Sequent Calculus as a Compiler Intermediate Language

Paul Downen Luke Maurer
Zena M. Ariola

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

Simon Peyton Jones

Microsoft Research Cambridge, UK
simonpj@microsoft.com

Abstract

The λ -calculus is popular as an intermediate language for practical compilers. But in the world of logic it has a lesser-known twin, born at the same time, called the *sequent calculus*. Perhaps that would make for a good intermediate language, too? To explore this question we designed Sequent Core, a practically-oriented core calculus based on the sequent calculus, and used it to re-implement a substantial chunk of the Glasgow Haskell Compiler.

Categories and Subject Descriptors D.3.4 [Programming Languages]: Processors—Compilers

Keywords Intermediate representations; Natural deduction; Sequent calculus; Compiler optimizations; Continuations; Haskell

1. Introduction

Steele and Sussman’s “Lambda the ultimate” papers [41, 42] persuasively argued that the λ -calculus is far more than a theoretical model of computation: it is an incredibly expressive and practical intermediate language for a compiler. The Rabbit compiler [40], its successors (e.g. Orbit [21]), and Appel’s book “Compiling with continuations” [1] all demonstrate the power and utility of the λ -calculus as a compiler’s intermediate language.

The typed λ -calculus arises canonically as the term language for a logic called *natural deduction* [14], using the *Curry-Howard isomorphism* [45]: the pervasive connection between logic and programming languages asserting that propositions are types and proofs are programs. Indeed, for many people, the λ -calculus is the living embodiment of Curry-Howard.

But natural deduction is not the only logic! Conspicuously, natural deduction has a twin, born in the very same paper [14], called the *sequent calculus*. Thanks to the Curry-Howard isomorphism, terms of the sequent calculus can also be seen as a programming language [9, 15, 44] with an emphasis on *control flow*.

This raises an obvious question: *does the language of the sequent calculus have merit as a practical compiler intermediate language, in the same way that the λ -calculus does?* What advantages and disadvantages might it have, compared to the existing λ -based technology? Curiously, we seem to be the first to address these

questions, and surprisingly the task was not as routine as we had expected. Specifically, our contributions are these:

- We describe a typed sequent calculus called Sequent Core with the same expressiveness as System $F\omega$, including `let`, algebraic data types, and `case` (Section 2).

The broad outline of the language is determined by the logic, but we made numerous choices driven by its role as a compiler intermediate representation (Section 2.2).

- Our language comes equipped with an operational semantics (Section 2.3), a type system (Section 2.4), and standard meta-theoretical properties. We also give direct-style translations to and from System $F\omega$ (Section 3).¹
- The proof of the pudding is in the eating. We have implemented our intermediate language as a plugin² for GHC, a state-of-the-art optimizing compiler for Haskell (Section 4). GHC’s intermediate language, called Core, is essentially System $F\omega$; our new plugin translates Core programs into Sequent Core, optimizes them, and translates them back. Moreover, we have re-implemented some of GHC’s Core-to-Core optimization passes, notably the simplifier, to instead use Sequent Core.
- From the implementation, we found a way that Sequent Core was qualitatively better than Core for optimization: the treatment of *join points*. Specifically, join points in Sequent Core are preserved during simplifications such as the ubiquitous case-of-case transformation (Sections 4.2 and 4.3). Further, we show how to recover the join points of Sequent Core programs, after they are lost in translation, using a lightweight version of a process known as *contification* [20] (Section 5).

So what kind of intermediate language do we get out of the sequent calculus? It turns out that the language resembles *continuation-passing style*, a common technique in the λ -calculus for representing control flow inside a program. The division between assumptions and conclusions in the logic gives us a divide between programs that yield results and continuations that observe those results in the language. Yet despite the surface similarity, Sequent Core is still quite different from continuation-passing style (Section 6). Perhaps most importantly, Sequent Core brings control flow and continuations to a compiler like GHC without stepping on its toes, allowing its extensive direct-style optimizations to still shine through. In the end, we get an intermediate language that lies somewhere in between direct and continuation-passing styles (Section 7), sharing some advantages of both.

In a sense, many of the basic ideas we present here have been re-discovered over the years as the tradition of Curry-Howard

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ICFP’16, September 18–24, 2016, Nara, Japan
ACM. 978-1-4503-4219-3/16/09...\$15.00
<http://dx.doi.org/10.1145/2951913.2951931>

¹ More details of the meta-theory can be found in the appendix: http://ix.cs.uoregon.edu/~pdownen/publications/scfp_ext.pdf

² Available at: <http://github.com/lukemaurer/sequent-core>

Shared syntax of kinds and types	
$a, b, c \in TyVar ::= \dots$	
$\kappa \in Kind ::= \star \mid \kappa \rightarrow \kappa$	
$\tau, \sigma \in Type ::= a \mid T \mid \sigma \tau \mid \sigma \rightarrow \tau \mid \forall a:\kappa.\tau \mid \exists \bar{a}:\bar{\kappa}.\langle \vec{\tau} \rangle$	
Syntax of Sequent Core	
$x, y, z \in Var ::= \dots$	$j \in Label ::= \dots$
$pgm \in Program ::= \overrightarrow{bind}$	
$bind \in Bind ::= bp \mid \mathbf{rec} \left\{ \overrightarrow{bp} \right\}$	
$bp \in BindPair ::= x:\tau = v \mid j:\tau = \tilde{\mu}[\bar{a}:\bar{\kappa}, \bar{x}:\bar{\sigma}].c$	
$v \in Term ::= \lambda x:\tau.v \mid \Lambda a:\kappa.v \mid x \mid K(\vec{\sigma}, \vec{v}) \mid \mu \mathbf{ret}.c$	
$c \in Command ::= \mathbf{let} \text{ bind } \mathbf{in} c \mid \langle v \parallel k \rangle \mid \mathbf{jump} j \vec{\sigma} \vec{v}$	
$k \in Kont ::= v \cdot k \mid \tau \cdot k \mid \mathbf{case} \text{ of } \overrightarrow{alt} \mid \mathbf{ret}$	
$alt \in Alternative ::= x:\tau \rightarrow c \mid K(\bar{a}:\bar{\kappa}, \bar{x}:\bar{\tau}) \rightarrow c$	
Syntax of Core	
$x, y, z \in Var ::= \dots$	
$pgm \in Program ::= \overrightarrow{bind}$	
$bind \in Bind ::= bp \mid \mathbf{rec} \left\{ \overrightarrow{bp} \right\}$	
$bp \in BindPair ::= x:\tau = v$	
$e \in Expression ::= \mathbf{let} \text{ bind } \mathbf{in} e$	
	$\mid \lambda x:\tau.e \mid \Lambda a:\kappa.e \mid x \mid K \vec{\sigma} \vec{e}$
	$\mid e e \mid e \tau \mid \mathbf{case} e \text{ of } \overrightarrow{alt}$
$alt \in Alternative ::= x:\tau \rightarrow e \mid K \bar{a}:\bar{\kappa} \bar{x}:\bar{\tau} \rightarrow e$	

Figure 1. Syntax

dictates [45]: first by a logician and later by computer scientists. Our goal is to put them together in a way that is useful for compilers. Our implementation demonstrates that Sequent Core is certainly up to the job: in short order, we achieved performance competitive with a highly mature optimizing compiler. While we are not ready to recommend that GHC change its intermediate language, we instead see Sequent Core as an illuminating new design point in the space of functional programming languages and laboratory for experiments on intermediate representation techniques. We hope that our work will bring the sequent calculus forth, Cinderella-like, out of the theory kitchen and into the arms of compiler writers.

2. Sequent Core

In this section we present the specifics of our new sequent-style intermediate language for functional programs, along with its type system and operational semantics. The language that comes out of the logic “for free” is more expressive [11] than the pure λ -calculus, since it naturally speaks of control flow as a first-class entity. Thus, our task is to find the sweet spot between the permission to express interesting control flow and the restriction to pure functions.

2.1 Overview

Figure 1 gives the syntax of Sequent Core. For comparison purposes, we also give the syntax of Core, GHC’s current intermediate language [38]. Both languages share the same syntax of types and kinds, also given in Figure 1. We omit two important features of

Core, namely casts and unboxed types; both are readily accommodated in Sequent Core, and our implementation does so, but they distract from our main point.

Here is a small program written in both representations:

Core

$plusOne : Int \rightarrow Int = \lambda x: Int. (+) x 1$

Sequent Core

$plusOne : Int \rightarrow Int = \lambda x: Int. \mu \mathbf{ret}. \langle (+) \parallel x \cdot 1 \cdot \mathbf{ret} \rangle$

Referring to Figure 1, we see that

- Just as in Core, a Sequent Core *program* is a set of top-level *bindings*; bindings can be non-recursive or (mutually) recursive.
- A binding in Sequent Core is either a *value binding* $x:\tau = v$, or a *continuation binding* $j:\tau = \tilde{\mu}[\bar{a}:\bar{\kappa}, \bar{x}:\bar{\sigma}].c$.³ We discuss the latter in Section 2.2.2.
- The right-hand side of a value binding is a *term* v , which begins with zero or more lambdas followed by a variable, a constructor application, or a *computation* $\mu \mathbf{ret}.c$.
- The computation term $\mu \mathbf{ret}.c$, where c is a *command*, means “run command c and return whatever is passed to $\mathbf{ret}$ as the result of this term.”
- All the interesting work gets done by commands. A command c is a collection of local bindings, wrapping either a *cut pair* $\langle v \parallel k \rangle$ or a *jump*. We discuss jumps in Section 2.2.2.
- Finally, cut pairs do some real computation. They are pairs $\langle v \parallel k \rangle$ of a term v and a *continuation* k . A continuation k is a *call stack* containing a sequence of applications (to types or terms) ending with either a *case* analysis or a return to $\mathbf{ret}$.

In *plusOne* as written in Sequent Core above, the calculation of the function is carried out by the command in its body: the term is the function $(+)$, while the continuation is $x \cdot 1 \cdot \mathbf{ret}$, meaning “apply to x , then apply to 1, then return.” With this reading, Sequent Core cut pairs closely resemble the *states* of many abstract machines (e.g., the CEK machine [12]), with a term v in the focus and a continuation or call stack k that describes how it is consumed.

