X:K. T) \equiv (\forall X:K. T')}$ $\frac{T \equiv T'}{(\lambda X:K. T) \equiv (\lambda X:K. T')}$ $(\lambda X:K. T_1) T_2 \equiv (T_1[X := T_2])$ $(\forall X:K. T_2) \equiv (\forall X':K. T_2[X := X'])$ $(\lambda X:K. T) \equiv (\lambda X':K. T[X := X'])$ $\text{Typecase } F_1 F_2 F_3 F_4 (T_1 \rightarrow T_2) \equiv F_1 T_1 T_2$ $\text{Typecase } F_1 F_2 F_3 F_4 (\mu T_1 T_2) \equiv F_4 T_1 T_2$ $X \notin \text{FV}(F_3)$ $\text{Typecase } F_1 F_2 F_3 F_4 (\forall X:K. T) \equiv F_2 (\forall X:K. F_3 T)$				
Term Formation				
$\frac{(x:T) \in \Gamma \quad \Gamma \vdash x : T}{\Gamma \vdash x : T}$ $\frac{\Gamma \vdash T_1 : * \quad \Gamma, (x:T_1) \vdash e : T_2}{\Gamma \vdash (\lambda x:T_1. e) : T_1 \rightarrow T_2}$ $\frac{\Gamma \vdash e_1 : T_2 \rightarrow T \quad \Gamma \vdash e_2 : T_2}{\Gamma \vdash e_1 e_2 : T}$ $\frac{\Gamma, (X:K) \vdash e : T}{\Gamma \vdash (\Lambda X:K. e) : (\forall X:K. T)}$ $\frac{\Gamma \vdash e : (\forall X:K. T_1) \quad \Gamma \vdash T_2 : K}{\Gamma \vdash e : T_1[X := T_2]}$ $\frac{\Gamma \vdash F : (* \rightarrow *) \rightarrow * \rightarrow * \quad \Gamma \vdash T : *}{\Gamma \vdash e : F(\mu F) T}$ $\frac{\Gamma \vdash \text{fold } F T e : F T}{\Gamma \vdash F : (* \rightarrow *) \rightarrow * \rightarrow * \quad \Gamma \vdash T : *}$ $\frac{\Gamma \vdash e : \mu F T}{\Gamma \vdash \text{fold } F T e : F T}$ $\frac{\Gamma \vdash F : (* \rightarrow *) \rightarrow * \rightarrow * \quad \Gamma \vdash T : *}{\Gamma \vdash e : \mu F T}$ $\frac{\Gamma \vdash \text{unfold } F T e : F(\mu F) T}{\Gamma \vdash e : T_1 \quad T_1 \equiv T_2 \quad \Gamma \vdash T_2 : *}$ $\Gamma \vdash e : T_2$				
Type Equivalence				
$(\lambda x:T. e) e_1 \rightarrow e[x := e_1]$ $(\Lambda X:K. e) T \rightarrow e[X := T]$ $\text{unfold } F T (\text{fold } F' T' e) \rightarrow e$ $e_1 \rightarrow e_2$ $e_1 e_3 \rightarrow e_2 e_3$ $e_3 e_1 \rightarrow e_3 e_2$ $e_1 T \rightarrow e_2 T$ $(\lambda x:T. e_1) \rightarrow (\lambda x:T. e_2)$ $(\Lambda X:K. e_1) \rightarrow (\Lambda X:K. e_2)$ $\text{fold } F T e_1 \rightarrow \text{fold } F T e_2$ $\text{unfold } F T e_1 \rightarrow \text{unfold } F T e_2$				
Reduction				

Fig. 14. Definition of $F_{\omega}^{\mu i}$

The self-interpreter unquote for the tagless-final HOAS representation is shown in Figure 17. We prove that our partial evaluator is Jones-optimal for unquote and a $time(-)$ function based on call-by-value type-erasure semantics.

Definition 6.1 (Type erasure). The type erasure of an $F_{\omega}^{\mu i}$ term e is defined recursively as follows:

```

decl PExp : (* → *) → * → * =
  λV : * → *. λA:*.
    (∀S:*. ∀T:*. (V S → V T) → V (S → T)) →
    (∀A:*. ∀B:*. V (A → B) → V A → V B) →
    (∀A:*. IsAll A → StripAll A → UnderAll A → All Id V A → V A) →
    (∀A:*. ∀B:*. IsAll A → Inst A B → V A → V B) →
    (∀F : (* → *) → * → *. ∀B : *. V (F (μ F) B) → V (μ F B)) →
    (∀F : (* → *) → * → *. ∀B : *. V (μ F B) → V (F (μ F) B)) →
    V A

decl Exp : * → * = λA:*. ∀V:* → *. PExp V A

decl Abs V = ∀A1:*. ∀A2:*. (V A1 → V A2) → V (A1 → A2)
decl App V = ∀A:*. ∀B:*. V (A → B) → V A → V B
decl TAbs V =
  ∀A:*. IsAll A → StripAll A → UnderAll A → All Id V A → V A
decl TApp V = ∀A:*. ∀B:*. IsAll A → Inst A B → V A → V B
decl Fold V =
  ∀F : (*→*) → * → *. ∀B:*. V (F (μ F) B) → V (μ F B)
decl Unfold V =
  ∀F : (*→*) → * → *. ∀B : *. V (μ F B) → V (F (μ F) B)

```

Fig. 15. Definitions of PExp and Exp

1. $te(x) = x$
2. $te(\lambda x:T. e) = \lambda x. te(e)$
3. $te(e_1 e_2) = te(e_1) te(e_2)$
4. $te(\lambda X:K. e) = te(e)$
5. $te(e \top) = te(e)$
6. $te(\text{fold } T_1 T_2 e) = te(e)$
7. $te(\text{unfold } T_1 T_2 e) = te(e)$

Definition 6.2. $time_{cbv}^{te}(e) = n$ if and only if $te(e) \rightarrow^n v$.

LEMMA 6.3. *If $\langle \rangle \vdash e : T$, then $te(\text{unquote } \bar{e}) \Longrightarrow_s te(e)$.*

THEOREM 6.4. *mix is Jones-optimal for $time_{cbv}^{te}$ and $\text{unquote}/\bar{\cdot}$.*

6.1 Implementation of mix

Our partial evaluator `mix` is implemented in two steps: first, we analyze the program and input for affine variable annotations, then we reduce the annotated term to specialization-safe normal form. Each of these steps is programmed by folds over the representation. To infer affine variable annotations, we do a fold that returns for each subterm 1) a tree data structure that indicates the annotation of each bound variable in the subterm, 2) how many times each free variable occurs, and 3) which free variables occur under a λ -abstraction. Items 2 and 3 are used to determine the annotation of each variable at its binding λ -abstraction. A second fold uses the annotation tree to annotate the term. A third and final fold reduces the annotated term to specialization-safe normal form. This step is based on modification of a Mogensen's partial evaluator that reduces to β -normal form [Mogensen 1995], which was also used as the basis for a partial evaluator by Carette et al. [Carette et al. 2009]. Our modifications check whether each β -redex is specialization-safe by

$$\begin{array}{c}
\frac{(x : T) \in \Gamma}{\Gamma \vdash x : T \triangleright x} \\
\frac{\Gamma \vdash T1 : * \quad \Gamma, (x:T1) \vdash e : T2 \triangleright q}{\Gamma \vdash (\lambda x:T1. e) : T1 \rightarrow T2 \triangleright \text{abs } T1 \ T2 \ (\lambda x:V \ T1. \ q)} \\
\frac{\Gamma \vdash e_1 : T_2 \rightarrow T \triangleright q_1 \quad \Gamma \vdash e_2 : T_2 \triangleright q_2}{\Gamma \vdash e_1 \ e_2 : T \triangleright \text{app } T_2 \ T \ q_1 \ q_2} \\
\frac{\Gamma, (X:K) \vdash e : T \triangleright q \quad \begin{array}{l} \text{isAll}_{X,K,T} = p \\ \text{stripAll}_K = s \\ \text{underAll}_{X,K,T} = u \end{array}}{\Gamma \vdash (\Lambda X:K. e) : (\forall X:K. T) \triangleright \text{tabs } (\forall X:K. T) \ p \ s \ u \ (\Lambda X:K. q)} \\
\frac{\Gamma \vdash e : (\forall X:K. T) \triangleright q \quad \begin{array}{l} \text{isAll}_{X,K,T} = p \\ \text{inst}_{X,K,T,A} = i \end{array} \quad \Gamma \vdash A : K}{\Gamma \vdash e \ A : T[X:=A] \triangleright \text{tapp } (\forall X:K. T) \ (T[X:=A]) \ p \ i \ q} \\
\frac{\Gamma \vdash F : (* \rightarrow *) \rightarrow * \rightarrow * \quad \Gamma \vdash T : * \quad \Gamma \vdash e : F (\mu F) T \triangleright q}{\Gamma \vdash \text{fold } F \ T \ e : \mu F \ T \triangleright \text{fld } F \ T \ q} \\
\frac{\Gamma \vdash F : (* \rightarrow *) \rightarrow * \rightarrow * \quad \Gamma \vdash T : * \quad \Gamma \vdash e : \mu F \ T \triangleright q}{\Gamma \vdash \text{unfold } F \ T \ e : F (\mu F) T \triangleright \text{unfld } F \ T \ q} \\
\frac{\langle \rangle \vdash e : T \triangleright q}{\bar{e} = \Lambda V : * \rightarrow *. \\ \lambda \text{abs} : \text{Abs } V. \ \lambda \text{app} : \text{App } V. \\ \lambda \text{tabs} : \text{TAbs } V. \ \lambda \text{tapp} : \text{TApp } V. \\ \lambda \text{fld} : \text{Fold } V. \ \lambda \text{unfld} : \text{Unfold } V. \\ q}
\end{array}$$

Fig. 16. Tagless-final quotation and pre-quotation of $F_{\omega}^{\mu i}$

inspecting the affine variable annotation of the λ -abstraction and the form of the argument. If it's a value or an unlimited variable, the β -redex is specialization-safe. If it's an application or an affine variable, the β -redex is not specialization-safe and is residualized.

Figure 18 shows some highlights of our Jones-optimal partial evaluator.

At its core is the function `ssnorm`, which reduces annotated representations of terms to (unannotated) representations of their specialization-safe normal form. The function `sem` is a fold that converts an `Exp A` to a semantic object of type `Sem V A`, and `reify` extracts a representation of the evaluated term from the semantic object. The semantic object is a pair containing the specialization-safe normal form of the term and a helper function. `sem` traverses the term and creates a semantic object for each subterm. When a subterm occurs in head position of a specialization-safe redex, its helper function is used to compute the reduct.

At the bottom of the figure is `mix` itself. It first builds `e`, a representation of the program applied to the input. Then it uses `forceExp` to check that the input `x` is available, in which case `forceExp` returns `e` to be annotated and normalized. This check is an important optimization for self-application, specifically when specializing `mix` to a single input. It prevents a blow up in the output code size caused by many copies of the traversal functions from `annotate` and `sem` being residualized throughout `f`.

```

decl Id : * → * = λA:*. A;

decl unAbs : Abs Id =
  ΛA:*. ΛB:*. λf : A → B. f;
decl unApp : App Id =
  ΛA:*. ΛB:*. λf : A → B. f;
decl unTAbs : TAbs Id =
  ΛA:*. λp:IsAll A. λs:StripAll A. λu:UnderAll A. λe:All Id Id A.
  unAll A p Id Id e;
decl unTApp : TApp Id =
  ΛA:*. ΛB:*. λp : IsAll A. λi : Inst A B. λx : A.
  i Id (sym (All Id Id A) A (unAll A p Id) Id x);
decl unFold : Fold Id =
  ΛF: (* → *) → * → *. ΛA:*. λx : F (μ F) A. fold F A x;
decl unUnfold : Unfold Id =
  ΛF: (* → *) → * → *. ΛA:*. λx : μ F A. unfold F A x;

decl unquote : ∀(A:*. Exp A → A) =
  ΛA:*. λe:Exp A. e Id unAbs unApp unTAbs unTApp unFold unUnfold;

```

Fig. 17. The $F_{\omega}^{\mu i}$ tagless-final HOAS self-interpreter unquote.

The type Sem of semantic objects is shown in Figure 19. It is an iso-recursive intensional type function, combining two features of $F_{\omega}^{\mu i}$: iso-recursive types and intensional type functions. The type $\text{Sem } V \text{ A}$ is equivalent to $(\text{SemF } V) \text{ A}$ and isomorphic (via (un)folding the recursion) to $\text{SemF } (\text{Sem } V) \text{ A}$.

There are two variants of semantic object: neutral objects that do not form redexes (specialization-safe or otherwise) when in head-position, and non-neutral objects that do. Neutral objects contain a boolean and a representation. The representation is used to reify the semantic object. The boolean is true if that representation is of a term that type-erases to an unlimited variable. If so, then any redex in which this term is in argument position is specialization-safe.

Non-neutral objects form redexes when in head-position. Their first two components are similar to those of neutral objects. The boolean is true if the representation erases to a λ -abstraction. The third component is used to compute the reduct of redexes with this object in head-position. These have types of the form $\text{SemF1 } (\text{Sem } V) \text{ A}$. SemF1 is an intensional type function that depends on the structure of A :

$$\begin{aligned}
\text{SemF1 } (\text{Sem } V) (A \rightarrow B) &\equiv \text{Pair Bool } (\text{Sem } V \text{ A} \rightarrow \text{Sem } V \text{ B}) \\
\text{SemF1 } (\text{Sem } V) (\forall X:K. T) &\equiv \forall X:K. \text{Sem } V \text{ T} \\
\text{SemF1 } (\text{Sem } V) (F \text{ A}) &\equiv \text{Sem } V (F (F) \text{ A})
\end{aligned}$$

A non-neutral semantic object of arrow type is a λ -abstraction. For these, the third component is a pair of boolean and function from semantic objects to semantic objects. The boolean indicates whether the λ -abstraction is affine. When this semantic object is in head position of a redex, the function is applied to the semantic object of the argument to compute the semantic object of the reduct.

If a non-neutral semantic object has a quantified type, then it is a Λ -abstraction. For these, the third component is a polymorphic semantic object. If a non-neutral semantic object has a recursive type, it is for a fold term $\text{fold } F \text{ A } e$, and the third component is the semantic object of the unfolded term e .

```

annotate : (∀V:* → *. ∀A : *. Exp A → ExpAnn Bool A)
sem       : (∀V:* → *. ∀A : *. ExpAnn Bool A → Sem V A)
reify    : (∀V:* → *. ∀A : *. Sem V A → PExp V A)
forceExp : (∀T:* . Exp T → (∀A:* . A → A))

decl ssnorm : (∀A : *. ExpAnn Bool A → Exp A) =
  ΛA:* . λe : Exp A. ΛV:* → *.
  reify V A (sem V A e);

decl mix : (∀A : *. ∀B : *. Exp (A → B) → Exp A → Exp B) =
  ΛA:* . ΛB:* . λf:Exp (A → B). λx:Exp A.
  let e : Exp B = (ΛV:* → *. app V A B (f V) (x V)) in
  ssnorm B (annotate B (forceExp A x (Exp B) e));

```

Fig. 18. Highlights of our Jones-optimal partial evaluator mix

```

decl SemF1 : (* → *) → * → * =
  λSem : * → *. λA:* .
  Typecase
    (λA1:* . λA2:* . Pair Bool (Sem A1 → Sem A2))
  Id Sem
  (λF : (* → *) → * → *. λB : *.
    Sem (F (μ F) B))
  A
;

decl SemF : (* → *) → (* → *) → * → * =
  λV : * → *. λSem : * → *. λA : *.
  Triple Bool (PExp V A) (Maybe (SemF1 Sem A))
;

decl Sem : (* → *) → * → * = λV : * → *. μ (SemF V)

```

Fig. 19. Internal types of our Jones-optimal partial evaluator.

Figure 20 shows the function `semApp`, the part of `sem` that constructs semantic objects for applications. Given semantic objects for a function `f` and an argument `x`, `semApp` first unpacks `f` and binds the names `fValueOrUnlimited`, `rep_f`, and `msem` to its components. `fValueOrUnlimited` is not used, `rep_f` is the representation of the specialization-safe normal form of `f`, and `msem` indicates whether `f` is neutral or non-neutral. If `f` is neutral, then the application is not a redex and is residualized. Otherwise, the application is a β -redex. To check if it's specialization-safe, we test whether `f`'s bound variable is affine, or `x` is either a value or an unlimited variable. If either case is true, the β -redex is specialization-safe and we reduce. Otherwise, we residualize. Whenever we residualize, we use `semNe` to construct a semantic object for the (neutral) residual term. The argument `false` to `semNe` indicates that result is not a value or an unlimited variable.

6.2 Type-checking the Futamura Projections

Figure 21 shows typed versions of the Futamura projections in the concrete syntax of F_{ω}^{μ} . Quotations are denoted using square brackets `[-]` rather than overlines, and so double-quotations

```

decl semApp : (∀V:*→*.∀B:*.∀A:*. Sem V (B → A) → Sem V B → Sem V A) =
  ΛV:* → *. ΛB:*. ΛA:*. λf:Sem V (B → A). λx:Sem V B.
  unfold (SemF V) (B → A) f
  (Sem V A)
  (λfValueOrUnlimited : Bool.
  λrep_f : PExp V (B → A)).
  λmsem : Maybe (Pair Bool (Sem V B → Sem V A)).
  msem (Sem V A)
  -- neutral
  (semNe V A false (app V B A rep_f (reify V B x)))
  -- not neutral
  (λp : Pair Bool (Sem V B → Sem V A).
  let affine : Bool = fst Bool (Sem V B → Sem V A) p in
  let sem_f : Sem V B → Sem V A = snd Bool (Sem V B → Sem V A) p in
  or affine (isValueOrUnlimited V B x)
  (Sem V A)
  -- redex is specialization-safe: reduce
  (sem_f x)
  -- redex is not safe: residualize
  (semNe V A false (app V B A rep_f (reify V B x))))))

```

Fig. 20. A key function from sem.

are denoted using nested square brackets $[[-]]$ rather than stacked overlines $\overline{\overline{\quad}}$. The first projection compiles the factorial function `fact` by specializing the self-interpreter `unquote` to it. The second projection generates a compiler by specializing `mix` to `unquote`. The third projection generates a compiler-generator by specializing `mix` to itself. The fourth projection generates the self-generating compiler-generator. It is also obtained by specializing `mix` to itself, but at different types.

This is the first time the Futamura projections have been type-checked for a partial evaluator that operates on typed representations. One consequence of this is that we can clearly see what the types of generated programs are. In particular, the type of each generated program is determined by the type of its representation.

The compiled factorial function `fact_compiled` has the expected type $\text{Nat} \rightarrow \text{Nat}$. The generated compiler `compile` has the polymorphic type $\forall A:*. \text{Exp} (\text{Exp } A) \rightarrow \text{Exp } A$. It compiles programs by mapping double representations to representations (like the First futamura projection does). The generated compiler generator `cogen` has the polymorphic type

$$\forall A:*. \forall B:*. \text{Exp} (\text{Exp} (A \rightarrow B)) \rightarrow \text{Exp} (\text{Exp } A \rightarrow \text{Exp } B).$$

If we instantiate A to $\text{Exp } C$ and B to C , we can derive the type:

$$\forall C:*. \text{Exp} (\text{Exp} (\text{Exp } C \rightarrow C)) \rightarrow \text{Exp} (\text{Exp} (\text{Exp } C) \rightarrow \text{Exp } C).$$

Then we can apply `cogen` to the (double) representation of a self-interpreter like `unquote` and generate a compiler. The self-generating compiler generator `selfgen` has the polymorphic type:

$$\begin{aligned} &\forall A:*. \forall B:*. \\ &\text{Exp} (\text{Exp} (\text{Exp} (A \rightarrow B) \rightarrow \text{Exp } A \rightarrow \text{Exp } B)) \rightarrow \\ &\text{Exp} (\text{Exp} (\text{Exp} (A \rightarrow B)) \rightarrow \text{Exp} (\text{Exp } A \rightarrow \text{Exp } B)) \end{aligned}$$

We can apply `selfgen` to the (double) representation of `mix`, and generate `cogen` again.

```

decl fact_compiled : Exp (Nat → Nat) =
  mix (Exp (Nat → Nat)) (Nat → Nat)
    [unquote (Nat → Nat)]
    [[fact]]

decl compile : ∀A:*. Exp (Exp (Exp A) → Exp A) =
  ΛA:*.
  mix (Exp (Exp A → A)) (Exp (Exp A) → Exp A)
    [mix (Exp A) A]
    [[unquote A]]

decl cogen : (∀A:*. ∀B:*. Exp (Exp (Exp (A → B)) →
                               Exp (Exp A → Exp B))) =
  ΛA:*. ΛB:*.
  mix (Exp (Exp (A → B) → Exp A → Exp B))
    (Exp (Exp (A → B)) → Exp (Exp A → Exp B))
    [mix (Exp (A → B)) (Exp A → Exp B)]
    [[mix A B]]

decl selfgen : (∀A:*. ∀B:*.
               Exp (Exp (Exp (Exp (A → B) → Exp A → Exp B)) →
                     Exp (Exp (Exp (A → B)) →
                               Exp (Exp A → Exp B)))) =
  ΛA:*. ΛB:*.
  mix (Exp (Exp (Exp (A → B) → Exp A → Exp B) →
            Exp (Exp (A → B)) →
            Exp (Exp A → Exp B)))
    (Exp (Exp (Exp (A → B) → Exp A → Exp B)) →
     Exp (Exp (Exp (A → B)) → Exp (Exp A → Exp B)))
    [mix (Exp (Exp (A → B) → Exp A → Exp B))
      (Exp (Exp (A → B)) → Exp (Exp A → Exp B))]
    [[mix (Exp (A → B)) (Exp A → Exp B)]]

```

Fig. 21. The four typed Futamura projections

7 EXPERIMENTAL RESULTS

We implemented $F_{\omega}^{\mu i}$ in Haskell, including a parser, a type checker, a β -equivalence checker, and three different *time* measures. We tested that our partial evaluator type checks and the Futamura projections type check with the types given in Figure 21. We used our β -equivalence checker to verify that the Futamura projections meet their specification. In addition to formally proving that our partial evaluator is Jones-optimal for call-by-value reduction, our experiments demonstrate Jones-optimality for call-by-value, normal-order, and memoized normal-order reduction. We also demonstrate that specialization by β -normalization is Jones-optimal for normal-order reduction, but not for call-by-value or memoized normal-order.

