

Compiling without Continuations

Luke Maurer Paul Downen
Zena M. Ariola

University of Oregon, USA
{maurerl,pdownen,ariola}@cs.uoregon.edu

Simon Peyton Jones

Microsoft Research, UK
simonpj@microsoft.com

Abstract

Many fields of study in compilers give rise to the concept of a *join point*—a place where different execution paths come together. Join points are often treated as functions or continuations, but we believe it is time to study them in their own right. We show that adding join points to a direct-style functional intermediate language is a simple but powerful change that allows new optimizations to be performed, including a significant improvement to list fusion. Finally, we report on recent work on adding join points to the intermediate language of the Glasgow Haskell Compiler.

Keywords Intermediate languages; CPS; ANF; Stream fusion; Haskell; GHC

CCS Concepts • **Software and its engineering** → **Compilers**; *Software performance*; *Formal language definitions*; Functional languages; Control structures; • **Theory of computation** → Equational logic and rewriting

1. Introduction

Consider this code, in a functional language:

```
if (if e1 then e2 else e3) then e4 else e5
```

Many compilers will perform a *commuting conversion* [13], which naively would produce:

```
if e1 then (if e2 then e4 else e5)
         else (if e3 then e4 else e5)
```

Commuting conversions are tremendously important in practice (Sec. 2), but there is a problem: the conversion duplicates e_4 and e_5 . A natural countermeasure is to name the offending expressions and duplicate the names instead:

```
let { j4 () = e4; j5 () = e5 }
in if e1 then (if e2 then j4 () else j5 ())
           else (if e3 then j4 () else j5 ())
```

We describe j_4 and j_5 as *join points*, because they say where execution of the two branches of the outer `if` joins up again. The duplication is gone, but a new problem has surfaced: the compiler may allocate closures for locally-defined functions like j_4 and j_5 . That is bad because allocation is expensive. And it is tantalizing because all we are doing here is encoding control flow: it is plain as a pikestaff that the “call” to j_4 should be no more than a jump, with no allocation anywhere. That’s what a C compiler would do! Some code generators can cleverly eliminate the closures, but perhaps not if further transformations intervene.

The reader of Appel’s inspirational book [1] may be thinking “*Just use continuation-passing style (CPS)!*” When expressed over CPS terms, many classic optimizations boil down to β -reduction (*i.e.*, function application), or arithmetic reductions, or variants thereof. And indeed it turns out that commuting conversions fall out rather naturally as well. But using CPS comes at a fairly heavy price: the intermediate language becomes more complicated, some transformations are harder or out of reach, and (unlike direct style) CPS commits to a particular evaluation order (Sec. 8).

Inspired by Flanagan *et al.* [10], the reader may now be thinking “*OK, just use administrative normal form (ANF)!*” That paper shows that many transformations achievable in CPS are equally accessible in direct style. ANF allows an optimizer to *exploit* CPS technology without needing to *implement* it. The motto is: *Think in CPS; work in direct style.*

But alas, a subsequent paper by Kennedy shows that there remain transformations that are inaccessible in ANF but fall out naturally in CPS [16]. So the obvious question is this: could we extend ANF in some way, to get all the goodness of direct style *and* the benefits of CPS? In this paper we say “yes!”, making the following contributions:

- We describe a modest extension to a direct-style λ -calculus intermediate language, namely adding join points (Sec. 3). We give the syntax, type system, and operational semantics, together with optimising transformations.

- We describe how to infer which ordinary bindings are in fact join points (Sec. 4). In a CPS setting this analysis is called *contification* [16], but it looks rather different here.
- We show that join points can be recursive, and that recursive join points open up a new and entirely unexpected (to us) optimization opportunity for fusion (Sec. 5). In particular, this insight fully resolves a long-standing tension between two competing approaches to fusion, namely stream fusion [6] and unfold/destroy fusion [26].
- We give some metatheory in Sec. 6, including type soundness and correctness of the optimizing transformations. We show the safety of adding jumps as a control effect by establishing an equivalence with System F.
- We demonstrate that our approach works at scale, in a state-of-the-art optimizing compiler for Haskell, GHC (Sec. 7). As hoped, adding join points turned out to be a very modest change, despite GHC's scale and complexity. Like any optimization, it does not make every program go faster, but it has a dramatic effect on some.

Overall, adding join points to ANF has an extremely good power-to-weight ratio, and we strongly recommend it to any direct-style compiler. Our title is somewhat tongue-in-cheek, but we now know of no optimizing transformation that is accessible to a CPS compiler but not to a direct-style one.

2. Motivation and Key Ideas

We review compilation techniques for commuting conversions, to expose the challenge that we tackle in this paper. For the sake of concreteness we describe the way things work in GHC. However, we believe that the whole paper is equally applicable to a call-by-value language.

The Case-of-Case Transformation Consider:

```
isNothing :: Maybe a -> Bool
isNothing x = case x of Nothing -> True
                    Just _ -> False
```

```
mHead :: [a] -> Maybe a
mHead ps = case ps of [] -> Nothing
                    (p:_) -> Just p
```

```
null :: [a] -> Bool
null as = isNothing (mHead as)
```

Here `null`¹ is a simple composition of the library functions `isNothing` and `mHead`. When the optimizer works on `null`, it will inline both `isNothing` and `mHead` to yield:

```
null as = case (case as of [] -> Nothing
                  (p:_) -> Just p) of
  { Nothing -> True; Just _ -> False }
```

Executed directly, this would be terribly inefficient; if the argument list is non-empty we would allocate a result `Just p` only to immediately decompose it. We want to move the outer case into the branches of the inner one, like this:

¹Haskell's standard `null` function returns whether a list is empty.

```
null as = case as of
  [] -> case Nothing of Nothing -> True
                    Just z -> False
  p:_ -> case Just p of Nothing -> True
                    Just _ -> False
```

This is a commuting conversion, specifically the *case-of-case transformation*. In this example, it now happens that both inner case expressions scrutinize a data constructor, so they can be simplified, yielding

```
null as = case as of { [] -> True; _:_ -> False }
```

which is exactly the code we would have written for `null` from scratch.

GHC does a tremendous amount of inlining, including across modules or even packages, so commuting conversions like this are very important in practice: they are the key that unlocks a cascade of further optimizations.

Join Points Commuting conversions have a problem, though: *they often duplicate the outer case*. In our example that was OK, but what about

```
case (case v of { p1 -> e1; p2 -> e2 }) of
  { Nothing -> BIG1; Just x -> BIG2 }
```

where `BIG1` and `BIG2` are big expressions? Duplicating these large expressions would risk bloating the compiled code, perhaps exponentially when case expressions are deeply nested [17]. It is easy to avoid this duplication by first introducing an auxiliary let binding:

```
let { j1 () = BIG1; j2 x = BIG2 } in
case (case v of { p1 -> e1; p2 -> e2 }) of
  { Nothing -> j1 (); Just x -> j2 x }
```

Now we can move the outer case expression into the arms of the inner case, without duplicating `BIG1` or `BIG2`, thus:

```
let { j1 () = BIG1; j2 x = BIG2 } in
case v of
  p1 -> case e1 of Nothing -> j1 ()
                    Just x -> j2 x
  p2 -> case e2 of Nothing -> j1 ()
                    Just x -> j2 x
```

Notice that `j2` takes as its parameter the variable bound by the pattern `Just x`, whereas `j1` has no parameters².

Compiling Join Points Efficiently We call `j1` and `j2` *join points* because you can think of them as places where control joins up again, but so far they are perfectly ordinary let-bound functions, and as such they will be allocated as closures in the heap. But that's ridiculous: all that is happening here is control flow splitting and joining up again. A C compiler would generate a jump to a label, not a call to a heap-allocated function closure!

So, right before code generation, GHC performs a simple analysis to identify bindings that can be compiled as join

²The dummy unit parameter is not necessary in a lazy language, but it is in a call-by-value language.

points. This identifies `let`-bound functions that will never be captured in a closure or thunk, and will only be tail-called with exactly the right number of arguments. (We leave the exact criteria for Sec. 4.) These join-point bindings do not allocate anything; instead a tail call to a join point simply adjusts the stack and jumps to the code for the join point.

The case-of-case transformation, including the idea of using `let` bindings to avoid duplication, is very old; for example, both are features of Steele’s Rabbit compiler for Scheme [24]. In Rabbit the transformation is limited to booleans, but the discussion above shows that it generalizes very naturally to arbitrary data types. In this more general form, it has been part of GHC for decades [19]. Likewise, the idea of generating more efficient code for non-escaping `let` bindings is well established in many other compilers [15, 23, 27], dating back to Rabbit and its `BIND-ANALYZE` routine [24].

Preserving and Exploiting Join Points So far so good, but there is a serious problem with recognizing join points only in the back end of the compiler. Consider this expression:

```
case (let j x = BIG in
      case v of { A -> j 1; B -> j 2; C -> True })
of { True  -> False; False -> True }
```

Here `j` is a join point. Now suppose we do case-of-case on this expression. Treating the binding for `j` as an ordinary `let` binding (as GHC does today), we move the outer case past the `let`, and duplicate it into the branches of the inner case:

```
let j x = BIG in
case v of
  A -> case j 1 of { True -> False; False -> True }
  B -> case j 2 of { True -> False; False -> True }
  C -> case True of { True -> False; False -> True }
```

The third branch simplifies nicely, but the first two do not. There are two distinct problems:

1. The binding for `j` is no longer a join point (it is not tail-called), so the super-efficient code generation strategy does not apply, and the compiler will allocate a closure for `j` at runtime. This happens in practice: we have cases in which GHC’s optimizer actually *increases* allocation because it inadvertently destroys a join point.
2. Even worse, the two copies of the outer case now scrutinize an uninformative call like `(j 1)`. So the extra code bloat from duplicating the outer case is entirely wasted.

And it’s a huge lost opportunity, as we shall see.

So *it is not enough to generate efficient code for join points; we must identify, preserve, and exploit them.* In our example, if the optimizer *knew* that the binding for `j` is a join point, it could exploit that knowledge to transform our original expression like this:

```
let j x = case BIG of True  -> False
                False -> True
in case v of
  A -> j 1
  B -> j 2
  C -> case True of { True -> False; False -> True }
```

This is much, much better than our previous attempt:

- The outer case has moved into the right-hand side of the join point, so it now scrutinizes `BIG`. That’s good, because `BIG` might be a data constructor or a case expression (which would expose another case-of-case opportunity). So the outer case now scrutinizes the actual result of the expression, rather than an uninformative join-point call. That solves problem (2).
- The `A` and `B` branches do not mention the outer case, because it has moved into the join point itself. So `j` is still tail-called and remains an efficiently-compiled join point. That solves problem (1).
- The outer case still scrutinizes the branches that do not finish with a join point call, e.g. the `C` branch.

The Key Idea Thus motivated, in the rest of this paper we explore the following very simple idea:

- Distinguish certain `let` bindings as *join-point bindings*, and their (tail-)call sites as *jumps*. This, by itself, is not new; see Section 9.
- Adjust the case-of-case transformation to take account of join-point bindings and jumps.
- In all the other transformations carried out by the compiler, ensure that join points remain join points.

Our key innovation is that, by recognising join points as a language construct, we both *preserve* join points through subsequent transformations and *exploit* them to make those transformations more effective. Next, we formalize this approach; subsequent sections develop the consequences.

3. System F_J : Join Points and Jumps

We now formalize the intuitions developed so far by describing System F_J , a small intermediate language with join points. F_J is an extension of GHC’s Core intermediate language [19]. We omit existentials, GADTs, and coercions [25], since they are largely orthogonal to join points.

Syntax System F_J is a simple λ -calculus language in the style of System F, with `let` expressions, data type constructors, and case expressions; its syntax is given in Fig. 1. System F_J is an explicitly-typed language, so all binders are typed, but in our presentation we will often drop the types.

The join-point extension is highlighted in the figure and consists of two new syntactic constructs:

- A **join** binding declares a (possibly-recursive) join point. Each join point has a name, a list of type parameters, a list of value parameters, and a body.
- A **jump** expression invokes a join point, passing all indicated arguments as well as an additional result-type argument (as discussed shortly, under “The type of a join point”).

Although we use curried syntax for **jumps**, join points are *polyadic*; partial application is not allowed.

Terms

x	\in	Term variables
j	\in	Label variables
e, u, v	$::=$	$x \mid l \mid \lambda x:\sigma.e \mid e u$
		$\Lambda a.e \mid e \varphi$ Type polymorphism
		$K \vec{\varphi} \vec{e}$ Data construction
		case e of \vec{alt} Case analysis
		let vb in v Let binding
		join jb in u Join-point binding
		jump $j \vec{\varphi} \vec{e} \tau$ Jump
alt	$::=$	$K \vec{x}:\vec{\sigma} \rightarrow u$ Case alternative

Value bindings and join-point bindings

vb	$::=$	$x:\tau = e$ Non-recursive value
		$\mathbf{rec} \vec{x}:\vec{\tau} = \vec{e}$ Recursive values
jb	$::=$	$j \vec{a} \vec{x}:\vec{\sigma} = e$ Non-recursive join point
		$\mathbf{rec} j \vec{a} \vec{x}:\vec{\sigma} = e$ Recursive join points

Answers

A	$::=$	$\lambda x:\sigma.e \mid \Lambda a.e \mid K \vec{\varphi} \vec{v}$
-----	-------	--

Types

a, b	\in	Type variables
τ, σ, φ	$::=$	a Variable
		T Datatype
		$\sigma \rightarrow \tau$ Function type
		$\tau \varphi$ Application
		$\forall a. \tau$ Polymorphic type

Frames, evaluation contexts, and stacks

F	$::=$	$\square v$ Applied function
		$\square \tau$ Instantiated polymorphism
		case \square of $\vec{p} \rightarrow \vec{u}$ Case scrutinee
		join jb in \square Join point
E	$::=$	$\square \mid F[E]$ Evaluation contexts
s	$::=$	$\varepsilon \mid F : s$ Stacks

Tail contexts

L	$::=$	\square Empty unary context
		case e of $\vec{p} \rightarrow \vec{L}$ Case branches
		let vb in L Body of let
		join $j \vec{a} \vec{x}:\vec{\sigma} = L$ in L' Join point, body
		join $\mathbf{rec} j \vec{a} \vec{x}:\vec{\sigma} = \vec{L}$ in L' Rec join points, body

Miscellaneous

C	\in	General single-hole term contexts
Σ	$::=$	$\cdot \mid \Sigma, x:\sigma = v$ Heap
c	$::=$	$\langle e; s; \Sigma \rangle$ Configuration

Figure 1: Syntax of System F_J .

Static Semantics The type system for System F_J is given in Fig. 2, where `typeof` gives the type of a constructor and `ctors` gives the set of constructors for a datatype.

The typing judgement carries two environments, Γ and Δ , with Δ binding join points. The environment Δ is extended by a **join** (rules JBIND and RJBIND) and consulted at a **jump**. Note that we rely on scoping conventions in some places: if $\Gamma; \Delta \vdash e : \tau$, then every variable (type or term) free in e or τ appears in Γ , and the symbols in Γ are unique. Similarly, every label free in e appears in Δ .

For jumps to truly be compiled as mere jumps, they must not occur in any subterm whose evaluation will take place in some unknown context. Otherwise, the jump would leave the stack in an unknown state. We enforce this invariant by resetting Δ to ε in every premise for such a subterm, for example in the premise for the body e in rule ABS.

The Type of a Join Point The type given to a join point deserves some attention. A join point that binds type variables \vec{a} and value arguments of types $\vec{\sigma}$ is given the type $\forall \vec{a}. \vec{\sigma} \rightarrow \forall r. r$ (rule JBIND). Dually, a **jump** applies a join point to some type arguments (to instantiate \vec{a}), some value arguments (to saturate the $\vec{\sigma}$), and a final type argument (to instantiate r) that specifies the type returned by the **jump**. We put the universal quantification of r at the end to indicate that the argument types $\vec{\sigma}$ do not (and must not³) mention this “return-type parameter.” Indeed, when we introduce the *abort* axiom (Sec. 3), it will need to change this type argument arbitrarily, which it can only safely do if the type is never actually used in the other parameters.

So a join point’s result type $\forall r$ does not reflect the value of its body. What then keeps a join point from returning arbitrary values? It is the JBIND rule (or its recursive variant) that checks the right-hand side of the join point, making sure it is the same as that of the entire **join** expression. Thus we cannot have

join $j = \text{"Gotcha!"}$ **in** b **then** **jump** j Int **else** 4

because j returns a *String* but the body of the **join** returns an *Int*. In short, the burden of typechecking has moved: whereas a function can be declared to return any type but can only be invoked in certain contexts, a join point can be invoked in any context but can only return a certain type.

Finally, the reader may wonder why join points are polymorphic (apart from the result type). In F_J as presented here, we could manage with monomorphic join points, but they become absolutely necessary when we add data constructors that bind existential type variables. We omitted existentials from this paper for simplicity, but they are very important in practice and GHC certainly supports them.

$$\boxed{\Gamma; \Delta \vdash e : \tau}$$

$$\frac{(x:\tau) \in \Gamma}{\Gamma; \Delta \vdash x : \tau} \text{VAR} \quad \frac{\text{typeof}(K) = \forall \vec{a}. \vec{\sigma} \rightarrow T \vec{a}}{\Gamma; \varepsilon \vdash u : \sigma\{\varphi/a\}} \text{CON} \quad \frac{\Gamma, (x:\sigma); \varepsilon \vdash e : \tau}{\Gamma; \Delta \vdash \lambda(x:\sigma).e : \sigma \rightarrow \tau} \text{ABS} \quad \frac{\Gamma, a; \varepsilon \vdash e : \tau}{\Gamma; \Delta \vdash \Lambda a.e : \forall a.\tau} \text{TABS}$$

$$\frac{\Gamma; \Delta \vdash e : \sigma \rightarrow \tau \quad \Gamma; \varepsilon \vdash u : \sigma}{\Gamma; \Delta \vdash e u : \tau} \text{APP} \quad \frac{\Gamma; \Delta \vdash e : \forall a.\tau}{\Gamma; \Delta \vdash e \varphi : \tau\{\varphi/a\}} \text{TAPP} \quad \frac{(j:\forall \vec{a}. \vec{\sigma} \rightarrow \forall r.r) \in \Delta \quad \Gamma; \varepsilon \vdash u : \sigma\{\varphi/a\}}{\Gamma; \Delta \vdash \text{jump } j \vec{\varphi} \vec{u} \tau : \tau} \text{JUMP}$$

$$\frac{\Gamma; \varepsilon \vdash u : \sigma \quad \Gamma, x:\sigma; \Delta \vdash e : \tau}{\Gamma; \Delta \vdash \text{let } x:\sigma = u \text{ in } e : \tau} \text{VBIND} \quad \frac{\Gamma, \vec{x}:\vec{\sigma}; \varepsilon \vdash u : \vec{\sigma} \quad \Gamma, \vec{x}:\vec{\sigma}; \Delta \vdash e : \tau}{\Gamma; \Delta \vdash \text{let rec } \vec{x}:\vec{\sigma} = \vec{u} \text{ in } e : \tau} \text{RVBIND}$$

$$\frac{\Gamma, \vec{a}, \vec{x}:\vec{\sigma}; \Delta \vdash u : \tau \quad \Gamma; \Delta, (j:\forall \vec{a}. \vec{\sigma} \rightarrow \forall r.r) \vdash e : \tau}{\Gamma; \Delta \vdash \text{join } j \vec{a} \vec{x}:\vec{\sigma} = u \text{ in } e : \tau} \text{JBIND}$$

$$\frac{\Gamma, \vec{a}, \vec{x}:\vec{\sigma}; \Delta, j:\forall \vec{a}. \vec{\sigma} \rightarrow \forall r.r \vdash u : \tau \quad \Gamma; \Delta, j:\forall \vec{a}. \vec{\sigma} \rightarrow \forall r.r \vdash e : \tau}{\Gamma; \Delta \vdash \text{join rec } j \vec{a} \vec{x}:\vec{\sigma} = u \text{ in } e : \tau} \text{RJBIND}$$

$$\frac{\Gamma; \Delta \vdash e : T \vec{\varphi} \quad \text{typeof}(K) = \forall \vec{a}. \vec{\sigma} \rightarrow T \vec{a} \quad \vec{v} = \vec{\sigma}\{\varphi/a\} \quad \Gamma, \vec{x}:\vec{v}; \Delta \vdash u : \tau \quad \text{ctors}(T) = \{\vec{K}\}}{\Gamma; \Delta \vdash \text{case } e \text{ of } K \vec{x}:\vec{v} \rightarrow u : \tau} \text{CASE}$$

Figure 2: Type system for System F_J .