Here is another example program, written in both representations, that drops the last element of a list:

Core

$init : \forall a.[a] \rightarrow [a]$
 $= \Lambda a. \lambda x:[a]. \mathbf{case} \text{ reverse } a \text{ xs of}$
 $\quad [] \rightarrow []$
 $\quad (y : ys) \rightarrow \text{reverse } a \text{ xs}$

Sequent Core

$init : \forall a.[a] \rightarrow [a]$
 $= \Lambda a. \lambda x:[a]. \mu \mathbf{ret}. \langle \text{reverse} \parallel a \cdot \text{xs} \cdot \mathbf{case} \text{ of}$
 $\quad [] \rightarrow []$
 $\quad (y : ys) \rightarrow \langle \text{reverse} \parallel a \cdot \text{ys} \cdot \mathbf{ret} \rangle$

As before, the outer structure is the same, but the *case*, which is so prominent in Core, appears in Sequent Core as the continuation of the call to *reverse*. Indeed, this highlights a key difference: in Sequent Core, the focus of evaluation is always “at the top”, whereas in Core it may be deeply buried [4]. In this example, the call to *reverse* is the first thing to happen, and it is visibly at the top of the body of the lambda. In this way, Sequent Core’s operational reading is somewhat more direct, a useful property for a compiler intermediate language.

³ We occasionally omit type annotations in examples for clarity.

2.2 The Language

Having seen how Sequent Core is a language resembling an abstract machine, let's look more closely at the new linguistic concepts that it introduces and how Sequent Core compares to Core. On closer inspection, Sequent Core can be seen as a nuanced variation on Core, separating the roles of distinct concepts of Core syntactically as part of the effort to split calculations across the two sides of a cut pair. More specifically, each construct of Core has an exact analogue in Sequent Core, but the single grammar of Core expressions e is divided among terms v , continuations k , and commands c in Sequent Core. Additionally, Sequent Core has special support for labels and direct jumps, which are not found in Core.

2.2.1 Terms and Continuations

Core expressions e , as shown in Figure 1, include a variety of values (more specifically *weak-head normal forms*) which require no further evaluation: lambdas (both small λ and big Λ) and applied constructors. Along with variables, these are all terms in Sequent Core, as they do not involve any work to be done and they immediately produce themselves as their result.

On the other hand, Core also includes expressions which *do* require evaluation: function applications $e\ e'$, polymorphic instantiations $e\ \tau$, and **case** expressions. Each of these expressions *uses* something to create the next result, and thus these are reflected as continuations k in Sequent Core. As usual, Sequent Core continuations represent evaluation contexts that receive an input which will be used immediately. For example, the application context $\square\ 1$, where “ $\square$ ” is the hole where the input is placed, corresponds to the call stack $1 \cdot \text{ret}$. Furthermore, we can apply the curried function $\lambda x.\lambda y.x$ to the arguments 1 and 2 by running it in concert with the stack $1 \cdot 2 \cdot \text{ret}$, as in:

$$\langle \lambda x.\lambda y.x \parallel 1 \cdot 2 \cdot \text{ret} \rangle = \langle \lambda y.1 \parallel 2 \cdot \text{ret} \rangle = \langle 1 \parallel \text{ret} \rangle$$

where ret signals a stop, so that the result 1 can be returned.

Since we are interested in modeling *lazy* functional languages, we also need to include the results of arbitrary deferred computations as terms in themselves. For example, when we perform the lazy function composition $f\ (g\ x)$ in Core, $g\ x$ is only computed when f demands it. This means we need the ability to inject computations into terms, which we achieve with μ -abstractions. A μ -abstraction extracts a result from a command by *binding* the occurrences of ret in that command, so that anything passed to ret is returned from the μ -abstraction. However, because we are only modeling purely functional programs, there is only ever one ret available at a time, making it a rather limited namespace. Thus, $\mu \text{ret}. \langle g \parallel x \cdot \text{ret} \rangle$ runs the underlying command, calling the function g with the argument x , so that whatever is returned by g pops out as the result of the term. So lazy function composition can be written in Sequent Core as $\langle f \parallel (\mu \text{ret}. \langle g \parallel x \cdot \text{ret} \rangle) \cdot \text{ret} \rangle$.

Notice that every closed command must give a result to ret if it ever stops at all. Another way of looking at this fact is that every (finite) continuation has ret “in the tail”; it plays the role of “ nil ” in a linked list. However, the return structure of continuations is more complex than a plain linked list, since the terminating ret of a continuation may occur in several places. By inspection, a continuation is a sequence of zero or more type or term applications, followed by either ret itself or by a **case** continuation. But in the latter case, each alternative has a command whose continuation must in turn have ret in the tail. Unfortunately, this analogy breaks down in the presence of local bindings, as we will see. Luckily, however, viewing ret as a static variable bound by μ -abstractions tells us exactly how to “chase the tail” of a continuation by following the normal rules of static scope. So we may still say that every closed computation $\langle v \parallel k \rangle$ eventually returns if it does not diverge.

2.2.2 Bindings and Jumps

There is one remaining Core expression to be sorted into the Sequent Core grammar: **let** bindings. In Sequent Core, **let** bindings are commands, as they set up an enclosing environment for another command to run in, forming an executable code block. In both representations, **let** bindings serve two purposes: to give a shared name to the result of some computation, and to express (mutual) recursion. Thus in the Sequent Core command c , we can share the results of terms through **let** $x = v$ **in** c and we can share a continuation through $\langle \mu \text{ret}.c \parallel k \rangle$. But something is missing. How can we give a shared label to a command (i.e., to a block of code) that we can go to during the execution of another command? This facility is critical for maintaining small code size, so that we are not forced to repeat the same command verbatim in a program.

For example, suppose we have the command

$$\langle z \parallel \text{case of } \text{Left}(x) \rightarrow c, \text{Right}(x) \rightarrow c \rangle$$

wherein the same c is repeated twice due to the **case** continuation. Now, how do we lift out and give a name to c , given that it contains the free variable x ? We would rather not use a lambda, as in $\lambda x.\mu \text{ret}.c$, since that introduces additional overhead compared to the original command. Instead, we would rather think of c as a sort of continuation whose input is named x during execution of c . In the syntax of $\bar{\lambda}\mu\tilde{\mu}$ [9], this would be written as $\tilde{\mu}x.c$, the *dual* of μ -abstractions. However, this is not like the other continuations we have seen so far! There is no guarantee that $\tilde{\mu}x.c$ uses its input immediately, or even at all. Thus, we are not dealing with an evaluation context, but rather an arbitrary context. Furthermore, we might (reasonably) want to name commands with multiple free variables, or even free type variables. So in actuality, we are looking for a representation of continuations taking *multiple values* as inputs of *polymorphic* types, corresponding to general contexts with multiple holes.

This need leads us to multiple-input continuations, which we write as $\tilde{\mu}[a_1, \dots, a_n, x_1, \dots, x_m].c$ in the style of $\bar{\lambda}\mu\tilde{\mu}$. These continuations accept several inputs (named $x_1 \dots x_m$), whose types are polymorphic over the choice of types for $a_1 \dots a_n$, in order to run a command c . Intuitively, we may also think of these multiple-input continuations as a sequence of lambdas $\Lambda a_1 \dots a_n.\lambda x_1 \dots x_m.c$, except that the body is a command because it does not return. The purpose of introducing multiple-input continuations was to lift out and name arbitrary commands, and so they appear as a Sequent Core binding. Specifically, all multiple-input continuations in Sequent Core are given a label j , as in the continuation binding $j = \tilde{\mu}[x, y]. \langle (+) \parallel x \cdot y \cdot \text{ret} \rangle$. These labeled continuations serve as *join points*: places where the control flow of several diverging branches of a program joins back up again.

In order to invoke a bound continuation, we can *jump* to it by providing the correct number of terms for the inputs, as well as explicitly specifying the instantiation of any polymorphic type in System $F\omega$ style. For example, the command

$$\begin{aligned} &\text{let } j = \tilde{\mu}[a:\star, x:a, f:a \rightarrow \text{Bool}]. \langle f \parallel x \cdot \text{ret} \rangle \\ &\text{in jump } j\ \text{Bool}\ \text{True}\ \text{not} \end{aligned}$$

will jump to the label j with the inputs *Bool*, *True*, and *not*, which results in $\langle \text{not} \parallel \text{True} \cdot \text{ret} \rangle$. So when viewing Sequent Core from the perspective of an abstract machine, its command language provides three instructions: (1) set a binding with **let**, (2) evaluate an expression with a cut pair, or (3) perform a direct jump.

Take note that a labeled continuation does not introduce a μ -binder. As a consequence, the occurrence of ret found in $j = \tilde{\mu}[x, y]. \langle (+) \parallel x \cdot y \cdot \text{ret} \rangle$ refers to the nearest surrounding μ , unlike the ret found in $f = \lambda x.\lambda y.\mu \text{ret}. \langle (+) \parallel x \cdot y \cdot \text{ret} \rangle$. Viewing ret as a statically bound variable means that labeled continuations participate in the “tail chasing” discussed previously in Section 2.2.1.

Thus, the `ret` structure of commands and continuations treats labeled continuations quite the same as `case` alternatives for free.

2.2.3 The Scope of Labels

There is one major restriction that we enforce to ensure that each term must have a *unique* exit point by which it returns its result, and so evaluating a term cannot cause an observable jump to some surrounding continuation binding. The intuition is:

Terms contain no free references to continuation variables,

where continuation variables can be labels j as well as `ret`. This restriction, similar to a restriction on CPS [20], makes sure that lambdas cannot close over labels available from their contexts, so that labels do not escape through returned lambdas. Thus, all jumps within the body of a lambda must be local. Likewise, in all computations $\mu\text{ret}.c$, the underlying command c has precisely one unique exit point from which the computation can return a result, denoted by `ret`. Therefore, all jumps made during the execution of c are internal to c , and unobservable during evaluation of $\mu\text{ret}.c$.