Reduction strategies. We measure call-by-value steps to a value, and normal-order and memoized normal-order steps to a β -normal form. Normal-order reduction is related to call-by-name reduction, but evaluates to a β -normal form instead of a value. Counting call-by-name reduction steps to a value is not a useful measure, because many of our benchmark programs reduce to a value in only a few steps of call-by-name reduction, obscuring the effect of specialization on the

Interpreter	Time	Slowdown	Overhead
Tagless-final PHOAS	Call-by-value	10.0	900%
	Normal order	10.9	990%
	Memoized normal order	9.9	890%
Mogensen-Scott PHOAS	Call-by-value	105.3	10430%
	Normal order	136.5	13550%
	Memoized normal order	68.8	6780%
Tagless-final deBruijn	Call-by-value	24.4	2340%
	Normal order	118.3	11730%
	Memoized normal order	24.8	2380%

Table 2. Average interpretational overhead of our self-interpreters

run time. Memoized normal-order reduction is similarly related to call-by-need (lazy) reduction. It reduces a β -redex without first reducing its parameter; parameters are only reduced if and when they need to be. When a parameter is reduced to a value, the value is saved in case it is needed again later. We do not memoize the β -normal form of parameters, so as we normalize a term, the work to normalize each parameter's value could be repeated.

Definition 7.1. $time_{no}^{te}(e) = n$ if $te(e)$ normalizes in n steps of normal-order reduction.

Definition 7.2. $time_{mno}^{te}(e) = n$ if $te(e)$ normalizes in n steps of memoized normal-order reduction.

Self-interpreters. We experiment with three typed self-interpreters for $F_{\omega}^{\mu_i}$ – essentially typed versions of the untyped λ -calculus self-interpreters in Figures 4, 12, and 13. Each of these interpreters has a different amount of interpretational overhead, which also depends on which *time* function is used. We measured the mean interpretational slowdown and overhead for the cube, factorial, and Ackermann functions on a range inputs. The tagless-final PHOAS interpreter is the most efficient, with an interpretational overhead of around 900% – corresponding to a 10X slowdown – for each evaluation strategy. The Mogensen-Scott PHOAS interpreter is the least efficient, with an overhead of from 6780% for memoized normal order and 13550% for normal order. The tagless-final deBruijn interpreter is in between, with an overhead of from 2340% for call-by-value and 11730% for normal order.

Jones-optimality. Our partial evaluator is Jones-optimal for all 9 combinations of our three self-interpreters and three evaluation strategies. Table 3 gives a summary of our Jones optimality tests for our three $F_{\omega}^{\mu_i}$ self-interpreters and three evaluation strategies. Each row aggregates the result of a set of benchmarks comparing the *time* cost of each subject program with the specialization of the self-interpreter to that program. Again, we use the cube, factorial, and Ackermann functions for our subject programs, and we test with a range of inputs for each. A speedup of 1.00 indicates that the subject program runs in exactly the same amount of time as the specialization. This usually happens when the two programs are exactly the same; when specialization removes all the interpretational overhead and does nothing else, recovering the original program. A speedup greater than 1 indicates that the specialized program is faster than the original, and a speedup less than 1 indicates a slowdown. The table demonstrates that our partial evaluator is Jones-optimal for each combination: in each case, the minimum speedup is 1. It also shows that we can sometimes achieve speedups beyond removing interpretational overhead – up to approximately 1.5X for our benchmarks.

Interpreter	Time	Min Speedup	Mean Speedup	Max Speedup
Tagless-final PHOAS	Call-by-value	1.00	1.06	1.36
	Normal order	1.00	1.14	1.41
	Memoized normal order	1.00	1.06	1.33
Mogensen-Scott PHOAS	Call-by-value	1.00	1.14	1.50
	Normal order	1.00	1.34	1.54
	Memoized normal order	1.00	1.12	1.46
Tagless-final deBruijn	Call-by-value	1.00	1.16	1.50
	Normal order	1.00	1.35	1.54
	Memoized normal order	1.00	1.44	1.46

Table 3. Summary of Jones Optimality Tests

Interpreter	Time	Min Speedup	Mean Speedup	Max Speedup
Tagless-final PHOAS	Call-by-value	0.03	0.74	1.89
	Normal order	1.00	1.55	1.92
	Memoized normal order	0.11	0.80	1.80
Mogensen-Scott PHOAS	Call-by-value	0.03	0.74	1.89
	Normal order	1.00	1.55	1.92
	Memoized normal order	0.11	0.80	1.80
Tagless-final deBruijn	Call-by-value	0.06	0.83	1.89
	Normal order	1.00	1.53	1.92
	Memoized normal order	0.17	0.89	1.80

Table 4. Demonstration of Non-Optimality of Specialization by β -Normalization

Interpreter	Reduction	Min Speedup	Mean Speedup	Max Speedup
Tagless-final PHOAS	Call-by-value	1.27	1.29	1.29
	Normal order	1.35	1.37	1.39
	Memoized normal order	1.26	1.27	1.27
Mogensen-Scott PHOAS	Call-by-value	1.07	1.08	1.09
	Normal order	1.29	1.30	1.30
	Memoized normal order	1.10	1.12	1.12
Tagless-final deBruijn	Call-by-value	1.32	1.33	1.34
	Normal order	1.38	1.39	1.39
	Memoized normal order	1.27	1.27	1.28

Table 5. Summary of Second Futamura Projection Tests

As a point of contrast, Table 4 shows the results of the same Jones-optimality tests for a specialization by normalization partial evaluator for F_{ω}^{hi} . The table demonstrates that specialization by β -normalization is only Jones optimal for normal order reduction, in which the value of a variable is always recomputed for each reference. Therefore, all β -reductions are specialization-safe for normal order reduction. Specialization by β -normalization can cause significant slowdowns for call-by-value or memoized normal order reduction – the 0.03X speedup for call-by-value corresponds to a 33X slowdown – the original program runs 33X faster than the specialized one.

Futamura Projections. The Jones optimality benchmarks in Table 3 show the speedups obtained by the first Futamura projection. We also measure the speedups produced by the second Futamura

Time	Min Speedup	Mean Speedup	Max Speedup
Call-by-value	1.16	1.17	1.19
Normal order	1.37	1.37	1.38
Memoized normal order	1.21	1.22	1.22

Table 6. Summary of Third Futamura Projection Benchmarks

Time	Speedup
Call-by-value	1.22
Normal order	1.35
Memoized normal order	1.24

Table 7. Fourth Futamura Projection Benchmarks

projection, which are listed in Table 5. For each self-interpreter and each *time* function, we use the second Futamura projection to generate a compiler, and compare its running time with that of the first Futamura projection. We compile each of the cube, factorial, and Ackermann functions.

We measure the speedup obtained by the third Futamura projection by comparing the *time* to generate compilers using *cogen* versus the second Futamura projection. For each evaluation strategy, we generate compilers using each of our self-interpreters. The results are summarized in Table 6.

We measure the speedup obtained by the fourth Futamura projection by comparing the *time* to generate a compiler generator using *selfgen* versus the third Futamura projection. The speedups are shown in Table 7.

8 RELATED WORK

The following table compares our work with five classical papers in the three main dimensions that we consider. After the table, we give a discussion of a variety of related work.

	Typed representation	Futamura projections	Jones optimality
Original mix [Jones et al. 1985]		√	
Lambda-mix [Gomard and Jones 1991]		√	
Similix [Bondorf and Danvy 1991]		√	√
Schism [Consel 1993]		√	
TDPE [Danvy 1996]		√	
This paper	√	√	√

Typed representation. We were inspired by the considerable body of work on typed self-representation and self-interpretation. The classical approach to meta-programming in a statically-typed language is to use a single *universal* type for all program representations. This incurs a tagging overhead and allows type errors in generated code. On the other hand, typed representation can eliminate tags and ensure that only well-typed programs can be represented. Further, it can provide strong correctness properties for meta-programs just by type checking them. For example, the type of our partial evaluator ensures that it always generates well-typed code of the correct type. At the core of this line of work is the challenge to define a typed self-representation. In their seminal paper, Pfenning and Lee [1991] showed how to represent System F in F_{ω} , and how to

represent F_ω in F_ω^+ , but fell short of a self-representation. Rendel et al. [2009] and Jay and Palsberg [2011] presented the first typed self-representations, for two calculi with undecidable type checking. Brown and Palsberg [2015, 2016] defined typed self-representations for two calculi with decidable type checking, the latter of which is also strongly normalizing.

Most closely related to this work is our previous work [Brown and Palsberg 2017] that defined typed self-representations and three typed self-evaluators. Each of those self-evaluators implements a particular evaluation strategy for the calculus itself. The key insight in that paper was a way to encode type equality proofs for a GADT-style self-representation. In this paper, we use the same language and some of the representation techniques first presented there, but we use a tagless-final self-representation instead of the GADT-style encoding, which is important for enabling our partial evaluator to generate the Futamura projections.

Futamura Projections. A variety of partial evaluators for untyped languages have enabled the three classical Futamura projections including the original mix [Jones et al. 1985], Lambda-mix [Gomard and Jones 1991], Schism [Consel 1993], and Similix [Bondorf and Danvy 1991]. Glück [1999, 2009] popularized the fourth Futamura projection and showed how to achieve it. Launchbury [1991] described a partial evaluator for LML, which operates on a universally typed representation and supports the first and second Futamura projections. Danvy [1996] presented type-directed partial evaluation for a simply-typed language and showed how it enables the first and second Futamura projections.

Carette et al. [2009] implemented a typed partial evaluator that operated on a typed tagless-final representation. Their object language was the simply-typed λ -calculus extended with integers, booleans, and a fixpoint operator, and their meta-languages was MetaOCaml. Their partial evaluator could not generate the Futamura projections and was not Jones-optimal.

Our partial evaluator is the first that operates on a typed representation and can be self-applied. Typed representation ensures that the partial evaluator always generates well-typed code of the correct type. This is a strong correctness properties guaranteed just by type checking.

Jones Optimality. Similix [Bondorf and Danvy 1991] achieved Jones optimality for an untyped language. Makhholm [2000] discussed the challenges for Jones optimality for simply-typed languages when the partial evaluator relies on a universal type for all program representations. Later, Taha et al. [2001], and also Danvy and López [2003], achieved Jones optimality for simply-typed languages. Gluck [2002, 2008] discussed the effect of semantics-preserving program transformations ahead of partial evaluation. He showed that such transformations can enable a Jones optimal *offline* partial evaluator to achieve the effect of an *online* partial evaluator.

Some papers [Barker et al. 2007; Feigin and Mycroft 2008; Gade and Gluck 2006; Makhholm 2000] work with first-order languages and define Jones optimality via the time-based comparison that we state in Definition 2.2. Other papers [Danvy and López 2003; Skalberg 1999; Taha et al. 2001] work with higher-order functional languages and define Jones Optimality by checking that the specialization of the self-interpreter to a program produces the program itself (up to α -equivalence). This rules out the possibility that specialization could cause a slowdown, but also rules out the possibility of a speedup.

Our proof of Jones optimality is similar in spirit to previous proofs [Gade and Gluck 2006; Gomard and Jones 1991]. In particular, we make similar assumptions in choosing steps of call-by-value reduction as a measure of time – that all reductions have an equal cost, and that other factors such as program size do not affect running time. The key to our proof of Jones optimality is the specialization-safety of specialization-safe reduction. We prove this by reordering a sequence of specialization-safe steps followed by call-by-value steps to align with a sequence of only call-by-value steps. This reordering can only add steps, and the pure call-by-value sequence establishes

a maximum number of steps. This technique is inspired by a proof by Alex Simpson about linear λ -calculus [Simpson 2005].

Terminating Partial Evaluation. Asai et al. [2014] proved various correctness properties of a partial evaluator, including termination. The partial evaluator is implemented in a language that is more expressive than the language being partially evaluated, so self-application is impossible. Our work differs in that we do achieve self-application while we use experiments to demonstrate termination for important cases.

Other Related Work. Bondorf and Dussart [1994], Birkedal and Welinder [1994], and Thiemann [1996] have shown how to write strikingly simple compiler generators by hand for untyped or simply-typed languages. This is a major alternative to the use of a partial evaluator to produce a compiler generator via the third Futamura projection. However, the task to write a polymorphically-typed compiler generator by hand is significantly harder. It is an open question whether one can write a polymorphically-typed compiler generator by hand that is simpler than the one we have generated automatically.

Researchers have developed techniques for proving the semantic correctness of partial evaluation, including the recent paper by Hirota and Asai [2014].

Shali and Cook [2011] described a simply-typed partial evaluator for Java, which enables the first Futamura projection. Brady and Hammond [2010] described a partial evaluator for a dependently typed language that enables the first Futamura projection. It is an open question whether those approaches can be extended to enable the other Futamura projections.

Feigin and Mycroft [2008] showed that a version of Jones optimality coincides with a known criterion for efficient virtualization. Barker, Leuschel, and Varea showed how to build a Jones optimal partial evaluator for an untyped logic programming language [Barker et al. 2007].

9 CONCLUSION

In this work we presented the first self-applicable partial evaluator that operates on typed representations, generates the Futamura projections, and is Jones optimal. Operating on typed representations avoids tagging overhead and guarantees that it produces type-correct code. We use a novel approach to a partial-evaluator that focuses on specialization-safe β -reduction that is guaranteed never to cause a slowdown at runtime. We prove that our partial evaluator is Jones optimal for one self-interpreter and a measure of running time that counts steps of call-by-value reduction. Our experiments demonstrate that it is also Jones optimal for two other self-interpreters and time measures based on normal-order and memoized normal-order reduction.

Our results open new challenges for future work, including extending specialization-safe reduction to higher-level languages with efficient base types and user-defined types, and going beyond β -reduction to include more complex specialization-safe transformations. Another direction is to explore a complex affine variable analysis, which might lead to a better partial evaluator. Our affine analysis is sufficient to achieve Jones optimality and also simple enough to enable the Futamura projections. In particular, it can be implemented without a fixpoint combinator, which helps make our partial evaluator strongly-normalizing so the Futamura projections terminate.

One limitation of our partial evaluator is that specialization doesn't always terminate. Also, that our partial evaluator is Jones-optimal does not imply that its speedups are maximal. We leave for future work the challenge to implement a typed self-applicable and Jones-optimal partial evaluator that terminates on all inputs, and the question of identifying more powerful specialization-safe reduction relations that can be used to define a typed self-applicable and Jones-optimal partial evaluator with greater speedups.

ACKNOWLEDGMENTS

We thank Nada Amin, Jacques Carette, William Cook, Olivier Danvy, Danko Ilik, Neil Jones, Oleg Kiselyov, Philip Wadler, and the POPL reviewers for helpful discussions and suggestions. This material is based upon work supported by the National Science Foundation under Grant Number 1219240.

REFERENCES

- Lars Ole Andersen. Self-applicable C program specialization. In *Proceedings of PEPM'92, Workshop on Partial Evaluation and Semantics-Based Program Manipulation*, pages 54–61, June 1992. (Technical Report YALEU/DCS/RR-909, Yale University).
- Kenichi Asai, Luminous Fennell, Peter Thiemann, and Yang Zhang. A type theoretic specification of partial evaluation. In *PPDP'14 Proceedings of the 16th International Symposium on Principles and Practice of Declarative Programming*, pages 57–68, 2014.
- Steve Barker, Michael Leuschel, and Mauricio Varea. Efficient and flexible access control via Jones-optimal logic program specialisation. *Higher-Order and Symbolic Computation*, pages 3–35, 2007.
- Lars Birkedal and Morten Welinder. Hand-writing program generator generators. In *Programming Language Implementation and Logic Programming*, pages 198–214, 1994.
- Anders Bondorf and Olivier Danvy. Automatic autoprojection of recursive equations with global variables and abstract data types. *Science of Computer Programming*, 16:151–195, 1991.
- Anders Bondorf and Dirk Dussart. Improving CPS-based partial evaluation: Writing cogen by hand. In *Proceedings of ACM SIGPLAN Workshop on Partial Evaluation and Semantics-Based Program Manipulation*, pages 1–10, 1994.
- Edwin C. Brady and Kevin Hammond. Scrapping your inefficient engine: Using partial evaluation to improve domain-specific language implementation. In *Proceedings of ICFP'10, ACM SIGPLAN International Conference on Functional Programming*, 2010.
- Matt Brown and Jens Palsberg. Self-Representation in Girard's System U. In *Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '15, pages 471–484, New York, NY, USA, 2015. ACM. ISBN 978-1-4503-3300-9. doi: 10.1145/2676726.2676988. URL <http://doi.acm.org/10.1145/2676726.2676988>.
- Matt Brown and Jens Palsberg. Breaking through the normalization barrier: A self-interpreter for f-omega. In *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL 2016, pages 5–17, New York, NY, USA, 2016. ACM. ISBN 978-1-4503-3549-2. doi: 10.1145/2837614.2837623. URL <http://doi.acm.org/10.1145/2837614.2837623>.
- Matt Brown and Jens Palsberg. Typed self-evaluation via intensional type functions. In *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages*, POPL 2017, pages 415–428, New York, NY, USA, 2017. ACM. ISBN 978-1-4503-4660-3. doi: 10.1145/3009837.3009853. URL <http://doi.acm.org/10.1145/3009837.3009853>.
- Matt Brown and Jens Palsberg. Webpage accompanying this paper. The webpage accompanying this paper is available at <http://compilers.cs.ucla.edu/pop18/>. The full paper with the appendix is available there, as is the source code for our implementation of System F_{ω}^{Hi} and its self-interpreters and partial evaluator., 2018.
- Jacques Carette, Oleg Kiselyov, and Chung-chieh Shan. Finally tagless, partially evaluated: Tagless staged interpreters for simpler typed languages. *Journal of Functional Programming*, 19(5):509–543, 2009.
- Adam Chlipala. Parametric higher-order abstract syntax for mechanized semantics. In *Proceedings of the 13th ACM SIGPLAN International Conference on Functional Programming*, ICFP '08, pages 143–156, New York, NY, USA, 2008. ACM. ISBN 978-1-59593-919-7. doi: 10.1145/1411204.1411226. URL <http://doi.acm.org/10.1145/1411204.1411226>.
- Charles Consel. A tour of Schism: A partial evaluation system for higher-order applicative languages. In *Proceedings of PEPM'93, Second ACM SIGPLAN Symposium on Partial Evaluation and Semantics-Based Program Manipulation*, pages 145–154, 1993.
- Olivier Danvy. Type-directed partial evaluation. In *Proceedings of POPL'96, 23rd Annual SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 242–257, 1996.
- Olivier Danvy and Pablo E. Martínez López. Tagging, encoding, and Jones optimality. In *Proceedings of ESOP'03, European Symposium on Programming*, pages 335–347. Springer-Verlag (LNCS), 2003.
- Boris Feigin and Alan Mycroft. Jones optimality and hardware virtualization: a report on work in progress. In *PEPM*, pages 169–175, 2008.
- Yoshihiko Futamura. Evaluation of computation process—an approach to a compiler-compiler. *Higher-Order and Symbolic Computation*, 12:381–391, 1999. First published in 1971 in *System.Computers.Controls*, Volume 2, Number 5, pages 45–50.
- Johan Gade and Robert Gluck. On Jones-optimal specializers: A case study using unmix. In *Proceedings of APLAS'06, Asian Symposium on Programming Languages and Systems*, pages 406–422. Springer-Verlag (LNCS 4279), 2006.
- Robert Glück. Is there a fourth Futamura projection? In *Proceedings of PEPM'99, Workshop on Partial Evaluation and Semantics-Based Program Manipulation*, pages 51–60, 1999.
- Robert Glück. Jones optimality, binding-time improvements, and the strength of program specializers. In *ASIA-PEPM*, pages 9–19, 2002.
- Robert Glück. An investigation of jones optimality and bti-universal specializers. *Higher-Order and Symbolic Computation*, 21(3):283–309, 2008.
- Robert Glück. An experiment with the fourth Futamura projection. In *Ershov Memorial Conference*, pages 135–150, 2009.