Managing Δ The typing of join points is a little bit more flexible than you might suspect. Consider this expression:

$$\left(\begin{array}{l} \text{join } j x = \text{RHS} \\ \text{in case } v \text{ of } A \rightarrow \text{jump } j \text{ True } C2C \\ \quad B \rightarrow \text{jump } j \text{ False } C2C \\ \quad C \rightarrow \lambda c.c \end{array} \right) 'x',$$

where $C2C = \text{Char} \rightarrow \text{Char}$. This is certainly well typed. A valid transformation is to move the application to 'x' into both the body and the right-hand side of the join, thus:

$$\text{join } j x = \text{RHS } 'x' \\
\text{in } \left(\begin{array}{l} \text{case } v \text{ of } A \rightarrow \text{jump } j \text{ True } C2C \\ \quad B \rightarrow \text{jump } j \text{ False } C2C \\ \quad C \rightarrow \lambda c.c \end{array} \right) 'x',$$

Now we can move the application into the branches:

$$\text{join } j x = \text{RHS } 'x' \\
\text{in case } v \text{ of } A \rightarrow (\text{jump } j \text{ True } C2C) 'x' \\
\quad B \rightarrow (\text{jump } j \text{ False } C2C) 'x' \\
\quad C \rightarrow (\lambda c.c) 'x'$$

Should this be well typed? The jumps to j are not exactly tail calls, but they can (and indeed must) discard their context—here the application to 'x'—and resume execution at j . We will see shortly how this program can be further transformed to remove the redundant applications to 'x', but the point here is that this intermediate program should be well typed. That point is reflected in the typing rules by the fact that Δ is *not* reset in the function part of an application (rule APP), or in the scrutinee of a case (rule CASE).

$$\boxed{\langle e; s; \Sigma \rangle \mapsto \langle e'; s'; \Sigma' \rangle}$$

$$\begin{array}{ll}
\langle F[e]; s; \Sigma \rangle \mapsto \langle e; F : s; \Sigma \rangle & (\text{push}) \\
\langle \lambda x.e; \square v : s; \Sigma \rangle \mapsto \langle e; s; \Sigma, x = v \rangle & (\beta) \\
\langle \Lambda a.e; \square \varphi : s; \Sigma \rangle \mapsto \langle e\{\varphi/a\}; s; \Sigma \rangle & (\beta_\tau) \\
\langle \text{let } vb \text{ in } e; s; \Sigma \rangle \mapsto \langle e; s; \Sigma, vb \rangle & (\text{bind}) \\
\langle x; s; \Sigma[x = v] \rangle \mapsto \langle v; s; \Sigma[x = v] \rangle & (\text{look})
\end{array}$$

$$\left\langle \begin{array}{l} K \vec{\varphi} \vec{v}; \\ \text{case } \square \text{ of } \text{alt} : s; \\ \Sigma \end{array} \right\rangle \mapsto \left\langle u; s; \Sigma, \vec{x} = \vec{v} \right\rangle \quad (\text{case})$$

if $(K \vec{x} \rightarrow u) \in \text{alt}$

$$\left\langle \begin{array}{l} \text{jump } j \vec{\varphi} \vec{v} \tau; \\ s' \text{ ++ } (\text{join } jb \text{ in } \square : s); \\ \Sigma \end{array} \right\rangle \mapsto \left\langle \begin{array}{l} u\{\varphi/a\}; \\ \text{join } jb \text{ in } \square : s; \\ \Sigma, \vec{x} = \vec{v} \end{array} \right\rangle \quad (\text{jump})$$

if $(j \vec{a} \vec{x} = u) \in jb$

$$\left\langle \begin{array}{l} A; \\ \text{join } jb \text{ in } \square : s; \\ \Sigma \end{array} \right\rangle \mapsto \langle A; s; \Sigma \rangle \quad (\text{ans})$$

Figure 3: Call-by-name operational semantics for System F_J .

Operational Semantics We give System F_J an operational semantics (Fig. 3) in the style of an abstract machine. A *configuration* of the machine is a triple $\langle e; s; \Sigma \rangle$ consisting of an expression e which is the current focus of execution; a *stack* s representing the current evaluation context (including join-point bindings); and a *heap* Σ of value bindings. The stack is a list of *frames*, each of which is an argument to apply, a case analysis to perform, or a bound join point (or recursive group). Each frame is moved to the stack via the *push* rule. Since we define evaluation contexts by composing frames (hence $F[E]$ in Fig. 1), the rule has a simple form. Most of the

rules are quite conventional. We describe only call-by-name evaluation here, as rule *look* shows; switching to call-by-need by pushing an update frame is absolutely standard.

Note that only *value* bindings are put in the heap. Join points are stack-allocated in a frame: they represent mere code blocks, not first-class function closures. As expected, a jump throws away its context (the *jump* rule); it does so by popping all the frames from the stack to the binding (as usual, $++$ stands for the concatenation of two stacks):

$$\begin{aligned} & \left\langle \begin{array}{l} \mathbf{join} \ j \ x = x \\ \mathbf{in\ case} \ (\mathbf{jump} \ j \ 2 \ (Int \rightarrow \ Bool)) \ 3 \ \mathbf{of} \ \dots; \ \varepsilon; \ \varepsilon \end{array} \right\rangle \\ & \mapsto^* \left\langle \begin{array}{l} \mathbf{jump} \ j \ 2 \ (Int \rightarrow \ Bool); \\ \square \ 3 : \mathbf{case} \ \square \ \mathbf{of} \ \dots : \mathbf{join} \ j \ x = x \ \mathbf{in} \ \square : \varepsilon; \end{array} \right\rangle \\ & \mapsto \langle x; \mathbf{join} \ j \ x = x \ \mathbf{in} \ \square : \varepsilon; x = 2 \rangle \end{aligned}$$

Here three frames are pushed onto the stack: the join-point binding, the case analysis, and finally the application of the jump to 3. Then the jump is evaluated, popping the latter two frames, replacing the term with the one from the join point, and binding the argument.

The *ans* rule removes a join-point binding from the context once an answer A (see Fig. 1) is computed; note that a well-typed answer cannot contain a jump, so at that point the binding must be dead code. Continuing our example:

$$\langle x; \mathbf{join} \ j \ x = x \ \mathbf{in} \ \square : \varepsilon; x = 2 \rangle \mapsto^* \langle 2; \varepsilon; x = 2 \rangle$$

Optimizing Transformations The operational semantics operates on closed configurations. An optimizing compiler, by contrast, must transform *open* terms. To describe possible optimizations, then, we separately develop a sound *equational theory* (Fig. 4), which lays down the “rules of the game” by which the optimizer is allowed to work. It is up to the optimizer to determine how to apply the rules to rewrite code. All the axioms carry the usual implicit scoping restrictions to avoid free-variable capture. We spell out the side conditions in *drop* and *jdrops* because these actually restrict when the rule can be applied rather than merely ensuring hygiene. (In these side conditions, the *bv* function stands for “bound variables”; for instance, $\text{bv}(x:\sigma = e) = \{x\}$.)

The β , β_τ , and *case* axioms are analogues of the similarly-named rules in the operational semantics. Since there is no heap, β and *case* create **let** expressions instead. Compile-time substitution, or *inlining*, is performed for values by *inline* and for join points by *jinline*. If a binding is inlined exhaustively, it becomes dead code and can be eliminated by the *drop* or *jdrops* axiom. Values may be substituted anywhere⁴, which we indicate using a general single-hole context C in *inline*. Inlining of join points is a bit more delicate. A jump indicates *both* that we should execute the join point *and* that we should throw out the evaluation context up to the join point’s declaration. Simply copying the body accomplishes the former but not the latter. For example:

$$\mathbf{join} \ j \ (x : Int) = x + 1 \ \mathbf{in} \ (\mathbf{jump} \ j \ 2 \ (Int \rightarrow Int)) \ 3$$

⁴For brevity, we have omitted rules allowing inlining a recursive definition into the definition itself (or another definition in the same recursive group).

If we naïvely inline j here, we end up with the ill-typed term:

$$\mathbf{join} \ j \ (x : Int) = x + 1 \ \mathbf{in} \ (2 + 1) \ 3$$

Inlining is safe, however, if the jump is a tail call, since then there is no extra evaluation context to throw away. To specify the allowable places to inline a join point, then, we use a syntactic notion called a *tail context*. A tail context L (see Fig. 1) is a multi-hole context describing the places where a term may return to its evaluation context. Since $\square \ 3$ is not a tail context, the *jinline* axiom fails for the above term.

The *casefloat*, *float*, *jfloat*, and *jfloat_{rec}* axioms perform commuting conversions. The former two are conventional, but *jfloat* and *jfloat_{rec}* exploit the new join-point construct to perform exactly the transformation we needed in Sec. 2 to avoid destroying a join point. The only difference between the two is that *jfloat_{rec}* acts on a recursive binding; the operation performed is the same.

Consider again the example at the beginning of Sec. 2. With our new syntax, we can write it as:

$$\mathbf{case} \ \left(\begin{array}{l} \mathbf{join} \ j \ x = \text{BIG} \\ \mathbf{in\ case} \ v \ \mathbf{of} \ A \rightarrow \mathbf{jump} \ j \ 1 \ \text{Bool} \\ \quad \quad \quad B \rightarrow \mathbf{jump} \ j \ 2 \ \text{Bool} \\ \quad \quad \quad C \rightarrow \text{True} \end{array} \right) \ \mathbf{of} \\ \{ \text{True} \rightarrow \text{False}; \text{False} \rightarrow \text{True} \}$$

We can use *jfloat* to move the outer **case** both into the right-hand side of the **join** binding and into its body; use *casefloat* to move the outer **case** into the branches of the inner **case**; use *abort* to discard the outer **case** where it scrutinizes a **jump**; and use *case* to simplify the C alternative. The result is just what we want:

$$\begin{aligned} & \mathbf{join} \ j \ x = \mathbf{case} \ \text{BIG} \ \mathbf{of} \ \{ \text{True} \rightarrow \text{False}; \text{False} \rightarrow \text{True} \} \\ & \mathbf{in\ case} \ v \ \mathbf{of} \ A \rightarrow \mathbf{jump} \ j \ 1 \ \text{Bool} \\ & \quad \quad \quad B \rightarrow \mathbf{jump} \ j \ 2 \ \text{Bool} \\ & \quad \quad \quad C \rightarrow \text{False} \end{aligned}$$

The commute Axiom The left-hand sides of axioms *float*, *jfloat*, *jfloat_{rec}*, and *casefloat* enumerate the forms of a tail context L (Figure 1). So the four axioms are all instances of a single equivalent form:

$$E[L[\vec{e}]] = L[\overrightarrow{E[e]}] \quad (\text{commute})$$

This rule *commute* moves the evaluation context E into *each* hole of the tail context L .⁵

We can also derive new axioms succinctly using tail contexts. For example, our commuting conversions as written risk quite a bit of code duplication by copying E arbitrarily many times (into each branch of a **case** and each join point). Of course, in a real implementation, we would prefer not to do this, so instead we might use a different axiom:

$$E[L[\vec{e}] : \tau] = \mathbf{join} \ j \ x = E[x] \ \mathbf{in} \ L[\overrightarrow{\mathbf{jump} \ j \ e \ \tau}]$$

⁵Since L has a hole wherever something is returned to E , *commute* “substitutes” the latter into the places where it is invoked. In fact, from a CPS standpoint, *commute* (in concert with *abort*) is *precisely* a substitution operation.

$$\boxed{e = e'}$$

$(\lambda x:\sigma.e) v$	$= \text{let } x:\sigma = v \text{ in } e$	(β)
$(\Lambda a.e) \varphi$	$= e\{\varphi/a\}$	(β_τ)
$\text{let } vb \text{ in } C[x]$	$= \text{let } vb \text{ in } C[v]$	if $(x:\sigma = v) \in vb$ $(inline)$
$\text{let } vb \text{ in } e$	$= e$	if $\text{bv}(vb) \cap \text{fv}(e) = \emptyset$ $(drop)$
$\text{join } jb \text{ in } L[\vec{e}, \text{jump } j \vec{\varphi} \vec{\tau}, \vec{e}']$	$= \text{join } jb \text{ in } L[\vec{e}, \text{let } \vec{x}:\vec{\sigma} = \vec{v} \text{ in } u\{\vec{\varphi}/\vec{a}\}, \vec{e}']$	if $(j \vec{a} \vec{x}:\vec{\sigma} = u) \in jb$ $(jinline)$
$\text{join } jb \text{ in } e$	$= e$	if $\text{bv}(jb) \cap \text{fv}(e) = \emptyset$ $(jdrop)$
$\text{case } K \vec{\varphi} \vec{v} \text{ of } \vec{alt}$	$= \text{let } \vec{x}:\vec{\sigma} = \vec{v} \text{ in } e$	if $(K \vec{x}:\vec{\sigma} \rightarrow e) \in \vec{alt}$ $(case)$
$E[\text{case } e \text{ of } K \vec{x} \rightarrow u]$	$= \text{case } e \text{ of } K \vec{x} \rightarrow E[u]$	$(casefloat)$
$E[\text{let } vb \text{ in } e]$	$= \text{let } vb \text{ in } E[e]$	$(float)$
$E[\text{join } j \vec{a} \vec{x} = u \text{ in } e]$	$= \text{join } j \vec{a} \vec{x} = E[u] \text{ in } E[e]$	$(jfloat)$
$E[\text{join rec } j \vec{a} \vec{x} = u \text{ in } e]$	$= \text{join rec } j \vec{a} \vec{x} = E[u] \text{ in } E[e]$	$(jfloat_{rec})$
$E[\text{jump } j \vec{\varphi} \vec{e} \tau] : \tau'$	$= \text{jump } j \vec{\varphi} \vec{e} \tau'$	$(abort)$

Figure 4: Common optimizations for System F_J.

This can be derived from *commute* by first applying *jdrop* and *jinline* backward.

4. Contification: Inferring Join Points

Not all join points originate from commuting conversions. Though the source language doesn't have join points or jumps, many let-bound functions can be converted to join points without changing the meaning of the program. In particular, if *every* call to a given function is a saturated tail call (i.e. appears only in an *L*-context), and we turn the calls into jumps, then whenever one of the jumps is executed, there will be nothing to drop from the evaluation context (the *s'* in *jump* will be empty).

The process is a form of *contification* [16] (or *continuation demotion*), which we formalize in Fig. 5, where $\text{fv}(e)$ means the set of free variables of e (and similarly $\text{fv}(L)$ for tail contexts), and $\text{dom}(\rho)$ means the domain of the environment ρ (to be described shortly).

The non-recursive version, *contify*, attempts to decompose the body of the let (i.e., the scope of f) into a tail context L and its arguments, where the arguments contain all the occurrences of f , then attempts to run the special partial function tail on each argument to the tail context. This function will only succeed if there are no non-tail calls to f .

The tail function takes an environment ρ mapping applications of contifiable variables f to jumps to corresponding join points j . For each expression that matches the form of a saturated call to such an f , then, tail turns the call into a jump to its j , *provided* that none of the arguments to the function contains a free occurrence of a variable being contified—an occurrence in argument position is disallowed by the typing rules. For any other expression, tail changes nothing but does check that no variable being contified appears; otherwise, tail fails, causing the *contify* axiom not to match.

There is one last proviso in the *contify* and *contify_{rec}* axioms, which is that the body of each function to be contified must have the same type as the body of the let. This can fail if some function f is polymorphic in its return type [8].

Finding bindings to which *contify* or *contify_{rec}* will apply is not difficult. Our implementation is essentially a free-variable analysis that also tracks whether each free variable has appeared *only* in the holes of tail contexts. This is much simpler than previous contification algorithms because we *only look for tail calls*. We invite the reader to compare to [11] or to Sec. 5 of [16], which both allow for more general calls to be dealt with. Yet we claim that, in concert with the Simplifier and the Float In pass, our algorithm covers most of the same ground.