Notice that this restriction on the scope of continuation variables, while not very complex, still manages to tell us something about the expressive capabilities of Sequent Core. For example, we syntactically permit value and continuation bindings within the same recursive block, but can they mutually call one another? It turns out that the scoping restriction disallows any sort of interesting mutual recursion between terms and continuations, because terms are *prevented* from referencing labels within their surrounding (or same) binding environment. Specifically, there is some additional structure implicit to `let` bindings:

- Continuation bindings can reference value bindings and other continuation bindings, but value bindings can only reference other value bindings.
- In any sequence of bindings, all value bindings can always be placed before all continuation bindings.
- Value and continuation bindings cannot be mutually recursive. Any minimal, mutually recursive `rec {  $\vec{bp}$  }` block will consist of only value bindings or only continuation bindings.

For example, consider the recursive bindings:

$$\text{rec } \{f = \lambda x.v, j = \tilde{\mu}[y].c\}$$

By the scoping rules, j may call f through c , but f cannot jump back to j in v because $\lambda x.v$ cannot contain a free reference to j . Therefore, since there is no true mutual recursion between f and j , we can break the recursive bindings into two separate blocks with the correct scope, placing the binding for f first:

$$\text{rec } \{f = \lambda x.v\}, \text{rec } \{j = \tilde{\mu}[y].c\}$$

While we do not syntactically enforce this separation, doing so would not cause any loss of expressiveness. Indeed, we could normalize all commands by gathering and partitioning all bindings into (1) first, the list of value bindings, Γ , and (2) second, the list of continuation bindings, Δ , so that commands have the form `let  $\Gamma$  in let  $\Delta$  in  $\langle v \parallel k \rangle$` . However, we do not enforce this normal form in Sequent Core.

2.3 Operational Semantics

A useful way to understand Sequent Core is through its operational semantics, given in Figure 2, which provides a high-level specification for reasoning about the correctness of program transformations. The rules for lambda (both small λ and big Λ) are self-explanatory. The rules for `case` are disambiguated by selecting the first match, so the order of alternatives matters. For a non-recursive `let`, we simply substitute, thus implementing call-by-name; implementing recursive

$$\begin{aligned} W \in WHNF &::= \lambda x:\tau.v \mid \Lambda a:\kappa.v \mid x \mid K(\vec{\sigma}, \vec{v}) \\ \langle \lambda x:\tau.v \parallel v' \cdot k \rangle &\mapsto \langle v \{v'/x\} \parallel k \rangle \\ \langle \Lambda a:\kappa.v \parallel \tau \cdot k \rangle &\mapsto \langle v \{\tau/a\} \parallel k \rangle \\ \langle K(\vec{\sigma}, \vec{v}) \parallel \text{case of } \vec{alt} \rangle &\mapsto c\{\vec{\sigma}/\vec{b}\}\{\vec{v}/x\} \quad K(\vec{b}:\kappa, \vec{x}:\vec{\tau}) \rightarrow c \in \vec{alt} \\ \langle W \parallel \text{case of } \vec{alt} \rangle &\mapsto c\{W/x\} \quad x:\tau \rightarrow c \in \vec{alt} \\ \langle \mu\text{ret}.c \parallel k \rangle &\mapsto c\{k/\text{ret}\} \\ \text{let } x:\tau = v \text{ in } c &\mapsto c\{v/x\} \\ \text{let } j:\tau' = \tilde{\mu}[\vec{a}:\vec{\kappa}, \vec{x}:\vec{\tau}].c' \text{ in } c &\mapsto c\{c'\{\vec{\sigma}/\vec{a}\}\{\vec{v}/x\}/\text{jump } j \vec{\sigma} \vec{v}\} \end{aligned}$$

Figure 2. Call-by-name operational semantics

`let` is only slightly harder, but we omit it here for simplicity. Note that the rule for continuation `lets` uses *structural substitution* [2, 28], which replaces every command matching the pattern `jump  $j \vec{\sigma} \vec{v}$` with the given command. Intuitively, we can think of this substitution as inlining the right-hand side for j everywhere, and then β -reducing the jump at each inline site.

Note that Figure 2 serves equally well as an abstract machine, since every rule applies to the top of a command without having to search for a redex. Figure 2 can also be extended to a reduction theory for Sequent Core by permitting the rules to apply in any context, and further to an equational theory by symmetry, thereby providing a specification for valid transformations that a compiler might apply. Thus, a call-by-name operational semantics and an abstract machine are the same for Sequent Core, and the difference between a reduction and an equational theory is the difference between reducing anywhere or reducing only at the top of a command.

The most interesting rule is the one for computations:

$$\langle \mu\text{ret}.c \parallel k \rangle \mapsto c\{k/\text{ret}\}$$

If the computation $\mu\text{ret}.c$ is consumed by continuation k , then we can just substitute k for all the occurrences of `ret` in c . From the point of view of a control calculus or continuation-passing style, `ret` can be seen as a static variable bound by μ -abstractions, providing the correct notion of substitution. Another way to think about it is that the substitution $c\{k/\text{ret}\}$ appends k to the continuation(s) of c , including those in labeled continuations but not under any intervening μ -abstractions. For example:

$$\langle f \parallel x \cdot (\mu\text{ret}.c) \cdot \text{ret} \rangle \{y \cdot \text{ret}/\text{ret}\} = \langle f \parallel x \cdot (\mu\text{ret}.c) \cdot y \cdot \text{ret} \rangle$$

Expressing call-by-need simply requires the addition of a Launchbury-style [22] heap, as shown in Figure 3, which gives a lower-level operational reading of the different language constructs and shows how Sequent Core can be efficiently implemented. Note that unless otherwise specified in the rules, all additions to the heap $\mathcal{H}$ and jump environment $\mathcal{J}$ are assumed to be fresh, using α -renaming as necessary to pick a fresh variable or label. Specifically, the force and update rules modify an existing binding in the heap, whereas all the other rules allocate new heap bindings. Also note that for simplicity of the `let` rules, we assume that value bindings, bind_v , are kept separate from continuation bindings, bind_k , which can always be done as described in Section 2.2.3.

The main thing to notice about this semantics is how the different components of the state are separated—the heap $\mathcal{H}$, the jump environment $\mathcal{J}$, and the linear continuation $\mathcal{R}$ —which is only possible because of our scope restrictions described in Section 2.2.3. Specifically, every rule that allocates in the heap uses the fact that terms cannot access the labels in the jump environment, and the μ and force rules use the fact that a μ -abstraction starts a fresh scope of labels. The scoping rules further allow these different components

$V \in \text{Value} ::= \lambda x:\tau.v \mid \Lambda a:\kappa.v \mid x \mid K(\vec{\sigma}, \vec{v})$		$W \in \text{WHNF} ::= V \mid K(\vec{\sigma}, \vec{v})$
$\mathcal{H} \in \text{Heap} ::= \varepsilon \mid \Gamma, x = v \mid \Gamma, x = \bullet$		$\mathcal{J} \in \text{JumpEnv} ::= \varepsilon \mid \mathcal{J}, j = \bar{\mu}[\vec{a}:\vec{\kappa}, \vec{x}:\vec{\tau}].c \quad \mathcal{R} \in \text{LinKont} ::= \varepsilon \mid (k, \mathcal{J}) : \mathcal{R} \mid \text{upd } x : \mathcal{R}$
$\boxed{\langle \Gamma; \mathcal{J}, \mathcal{R}; c \rangle \rightsquigarrow \langle \Gamma; \mathcal{J}, \mathcal{R}; c \rangle}$		
$(\beta^{\rightarrow})$	$\langle \mathcal{H}; \mathcal{J}, \mathcal{R}; \langle \lambda x:\tau.v_1 \mid v_2 \cdot k \rangle \rangle \rightsquigarrow \langle \mathcal{H}, x = v_2; \mathcal{J}, \mathcal{R}; \langle v_1 \mid k \rangle \rangle$	
$(\beta^{\forall})$	$\langle \mathcal{H}; \mathcal{J}, \mathcal{R}; \langle \Lambda a:\kappa.v \mid \tau \cdot k \rangle \rangle \rightsquigarrow \langle \mathcal{H}; \mathcal{J}, \mathcal{R}; \langle v \{ \tau/a \} \mid k \rangle \rangle$	
$(\text{case}_{\text{cons}})$	$\langle \mathcal{H}; \mathcal{J}, \mathcal{R}; \langle K(\vec{\sigma}, \vec{v}) \parallel \text{case of } \vec{alt} \rangle \rangle \rightsquigarrow \langle \mathcal{H}, \vec{x} = \vec{v}; \mathcal{J}, \mathcal{R}; c \{ \vec{\sigma}/\vec{a} \} \rangle$	$K(\vec{a}:\vec{\kappa}, \vec{x}:\vec{\tau}) \rightarrow c \in \vec{alt}$
$(\text{case}_{\text{def}})$	$\langle \mathcal{H}; \mathcal{J}, \mathcal{R}; \langle W \parallel \text{case of } \vec{alt} \rangle \rangle \rightsquigarrow \langle \mathcal{H}, x = W; \mathcal{J}, \mathcal{R}; c \rangle$	$x \rightarrow c \in \vec{alt}$
(μ)	$\langle \mathcal{H}; \mathcal{J}, \mathcal{R}; \langle \mu \text{ret}.c \mid k \rangle \rangle \rightsquigarrow \langle \mathcal{H}; \varepsilon, (k, \mathcal{J}) : \mathcal{R}; c \rangle$	
(jump)	$\langle \mathcal{H}; \mathcal{J}, \mathcal{R}; \text{jump } j \vec{\sigma} \vec{v} \rangle \rightsquigarrow \langle \mathcal{H}, \vec{x} = \vec{v}; \mathcal{J}, \mathcal{R}; c \{ \vec{\sigma}/\vec{a} \} \rangle$	$j = \bar{\mu}[\vec{a}:\vec{\kappa}, \vec{x}:\vec{\tau}].c \in \mathcal{J}$
(lookup)	$\langle \mathcal{H}; \mathcal{J}, \mathcal{R}; \langle x \mid k \rangle \rangle \rightsquigarrow \langle \mathcal{H}; \mathcal{J}, \mathcal{R}; \langle V \mid k \rangle \rangle$	$x = V \in \mathcal{H}$
(lazysubst)	$\langle \mathcal{H}; \mathcal{J}, \mathcal{R}; \langle x \mid k \rangle \rangle \rightsquigarrow \langle \mathcal{H}, \vec{y} = \vec{v}, x = K(\vec{\sigma}, \vec{y}); \mathcal{J}, \mathcal{R}; \langle K(\vec{\sigma}, \vec{y}) \mid k \rangle \rangle$	$x = K(\vec{\sigma}, \vec{v}) \in \mathcal{H}$
(force)	$\langle \mathcal{H}; \mathcal{J}, \mathcal{R}; \langle x \mid k \rangle \rangle \rightsquigarrow \langle \mathcal{H}, x = \bullet; \varepsilon, \text{upd } x : (k, \mathcal{J}) : \mathcal{R}; c \rangle$	$x = \mu \text{ret}.c \in \mathcal{H}$
(update)	$\langle \mathcal{H}; \mathcal{J}, \text{upd } x : \mathcal{R}; \langle W \parallel \text{ret} \rangle \rangle \rightsquigarrow \langle \mathcal{H}, x = W; \mathcal{J}, \mathcal{R}; \langle W \parallel \text{ret} \rangle \rangle$	$x = \bullet \in \mathcal{H}$
(ret)	$\langle \mathcal{H}; \mathcal{J}, (k', \mathcal{J}') : \mathcal{R}; \langle W \parallel \text{ret} \rangle \rangle \rightsquigarrow \langle \mathcal{H}; \mathcal{J}', \mathcal{R}; \langle W \mid k' \rangle \rangle$	
$(\text{let}_{\text{val}})$	$\langle \mathcal{H}; \mathcal{J}, \mathcal{R}; \text{let } \text{bind}_v \text{ in } c \rangle \rightsquigarrow \langle \mathcal{H}, \text{bind}_v; \mathcal{J}, \mathcal{R}; c \rangle$	$\text{Dom}(\mathcal{H}) \cap \text{Dom}(\text{bind}_v) = \emptyset$
$(\text{let}_{\text{cont}})$	$\langle \mathcal{H}; \mathcal{J}, \mathcal{R}; \text{let } \text{bind}_k \text{ in } c \rangle \rightsquigarrow \langle \mathcal{H}; \mathcal{J}, \text{bind}_k; \mathcal{R}; c \rangle$	$\text{Dom}(\mathcal{J}) \cap \text{Dom}(\text{bind}_k) = \emptyset$