- Robert Glück. Is there a fourth futamura projection? In *Proceedings of the 2009 ACM SIGPLAN Workshop on Partial Evaluation and Program Manipulation*, PEPM '09, pages 51–60, New York, NY, USA, 2009. ACM. ISBN 978-1-60558-327-3. doi: 10.1145/1480945.1480954. URL <http://doi.acm.org/10.1145/1480945.1480954>.
- Carsten K. Gomard and Neil D. Jones. A partial evaluator for the untyped lambda-calculus. *Journal of Functional Programming*, 1(1):21–69, 1991.
- Noriko Hirota and Kenichi Asai. Formalizing a correctness property of a type-directed partial evaluator. In *PLPV'14 Proceedings of the ACM SIGPLAN 2014 Workshop on Programming Languages meets Program Verification*, pages 41–46, 2014.
- Barry Jay and Jens Palsberg. Typed self-interpretation by pattern matching. In *Proceedings of ICFP'11, ACM SIGPLAN International Conference on Functional Programming*, pages 247–258, Tokyo, September 2011.
- Neil D. Jones, Peter Sestoft, and Harald Søndergaard. An experiment in partial evaluation: The generation of a compiler generator. In J.-P. Jouannaud, editor, *Proceedings of Rewriting Techniques and Applications*, pages 225–282. Springer-Verlag (LNCS 202), 1985.
- Neil D. Jones, Carsten K. Gomard, and Peter Sestoft. *Partial Evaluation and Automatic Program Generation*. Prentice-Hall International, 1993.
- Oleg Kiselyov. *Typed Tagless Final Interpreters*, pages 130–174. Springer Berlin Heidelberg, Berlin, Heidelberg, 2012. ISBN 978-3-642-32202-0. doi: 10.1007/978-3-642-32202-0_3. URL http://dx.doi.org/10.1007/978-3-642-32202-0_3.
- John Launchbury. A strongly-typed self-applicable partial evaluator. In *Proceedings of FPCA'91, Sixth ACM Conference on Functional Programming Languages and Computer Architecture*, pages 145–164, 1991.
- Henning Makholm. On jones-optimal specialization for strongly typed languages. In *SAIG*, pages 129–148, 2000.
- Torben A. Mogensen. Efficient self-interpretations in lambda calculus. *Journal of Functional Programming*, 2(3):345–363, 1992. See also DIKU Report D–128, Sep 2, 1994.
- Torben A. Mogensen. Self-applicable online partial evaluation of the pure lambda calculus. In *Proceedings of the 1995 ACM SIGPLAN Symposium on Partial Evaluation and Semantics-based Program Manipulation*, PEPM '95, pages 39–44, New York, NY, USA, 1995. ACM. ISBN 0-89791-720-0. doi: 10.1145/215465.215469. URL <http://doi.acm.org/10.1145/215465.215469>.
- Tobias Nipkow. Lambda calculus. <https://www21.in.tum.de/teaching/logik/SS13/lambda-en.pdf>, (accessed July 7, 2017), August 2, 2012.
- Frank Pfenning and Peter Lee. Metacircularity in the polymorphic λ -calculus. *Theoretical Computer Science*, 89(1):137–159, 1991.
- Tillmann Rendel, Klaus Ostermann, and Christian Hofer. Typed self-representation. In *Proceedings of PLDI'09, ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 293–303, June 2009.
- Amin Shali and William Cook. Hybrid partial evaluation. In *Proceedings of OOPSLA'11, ACM SIGPLAN Conference on Object-Oriented Programming Systems, Languages and Applications*, pages 375–390, 2011.
- Alex Simpson. Reduction in a linear lambda-calculus with applications to operational semantics. In *Proceedings of the 16th International Conference on Term Rewriting and Applications*, RTA'05, pages 219–234, Berlin, Heidelberg, 2005. Springer-Verlag. ISBN 3-540-25596-6, 978-3-540-25596-3. doi: 10.1007/978-3-540-32033-3_17. URL http://dx.doi.org/10.1007/978-3-540-32033-3_17.
- Sebastian Skalberg. Mechanical proof of the optimality of a partial evaluator. M.S. Thesis; DIKU, University of Copenhagen, 1999.
- Walid Taha, Henning Makholm, and John Hughes. Tag elimination and Jones-optimality. In *Proceedings of PADO'01, Programs as Data Objects, Second Symposium*, pages 257–275, 2001.
- Peter Thiemann. Cogen in six lines. In *Proceedings of ICFP'96, ACM International Conference on Functional Programming*, pages 180–189, 1996.
- Geoffrey Washburn and Stephanie Weirich. Boxes go bananas: Encoding higher-order abstract syntax with parametric polymorphism. In *Proceedings of the Eighth ACM SIGPLAN International Conference on Functional Programming*, ICFP '03, pages 249–262, New York, NY, USA, 2003. ACM. ISBN 1-58113-756-7. doi: 10.1145/944705.944728. URL <http://doi.acm.org/10.1145/944705.944728>.

$$\begin{array}{c}
\text{NP-1 } \frac{}{e > e} \\
\text{NP-2 } \frac{e > e'}{\lambda x^* . e > \lambda x^* . e'} \\
\text{NP-3 } \frac{a > a' \quad b > b'}{a b > a' b'} \\
\text{NP-4 } \frac{e > e' \quad v > v'}{(\lambda x . e)v > e'[x:=v']} \\
\text{NP-5 } \frac{e > e'}{(\lambda x . e)y > e'[x:=y]} \\
\text{NP-6 } \frac{a > a' \quad b > b'}{(\lambda x^\circ . a)b > a'[x^\circ:=b']}
\end{array}$$

A CONFLUENCE

Following the proof technique in Nipkow's lecture notes, we define the nested parallel strong optimization relation $a > b$.

Lemma A.1. *If $e1 \rightarrow_s e2$, then $e1 > e2$*

PROOF. By induction on the derivation of $e1 \rightarrow_s e2$.

Suppose $e1 \rightarrow_s e2$ is by SSR-1. Then $e1 = (\lambda x . e)v$ and $e2 = e[x:=v]$, and $e1 > e2$ is derivable by NP-4.

Suppose $e1 \rightarrow_s e2$ is by SSR-2. Then $e1 = a1 b$ and $e2 = a2 b$ and $a1 \rightarrow_s a2$. By induction, $a1 > a2$. Therefore, $e1 > e2$ is derivable by NP-3.

Suppose $e1 \rightarrow_s e2$ is by SSR-3. Then $e1 = a b1$ and $e2 = a b2$ and $b1 \rightarrow_s b2$. By induction, $b1 > b2$. Therefore, $e1 > e2$ is derivable by NP-3.

Suppose $e1 \rightarrow_s e2$ is by SSR-4. Then $e1 = (\lambda x^\circ . a)b$ and $e2 = a[x^\circ:=b]$. Therefore, $e1 > e2$ is derivable by NP-6.

Suppose $e1 \rightarrow_s e2$ is by SSR-5. Then $e1 = (\lambda x . e)y$ and $e2 = e[x:=y]$. Therefore, $e1 > e2$ is derivable by NP-5.

Suppose $e1 \rightarrow_s e2$ is by SSR-6. Then $e1 = (\lambda x . a1)$ and $e2 = (\lambda x . a2)$ and $a1 \rightarrow_s a2$. By induction, $a1 > a2$. Therefore, $e1 > e2$ is derivable by NP-2.

Suppose $e1 \rightarrow_s e2$ is by SSR-7. Then $e1 = (\lambda x^\circ . a1)$ and $e2 = (\lambda x^\circ . a2)$ and $a1 \rightarrow_s a2$. By induction, $a1 > a2$. Therefore, $e1 > e2$ is derivable by NP-2. □

Lemma A.2. *If $e1 > e2$, then $e1 \rightarrow_s^* e2$*

PROOF. By induction on the derivation of $e1 > e2$.

Suppose $e1 > e2$ is by NP-1. Then $e1 = e2$, and $e1 \rightarrow_s^0 e2$.

Suppose $e1 > e2$ is by NP-2. Then $e1 = (\lambda x^* . e)$ and $e2 = (\lambda x^* . e')$ and $e > e'$. By induction, $e \rightarrow_s^* e'$. Therefore, $e1 \rightarrow_s^* e2$ by multiple uses of SSR-6 or SSR-7.

Suppose $e1 > e2$ is by NP-3. Straightforward induction.

Suppose $e1 > e2$ is by NP-4. Then $e1 = (\lambda x . e)v$ and $e2 = e'[x:=v']$ and $e > e'$ and $v > v'$. By induction, $e \rightarrow_s^* e'$ and $v \rightarrow_s^* v'$. Therefore, $e1 = (\lambda x . e)v \rightarrow_s^* (\lambda x . e')v' \rightarrow_s e'[x:=v'] = e2$.

Suppose $e1 > e2$ is by NP-5. Then $e1 = (\lambda x . e)y$ and $e2 = e'[x:=y]$ and $e > e'$. By induction $e \rightarrow_s^* e'$. Therefore, $e1 = (\lambda x . e)y \rightarrow_s^* (\lambda x . e')y \rightarrow_s e'[x:=y] = e2$.

Suppose $e1 > e2$ is by NP-6. Then $e1 = (\lambda x^\circ . a)b$ and $e2 = a'[x^\circ:=b']$ and $a > a'$ and $b > b'$. By induction, $a \rightarrow_s^* a'$ and $b \rightarrow_s^* b'$. Therefore, $e1 = (\lambda x^\circ . a)b \rightarrow_s^* (\lambda x^\circ . a')b' \rightarrow_s a'[x^\circ:=b'] = e2$.

□

Lemma A.3. *If $(\lambda x^*. a) > b$, then there exists a' such that $b = (\lambda x^*. a')$ and $a > a'$*

PROOF. Straightforward. □

Lemma A.4. *If $e > e'$, then $e[x:=y] > e'[x:=y]$.*

PROOF. By straightforward induction on e . □

Lemma A.5. *If $e > e'$ and $v > v'$, then $e[x:=v] > e'[x:=v']$.*

PROOF. By induction on e .

Suppose $e = x$. Then $e' = x$. We have that $e[x:=v] = v$ and $e'[x:=v'] = v'$. Since $v > v'$, $e[x:=v] > e'[x:=v']$.

Suppose $e = y^* \neq x$. Then $e[x:=v] = e$ and $e'[x:=v'] = e'$. Since $e > e'$, $e[x:=v] > e'[x:=v']$.

Suppose $e = (\lambda y^*. e1)$. Without loss of generality, assume $y^* \neq x$ and $y^* \notin FV(v) \cup FV(v')$. By Lemma A.3, $e' = (\lambda y^*. e1')$ and $e1 > e1'$. By induction, $e1[x:=v] > e1'[x:=v']$. Therefore, $\lambda y^*. e1[x:=v] > \lambda y^*. e1'[x:=v']$. Since $\lambda y^*. e1[x:=v] = (\lambda y^*. e1)[x:=v] = e[x:=v]$ and $\lambda y^*. e1'[x:=v'] = (\lambda y^*. e1')[x:=v'] = e'[x:=v']$, we have $e[x:=v] > e'[x:=v']$ as required.

Suppose e is an application. We proceed by inspecting the rule used to derive $e > e'$.

Suppose that $e > e'$ is by NP-3. Then $e = e1 e2$ and $e' = e1' e2'$ and $e1 > e1'$ and $e2 > e2'$. By induction, $e1[x:=v] > e1'[x:=v']$ and $e2[x:=v] > e2'[x:=v']$. Therefore $e1[x:=v] e2[x:=v] > e1'[x:=v'] e2'[x:=v']$. Since $e1[x:=v] e2[x:=v] = (e1 e2)[x:=v] = e[x:=v]$, and $e1'[x:=v'] e2'[x:=v'] = (e1' e2')[x:=v'] = e'[x:=v']$, NP-3 derives $e[x:=v] > e'[x:=v']$ as required.

Suppose that $e > e'$ is by NP-4. Then $e = (\lambda y. e1)v1$ and $e' = e1'[y:=v1']$ and $e1 > e1'$ and $v1 > v1'$. Without loss of generality, assume $y \neq x$ and $y \notin FV(v)$. By induction, $e1[x:=v] > e1'[x:=v']$ and $v1[x:=v] > v1'[x:=v']$. It must be the case that $v1[x:=v]$ is a value. Therefore, NP-4 derives $(\lambda y. e1[x:=v])(v1[x:=v]) > e1'[x:=v'] [y:=v1'[x:=v']]$. Since $(\lambda y. e1[x:=v])(v1[x:=v]) = ((\lambda y. e1)v1)[x:=v] = e[x:=v]$ and $e1'[x:=v'] [y:=v1'[x:=v']] = (e1'[y:=v1'])[x:=v'] = e'[x:=v']$, we have that $e[x:=v] > e'[x:=v']$ as required.

Suppose that $e > e'$ is by NP-5. Then $e = (\lambda y. e1)z$ and $e' = e1'[y:=z]$ and $e1 > e1'$. Without loss of generality, assume $y \neq x$ and $y \notin FV(v)$. By induction, $e1[x:=v] > e1'[x:=v']$.

This step is key! A NP-5 step turns into a NP-4 step. If $z = x$, then $e[x:=v] = ((\lambda y. e1)z)[x:=v] = (\lambda y. e1[x:=v])v$ and $e'[x:=v'] = (e1'[y:=z])[x:=v'] = (e1'[y:=v'])[x:=v'] = (e1'[x:=v'])[y:=v']$. Since $e1[x:=v] > e1'[x:=v']$ and $v > v'$, NP-4 derives $(\lambda y. e1[x:=v])v > (e1'[x:=v'])[y:=v']$. Therefore, $e[x:=v] > e'[x:=v']$ as required.

If $z \neq x$, then $e[x:=v] = ((\lambda y. e1)z)[x:=v] = (\lambda y. e1[x:=v])z$ and $e'[x:=v'] = (e1'[y:=z])[x:=v'] = (e1'[x:=v'])[y:=z]$. Since $e1[x:=v] > e1'[x:=v']$, NP-5 derives $(\lambda y. e1[x:=v])z > (e1'[x:=v'])[y:=z]$. Therefore, $e[x:=v] > e'[x:=v']$ as required.

Suppose that $e > e'$ is by NP-6. Then $e = (\lambda y^\circ. a)b$ and $e' = a'[y^\circ:=v']$ and $a > a'$ and $b > b'$. Without loss of generality, assume $y^\circ \notin FV(v)$. By induction, $a[x:=v] > a'[x:=v']$ and $b[x:=v] > b'[x:=v']$. Therefore, NP-6 derives $(\lambda y^\circ. a[x:=v])(b[x:=v]) > a'[x:=v'] [y^\circ:=b'[x:=v']]$. Since $(\lambda y^\circ. a[x:=v])(b[x:=v]) = ((\lambda y^\circ. a)b)[x:=v] = e[x:=v]$ and $a'[x:=v'] [y^\circ:=b'[x:=v']] = (a'[y^\circ:=b'])[x:=v'] = e'[x:=v']$, $e[x:=v] > e'[x:=v']$ as required. □

Lemma A.6. *If $a > a'$ and $b > b'$, then $a[x^\circ:=b] > a[x^\circ:=b']$.*

Note: the proof is identical to that of Lemma A.5, except that in the case that $a > a'$ is by NP-5, the argument z cannot be x° .

PROOF. By induction on a .

Suppose $a = x^\circ$. Then $a' = x^\circ$, and $a[x^\circ := b] = b$ and $a'[x^\circ := b'] = b'$. Therefore $b > b'$ implies $a[x^\circ := b] > a'[x^\circ := b']$.

Suppose $a = y^* \neq x^\circ$. $a[x^\circ := b] = a$ and $a'[x^\circ := b'] = a'$. Therefore $a > a'$ implies $a[x^\circ := b] > a'[x^\circ := b']$.

Suppose $a = (\lambda y^*. a1)$. Without loss of generality, assume $y^* \neq x^\circ$ and $y^* \notin \text{FV}(b) \cup \text{FV}(b')$. By Lemma A.3, $a' = (\lambda y^*. a1')$ and $a1 > a1'$. By induction, $a1[x^\circ := b] > a1'[x^\circ := b']$. Therefore, $\lambda y^*. a1[x^\circ := b] > \lambda y^*. a1'[x^\circ := b']$. Since $\lambda y^*. a1[x^\circ := b] = (\lambda y^*. a1)[x^\circ := b] = a[x^\circ := b]$ and $\lambda y^*. a1'[x^\circ := b'] = (\lambda y^*. a1')[x^\circ := b'] = a'[x^\circ := b']$, we have $a[x^\circ := b] > a'[x^\circ := b']$ as required.

Suppose a is an application. We proceed by inspecting the rule used to derive $a > a'$.

Suppose that $a > a'$ is by NP-3. Then $a = a1 a2$ and $a' = a1' a2'$ and $a1 > a1'$ and $a2 > a2'$. By induction, $a1[x^\circ := b] > a1'[x^\circ := b']$ and $a2[x^\circ := b] > a2'[x^\circ := b']$. Therefore $a1[x^\circ := b] a2[x^\circ := b] > a1'[x^\circ := b'] a2'[x^\circ := b']$. Since $a1[x^\circ := b] a2[x^\circ := b] = (a1 a2)[x^\circ := b] = a[x^\circ := b]$, and $a1'[x^\circ := b'] a2'[x^\circ := b'] = (a1' a2')[x^\circ := b'] = a'[x^\circ := b']$, NP-3 derives $a[x^\circ := b] > a'[x^\circ := b']$ as required.

Suppose that $a > a'$ is by NP-4. Then $a = (\lambda y. a1)v$ and $a' = a1'[y := v']$ and $a1 > a1'$ and $v > v'$. Without loss of generality, assume $y \neq x^\circ$ and $y \notin \text{FV}(b)$. By induction, $a1[x^\circ := b] > a1'[x^\circ := b']$ and $v[x^\circ := b] > v'[x^\circ := b']$. It must be the case that $v[x^\circ := b]$ is a value. Therefore, NP-4 derives $(\lambda y. a1[x^\circ := b])(v[x^\circ := b]) > a1'[x^\circ := b'] [y := v'[x^\circ := b']]$. Since $(\lambda y. a1[x^\circ := b])(v[x^\circ := b]) = ((\lambda y. a1)v)[x^\circ := b] = a[x^\circ := b]$ and $a1'[x^\circ := b'] [y := v'[x^\circ := b']] = (a1'[y := v'])[x^\circ := b'] = a'[x^\circ := b']$, we have that $a[x^\circ := b] > a'[x^\circ := b']$ as required.

Suppose that $a > a'$ is by NP-5. Then $a = (\lambda y. a1)z$ and $a' = a1[y := z]$ and $a1 > a1'$. Without loss of generality, assume $y \neq x^\circ$ and $y \notin \text{FV}(b)$. By induction, $a1[x^\circ := b] > a1'[x^\circ := b']$.

Since $z \neq x^\circ$, we have that $a[x^\circ := b] = ((\lambda y. a1)z)[x^\circ := b] = (\lambda y. a1[x^\circ := b])z$ and $a'[x^\circ := b'] = (a1'[y := z])[x^\circ := b'] = (a1'[x^\circ := b']) [y := z]$. Since $a1[x^\circ := b] > a1'[x^\circ := b']$, NP-5 derives $(\lambda y. a1[x^\circ := b])z > (a1'[x^\circ := b']) [y := z]$. Therefore, $a[x^\circ := b] > a'[x^\circ := b']$ as required.

Suppose that $a > a'$ is by NP-6. Then $a = (\lambda y^\circ. a)b1$ and $a' = a'[y^\circ := b']$ and $a > a'$ and $b1 > b1'$. Without loss of generality, assume $y^\circ \notin \text{FV}(b)$. By induction, $a[x^\circ := b] > a'[x^\circ := b']$ and $b1[x^\circ := b] > b1'[x^\circ := b']$. Therefore, NP-6 derives $(\lambda y^\circ. a[x^\circ := b])(b1[x^\circ := b]) > a'[x^\circ := b'] [y^\circ := b1'[x^\circ := b']]$. Since $(\lambda y^\circ. a[x^\circ := b])(b1[x^\circ := b]) = ((\lambda y^\circ. a)b1)[x^\circ := b] = a[x^\circ := b]$ and $a'[x^\circ := b'] [y^\circ := b1'[x^\circ := b']] = (a'[y^\circ := b1'])[x^\circ := b'] = a'[x^\circ := b']$, $a[x^\circ := b] > a'[x^\circ := b']$ as required. \square

Lemma A.7 (Diamond property for $>$). *If $a > b1$ and $a > b2$, then there exists a term c such that $b1 > c$ and $b2 > c$.*

PROOF. By induction on a .

Suppose $a = x^*$. Then $b1 = b2 = x^*$. Holds with $c = x^*$.

Suppose $a = (\lambda x^*. a')$. By Lemma A.3, $b1 = (\lambda x^*. b1')$ and $b2 = (\lambda x^*. b2')$ and $a' > b1'$ and $a' > b2'$. By induction, there exists a c' such that $b1' > c'$ and $b2' > c'$. Therefore, NP-2 derives $(\lambda x^*. b1') > (\lambda x^*. c')$ and $(\lambda x^*. b2') > (\lambda x^*. c')$. Holds with $c = (\lambda x^*. c')$.