To demonstrate, consider the *local CPS transformation* in Moby [23], which produces mutually tail-recursive functions to improve code generation in much the same way GHC does. Moby uses a direct-style intermediate representation, but its contification pass is expressed in terms of a CPS transform, which turns

$$\text{let } f \ x = \dots \text{ in } E[\dots f \ y \ \dots f \ z \ \dots]$$

(where the calls to f are tail calls within E) into

$$\text{let } \{ j \ x = E[x]; f \ x = j \ \langle \text{rhs} \rangle \}$$

$$\text{in } \dots f \ y \ \dots f \ z \ \dots$$

where the tail calls to f are now compiled as jumps. Note that f now matches the *contify* axiom, but it did not before due to the E in the way. Nonetheless, our extended GHC achieves the same effect, only in stages. Starting with:

$$\text{let } f \ x = \text{rhs in } E[\dots f \ y \ \dots f \ z \ \dots]$$

First, applying *float* from right to left floats f inward:

$$E[\text{let } f \ x = \text{rhs in } \dots f \ y \ \dots f \ z \ \dots]$$

Next, *contify* applies, since the calls to f are now tail calls:

$$E[\text{join } f \ x = \text{rhs in } \dots \text{jump } f \ y \ \tau \ \dots \text{jump } f \ z \ \tau \ \dots]$$

$$\begin{array}{l}
\boxed{e = e'} \\
\text{let } f = \Lambda \vec{a}. \lambda \vec{x}. u \text{ in } L[\vec{e}] : \tau = \text{join } j \vec{a} \vec{x} = u \text{ in } L[\overrightarrow{\text{tail}_\rho(e)}] \quad (\text{contify}) \\
\quad \text{if } \rho(f \vec{a} \vec{x}) = \text{jump } j \vec{a} \vec{x} \tau \\
\quad \text{and } f \notin \text{fv}(L), u : \tau \\
\text{let rec } \overrightarrow{f = \Lambda \vec{a}. \lambda \vec{x}. L[\vec{u}]} \text{ in } L'[\vec{e}] : \tau = \text{join rec } j \vec{a} \vec{x} = L[\overrightarrow{\text{tail}_\rho(u)}] \text{ in } L'[\overrightarrow{\text{tail}_\rho(e)}] \quad (\text{contify}_{\text{rec}}) \\
\quad \text{if } \rho(f \vec{a} \vec{x}) = \text{jump } j \vec{a} \vec{x} \tau \\
\quad \text{and } f \notin \text{fv}(\vec{L}), f \notin \text{fv}(L'), L[\vec{u}] : \tau \\
\text{tail}_\rho(f \vec{\sigma} \vec{u}) \triangleq e\{\overrightarrow{\sigma/a}\}\{\overrightarrow{u/x}\} \quad \text{if } \rho(f \vec{a} \vec{x}) = e \text{ and } \text{dom}(\rho) \cap \text{fv}(\vec{u}) = \emptyset \\
\text{tail}_\rho(e) \triangleq e \quad \text{if } \text{dom}(\rho) \cap \text{fv}(e) = \emptyset \\
\text{tail}_\rho(e) \triangleq \text{undefined} \quad \text{otherwise}
\end{array}$$

Figure 5: Contification as a source-to-source transformation.

And now *jfloat* pushes E into the join point f and the body:

$\text{join } f x = E[\text{rhs}] \text{ in } \dots E[\text{jump } f y \tau] \dots E[\text{jump } f z \tau] \dots$

From here, *abort* removes E from the jumps, and we can abstract E by running *jdrops* and *jinlines* backward:

$\text{join } \{j x = E[x]; f x = \text{jump } j \text{ rhs } \tau\} \text{ in } \dots f y \dots f z \dots$

Thus we achieve the same result without any extra effort⁶.

Naturally, contification is more routine and convenient in CPS-based compilers [11, 16]. The ability to handle an intervening context comes nearly “for free” since contexts already have names. Notably, it is still possible to name contexts in direct style (the Moby paper [23] does so using labelled expressions), so it is *only* a matter of convenience.

5. Recursive Join Points and Fusion

We have mentioned, without stressing the point, that join points can be recursive. We have also shown that it is rather easy to identify let-bindings that can be re-expressed (more efficiently) as join points. To our complete surprise, we discovered that the combination of these two features allowed us to solve a long-standing problem with stream fusion.

Recursive Join Points Consider this program, which finds the first element of a list that satisfies a predicate p :

$$\begin{array}{l}
\text{find} = \Lambda a. \lambda (p : a \rightarrow \text{Bool})(xs : [a]). \\
\text{let } go \text{ xs} = \text{case } xs \text{ of} \\
\quad x : xs' \rightarrow \text{if } p x \text{ then } Just x \\
\quad \quad \quad \text{else } go \text{ xs}' \\
\quad [] \rightarrow \text{Nothing} \\
\text{in } go \text{ xs}_0
\end{array}$$

Programmers quite often write loops like this, with a local definition for go , perhaps to allow $find$ to be inlined at a call site. Our first observation is this: go is a (recursive) join point!

⁶The parts of this sequence not specifically to do with join points were already implemented before in GHC: The Float In pass applies *float* in reverse, and the Simplifier regularly creates join points to share evaluation contexts (except that previously they were ordinary let bindings).

The contification transformation will identify go as a join point, and will transform the **let** (which allocates) to a **join** (which does not), and each call to go into an efficient **jump**.

But it gets better! Because go is a join point, it can participate in a commuting conversion. Suppose, for example, that $find$ is called from any like this:

$$\begin{array}{l}
any = \Lambda a. \lambda (p : a \rightarrow \text{Bool})(xs : [a]). \\
\quad \text{case } find \text{ p } xs \text{ of } Just _ \rightarrow \text{True} \\
\quad \quad \quad \text{Nothing} \rightarrow \text{False}
\end{array}$$

The call to $find$ can be inlined:

$$\begin{array}{l}
any = \Lambda a. \lambda (p : a \rightarrow \text{Bool})(xs : [a]). \\
\quad \text{case } \left(\begin{array}{l} \text{join } go \text{ xs} = \text{case } xs \text{ of} \\ \quad x : xs' \rightarrow \text{if } p x \text{ then } Just x \\ \quad \quad \quad \text{else } \text{jump } go \text{ xs}' (Maybe a) \\ \quad [] \rightarrow \text{Nothing} \\ \quad \text{in } \text{jump } go \text{ xs} (Maybe a) \end{array} \right) \text{ of} \\
\quad \{Just _ \rightarrow \text{True}; \text{Nothing} \rightarrow \text{False}\}
\end{array}$$

Now, we have a **case** scrutinizing a **join** so we can apply axiom *jfloat* from Figure 4. After some easy further transformations, we get

$$\begin{array}{l}
any = \Lambda a. \lambda (p : a \rightarrow \text{Bool})(xs : [a]). \\
\quad \text{join } go \text{ xs} = \text{case } xs \text{ of} \\
\quad \quad x : xs' \rightarrow \text{if } p x \text{ then } \text{True} \\
\quad \quad \quad \text{else } \text{jump } go \text{ xs}' \text{ Bool} \\
\quad \quad [] \rightarrow \text{False} \\
\quad \text{in } \text{jump } go \text{ xs} \text{ Bool}
\end{array}$$

Look carefully at what has happened here: the consumer (any) of a recursive loop (go) has moved *all the way to the return point of the loop*, so that we were able to cancel the **case** in the consumer with the data constructor returned at the conclusion of the loop.

Stream Fusion It turns out that this new ability to move a consumer all the way to the return points of a tail-recursive loop has direct implications for a very widely used transformation: stream fusion. The key idea of stream fusion is to represent a list (or array, or other sequence) by a pair of a *state* and a *stepper function*, thus:⁷


```
data Stream a where
  MkStream :: s -> (s -> Step s a) -> Stream a
```

There are two competing approaches to the `Step` type. In `unfold/destroy fusion` (Svenningsson [26]), we have:

```
data Step s a = Done | Yield s a
```

Hence a stepper function takes an incoming state and either yields an element and a new state or signals the end. Now a pipeline of list processors can be rewritten as a pipeline of stepper functions, each of which produces and consumes elements one by one. A typical stepper function for a stream transformer looks like:

```
next s = case <incoming step> of
  Yield s' a -> <process element>
  Done       -> <process end of stream>
```

When composed together and inlined, the stepper functions become a nest of cases, each scrutinizing the output of the previous stepper. It is crucial for performance that each `Yield` or `Done` expression be matched to a case, much as we did with `Just` and `Nothing` in the example that began Sec. 2. Fortunately, case-of-case and the other commuting conversions that GHC performs are usually up to the task.

Alas, this approach requires a recursive stepper function when implementing `filter`, which must loop over incoming elements until it finds a match. This breaks up the chain of cases by putting a loop in the way, much as our `any` above becomes a `case` on a loop. Hence until now, recursive stepper functions have been un-fusible. Coutts *et al.* [6] suggested adding a `Skip` constructor to `Step`, thus:

```
data Step s a = Done | Yield s a | Skip s
```

Now the stepper function can say to update the state and call again, obviating the need for a loop of its own. This makes `filter` fusible, but it complicates everything else! Everything gets three cases instead of two, leading to more code and more runtime tests; and functions like `zip` that consume two lists become more complicated and less efficient.

But with join points, just as with `any`, Svenningsson’s original `Skip`-less approach fuses just fine! Result: simpler code, less of it, and faster to execute. It’s a straight win.

6. Metatheory of F_J

Proofs can be found in the extended version of this paper⁸.

Correctness and Type Safety The way to “run” a program on our abstract machine is to initialize the machine with an empty stack and an empty store. Type safety, then, says that once we start the machine, the program either runs forever or successfully returns an answer.

Proposition 1 (Type safety). *If $\varepsilon; \varepsilon \vdash e : \tau$, then either:*

⁷Note that `Stream` is an existential type, so as to abstract the internal state type `s` as an implementation detail of the stream.

⁸<https://www.microsoft.com/en-us/research/publication/compiling-without-continuations/>

1. The initial configuration $\langle e; \varepsilon; \varepsilon \rangle$ diverges, or
2. $\langle e; \varepsilon; \varepsilon \rangle \mapsto^* \langle A; \varepsilon; \Sigma \rangle$, for some store Σ and answer A .

To establish the correctness of our rewriting axioms, we first define a notion of *observational equivalence*.

Definition 2. *Two terms e and e' are observationally equivalent, written $e \cong e'$, if, given any context C , $\langle C[e]; \varepsilon; \varepsilon \rangle$ diverges if and only if $\langle C[e']; \varepsilon; \varepsilon \rangle$ diverges.*

The equational theory is sound with respect to \cong :

Proposition 3. *If $e = e'$, then $e \cong e'$.*

Equivalence to System F The best way to be sure that F_J can be implemented without headaches is to show that it is equivalent to GHC’s existing System F-based language. This would suggest that join points do not allow us to write any *new* programs, only to implement existing programs more efficiently. To prove the equivalence, we establish an *erasure* procedure that removes all join points from an F_J term, leaving an equivalent System F term.

To erase the join points, we want to apply the *contify* axiom (or its recursive variant) from right to left. However, we cannot necessarily do so immediately for each join point, since *contify* only applies when all invocations j here are in tail position. For example, we cannot de-contify j here:

$$\mathbf{join} \ j \ x = x + 1 \ \mathbf{in} \ (\mathbf{jump} \ j \ 1 \ (Int \rightarrow Int)) \ 2$$

Simply rewriting the join point as a function and the jump as a function call would change the meaning of the program—in fact, it would not even be well-typed:

$$\mathbf{let} \ f = \lambda x. x + 1 \ \mathbf{in} \ f \ 1 \ 2$$

However, if we apply *abort* first:

$$\mathbf{join} \ j \ x = x + 1 \ \mathbf{in} \ \mathbf{jump} \ j \ 1 \ Int$$

Now the jump is a tail call, so *contify* applies.

The *abort* axiom is not enough on its own, since the jump may be buried inside a tail context:

$$\mathbf{join} \ j \ x = x + 1 \ \mathbf{in} \ \left(\begin{array}{l} \mathbf{case} \ b \ \mathbf{of} \\ \quad True \rightarrow \mathbf{jump} \ j \ 1 \ (Int \rightarrow Int) \\ \quad False \rightarrow \mathbf{jump} \ j \ 3 \ (Int \rightarrow Int) \end{array} \right) \ 2$$

However, this can be handled by a commuting conversion:

$$\mathbf{join} \ j \ x = x + 1 \ \mathbf{in} \ \mathbf{case} \ b \ \mathbf{of} \\ \quad True \rightarrow (\mathbf{jump} \ j \ 1 \ (Int \rightarrow Int)) \ 2 \\ \quad False \rightarrow (\mathbf{jump} \ j \ 3 \ (Int \rightarrow Int)) \ 2$$

And now *abort* applies twice and j can be de-contified.

Lemma 4. *For any well-typed term e , there is an e' such that $e' = e$ and every jump in e' is in tail position.*

By “tail position,” we mean one of the holes in a tail context that starts with the binding for the join point being called. In other words, given a term

$$\mathbf{join} \ j \ \vec{a} \ \vec{x} = u \ \mathbf{in} \ L[\vec{e}],$$

the terms \vec{e} are in tail position for j .

The proof of Lemma 4 relies on the observation that the places in a term that may contain free occurrences of labels are precisely those appearing in the hole of either an evaluation or a tail context. For example, the CASE typing rule propagates Δ into both the scrutinee and the branches; note that $\text{case } \square \text{ of } \vec{alt}$ is an evaluation context and $\text{case } e \text{ of } p \rightarrow \square$ is a tail context. But $e \square$ is (in call-by-name) neither an evaluation context *nor* a tail context, and APP does *not* propagate Δ into the argument.

Thus *any* expression can be written as:

$$L[E[L'[E'[\dots[L^{(n)}[E^{(n)}[e]]]\dots]]], \quad (1)$$

which is to say a tree of tail contexts alternating with evaluation contexts, where all free occurrences of join points are at the leaves. By iterating *commute* and *abort*, we can flatten the tree, rewriting (1) to say that any expression can be written $L[\vec{e}]$, where each e_i is a leaf from the tree in (1). In this form, no e_i can be expressed as $E[L[\dots]]$ for non-trivial, non-binding⁹ E and nontrivial L , and every jump to a free occurrence of a label is some e_i . Let us say a term in the above form is in *commuting-normal form*. (Note that ANF is simply commuting-normal form with named intermediate values.) By *commute* and *abort*, every term has a commuting-normal form, and by construction, every jump in a commuting-normal form is a tail call. Thus *every label can be decontified*, and we have:

Theorem 5 (Erasure). *For any closed, well-typed F_J term e , there is a System F term e' such that $e' = e$.*

7. Join Points in Practice

It is one thing to define a calculus, but quite another to use it in a full-scale optimising compiler. In this section we report on our experience of doing so in GHC.

Implementing Join Points in GHC We have implemented System F_J as an extension to the Core language in GHC. As a representation choice, instead of adding two new data constructors for **join** and **jump** to the Core data type, we instead re-use ordinary **let**-bindings and function applications, distinguishing join points only by a flag *on the identifier itself*.

Thus, with no code changes, GHC treats join-point identifiers identically to other identifiers, and join-point bindings identically to ordinary **let** bindings. This is extremely convenient in practice. For example, all the code that deals with dropping dead bindings, inlining a binding that occurs just once, inlining a binding whose right-hand side is small, and so on, all works automatically for join points too.

With the modified Core language in hand, we had three tasks. First, GHC has an internal typechecker, called Core Lint, that (optionally) checks the type-correctness of the

⁹ A **join** can be treated as either an evaluation context or a tail context; using *commute* to push a **join** inward is not necessarily helpful, however.

intermediate program after each pass. We augmented Core Lint for F_J according to the rules of Fig. 2.

Second, we added a simple new contification analysis to identify **let**-bindings that can be converted into join points (see Sec. 4). Since the analysis is simple, we run it frequently, whenever the so-called occurrence analyzer runs.

Finally, the new Core Lint forensically identified several existing Core-to-Core passes that were “destroying” join points (see Sec. 2). Destroying a join point de-optimizes the program, so it is wonderful now to have a way to nail such problems at their source. Moreover, once Core Lint flagged a problem, it was never difficult to alter the Core-to-Core transformation to make it preserve join points. Here are some of the specifics about particular passes:

The Simplifier is a sort of partial evaluator responsible for many local transformations, including commuting conversions and inlining [19]. The Simplifier is implemented as a tail-recursive traversal that builds up a representation of the evaluation context as it goes; as such, implementing the *jfloat* and *abort* axioms (Sec. 3) requires only two new behaviors:

- (*jfloat*) When traversing a join-point binding, copy the evaluation context into the right-hand side.
- (*abort*) When traversing a jump, throw away the evaluation context.

The Float Out pass moves **let** bindings outwards [20]. Moving a **join** binding outwards, however, risks destroying the join point, so we modified Float Out to leave **join** bindings alone in most cases.

The Float In pass moves **let** bindings inwards. It too can destroy join points by un-saturating them. For example, given `let j x y = ... in j 1 2`, the Float In pass wants to narrow j 's scope as much as possible: `(let j x y = ... in j) 1 2`. We modified Float In so that it never un-saturates a join point.

Strictness analysis is as useful for join points as it is for ordinary **let** bindings, so it is convenient that **join** bindings are, by default, treated identically to ordinary **let** bindings. In GHC, the results of strictness analysis are exploited by the so-called worker/wrapper transform [12, 19]. We needed to modify this transform so that the generated worker and wrapper are both join points. We found that GHC's *constructed product result* (CPR) analysis [3] caused the wrapper to invoke the worker inside a **case** expression, thus preventing the worker from being a join point. We simply disable CPR analysis for join points; it turns out that the commuting conversions for join points do a better job anyway.

Benchmarks The reason for adding join points is to improve performance; expressiveness is unchanged (Sec. 6). So does performance improve? Table 1 presents benchmark data on allocations, collected from the standard `spectral`, `real`

spectral		real	
Program	Allocs	Program	Allocs
fibheaps	-1.1%	anna	+0.5%
ida	-1.4%	cacheprof	-0.5%
nucleic2	+0.2%	fem	+3.6%
para	-4.3%	gamteb	-1.4%
primetest	-3.6%	hpg	-2.1%
simple	-0.9%	parser	+1.2%
solid	-8.4%	rsa	-4.7%
sphere	-3.3%	(18 others)	
transform	+1.1%	Min	-4.7%
(45 others)		Max	+3.6%
Min	-8.4%	Geo. Mean	-0.2%
Max	+1.1%		
Geo. Mean	-0.4%		

shootout	
Program	Allocs
k-nucleotide	-85.9%
n-body	-100.0%
spectral-norm	-0.8%
(5 others)	
Min	-100.0%
Max	+0.0%
Geo. Mean	n/a

Table 1: Benchmarks from the `spectral`, `real`, and `shootout` NoFib suites.

and `shootout` NoFib benchmark suites¹⁰. We ran the tests on our modified GHC branch, and compared them to the GHC baseline to which our modifications were applied. Remember, the baseline compiler *already* recognises join points in the back end and compiles them efficiently (Sec. 2); the performance changes here come from preserving and exploiting join points during optimization.

We report only heap allocations because they are a repeatable proxy for runtime; the latter is much harder to measure reliably. All tests omitted from the tables had an improvement in allocations, but less than 0.3%.

There are some startling figures: using join points eliminated *all* allocations in `n-body` and 85.9% in `k-nucleotide`. We caution that these are highly atypical programs, already hand-crafted to run fast. Still, it seems that our work may make it easier for performance-hungry authors to squeeze more performance out of their inner loops.