Figure 3. Call-by-need operational semantics

to embody different commonplace run-time entities. The heap $\mathcal{H}$ is of course implemented with a random-access mutable heap as usual. The linear continuation $\mathcal{R}$ is a mutable stack, since each element is accessed exactly once before disappearing. Contrastingly, the jump environment $\mathcal{J}$ serves only as syntactic bookkeeping for statically allocated code. Because of the scope restrictions, the binding for each label can be determined before execution, which is evident from the fact the let_{cont} side condition is guaranteed to hold whenever the initial program has distinct let -bound labels, making dynamic allocation unnecessary. So during execution the jump rule is a direct jump and the let_{cont} rule is nothing at all!

Even though the operational semantics of Figure 2 and Figure 3 implement different evaluation strategies—call-by-name and call-by-need, respectively—the two still produce the same answers. In particular, the abstract machine terminates if and only if the operational semantics does, which is enough to guarantee that the two semantics agree [24, 33].

Proposition 1 (Termination equivalence). *For any closed command c , $c \mapsto^* c_1 \not\mapsto$ if and only if $\langle \varepsilon; \varepsilon, \varepsilon; c \rangle \rightsquigarrow^* \langle \mathcal{H}; \mathcal{J}, \mathcal{R}; c_2 \rangle \not\rightsquigarrow$.*

This should not be a surprise, as Sequent Core is intended for representing pure functional programs, and for the pure λ -calculus the two evaluation strategies agree [3].

2.4 Type System

Another way to understand Sequent Core is through its type system, which is given in Figure 4. Unsurprisingly, the type system for Sequent Core is based on the sequent calculus by reflecting programs as right rules and continuations as left rules. In particular, it is an extension (as well as a restriction) of the type system for the $\bar{\lambda}\mu\tilde{\mu}$ -calculus [9, 16], and likewise we share the same unconventional notation for typing judgements to maintain a visible connection with the sequent calculus. More specifically, the typing judgements for the different syntactic categories of Sequent Core are written with the following sequents:

- The type of a command is described by $c : (\Gamma \vdash \Delta)$, which says that c is a well-typed command and may contain free variables described by Γ and Δ . Note that the command itself does not

have a type directly, as terms and continuations do, because it does not take input or produce output directly. Rather, the “return” type of a command is in Δ , as in $\langle 1 \parallel \text{ret} \rangle : (\varepsilon \vdash \text{ret} : \text{Int})$.

- The type of a term is described by $\Gamma \vdash v : \tau$, which has the usual reading for typing terms of System F ω . In particular, it says that v returns a result of type τ and may contain free variables (for either values or types) with the types described by Γ .
- The type of a continuation is described by $\Gamma \mid k : \tau \vdash \Delta$, which says that k consumes an input of type τ and may contain free variables with the types described by Γ and Δ .
- The type of a binding is described by $\text{bind} : (\Gamma \mid \Delta' \vdash \Gamma' \mid \Delta)$, which is the most complex form of sequent. In essence, it says that bind binds the variables in Γ' and Δ' and may contain references to free variables from Γ and Δ . For example, we have $(x:\text{Int} = z + 1) : (z : \text{Int} \mid \varepsilon \vdash x : \text{Int} \mid \text{ret} : \text{Bool})$ and $(j:\text{Int} = \bar{\mu}[x:\text{Int}]. \langle x \parallel \text{ret} \rangle) : (\varepsilon \mid j : \text{Int} \vdash \varepsilon \mid \text{ret} : \text{Int})$.

One important detail to note is the careful treatment of the continuation environment Δ in the rules of Figure 4. In particular, the term-typing judgement is missing Δ , which enforces the scoping restriction of continuation variables discussed in Section 2.2.3. A consequence of this fact is that Δ is treated linearly in the type system; it is only duplicated across the multiple alternatives of a case or in continuation bindings. Type variables in the environment $(\Gamma, a : \kappa, \Gamma')$ additionally scope over the remainder of the environment (Γ') as well as the entire conclusion (either Δ or $v : \tau$) and obey static scoping rules. Thus, the $\forall R$, TL , and $Label$ rules only apply if they do not violate the static scopes of type variables.

The unusual notation we used for the type system makes it easy to bridge the gap between Sequent Core and the sequent calculus. In particular, if we drop all expressions (commands, *etc.*) and vertical bars from Figure 4, we end up with a corresponding logic of the sequent calculus, as shown in Figure 5 (where the type kinding is identical to Figure 4). All the similarly named rules come directly from erasing extra information from Figure 4, and additionally Ax represents both Var and Ret , $MultiCut$ represents Let , WR represents $Name$, and the rest of the typing rules do nothing in