Suppose a is an application. Consider the derivation of $a > b1$ and $a > b2$.

Suppose $a > b1$ is by NP-3 and $a > b2$ is by NP-3. Then $a = d e$ and $b1 = d1 e1$ and $b2 = d2 e2$ and $d > d1$ and $d > d2$ and $e > e1$ and $e > e2$. By induction, there exist $d3$ and $e3$ such that $d1 > d3$ and $d2 > d3$ and $e1 > e3$ and $e2 > e3$. Let $c = d3 e3$. NP-3 derives $b1 = d1 e1 > d3 e3 = c$ and $b2 = d2 e2 > d3 e3 = c$.

Suppose $a > b1$ is by NP-4 and $a > b2$ is by NP-4. Then $a = (\lambda x. e)v$ and $b1 = e1[x := v1]$ and $b2 = e2[x := v2]$ and $e > e1$ and $e > e2$ and $v > v1$ and $v > v2$. By induction, there exist terms $e3$ and $v3$ such that $e1 > e3$ and $e2 > e3$ and $v1 > v3$ and $v2 > v3$. Let $c = e3[x := v3]$. By Lemma A.5, $b1 = e1[x := v1] > e3[x := v3] = c$ and $b2 = e2[x := v2] > e3[x := v3] = c$.

$$\begin{array}{c}
\text{WOPT-1} \frac{}{(\lambda x. e)v \rightarrow_w e[x:=v]} \\
\text{WOPT-2} \frac{e1 \rightarrow_w e1'}{e1 e2 \rightarrow_w e1' e2} \\
\text{WOPT-3} \frac{e1 \rightarrow_w e1'}{(\lambda x. e)e1 \rightarrow_w (\lambda x. e)e1'} \\
\text{WOPT-4} \frac{}{(\lambda x^\circ. e)e' \rightarrow_w e[x^\circ:=e']} \\
\text{WOPT-5} \frac{e1 \rightarrow_w e1'}{(\lambda x^\circ. e)e1 \rightarrow_w (\lambda x^\circ. e)e1'}
\end{array}$$

Fig. 22. Weak Optimization

Suppose $a > b1$ is by NP-5 and $a > b2$ is by NP-5. Then $a = (\lambda x. e)y$ and $b1 = e1[x:=y]$ and $b2 = e2[x:=y]$ and $e > e1$ and $e > e2$. By induction, there exists a term $e3$ such that $e1 > e3$ and $e2 > e3$. Let $c = e3[x:=y]$. By Lemma A.4, $b1 = e1[x:=y] > e3[x:=y] = c$ and $b2 = e2[x:=y] > e3[x:=y] = c$.

Suppose $a > b1$ is by NP-6 and $a > b2$ is by NP-6. Then $a = (\lambda x^\circ. d)e$ and $b1 = d1[x^\circ:=e1]$ and $b2 = d2[x^\circ:=e2]$ and $d > d1$ and $d > d2$ and $e > e1$ and $e > e2$. By induction, there exist terms $d3$ and $e3$ such that $d1 > d3$ and $d2 > d3$ and $e1 > e3$ and $e2 > e3$. Let $c = d3[x^\circ:=e3]$. By Lemma A.5, $b1 = d1[x^\circ:=e1] > d3[x^\circ:=e3] = c$ and $b2 = d2[x^\circ:=e2] > d3[x^\circ:=e3] = c$.

Suppose $a > b1$ is by NP-3 and $a > b2$ is by NP-4. Then $a = (\lambda x. e)v$ and $b1 = (\lambda x. e1)v1$ and $e > e1$ and $v > v1$, and $b2 = e2[x:=v2]$ and $e > e2$ and $v > v2$. Since v is a value and $v > v1$, $v1$ must be a value. By induction, there is term $v3$ such that $v1 > v3$ and $v2 > v3$. Since $v1$ and $v2$ are values, $v3$ is a value. By induction again, there is a term $e3$ such that $e1 > e3$ and $e2 > e3$. Let $c = e3[x:=v3]$. By rule NP-4, $b1 = (\lambda x. e1)v1 > e3[x:=v3]$. By Lemma A.5, $b2 = e2[x:=v2] > e3[x:=v3]$.

Suppose $a > b1$ is by NP-3 and $a > b2$ is by NP-5. Then $a = (\lambda x. e)y$ and $b1 = (\lambda x. e1)y$ and $b2 = e2[x:=y]$ and $e > e1$ and $e > e2$. By induction, there exists an $e3$ such that $e1 > e3$ and $e2 > e3$. Let $c = e3[x:=y]$. By rule NP-5, $b1 = (\lambda x. e1)y > e3[x:=y] = c$. By Lemma A.4, $b2 = e2[x:=y] > e3[x:=y] = c$.

Suppose $a > b1$ is by NP-3 and $a > b2$ is by NP-6. Then $a = (\lambda x^\circ. e)e'$ and $b1 = (\lambda x^\circ. e1)e1'$ and $e > e1$ and $e > e1'$ and $b2 = e2[x^\circ:=e2']$ and $e > e2$ and $e > e2'$. By induction, there exist $e3$ and $e3'$ such that $e1 > e3$ and $e2 > e3$, and $e1' > e3'$ and $e2' > e3'$. Let $c = e3[x^\circ:=e3']$. By rule NP-6, $b1 = (\lambda x^\circ. e1)e1' > e3[x^\circ:=e3'] = c$. By Lemma A.6, $b2 = e2[x^\circ:=e2'] > e3[x^\circ:=e3'] = c$. \square

THEOREM 4.3. *Specialization-safe reduction is confluent.*

PROOF. By Nipkow's Lemma A.2.5, and Lemmas A.1, A.2, and A.7 \square

COROLLARY 4.4. *Specialization-safe normal forms are unique.*

B SPECIALIZATION SAFETY

Figure 7 defines call-by-value reduction on annotated terms. Figure 22 defines weak optimization. It extends call-by-value reduction with the rule WOPT-4, which reduces β -redexes without evaluating the argument when the λ -abstraction abstracts over a top-level linear variable.

The following lemma allows us to reason about multiple steps of call-by-value reduction in a big-step-style.

Lemma B.1. *For any n , $e_1 \ e_2 \rightarrow^n v$ if and only if there exist integers n_1 , n_2 , and n_3 , and terms $(\lambda x^*. e_1')$ and v_2 , such that $e_1 \rightarrow^{n_1} (\lambda x^*. e_1')$, and $e_2 \rightarrow^{n_2} v_2$, and $e_1'[x^* := v_2] \rightarrow^{n_3} v$, and $n = n_1 + n_2 + n_3 + 1$.*

We prove each direction separately.

PROOF. Suppose $e_1 \ e_2 \rightarrow^n v$. We want to prove there exist integers n_1 , n_2 , and n_3 , and terms $(\lambda x^*. e_1')$ and v_2 , such that $e_1 \rightarrow^{n_1} (\lambda x^*. e_1')$, and $e_2 \rightarrow^{n_2} v_2$, and $e_1'[x^* := v_2] \rightarrow^{n_3} v$, and $n = n_1 + n_2 + n_3 + 1$.

We proceed by induction on the derivation of $e_1 \ e_2 \rightarrow^n v$.

Suppose $n = 0$. Then $e_1 \ e_2 = v$. This is a contradiction, since $e_1 \ e_2$ is not a λ -abstraction.

Suppose $n > 0$. Then the sequence $e_1 \ e_2 \rightarrow^n v$ contains at least one step. Consider the first step in the sequence.

If the first step is by CBV-1 or CBV-4, then we have that $e_1 = (\lambda x^*. e_1')$ and $e_2 = v_2$ and $e_1'[x^* := v_2] \rightarrow^{n'} v$. The result holds with $n_1 = 0$ and $n_2 = 0$ and $n_3 = n'$.

If the first step is by CBV-2, then $e_1 \rightarrow e_1'$ and $e_1' \ e_2 \rightarrow^{n'} v$. By induction, $e_1' \rightarrow^{n_1'} \lambda x. e$ and $e_2 \rightarrow^{n_2'} v_2$ and $e[x := v_2] \rightarrow^{n_3'} v$ and $n' = n_1' + n_2' + n_3' + 1$. Then $e_1 \rightarrow^{n_1'+1} \lambda x. e$. The result holds with $n_1 = n_1' + 1$.

If the first step is by CBV-3 or CBV-5, then $e_1 = (\lambda x^*. e_1')$ and $e_2 \rightarrow e_2'$ and $(\lambda x^*. e_1') \ e_2' \rightarrow^{n'} v$. Since $(\lambda x^*. e_1')$ is a value, it doesn't step. Therefore, the induction hypothesis states that $e_2' \rightarrow^{n_2'} v_2$ and $e_1'[x^* := v_2] \rightarrow^{n_3'} v$ and $n' = n_2' + n_3' + 1$. Therefore, $e_2 \rightarrow^{n_2'+1} v_2$. The result holds with $n_1 = 0$ and $n_2 = n_2' + 1$. \square

PROOF. Suppose $e_1 \rightarrow^{n_1} (\lambda x^*. e_1')$, and $e_2 \rightarrow^{n_2} v_2$, and $e_1'[x^* := v_2] \rightarrow^{n_3} v$. We want to prove $e_1 \ e_2 \rightarrow^{n_1+n_2+n_3+1} v$.

We proceed by induction on $n_1 + n_2$.

Suppose $n_1 + n_2 = 0$. Then $n_1 = 0$ and $n_2 = 0$. Therefore, $e_1 = (\lambda x^*. e_1')$ and $e_2 = v_2$. We have $e_1 \ e_2 \rightarrow e_1'[x^* := v_2]$ and since $e_1'[x^* := v_2] \rightarrow^{n_3} v$, $e_1 \ e_2 \rightarrow^{n_3+1} v$. The conclusion holds because $n_3 + 1 = n_1 + n_2 + n_3 + 1$.

Suppose $n_1 + n_2 > 0$. Then either $n_1 > 0$, or $n_1 = 0$ and $n_2 > 0$.

Suppose $n_1 > 0$. In particular, $n_1 = n_1' + 1$. Then there is a term e such that $e_1 \rightarrow e$ and $e \rightarrow^{n_1'} (\lambda x^*. e_1')$. By induction, $e \ e_2 \rightarrow^{n_1'+n_2+n_3+1} v$. Therefore, $e_1 \ e_2 \rightarrow^{1+n_1'+n_2+n_3+1} v$. The conclusion holds since $n_1 = n_1' + 1$.

Suppose $n_1 = 0$ and $n_2 > 0$. In particular, $n_2 = n_2' + 1$. Then $e_1 = (\lambda x^*. e_1')$ and there exists a term e such that $e_2 \rightarrow e$ and $e \rightarrow^{n_2'} v_2$. By induction, $e_1 \ e \rightarrow^{n_1+n_2'+n_3+1} v$. Therefore, $e_1 \ e_2 \rightarrow^{1+n_1+n_2'+n_3+1} v$. The conclusion holds since $n_2 = n_2' + 1$. \square

Lemma B.2 (Weakening of unlimited context). *If $(\Gamma_1, \Gamma_2); \Sigma \vdash e$, then $(\Gamma_1, x, \Gamma_2); \Sigma \vdash e$.*

PROOF. By induction on the derivation of $(\Gamma_1, \Gamma_2); \Sigma \vdash e$.

Suppose e is an unlimited variable $y \neq x$. Immediate.

Suppose e is an affine variable y° . Immediate.

Suppose $e = e_1 \ e_2$. Then $(\Gamma_1, \Gamma_2); \Sigma_1 \vdash e_1$ and $(\Gamma_1, \Gamma_2); \Sigma_2 \vdash e_2$ and $\Sigma = \Sigma_1 \uplus \Sigma_2$. By induction, $(\Gamma_1, x, \Gamma_2); \Sigma_1 \vdash e_1$, and $(\Gamma_1, x, \Gamma_2); \Sigma_2 \vdash e_2$, so $(\Gamma_1, x, \Gamma_2); \Sigma \vdash e_1 \ e_2$ as required.

Suppose $e = (\lambda y. e')$. Then $(\Gamma_1, \Gamma_2, y); \langle \rangle \vdash e'$. By induction, $(\Gamma_1, x, \Gamma_2, y); \langle \rangle \vdash e'$, so $(\Gamma_1, x, \Gamma_2); \Sigma \vdash (\lambda y. e')$ as required. \square

Lemma B.3 (Weakening of affine context). *If $\Gamma; \langle \rangle \vdash e$, then $\Gamma; \Sigma \vdash e$.*

PROOF. By induction on the derivation of $\Gamma; \langle \rangle \vdash e$.

Suppose e is an unlimited variable x . Immediate.

Suppose e is an affine variable x° . Contradiction.

Suppose $e = e_1 e_2$. Then $\Gamma; \langle \rangle \vdash e_1$ and $\Gamma; \langle \rangle \vdash e_2$. By induction, $\Gamma; \Sigma \vdash e_1$, so $\Gamma; \Sigma \vdash e_1 e_2$ as required.

Suppose $e = (\lambda x^*. e')$. Immediate. \square

Lemma B.4. *If $(\Gamma_1, x, \Gamma_2); \Sigma \vdash e_1$ and $(\Gamma_1, \Gamma_2); \langle \rangle \vdash e_2$, then $(\Gamma_1, \Gamma_2); \Sigma \vdash e_1[x := e_2]$.*

PROOF. By induction on the derivation of $(\Gamma_1, x, \Gamma_2); \Sigma \vdash e_1$.

Suppose $e_1 = x$. Then $e_1[x := e_2] = e_2$, and the conclusion holds.

Suppose $e_1 = y \neq x$. Then $e_1[x := e_2] = y$, and the conclusion holds.

Suppose $e_1 = y^\circ$. Since $x \neq y^\circ$, $e_1[x := e_2] = y^\circ$, and the conclusion holds.

Suppose $e_1 = a b$. Then $(\Gamma_1, x, \Gamma_2); \Sigma_1 \vdash a$ and $(\Gamma_1, x, \Gamma_2); \Sigma_2 \vdash b$ and $\Sigma = \Sigma_1 \uplus \Sigma_2$. By induction, $(\Gamma_1, \Gamma_2); \Sigma_1 \vdash a[x := e_2]$ and $(\Gamma_1, \Gamma_2); \Sigma_2 \vdash b[x := e_2]$. Therefore $(\Gamma_1, \Gamma_2); \Sigma \vdash (a[x := e_2]) (b[x := e_2])$. The conclusion holds since $e_1[x := e_2] = (a b)[x := e_2] = (a[x := e_2]) (b[x := e_2])$.

Suppose $e_1 = (\lambda y. a)$. Then $(\Gamma_1, x, \Gamma_2, y); \langle \rangle \vdash a$. Without loss of generality, assume $x \neq y$ and y does not occur free in e_2 . By induction, $(\Gamma_1, \Gamma_2, y); \langle \rangle \vdash a[x := e_2]$. Therefore, $(\Gamma_1, \Gamma_2); \langle \rangle \vdash (\lambda y. a[x := e_2])$. The conclusion holds since $e_1[x := e_2] = (\lambda y. a)[x := e_2] = (\lambda y. a[x := e_2])$.

Suppose $e_1 = (\lambda y^\circ. a)$. This case is similar to the previous case. \square

Lemma B.5. *If $\Gamma; x^\circ \vdash e_1$ and $\Gamma; \Sigma \vdash e_2$, then $\Gamma; \Sigma \vdash e_1[x^\circ := e_2]$.*

PROOF. By induction on the derivation of $\Gamma; x^\circ \vdash e_1$.

Suppose $e_1 = y \in \Gamma$. Then $e_1[x^\circ := e_2] = y$, and the conclusion holds.

Suppose $e_1 = x^\circ$. Then $e_1[x^\circ := e_2] = e_2$, and the conclusion holds.

Suppose $e_1 = a b$. Then $\Gamma; \Sigma_1 \vdash a$ and $\Gamma; \Sigma_2 \vdash b$ and $x^\circ = \Sigma_1 \uplus \Sigma_2$. Therefore, either $\Sigma_1 = x^\circ$ and $\Sigma_2 = \langle \rangle$, or else $\Sigma_1 = \langle \rangle$ and $\Sigma_2 = x^\circ$. The two cases are similar, so we will only consider the first. Suppose $\Sigma_1 = x^\circ$ and $\Sigma_2 = \langle \rangle$. Then by induction, we have $\Gamma; \Sigma \vdash a[x^\circ := e_2]$. Since $\Gamma; \langle \rangle \vdash b$, x° does not occur free in b , so $b[x^\circ := e_2] = b$. Therefore, $e_1[x^\circ := e_2] = (a b)[x^\circ := e_2] = (a[x^\circ := e_2]) b$. Since $\Gamma; \Sigma \vdash a[x^\circ := e_2]$ and $\Gamma; \langle \rangle \vdash b$, we have $\Gamma; \Sigma \vdash (a[x^\circ := e_2]) b$ as required.

Suppose $e_1 = (\lambda y. a)$. Then $(\Gamma, y); \langle \rangle \vdash a$, so x° does not occur free in a . Therefore, x° does not occur free in e_1 , so $e_1[x^\circ := e_2] = e_1$. Finally, since $(\Gamma, y); \langle \rangle \vdash a$, $\Gamma; \Sigma \vdash (\lambda y. a)$ as required.

Suppose $e_1 = (\lambda y^\circ. a)$. This case is similar to the previous case. \square

Lemma B.6. *If $\Gamma; \Sigma \vdash v$, then $\Gamma; \langle \rangle \vdash v$.*

PROOF. We have either $v = (\lambda x. e)$ or $v = (\lambda x^\circ. e)$. If $v = (\lambda x. e)$, then $(\Gamma, x); \langle \rangle \vdash e$, so $\Gamma; \langle \rangle \vdash (\lambda x. e)$ as required. If $v = (\lambda x^\circ. e)$, then $\Gamma; x^\circ \vdash e$, so $\Gamma; \langle \rangle \vdash (\lambda x^\circ. e)$ as required. \square

THEOREM 4.2. *If $\Gamma; \Sigma \vdash e$ and $e \rightarrow_s e'$, then $\Gamma; \Sigma \vdash e'$.*

PROOF. By induction on the derivation of $e \rightarrow_s e'$.

Suppose $e \rightarrow_s e'$ is by SSR-1. Then $e = (\lambda x. a)v$ and $e' = a[x := v]$. We have that $(\Gamma, x); \langle \rangle \vdash a$, and $\Gamma; \Sigma \vdash v$. By Lemma B.6, $\Gamma; \Sigma \vdash v$. By Lemma B.4, $\Gamma; \langle \rangle \vdash a[x := v]$. By Lemma B.3, $\Gamma; \Sigma \vdash a[x := v]$ as required.

Suppose $e \rightarrow_s e'$ is by SSR-2. Then $e = e_1 e_2$ and $e' = e_1' e_2$ and $e_1 \rightarrow_s e_1'$. We have $\Gamma, \Sigma_1 \vdash e_1$ and $\Gamma, \Sigma_2 \vdash e_2$ and $\Sigma = \Sigma_1 \uplus \Sigma_2$. By induction, $\Gamma, \Sigma_1 \vdash e_1'$, so $\Gamma, \Sigma \vdash e_1' e_2$ as required.

Suppose $e \rightarrow_s e'$ is by SSR-3. This case is similar to the previous case.

Suppose $e \rightarrow_s e'$ is by SSR-4. Then $e = (\lambda x^\circ. a)b$ and $e' = a[x^\circ := b]$. We have $\Gamma; x^\circ \vdash a$ and $\Gamma; \Sigma \vdash b$. By Lemma B.5, $\Gamma; \Sigma \vdash a[x^\circ := b]$ as required.

Suppose $e \rightarrow_s e'$ is by SSR-5. Then $e = (\lambda x. a)y$ and $e' = e[x:=y]$. We have $(\Gamma, x); \langle \rangle \vdash a$ and $\Gamma; \Sigma \vdash y$. Since y is unlimited, $y \notin \Sigma$, so $\Gamma; \langle \rangle \vdash y$. Therefore, by Lemma B.4, $\Gamma; \langle \rangle \vdash a[x:=y]$ as required.

Suppose $e \rightarrow_s e'$ is by SSR-6. Then $e = (\lambda x. a)$ and $e' = (\lambda x. a')$ and $a \rightarrow_s a'$. We have that $(\Gamma, x); \langle \rangle \vdash a$, so by induction $(\Gamma, x); \langle \rangle \vdash a'$. Therefore, $\Gamma; \Sigma \vdash (\lambda x. a')$ as required.

Suppose $e \rightarrow_s e'$ is by SSR-7. Then $e = (\lambda x^\circ. a)$ and $e' = (\lambda x^\circ. a')$ and $a \rightarrow_s a'$. We have that $\Gamma; x^\circ \vdash a$, so by induction $\Gamma; x^\circ \vdash a'$. Therefore, $\Gamma; \Sigma \vdash (\lambda x^\circ. a')$ as required. \square

Lemma B.7. *Suppose $\langle \rangle; x^\circ \vdash e1$ and $e2 \rightarrow^j v2$. Then, if $e1[x^\circ:=v2] \rightarrow^k v$ then $e1[x^\circ:=e2] \rightarrow^{k'}$ v for some $k' \leq j + k$.*

PROOF. By induction on the derivation of $\langle \rangle; x^\circ \vdash e1$.