The complex interaction between inlining and other transformations makes it impossible to *guarantee* improvements. For example, improving a function f might make it small enough to inline into g , but this may cause g to become too *large* to inline elsewhere, and that in turn may lose the optimization opportunities previously exposed by inlining g . GHC’s approach is heuristic, aiming to make losses unlikely,

¹⁰ The `imaginary` suite had no interesting cases. We believe this is because join points tend to show up only in fairly large functions, and the `imaginary` tests are all micro-benchmarks.

but they do occur, including a 1.1% increase in allocations in `spectral/transform` and a 3.6% increase in `real/fem`.

Beyond Benchmarks These benchmarks show modest but fairly consistent improvements for existing, unmodified programs. But we believe that the systematic addition of join points may have a more significant effect on programming patterns. Our discussion of fusion in Sec. 5 is a case in point: with join points we can use skip-less `unfoldr/destroy` streams without sacrificing fusion. That knowledge in turn affects the way libraries are written: they can be smaller and faster.

Moreover, the transformation pipeline becomes more robust. In GHC today, if a “join point” is inlined we get good fusion behavior, but if its size grows to exceed the (arbitrary) inlining threshold, suddenly behavior becomes much worse. An innocuous change in the source program can lead to a big change in execution time. That step-change problem disappears when we formally add join points.

8. Why Not Use Continuation-Passing Style?

Our join points are, of course, nothing more than continuations, albeit second-class continuations that do not escape, and thus can be implemented efficiently. So why not just use CPS? Kennedy’s work makes a convincing argument for CPS as a language in which to perform optimization [16].

There are many similarities between Kennedy’s work and ours. Notably, Kennedy distinguishes ordinary bindings (`let`) from continuation bindings (`letcont`), just as we distinguish ordinary bindings from join points (`join`); similarly, he distinguishes continuation invocations (i.e. jumps) from ordinary function calls, and we follow suit. But there are a number of reasons to prefer direct style, if possible:

- Direct style is, well, more direct. Programs are simply easier to understand, and the compiler’s optimizations are easier to follow. Although it sounds superficial, in practice it is a significant advantage of direct style; for example Haskell programmers often pore over the GHC’s Core dumps of their programs.
- The translation into CPS encodes a particular order of evaluation, whereas direct style does not. That dramatically inhibits code-motion transformations. For example, GHC does a great deal of “`let floating`” [20], in which a `let` binding is floated outwards or inwards, which is valid for pure (effect-free) bindings. This becomes harder or impossible in CPS, where the order of evaluation is prescribed.
- Fixing the order of evaluation is a particular issue when compiling a call-by-need language, since the known call-by-need CPS transform [18] is quite involved.
- Some transformations are much harder in CPS. For example, consider common sub-expression elimination (CSE). In $f (g\ x) (g\ x)$, the common sub-expression is easy to see. But it is much harder to find in the CPS version:

```
letcont k1 xv = letcont k2 yv = f k xv yv
                in g k2 x
in g k1 x
```

- GHC makes extensive use of user-written rewrite rules as optimizing transformations [22]. For example, stream fusion relies on the following rule, which states that turning a stream into a list and back does nothing [6]:

```
{-# RULES "stream/unstream"
   forall s. stream (unstream s) = s #-}
```

In CPS, these nested function applications are more difficult to spot. Also, rule matching is simply easier to reason about when the rules are written in more-or-less the same syntax as the intermediate language, but CPS makes radical changes compared to the source language.

9. Related Work

Join Points and Commuting Conversions Join points have been around for a long time in practice [27], but they have lacked a formal treatment until now. By introducing join points at the level at which common optimizations are applied, we’re able to exploit them more fully. For example, stream fusion as discussed in Sec. 5 depends on several algorithms working in concert, including commuting conversions, inlining, user-specified rewrite rules [22], and call-pattern specialization [21].

Fluet and Weeks [11] describe MLton’s intermediate language, whose syntax is much like ours (only first-order). However, it requires that non-tail calls be written so as to pass the result to a named continuation (what we would call a join point). As the authors note, however, this is only a minor syntactic change from passing the continuation as a parameter, and so the language has more in common with CPS than with direct style.

Commuting conversions are also discussed by Benton *et al.* in a call-by-value setting [4]. Consider:

```
let z = let y = case a of { A -> e1; B -> e2 }
        in e3
in e4
```

They show how to apply commuting conversions from the inside outward, creating functions to share code, getting:

```
let z = let j2 y = e3
        in case a of { A -> j2 e1; B -> j2 e2 }
in e4
```

and then:

```
let { j1 z = e4; j2 y = e3 }
in case a of { A -> j1 (j2 e1); B -> j1 (j2 e2) }
```

They call j_1 a “useless function”: it is only applied to the result of j_2 . It would be better to combine j_1 with j_2 to save a function call. Their solution is to be careful about the order of commuting conversions, since the problem does not occur if one goes from the outside inward instead. However, with join points, the order does not matter! If we make j_2 a join point, then the second step is instead

```
join j2 y = let z = e3 in e4
in case a of { A -> j2 e1; B -> j2 e2 }
```

which is the same result one gets starting from the outside. So our approach is more robust to the order in which transformations are applied.

SSA The majority of current commercial and open-source compilers (including GCC, LLVM, Mozilla JavaScript) and compiler frameworks use the Static Single Assignment (SSA) form [7], which for an assembly-like language means that variables are assigned only once. If a variable might have different values, it is defined by a ϕ -node, which chooses a value depending on control flow. This makes data flow explicit, which helps to simplify some optimizations.

As it happens, SSA is inter-derivable with CPS [2] or ANF [5]. Code blocks in SSA become mutually-recursive continuations in CPS or functions in ANF, and ϕ -nodes indicate the parameters at the different call sites. In fact, in ANF, the functions representing blocks are always tail-called, so adding join points to ANF gives a closer correspondence to SSA code—functions correspond to functions and join points correspond to blocks. Indeed the Swift Intermediate Language SIL appears to have adopted the idea of “basic blocks with arguments” instead of ϕ -nodes [14].

Sequent Calculus Our previous work [8] showed how to define an intermediate language, called Sequent Core, which sits in between direct style and CPS. Sequent Core disentangles the concepts of “context” and “evaluation order”—contexts are invaluable, but Haskell has no fixed evaluation order, a fact which GHC exploits ruthlessly. The inspiration for our language’s design came from logic, namely the *sequent calculus*. The sequent calculus is the twin brother of *natural deduction*, the foundation of direct-style representations. In this paper, we use Sequent Core as our inspiration much as Flanagan *et al.* [10] used CPS, with a new motto: *Think in sequent calculus; work in λ -calculus*.

Relation to a Language with Control Since F_j has a notion of control, it becomes natural to relate it to known control theories such as the one developed to reason about `callcc` in Scheme [9]. In fact, our language can encode `callcc v` as `join j x = x in $\llbracket v \rrbracket$ ($\lambda y. \mathbf{jump} j y$)`. By design, this encoding does not type in our system since the continuation variable j is free in a lambda-abstraction. This has repercussions on the semantics: join points can no longer be saved in the stack but need to be stored in the heap, which is precisely what is needed to implement `callcc`.

10. Reflections

Based on our experience in a mature compiler for a statically-typed functional language, the use of F_j as an intermediate language seems very attractive. Compared to the baseline of System F, F_j is a rather small change; other transformations are barely affected; the new commuting conversions are valu-

able in practice; and they make the transformation pipeline more robust.

Although we have presented F_J as a lazy language, *everything in this paper applies equally to a call-by-value language*. All one needs to do is to change the evaluation context, the notion of what is substitutable, and a few typing rules (as described in Sec. 6).

Acknowledgments

We gratefully acknowledge helpful feedback from Arjan Boeijink, Joachim Breitner, Gabor Greif, Andrew Kennedy, Oleg Kiselyov, Simon Marlow, and Carter Schonwald. This material is based upon work supported by the National Science Foundation under Grant No. CCF-1423617.

References

- [1] A. W. Appel. *Compiling with Continuations*. Cambridge University Press, 1992. ISBN 0-521-41695-7.
- [2] A. W. Appel. SSA is functional programming. *SIGPLAN Notices*, 33(4):17–20, 1998.
- [3] C. A. Baker-Finch, K. Glynn, and S. L. Peyton Jones. Constructed product result analysis for Haskell. *J. Funct. Program.*, 14(2):211–245, 2004.
- [4] N. Benton, A. Kennedy, S. Lindley, and C. V. Russo. Shrinking reductions in SML.NET. In *Implementation and Application of Functional Languages, 16th International Workshop, IFL 2004, Lübeck, Germany, September 8-10, 2004, Revised Selected Papers*, pages 142–159, 2004.
- [5] M. M. T. Chakravarty, G. Keller, and P. Zadarnowski. A functional perspective on SSA optimisation algorithms. *Electr. Notes Theor. Comput. Sci.*, 82(2):347–361, 2003.
- [6] D. Coutts, R. Leshchinskiy, and D. Stewart. Stream fusion: from lists to streams to nothing at all. In *ACM SIGPLAN International Conference on Functional Programming (ICFP’07)*, Freiburg, Germany, Oct. 2007. ACM.
- [7] R. Cytron, J. Ferrante, B. K. Rosen, M. N. Wegman, and F. K. Zadeck. Efficiently computing static single assignment form and the control dependence graph. *ACM Transactions on Programming Languages and Systems*, 13(4):451–490, 1991.
- [8] P. Downen, L. Maurer, Z. M. Ariola, and S. Peyton Jones. Sequent calculus as a compiler intermediate language. In *Proceedings of the 21st ACM SIGPLAN International Conference on Functional Programming, ICFP 2016, Nara, Japan, September 18-22, 2016*, pages 74–88, 2016.
- [9] M. Felleisen and R. Hieb. A revised report on the syntactic theories of sequential control and state. *Theoretical Computer Science*, 103(2):235–271, 1992.
- [10] C. Flanagan, A. Sabry, B. F. Duba, and M. Felleisen. The essence of compiling with continuations. In *Proceedings of the ACM SIGPLAN’93 Conference on Programming Language Design and Implementation (PLDI)*, Albuquerque, New Mexico, USA, June 23-25, 1993, pages 237–247, 1993.
- [11] M. Fluet and S. Weeks. Contification using dominators. In *Proceedings of the Sixth ACM SIGPLAN International Conference on Functional Programming (ICFP ’01)*, Firenze (Florence), Italy, September 3-5, 2001., pages 2–13, 2001.
- [12] A. Gill and G. Hutton. The worker/wrapper transformation. *J. Funct. Program.*, 19(2):227–251, 2009.
- [13] J.-Y. Girard, P. Taylor, and Y. Lafont. *Proofs and types*, volume 7. Cambridge University Press Cambridge, 1989.
- [14] J. Groff and C. Lattner. Swift intermediate language. LLVM Developers Meeting <http://www.llvm.org/devmtg/2015-10/>, 2015.
- [15] A. Keep, A. Hearn, and R. Dybvig. Optimizing closures in O(0) time. In *Annual Workshop on Scheme and Functional Programming*. ACM, 2012.
- [16] A. Kennedy. Compiling with continuations, continued. In *Proceedings of the 12th ACM SIGPLAN International Conference on Functional Programming, ICFP 2007, Freiburg, Germany, October 1-3, 2007*, pages 177–190, 2007.
- [17] S. Lindley. *Normalisation by evaluation in the compilation of typed functional programming languages*. PhD thesis, University of Edinburgh, College of Science and Engineering, School of Informatics, 2005.
- [18] C. Okasaki, P. Lee, and D. Tarditi. Call-by-need and continuation-passing style. *Lisp and Symbolic Computation*, 7(1):57–82, 1994.
- [19] S. Peyton Jones and A. Santos. A transformation-based optimiser for Haskell. *Science of Computer Programming*, 32(1-3):3–47, Sept. 1998.
- [20] S. Peyton Jones, W. Partain, and A. Santos. Let-floating: moving bindings to give faster programs. In *ACM SIGPLAN International Conference on Functional Programming (ICFP’96)*, Philadelphia, May 1996. ACM.
- [21] S. L. Peyton Jones. Call-pattern specialisation for Haskell programs. In *Proceedings of the 12th ACM SIGPLAN International Conference on Functional Programming, ICFP 2007, Freiburg, Germany, October 1-3, 2007*, pages 327–337, 2007.
- [22] S. L. Peyton Jones, A. Tolmach, and T. Hoare. Playing by the rules: rewriting as a practical optimisation technique in GHC. In R. Hinze, editor, *2001 Haskell Workshop*. ACM, September 2001.
- [23] J. H. Reppy. Optimizing nested loops using local CPS conversion. *Higher-Order and Symbolic Computation*, 15(2-3):161–180, 2002.
- [24] G. L. Steele, Jr. RABBIT: A compiler for SCHEME. Technical Report AITR-474, Massachusetts Institute of Technology, Artificial Intelligence Laboratory, 1978.
- [25] M. Sulzmann, M. Chakravarty, S. P. Jones, and K. Donnelly. System F with type equality coercions. In *ACM SIGPLAN International Workshop on Types in Language Design and Implementation (TLDI’07)*, pages 53–66. ACM, January 2007. ISBN 1-59593-393-X.
- [26] J. Svenningsson. Shortcut fusion for accumulating parameters & zip-like functions. In *Proceedings of the Seventh ACM SIGPLAN International Conference on Functional Programming (ICFP ’02)*, Pittsburgh, Pennsylvania, USA, October 4-6, 2002., pages 124–132, 2002.
- [27] A. P. Tolmach and D. Oliva. From ML to Ada: Strongly-typed language interoperability via source translation. *J. Funct. Program.*, 8(4):367–412, 1998.

A. Proof of type safety

We did not give a type system for configurations, so the off-the-shelf proof of progress and preservation is not quite applicable. However, we can adapt easily enough by annotating each configuration with a well-typed term that corresponds to it. Write $\langle e/s; \Sigma; e \rangle$ (or $\langle e/c \rangle$) for an *annotated configuration*. We will need to track the connection between e and c , for which we need a few tools. Let B be a *binding context*, that is, series of **let** bindings surrounding a hole. Then write $\llbracket B \rrbracket$ for the store containing those same bindings (but with recursive groups flattened). Also, let $\llbracket E \rrbracket$ translate the evaluation context E to a stack (which is of course just another syntax for the same structure). Then let \sim relate terms to configurations such that

$$B[E[e]] \sim \langle e; \llbracket E \rrbracket; \llbracket B \rrbracket \rangle.$$

Now, write

$$\langle e/c \rangle : \tau$$

when $e \sim c$ and $\varepsilon; \varepsilon \vdash e : \tau$, and write

$$\langle e/c \rangle \mapsto \langle e'/c' \rangle$$

if $c \mapsto c'$.

It will also be convenient to consider the types of binding contexts. Let

$$\Gamma \vdash B : \Gamma'$$

denote that B binds the symbols in Γ' .

We need a few utilities before we tackle the proof.

Proposition 6 (Substitution). *1. If $\Gamma, x : \sigma; \Delta \vdash e : \tau$ and $\Gamma; \Delta \vdash v : \sigma$, then $\Gamma, \Delta \vdash e\{v/x\} : \tau$.
2. If $\Gamma, a; \Delta \vdash e : \tau$, then $\Gamma; \Delta \vdash e\{\sigma/a\} : \tau\{\sigma/a\}$.*

Lemma 7. *If $\Gamma; \Delta \vdash E[B[e]] : \tau$ and variables bound by B aren't free in E , then $\Gamma; \Delta \vdash B[E[\tau]]$.*

Proof. By induction on E and then B . □

Lemma 8. $\Gamma \vdash B[e] : \tau$ if and only if there exists Γ' with $\Gamma, \Gamma' \vdash e : \tau$ and $\Gamma \vdash B : \Gamma'$.

Proof. By induction on B . □

Now we are ready:

Lemma 9 (Progress and preservation). *If $\langle e/c \rangle : \tau$, then either*

1. $c \equiv \langle A; \varepsilon; \Sigma \rangle$, where A is an answer, or
2. $\langle e/c \rangle \mapsto \langle e'/c' \rangle$ for some e' and c' with $\langle e'/c' \rangle : \tau$.

Proof. Since $e \sim c$, we have $e \equiv B[E[e_0]]$ and $c \equiv \langle e_0; \llbracket E \rrbracket; \llbracket B \rrbracket \rangle$ for some B , E , and e_0 . Proceed by case analysis on e_0 .

- For $e_0 \equiv F[e_1]$, where F is any frame, the *push* rule applies. From the term's perspective, *push* does nothing,

since it only shuffles part of the configuration around, so we can take $e' \sim e$ and $e' : \tau$ holds by assumption.

$$\begin{aligned} B[E[e_0]] &\equiv B[E[F[e_1]]] \sim \langle F[\llbracket e_1 \rrbracket]; \llbracket E \rrbracket; \llbracket B \rrbracket \rangle \\ &\mapsto \langle e_1; F : \llbracket E \rrbracket; \llbracket B \rrbracket \rangle \\ &\sim B[E[F[e_1]]] : \tau \end{aligned}$$

- For $e_0 \equiv \mathbf{let} \ vb \ \mathbf{in} \ e_1$, the *bind* rule applies, and we finish with Lemmas 7 and 8.

$$\begin{aligned} B[E[e_0]] &\equiv B[E[\mathbf{let} \ vb \ \mathbf{in} \ e_1]] \\ &\sim \langle \mathbf{let} \ vb \ \mathbf{in} \ e_1; \llbracket E \rrbracket; \llbracket B \rrbracket \rangle \\ &\mapsto \langle e_1; \llbracket E \rrbracket; \llbracket B \rrbracket, \llbracket vb \rrbracket \rangle \\ &\sim B[\mathbf{let} \ vb \ \mathbf{in} \ E[e_1]] : \tau \text{ (by Lemma 7, 8)} \end{aligned}$$

- For $e_0 \equiv \mathbf{jump} \ j \ \vec{\varphi} \ \vec{v} \ \tau'$, note that there must be a matching **join** in s (provable by induction on E). In other words,

$$E \equiv E_1[\mathbf{join} \ jb \ \mathbf{in} \ E_2[\mathbf{jump} \ j \ \vec{\varphi} \ \vec{v} \ \tau']]$$

where some $j \ \vec{a} \ \vec{x} = u \in jb$. Then the body u must have the same type as $E_2[\mathbf{jump} \ j \ \vec{\varphi} \ \vec{v} \ \tau']$, and we finish by Prop. 6.

$$\begin{aligned} B[E[e_0]] &\equiv B[E[\mathbf{jump} \ j \ \vec{\varphi} \ \vec{v} \ \tau']] \\ &\equiv B[E_1[\mathbf{join} \ j \ \vec{a} \ \vec{x} = u \ \mathbf{in} \ E_2[\mathbf{jump} \ j \ \vec{\varphi} \ \vec{v} \ \tau']]] \\ &\sim \left\langle \llbracket E_1 \rrbracket ++ \left(\begin{array}{c} \mathbf{jump} \ j \ \vec{\varphi} \ \vec{v} \ \tau'; \\ \llbracket B \rrbracket \end{array} \right) \right\rangle \\ &\mapsto \left\langle \llbracket E_1 \rrbracket ++ \left(\begin{array}{c} u\{\vec{\varphi}/\vec{a}\}\{\vec{v}/\vec{x}\}; \\ \mathbf{join} \ j \ \vec{a} \ \vec{x} = u \ \mathbf{in} \ \square; \\ \llbracket B \rrbracket \end{array} \right) \right\rangle \\ &\sim B[E_1[\mathbf{join} \ j \ \vec{a} \ \vec{x} = u \ \mathbf{in} \ u\{\vec{\varphi}/\vec{a}\}\{\vec{v}/\vec{x}\}]] : \tau \\ &\text{(by Prop. 6)} \end{aligned}$$

This development assumes a non-recursive join point; nothing changes for a recursive one.