$$\begin{array}{c}
\Gamma \in \text{Environment} ::= \varepsilon \mid \Gamma, x : \tau \mid \Gamma, a : \kappa \mid \Gamma, K : \tau \mid \Gamma, T : \kappa \qquad \Delta \in \text{CoEnvironment} ::= \varepsilon \mid \text{ret} : \tau \mid \Delta, j : \tau \\
\\
\text{Type kinding: } \boxed{\Gamma \vdash \tau : \kappa} \\
\frac{}{\Gamma, a : \kappa \vdash a : \kappa} \text{TyVar} \quad \frac{}{\Gamma, T : \kappa \vdash T : \kappa} \text{TyCon} \quad \frac{\Gamma \vdash \sigma : \kappa' \rightarrow \kappa \quad \Gamma \vdash \tau : \kappa'}{\Gamma \vdash \sigma \tau : \kappa} \text{TyApp} \quad \frac{\Gamma, a : \kappa \vdash \tau : \star}{\Gamma \vdash \forall a : \kappa. \tau : \star} \forall \quad \frac{\overrightarrow{\Gamma, \bar{a} : \bar{\kappa} \vdash \tau : \star}}{\Gamma \vdash \exists \bar{a} : \bar{\kappa}. (\vec{\tau}) : \star} \exists \times \\
\\
\text{Command typing: } \boxed{c : (\Gamma \vdash \Delta)} \\
\frac{\text{bind} : (\Gamma \mid \Delta' \vdash \Gamma' \mid \Delta) \quad c : (\Gamma, \Gamma' \vdash \Delta', \Delta)}{\text{let bind in } c : (\Gamma \vdash \Delta)} \text{Let} \quad \frac{\Gamma \vdash v : \tau \quad \Gamma \mid k : \tau \vdash \Delta}{\langle v \parallel k \rangle : (\Gamma \vdash \Delta)} \text{Cut} \quad \frac{\overrightarrow{\Gamma \vdash \sigma : \kappa} \quad \overrightarrow{\Gamma \vdash v : \tau \{ \sigma / a \}}}{\text{jump } j \vec{\sigma} \vec{v} : (\Gamma \vdash j : \exists \bar{a} : \bar{\kappa}. (\vec{\tau}), \Delta)} \text{Jump} \\
\\
\text{Term typing: } \boxed{\Gamma \vdash v : \tau} \\
\frac{}{\Gamma, x : \tau \vdash x : \tau} \text{Var} \quad \frac{c : (\Gamma \vdash \text{ret} : \tau)}{\Gamma \vdash \mu \text{ret}. c : \tau} \text{Act} \quad \frac{\Gamma, x : \sigma \vdash v : \tau}{\Gamma \vdash \lambda x : \sigma. v : \sigma \rightarrow \tau} \rightarrow R \quad \frac{\Gamma, a : \kappa \vdash v : \tau}{\Gamma \vdash \Lambda a : \kappa. v : \forall a : \kappa. \tau} \forall R \\
\\
\frac{K : \forall \bar{a} : \bar{\kappa}. \forall \bar{b} : \bar{\kappa}'. \vec{\tau} \rightarrow T \vec{a} \in \Gamma \quad \overrightarrow{\Gamma \vdash \sigma : \kappa'} \quad \overrightarrow{\Gamma \vdash v : \tau' \{ \sigma / a \} \{ \sigma / b \}}}{\Gamma \vdash K(\vec{\sigma}, \vec{v}) : T \vec{\tau}} \text{TRK} \\
\\
\text{Continuation typing: } \boxed{\Gamma \mid k : \tau \vdash \Delta} \text{ and Alternative typing: } \boxed{\Gamma \mid \text{alt} : \tau \vdash \Delta} \\
\frac{}{\Gamma \mid \text{ret} : \tau \vdash \text{ret} : \tau, \Delta} \text{Ret} \quad \frac{\Gamma \vdash v : \sigma \quad \Gamma \mid k : \tau \vdash \Delta}{\Gamma \mid v : k : \sigma \rightarrow \tau \vdash \Delta} \rightarrow L \quad \frac{\Gamma \vdash \sigma : \kappa \quad \Gamma \mid k : \tau \{ \sigma / a \} \vdash \Delta}{\Gamma \mid \sigma : k : \forall a : \kappa. \tau \vdash \Delta} \forall L \quad \frac{\Gamma \mid \text{alt} : \tau \vdash \Delta}{\Gamma \mid \text{case of alt} : \tau \vdash \Delta} \text{Case} \\
\\
\frac{c : (\Gamma, x : \tau \vdash \Delta)}{\Gamma \mid x : \tau \rightarrow c : \tau \vdash \Delta} \text{Deflt} \quad \frac{K : \forall \bar{a} : \bar{\kappa}'. \forall \bar{b} : \bar{\kappa}. \vec{\sigma} \rightarrow T \vec{a} \in \Gamma \quad c : (\Gamma, \bar{b} : \bar{\kappa}, x : \sigma \{ \tau / a \} \vdash \Delta)}{\Gamma \mid K(\bar{b} : \bar{\kappa}, \bar{x} : \vec{\sigma}) \rightarrow c : T \vec{\tau} \vdash \Delta} \text{TLK} \\
\\
\text{Binding typing: } \boxed{\text{bind} : (\Gamma \mid \Delta' \vdash \Gamma' \mid \Delta)} \text{ and } \boxed{bp : (\Gamma \mid \Delta' \vdash \Gamma' \mid \Delta)} \\
\frac{\Gamma \vdash v : \tau}{(x : \tau = v) : (\Gamma \mid \varepsilon \vdash x : \tau \mid \Delta)} \text{Name} \quad \frac{c : (\Gamma, \bar{a} : \bar{\kappa}, \bar{x} : \vec{\tau} \vdash \Delta)}{(j : \exists \bar{a} : \bar{\kappa}. (\vec{\tau}) = \bar{\mu}[\bar{a} : \bar{\kappa}, \bar{x} : \vec{\tau}]. c) : (\Gamma \mid j : \exists \bar{a} : \bar{\kappa}. (\vec{\tau}) \vdash \varepsilon \mid \Delta)} \text{Label} \\
\\
\frac{\Gamma' = \vec{\Gamma}'' \quad \Delta' = \vec{\Delta}'' \quad bp : (\Gamma, \Gamma' \mid \Delta'' \vdash \Gamma'' \mid \Delta', \Delta)}{\text{rec } \{ \vec{bp} \} : (\Gamma \mid \Delta' \vdash \Gamma' \mid \Delta)} \text{Rec}
\end{array}$$

Figure 4. Type System

$$\begin{array}{c}
\Gamma \in \text{Assumption} ::= \varepsilon \mid \Gamma, \tau \mid \Gamma, K : \tau \mid \Gamma, a : \kappa \mid \Gamma, T : \kappa \qquad \Delta \in \text{Conclusion} ::= \varepsilon \mid \Delta, \tau \\
\\
\text{Structural rules:} \\
\frac{}{\Gamma, \tau \vdash \tau, \Delta} \text{Ax} \quad \frac{\Gamma \vdash \tau \quad \Gamma, \tau \vdash \Delta}{\Gamma \vdash \Delta} \text{Cut} \quad \frac{\Gamma, \Delta' \vdash \Gamma', \Delta \quad \Gamma, \Gamma' \vdash \Delta', \Delta}{\Gamma \vdash \Delta} \text{MultiCut} \quad \frac{\Gamma \vdash \tau}{\Gamma \vdash \tau, \Delta} \text{WR} \\
\\
\frac{\overrightarrow{\Gamma, \tau \vdash \Delta}}{\Gamma, \tau \vdash \Delta} \text{Case} \quad \frac{\Gamma' = \vec{\Gamma}'' \quad \Delta' = \vec{\Delta}'' \quad \overrightarrow{\Gamma, \Gamma', \Delta'' \vdash \Gamma'', \Delta', \Delta}}{\Gamma, \Delta' \vdash \Gamma', \Delta} \text{Rec} \\
\\
\text{Logical rules:} \\
\frac{\Gamma, \sigma \vdash \tau}{\Gamma \vdash \sigma \rightarrow \tau} \rightarrow R \quad \frac{\Gamma \vdash \sigma \quad \Gamma, \tau \vdash \Delta}{\Gamma, \sigma \rightarrow \tau \vdash \Delta} \rightarrow L \quad \frac{\Gamma, a : \kappa \vdash \tau}{\Gamma \vdash \forall a : \kappa. \tau} \forall R \quad \frac{\Gamma \vdash \sigma : \kappa \quad \Gamma, \tau \{ \sigma / a \} \vdash \Delta}{\Gamma, \forall a : \kappa. \tau \vdash \Delta} \forall L \\
\\
\frac{\overrightarrow{\Gamma \vdash \sigma : \kappa} \quad \overrightarrow{\Gamma \vdash \tau \{ \sigma / a \}}}{\Gamma \vdash \exists \bar{a} : \bar{\kappa}. (\vec{\tau}), \Delta} \text{Jump} \quad \frac{\Gamma, \bar{a} : \bar{\kappa}, \vec{\tau} \vdash \Delta}{\Gamma, \exists \bar{a} : \bar{\kappa}. (\vec{\tau}) \vdash \Delta} \text{Label} \\
\\
\frac{K : \forall \bar{a} : \bar{\kappa}. \forall \bar{b} : \bar{\kappa}'. \vec{\tau} \rightarrow T \vec{a} \in \Gamma \quad \overrightarrow{\Gamma \vdash \sigma : \kappa'} \quad \overrightarrow{\Gamma \vdash \tau' \{ \sigma / a \} \{ \sigma / b \}}}{\Gamma \vdash T \vec{\tau}} \text{TRK} \quad \frac{K : \forall \bar{a} : \bar{\kappa}. \forall \bar{b} : \bar{\kappa}'. \vec{\sigma} \rightarrow T \vec{a} \in \Gamma \quad \overrightarrow{\Gamma, \bar{b} : \bar{\kappa}', \sigma \{ \tau / a \} \vdash \Delta}}{\Gamma, T \vec{\tau} \vdash \Delta} \text{TLK}
\end{array}$$

Figure 5. Logical interpretation

the logic. This logic resembles Gentzen’s original sequent calculus LK [14], but there are still some differences—besides just choice of connectives for types and kinds—that come from its application as an intermediate language for functional programs:

- LK has explicit inference rules for the structural properties of sequents: weakening (adding extraneous assumptions or conclusions), contraction (merging duplicate assumptions or conclusions), and exchange (swapping two assumptions or conclusions). Instead, the logic of Sequent Core makes these properties (except for weakening on the right, WR) implicit by generalizing the initial Ax rule to allow for extraneous assumptions and conclusions, and by duplicating the assumptions and conclusions in rules like Cut, MultiCut, and $\rightarrow$ L. This comes from the interpretation of static variables in the language as side assumptions and side conclusions. However, implicit and explicit presentations of structural properties are logically equivalent to one another.
- LK is a *classical logic*, which is achieved by allowing for any number of extra side conclusions in all the left (L) and right (R) rules. Contrarily, Gentzen’s also introduced an *intuitionistic logic* LJ [14] as a variant of LK that just limits all rules to only allow exactly one conclusion at all times. Instead, the logic that comes out of Sequent Core lies in between these two systems: sometimes there can only be one conclusion, and sometimes there can be many. Furthermore, it happens that the well-typed terms of Sequent Core correspond to well-typed expressions of Core (see Theorem 4 in the following Section 3), so Sequent Core still captures a similar notion of purity from the λ -calculus. This demonstrates that there is room for more subtle variations on intuitionistic logic that lie between the freedom of LK and the purity of LJ.
- Unlike LK, Figure 5 is not logically consistent as is, corresponding to the fact that the presented type system for Sequent Core allows for non-terminating programs and does not force exhaustiveness of case analysis. In particular, the Rec rule can prove anything and would of course be left out of a consistent subset of the logic. Similarly, the Case and TLK rules should only be used to build proofs by exhaustive case analysis at a type, as in the (informally) derived TL rule:

$$\frac{\frac{\Gamma, \overrightarrow{a} : \kappa, \overrightarrow{b} : \kappa, \overrightarrow{\sigma} \rightarrow T \overrightarrow{\tau}}{\Gamma, T \overrightarrow{\tau} \vdash \Delta} \text{Case} \quad \frac{\Gamma, \overrightarrow{b} : \kappa, \overrightarrow{\sigma} \{ \tau / a \} \vdash \Delta}{\Gamma, T \overrightarrow{\tau} \vdash \Delta} \text{TLK}}{\Gamma, T \overrightarrow{\tau} \vdash \Delta}$$

So a consistent subset of Figure 5 is attainable by further restricting the rules and the types along these lines.⁴

Perhaps the most complex rules in Figure 4 are Jump and Label for multiple-input continuations. It may seem a bit bizarre that the polymorphism is pronounced “exists” instead of “forall,” but luckily the above Curry-Howard reading in Figure 5 helps explain what’s going on. There are two instances of the Jump and Label rules that are helpful to consider. The first is when we have exactly two monomorphic inputs:

$$\frac{\Gamma \vdash \tau_1 \quad \Gamma \vdash \tau_2}{\Gamma \vdash (\tau_1, \tau_2)} \text{Jump} \quad \frac{\Gamma, \tau_1, \tau_2 \vdash \Delta}{\Gamma, (\tau_1, \tau_2) \vdash \Delta} \text{Label}$$

⁴ Besides removing Rec and merging Case with TLK, both recursive types (implicitly available from assumptions $K : \tau$) and the MultiCut would also need to be restricted for consistency. Interestingly, the Γ - Δ partitioning of let bindings described in Section 2.2.3 provides a sufficient consistency criterion for MultiCut. So maintaining the distinction between active conclusions (*i.e.*, terms) and assumptions (*i.e.*, continuations) as presented in Figure 4 in the logic is enough to tame MultiCut.