Suppose $e1 = x^\circ$. Then $e1[x^\circ:=e2] = e2$ and $e1[x^\circ:=v2] = v2$. Since $e2 \rightarrow^j v2$, we have that $e1[x^\circ:=e2] \rightarrow^j v2$. Since $v2$ is a value, $v2 \rightarrow^0 v2$ so $e1[x^\circ:=v2] \rightarrow^0 v2$. The conclusion holds with $v = v2$, $i = j$, $k = 0$. The inequality $i \leq j + k$ simplifies to $j \leq j$.

Suppose $e1 = y^*$ where $y^* \neq x^\circ$. Then we have $\langle \rangle; x^\circ \vdash y^*$, a contradiction.

Suppose $e1 = (\lambda y. e1')$. Then we have $\langle \rangle; x^\circ \vdash (\lambda y. e1')$, which implies $(\langle \rangle, y); \langle \rangle \vdash e1'$. Therefore, x° does occur free in $e1$. Therefore, $e1[x^\circ:=e2] = e1[x^\circ:=v2] = (\lambda y. e1')$. Let $v = (\lambda y. e1')$. Then $e1[x^\circ:=e2] \rightarrow^0 v$ and $e1[x^\circ:=v2] \rightarrow^0 v$, so $i = k = 0$ and $i \leq j + k$.

Suppose $e1 = (\lambda y^\circ. e1')$. Then we have $\langle \rangle; x^\circ \vdash (\lambda y^\circ. e1')$, which implies $\langle \rangle; y^\circ \vdash e1'$. Therefore, x° does occur free in $e1$. Therefore, $e1[x^\circ:=e2] = e1[x^\circ:=v2] = (\lambda y^\circ. e1')$. Let $v = (\lambda y^\circ. e1')$. Then $e1[x^\circ:=e2] \rightarrow^0 v$ and $e1[x^\circ:=v2] \rightarrow^0 v$, so $i = k = 0$ and $i \leq j + k$.

Suppose $e1 = e3 e4$. Then either $\langle \rangle; x^\circ \vdash e3$ and $\langle \rangle; \langle \rangle \vdash e4$, or $\langle \rangle; \langle \rangle \vdash e3$ and $\langle \rangle; x^\circ \vdash e4$.

Suppose $\langle \rangle; x^\circ \vdash e3$ and $\langle \rangle; \langle \rangle \vdash e4$. Then $e1[x^\circ:=e2] = (e3 e4)[x^\circ:=e2] = e3[x^\circ:=e2] e4$, and $e1[x^\circ:=v2] = (e3 e4)[x^\circ:=v2] = e3[x^\circ:=v2] e4$. Since $e1[x^\circ:=v2] \rightarrow^k v$ Lemma B.1 states that $e3[x^\circ:=v2] \rightarrow^{k1} (\lambda y^*. e3')$, and $e4 \rightarrow^{k2} v4$, and $e3'[y^*:=v4] \rightarrow^{k3} v$, and $(e3 e4)[x^\circ:=v2] \rightarrow^{k1+k2+k3+1} v$, and $k = k1 + k2 + k3 + 1$. By induction, $e3[x^\circ:=e2] \rightarrow^{k1'} (\lambda y^*. e3')$ for some $k1' \leq j + k1$. We have that $e3[x^\circ:=e2] \rightarrow^{k1'} (\lambda y^*. e3')$, and $e4 \rightarrow^{k2} v4$, and $e3'[y^*:=v4] \rightarrow^{k3} v$. Therefore, Lemma B.1 states that $e3[x^\circ:=e2] e4 \rightarrow^{k1'+k2+k3+1} v$. Since $e1[x^\circ:=e2] = (e3 e4)[x^\circ:=e2] = e3[x^\circ:=e2] e4$, we also have that $e1[x^\circ:=e2] \rightarrow^{k1'+k2+k3+1} v$. Let $k' = k1' + k2 + k3 + 1 \leq j + k1 + k2 + k3 + 1 = j + k$. The conclusion holds since $k' \leq j + k$.

Suppose $\langle \rangle; \langle \rangle \vdash e3$ and $\langle \rangle; x^\circ \vdash e4$. Then $e1[x^\circ:=e2] = (e3 e4)[x^\circ:=e2] = e3 (e4[x^\circ:=e2])$, and $e1[x^\circ:=v2] = (e3 e4)[x^\circ:=v2] = e3 (e4[x^\circ:=v2])$. Since $e1[x^\circ:=v2] \rightarrow^k v$, Lemma B.1 states that $e3 \rightarrow^{k1} (\lambda y^*. e3')$, and $e4[x^\circ:=v2] \rightarrow^{k2} v4$, and $e3'[y^*:=v4] \rightarrow^{k3} v$, and $k = k1 + k2 + k3 + 1$. By induction, $e4[x^\circ:=e2] \rightarrow^{k2'} v4$ for some $k2' \leq j + k2$. Therefore, we have $e3 \rightarrow^{k1} (\lambda y^*. e3')$, and $e4[x^\circ:=e2] \rightarrow^{k2'} v4$, and $e3'[y^*:=v4] \rightarrow^{k3} v$, so Lemma B.1 states that $e3 (e4[x^\circ:=e2]) \rightarrow^{k1+k2'+k3+1} v$. Since $e1[x^\circ:=e2] = (e3 e4)[x^\circ:=e2] = e3 (e4[x^\circ:=e2])$, we also have that $e1[x^\circ:=e2] \rightarrow^{k1+k2'+k3+1} v$. Let $k' = k1 + k2' + k3 + 1 \leq j + k1 + k2 + k3 + 1 = j + k$. The conclusion holds since $k' \leq j + k$. \square

Lemma B.8. *If $\Gamma; x^\circ \vdash e$ and $e[x^\circ:=v1] \rightarrow^* v$ and $e1 \rightarrow^* v1$, then $e[x^\circ:=e1] \rightarrow^* v$.*

PROOF. By induction on e .

Suppose $e = x^\circ$. Then $e[x^\circ:=v1] = v1$ and $v = v1$. But then $e[x^\circ:=e1] \rightarrow^* v1 = v$ as required.

Suppose $e = y^\circ \neq x^\circ$. Contradiction: $\Gamma; x^\circ \vdash e$ cannot be derived.

Suppose $e = y$. Then $e[x^\circ:=v1] = e[x^\circ:=e1] = e$. Trivial.

Suppose $e = \lambda y^*. e'$. Since $\Gamma; x^\circ \vdash e$, $x^\circ \notin \text{FV}(e')$, so $e[x^\circ:=v1] = e[x^\circ:=e1] = e$. Trivial.

Suppose $e = a b$. Then either $\Gamma; x^\circ \vdash a$ and $\Gamma; \langle \rangle \vdash b$, or else $\Gamma; \langle \rangle \vdash a$ and $\Gamma; x^\circ \vdash b$. The two cases are similar, so we consider on the first. Suppose $\Gamma; x^\circ \vdash a$ and $\Gamma; \langle \rangle \vdash b$. Then $x^\circ \notin \text{FV}(b)$, so $e[x^\circ:=v1]$

$= a[x^\circ := v1] b$. Now consider $e[x^\circ := v1] \rightarrow^* v$. By Lemma B.1, $a[x^\circ := v1] \rightarrow^* \lambda y^*. a'$ and $b \rightarrow^* v2$ and $a'[y^* := v2] \rightarrow^* v$. By induction, $a[x^\circ := e1] \rightarrow^* \lambda y^*. a'$, so the conclusion holds. \square

Lemma B.9. *If $e1 \rightarrow_s e2$, then $e1[x^* := v] \rightarrow_s e2[x^* := v]$.*

PROOF. By induction on the derivation of $e1 \rightarrow_s e2$.

If $e1 \rightarrow_s e2$ is by SSR-1, then $e1 = (\lambda y. e)v1$ and $e2 = e[y := v1]$. Without loss of generality, $y \neq x^*$ and $y \notin FV(v)$, so $e1[x^* := v] = (\lambda y. e[x^* := v])(v1[x^* := v])$, and $e2[x^* := v] = e[y := v1][x^* := v] = e[x^* := v][y := v1[x^* := v]]$. Since $v1$ is a value, $v1[x^* := v]$ is a value. Therefore, $e1[x^* := v] = (\lambda y. e[x^* := v])(v1[x^* := v]) \rightarrow_s e[x^* := v][y := v1[x^* := v]] = e2[x^* := v]$ by SSR-1.

If $e1 \rightarrow e2$ is by SSR-2 or SSR-3, the result follows by straightforward induction.

If $e1 \rightarrow_s e2$ is by SSR-4, then $e1 = (\lambda y^\circ. a)b$. Without loss of generality, $x^* \neq y^\circ$ and $y^\circ \notin FV(a)$, so $e1[x^* := v] = (\lambda y^\circ. a[x^* := v])(b[x^* := v])$. Therefore, $e1[x^* := v] = (\lambda y^\circ. a[x^* := v])(b[x^* := v]) \rightarrow_s a[x^* := v][y := b[x^* := v]] = a[y := b][x^* := v] = e2[x^* := v]$ is derivable by SSR-4.

If $e1 \rightarrow_s e2$ is by SSR-5, then $e1 = (\lambda z. e)y$ and $e2 = e[z := y]$. Without loss of generality, assume $z \notin FV(v)$. If $x^* = y$, then $e1[x^* := v] = (\lambda z. e[x^* := v])v$ and $e2[x^* := v] = e[z := x^*][x^* := v] = e[x^* := v][z := v]$. Therefore, $e1[x^* := v] \rightarrow_s e2[x^* := v]$ is derivable using SSR-1. If $x \neq y$, then $e1[x^* := v] = (\lambda z. e[x^* := v])y$ and $e2[x^* := v] = e[z := y][x^* := v] = e[x^* := v][z := y]$. Therefore, $e1[x^* := v] \rightarrow_s e2[x^* := v]$ is derivable using SSR-5.

If $e1 \rightarrow_s e2$ is by SSR-6, then $e1 = (\lambda y. e)$ and $e2 = (\lambda y. e')$ and $e \rightarrow_s e'$. By induction, $e[x^* := v] \rightarrow_s e'[x^* := v]$. Therefore, $e1[x^* := v] = (\lambda y. e[x^* := v]) \rightarrow_s (\lambda y. e'[x^* := v]) = e2[x^* := v]$.

The case for when $e1 \rightarrow_s e2$ is by SSR-7 is similar to that for SSR-6. \square

Lemma B.10. *If $e1 \rightarrow_s e2$, then $e[x^* := e1] \rightarrow_s^* e[x^* := e2]$.*

PROOF. By induction on e .

If $e = x^*$, then $e[x^* := e1] = e1$ and $e[x^* := e2] = e2$, and $e1 \rightarrow_s e2$ implies $e[x^* := e1] \rightarrow_s^* e[x^* := e2]$.

If $e = y^* \neq x^*$, then $e[x^* := e1] = y^* = e[x^* := e2]$, and $e[x^* := e1] \rightarrow_s^* e[x^* := e2]$ is derivable by reflexivity.

If $e = (\lambda y. e')$, then $e[x^* := e1] = (\lambda y. e'[x^* := e1])$ and $e[x^* := e2] = (\lambda y. e'[x^* := e2])$. By induction, $e'[x^* := e1] \rightarrow_s^* e'[x^* := e2]$. Therefore, $(\lambda y. e'[x^* := e1]) \rightarrow_s^* (\lambda y. e'[x^* := e2])$ can be derived by repeated uses of SSR-6.

If $e = (\lambda y^\circ. e')$, then $e[x^* := e1] = (\lambda y^\circ. e'[x^* := e1])$ and $e[x^* := e2] = (\lambda y^\circ. e'[x^* := e2])$. By induction, $e'[x^* := e1] \rightarrow_s^* e'[x^* := e2]$. Therefore, $(\lambda y^\circ. e'[x^* := e1]) \rightarrow_s^* (\lambda y^\circ. e'[x^* := e2])$ can be derived by repeated uses of SSR-7.

If $e = a1 a2$, then $e[x^* := e1] = (a1[x^* := e1]) (a2[x^* := e1])$ and $e[x^* := e2] = (a1[x^* := e2]) (a2[x^* := e2])$. By induction, $a1[x^* := e1] \rightarrow_s^* a1[x^* := e2]$ and $a2[x^* := e1] \rightarrow_s^* a2[x^* := e2]$. Then $e[x^* := e1] \rightarrow_s^* e[x^* := e2]$ can be derived by repeated uses of SSR-2 and SSR-3. \square

Definition B.1. *A step $e1 \rightarrow_s e2$ is called a **weak step** if $e1 \rightarrow_w e2$ is derivable. Otherwise, it is called a **strong step**.*

Lemma B.11. *If $e \rightarrow_s v$ and e is closed, then either $e \rightarrow_s v$ is a weak step, or e is a value.*

PROOF. By induction on the derivation of $e \rightarrow_s v$.

If $e \rightarrow_s v$ is by SSR-1, then $e \rightarrow_w v$ is derivable by WSSR-1. If $e \rightarrow_s v$ is by SSR-2 or SSR-3, then v is an application. Contradiction. If $e \rightarrow_s v$ is by SSR-5, then e is not closed. Contradiction. If $e \rightarrow_s v$ is by SSR-6, then $e = \lambda x. e'$, so e is a value. If $e \rightarrow_s v$ is by SSR-7, then $e = \lambda x^\circ. e'$, so e is a value. \square

Lemma B.12. *If $e \rightarrow_s^* v$ and e is closed, then either e is a value or $e \rightarrow_s^* v$ contains at least one weak step.*

PROOF. By induction on the number n of steps in $e \rightarrow_s^* v$.

Suppose $n = 0$. Then $e = v$.

Suppose $n = n' + 1$. Then $e \rightarrow_s^{n'} e' \rightarrow_s v$. Since e is closed, so is e' . By Lemma B.11 on $e' \rightarrow_s v$, either $e' \rightarrow_s v$ is a weak step, or else e' is a value. If $e' \rightarrow_s v$ is a weak step, then $e \rightarrow_s^* v$ contains at least one weak step. If e' is a value, then by induction on $e \rightarrow_s^{n-1} e'$, either e is a value or $e \rightarrow_s^{n-1} e'$ contains at least one weak step. If $e \rightarrow_s^{n-1} e'$ contains at least one weak step, then so does $e \rightarrow_s^n v$. \square

Lemma B.13. *Let e, e_1 , and e' be closed terms. If $e \rightarrow_s e_1 \rightarrow_w e'$, then there exist closed terms e_2 and e_3 such that $e \rightarrow_w e_2 \rightarrow_w^* e_3 \rightarrow_s^* e'$.*

PROOF. **Suppose $e_1 \rightarrow_w e'$ is by WSSR-1.**

Then $e_1 = (\lambda x. a)v$ and $e' = a[x:=v]$. Now, consider the step $e \rightarrow_s e_1$.

If $e \rightarrow_s e_1$ is by SSR-1, then $e \rightarrow_w e_1$ and we have $e \rightarrow_w e_1 \rightarrow_w e'$ as required.

If $e \rightarrow_s e_1$ is by SSR-2, then $e = e_1' v$ and $e_1' \rightarrow_s (\lambda x. a)$. By Lemma B.11, then either 1) $e_1' \rightarrow_w (\lambda x. a)$, or 2) e_1' is a value. If 1), we have $e_1' v \rightarrow_w (\lambda x. a) v \rightarrow_w a[x:=v]$ as required. If 2), then $e_1' = (\lambda x. a')$ and it must be that that $a' \rightarrow_s a$. By Lemma B.9, $a'[x:=v] \rightarrow_s a[x:=v]$. Therefore, $e = (\lambda x. a')v \rightarrow_w a'[x:=v] \rightarrow_s a[x:=v] = e'$.

If $e \rightarrow_s e_1$ is by SSR-3, then $e = (\lambda x. a)b$ and $b \rightarrow_s v$. By Lemma B.11, $b \rightarrow_w v$ or else b is a value v' . In the first case, we have $(\lambda x. a)b \rightarrow_w (\lambda x. a)v \rightarrow_w a[x:=v]$ as required. In the second case, we have $v' \rightarrow_s v$, so by Lemma B.10, $a[x:=v'] \rightarrow_s^* a[x:=v]$. Therefore, $e = (\lambda x. a)v' \rightarrow_w a[x:=v'] \rightarrow_s^* a[x:=v] = e'$.

If $e \rightarrow_s e_1$ is by SSR-4, then $e \rightarrow_w e_1$ by WSSR-4, so $e \rightarrow_w e_1 \rightarrow_w e'$ and the conclusion holds.

If $e \rightarrow_s e_1$ is by SSR-5, then $e = (\lambda x. a)y$, which is not closed. Contradiction.

If $e \rightarrow_s e_1$ is by SSR-6, then $e_1 = (\lambda x. e_1')$, which contradicts that $e_1 = (\lambda x. a)v$.

If $e \rightarrow_s e_1$ is by SSR-7, then $e_1 = (\lambda x^\circ. e_1')$, which contradicts that $e_1 = (\lambda x. a)v$.

Suppose $e_1 \rightarrow_w e'$ is by WSSR-4. This case is similar to the case for WR-1.

Suppose $e_1 \rightarrow_w e'$ is by WSSR-2. Then $e_1 = a_1 b$ and $e' = a_2 b$ and $a_1 \rightarrow_w a_2$. Now consider $e \rightarrow_s e_1$.

If $e \rightarrow_s e_1$ is by SSR-1, then $e \rightarrow_w e_1 \rightarrow_w^* e$.

If $e \rightarrow_s e_1$ is by SSR-2, then $e = a_0 b$ and $a_0 \rightarrow_s a_1$. By induction, there exist terms a_3, a_4 such that $a_0 \rightarrow_w a_3 \rightarrow_w^* a_4 \rightarrow_s^* a_2$. Therefore, $e = a_0 b \rightarrow_w a_3 b \rightarrow_w^* a_4 b \rightarrow_s^* a_2 b = e'$.

If $e \rightarrow_s e_1$ is by SSR-3, then $e = a_1 b_0$ and $b_0 \rightarrow_s b$. Then $e = a_1 b_0 \rightarrow_w a_2 b_0 \rightarrow_s a_2 b = e'$.

If $e \rightarrow_s e_1$ is by SSR-4, then $e \rightarrow_w e_1$ by rule WSSR-4, so $e \rightarrow_w e_1 \rightarrow_w e'$ and the conclusion holds.

If $e \rightarrow_s e_1$ is by SSR-5, then e is not closed. Contradiction.

If $e \rightarrow_s e_1$ is by SSR-6 or SSR-7, then e_1 is a λ -abstraction. Contradicts $e_1 = a_1 b$.

Suppose $e_1 \rightarrow_w e'$ is by WSSR-3. Then $e_1 = (\lambda x. a)b_1$ and $e' = (\lambda x. a)b_2$ and $b_1 \rightarrow_w b_2$. Now consider $e \rightarrow_s e_1$.

If $e \rightarrow_s e_1$ is by SSR-1, then $e \rightarrow_w e_1 \rightarrow_w^* e$.

If $e \rightarrow_s e_1$ is by SSR-2, then $e = a_0 b_1$ and $a_0 \rightarrow_s (\lambda x. a)$. By Lemma B.11, either 1) $a_0 \rightarrow_w (\lambda x. a)$, or 2) $a_0 = (\lambda x. a')$. If 1), we have $a_0 b_1 \rightarrow_w (\lambda x. a) b_1 \rightarrow_w (\lambda x. a) b_2$ as required. If 2), we have $a_0 = (\lambda x. a')$, so $a' \rightarrow_s a$. Therefore, $e = (\lambda x. a') b_1 \rightarrow_w (\lambda x. a') b_2 \rightarrow_s (\lambda x. a) b_2 = e'$.

If $e \rightarrow_s e_1$ is by SSR-3, then $e = (\lambda x. a) b_0$ and $b_0 \rightarrow_s b_1$. By induction, there exist closed terms b_3, b_4 such that $b_0 \rightarrow_w b_3 \rightarrow_w^* b_4 \rightarrow_s^* b_2$. Therefore, $e = (\lambda x. a) b_0 \rightarrow_w (\lambda x. a) b_3 \rightarrow_w^* (\lambda x. a) b_4 \rightarrow_s^* (\lambda x. a) b_2 = e'$.

If $e \rightarrow_s e1$ is by SSR-4, then $e \rightarrow_w e1$ by rule WSSR-4, so $e \rightarrow_w e1 \rightarrow_w e'$ and the conclusion holds.

If $e \rightarrow_s e1$ is by SSR-5, then e is not closed. Contradiction.

If $e \rightarrow_s e1$ is by SSR-6 or SSR-7, then $e1$ is a λ -abstraction. Contradicts $e1 = (\lambda x. a)b2$.

Suppose $e1 \rightarrow_w e'$ **is by WSSR-5.** This case is similar to the case for WR-3.

□

Lemma B.14. *Suppose $e1 \rightarrow_w e2$ is derived without using WSSR-4. Then $e1 \rightarrow e2$.*

PROOF. By induction on the derivation of $e1 \rightarrow_w e2$.