- For $e_0 \equiv A$, examine E :
 - If $E \equiv \square$, we are done (case 1).
 - If $E \equiv E'[\mathbf{join} \ jb \ \mathbf{in} \ \square]$, then *ans* applies. The reduct typechecks by a standard strengthening lemma, since no label can appear free in an answer.

$$\begin{aligned} B[E[e_0]] &\equiv B[(E'[\mathbf{join} \ jb \ \mathbf{in} \ \square])[e_0]] \\ &\equiv B[E'[\mathbf{join} \ jb \ \mathbf{in} \ e_0]] \\ &\sim \langle e_0; \llbracket E' \rrbracket; \mathbf{join} \ jb \ \mathbf{in} \ \square; \llbracket B \rrbracket \rangle \\ &\mapsto \langle e_0; \llbracket E' \rrbracket; \llbracket B \rrbracket \rangle \\ &\sim B[E'[e_0]] : \tau \end{aligned}$$

- Otherwise, the outermost frame must be of the correct form according to the type of e_0 , so one of β , β_τ , or *case* applies. In each case we finish with either Lemma 7 or Prop. 6.

- For $e_0 \equiv x$, by Lemma 8 we must have that $B : \Gamma$ and because $x \in \text{dom } \Gamma$, bind applies; then the RHS of x has the same type that x typechecked under,

$$\begin{aligned} B[E[e_0]] &\equiv B[E[x]] \sim \langle x; \llbracket E \rrbracket; \llbracket B \rrbracket \rangle \\ &\mapsto \langle u; \llbracket E \rrbracket; \llbracket B \rrbracket \rangle \\ &\sim B[E[u]] : \tau \end{aligned}$$

□

Proof of Thm. 1. Generalize from the initial configuration to any $\langle e/c \rangle : \tau$, since clearly $e \sim \langle e; \varepsilon; \varepsilon \rangle$ and hence $\langle e/\varepsilon; \varepsilon; e \rangle : \tau$. Proceed by coinduction. By Lemma 9, either c is an answer configuration (proving case 2) or $\langle e/c \rangle \mapsto \langle e'/c' \rangle$ where $\langle e'/c' \rangle : \tau$. This may proceed forever, proving case 1, or else eventually there must be an answer. □

B. Proof of soundness of equational axioms

We reuse the notation from App. A.

The first order of business is to reconcile the two semantics—the operational semantics relates configurations, but the rewrite rules relate terms. Thus we create an alternative operational semantics that relates terms (see Fig. 6). Note that this operational semantics is not strictly deterministic because there may be several bind redexes within a term. The bind rule can be made deterministic by either favoring the outermost bind redex over all others, or delaying bind reduction until no other rules apply. However, doing so is not necessary for our goal of proving soundness of the equational theory (and in fact it makes the task considerably more difficult without gain), and we already have a deterministic operational reduction relation based on configurations.

To formulate the correspondence, we divide the rules into *external* and *internal* categories. For the original configuration-based semantics, the push and bind rules are internal and the rest are external; for the new one, only bind is internal (there is no push).

Our \sim relation from last section can now be more generally characterized for the looser operational semantics which uses terms of the form $S[e]$ as

$$e \sim c \text{ iff } \langle e; \varepsilon; \varepsilon \rangle \mapsto_i^* c.$$

which effectively divides S into a B and E which are translated as before. From this understanding, \sim is now a *bisimulation*.

Proposition 10. *If $e \sim c$, then:*

1. If $e \mapsto_i^* \mapsto_e e'$, then $c \mapsto_i^* \mapsto_e c'$ with $e' \sim c'$.
2. If $c \mapsto_i^* \mapsto_e c'$, then $e \mapsto_i^* \mapsto_e e'$ with $e' \sim c'$.

Proof. Before we demonstrate the bisimulation, observe that for every S , there are (unique) Σ and s such that $\langle S[e]; \varepsilon; \varepsilon \rangle \mapsto_i^* \langle e; s; \Sigma \rangle$, which follows by induction on S . Additionally, if $\langle e; \varepsilon; \varepsilon \rangle \mapsto_i^* \langle e'; s; \Sigma \rangle$ then there is some S such that $e \equiv S[e'] \mapsto_i^* \llbracket \Sigma \rrbracket \llbracket \llbracket s \rrbracket \llbracket e' \rrbracket \rrbracket$ (where $\llbracket \Sigma \rrbracket$ and $\llbracket s \rrbracket$ is

the reverse translation from heaps to binding contexts and from stacks to evaluation contexts), which follows by induction on \mapsto_i^* and performing bind reductions as necessary.

1. First, we show that if $e \mapsto_i e'$ then $c \mapsto_i^* c'$ with $e' \sim c'$. In particular, there is only the one internal reduction bind ($S[F[\text{let } vb \text{ in } e]] \mapsto_i S[\text{let } vb \text{ in } F[e]]$). Before the bind reduction, we have for some Σ and s

$$\begin{aligned} \langle S[F[\text{let } vb \text{ in } e]]; \varepsilon; \varepsilon \rangle &\mapsto_i^* \langle F[\text{let } vb \text{ in } e]; s; \Sigma \rangle \\ &\mapsto_i \langle \text{let } vb \text{ in } e; F : s; \Sigma \rangle \\ &\mapsto_i \langle e; F : s; \Sigma, vb \rangle \end{aligned}$$

whereas after the bind reduction, we have

$$\begin{aligned} \langle S[\text{let } vb \text{ in } F[e]]; \varepsilon; \varepsilon \rangle &\mapsto_i^* \langle \text{let } vb \text{ in } F[e]; s; \Sigma \rangle \\ &\mapsto_i \langle F[e]; s; \Sigma, vb \rangle \\ &\mapsto_i \langle e; F : s; \Sigma, vb \rangle \end{aligned}$$

Either $c \mapsto_i^* \langle e; F : s; \Sigma, vb \rangle$ —in which case we can take $c' \equiv \langle e; F : s; \Sigma, vb \rangle$ —or $\langle e; F : s; \Sigma, vb \rangle \mapsto_i^* c'$ —in which case, we can take $c' \equiv c$.

Second, we show that if $e \mapsto_e e'$ then $c \mapsto_i^* \mapsto_e c'$ with $e' \sim c'$ by cases on the possible reductions for $e \mapsto_e e'$. Each case is similar to β ($S[(\lambda x:\sigma.e) u] \mapsto_e S[\text{let } x:\sigma = u \text{ in } e]$). Before β reduction, we have for some Σ and s

$$\begin{aligned} \langle S[(\lambda x:\sigma.e) u]; \varepsilon; \varepsilon \rangle &\mapsto_i^* \langle (\lambda x:\sigma.e) u; s; \Sigma \rangle \\ &\mapsto_i \langle \lambda x:\sigma.e; \square u : s; \Sigma \rangle \\ &\mapsto_e \langle e; s; \Sigma, x:\sigma = u \rangle \end{aligned}$$

whereas after the β reduction, we have

$$\begin{aligned} \langle S[\text{let } x:\sigma = u \text{ in } e]; \varepsilon; \varepsilon \rangle &\mapsto_i^* \langle \text{let } x:\sigma = u \text{ in } e; s; \Sigma \rangle \\ &\mapsto_i \langle e; s; \Sigma, x:\sigma = u \rangle \end{aligned}$$

Since $\langle \lambda x:\sigma.e; \square u : s; \Sigma \rangle \not\mapsto_i$, it must be that $c \mapsto_i^* \langle \lambda x:\sigma.e; \square u : s; \Sigma \rangle$, so we can take $c' \equiv \langle e; s; \Sigma, x:\sigma = u \rangle$.

Therefore, the result follows by induction on \mapsto_i^* using the above two facts.

2. First, we show that if $c \mapsto_i c'$ then with $e \sim c'$. Since $e \sim c$ we know $e \mapsto_i^* c \mapsto_i c'$ by definition, and so $e \sim c'$. Second, we show that if $c \mapsto_e c'$ then $e \mapsto_i^* \mapsto_e e'$ with $e' \sim c'$ by cases on the possible reductions for $c \mapsto_e c'$. Each case is similar to β ($\langle \lambda x:\sigma.u; \square v : s; \Sigma \rangle \mapsto_e \langle u; s; \Sigma, x:\sigma = u \rangle$). Because $e \sim \langle \lambda x:\sigma.u; \square v : s; \Sigma \rangle$ we have $e \mapsto_i^* \llbracket \Sigma \rrbracket \llbracket \llbracket s \rrbracket \llbracket (\lambda x:\sigma.u) v \rrbracket \rrbracket$, so that

$$\begin{aligned} e \mapsto_i^* \llbracket \Sigma \rrbracket \llbracket \llbracket s \rrbracket \llbracket (\lambda x:\sigma.u) v \rrbracket \rrbracket &\mapsto_e \llbracket \Sigma \rrbracket \llbracket \llbracket s \rrbracket \llbracket \text{let } x:\sigma = v \text{ in } u \rrbracket \rrbracket \\ &\sim \langle u; s; \Sigma, x:\sigma = u \rangle \end{aligned}$$

Therefore, the result follows by induction on \mapsto_i^* using the above two facts. □

$$\begin{array}{l}
B ::= \square \mid \text{let } vb \text{ in } B \\
S ::= \square \mid \text{let } vb \text{ in } S \mid F[S] \\
\\
\boxed{e \mapsto_e e', e \mapsto_i e', e \mapsto e'} \\
(\lambda x:\sigma.e)u \mapsto_e \text{let } x:\sigma = u \text{ in } e \quad (\beta) \\
(\Lambda a.e)\varphi \mapsto_e e\{\varphi/a\} \quad (\beta_\tau) \\
\text{let } vb \text{ in } S[x] \mapsto_e \text{let } vb \text{ in } S[x] \quad (\text{look}) \\
\text{case } K \vec{\varphi} \vec{u} \text{ of } \vec{alt} \mapsto_e \text{let } \overline{x:\sigma} = \vec{u} \text{ in } e \quad (\text{case}) \\
\text{if } (x:\sigma = v) \in vb, x \notin \text{bv}(S) \\
\text{if } (K \overline{x:\sigma} \rightarrow e) \in \vec{alt} \\
\text{join } jb \text{ in } E[\text{jump } j \vec{\varphi} \vec{v} \tau] \mapsto_e \text{let } \overline{x:\sigma} = \vec{v} \text{ in join } jb \text{ in } e\{\vec{\varphi}/\vec{a}\} \quad (\text{jump}) \\
\text{if } (j \vec{a} \overline{x:\sigma} = e) \in jb, j \notin \text{bv}(E) \\
\text{join } jb \text{ in } A \mapsto_e A \quad (\text{ans}) \\
\\
F[\text{let } vb \text{ in } e] \mapsto_i \text{let } vb \text{ in } F[e] \quad (\text{bind}) \\
\\
\frac{e \mapsto_e e'}{S[e] \mapsto_e S[e']} \quad \frac{e \mapsto_i e'}{S[e] \mapsto_i S[e']} \quad \frac{e \mapsto_e e'}{e \mapsto e'} \quad \frac{e \mapsto_i e'}{e \mapsto e'} \quad \frac{e \mapsto e'}{S[e] \mapsto S[e']}
\end{array}$$

Figure 6: Term-based operational semantics for System F_J .

We can now restate observational equivalence in terms of standard reductions on terms.

Proposition 11. $e \cong e'$ if and only if for all C , $C[e]$ diverges if and only if $C[e']$ diverges.

Proof. First, we show that given $e \sim c$, e diverges if and only if c diverges. Note that the internal reductions of both operational semantics are strongly normalizing. Therefore, any infinite reduction sequence of e contains an infinite number of external reductions, and similarly for c . The correspondence of divergence then follows from the bisimulation in Lemma 10.

Now, suppose that $e \cong e'$. From the above, if $C[e]$ diverges according to the term-based operational semantics, then

- $\langle C[e]; \varepsilon; \varepsilon \rangle$ diverges because $C[e] \sim \langle C[e]; \varepsilon; \varepsilon \rangle$,
- $\langle C[e']; \varepsilon; \varepsilon \rangle$ diverges because $e \cong e'$, and
- $C[e']$ diverges because $C[e'] \sim \langle C[e']; \varepsilon; \varepsilon \rangle$.

So for all C , $C[e]$ diverges if and only if $C[e']$ diverges. Going the other way, suppose that for all C , $C[e]$ diverges if and only if $C[e']$ diverges. Similarly, if $\langle C[e]; \varepsilon; \varepsilon \rangle$ diverges, then

- $C[e]$ diverges because $C[e] \sim \langle C[e]; \varepsilon; \varepsilon \rangle$,
- $C[e']$ diverges because $C[e]$ diverges, and
- $\langle C[e']; \varepsilon; \varepsilon \rangle$ because $C[e'] \sim \langle C[e']; \varepsilon; \varepsilon \rangle$.

So $e \cong e'$. \square

Now that we have our footing, we demonstrate Theorem 3 via two common properties of reduction relations: *confluence*

and *standardization*. Let \longrightarrow be defined by the rules for $=$ read as-is from left to right and \rightarrow be the compatible closure of \longrightarrow . See Fig. 7 for the precise definitions, and note the presence of two extra rules *confitydrop* and *confitydrop_{rec}* which are variants of *confity* and *confity_{rec}*, as well as the *letcomm* which commutes let bindings. The two extra confity rules do not extend the reduction theory (i.e., \rightarrow^* , the reflexive-transitive closure of \rightarrow), since they are simulated by *confity* and *confity_{rec}* with the help of *jdrops*, however they will be helpful bigger-step-reductions for the purposes of demonstrating standardization.

Confluence can be shown directly using the available powerful methodologies, in particular the *decreasing diagrams* technique for establishing confluence.

Theorem 12 (Confluence). *If $e \rightarrow^* e_1$ and $e \rightarrow^* e_2$ then $e_1 \rightarrow^* e'$ and $e_2 \rightarrow^* e'$ for some e' .*

Proof. Confluence follows by a decreasing diagrams argument using the following ordering among the reduction rules:

$$\begin{array}{l}
float < casefloat < jfloat \\
< \beta < \beta_\tau < case \\
< inline < jinlincomm < letcomm \\
< confity < confity_{rec}
\end{array}$$

Note that we do not need to consider *confitydrop* and *confitydrop_{rec}* since they are simulated by the other reduction rules and thus are subsumed by them in the reflexive-transitive

$$\boxed{e \rightarrow e', e \rightarrow e'}$$

$$\frac{}{(\lambda x:\sigma.e)u \rightarrow \text{let } x:\sigma = u \text{ in } e} \beta \quad \frac{}{(\Lambda a.e)\varphi \rightarrow e\{\varphi/a\}} \beta_\tau \quad \frac{K \vec{x}:\vec{\sigma} \rightarrow e \in \overline{\text{alt}}}{\text{case } K \vec{\varphi} \vec{u} \text{ of } \overline{\text{alt}} \rightarrow \text{let } \vec{x}:\vec{\sigma} = \vec{u} \text{ in } e} \text{case}$$

$$\frac{x:\sigma = v \in vb}{\text{let } vb \text{ in } C[x] \rightarrow \text{let } vb \text{ in } C[v]} \text{inline} \quad \frac{\text{bv}(vb) \cap \text{fv}(e) = \emptyset}{\text{let } vb \text{ in } e \rightarrow e} \text{drop} \quad \frac{\text{bv}(vb) \cap \text{fv}(vb') = \emptyset = \text{bv}(vb') \cap \text{fv}(vb)}{\text{let } vb \text{ in let } vb' \text{ in } e \rightarrow \text{let } vb' \text{ in let } vb \text{ in } e} \text{letcomm}$$

$$\frac{(j \vec{a} \vec{x}:\vec{\sigma} = u) \in jb}{\text{join } jb \text{ in } L[\text{jump } j \vec{\varphi} \vec{v} \tau] \rightarrow \text{join } jb \text{ in } L[\text{let } \vec{x}:\vec{\sigma} = \vec{v} \text{ in } u\{\varphi/a\}]} \text{jinline} \quad \frac{\text{bv}(jb) \cap \text{fv}(e) = \emptyset}{\text{join } jb \text{ in } e \rightarrow e} \text{jdrop}$$

$$\frac{\text{bv}(vb) \cap \text{fv}(E) = \emptyset}{E[\text{let } vb \text{ in } e] \rightarrow \text{let } vb \text{ in } E[e]} \text{float} \quad \frac{\text{bv}(jb) \cap \text{fv}(E) = \emptyset}{E[\text{join } jb \text{ in } e] \rightarrow \text{join } E[jb] \text{ in } E[e]} \text{jfloat}$$

$$\frac{(\bigcup \vec{x}) \cap \text{fv}(E) = \emptyset}{E[\text{case } e \text{ of } \overline{K \vec{x} \rightarrow u}] \rightarrow \text{case } e \text{ of } \overline{K \vec{x} \rightarrow E[u]}} \text{casefloat} \quad \frac{E[\text{jump } j \vec{\varphi} \vec{e} \tau] : \tau'}{E[\text{jump } j \vec{\varphi} \vec{e} \tau] \rightarrow \text{jump } j \vec{\varphi} \vec{e} \tau'} \text{abort}$$

$$\frac{\text{let } f = \Lambda \vec{a}.\lambda \vec{x}.u \text{ in } L[\vec{e}] : \tau \quad u : \tau \quad f \notin \text{fv}(L) \quad \rho(f \vec{a} \vec{x}) = \text{jump } j \vec{a} \vec{x} \tau}{\text{let } f = \Lambda \vec{a}.\lambda \vec{x}.u \text{ in } L[\vec{e}] \rightarrow \text{join } j \vec{a} \vec{x} = u \text{ in } L[\text{tail}_\rho(e)]} \text{contify}$$

$$\frac{\text{let rec } f = \Lambda \vec{a}.\lambda \vec{x}.L[\vec{u}] \text{ in } L'[\vec{e}] : \tau \quad L[\vec{u}] : \tau \quad (\bigcup \vec{f}) \cap (\bigcup \text{fv}(L) \cup \text{fv}(L')) = \emptyset \quad \rho(f \vec{a} \vec{x}) = \text{jump } j \vec{a} \vec{x} \tau}{\text{let rec } f = \Lambda \vec{a}.\lambda \vec{x}.L[\vec{u}] \text{ in } L'[\vec{e}] \rightarrow \text{join rec } j \vec{a} \vec{x} = L[\text{tail}_\rho(u)] \text{ in } L'[\text{tail}_\rho(e)]} \text{contify}_{\text{rec}}$$

$$\frac{E \in \text{Tail Contexts} \quad f \notin \text{fv}(E[e]) \quad \text{bv}(E) \cap \text{fv}(e) = \emptyset}{\text{let } f = \Lambda \vec{a}.\lambda \vec{x}.u \text{ in } E[e] \rightarrow \text{join } j \vec{a} \vec{x} = u \text{ in } e} \text{contifydrop}$$

$$\frac{L[\vec{u}] : \tau \quad \emptyset = \text{bv}(E) \cap \text{fv}(e) \quad \rho(f \vec{a} \vec{x}) = \text{jump } j \vec{a} \vec{x} \tau}{\text{let rec } f = \Lambda \vec{a}.\lambda \vec{x}.L[\vec{u}] \text{ in } E[L'[e]] : \tau \quad \emptyset = (\bigcup \vec{f}) \cap (\bigcup \text{fv}(L) \cup \text{fv}(E[L']))) \quad \rho(f \vec{a} \vec{x}) = \text{jump } j \vec{a} \vec{x} \tau} \text{contifydrop}_{\text{rec}}$$

$$\frac{}{\text{let rec } f = \Lambda \vec{a}.\lambda \vec{x}.L[\vec{u}] \text{ in } E[L'[\vec{e}]] \rightarrow \text{join rec } j \vec{a} \vec{x} = L[\text{tail}_\rho(u)] \text{ in } L'[\text{tail}_\rho(e)]}$$

$$\boxed{e \rightarrow e', e \rightsquigarrow e', e \rightsquigarrow e', e \rightsquigarrow_B e'}$$

$$\frac{e \rightarrow e'}{C[e] \rightarrow C[e']} \quad \frac{e \rightarrow e' \quad e \not\rightsquigarrow e'}{e \rightsquigarrow e'} \quad \frac{e \rightarrow e' \quad e \not\rightsquigarrow e'}{e \rightsquigarrow e'} \quad \frac{e \rightsquigarrow e'}{B[e] \rightsquigarrow_B B[e']}$$