These are the right and left rules for the (tensor) product type in the sequent calculus—the right Jump rule allows only one conclusion like LJ and the left Label rule allows multiple conclusions like LK—illustrating the (tuple) product nature of multiple-input continuations and jumps. Second is when we have one polymorphic type quantified over one input:

$$\frac{\Gamma \vdash \sigma : \kappa \quad \Gamma \vdash \tau \{ \sigma / a \}}{\Gamma \vdash \exists a : \kappa. \tau} \text{Jump} \quad \frac{\Gamma, a : \kappa, \tau \vdash \Delta}{\Gamma, \exists a : \kappa. \tau \vdash \Delta} \text{Label}$$

These are exactly the right and left rules for existential types in the sequent calculus—with the same comparison to the right rule of LJ and the left rule of LK—which justifies the use of $\exists$ for the types of labels and jumps. Notice that this special case giving existential quantification is formally dual⁵ to the rules $\forall$ L and $\forall$ R for universal quantification by transporting conclusions on the right of the turnstile ($\vdash$) to the left and vice versa (except for the quantified σ and a , which stay put).

Furthermore, notice that the static scope of type variables uniformly gives the correct logical provisos for the quantifiers in the $\forall$ R, Label, and TLK rules. For example, the following command from Section 2.2.2

```
let j:∃a:*. (a, a → Bool) = μ[a:*, x:a, f:a → Bool]. ⟨f || x · ret⟩
in jump j Bool True not
```

is typable by the sequent ($\varepsilon \vdash \text{ret} : \text{Bool}$) because the type of the free variable `ret` does not reference the locally quantified a . However, the seemingly similar command

```
let j:∃a:*. (a, a → a) = μ[a:*, x:a, f:a → a]. ⟨f || x · ret⟩
in jump j Bool True not
```

is not typable by the sequent ($\varepsilon \vdash \text{ret} : a$)—or any other one—because a is local to the definition of j , and thus cannot escape in the type of `ret`. Pattern matching on existential data types in Haskell follows restrictions that are analogous to these scoping rules.

Finally, observe that the type system of Figure 4 is enough to ensure that well-typed programs don’t go wrong according to the operational semantics of Figure 2. In particular, type safety follows from standard lemmas for *progress* (every well-typed command is either a final state or can take a step) and *preservation* (every well-typed command is still well-typed after taking a step).

Proposition 2 (Type safety). 1. *Progress*: If $c : (\varepsilon \vdash \text{ret} : \tau)$ without Rec, then $c \mapsto c'$ or c has one of the following forms: $\langle W \parallel \text{ret} \rangle$ or $\langle K(\overrightarrow{\sigma}, \overrightarrow{v}) \parallel \text{case of } \overrightarrow{\text{alt}} \rangle$ where neither $K(b:\kappa, \overrightarrow{x}:\overrightarrow{\sigma}) \rightarrow c$ nor $x : \tau \rightarrow c$ are in $\overrightarrow{\text{alt}}$.

2. *Preservation*: If $c : (\Gamma \vdash \Delta)$ and $c \mapsto c'$ then $c' : (\Gamma \vdash \Delta)$.

Note that an unknown `case` is a possibility allowed by the type system as is, and just like with Core, ensuring exhaustiveness of case analysis rules this final state out so that the only result is $\langle W \parallel \text{ret} \rangle$. Also, because Figure 2 does not account for recursive lets, neither does the above progress proposition. Recursion is not a problem for progress, but it does take some additional care to treat explicitly.

3. Translating to and from Core

Core and Sequent Core are equally expressive, and it is illuminating to give translations between them, in both directions. In practical terms these translations were useful in our implementation, because in order to fit Sequent Core into the compiler pipeline, we need to translate from Core to Sequent Core and back.

⁵ The Jump and Label are not exactly dual to $\forall$ L and $\forall$ R in the classical sense due to the restriction of right rules to only one conclusion. Lifting this restriction to get the more LK-like logic makes these rules classically dual.

$$\begin{aligned}
S \llbracket \lambda x:\tau. e \rrbracket &= \lambda x:\tau. S \llbracket e \rrbracket & S \llbracket x \rrbracket &= x \\
S \llbracket \Lambda a:\kappa. e \rrbracket &= \Lambda a:\kappa. S \llbracket e \rrbracket & S \llbracket K \vec{\sigma} \vec{e} \rrbracket &= K(\vec{\sigma}, \overrightarrow{S \llbracket e \rrbracket}) \\
S \llbracket \text{let } bind \text{ in } e \rrbracket &= \mu ret. \text{let } S \llbracket bind \rrbracket \text{ in } \langle S \llbracket e \rrbracket \parallel ret \rangle \\
S \llbracket e e' \rrbracket &= \mu ret. \langle S \llbracket e \rrbracket \parallel S \llbracket e' \rrbracket \cdot ret \rangle \\
S \llbracket e \tau \rrbracket &= \mu ret. \langle S \llbracket e \rrbracket \parallel \tau \cdot ret \rangle \\
S \llbracket \text{case } e \text{ of } \overrightarrow{alt} \rrbracket &= \mu ret. \langle S \llbracket e \rrbracket \parallel \text{case of } \overrightarrow{S \llbracket alt \rrbracket} \rangle \\
S \llbracket x:\tau = e \rrbracket &= x:\tau = S \llbracket e \rrbracket \\
S \llbracket \text{rec } \{ \overrightarrow{x:\tau} \Rightarrow \vec{e} \} \rrbracket &= \text{rec } \{ \overrightarrow{x:\tau} = S \llbracket e \rrbracket \} \\
S \llbracket x:\tau \rightarrow e \rrbracket &= x:\tau \rightarrow \langle S \llbracket e \rrbracket \parallel ret \rangle \\
S \llbracket K \overrightarrow{a:\kappa} \overrightarrow{x:\tau} \rightarrow e \rrbracket &= K(\overrightarrow{a:\kappa}, \overrightarrow{x:\tau}) \rightarrow \langle S \llbracket e \rrbracket \parallel ret \rangle
\end{aligned}$$

Figure 6. Definitional translation from Core to Sequent Core

Core to Sequent Core	
$S_a \llbracket \lambda x:\tau. e \rrbracket = \lambda x:\tau. S_a \llbracket e \rrbracket$	$S_a \llbracket x \rrbracket = x$
$S_a \llbracket \Lambda a:\kappa. e \rrbracket = \Lambda a:\kappa. S_a \llbracket e \rrbracket$	$S_a \llbracket K \vec{\sigma} \vec{e} \rrbracket = K(\vec{\sigma}, \overrightarrow{S_a \llbracket e \rrbracket})$
$S_a \llbracket e \rrbracket = \mu ret. S_a \llbracket e \rrbracket ret$	where e is a computation
$S_a \llbracket \text{let } bind \text{ in } e \rrbracket k = \text{let } S_a \llbracket bind \rrbracket \text{ in } S_a \llbracket e \rrbracket k$	
$S_a \llbracket e e' \rrbracket k = S_a \llbracket e \rrbracket (S_a \llbracket e' \rrbracket \cdot k)$	
$S_a \llbracket e \tau \rrbracket k = S_a \llbracket e \rrbracket (\tau \cdot k)$	
$S_a \llbracket \text{case } e \text{ of } \overrightarrow{alt} \rrbracket k = \text{let } \overrightarrow{bind} \text{ in } S_a \llbracket e \rrbracket (\text{case of } \overrightarrow{S_a \llbracket alt \rrbracket} k')$	
	where $(\overrightarrow{bind}, k') = \text{shrink}(k)$
$S_a \llbracket e \rrbracket k = \langle S_a \llbracket e \rrbracket \parallel k \rangle$	where e is a value
$S_a \llbracket x:\tau = e \rrbracket = x:\tau = S_a \llbracket e \rrbracket$	
$S_a \llbracket \text{rec } \{ \overrightarrow{x:\tau} \Rightarrow \vec{e} \} \rrbracket = \text{rec } \{ \overrightarrow{x:\tau} = S_a \llbracket e \rrbracket \}$	
$S_a \llbracket x:\tau \rightarrow e \rrbracket k = x:\tau \rightarrow S_a \llbracket e \rrbracket k$	
$S_a \llbracket K \overrightarrow{a:\kappa} \overrightarrow{x:\tau} \rightarrow e \rrbracket k = K(\overrightarrow{a:\kappa}, \overrightarrow{x:\tau}) \rightarrow S_a \llbracket e \rrbracket k$	
Sequent Core to Core	
$D \llbracket \lambda x:\tau. v \rrbracket = \lambda x:\tau. D \llbracket v \rrbracket$	$D \llbracket x \rrbracket = x$
$D \llbracket \Lambda a:\kappa. v \rrbracket = \Lambda a:\kappa. D \llbracket v \rrbracket$	$D \llbracket K(\vec{\sigma}, \vec{v}) \rrbracket = K \vec{\sigma} \overrightarrow{D \llbracket v \rrbracket}$
$D \llbracket \mu ret. c \rrbracket = D \llbracket c \rrbracket$	
$D \llbracket \text{let } bind \text{ in } c \rrbracket = \text{let } D \llbracket bind \rrbracket \text{ in } D \llbracket c \rrbracket$	
$D \llbracket \langle v \parallel k \rangle \rrbracket = D \llbracket k \rrbracket [D \llbracket v \rrbracket]$	
$D \llbracket \text{jump } j \vec{\sigma} \vec{v} \rrbracket = j \vec{\sigma} \overrightarrow{D \llbracket v \rrbracket}$	
$D \llbracket v \cdot k \rrbracket = D \llbracket k \rrbracket [\square D \llbracket v \rrbracket]$	$D \llbracket ret \rrbracket = \square$
$D \llbracket \sigma \cdot k \rrbracket = D \llbracket k \rrbracket [\square \sigma]$	$D \llbracket \text{case of } \overrightarrow{alt} \rrbracket = \text{case } \square \text{ of } \overrightarrow{D \llbracket alt \rrbracket}$
$D \llbracket x:\tau = v \rrbracket = (x:\tau = D \llbracket v \rrbracket)$	
$D \llbracket j:\exists \overrightarrow{a:\kappa}. (\vec{\tau}) = \tilde{\mu}[\overrightarrow{a:\kappa}, \overrightarrow{x:\tau}]. c \rrbracket = (j:\forall \overrightarrow{a:\kappa}. \vec{\tau} \rightarrow \sigma = \Lambda \overrightarrow{a:\kappa}. \lambda \overrightarrow{x:\tau}. D \llbracket c \rrbracket)$	where $ret : \sigma$
$D \llbracket \text{rec } \{ \overrightarrow{bp} \} \rrbracket = \text{rec } \{ \overrightarrow{D \llbracket bp \rrbracket} \}$	
$D \llbracket x:\sigma \rightarrow c \rrbracket = x:\sigma \rightarrow D \llbracket c \rrbracket$	
$D \llbracket K(\overrightarrow{a:\kappa}, \overrightarrow{x:\tau}) \rightarrow c \rrbracket = K \overrightarrow{a:\kappa} \overrightarrow{x:\tau} \rightarrow D \llbracket c \rrbracket$	