If $e1 \rightarrow_w e2$ is by WSSR-1, then $e1 \rightarrow e2$ is derivable by CBV-1.

If $e1 \rightarrow_w e2$ is by WSSR-2, then $e1 = a1 b$ and $e2 = a2 b$ and $a1 \rightarrow_w a2$ is derived without using WSSR-4. By induction, $a1 \rightarrow a2$, so $e1 = a1 b \rightarrow a2 b = e2$ is derivable by CBV-2.

If $e1 \rightarrow_w e2$ is by WSSR-3, then $e1 = (\lambda x. e) a1$ and $e2 = (\lambda x. e) a2$ and $a1 \rightarrow_w a2$ is derived without using WSSR-4. By induction, $a1 \rightarrow a2$, so $e1 = (\lambda x. e) a1 \rightarrow (\lambda x. e) a2 = e2$ is derivable by CBV-3.

If $e1 \rightarrow_w e2$ is by WSSR-5, then $e1 = (\lambda x^\circ. e) a1$ and $e2 = (\lambda x^\circ. e) a2$ and $a1 \rightarrow_w a2$ is derived without using WSSR-4. By induction, $a1 \rightarrow a2$, so $e1 = (\lambda x^\circ. e) a1 \rightarrow (\lambda x^\circ. e) a2 = e2$ is derivable by CBV-5.

□

Lemma B.15. *Suppose $e1 \rightarrow_w^i e2$ contains no uses of WSSR-4. Then $e1 \rightarrow^i e2$.*

PROOF. By induction on the length of $e1 \rightarrow_w^i e2$.

Suppose $i = 0$. Trivial.

Suppose $i = i' + 1$. Then we have $e1 \rightarrow_w e1' \rightarrow_w^{i'} e2$. By induction $e1' \rightarrow^{i'} e2$. Since $e1 \rightarrow_w e1'$ is derived without using WSSR-4, Lemma B.14 states $e1 \rightarrow e1'$. Therefore, $e1 \rightarrow^{i'+1} e2$. □

Lemma B.16. *If $\Gamma; \Sigma \vdash e$ and $e \rightarrow_w^i e'$ and $e \rightarrow^j v$, then there exists an i' such that $e' \rightarrow^{i'} v$ and $i + i' \leq j$.*

PROOF. By induction on the number of WSSR-4 reductions in $e \rightarrow_w^i e'$.

If $e \rightarrow_w^i e'$ contains no WSSR-4 reductions, then by Lemma B.15, $e \rightarrow^i e'$. Since call-by-value reduction is deterministic, $e \rightarrow^i e'$ must be a prefix of $e \rightarrow^j v$.

If $e \rightarrow_w^i e'$ contains $n > 0$ WSSR-4 reductions, then there exist $e1, e2, i1,$ and $i2$ such that $e \rightarrow^{i1} e1 \rightarrow_w e2 \rightarrow_w^{i2} e'$, where $i = i1 + i2 + 1$ and $e1 \rightarrow_w e2$ is the first WSSR-4 reduction in the sequence. So $e1 = (\lambda x^\circ. a)b$, and $e2 = a[x^\circ := b]$.

Since $e \rightarrow^{i1} (\lambda x^\circ. a)b$, the derivation of $e \rightarrow^j v$ is of the form $e \rightarrow^{i1} (\lambda x^\circ. a)b \rightarrow^{j-i1} v$. By Lemma B.1, there exists a $v1$ such that $b \rightarrow^{j1} v1$ and $a[x^\circ := v1] \rightarrow^{j-i1-j1-1} v$. By Lemma B.8, $a[x^\circ := b] \rightarrow^{j2} v$ for some $j2$.

We have that $a[x^\circ := b] \rightarrow_w^{i2} e'$ contains $n-1$ WSSR-4 reductions, and $a[x^\circ := b] \rightarrow^{j2} v$. Therefore, it follows by induction that $e' \rightarrow^{i'} v$ and $i2 + i' \leq j2$.

Since $b \rightarrow^{j1} v1$, and $a[x^\circ := v1] \rightarrow^{j-i1-j1-1} v$, and $a[x^\circ := b] \rightarrow^{j2} v$, it follows by Lemma B.7 that $j2 \leq j1 + j - i1 - j1 - 1 = j - i1 - 1$. Therefore, $i2 + i' \leq j - i1 - 1$. Therefore, $i + i' = i1 + i2 + 1 + i' \leq i1 + 1 + (j - i1 - 1) = j$. □

Corollary B.1. *If $e \rightarrow_w^i v1$ and $e \rightarrow^j v2$, then $i \leq j$ and $v1 = v2$.*

PROOF. Suppose $e \rightarrow_w^i v1$ and $e \rightarrow^j v2$. By Lemma B.16, there exists an i' such that $v1 \rightarrow^{i'} v2$ and $i + i' \leq j$. But i' must equal 0, so $i \leq j$ and $v1 = v2$. □

Lemma B.17. *Suppose $\langle \cdot \rangle; \langle \cdot \rangle \vdash a$, and $a \rightarrow_s^* b$ is contains i weak steps, and $a \rightarrow^n v$. Then there exists a reduction sequence $a \rightarrow_w^j a' \rightarrow_s^* b$ such that $i \leq j$ and $a' \rightarrow_s^* b$ contains no weak steps.*

PROOF. For a given sequence of steps $e1 \rightarrow_s^* e2$, let p be the number of weak steps before the first strong step; and let q be the length of the first subsequence of strong steps. By Lemma B.16, $p \leq n$, so $n - p \geq 0$. We proceed by induction on the lexicographic order of $(n - p, q)$.

Suppose $n - p = 0$, so $p = n$. Then there is an a' such that $a \rightarrow_w^n a' \rightarrow_s^* b$. Since $a \rightarrow_w^n a'$ and $a \rightarrow^n v$, Lemma B.16 states that there is an n' such that $a' \rightarrow^{n'} v$ and $n + n' \leq n$. Then $n' = 0$, so $a' = v$. We have $v \rightarrow_s^* b$, which contains no weak steps because v is a value. Therefore $i = n$. Let $j = n$, so $j \leq i$. Since $a \rightarrow_w^j v \rightarrow_s^* b$ and $j \leq i$ and $v \rightarrow_s^* b$ contains no weak steps, the conclusion holds.

Suppose $n - p > 0$, so $p < n$. If $q = 0$, then we have $a \rightarrow_w^p b$, and the conclusion holds. If $q \geq 1$, we have $a \rightarrow_w^p a1 \rightarrow_s^q b$. If $a1 \rightarrow_s^q b$ contains only strong steps, then the conclusion holds with $a' = a1$. Otherwise, we have $a \rightarrow_w^p a1 \rightarrow_s^{q-1} a2 \rightarrow_s a3 \rightarrow_w a4 \rightarrow_s^* b$, where $a3 \rightarrow_w a4$ is the first weak reduction step that follows the first sequence of strong reduction steps.

Since $a2 \rightarrow_s a3 \rightarrow_w a4$, Lemma B.13 states that there exist terms $c1$ and $c2$ such that $a2 \rightarrow_w c1 \rightarrow_w^* c2 \rightarrow_s^* a4$.

Replace $a2 \rightarrow_s a3 \rightarrow_w a4$ with $a2 \rightarrow_w c1 \rightarrow_w^* c2 \rightarrow_s^* a4$ in the full sequence. If the original sequence had i weak steps, the new sequence has $i' \geq i$ weak steps.

Suppose now that $q = 1$. Then we added $k \geq 1$ reduction steps to the initial sequence of weak reduction steps, so the new sequence will have measure $(n - (p + k), l)$ for some $l \geq 0$. Since $n - (p + k) < n - p$, we have that $(n - (p + k), l) < (n, q)$. Therefore, by induction, there exists a weak reduction sequence $a \rightarrow_w^j a' \rightarrow_s^* b$ such that $j \geq i' \geq i$ and $a' \rightarrow_s^* b$ contains no weak steps.

Suppose instead that $q > 1$. Then the new sequence will have measure $(n - p, q - 1)$. Since $q - 1 < q$, $(n - p, q - 1) < (n - p, q)$. Therefore, by induction, there exists a weak reduction sequence $a \rightarrow_w^j a' \rightarrow_s^* b$ such that $j \geq i' \geq i$ and $a' \rightarrow_s^* b$ contains no weak steps. \square

Lemma B.18 (Preservation of termination behavior). *Suppose $\langle \cdot \rangle; \langle \cdot \rangle \vdash a$, and Suppose $a \rightarrow^* v$ and $a \rightarrow_s^* a1$. Then there exists a value $v1$ such that $a1 \rightarrow^* v1$.*

PROOF. By contradiction.

Suppose there is no such value $v1$. Then for any n , there exists a term $a2$ such that $a1 \rightarrow^n a2$. Now suppose $a \rightarrow^* v$ contains i steps. Then $a \rightarrow_s^* a1 \rightarrow_s^{i+1} a2$ contains at least $i + 1$ steps of weak reduction. By Lemma B.17, there exists a reduction sequence $a \rightarrow_w^j a3 \rightarrow_s^* a2$ such that $i + 1 \leq j$. But by Lemma B.16, $a \rightarrow_w^j a3$ and $a \rightarrow^i v$ implies $a3 \rightarrow^k v$ for some k such that $j + k \leq i$. Therefore, we have $j \leq j + k \leq i < i + 1 \leq j$. Contradiction. \square

THEOREM 4.5. *If $\langle \cdot \rangle; \langle \cdot \rangle \vdash e$ and $e \rightarrow_s^* e'$ and $e \rightarrow^i v$, then there exists an $i' \leq i$ and a value v' such that $e' \rightarrow^{i'} v'$, and $v \rightarrow_s^* v'$.*

PROOF. By Lemma B.18, there exists an i' and a value v' such that $e' \rightarrow^{i'} v'$.

Therefore, we have that $e \rightarrow_s^* e' \rightarrow^{i'} v'$, and since each \rightarrow step is also a \rightarrow_w step, we have $e \rightarrow_s^* e' \rightarrow_w^{i'} v'$. Suppose $e \rightarrow_s^* e' \rightarrow_w^{i'} v'$ contains $i1$ steps of weak reduction. It is clear that $i' \leq i1$. By Lemma B.17, there exists an a' such that $e \rightarrow_w^j a' \rightarrow_s^* v'$ for some j , and $i1 \leq j$, and $a' \rightarrow_s^* v'$ contains no weak steps. Since e is closed, a' is closed, so by Lemma B.12 on $a' \rightarrow_s^* v'$, a' is a value v'' . By Corollary B.1, $j \leq i$ and $v'' = v$. Therefore, $i' \leq i1 \leq j \leq i$, and $v \rightarrow_s^* v'$. \square

LEMMA 5.2. *For any self-interpreter u for a representation function $\widehat{\cdot}$, and for any term p , if $u \widehat{p} \Rightarrow_s p$, then mix is Jones-optimal for u and time_{cbv} .*




Fig. 23. Illustration of Theorem 4.5

$$\begin{array}{c}
 \overline{x^* \triangleright x^*} \\
 e \triangleright q \\
 \hline
 \lambda x^*. e \triangleright \text{abs} (\lambda x^*. q) \\
 e1 \triangleright q1 \quad e2 \triangleright q2 \\
 \hline
 e1 e2 \triangleright \text{app} q1 q2 \\
 e \triangleright q \\
 \hline
 \bar{e} = \lambda \text{abs}. \lambda \text{app}. q
 \end{array}$$

$$u = \lambda e^\circ. e^\circ (\lambda x^\circ. x^\circ) (\lambda x^\circ. x^\circ)$$

Fig. 24. Mogensen's Church-encoded representation self-interpreter with affine variable annotations.

PROOF. Let u be an unannotated self-interpreter, let p be an unannotated term, and suppose $u \bar{p} \Rightarrow_s p$. We want to show that for any unannotated term d , $\text{time}_{cbv}(p' d) \leq \text{time}_{cbv}(p d)$, where $\text{mix } \bar{u} \bar{p} \equiv_\beta \bar{p}'$. By the specification of mix , we have that $\bar{p}' \equiv_\beta \text{erase}(\text{NF}_s(\text{annotate}(u \bar{p})))$. By the definition of $\bar{\cdot}$, $p' \equiv_\alpha \text{erase}(\text{NF}_s(\text{annotate}(u \bar{p})))$. Let p'' be the annotated term $\text{NF}_s(\text{annotate}(u \bar{p}))$. Then $\text{erase}(p'') = p'$. Since $u \bar{p} \Rightarrow_s p$, we have that $\text{annotate}(u \bar{p}) \rightarrow_s^* p'''$ for some annotated term p''' such that $\text{erase}(p''') = p$. Since \rightarrow_s is confluent, $p'' = \text{NF}_s(\text{annotate}(u \bar{p})) = \text{NF}_s(p''')$. Therefore, $p'' \rightarrow_s^* p'''$. By Corollary 4.6, $\text{time}_{cbv}(p'' \text{annotate}(d)) \leq \text{time}_{cbv}(p''' \text{annotate}(d))$. By Lemma 5.1, $\text{time}_{cbv}(\text{erase}(p'' \text{annotate}(d))) \leq \text{time}_{cbv}(p''' \text{annotate}(d))$. Therefore, $\text{time}_{cbv}(p' d) \leq \text{time}_{cbv}(p d)$ as required. \square

C JONES OPTIMALITY FOR ULC

C.1 Church encoding

Figure 24 shows Mogensen's Church encoding and self-interpreter u with affine variable annotations. Our theorems do not depend on how the term is annotated before quotation – it can be unannotated or maximally annotated, and the theorems hold.

Lemma C.1. *If $\langle \rangle; \langle \rangle \vdash e$ and $e \triangleright q$, then $\langle \rangle; \langle \rangle, \text{abs}, \text{app}; \langle \rangle \vdash q$*

Note: abs and app may be affine or not, depending on how many abstractions or applications are in e . Optimality does not depend on whether abs and app are affine or universal. We assume they are universal.

PROOF. Straightforward induction. □

Lemma C.2. *If $\langle \rangle; \langle \rangle \vdash e$, $\langle \rangle; \langle \rangle \vdash \bar{e}$*

PROOF. Straightforward by Lemma C.1. □

Lemma C.3. $\langle \rangle; \langle \rangle \vdash \text{unquote}$

PROOF. Trivial. □

Lemma C.4. *If $e \ q$, then $q[\text{abs} := (\lambda x^\circ. x^\circ), \text{app} := (\lambda x^\circ. x^\circ)] \rightarrow_s^* e$.*

PROOF. By induction on e .

Suppose $e = x^*$. Then $q = x^*$, so $q[\text{abs} := (\lambda x^\circ. x^\circ), \text{app} := (\lambda x^\circ. x^\circ)] = e$.

Suppose $e = (\lambda y^*. e')$. Then $q = \text{abs } (\lambda y^*. q')$ and $e' \ q'$. By induction, $q'[\text{abs} := (\lambda x^\circ. x^\circ), \text{app} := (\lambda x^\circ. x^\circ)] \rightarrow_s^* e'$. Therefore, we can derive:

$$\begin{aligned} & q[\text{abs} := (\lambda x^\circ. x^\circ), \text{app} := (\lambda x^\circ. x^\circ)] \\ &= (\text{abs } (\lambda y^*. q'))[\text{abs} := (\lambda x^\circ. x^\circ), \text{app} := (\lambda x^\circ. x^\circ)] \\ &= (\lambda x^\circ. x^\circ) (\lambda y^*. q'[\text{abs} := (\lambda x^\circ. x^\circ), \text{app} := (\lambda x^\circ. x^\circ)]) \\ &\rightarrow_s^* (\lambda x^\circ. x^\circ) (\lambda y^*. e') \\ &\rightarrow_s (\lambda y^*. e') \\ &= e \end{aligned}$$

Suppose $e = e_1 \ e_2$. Then $q = \text{app } q_1 \ q_2$ and $e_1 \ q_1$ and $e_2 \ q_2$. By induction, $q_1[\text{abs} := (\lambda x^\circ. x^\circ), \text{app} := (\lambda x^\circ. x^\circ)] \rightarrow_s^* e_1$ and $q_2[\text{abs} := (\lambda x^\circ. x^\circ), \text{app} := (\lambda x^\circ. x^\circ)] \rightarrow_s^* e_2$. Therefore, we can derive:

$$\begin{aligned} & q[\text{abs} := (\lambda x^\circ. x^\circ), \text{app} := (\lambda x^\circ. x^\circ)] \\ &= (\text{app } q_1 \ q_2)[\text{abs} := (\lambda x^\circ. x^\circ), \text{app} := (\lambda x^\circ. x^\circ)] \\ &= (\lambda x^\circ. x^\circ) (q_1[\text{abs} := (\lambda x^\circ. x^\circ), \text{app} := (\lambda x^\circ. x^\circ)]) (q_2[\text{abs} := (\lambda x^\circ. x^\circ), \text{app} := (\lambda x^\circ. x^\circ)]) \\ &\rightarrow_s^* (\lambda x^\circ. x^\circ) e_1 (q_2[\text{abs} := (\lambda x^\circ. x^\circ), \text{app} := (\lambda x^\circ. x^\circ)]) \\ &\rightarrow_s^* (\lambda x^\circ. x^\circ) e_1 \ e_2 \\ &\rightarrow_s e_1 \ e_2 \\ &= e \end{aligned}$$

□

LEMMA 5.3. *For any closed term $e, u \ \bar{e} \Longrightarrow_s e$.*

PROOF. By definition, $\bar{e} = \lambda \text{abs}. \lambda \text{app}. q$, where $e \ q$. By Lemma C.4, $q[\text{abs} := (\lambda x^\circ. x^\circ), \text{app} := (\lambda x^\circ. x^\circ)] \rightarrow_s^* e$. Therefore, we can derive:

$$\begin{aligned} & \text{unquote } \bar{e} \\ &= (\lambda e^\circ. e^\circ (\lambda x^\circ. x^\circ) (\lambda x^\circ. x^\circ)) \bar{e} \\ &\rightarrow_s \bar{e} (\lambda x^\circ. x^\circ) (\lambda x^\circ. x^\circ) \\ &= (\lambda \text{abs}. \lambda \text{app}. q) (\lambda x^\circ. x^\circ) (\lambda x^\circ. x^\circ) \\ &\rightarrow_s^2 q[\text{abs} := (\lambda x^\circ. x^\circ), \text{app} := (\lambda x^\circ. x^\circ)] \\ &\rightarrow_s^* e \end{aligned}$$

□

C.2 Mogensen-Scott encoding

In this section we prove Lemma 5.5, which states that for any term e , $\text{msint } \widehat{e} \Longrightarrow_s e$. We proceed by defining annotated versions of msint and then proving for the annotated version that $\text{msint } \widehat{e} \rightarrow_s^* e$.

Define:

$$\text{fix}_f = (\lambda x. f (\lambda y^\circ. x x y^\circ)) (\lambda x. f (\lambda y^\circ. x x y^\circ))$$

$$\text{fix} = \lambda f. \text{fix}_f$$

$$u = \lambda u. \lambda e^\circ. e^\circ (\lambda x^\circ. x^\circ) (\lambda f. \lambda x. u (f ((\lambda y^\circ. x)x))) (\lambda f. \lambda x^\circ. u f (u x^\circ))$$

$$\text{msint} = \text{fix } u$$

Lemma C.5. $\langle \rangle; \langle \rangle \vdash \text{fix}$

PROOF. Straightforward. □

Lemma C.6. $\text{msint} \rightarrow_s \text{fix}_u$

PROOF.

$$\text{msint}$$

$$= \text{fix } u$$

$$\rightarrow_s (\lambda x. u (\lambda y^\circ. x x y^\circ)) (\lambda x. u (\lambda y^\circ. x x y^\circ))$$

$$= \text{fix}_u$$

□

Lemma C.7. $\text{fix}_u \rightarrow_s u (\lambda y^\circ. \text{fix}_u y^\circ)$

PROOF.

$$\text{fix}_u$$

$$= (\lambda x. u (\lambda y^\circ. x x y^\circ)) (\lambda x. u (\lambda y^\circ. x x y^\circ))$$

$$\rightarrow_s u (\lambda y^\circ. (\lambda x. u (\lambda y^\circ. x x y^\circ)) (\lambda x. u (\lambda y^\circ. x x y^\circ)) y^\circ)$$

$$= u (\lambda y^\circ. \text{fix}_u y^\circ)$$

□

Lemma C.8. For any e , $\text{fix}_u \widehat{e} \rightarrow_s^* e$.

PROOF. By induction on e .

The annotations on \widehat{e} don't affect this proof, so we will leave all variables its variables unannotated.