Figure 7: Single-step reduction relations for System F_J .

reduction relation. Additionally, we will replace the *jinline* reduction rule with the following more general *jinlinecomm* rule where \overrightarrow{EL} is a composition of many evaluation and tail contexts:

$$\begin{aligned}
& \text{join } jb \text{ in } \overrightarrow{EL}[\vec{e}, \text{jump } j \vec{\varphi} \vec{v} \tau, \vec{e}'] \\
& \rightarrow \text{join } jb \text{ in } \overrightarrow{L}[\overrightarrow{E}[e], \text{let } \vec{x} = \vec{v} \text{ in } u\{\varphi/a\}, \overrightarrow{E}[e']] \\
& \text{if } (j \vec{a} \vec{x} = u) \in jb, j \notin \text{bv}(E)
\end{aligned}$$

which does not change the reflexive-transitive reduction relation since *jinlinecomm* is simulated by the *jinline*, *float*, *casefloat*, *jfloat*, and *abort* reduction rules. The non-trivial critical pairs (ones that join in more than one step or using other rules) are then:

- The *case-casefloat* critical-pair:

$$\begin{aligned}
& E[\text{case } K_i \vec{\varphi} \vec{u} \text{ of } \overline{K_i \vec{x} \rightarrow e_i}] \rightarrow_{\text{case}} E[\text{let } \vec{x}:\vec{\sigma} = \vec{u} \text{ in } e_i] \\
& E[\text{case } K_i \vec{\varphi} \vec{u} \text{ of } \overline{K_i \vec{x} \rightarrow e_i}]
\end{aligned}$$

$$\rightarrow_{\text{casefloat}} \text{case } K_i \vec{\varphi} \vec{u} \text{ of } \overline{K_i \vec{x} \rightarrow E[e_i]}^i$$

joins to $\text{let } \vec{x}:\vec{\sigma} = \vec{u} \text{ in } E[e_i]$ as follows:

$$\begin{aligned}
& E[\text{let } \vec{x}:\vec{\sigma} = \vec{u} \text{ in } e_i] \xrightarrow{\text{float}} \text{let } \vec{x}:\vec{\sigma} = \vec{u} \text{ in } E[e_i] \\
& \text{case } K_i \vec{\varphi} \vec{u} \text{ of } \overline{K_i \vec{x} \rightarrow E[e_i]}^i \rightarrow_{\text{case}} \text{let } \vec{x}:\vec{\sigma} = \vec{u} \text{ in } E[e_i]
\end{aligned}$$

which is decreasing because *float* < *case*.

- The *jinlinecomm-jfloat* critical pair:

$$\begin{aligned}
& E'[\text{join } jb \text{ in } \overrightarrow{EL}[\vec{e}, \text{jump } j \vec{\varphi} \vec{v} \tau, \vec{e}']] \\
& \xrightarrow{\text{jinlinecomm}} E'[\text{join } jb \text{ in } \overrightarrow{L}[\overrightarrow{E}[e], \text{let } \vec{x}:\vec{\sigma} = \vec{v} \text{ in } u\{\varphi/a\}, \overrightarrow{E}[e']]] \\
& \quad \text{where } (j \vec{a} \vec{x} = u) \in jb \\
& E'[\text{join } jb \text{ in } \overrightarrow{EL}[\vec{e}, \text{jump } j \vec{\varphi} \vec{v} \tau, \vec{e}']] \\
& \xrightarrow{\text{jfloat}} \text{join } E'[jb] \text{ in } E'[\overrightarrow{EL}[\vec{e}, \text{jump } j \vec{\varphi} \vec{v} \tau, \vec{e}']]
\end{aligned}$$

joins to

$$\begin{array}{l} \mathbf{join} \\ E'[jb] \\ \mathbf{in} \overrightarrow{L}[E'[\overrightarrow{E}[e]], \mathbf{let} \overline{x:\sigma} = \overline{v} \mathbf{in} E'[u]\{\overline{\varphi/a}\}, \overrightarrow{E}'[e']] \end{array}$$

as follows:

$$\begin{array}{l} E'[\mathbf{join} \, jb \mathbf{in} \overrightarrow{L}[\overrightarrow{E}[e], \mathbf{let} \overline{x:\sigma} = \overline{v} \mathbf{in} u\{\overline{\varphi/a}\}, \overrightarrow{E}'[e']]] \\ \xrightarrow{*}_{float, casefloat, jfloat} \\ \mathbf{join} \\ E'[jb] \\ \mathbf{in} \overrightarrow{L}[E'[\overrightarrow{E}[e]], \mathbf{let} \overline{x:\sigma} = \overline{v} \mathbf{in} E'[u]\{\overline{\varphi/a}\}, \overrightarrow{E}'[e']] \\ \mathbf{join} E'[jb] \mathbf{in} E'[\overrightarrow{EL}[\overline{e}, \mathbf{jump} \, j \, \overline{\varphi} \, \overline{v} \, \tau, \overline{e'}]] \\ \xrightarrow{jinlinecomm} \\ \mathbf{join} \\ E'[jb] \\ \mathbf{in} \overrightarrow{L}[E'[\overrightarrow{E}[e]], \mathbf{let} \overline{x:\sigma} = \overline{v} \mathbf{in} E[u]\{\overline{\varphi/a}\}, \overrightarrow{E}'[e']] \end{array}$$

which is decreasing because $float < casefloat < jfloat < jinlinecomm$.

- The *inline-contify* critical pair:

$$\begin{array}{l} \mathbf{let} \, f = \Lambda \overline{a} . \lambda \overline{x} . u \mathbf{in} L[\overline{e}, f \, \overline{\varphi} \, \overline{v}, \overline{e'}] \\ \xrightarrow{inline} \mathbf{let} \, f = \Lambda \overline{a} . \lambda \overline{x} . u \mathbf{in} L[\overline{e}, (\Lambda \overline{a} . \lambda \overline{x} . u) \, \overline{\varphi} \, \overline{v}, \overline{e'}] \\ \mathbf{let} \, f = \Lambda \overline{a} . \lambda \overline{x} . u \mathbf{in} L[\overline{e}, f \, \overline{\varphi} \, \overline{v}, \overline{e'}] \\ \xrightarrow{contify} \mathbf{join} \\ j \, \overline{a} \, \overline{x} = u \\ \mathbf{in} L[\overrightarrow{tail}_\rho(e), \mathbf{jump} \, j \, \overline{\varphi} \, \overline{v} \, \tau, \overrightarrow{tail}_\rho(e')] \end{array}$$

joins to

$$\mathbf{join} \, j \, \overline{a} \, \overline{x} = u \mathbf{in} L[\overrightarrow{tail}_\rho(e), \mathbf{let} \overline{x} = \overline{v} \mathbf{in} u\{\overline{\varphi/a}\}, \overrightarrow{tail}_\rho(e')]$$

as follows:

$$\begin{array}{l} \mathbf{let} \, f = \Lambda \overline{a} . \lambda \overline{x} . u \mathbf{in} L[\overline{e}, (\Lambda \overline{a} . \lambda \overline{x} . u) \, \overline{\varphi} \, \overline{v}, \overline{e'}] \\ \xrightarrow{*}_{\beta_\tau} \mathbf{let} \, f = \Lambda \overline{a} . \lambda \overline{x} . u \mathbf{in} L[\overline{e}, (\lambda \overline{x} . u\{\overline{\varphi/a}\}) \, \overline{v}, \overline{e'}] \\ \xrightarrow{*}_{\beta, float} \mathbf{let} \, f = \Lambda \overline{a} . \lambda \overline{x} . u \mathbf{in} L[\overline{e}, \mathbf{let} \overline{x} = \overline{v} \mathbf{in} u\{\overline{\varphi/a}\}, \overline{e'}] \\ \xrightarrow{contify} \mathbf{join} \\ j \, \overline{a} \, \overline{x} = u \\ \mathbf{in} L[\overrightarrow{tail}_\rho(e), \mathbf{let} \overline{x} = \overline{v} \mathbf{in} u\{\overline{\varphi/a}\}, \overrightarrow{tail}_\rho(e')] \\ \mathbf{join} \, j \, \overline{a} \, \overline{x} = u \mathbf{in} L[\overrightarrow{tail}_\rho(e), \mathbf{jump} \, j \, \overline{\varphi} \, \overline{v} \, \tau, \overrightarrow{tail}_\rho(e')] \\ \xrightarrow{jinlinecomm} \mathbf{join} \\ j \, \overline{a} \, \overline{x} = u \\ \mathbf{in} L[\overrightarrow{tail}_\rho(e), \mathbf{let} \overline{x} = \overline{v} \mathbf{in} u\{\overline{\varphi/a}\}, \overrightarrow{tail}_\rho(e')] \end{array}$$

which is decreasing because $float < \beta < \beta_\tau < inline$ and $jinlinecomm < contify$.

- The *inline-contify_{rec}* critical pair joins similarly to the previous pair, which is also decreasing because $float < \beta < \beta_\tau < inline$ and $jinline < contify_{rec}$.

- The first *letcomm-contify* critical pair:

$$\begin{array}{l} \mathbf{let} \, vb \mathbf{in} \mathbf{let} \, f = \Lambda \overline{a} . \lambda \overline{x} . u \mathbf{in} L[\overline{e}] \\ \xrightarrow{letcomm} \mathbf{let} \, f = \Lambda \overline{a} . \lambda \overline{x} . u \mathbf{in} \mathbf{let} \, vb \mathbf{in} L[\overline{e}] \\ \mathbf{let} \, vb \mathbf{in} \mathbf{let} \, f = \Lambda \overline{a} . \lambda \overline{x} . u \mathbf{in} L[\overline{e}] \\ \xrightarrow{contify} \mathbf{let} \, vb \mathbf{in} \mathbf{join} \, j \, \overline{a} \, \overline{x} = u \mathbf{in} L[\overrightarrow{tail}_\rho(e)] \end{array}$$

joins to $\mathbf{let} \, vb \mathbf{in} \mathbf{join} \, j \, \overline{a} \, \overline{x} = u \mathbf{in} L[\overrightarrow{tail}_\rho(e)]$ as follows:

$$\begin{array}{l} \mathbf{let} \, f = \Lambda \overline{a} . \lambda \overline{x} . u \mathbf{in} \mathbf{let} \, vb \mathbf{in} L[\overline{e}] \\ \xrightarrow{contify} \mathbf{join} \, j \, \overline{a} \, \overline{x} = u \mathbf{in} \mathbf{let} \, vb \mathbf{in} L[\overrightarrow{tail}_\rho(e)] \\ \xrightarrow{float} \mathbf{let} \, vb \mathbf{in} \mathbf{join} \, j \, \overline{a} \, \overline{x} = u \mathbf{in} L[\overrightarrow{tail}_\rho(e)] \end{array}$$

which is decreasing because $float < letcomm < contify$. The second *letcomm-contify* critical pair:

$$\begin{array}{l} \mathbf{let} \, f = \Lambda \overline{a} . \lambda \overline{x} . u \mathbf{in} \mathbf{let} \, vb \mathbf{in} L[\overline{e}] \\ \xrightarrow{letcomm} \mathbf{let} \, vb \mathbf{in} \mathbf{let} \, f = \Lambda \overline{a} . \lambda \overline{x} . u \mathbf{in} L[\overline{e}] \\ \mathbf{let} \, f = \Lambda \overline{a} . \lambda \overline{x} . u \mathbf{in} \mathbf{let} \, vb \mathbf{in} L[\overline{e}] \\ \xrightarrow{contify} \mathbf{join} \, j \, \overline{a} \, \overline{x} = u \mathbf{in} \mathbf{let} \, vb \mathbf{in} L[\overrightarrow{tail}_\rho(e)] \end{array}$$

joins to $\mathbf{let} \, vb \mathbf{in} \mathbf{join} \, j \, \overline{a} \, \overline{x} = u \mathbf{in} L[\overrightarrow{tail}_\rho(e)]$ as follows:

$$\begin{array}{l} \mathbf{let} \, vb \mathbf{in} \mathbf{let} \, f = \Lambda \overline{a} . \lambda \overline{x} . u \mathbf{in} L[\overline{e}] \\ \xrightarrow{contify} \mathbf{let} \, vb \mathbf{in} \mathbf{join} \, j \, \overline{a} \, \overline{x} = u \mathbf{in} L[\overrightarrow{tail}_\rho(e)] \\ \mathbf{join} \, j \, \overline{a} \, \overline{x} = u \mathbf{in} \mathbf{let} \, vb \mathbf{in} L[\overrightarrow{tail}_\rho(e)] \\ \xrightarrow{float} \mathbf{let} \, vb \mathbf{in} \mathbf{join} \, j \, \overline{a} \, \overline{x} = u \mathbf{in} L[\overrightarrow{tail}_\rho(e)] \end{array}$$

which is decreasing because $float < contify$.

- The *letcomm-contify_{rec}* critical pairs join similarly to the previous pair, which is also decreasing because $float < letcomm < contify_{rec}$. □

Standardization—which states that for any result (*i.e.*, a term without a standard redex written $e \not\rightarrow$) reachable by the general \rightarrow , there is a similar expanded result that is reachable by the standard \mapsto relation—is harder. We will take the approach of postponing all non-standard reductions after standard ones. Let \succ be defined as any \rightarrow that is not standard, \succ be any \rightarrow that is not standard, and \succ_B be any \succ within a B context, as shown in 7. Sorting a reduction sequence to put the standard reductions first works as a standardization technique since non-standard reduction only relates results to other results.

- Lemma 13.** 1. If $e_1 \rightarrow e_2$ and $e_1 \not\rightarrow$ then $e_2 \mapsto_i^* e'_2 \not\rightarrow$.
2. If $e_1 \succ e_2$ and $e_1 \mapsto e'_1$ then $e_2 \mapsto e'_2$ for some e'_2 . In other words, if $e_1 \succ e_2 \not\rightarrow$, then $e_1 \not\rightarrow$.

Proof. • The only (non-standard) reduction that can create a standard redex are the *contify* family, which have the ability to create a *bind* redex by surrounding a let binding with a join binding. However, in this case any let bindings within the newly-created join point can be floated out with the *bind* reduction.

- The first fact holds because non-standard reductions cannot destroy standard redexes. The second fact follows from the first, since it is impossible to have a term $e_1 \mapsto e'_1$ such that $e_1 \mapsto e_2 \not\mapsto e'_1$.

□

We now seek to postpone non-standard reductions after standard ones. In general, this task is rather difficult to approach directly due to the possibility of duplicating reductions (which defeats attempts to generalize the single-step process to multiple steps), so first we only consider the subset of non-standard reductions which occur within the eye of a chain of let bindings, as described by the \mapsto_B reduction relation. These specific non-standard reductions can be postponed without duplicating the non-standard reduction, but at the cost of duplication standard reductions and introducing many uses of *letcomm* (denoted by $=_{\text{letcomm}}$ since it is inherently symmetric).

Lemma 14. *If $e \mapsto_B e_1 \mapsto e'$ then $e \mapsto^* e_2 =_{\text{letcomm}} e'_2 \mapsto_B^{0/1} e'$ for some e_2 and e'_2 .*

Proof. By cases on the possible reductions $e \mapsto_B e_1 \mapsto e'$. When the two reductions are non-overlapping they commute immediately in one step, as in

$$\begin{aligned} B[\text{let } vb \text{ in } e] &\mapsto_B B[e] \mapsto B[e'] \\ B[\text{let } vb \text{ in } e] &\mapsto B[\text{let } vb \text{ in } e'] \mapsto_B B[e'] \end{aligned}$$

for example. The other cases where the two reductions overlap in such a way that necessitates a more complex commutation are as follows.