Figure 7. Round-trip translations from Core to Sequent Core (S_a) and back to Core (D)

3.1 The Definitional Translation

Luckily, we can leverage the relationship between natural deduction and sequent calculus [9] to come up with a definitional translation from Sequent Core to Core, as shown in Figure 6. Notice that value expressions—lambdas and applied constructors, corresponding to introductory forms of natural deduction, as well as variables—translate one-for-one as terms of Sequent Core. Dealing with computations—applications and case expressions, corresponding to elimination forms, as well as let expressions—requires a μ -abstraction, since they are about introducing a continuation or setting up a binding in Sequent Core. Note how the introduction of μ -abstractions turns the focus on the computation in an expression: the operator of an application or the discriminant of a case becomes the term in a command, where the rest becomes a continuation waiting for its result.

3.2 A More Efficient Translation

Unfortunately, while the definitional translation into Sequent Core is straightforward, it is not very practical due to an excessive number of μ -abstractions. For example, the simple application $f \ 1 \ 2 \ 3$ is translated as

$$S \llbracket ((f \ 1) \ 2) \ 3 \rrbracket = \mu ret. \langle \mu ret. \langle \mu ret. \langle f \parallel 1 \cdot ret \rangle \parallel 2 \cdot ret \rangle \parallel 3 \cdot ret \rangle$$

instead of the direct $\mu ret. \langle f \parallel 1 \cdot 2 \cdot 3 \cdot ret \rangle$. These μ -abstractions are analogous to “administrative λ -abstractions” generated by CPS translations, since they both represent some bookkeeping that needs to be cleaned up before we get to the interesting part of a program. Thus, we want the analog of an administrative-free translation [32] for the sequent calculus, S_a , which we achieve by aggressively performing μ reductions during translation as shown in Figure 7 to give a reduced program in Sequent Core.

There is one caveat with the reduced translation into Sequent Core, though. One case during translation has a chance to duplicate the continuation by μ reduction: specifically, with case. Naïvely, the μ -reduced translation for case would be:

$$\llbracket \text{case } e \text{ of } \overrightarrow{alt} \rrbracket k = \llbracket e \rrbracket (\text{case of } \overrightarrow{\llbracket alt \rrbracket} k)$$

Because a case might have multiple alternatives, each alternative gets its own embedded copy of the given continuation. Simply copying the continuation is not good enough, as many programs can cause chain reactions of duplication, unacceptably inflating the size of the program (see Section 4.2). In practice we need to ensure that the continuation is small enough before duplicating it. Specifically, we force the continuation to be small by shrinking it, which we achieve by introducing extra value bindings for large arguments in a call stack and continuation bindings for the alternatives of a case. For example, the (large) call stack $v_1 \cdot v_2 \cdot ret$ can be shrunk to $x \cdot y \cdot ret$ along with the bindings $x = v_1$ and $y = v_2$. Additionally, the (large) case continuation

$$\text{case of } K_1(x, y) \rightarrow c_1; K_2(a, b, z) \rightarrow c_2$$

can be shrunk down to

$$\text{case of } K_1(x, y) \rightarrow \text{jump } j_1 x y; K_2(a, b, z) \rightarrow \text{jump } j_2 a b z$$

along with the bindings $j_1 = \tilde{\mu}[x, y]. c_1$ and $j_2 = \tilde{\mu}[a, b, z]. c_2$. Thus, when translating a case expression, we first shrink the given continuation, set up any bindings that the shrinking process created, and then copy the shrunken continuation in each alternative, as shown in Figure 7.

3.3 Translating Back to Core

But we don’t just want to translate one way, we also want the ability to come back to Core. That way, we can compose together both Core and Sequent Core passes by translating to and fro. Since Sequent

Core contains continuations, an obvious way to translate back to Core would be through a CPS translation. However, we want to make sure that a round trip through translations gives us back a similar program to what we originally had. This added round-trip stipulation means that CPS is right out: the program we would get from a CPS translation would be totally different from the one we started with. Even worse, if we were to iterate these round trips, it would compound the problem.

Thus, we look instead for a direct-style translation from Sequent Core to Core, which essentially reverses the translation into Sequent Core. This translation does not rely on types, but it does require properly scoped labels as described in Section 2.2.3. The scope restriction ensures that a Sequent Core program can be directly interpreted as a purely functional Core program without the use of any control effects.

3.4 Round Trips

The question remains: does a round-trip translation yield a similar program? To be more precise, we should expect, as a minimum criterion, that the round-trip translation respects *observational equivalence*: the same behavior in all contexts. We consider two Core expressions observationally equivalent, $e_1 \cong e_2$, whenever $C[e_1]$ terminates if and only if $C[e_2]$ does for all contexts C . Observational equivalence of two Sequent Core terms is similar. The answer is yes: the round-trip translation starting from Core may introduce some bindings and perform some commutative conversions, but the result is still observationally equivalent to where we started. Likewise, all round trips starting from Sequent Core produce observationally equivalent programs: The difference here is that some μ reductions may be performed by the round trip and, unfortunately, all continuation bindings are converted to value bindings.

Proposition 3 (Round trip). $D[S_a[e]] \cong e$ and $S_a[D[v]] \cong v$.

Also note that both directions of translation are type-preserving, as expected: the outputs of S_a and D are well-typed whenever their inputs are. However, unlike with a CPS transformation, the types are (largely) unchanged. In particular, the type of a Core expression is not changed by translation, which is evident by the fact that S_a doesn't change the types of bindings. Going the other way, D only changes the types of labels and nothing else, which is again evident by the translation of bindings. So the type of Sequent Core terms does not change by translation either.

Proposition 4 (Type preservation). *If $\Gamma \vdash e : \tau$ in Core then $\Gamma \vdash S_a[e] : \tau$. If $\Gamma \vdash v : \tau$ in Sequent Core then $\Gamma \vdash D[v] : \tau$.*

It is unfortunate that continuation bindings are lost en route during a round-trip translation starting from Sequent Core, as observed above. We had those continuation bindings for a reason, and they should not be erased. Fortunately though, there is a program transformation known as *contification* (described later in Section 5) which can recover the lost continuation bindings (and more) from the soup of value bindings, effectively *re-contifying* them. That means we can move between Core and Sequent Core with wild abandon without losing anything.

4. From Theory to Practice

To find out whether Sequent Core is a practical intermediate language, we implemented a plugin for the Glasgow Haskell Compiler (GHC), a mature, production-quality compiler for Haskell. In this section we reflect what we learned from this experience.

4.1 Sequent Core in GHC

GHC's simplifier is *the* central piece of GHC's optimization pipeline, comprising around 5,000 lines of Haskell code that has been refined over two decades. It applies a large collection of optimizing transformations, including inlining, β -reduction, η -conversion, let floating, case-of-case, case-of-known-constructor, etc.

We implemented a Sequent Core plug-in that can be used with an unmodified GHC. The plug-in inserts itself into the Core-to-Core optimization pipeline and replaces each invocation of GHC's simplifier by the following procedure: convert from Core to Sequent Core; apply the same optimizing transformations as the existing simplifier; and convert back to Core. Here is what we learned:

- Sequent Core is clearly up to the job. In a few months we were able to replicate all of the cleverness that GHC's Simplifier embodies, and our experiments confirm that the performance of the resulting code is essentially identical (see Section A in the appendix for details). Considering the maturity of GHC's existing Simplifier, this is a good result.
- We originally anticipated that Sequent Core could simplify GHC's Simplifier, since the latter accumulates and transforms an explicit continuation (called a *SimplCont*) in the style of a zipper [18], which is closely analogous to Sequent Core's continuations. Thus, we could say that Sequent Core gives a language to the logic that lies in the heart of GHC's Simplifier, providing a more direct representation of GHC optimizations.