Suppose $e = x$. Then $\widehat{e} = \lambda \text{var}. \lambda \text{abs}. \lambda \text{app}. \text{var } x$. Therefore, we can derive:

$$\text{fix}_u \widehat{e}$$

$$\rightarrow_s u (\lambda y^\circ. \text{fix}_u y^\circ) \widehat{e}$$

$$\rightarrow_s^2 \widehat{e} (\lambda x^\circ. x^\circ)$$

$$(\lambda f. \lambda x. (\lambda y^\circ. \text{fix}_u y^\circ) (f ((\lambda y^\circ. x)x)))$$

$$(\lambda f. \lambda x^\circ. (\lambda y^\circ. \text{fix}_u y^\circ) f ((\lambda y^\circ. \text{fix}_u y^\circ) x^\circ))$$

$$\rightarrow_s^3 (\lambda x^\circ. x^\circ) x$$

$$\rightarrow_s x$$

$$= e$$

Suppose $e = \lambda x. e'$. Then $\widehat{e} = \lambda \text{var}. \lambda \text{abs}. \lambda \text{app}. \text{abs} (\lambda x. \widehat{e}')$. By induction, $\text{fix}_u \widehat{e}' \rightarrow_s^* e'$. Therefore, can derive:

$$\begin{aligned}
& \text{fix}_u \widehat{e} \\
\rightarrow_s & u (\lambda y^\circ. \text{fix}_u y^\circ) \widehat{e} \\
\rightarrow_s^2 & \widehat{e} (\lambda x^\circ. x^\circ) \\
& (\lambda f. \lambda x. (\lambda y^\circ. \text{fix}_u y^\circ) (f ((\lambda y^\circ. x) x))) \\
& (\lambda f. \lambda x^\circ. (\lambda y^\circ. \text{fix}_u y^\circ) f ((\lambda y^\circ. \text{fix}_u y^\circ) x^\circ)) \\
\rightarrow_s^3 & (\lambda f. \lambda x. (\lambda y^\circ. \text{fix}_u y^\circ) (f ((\lambda y^\circ. x) x))) (\lambda x. \widehat{e}') \\
\rightarrow_s & \lambda x. (\lambda y^\circ. \text{fix}_u y^\circ) ((\lambda x. \widehat{e}') ((\lambda y^\circ. x) x)) \\
\rightarrow_s & \lambda x. (\lambda y^\circ. \text{fix}_u y^\circ) ((\lambda x. \widehat{e}') x) \\
\rightarrow_s & \lambda x. (\lambda y^\circ. \text{fix}_u y^\circ) \widehat{e}' \\
\rightarrow_s & \lambda x. \text{fix}_u \widehat{e}' \\
\rightarrow_s^* & \lambda x. e' \\
& = e
\end{aligned}$$

Suppose $e = e_1 e_2$. Then $\widehat{e} = \lambda \text{var}. \lambda \text{abs}. \lambda \text{app}. \text{app} \widehat{e}_1 \widehat{e}_2$. By induction, $\text{fix}_u \widehat{e}_1 \rightarrow_s^* e_1$ and $\text{fix}_u \widehat{e}_2 \rightarrow_s^* e_2$. Therefore, we can derive:

$$\begin{aligned}
& \text{fix}_u \widehat{e} \\
\rightarrow_s & u (\lambda y^\circ. \text{fix}_u y^\circ) \widehat{e} \\
\rightarrow_s^2 & \widehat{e} (\lambda x^\circ. x^\circ) \\
& (\lambda f. \lambda x. (\lambda y^\circ. \text{fix}_u y^\circ) (f ((\lambda y^\circ. x) x))) \\
& (\lambda f. \lambda x^\circ. (\lambda y^\circ. \text{fix}_u y^\circ) f ((\lambda y^\circ. \text{fix}_u y^\circ) x^\circ)) \\
\rightarrow_s^3 & (\lambda f. \lambda x^\circ. (\lambda y^\circ. \text{fix}_u y^\circ) f ((\lambda y^\circ. \text{fix}_u y^\circ) x^\circ)) \widehat{e}_1 \widehat{e}_2 \\
\rightarrow_s^2 & (\lambda y^\circ. \text{fix}_u y^\circ) \widehat{e}_1 ((\lambda y^\circ. \text{fix}_u y^\circ) \widehat{e}_2) \\
\rightarrow_s^2 & \text{fix}_u \widehat{e}_1 (\text{fix}_u \widehat{e}_2) \\
\rightarrow_s^* & e_1 (\text{fix}_u \widehat{e}_2) \\
\rightarrow_s^* & e_1 e_2 \\
& = e
\end{aligned}$$

□

LEMMA 5.5. For any term e , $\text{msint } \widehat{e} \Longrightarrow_s e$.

PROOF. By Lemma C.6, $\text{msint } \widehat{e} \rightarrow_s \text{fix}_u \widehat{e}$, and by Lemma C.8, $\text{fix}_u \widehat{e} \rightarrow_s^* e$. □

C.3 deBruijn Indices

Figure 13 defines our tagless-final representation using deBruijn indices to represent variables.

Here are annotated versions of `fst` and `snd`:

$$\begin{aligned}
\text{fst} &= (\lambda x^\circ. x^\circ (\lambda y^\circ. \lambda z^\circ. y)) \\
\text{snd} &= (\lambda x^\circ. x^\circ (\lambda y^\circ. \lambda z^\circ. z^\circ))
\end{aligned}$$

Define the mapping $\overline{\cdot}$ from contexts of unannotated variables to annotated terms as follows:

$$\begin{aligned}
\overline{\langle \rangle} &= (\lambda y^\circ. y^\circ) \\
\overline{\Gamma, x} &= \lambda f^\circ. f^\circ x \overline{\Gamma}
\end{aligned}$$

Note that every unannotated variable x becomes unlimited in the annotated term.

LEMMA C.1. If $\Gamma \vdash x \blacktriangleright q$, then $q \overline{\Gamma} \rightarrow_s^* x$.

PROOF. By induction on the derivation of $\Gamma \vdash x \blacktriangleright q$.

Suppose $\Gamma \vdash x \blacktriangleright q$ is by the first rule. Then $\Gamma = \Gamma', x$ and $q = \text{fst}$. We have:

$$\begin{aligned}
& q \bar{\Gamma} \\
&= \text{fst} (\lambda f^\circ. f^\circ x_i \bar{\Gamma}') \\
&= (\lambda x^\circ. x^\circ (\lambda y. \lambda z^\circ. y)) \\
&\quad (\lambda f^\circ. f^\circ x_i \bar{\Gamma}') \\
&\rightarrow_s (\lambda f^\circ. f^\circ x_i \bar{\Gamma}') (\lambda y. \lambda z^\circ. y) \\
&\rightarrow_s (\lambda y. \lambda z^\circ. y) x_i \bar{\Gamma}' \\
&\rightarrow_s^2 x_i
\end{aligned}$$

where the last step relies on that x_i is unlimited.

Suppose $\Gamma \vdash x \blacktriangleright q$ is by the second rule. Then $\Gamma = \Gamma', y$, and Then $\Gamma' \vdash x \blacktriangleright q'$, and $q = (\lambda p^\circ. q' (\text{snd } p^\circ))$. By induction, $q' \bar{\Gamma}' \rightarrow_s^* x$. Therefore, we have:

$$\begin{aligned}
& q \bar{\Gamma} \\
&= (\lambda p^\circ. q' (\text{snd } p^\circ)) \\
&\quad (\lambda f^\circ. f^\circ x_i \bar{\Gamma}') \\
&\rightarrow_s q' (\text{snd } (\lambda f^\circ. f^\circ x_i \bar{\Gamma}')) \\
&\rightarrow_s q' ((\lambda f^\circ. f^\circ x_i \bar{\Gamma}') (\lambda y^\circ. \lambda z^\circ. z^\circ)) \\
&\rightarrow_s q' ((\lambda y^\circ. \lambda z^\circ. z^\circ) x_i \bar{\Gamma}') \\
&\rightarrow_s^2 q' \bar{\Gamma}' \\
&\rightarrow_s^* x
\end{aligned}$$

□

LEMMA C.2. *If $\Gamma \vdash e \triangleright q$, then there exists a value v such that $q[\text{var} := \text{dbVar}, \text{abs} := \text{dbAbs}, \text{app} := \text{dbApp}] \rightarrow_s^* v$ and $v \bar{\Gamma} \rightarrow_s^* e$.*

PROOF. By induction on the derivation of $\Gamma \vdash e \triangleright q$.

Suppose e is a variable x . Then $\Gamma \vdash e \blacktriangleright q'$, and $q = \text{var } q'$. Let $v = q'$. We have that:

$$\begin{aligned}
& q[\text{var} := \text{dbVar}, \text{abs} := \text{dbAbs}, \text{app} := \text{dbApp}] \\
&= \text{dbVar } q' \\
&= (\lambda f. \lambda e^\circ. f e) q' \\
&\rightarrow_s q'
\end{aligned}$$

where the last step is because q' is a value. The conclusion follows from Lemma C.1, which says $q' \bar{\Gamma} \rightarrow_s^* x$.

Suppose $e = (\lambda x. e')$. Then $\Gamma, x \vdash e' \triangleright q'$, and $q = \text{abs } q'$. By induction,

$q'[\text{var} := \text{dbVar}, \text{abs} := \text{dbAbs}, \text{app} := \text{dbApp}] \rightarrow_s^* v'$, and $v' \bar{\Gamma}, x \rightarrow_s^* e'$. Therefore, we have:

$$\begin{aligned}
& q[\text{var} := \text{dbVar}, \text{abs} := \text{dbAbs}, \text{app} := \text{dbApp}] \\
&= \text{dbAbs } (q'[\text{var} := \text{dbVar}, \text{abs} := \text{dbAbs}, \text{app} := \text{dbApp}]) \\
&\rightarrow_s^* \text{dbAbs } v' \\
&= (\lambda b. \lambda e. \lambda x. b (\lambda f^\circ. f^\circ x e)) v' \\
&\rightarrow_s (\lambda e. \lambda x. v' (\lambda f^\circ. f^\circ x e))
\end{aligned}$$

Let $v = (\lambda e. \lambda x. v' (\lambda f^\circ. f^\circ x e))$. Then, we have:

$$\begin{aligned}
& v \bar{\Gamma} \\
&= (\lambda e. \lambda x. v' (\lambda f^\circ. f^\circ x e)) \bar{\Gamma} \\
&\rightarrow_s (\lambda x. v' (\lambda f^\circ. f^\circ x \bar{\Gamma})) \quad (\bar{\Gamma} \text{ is a value}) \\
&= (\lambda x. v' \bar{\Gamma}') \quad (x \text{ is unlimited}) \\
&\rightarrow_s^* (\lambda x. e') \\
&= e
\end{aligned}$$

Suppose $e = e_1 e_2$. Then $\Gamma \vdash e_1 \triangleright q_1$, and $\Gamma \vdash e_2 \triangleright q_2$, and $q = \text{app } q_1 q_2$. By induction, $q_1[\text{var} := \text{dbVar}, \text{abs} := \text{dbAbs}, \text{app} := \text{dbApp}] \rightarrow_s^* v_1$ and $v_1 \bar{\Gamma} \rightarrow_s^* e_1$, and

$q2[\text{var} := \text{dbVar}, \text{abs} := \text{dbAbs}, \text{app} := \text{dbApp}] \rightarrow_s^* v2$ and $v2 \bar{\Gamma} \rightarrow_s^* e2$. Therefore, we have:

$$\begin{aligned} & q[\text{var} := \text{dbVar}, \text{abs} := \text{dbAbs}, \text{app} := \text{dbApp}] \\ &= \text{dbApp} (q1[\text{var} := \text{dbVar}, \text{abs} := \text{dbAbs}, \text{app} := \text{dbApp}]) \\ & \quad (q2[\text{var} := \text{dbVar}, \text{abs} := \text{dbAbs}, \text{app} := \text{dbApp}]) \\ &\rightarrow_s^* \text{dbApp } v1 \ v2 \\ &= (\lambda a. \lambda b. \lambda e. a \ e \ (b \ e)) \ v1 \ v2 \\ &\rightarrow_s^2 (\lambda e. v1 \ e \ (v2 \ e)) \end{aligned}$$

Let $v = (\lambda e. v1 \ e \ (v2 \ e))$. Then we have that:

$$\begin{aligned} & v \ \bar{\Gamma} \\ &= (\lambda e. v1 \ e \ (v2 \ e)) \ \bar{\Gamma} \\ &\rightarrow_s^2 v1 \ \bar{\Gamma} \ (v2 \ \bar{\Gamma}) \\ &\rightarrow_s^* e1 \ e2 \end{aligned} \quad \square$$

LEMMA 5.7. *For any closed term e , $\text{dbint } \widehat{e} \Longrightarrow_s e$.*

PROOF. We have that $\widehat{e} = (\lambda \text{var}. \lambda \text{abs}. \lambda \text{app}. q)$ and $\langle \rangle \vdash e \triangleright q$. By Lemma C.2,

$$\begin{aligned} & q[\text{var} := \text{dbVar}, \text{abs} := \text{dbAbs}, \text{app} := \text{dbApp}] \langle \rangle \rightarrow_s^* e. \text{ Therefore, we have:} \\ & \quad \text{dbint } \widehat{e} \\ &= (\lambda q^\circ. q \ \text{dbVar} \ \text{dbAbs} \ \text{dbApp} \ (\lambda x^\circ. x^\circ)) \ \widehat{e} \\ &\rightarrow_s (\widehat{e} \ \text{dbVar} \ \text{dbAbs} \ \text{dbApp} \ (\lambda x^\circ. x^\circ)) \\ &\rightarrow_s^3 q[\text{var} := \text{dbVar}, \text{abs} := \text{dbAbs}, \text{app} := \text{dbApp}] \ (\lambda x^\circ. x^\circ) \\ &= q[\text{var} := \text{dbVar}, \text{abs} := \text{dbAbs}, \text{app} := \text{dbApp}] \langle \rangle \\ &\rightarrow_s^* e \end{aligned} \quad \square$$

D REPRESENTATION THEOREMS FOR $F_\omega^{\mu i}$

In this section, we prove that every closed and well-typed $F_\omega^{\mu i}$ term has a representation in $F_\omega^{\mu i}$, and that all representations produced by our quoter are β -normal forms.

The following definition relates the environment used to typecheck a term with the environment used to typecheck its pre-representation.

Definition D.1 (Environment mapping for pre-representations).

$$\begin{aligned} \overline{\langle \rangle} &= \langle \rangle \\ \overline{\Gamma, X:K} &= \bar{\Gamma}, X:K \\ \overline{\Gamma, x^*:T} &= \bar{\Gamma}, x^*:V \ T \end{aligned}$$

Lemma D.1. *If $\Gamma \vdash T : K$, then $\bar{\Gamma} \vdash T : K$*

PROOF. Straightforward, since $\bar{\cdot}$ does not affect the presence, order, or kinds of type variables in the environment. \square

Lemma D.2. *If $\Gamma \vdash e : T$ and e contains no free term variables, then $\bar{\Gamma} \vdash e : T$*

PROOF. Straightforward, using Lemma D.1. \square

Lemma D.3. *If $\Gamma \vdash (\forall X:K. T) : *$, then*

- (1) $\Gamma \vdash \text{tcAll}_{X,K,T} : \text{TcAll} \ (\forall X:K. T)$
- (2) $\Gamma \vdash \text{unAll}_{X,K,T} : \text{UnAll} \ (\forall X:K. T)$
- (3) $\Gamma \vdash \text{isAll}_{X,K,T} : \text{IsAll} \ (\forall X:K. T)$

PROOF. Suppose $\Gamma \vdash (\forall X:K. T) : *$. Then $\Gamma, X:K \vdash T : *$.

1) $\text{tcAll}_{X,K,T} = \Delta \text{Arr}:* \rightarrow * \rightarrow *$. $\Delta \text{Out}:* \rightarrow *$. $\Delta \text{In}:* \rightarrow *$. $\Delta \text{Mu}:((* \rightarrow *) \rightarrow * \rightarrow *) \rightarrow * \rightarrow *$. $\text{refl}(\text{Out}(\forall X:K. \text{In } T))$. It is easily checked that $\text{tcAll}_{X,K,T}$ has the type $\text{Eq}(\text{Out}(\forall X:K. \text{In } T))(\text{Out}(\forall X:K. \text{In } T))$. But $\text{Typecase Arr Out In Mu}(\forall X:K. T) \equiv \text{Out}(\forall X:K. \text{In } T)$, and also $\text{All Out In}(\forall X:K. T) \equiv \text{Out}(\forall X:K. \text{In } T)$, so we have that $\text{tcAll}_{X,K,T}$ has the type $\text{TcAll}(\forall X:K. T)$ as required.

2) Follows from reasoning similar to 1.

3) Follows from 1 and 2. \square

Lemma D.4. *If $\Gamma \vdash (\forall X:K. T) : *$ and $\Gamma \vdash S : K$, then*

- (1) $\bar{\Gamma} \vdash \text{underAll}_{X,K,T} : \text{UnderAll}(\forall X:K. T)$
- (2) $\bar{\Gamma} \vdash \text{stripAll}_K : \text{StripAll}(\forall X:K. T)$
- (3) $\bar{\Gamma} \vdash \text{inst}_{X,K,T,S} : \text{Inst}(\forall X:K. T)(T[X:=S])$

PROOF. First, note that $\text{underAll}_{X,K,T}$, stripAll_K , and $\text{inst}_{X,K,T,S}$ contain no free term variables. Therefore, by Lemma D.2 it is sufficient to show

- (1) $\Gamma \vdash \text{underAll}_{X,K,T} : \text{UnderAll}(\forall X:K. T)$
- (2) $\Gamma \vdash \text{stripAll}_K : \text{StripAll}(\forall X:K. T)$
- (3) $\Gamma \vdash \text{inst}_{X,K,T,S} : \text{Inst}(\forall X:K. T)(T[X:=S])$

For 1, it is easily checked that $\Gamma \vdash \text{underAll}_{X,K,T} : (\forall F1:* \rightarrow *. \forall F2:* \rightarrow *. (\forall A:*. F1 A \rightarrow F2 A) \rightarrow (\forall X:K. F1 T) \rightarrow (\forall X:K. F2 T))$. The result follows from the type equivalences $\text{All Id } F1(\forall X:K. T) \equiv (\forall X:K. F1 T)$ and $\text{All Id } F2(\forall X:K. T) \equiv (\forall X:K. F2 T)$.

The cases for 2 and 3 are similar. \square

Lemma D.5. *If $\Gamma \vdash e : T$, then there exists a unique q such that $\Gamma \vdash e : T \ q$ and $\bar{\Gamma} \vdash q : V T$.*

PROOF. By straightforward induction on the derivation of $\Gamma \vdash e : T$, using the types of the constructors, Lemma D.3, and Theorem D.4.

Suppose $\Gamma \vdash e : T$ is by the rule for variables. Then $e = x^*$ and $(x^*:T) \in \Gamma$ and $\Gamma \vdash x^* : T \ x^*$. By the definition of $\bar{\cdot}$, $(x^* : V T) \in \bar{\Gamma}$. Therefore, $\bar{\Gamma} \vdash x^* : V T$ as required.

Suppose $\Gamma \vdash e : T$ is by the rule for λ -abstractions. Then $e = (\lambda^*x:T1. e1)$ and $\Gamma, x^*:T1 \vdash e1 : T2$ and $T = T1 \rightarrow T2$. Also, $\Gamma, x^*:T1 \vdash e1 : T2 \ q1$ and $\Gamma \vdash (\lambda x^*:T1. e1) \ \text{abs } V T1 T2 (\lambda x^*:V T1. q1)$ By induction, $q1$ is unique and $\bar{\Gamma}, x^*:T1 \vdash q1 : V T2$. But $\bar{\Gamma}, x^*:T1 = \bar{\Gamma}, x:V T1$, so $\bar{\Gamma} \vdash (\lambda x:V T1. q1) : V T1 \rightarrow V T2$. Therefore, $q = \text{abs } V T1 T2 (\lambda x:PE\text{xp } V T1. q1)$ is unique and by the type of abs , it is easily checked that $\bar{\Gamma} \vdash \text{abs } V T1 T2 (\lambda x:V T1. q1) : V (T1 \rightarrow T2)$.

The cases for applications, and fold and unfold expressions are similar.

Suppose $\Gamma \vdash e : T$ is by the rule for type abstractions. Then $e = (\Lambda X:K. e1)$ and $\Gamma, X:K \vdash e1 : T1$ and $T = (\forall X:K. T1)$. Also, $\Gamma, X:K \vdash e1 : T1 \ q1$ and $\Gamma \vdash (\Lambda X:K. e1) \ \text{tabs } V (\forall X:K. T1) \ p \ s \ u (\Lambda X:K. q1)$, where $p = \text{isAll}_{X,K,T}$, $s = \text{stripAll}_K = s$, and $u = \text{underAll}_{X,K,T}$. By induction, $q1$ is unique and $\bar{\Gamma}, X:K \vdash q1 : V T1$. Since $\bar{\Gamma}, X:K = \bar{\Gamma}, X:K$, we have that $\bar{\Gamma} \vdash (\Lambda X:K. q1) : (\forall X:K. V T1)$. It follows from $(\text{All Id } V (\forall X:K. T1)) \equiv (\forall X:K. V T1)$ that $\bar{\Gamma} \vdash (\Lambda X:K. q1) : (\text{All Id } V (\forall X:K. T1))$. By Theorem D.3, $\Gamma \vdash p : \text{IsAll}(\forall X:K. T)$. But p does not contain any free term variables, so by Lemma D.2 we have that $\bar{\Gamma} \vdash p : \text{IsAll}(\forall X:K. T)$. By Theorem D.4, we have that $\bar{\Gamma} \vdash s : \text{StripAll}(\forall X:K. T)$ and $\bar{\Gamma} \vdash u : \text{UnderAll}(\forall X:K. T)$. Therefore, $q = \text{tabs } V (\forall X:K. T1) \ p \ s \ u (\Lambda X:K. q1)$ is unique and by the type of tabs , we have that $\bar{\Gamma} \vdash \text{tabs } V (\forall X:K. T1) \ p \ s \ u (\Lambda X:K. q1) : V (\forall X:K. T1)$ as required.