For *casefloat* followed by *case*, we have:

$$\begin{aligned} B \left[E[\text{case } K_i \vec{\varphi} \vec{u} \text{ of } \overline{K_i x:\vec{\sigma}} \rightarrow e_i^i] \right] \\ \mapsto_B B \left[\text{case } K_i \vec{\varphi} \vec{u} \text{ of } \overline{K_i x:\vec{\sigma}} \rightarrow E[e_i^i] \right] \\ \mapsto B [\text{let } \overline{x:\vec{\sigma}} = \vec{u} \text{ in } E[e_i]] \\ B \left[E[\text{case } K_i \vec{\varphi} \vec{u} \text{ of } \overline{K_i x:\vec{\sigma}} \rightarrow e_i^i] \right] \\ \mapsto B [E[\text{let } \overline{x:\vec{\sigma}} = \vec{u} \text{ in } e_i]] \\ \mapsto^* B [\text{let } \overline{x:\vec{\sigma}} = \vec{u} \text{ in } E[e_i]] \end{aligned}$$

For *jfloat* followed by *jump*, we have

$$\begin{aligned} B [E[\text{join } jb \text{ in } E'[\text{jump } j \vec{u} \tau \vec{\varphi}]]] \\ \text{where } j \vec{a} \overline{x:\vec{\sigma}} = e \in jb \\ \mapsto_B B [\text{join } E[jb] \text{ in } E[E'[\text{jump } j \vec{u} \tau \vec{\varphi}]]] \\ \mapsto B [\text{let } \overline{x:\vec{\sigma}} = \vec{u} \text{ in join } E[jb] \text{ in } E[e\{\varphi/a\}]] \end{aligned}$$

$$\begin{aligned} B [E[\text{join } jb \text{ in } E'[\text{jump } j \vec{u} \tau \vec{\varphi}]]] \\ \text{where } j \vec{a} \overline{x:\vec{\sigma}} = e \in jb \\ \mapsto B [E[\text{let } \overline{x:\vec{\sigma}} = \vec{u} \text{ in join } jb \text{ in } e\{\varphi/a\}]] \\ \mapsto^* B [\text{let } \overline{x:\vec{\sigma}} = \vec{u} \text{ in } E[\text{join } jb \text{ in } e\{\varphi/a\}]] \\ \mapsto_B B [\text{let } \overline{x:\vec{\sigma}} = \vec{u} \text{ in join } E[jb] \text{ in } E[e\{\varphi/a\}]] \end{aligned}$$

For *contify* followed by *bind*, we have:

$$\begin{aligned} B [\text{let } f = \Lambda \vec{a} . \lambda \vec{x} . u \text{ in let } vb \text{ in } L[\vec{e}]] \\ \mapsto_B B [\text{join } j \vec{a} \vec{x} = u \text{ in let } vb \text{ in } L[\overline{\text{tail}_\rho(e)}]] \\ \text{where } \rho(f \vec{a} \vec{x}) = \text{jump } j \vec{a} \vec{x} \tau \\ \mapsto B [\text{let } vb \text{ in join } j \vec{a} \vec{x} = u \text{ in } L[\overline{\text{tail}_\rho(e)}]] \\ B [\text{let } f = \Lambda \vec{a} . \lambda \vec{x} . u \text{ in let } vb \text{ in } L[\vec{e}]] \\ =_{\text{letcomm}} B [\text{let } vb \text{ in let } f = \Lambda \vec{a} . \lambda \vec{x} . u \text{ in } L[\vec{e}]] \\ \mapsto_B B [\text{let } vb \text{ in join } j \vec{a} \vec{x} = u \text{ in } L[\overline{\text{tail}_\rho(e)}]] \\ \text{where } \rho(f \vec{a} \vec{x}) = \text{jump } j \vec{a} \vec{x} \tau \end{aligned}$$

For *contify* followed by *jump*, we have:

$$\begin{aligned} B [\text{let } f = \Lambda \vec{a} . \lambda \vec{x} . u \text{ in } L[\vec{e}, f \vec{\varphi} \vec{v}, \vec{e}']] : \tau \\ \mapsto_B B \left[\begin{array}{c} \text{join} \\ j \vec{a} \vec{x} = u \\ \text{in } L[\overline{\text{tail}_\rho(e)}, \text{jump } j \vec{\varphi} \vec{v} \tau, \overline{\text{tail}_\rho(e')}] \end{array} \right] \\ \text{where } \rho(f \vec{a} \vec{x}) = \text{jump } j \vec{a} \vec{x} \tau \\ L[\vec{e}, \square, \vec{e}'] \in \text{Evaluation Contexts} \\ \mapsto B [\text{let } \overline{x} = \vec{v} \text{ in join } j \vec{a} \vec{x} = u \text{ in } u\{\varphi/a\}] \\ B [\text{let } f = \Lambda \vec{a} . \lambda \vec{x} . u \text{ in } L[\vec{e}, f \vec{\varphi} \vec{v}, \vec{e}']] : \tau \\ \mapsto B [\text{let } f = \Lambda \vec{a} . \lambda \vec{x} . u \text{ in } L[\vec{e}, (\Lambda \vec{a} . \lambda \vec{x} . u) \vec{\varphi} \vec{v}, \vec{e}']] \\ \mapsto^* B [\text{let } f = \Lambda \vec{a} . \lambda \vec{x} . u \text{ in } L[\vec{e}, (\lambda \vec{x} . u\{\varphi/a\}) \vec{v}, \vec{e}']] \\ \mapsto^* =_{\text{letcomm}} \\ B [\text{let } \overline{x} = \vec{v} \text{ in let } f = \Lambda \vec{a} . \lambda \vec{x} . u \text{ in } L[\vec{e}, u\{\varphi/a\}, \vec{e}']] \\ \mapsto_B B [\text{let } \overline{x} = \vec{v} \text{ in join } j \vec{a} \vec{x} = u \text{ in } u\{\varphi/a\}] \end{aligned}$$

For *contifydrop_{rec}* followed by *bind*, we have:

$$\begin{aligned} B [\text{let rec } \overline{f = \Lambda \vec{a} . \lambda \vec{x} . L[\vec{u}]} \text{ in } E' [E[\text{let } vb \text{ in } L'[\vec{e}]]]] : \tau \\ \mapsto_B B \left[\begin{array}{c} \text{join rec} \\ j \vec{a} \vec{x} = L[\overline{\text{tail}_\rho(u)}] \\ \text{in } E[\text{let } vb \text{ in } L'[\overline{\text{tail}_\rho(e)}]] \end{array} \right] \\ \text{where } \rho(f \vec{a} \vec{x}) = \text{jump } j \vec{a} \vec{x} \tau \\ E' [E] \in \text{Tail Contexts} \\ \mapsto B \left[\begin{array}{c} \text{let} \\ vb \\ \text{in join rec} \\ j \vec{a} \vec{x} = L[\overline{\text{tail}_\rho(u)}] \\ \text{in } E[L'[\overline{\text{tail}_\rho(e)}]] \end{array} \right] \end{aligned}$$

$$\begin{aligned}
& B \left[\overrightarrow{\text{let rec } f = \Lambda \vec{a} . \lambda \vec{x} . L[\vec{u}]} \text{ in } E'[E[\text{let } vb \text{ in } L'[\vec{e}]]] \right] : \tau \\
& \mapsto^{0/1} B \left[\overrightarrow{\text{let rec } f = \Lambda \vec{a} . \lambda \vec{x} . L[\vec{u}]} \text{ in let } vb \text{ in } E'[E[L'[\vec{e}]]] \right] \\
& =_{\text{letcomm}} \\
& B \left[\overrightarrow{\text{let } vb \text{ in let rec } f = \Lambda \vec{a} . \lambda \vec{x} . L[\vec{u}]} \text{ in } E'[E[L'[\vec{e}]]] \right] \\
& \rightsquigarrow_B B \left[\begin{array}{c} \text{let} \\ vb \\ \text{in join rec} \\ \overrightarrow{j \vec{a} \vec{x} = L[\text{tail}_\rho(u)]} \\ \text{in } E[L'[\text{tail}_\rho(e)]] \end{array} \right] \\
& \text{where } \rho(f \vec{a} \vec{x}) = \text{jump } j \vec{a} \vec{x} \tau \\
& \quad E[E'] \in \text{Tail Contexts}
\end{aligned}$$

For *confitydrop_{rec}* followed by *jump*, we have:

$$\begin{aligned}
& B \left[\overrightarrow{\text{let rec } f = \Lambda \vec{a} . \lambda \vec{x} . L[\vec{u}]} \text{ in } E[L'[\vec{e}, f \vec{\varphi} \vec{v}, \vec{e}']] \right] : \tau \\
& \rightsquigarrow_B B \left[\begin{array}{c} \text{join rec} \\ \overrightarrow{j \vec{a} \vec{x} = L[\text{tail}_\rho(u)]} \\ \text{in } E[L'[\text{tail}_\rho(e), \text{jump } j \vec{\varphi} \vec{v} \tau, \text{tail}_\rho(e)]] \end{array} \right] \\
& \text{where } \rho(f \vec{a} \vec{x}) = \text{jump } j \vec{a} \vec{x} \tau \\
& \quad E \in \text{Tail Contexts} \\
& \mapsto_B \left[\begin{array}{c} \text{let} \\ \vec{x} = \vec{v} \\ \text{in join rec} \\ \overrightarrow{j \vec{a} \vec{x} = L[\text{tail}_\rho(u)]} \\ \text{in } L[\text{tail}_\rho(u)]\{\vec{\varphi}/a\} \end{array} \right] \\
& \text{where } E[L'[\vec{e}, \square, \vec{e}']] \in \text{Evaluation Contexts}
\end{aligned}$$

$$\begin{aligned}
& B \left[\overrightarrow{\text{let rec } f = \Lambda \vec{a} . \lambda \vec{x} . L[\vec{u}]} \text{ in } E[L'[\vec{e}, f \vec{\varphi} \vec{v}, \vec{e}']] \right] : \tau \\
& \mapsto_B \left[\begin{array}{c} \text{let rec} \\ \overrightarrow{f = \Lambda \vec{a} . \lambda \vec{x} . L[\vec{u}]} \\ \text{in } E[L'[\vec{e}, (\Lambda \vec{a} . \lambda \vec{x} . L[\vec{u}]) \vec{\varphi} \vec{v}, \vec{e}']] \end{array} \right] \\
& \text{where } E[L'[\vec{e}, \square, \vec{e}']] \in \text{Evaluation Contexts} \\
& \mapsto^* B \left[\begin{array}{c} \text{let rec} \\ \overrightarrow{f = \Lambda \vec{a} . \lambda \vec{x} . L[\vec{u}]} \\ \text{in } E[L'[\vec{e}, (\lambda \vec{x} . L[\vec{u}])\{\vec{\varphi}/a\}] \vec{v}, \vec{e}']] \end{array} \right] \\
& \mapsto^* =_{\text{letcomm}} B \left[\begin{array}{c} \text{let} \\ \vec{x} = \vec{v} \\ \text{in let rec} \\ \overrightarrow{f = \Lambda \vec{a} . \lambda \vec{x} . L[\vec{u}]}, \\ \text{in } E[L'[\vec{e}, L[\vec{u}]\{\vec{\varphi}/a\}, \vec{e}']] \end{array} \right] \\
& \rightsquigarrow_B B \left[\begin{array}{c} \text{let} \\ \vec{x} = \vec{v} \\ \text{in join} \\ \overrightarrow{j \vec{a} \vec{x} = L[\text{tail}_\rho(u)]}, \\ \text{in } L[\text{tail}_\rho(u)]\{\vec{\varphi}/a\} \end{array} \right] \\
& \text{where } \rho(f \vec{a} \vec{x}) = \text{jump } j \vec{a} \vec{x} \tau
\end{aligned}$$

The two cases involving *confity_{rec}* (*confity_{rec}* followed by *bind*, and *confity_{rec}* followed by *jump*) are special cases of the above two cases involving *confitydrop_{rec}*. \square

The reliance on *letcomm* does not cause us trouble, however, since it is a simple enough reduction that $=_{\text{letcomm}}$ can be postponed after standard reduction.

Lemma 15. *If $e =_{\text{letcomm}} e_1 \mapsto e'$ then $e \mapsto e_2 =_{\text{letcomm}} e'$ for some e_2 .*

Proof. Since $e =_{\text{letcomm}} e_1$, then e and e_1 are the same up to permuted let bindings. The fact that there is an e_2 such that $e \mapsto e_2 =_{\text{letcomm}} e'$ follows by cases on $e_1 \mapsto e$ and permuting let bindings afterward as necessary. \square

Both of these postponements can be combined together to form the (composable) postponement of $=_{\text{letcomm}} \rightsquigarrow_B^{0/1}$ after \mapsto^* .

Lemma 16. *If $e =_{\text{letcomm}} e_1 \rightsquigarrow_B^{0/1} e'_1 \mapsto^* e'$ then $e \mapsto^* e_2 =_{\text{letcomm}} e'_2 \rightsquigarrow_B^{0/1} e'$ for some e_2 and e'_2 .*

Proof. First, consider the simpler case when $e'_1 \mapsto^* e'$. If $e_1 \equiv e'_1$ then $e \mapsto e_2 =_{\text{letcomm}} e' \rightsquigarrow_B^{0/1} e'$ for some e_2 by Lemma 15. Otherwise $e_1 \rightsquigarrow_B e'_1$ so $e =_{\text{letcomm}} e_1 \mapsto^* e_2 =_{\text{letcomm}} e'_2 \rightsquigarrow_B^{0/1} e'$ for some e_2 and e'_2 by Lemma 14, and $e \mapsto^* e_3 =_{\text{letcomm}} e_2 =_{\text{letcomm}} e'_2 \rightsquigarrow_B^{0/1} e'$ for some e_3 by Lemma 15 and induction on $e_1 \mapsto^* e_2$.

Finally, the more general case when $e'_1 \mapsto^* e'$ by any number of steps follows by induction on the standard reduction sequence. \square

Generalizing the above postponement procedure to account for *any* non-standard reduction is difficult because, in general, a non-standard reduction on the right-hand-side of a binding can be duplicated by a standard reduction which copies and inlines that right-hand-side. To counter this complexity we introduce a specific form of *grand reduction* (also known as *parallel reduction*) which allows for many single reduction steps to happen simultaneously. We write $e \rightsquigarrow_i e'$ for the non-standard grand reduction internally within the structure of a term, defined inductively in Fig. 8, $e \rightsquigarrow e'$ for postfixing $e \rightsquigarrow_i e'$ with $=_{\text{letcomm}} \rightsquigarrow_B$, and $e \Rightarrow e'$ for prefixing $e \rightsquigarrow e'$ with zero or more standard reductions. In addition, we use $F \rightsquigarrow_i F'$ for the grand internal non-standard reduction within frames, also defined inductively in Fig. 8, as well as $vb \Rightarrow vb'$, $jb \Rightarrow jb'$, and $alt \Rightarrow alt'$ which are just defined pointwise by allowing $e \Rightarrow e'$ on their immediate subterms.

The grand reduction relation satisfies some important basic properties: \rightsquigarrow is reflexive, compatible, and lies between a single step and multiple steps of \rightsquigarrow .

Lemma 17. *1. $e \rightsquigarrow e$ and $e \rightsquigarrow_i e$.
2. $e \rightsquigarrow e$ implies $C[e] \rightsquigarrow C[e']$.*

$$\begin{array}{c}
\boxed{e \mapsto e', e \rightsquigarrow e', e \rightsquigarrow_i e', F \rightsquigarrow_i F'} \\
\frac{e \mapsto^* e'' \rightsquigarrow e'}{e \mapsto e'} \quad \frac{e \rightsquigarrow_i e'' =_{\text{letcomm}} e''' \succrightarrow_B^{\text{O}1} e'}{e \rightsquigarrow e'} \quad \frac{F \rightsquigarrow_i F' \quad e \rightsquigarrow e'}{F[e] \rightsquigarrow_i F'[e']} \\
\frac{}{x \rightsquigarrow_i x} \quad \frac{}{l \rightsquigarrow_i l} \quad \frac{e \mapsto e'}{\lambda x:\sigma.e \rightsquigarrow_i \lambda x:\sigma.e'} \quad \frac{e \mapsto e'}{\Lambda a.e \rightsquigarrow_i \Lambda a.e'} \quad \frac{u \mapsto u'}{K \vec{\varphi} \vec{u} \rightsquigarrow_i K \vec{\varphi} \vec{u}'} \quad \frac{u \mapsto u'}{\mathbf{jump} \ j \ \vec{\varphi} \ \vec{u} \ \tau \rightsquigarrow_i \mathbf{jump} \ j \ \vec{\varphi} \ \vec{u}' \ \tau} \\
\frac{e \rightsquigarrow e' \quad vb \mapsto vb'}{\mathbf{let} \ vb \ \mathbf{in} \ e \rightsquigarrow_i \mathbf{let} \ vb' \ \mathbf{in} \ e'} \\
\frac{u \mapsto u'}{\square \ u \rightsquigarrow_i \square \ u'} \quad \frac{}{\square \ \varphi \rightsquigarrow_i \square \ \varphi} \quad \frac{\overrightarrow{\text{alt}} \mapsto \overrightarrow{\text{alt}'}}{\mathbf{case} \ \square \ \mathbf{of} \ \overrightarrow{\text{alt}} \rightsquigarrow_i \mathbf{case} \ \square \ \mathbf{of} \ \overrightarrow{\text{alt}'}} \quad \frac{jb \mapsto jb'}{\mathbf{join} \ jb \ \mathbf{in} \ \square \rightsquigarrow_i \mathbf{join} \ jb' \ \mathbf{in} \ \square}
\end{array}$$

Figure 8: Non-standard and grand reduction relations for System F_J .

3. $e \rightsquigarrow e'$ implies $e \rightsquigarrow e'$ implies $e \mapsto^* e'$.

Proof. The first fact follows by induction on \rightsquigarrow . The second fact follows from the first by induction on C . And the third fact follows from the second (to show \succrightarrow is included in \rightsquigarrow) as well as by induction on \rightsquigarrow (to show \rightsquigarrow is included in \mapsto^*). \square

We are now ready to demonstrate the postponement of \rightsquigarrow after \mapsto , which relies on the following ability to read \rightsquigarrow “backwards.”

- Lemma 18.** 1. If $e \rightsquigarrow A'$ then $e \mapsto^* B[A]$ for some B such that $A \rightsquigarrow_i A'$ and $B[u] \rightsquigarrow u'$ for all $u \rightsquigarrow u$.
2. If $e \rightsquigarrow \mathbf{let} \ vb' \ \mathbf{in} \ e'$ then $e \mapsto^* B[\mathbf{let} \ vb \ \mathbf{in} \ E[e]]$ for some B and E such that $\mathbf{let} \ vb \ \mathbf{in} \ e \rightsquigarrow_i \mathbf{let} \ vb' \ \mathbf{in} \ e'$, $B[u] \rightsquigarrow u$ for all $u \rightsquigarrow u'$, and $E[u] \rightsquigarrow u'$ for all $u \rightsquigarrow u'$.
3. If $e \rightsquigarrow E'[\mathbf{jump} \ j \ \vec{\varphi} \ \vec{u}' \ \tau]$ then $e \mapsto^* B[E[\mathbf{jump} \ j \ \vec{\varphi} \ \vec{u} \ \tau]]$ for some B and E such that $\mathbf{jump} \ j \ \vec{\varphi} \ \vec{u} \ \tau \rightsquigarrow_i \mathbf{jump} \ j \ \vec{\varphi} \ \vec{u}' \ \tau$, $B[u] \rightsquigarrow u'$ for all $u \rightsquigarrow u'$, and $j \notin \text{bv}(E)$.
4. If $e \rightsquigarrow S'[x]$ then $e \mapsto^* S[x]$ for some S such that $S[u] \rightsquigarrow S'[u']$ for all $u \rightsquigarrow u'$.

Proof. By induction on $e \rightsquigarrow e'$ in each case. The general pattern is that each \rightsquigarrow expansion can create unreferenced **let** and **join** bindings via the nonstandard *drop* and *jdrops* reductions. Note that for cases 3 and 4, the contexts E and S may additionally differ from E' and S' by use of the nonstandard *casefloat* and *jfloat* reductions, but this cannot bind additional labels in E nor take x out of the eye of S . \square

- Lemma 19.** 1. If $e \rightsquigarrow e_1 \mapsto e'$ then $e \mapsto^* e_2 \rightsquigarrow e'$ for some e_2 .
2. If $e \rightsquigarrow_i e_1 \mapsto e'$ then $e \mapsto^* e_2 \rightsquigarrow_i e'$ for some e_2 .

Proof. By mutual induction on $e \rightsquigarrow e_1$ and $e \rightsquigarrow_i e_1$. The postponement of \rightsquigarrow after \mapsto^* follows immediately from the inductive hypothesis, Lemma 16, and an induction on the \mapsto^* standard reduction sequence that comes

from Lemma 16. Furthermore, none of $x, l, \lambda x:\sigma.e, \Lambda a.e, K \vec{\varphi} \vec{u}$, or $\mathbf{jump} \ j \ \vec{\varphi} \ \vec{u} \ \tau$ have a standard reduction, so the cases for the corresponding rules of \rightsquigarrow_i never happen. The remaining cases are for reductions of let bindings, $\mathbf{let} \ vb \ \mathbf{in} \ e \rightsquigarrow_i \mathbf{let} \ vb_1 \ \mathbf{in} \ e_1$, and reductions within frames, $F[e] \rightsquigarrow_i F_1[e_1]$.

In the case of a let binding, we have

$$\frac{vb \mapsto vb_1 \quad e \rightsquigarrow e_1}{\mathbf{let} \ vb \ \mathbf{in} \ e \rightsquigarrow_i \mathbf{let} \ vb_1 \ \mathbf{in} \ e_1}$$

followed by one of two possible standard reductions:

- $\mathbf{let} \ vb_1 \ \mathbf{in} \ e_1 \mapsto \mathbf{let} \ vb_1 \ \mathbf{in} \ e'$ because $e_1 \mapsto e'$: then by the inductive hypothesis we have an e_2 such that $e \mapsto^* e_2 \rightsquigarrow e'$ and therefore
- $$\mathbf{let} \ vb \ \mathbf{in} \ e \mapsto^* \mathbf{let} \ vb \ \mathbf{in} \ e_2 \quad \frac{vb \mapsto vb_1 \quad e_2 \rightsquigarrow e'}{\mathbf{let} \ vb \ \mathbf{in} \ e_2 \rightsquigarrow_i \mathbf{let} \ vb_1 \ \mathbf{in} \ e'}$$
- $\mathbf{let} \ vb_1 \ \mathbf{in} \ S_1[x] \mapsto \mathbf{let} \ vb_1 \ \mathbf{in} \ S_1[u_1]$ because $e_1 \equiv S_1[x]$ and $x:\sigma = u_1 \in vb_1$: then by Lemma 18 we have $e \mapsto^* S_2[x]$ such that $S_2[v] \rightsquigarrow S_1[v']$ for all $v \rightsquigarrow v'$. Furthermore, because there is a $x:\sigma = u \in vb$ such that $u \mapsto^* u'_1 \rightsquigarrow u_1$ we have $S_2[u] \mapsto^* S_2[u'_1] \rightsquigarrow S_1[u_1]$. Therefore

$$\mathbf{let} \ vb \ \mathbf{in} \ e \mapsto^* \mathbf{let} \ vb \ \mathbf{in} \ S_2[x] \mapsto \mathbf{let} \ vb \ \mathbf{in} \ S_2[u] \mapsto^* \mathbf{let} \ vb \ \mathbf{in} \ S_2[u'_1]$$

$$\frac{vb \mapsto vb_1 \quad S_2[u'_1] \rightsquigarrow S_1[u_1]}{\mathbf{let} \ vb \ \mathbf{in} \ S_2[u'_1] \rightsquigarrow_i \mathbf{let} \ vb_1 \ \mathbf{in} \ S_1[u_1]}$$

In the case of a frame, we have

$$\frac{F \rightsquigarrow_i F_1 \quad e \rightsquigarrow e_1}{F[e] \rightsquigarrow F_1[e_1]}$$

followed by one of the following possible standard reductions:

- $F_1[e_1] \mapsto F_1[e']$ because $e_1 \mapsto e'$: then by the inductive hypothesis we have an e_2 such that $e \mapsto^* e_2 \Rightarrow e'$ and therefore

$$F[e] \mapsto^* F[e_2] \quad \frac{F \Rightarrow_i F_1 \quad e_2 \Rightarrow e'}{F[e_2] \Rightarrow F_1[e']}$$

- $(\lambda x:\sigma.v_1) u_1 \mapsto \mathbf{let} x:\sigma = u_1 \mathbf{in} v_1$ because $e_1 \equiv (\lambda x:\sigma.v_1)$ and $F_1 \equiv \square u_1$: since $F \Rightarrow_i \square u_1$ we know that $F \equiv \square u$ such that $u \mapsto u_1$. Furthermore, since $e \Rightarrow \lambda x:\sigma.v_1$, we know that $e \mapsto^* B[\lambda x:\sigma.v]$ such that $v \mapsto^* v'_1 \Rightarrow v_1$, and $B[u] \Rightarrow u'$ for all $u \Rightarrow u'$ by Lemma 18. Therefore

$$F[e] \mapsto^* B[\lambda x:\sigma.v] u \mapsto^* B[(\lambda x:\sigma.v) u] \\ \mapsto B[\mathbf{let} x:\sigma = u \mathbf{in} v] \mapsto^* B[\mathbf{let} x:\sigma = u \mathbf{in} v'_1]$$

$$\frac{x:\sigma = u \mapsto x:\sigma = u_1 \quad v'_1 \Rightarrow v_1}{\mathbf{let} x:\sigma = u \mathbf{in} v'_1 \Rightarrow \mathbf{let} x:\sigma = u_1 \mathbf{in} v_1} \\ B[\mathbf{let} x:\sigma = u \mathbf{in} v'_1] \Rightarrow \mathbf{let} x:\sigma = u_1 \mathbf{in} v_1$$

- $(\Lambda a.e'_1) \varphi \mapsto v_1\{\varphi/a\}$ because $e_1 \equiv (\Lambda a.v_1)$ and $F_1 \equiv \square \varphi$: similar to the previous case.
- case $K \vec{\varphi} \vec{v}_1$ of $\vec{alt}_1 \mapsto \mathbf{let} \vec{x}:\vec{\sigma} = \vec{v}_1 \mathbf{in} u_1$ because $e_1 \equiv K \vec{\varphi} \vec{v}_1$, $F_1 \equiv \mathbf{case} \square \mathbf{of} \vec{alt}_1$, and $(K \vec{a}, \vec{x}:\vec{\sigma} \rightarrow u_1) \in \vec{alt}_1$: since $F \Rightarrow_i \mathbf{case} \square \mathbf{of} \vec{alt}_1$ we know that $F \equiv \mathbf{case} \square \mathbf{of} \vec{alt}$ such that $\vec{alt} \mapsto^* \vec{alt}'_1 \Rightarrow \vec{alt}_1$. Furthermore, since $e \Rightarrow K \vec{\varphi} \vec{v}_1$ we know that $e \mapsto^* B[K \vec{\varphi} \vec{v}]$ such that $\vec{v} \mapsto^* \vec{v}'_1$, and $B[u] \Rightarrow u'$ for all $u \Rightarrow u'$ by Lemma 18. Therefore

$$F[e] \mapsto^* \mathbf{case} B[K \vec{\varphi} \vec{v}] \mathbf{of} \vec{alt} \mapsto^* B[\mathbf{case} K \vec{\varphi} \vec{v} \mathbf{of} \vec{alt}] \\ \mapsto B[\mathbf{let} \vec{x}:\vec{\sigma} = \vec{v} \mathbf{in} u] \mapsto^* B[\mathbf{let} \vec{x}:\vec{\sigma} = \vec{v} \mathbf{in} u'_1]$$

$$\frac{\vec{x}:\vec{\sigma} = v \mapsto \vec{x}:\vec{\sigma} = \vec{v}_1 \quad u'_1 \Rightarrow u_1}{\mathbf{let} \vec{x}:\vec{\sigma} = \vec{v} \mathbf{in} u'_1 \Rightarrow \mathbf{let} \vec{x}:\vec{\sigma} = \vec{v}_1 \mathbf{in} u_1} \\ B[\mathbf{let} \vec{x}:\vec{\sigma} = \vec{v} \mathbf{in} u'_1] \Rightarrow \mathbf{let} \vec{x}:\vec{\sigma} = \vec{v}_1 \mathbf{in} u_1$$

- $\mathbf{join} jb_1 \mathbf{in} E_1[\mathbf{jump} j \vec{\varphi} \vec{v}_1 \tau]$
 $\mapsto \mathbf{let} \vec{x}:\vec{\sigma} = \vec{v}_1 \mathbf{in} \mathbf{join} jb_1 \mathbf{in} u_1\{\varphi/a\}$
 because $e_1 \equiv E_1[\mathbf{jump} j \vec{\varphi} \vec{v}_1 \tau]$, $F_1 \equiv \mathbf{join} jb_1 \mathbf{in} \square$, and $(j \vec{a} x:\sigma = u_1) \in jb_1$: since $F \Rightarrow_i \mathbf{join} jb_1 \mathbf{in} \square$ we know that $F \equiv \mathbf{join} jb \mathbf{in} \square$ such that $jb \mapsto^* jb'_1 \Rightarrow jb_1$. Furthermore, since $e \Rightarrow E_1[\mathbf{jump} j \vec{\varphi} \vec{v}_1 \tau]$ we know $e \mapsto^* B_2[E_2[\mathbf{jump} j \vec{\varphi} \vec{v} \tau]]$ such that $j \notin \text{bv} E_2$, $\vec{v} \mapsto^* \vec{v}'_1$, and $B_2[u] \Rightarrow u'$ for all $u \Rightarrow u'$ by Lemma 18. Therefore

$$F[e] \mapsto^* \mathbf{join} jb_1 \mathbf{in} B_2[E_2[\mathbf{jump} j \vec{\varphi} \vec{v} \tau]] \\ \mapsto^* B_2[\mathbf{join} jb_1 \mathbf{in} E_2[\mathbf{jump} j \vec{\varphi} \vec{v} \tau]] \\ \mapsto B_2[\mathbf{let} \vec{x}:\vec{\sigma} = \vec{v} \mathbf{in} \mathbf{join} jb \mathbf{in} u\{\varphi/a\}] \\ \mapsto^* B_2[\mathbf{let} \vec{x}:\vec{\sigma} = \vec{v} \mathbf{in} \mathbf{join} jb \mathbf{in} u'_1\{\varphi/a\}]$$

$$\frac{jb_1 \mapsto jb_1 \quad u'_1\{\varphi/a\} \Rightarrow u_1\{\varphi/a\}}{\mathbf{join} jb \mathbf{in} u'_1\{\varphi/a\} \Rightarrow \mathbf{join} jb_1 \mathbf{in} u_1\{\varphi/a\}} \\ \frac{x:\sigma = v \mapsto x:\sigma = v_1}{\mathbf{let} \vec{x}:\vec{\sigma} = \vec{v} \mathbf{in} \mathbf{join} jb \mathbf{in} u'_1\{\varphi/a\} \Rightarrow \mathbf{let} \vec{x}:\vec{\sigma} = \vec{v}_1 \mathbf{in} \mathbf{join} jb_1 \mathbf{in} u_1\{\varphi/a\}} \\ B_2[\mathbf{let} \vec{x}:\vec{\sigma} = \vec{v} \mathbf{in} \mathbf{join} jb \mathbf{in} u'_1\{\varphi/a\}] \Rightarrow \mathbf{let} \vec{x}:\vec{\sigma} = \vec{v}_1 \mathbf{in} \mathbf{join} jb_1 \mathbf{in} u_1\{\varphi/a\}]$$

- $\mathbf{join} jb_1 \mathbf{in} A_1 \mapsto A_1$ because $e_1 \equiv A_1$ and $F_1 \equiv \mathbf{join} jb_1 \mathbf{in} \square$: since $F \Rightarrow_i \mathbf{join} jb_1 \mathbf{in} \square$ we know that $F \equiv \mathbf{join} jb \mathbf{in} \square$ such that $jb \mapsto jb_1$. Furthermore, since $e \Rightarrow A_1$ we know $e \mapsto^* B[A]$ such that $A \Rightarrow_i A'$, and $B[u] \Rightarrow u'$ for all $u \Rightarrow u'$ by Lemma 18. Therefore

$$F[e] \mapsto^* \mathbf{join} jb \mathbf{in} B[A] \mapsto^* B[\mathbf{join} jb \mathbf{in} A] \mapsto B[A]$$

$$\frac{A \Rightarrow A_1}{B[A] \Rightarrow A_1}$$

- $F_1[\mathbf{let} vb_1 \mathbf{in} v_1] \mapsto \mathbf{let} vb_1 \mathbf{in} F_1[v_1]$ because $e_1 \equiv \mathbf{let} vb_1 \mathbf{in} v_1$: since $e \Rightarrow \mathbf{let} vb_1 \mathbf{in} v_1$ we know that $e \mapsto^* B_2[\mathbf{let} vb \mathbf{in} E_2[v]]$ such that $vb \Rightarrow vb_1$, $v \Rightarrow v_1$, and $B_2[u] \Rightarrow u'$ for all $u \Rightarrow u'$, and $E_2[u] \Rightarrow u'$ for all $u \Rightarrow u'$ by Lemma 18. Therefore

$$F[e] \mapsto^* F[B_2[\mathbf{let} vb \mathbf{in} E_2[v]]] \mapsto^* B_2[\mathbf{let} vb \mathbf{in} F[E_2[v]]]$$

$$\frac{v \Rightarrow v_1}{F \Rightarrow_i F_1 \quad E_2[v] \Rightarrow v_1} \\ \frac{vb \mapsto vb_1 \quad F[E_2[v]] \Rightarrow_i F_1[v_1]}{\mathbf{let} vb \mathbf{in} F[E_2[v]] \Rightarrow_i \mathbf{let} vb_1 \mathbf{in} F_1[v_1]} \\ B_2[\mathbf{let} vb \mathbf{in} F[E_2[v]]] \Rightarrow \mathbf{let} vb_1 \mathbf{in} F_1[v_1]$$

□

From the postponement of a single \Rightarrow after \mapsto , we can easily derive the postponement of the many-step \mapsto^* after \mapsto^* .

Lemma 20. *If $e \mapsto^* e_1 \mapsto^* e'$ then $e \mapsto^* e_2 \mapsto^* e'$ for some e_2 .*

Proof. First, note that Lemma 19 can be generalized to the fact that if $e \Rightarrow e_1 \mapsto^* e'$ then $e \mapsto^* e_2 \Rightarrow e'$ for some e_2 by induction on $e_1 \mapsto^* e'$.

Now, since \mapsto implies \Rightarrow (Lemma 17), we have $e \Rightarrow^* e_1 \mapsto^* e'$. By the above generalization of Lemma 19 and induction on $e \Rightarrow^* e_1$, we have $e \mapsto^* e_2 \Rightarrow^* e'$ for some e_2 . Finally, since \Rightarrow implies \mapsto^* , we have $e \mapsto^* e_2 \mapsto^* e'$. □

This final many-step postponement is powerful enough to let us sort between standard and non-standard reductions in a general reduction sequence, putting the standard ones first and thus achieving the standardization procedure.

Theorem 21 (Standardization). *If $e \rightarrow^* e' \not\vdash^*$ then $e \mapsto^* e'' \rightarrow^* e'$ for some $e' \not\vdash^*$.*

Proof. Note that every \rightarrow is either \mapsto or \succmapsto by definition. Therefore, we can proceed by induction on the reduction sequence $e \rightarrow^* e'$ to show that $e \mapsto^* e'' \succmapsto^* e'$ for some e'' :

- If $e \equiv e'$ then the result is immediate.
- If $e \mapsto e''' \rightarrow^* e'$ then we have $e \mapsto e''' \mapsto^* e'' \succmapsto^* e'$ for some e'' by the inductive hypothesis.
- If $e \succmapsto e''' \rightarrow^* e'$ then we have $e \succmapsto e''' \mapsto^* e'' \succmapsto^* e'$ for some e'' by the inductive hypothesis, and so $e \mapsto e'''' \mapsto^* e'' \succmapsto^* e'$ for some e'''' by Lemma 20.


Finally, because results are preserved by \succmapsto expansion (Lemma 18), we have that $e' \not\vdash^*$ implies that $e'' \not\vdash^*$ by induction on $e'' \succmapsto^* e'$. \square

With both confluence and standardization at hand, we can now justify the correctness of the equational axioms.

Theorem 22 (Correctness). *If $e = e'$ then $e \cong e'$.*

Proof. Let C be any context, and (without loss of generality) suppose that $C[e]$ converges to $C[e] \mapsto^* e_1 \not\vdash^*$. Since $e = e'$, we have $C[e] = C[e']$ by the compatibility of $=$. We have

that $C[e] \rightarrow^* e_2 \leftarrow^* C[e']$ for some e_2 by the Church-Rosser property (which is known to follow from confluence (Theorem 12) by induction on $=$ as the reflexive-transitive closure of $\leftarrow^* \rightarrow^*$). Since $e_2 \leftarrow^* C[e] \mapsto^* e_1 \not\vdash^*$, we have that $e_2 \rightarrow^* e_3 \leftarrow^* e_1$ by confluence (since \mapsto is included in \rightarrow^*) for some e_3 , where $e_3 \mapsto_i^* e_4 \not\vdash^*$ for some e_4 by Lemma 13 since $e_1 \not\vdash^*$. Putting it all together, we get $C[e'] \rightarrow^* e_2 \rightarrow^* e_3 \mapsto_i^* e_4 \not\vdash^*$. Therefore, by standardization (Theorem 21) we have an e_5 such that $C[e'] \mapsto^* e_5 \not\vdash^*$. Summarizing, Theorem 12, Theorem 21, and Lemma 13 together produce the following diagram:



In conclusion, if $e = e'$ then for all C , $C[e]$ converges if and only if $C[e']$ does, which means that $C[e]$ diverges if and only if $C[e']$ does. By Lemma 11, this means that $e = e'$ implies $e \cong e'$. \square