The case for type applications is similar.

Suppose $\Gamma \vdash e : T$ is by the rule for type conversion. The $\Gamma \vdash e : S$ and $S \equiv T$. By the induction hypothesis, there exists a unique q such that $\Gamma \vdash e : S \rightarrow q$ and $\bar{\Gamma} \vdash q : \text{PExp } V \ S$. Since $V \ S \equiv V \ T$, we have that $\bar{\Gamma} \vdash q : V \ T$ as required. \square

THEOREM 6.1. *If $\langle \rangle \vdash e : T$, then $\langle \rangle \vdash \bar{e} : \text{Exp } T$.*

PROOF. Follows straightforwardly from Theorem D.5. \square

Lemma D.6. *For any type T , $\text{refl } T$ is β -normal.*

PROOF. Straightforward. \square

Lemma D.7. *If $\Gamma \vdash (\forall X:K. T) : *$, then $\text{isAll}_{X,K,T}$ is β -normal.*

PROOF. Straightforward using Lemma D.6, which states that $\text{tcAll}_{X,K,T}$ and $\text{unAll}_{X,K,T}$ are β -normal. \square

Lemma D.8. *If $\Gamma \vdash (\forall X:K. T) : *$ and $\Gamma \vdash S : K$, then $\text{inst}_{X,K,T,S}$ is β -normal.*

PROOF. Straightforward. \square

Lemma D.9. *For any kind K , stripAll_K is β -normal.*

PROOF. Straightforward. \square

Lemma D.10. *If $\Gamma \vdash (\forall X:K. T) : *$, then $\text{underAll}_{X,K,T}$ is β -normal.*

PROOF. Straightforward. \square

Lemma D.11. *If $\Gamma \vdash e : T \rightarrow q$, then q is β -normal.*

PROOF. By straightforward induction on the derivation of $\Gamma \vdash e : T \rightarrow q$, using Lemmas D.7, D.8, D.9, and D.10. \square

THEOREM 6.2. *If $\langle \rangle \vdash e : T$, then \bar{e} is β -normal.*

PROOF. We have that $\bar{e} = \Lambda V : * \rightarrow *. q$ and $\langle \rangle \vdash e : T \rightarrow q$. By Lemma D.11, q is β -normal. Therefore \bar{e} is also. \square

E JONES OPTIMALITY FOR $F_{\omega}^{\mu i}$

In this section, we prove that a strong optimization under type-erasure can reduce all interpretation work for $F_{\omega}^{\mu i}$. We define the original and type-erased versions of the helper functions used by the quoter, quotation itself, and the original and type-erased versions of the self-interpreter `unquote`. The representation and interpreter are tagless-final-style – essentially a typed version of the Church encoding for ULC from Section C.1.

Define:

$$\theta_u(e) = e[\text{abs} := \text{unAbs}, \text{app} := \text{unApp}, \\ \text{tabs} := \text{unTAbs}, \text{tapp} := \text{unTApp}, \\ \text{fld} := \text{unUnfold}, \text{unfld} := \text{unUnfold}]$$

Lemma E.1. *If $\Gamma \vdash e : T \rightarrow q$, then $te(\theta_u(q)) \rightarrow_s^* te(e)$.*


```

decl Eq : * → * =
  λA:*. λB:*. ∀F:* → *. F A → F B;

decl refl : (∀A:*. Eq A A) =
  ΛA:*. ΔF:* → *. λx° : F A. x°;

decl sym : (∀A:*. ∀B:*. Eq A B → Eq B A) =
  ΛA:*. ΛB:*. λeq° : Eq A B.
  eq° (λT:*. Eq T A) (refl A);

decl trans : (∀A:*. ∀B:*. ∀C:*.
  Eq A B → Eq B C → Eq A C) =
  ΛA:*. ΛB:*. ΛC:*. λeqAB:Eq A B. λeqBC:Eq B C.
  ΔF:* → *. λx°:F A. eqBC F (eqAB F x°);

decl eqApp : (∀A:*. ∀B:*. ∀F:* → *.
  Eq A B → Eq (F A) (F B)) =
  ΛA:*. ΛB:*. ΔF:* → *. λeq° : Eq A B.
  eq° (λT:*. Eq (F A) (F T)) (refl (F A));

decl arrL : (∀A1:*. ∀A2:*. ∀B1:*. ∀B2:*.
  Eq (A1 → A2) (B1 → B2) →
  Eq A1 B1) =
  ΛA1 A2 B1 B2. eqApp (A1→A2) (B1→B2) ArrL;

decl arrR : (∀A1:*. ∀A2:*. ∀B1:*. ∀B2:*.
  Eq (A1 → A2) (B1 → B2) →
  Eq A2 B2) =
  ΛA1 A2 B1 B2. eqApp (A1→A2) (B1→B2) ArrR;

decl coerce : (∀A:*. ∀B:*. Eq A B → A → B) =
  ΛA:*. ΛB:*. λeq°:Eq A B. eq° Id;

```

Fig. 25. Implementation of type equality proofs in $F_{\omega}^{\mu i}$.

```

decl reflte = λx°. x°;
decl symte = λeq°. eq° reflte;
decl transte = λeqAB. λeqBC. λx°. eqBC (eqAB x°);
decl eqAppte = λeq°. eq° reflte;
decl arrLte = eqAppte;
decl arrRte = eqAppte;
decl coercete = λeq°. eq°;

```

Fig. 26. The type-erasure of type equality proofs.

PROOF. By induction on the derivation of $\Gamma \vdash e : T \ q$.

Suppose $e = x^*$. Then $te(e) = x^*$, and $q = te(q) = x^*$, and $\theta_u(te(q)) = x^* = te(e)$.

```

decl TcAll : * → * =
  λT. VArr. VOut. VIn. VMu.
  Eq (Typecase Arr Out In Mu T) (All Out In T);
decl UnAll : * → * =
  λT. VOut. Eq (All Out (λA:*. A) T) (Out T);
decl IsAll : * → * = λT. VA. TcAll T → UnAll T → A;

decl tcAll : IsAll T → TcAll T =
  λp° : IsAll T. p° (TcAll T) (λx:TcAll T. λy°:UnAll T. x);
decl unAll : IsAll T → UnAll T =
  λp° : IsAll T. p° (UnAll T) (λx°:TcAll T. λy°:UnAll T. y°);

tcAllX,K,T = ΛArr. ΛOut. ΛIn. ΛMu. refl (Out (VX:K. In T))
unAllX,K,T = ΛOut. refl (Out (VX:K. T))
isAllX,K,T =
  ΛA. λf°:TcAll (VX:K. T) → UnAll (VX:K. T) → A.
  f° tcAllX,K,T unAllX,K,T

```

Fig. 27. IsAll proofs.

```

decl tcAllte = λp°. p° (λx. λy°. x);
decl unAllte = λp°. p° (λx°. λy°. y°);

tcAllteX,K,T = reflte
unAllteX,K,T = reflte
isAllteX,K,T = λf°. f° tcAllteX,K,T unAllteX,K,T

```

Fig. 28. Type-erasure of IsAll proofs.

```

decl Id : * → * = λA:*. A;

decl UnderAll : * → * =
  λT:*. VF1:* → *. VF2:* → *.
  (VA:*. F1 A → F2 A) →
  All Id F1 A → All Id F2 A;
decl StripAll : * → * =
  λT:*. VA:*. All Id (λB:*. A) T → A;
decl Inst : * → * → * =
  λA:*. λB:*. VF:* → *. All Id F A → F B;

underAllX,K,T =
  ΔF1. ΔF2. λf : (VA:*. F1 A → F2 A).
  λe° : (VX:K. F1 T). ΔX:K. f T (e° X)
stripAllK = ΛA. λe°:(VX:K. A). e° TK
instX,K,T,S = ΔF. λf°:(VX:K. F T). f° S

```

Fig. 29. underAll, stripAll, and inst functions

```

underAllx,K,Tte = λf. λe°. f e°
stripAllKte = λe°. e°
instx,K,T,Ste = λf°. f°

```

Fig. 30. Type-erasure of underAll, stripAll, and inst functions

```

decl unAbste = λf°. f°;
decl unAppte = λf°. f°;
decl unTAbste = λp. λs. λu. λe°. unAllte p e°;
decl unTAppte = λp. λi. λx°. i (symte (unAllte p) x°);
decl unFoldte = λx°. x°;
decl unUnfoldte = λx°. x°;

decl unquotete =
  λe°. e° unAbste unAppte unTAbste unTAppte unFoldte unUnfoldte

```

Fig. 31. The type-erasure of unquote.

Suppose $e = (\lambda x^*: T1. e1)$. Then $q = \text{abs } T1 \ T2 \ (\lambda x^*: V \ T1. q1)$, and $\Gamma, (x^*: T1) \vdash e1 : T2 \ q1$. By induction, $te(\theta_u(q1)) \rightarrow_s^* te(e1)$. Therefore, we can derive:

$$\begin{aligned}
& te(\theta_u(q)) \\
&= te(\text{unAbs } T1 \ T2 \ (\lambda x^*: V \ T1. \theta_u(q1))) \\
&= \text{unAbs}^{te} (\lambda x^*. te(\theta_u(q1))) \\
&= (\lambda f^\circ. f^\circ) (\lambda x^*. te(\theta_u(q1))) \\
&\rightarrow_s (\lambda x^*. te(\theta_u(q1))) \\
&\rightarrow_s^* (\lambda x^*. te(e1)) \\
&= te(e)
\end{aligned}$$

Suppose $e = e1 \ e2$. Then $q = \text{app } T2 \ T \ q1 \ q2$ and $\Gamma \vdash e1 : T2 \rightarrow T \ q1$ and $\Gamma \vdash e2 : T2 \ q2$. By induction, $te(\theta_u(q1)) \rightarrow_s^* te(e1)$ and $te(\theta_u(q2)) \rightarrow_s^* te(e2)$. Therefore, we can derive:

$$\begin{aligned}
& te(\theta_u(q)) \\
&= te(\text{unApp } T2 \ T \ \theta_u(q1) \ \theta_u(q2)) \\
&= \text{unApp}^{te} te(\theta_u(q1)) \ te(\theta_u(q2)) \\
&= (\lambda f^\circ. f^\circ) te(\theta_u(q1)) \ te(\theta_u(q2)) \\
&\rightarrow_s te(\theta_u(q1)) \ te(\theta_u(q2)) \\
&\rightarrow_s^* te(e1) \ te(e2) \\
&= te(e1 \ e2) \\
&= te(e)
\end{aligned}$$

Suppose $e = \Lambda X:K. e1$. Then $q = \text{tabs } (\forall X:K. T) \text{ isAll}_{X,K,T} \text{ stripAll}_K \text{ underAll}_{X,K,T} (\Lambda X:K. q1)$ and $\Gamma, (X:K) \vdash e1 : T \ q1$. By induction, $te(\theta_u(q1)) \rightarrow_s^* te(e1)$. Therefore, we can derive:

$$\begin{aligned}
& te(\theta_u(q)) \\
&= te(\text{unTABS } (\forall X:K. T) \text{ isAll}_{X,K,T} \text{ stripAll}_K \text{ underAll}_{X,K,T} (\Lambda X:K. \theta_u(q1))) \\
&= \text{unTABS}^{te} \text{ isAll}_{X,K,T}^{te} \text{ stripAll}_K^{te} \text{ underAll}_{X,K,T}^{te} te(\theta_u(q1)) \\
&= (\lambda p. \lambda s. \lambda u. \lambda e^\circ. \text{unAll}^{te} p e^\circ) \\
&\quad \text{isAll}_{X,K,T}^{te} \text{ stripAll}_K^{te} \text{ underAll}_{X,K,T}^{te} te(\theta_u(q1)) \\
&\rightarrow_s^4 \text{unAll}^{te} \text{ isAll}_{X,K,T}^{te} te(\theta_u(q1)) \\
&= (\lambda p^\circ. p^\circ (\lambda x^\circ. \lambda y^\circ. y^\circ)) \text{ isAll}_{X,K,T}^{te} te(\theta_u(q1)) \\
&\rightarrow_s \text{isAll}_{X,K,T}^{te} (\lambda x^\circ. \lambda y^\circ. y^\circ) te(\theta_u(q1)) \\
&= (\lambda f^\circ. f^\circ \text{tcAll}_{X,K,T}^{te} \text{unAll}_{X,K,T}^{te} (\lambda x^\circ. \lambda y^\circ. y^\circ) te(\theta_u(q1))) \\
&\rightarrow_s (\lambda x^\circ. \lambda y^\circ. y^\circ) \text{tcAll}_{X,K,T}^{te} \text{unAll}_{X,K,T}^{te} te(\theta_u(q1)) \\
&\rightarrow_s^2 \text{unAll}_{X,K,T}^{te} te(\theta_u(q1)) \\
&= \text{refl}^{te} te(\theta_u(q1)) \\
&= (\lambda x^\circ. x^\circ) te(\theta_u(q1)) \\
&\rightarrow_s te(\theta_u(q1)) \\
&\rightarrow_s^* te(e1) \\
&= te(\Lambda X:K. e1) \\
&= te(e)
\end{aligned}$$

Suppose $e = e1 \ A$. Then $q = \text{tapp } (\forall X:K. T) (\top[X:=A]) \text{ isAll}_{X,K,T} \text{ inst}_{X,K,T,A} q1$ and $\Gamma \vdash e1 : (\forall X:K. T) \ q1$. By induction, $te(\theta_u(q1)) \rightarrow_s^* e1$. Therefore, we can derive:

$$\begin{aligned}
& te(\theta_u(q)) \\
&= te(\text{unTApp } (\forall X:K. T) (\top[X:=A]) \text{ isAll}_{X,K,T} \text{ inst}_{X,K,T,A} \theta_u(q1)) \\
&= \text{unTApp}^{te} \text{ isAll}_{X,K,T}^{te} \text{ inst}_{X,K,T,A}^{te} te(\theta_u(q1)) \\
&= (\lambda p. \lambda i. \lambda x^\circ. i (\text{sym}^{te} (\text{unAll}^{te} p) x^\circ)) \\
&\quad \text{isAll}_{X,K,T}^{te} \text{ inst}_{X,K,T,A}^{te} te(\theta_u(q1)) \\
&\rightarrow_s^3 \text{inst}_{X,K,T,A}^{te} (\text{sym}^{te} (\text{unAll}^{te} \text{ isAll}_{X,K,T}^{te}) te(\theta_u(q1))) \\
&= (\lambda f^\circ. f^\circ) (\text{sym}^{te} (\text{unAll}^{te} \text{ isAll}_{X,K,T}^{te}) te(\theta_u(q1))) \\
&\rightarrow_s \text{sym}^{te} (\text{unAll}^{te} \text{ isAll}_{X,K,T}^{te}) te(\theta_u(q1)) \\
&= (\lambda eq^\circ. eq^\circ \text{refl}^{te}) (\text{unAll}^{te} \text{ isAll}_{X,K,T}^{te}) te(\theta_u(q1)) \\
&\rightarrow_s \text{unAll}^{te} \text{ isAll}_{X,K,T}^{te} \text{refl}^{te} te(\theta_u(q1)) \\
&= (\lambda p^\circ. p^\circ (\lambda x^\circ. \lambda y^\circ. y^\circ)) \text{ isAll}_{X,K,T}^{te} \text{refl}^{te} te(\theta_u(q1)) \\
&\rightarrow_s \text{isAll}_{X,K,T}^{te} (\lambda x^\circ. \lambda y^\circ. y^\circ) \text{refl}^{te} te(\theta_u(q1)) \\
&= \text{isAll}_{X,K,T}^{te} (\lambda x^\circ. \lambda y^\circ. y^\circ) \text{refl}^{te} te(\theta_u(q1)) \\
&= (\lambda f^\circ. f^\circ \text{tcAll}_{X,K,T}^{te} \text{unAll}_{X,K,T}^{te} (\lambda x^\circ. \lambda y^\circ. y^\circ) \text{refl}^{te} te(\theta_u(q1))) \\
&\rightarrow_s (\lambda x^\circ. \lambda y^\circ. y^\circ) \text{tcAll}_{X,K,T}^{te} \text{unAll}_{X,K,T}^{te} \text{refl}^{te} te(\theta_u(q1)) \\
&\rightarrow_s^2 \text{unAll}_{X,K,T}^{te} \text{refl}^{te} te(\theta_u(q1)) \\
&= \text{refl}^{te} \text{refl}^{te} te(\theta_u(q1)) \\
&= (\lambda x^\circ. x^\circ) (\lambda x^\circ. x^\circ) te(\theta_u(q1)) \\
&\rightarrow_s^2 te(\theta_u(q1)) \\
&\rightarrow_s^* te(e1) \\
&= te(e1 \ A) \\
&= te(e)
\end{aligned}$$

Suppose $e = \text{fold } F \ T \ e1$. Then $q = \text{fld } F \ T \ q1$ and $\Gamma \vdash e1 : F \ (F) \ T \ q1$. By induction $te(\theta_u(q1)) \rightarrow_s^* te(e1)$. Therefore, we can derive:

$$\begin{aligned}
& te(\theta_u(q)) \\
&= te(\theta_u(\text{fld } F \ T \ q1)) \\
&= te(\text{unFold } F \ T \ \theta_u(q1)) \\
&= \text{unFold}^{te} te(\theta_u(q1)) \\
&= (\lambda x^\circ. x^\circ) te(\theta_u(q1)) \\
&\rightarrow_s te(\theta_u(q1)) \\
&\rightarrow_s^* te(e1) \\
&= te(\text{fold } F \ T \ e1) \\
&= te(e)
\end{aligned}$$

Suppose $e = \text{unfold } F \ T \ e1$. Then $q = \text{unfld } F \ T \ q1$ and $\Gamma \vdash e1 : F \ T \ q1$. By induction $te(\theta_u(q1)) \rightarrow_s^* te(e1)$. Therefore, we can derive:

$$\begin{aligned}
& te(\theta_u(q)) \\
&= te(\theta_u(\text{unfld } F \ T \ q1)) \\
&= te(\text{unUnfold } F \ T \ \theta_u(q1)) \\
&= \text{unUnfold}^{te} te(\theta_u(q1)) \\
&= (\lambda x^\circ. x^\circ) te(\theta_u(q1)) \\
&\rightarrow_s te(\theta_u(q1)) \\
&\rightarrow_s^* te(e1) \\
&= te(\text{unfold } F \ T \ e1) \\
&= te(e)
\end{aligned}$$

□

LEMMA 6.3. *If $\langle \rangle \vdash e : T$, then $te(\text{unquote } \bar{e}) \Longrightarrow_s te(e)$.*

PROOF. Suppose $\langle \rangle \vdash e : T$. Then $\langle \rangle \vdash e : T \ q$ and $\bar{e} = \Lambda V : * \rightarrow *$.

$$\begin{aligned}
& \lambda \text{abs} : \text{Abs } V. \ \lambda \text{app} : \text{App } V. \\
& \lambda \text{tabs} : \text{TAbs } V. \ \lambda \text{tapp} : \text{TApp } V. \\
& \lambda \text{fld} : \text{Fold } V. \ \lambda \text{unfld} : \text{Unfold } V. \\
& q
\end{aligned}$$

Therefore, $te(\bar{e}) = \lambda \text{abs}. \lambda \text{app}. \lambda \text{tabs}. \lambda \text{tapp}. \lambda \text{fld}. \lambda \text{unfld}. te(q)$.

By induction, $te(\theta_u(q)) \rightarrow_s^* te(e)$.

Therefore, we can derive:

$$\begin{aligned}
& te(\text{unquote } \bar{e}) \\
&= \text{unquote}^{te} te(\bar{e}) \\
&= \text{unquote}^{te} (\lambda \text{abs}. \lambda \text{app}. \lambda \text{tabs}. \lambda \text{tapp}. \lambda \text{fld}. \lambda \text{unfld}. te(q)) \\
&= (\lambda e^\circ. e^\circ \text{unAbs}^{te} \text{unApp}^{te} \text{unTAbs}^{te} \text{unTApp}^{te} \text{unFold}^{te} \text{unUnfold}^{te}) \\
&\quad (\lambda \text{abs}. \lambda \text{app}. \lambda \text{tabs}. \lambda \text{tapp}. \lambda \text{fld}. \lambda \text{unfld}. te(q)) \\
&\rightarrow_s^7 te(q) [\text{abs} := \text{unAbs}^{te}, \text{app} := \text{unApp}^{te}, \text{tabs} := \text{unTAbs}^{te}, \text{tapp} := \text{unTApp}^{te}, \text{fld} := \text{unFold}^{te}, \text{unfld} := \text{unUnfold}^{te}] \\
&= te(\theta_u(q)) \\
&\rightarrow_s^* te(e)
\end{aligned}$$

□