In practice, we found that Sequent Core did not dramatically reduce the lines of code of the Simplifier. The syntax of Sequent Core jumps straight to the interesting action, avoiding the need to accumulate a continuation. However, the Simplifier is complex enough, and requires enough auxiliary information, that the lines of code saved here were a drop in the bucket. The savings were further offset by functions need to traverse the additional syntactic structures of Sequent Core.

So on the practical front, we are not yet ready to abandon Core in favor of Sequent Core in GHC. However, we did find an aspect of optimization for which Sequent Core was qualitatively better than Core: the treatment of join points, to which we turn next.

4.2 Join Points and Case-of-Case

Optimizing compilers commonly push code down into the branches of conditional constructs when possible, bringing *ifs* and *cases* to the top level [39, 40]. Besides clarifying the possible code paths, this tends to put intermediate results in their proper context, which enables further optimizations.

GHC performs such code motion aggressively. The most ambitious example is the *case-of-case transform* [30, 38]. Consider:

half $x = \text{if even } x \text{ then Just}(x \text{ 'div' } 2) \text{ else Nothing}$

After desugaring and inlining, *half* is written in Core as:

half $= \lambda x. \text{case } (\text{case } x \text{ 'mod' } 2 \text{ of } 0 \rightarrow \text{True} \text{ } _ \rightarrow \text{False}) \text{ of}$
 $\text{True} \rightarrow \text{Just}(x \text{ 'div' } 2)$
 $\text{False} \rightarrow \text{Nothing}$

Notice how the outer boolean *case* will receive a *True* or *False* value, depending on the result of $(x \text{ 'mod' } 2)$. We can make this fact clearer by applying the case-of-case transform, bringing the whole outer *case* inside each branch of the inner *case*:

half $= \lambda x. \text{case } x \text{ 'mod' } 2 \text{ of}$
 $0 \rightarrow \text{case True of True} \rightarrow \text{Just}(x \text{ 'div' } 2)$
 $\text{False} \rightarrow \text{Nothing}$
 $_ \rightarrow \text{case False of True} \rightarrow \text{Just}(x \text{ 'div' } 2)$
 $\text{False} \rightarrow \text{Nothing}$

Happily, case-of-case has revealed an easy simplification, giving:

$$\begin{aligned} \text{half} &= \lambda x. \text{case } x \text{ 'mod' } 2 \text{ of } 0 \rightarrow \text{Just}(x \text{ 'div' } 2) \\ &\quad - \rightarrow \text{Nothing} \end{aligned}$$

Of course, this is the ideal outcome, because the outer **case** ultimately vanished entirely. But if case-of-case did not reveal further simplifications, we would have duplicated the outer case, whose alternatives might be of arbitrary size.⁶

Like many compilers, including the original ANF implementation [13], GHC avoids excessive code duplication by abstracting large **case** alternatives into named functions that serve as join points. A typical result looks like this:

$$\begin{aligned} f : \text{Int} \rightarrow \text{Int} &= \lambda x. \text{let } j : \text{Maybe Int} \rightarrow \text{Int} = \lambda w. \dots \\ &\quad \text{in case } g \ x \text{ of } \text{Left } y \rightarrow j \ (\text{Just } y) \\ &\quad \quad \quad \text{Right } _ \rightarrow j \ \text{Nothing} \end{aligned}$$

It may appear that we have introduced extra overhead—allocating a closure for j (at the **let**) and calling the function. But a function like j has special properties: it is only tail-called, and it is never captured in a closure. GHC’s code generator takes advantage of these properties and compiles the tail call to j into two instructions: adjust the stack pointer and jump. Apart from this special treatment in the code generator, GHC’s Core Simplifier does not treat join points specially: they are just local function bindings, subject to the usual optimizations. Collapsing multiple concepts into one can be a strength—but as we see next, it can also be a weakness.

4.3 Losing Join Points

Continuing the example of the previous section, suppose f is called in the following way: **case** $f \ x$ **of** $0 \rightarrow \text{False}; _ \rightarrow \text{True}$. If f is inlined at this call site, another case-of-case transformation will occur, and after some further simplifications, we get this:

$$\begin{aligned} \text{let } j : \text{Maybe Int} \rightarrow \text{Int} &= \lambda w. \dots \\ \text{in case } g \ x \text{ of } \text{Left } y &\rightarrow \text{case } j \ (\text{Just } y) \text{ of } 0 \rightarrow \text{False} \\ &\quad - \rightarrow \text{True} \\ \text{Right } _ &\rightarrow \text{case } j \ \text{Nothing} \text{ of } 0 \rightarrow \text{False} \\ &\quad - \rightarrow \text{True} \end{aligned}$$

Now j is no longer tail-called and must be compiled as a regular function, with all the overhead entailed. Case-of-case has ruined a perfectly good join point!

This does not happen in Sequent Core. Here is the same function f in Sequent Core:

$$\begin{aligned} f : \text{Int} \rightarrow \text{Int} &= \lambda x. \mu \text{ret}. \\ \text{let } j : \text{Maybe Int} &= \tilde{\mu}[w]. \dots \text{ret} \dots \\ \text{in } \langle g \parallel x \cdot \text{case of } &\text{Left } y \rightarrow \text{jump } j \ \text{Just}(y) \\ &\quad \text{Right } _ \rightarrow \text{jump } j \ \text{Nothing} \rangle \end{aligned}$$

This time, j is represented by a labeled continuation accepting a *Maybe Int*. Moreover, observe that the body of j refers to the **ret** bound the surrounding μ . In Sequent Core, the case-of-case transformation is implemented by a μ reduction, which substitutes a **case** continuation for **ret** in a computation. For example, inlining f into the command $\langle f \parallel x \cdot \text{case of } 0 \rightarrow \text{False}; _ \rightarrow \text{True} \rangle$ followed by routine Sequent Core simplification instead gives us:

$$\begin{aligned} \text{let } j : \text{Maybe Int} &= \tilde{\mu}[w]. \dots \text{case of } 0 \rightarrow \text{False}; _ \rightarrow \text{True} \dots \\ \text{in } \langle g \parallel x \cdot \text{case of } &\text{Left } y \rightarrow \text{jump } j \ \text{Just}(y) \\ &\quad \text{Right } _ \rightarrow \text{jump } j \ \text{Nothing} \rangle \end{aligned}$$

Notice what happened here: just by substituting for **ret**, we did not push the continuation into the alternatives of the **case**, but instead it naturally flowed into the body of j . Jumps are *stable* in Sequent

Core, and the operational semantics of Sequent Core has showed us how to perform case-of-case without ruining any join points!

Could we do the same in Core? Well, yes: the case-of-case transform should (somehow) push the outer **case** into the join point j itself, rather than wrapping it around the *calls* to j , as in

$$\begin{aligned} \text{let } j : \text{Maybe Int} \rightarrow \text{Bool} &= \lambda w. \dots \text{case } e \text{ of } 0 \rightarrow \text{False}; _ \rightarrow \text{True} \dots \\ \text{in case } g \ x \text{ of } \text{Left } y &\rightarrow j \ (\text{Just } y) \\ &\quad \text{Right } _ \rightarrow j \ \text{Nothing} \end{aligned}$$

where the **case** e **of** $0 \rightarrow \text{False}; _ \rightarrow \text{True}$ means to wrap every expression e that returns from j with the **case** analysis. This effect is hard to achieve in Core as it stands, because join points are not distinguished and so the above substitution is not obvious. However, it is natural and straightforward in Sequent Core.

5. Contification

As we saw in Section 3, the translation from Sequent Core to Core is lossy. Sequent Core maintains a distinction between multiple-input continuations (join points) and ordinary functions because they have a different operational and logical reading, but Core only has functions. Converting Core to Sequent Core produces a program with no join points; and even if the Sequent Core Simplifier creates some, they will be lost in the next round trip through Core.

CPS-based compilers often employ a demotion technique called *contification* [20] that turns ordinary functions into continuations. Since direct jumps are faster than function calls, this operation is useful in its own right, but for us it is essential to make the round trip from Sequent Core to Core and back behave like an identity function. So, whenever we translate to Sequent Core, we also perform a simple contification pass that is just thorough enough to find and restore any continuation bindings that could have been lost in translation. In other words, contification picks up what we dropped while going back and forth between the two representations, and hence we are “re-contifying.” But we may also discover, and then exploit, join points that happened to be written by the user (Section 5.2).

The mechanics of contification are straightforward. In essence, contification converts function calls (which need to return) into direct jumps (which don’t) by baking the calling context into the body of the function. For example, suppose we have this code:

$$\begin{aligned} \text{let } f &= \lambda y. \mu \text{ret}. c \\ \text{in } \langle g \parallel x \cdot \text{case of } &A \ z \rightarrow \langle f \parallel z \cdot \text{ret} \rangle \\ &B \ _ \rightarrow \langle \text{True} \parallel \text{ret} \rangle \\ &C \ _ \rightarrow \langle f \parallel \text{True} \cdot \text{ret} \rangle \rangle \end{aligned}$$

Here f is an ordinary function, bound by the **let** and called in two of the three branches of the **case**. Moreover, both its calls are saturated tail calls, and f is not captured inside a thunk or function closure. Under these circumstances, it is semantics-preserving to replace f with a join point and replace its calls with more efficient direct jumps, thus:

$$\begin{aligned} \text{let } j &= \tilde{\mu}[y]. c \\ \text{in } \langle g \parallel x \cdot \text{case of } &A \ z \rightarrow \text{jump } j \ z \\ &B \ _ \rightarrow \langle \text{True} \parallel \text{ret} \rangle \\ &C \ _ \rightarrow \text{jump } j \ \text{True} \rangle \end{aligned}$$

This transformation is sound even if the binding of f is recursive, provided all the recursive calls obey the same rules.

We saw in Section 4.2 that GHC already performs a similar analysis during code generation to identify tail calls that can be converted to jumps. However, this conversion happens just *once* and only *after* all Core-to-Core optimizations are finished. Here, we bring contification forward as a pass that can happen in the midst of the main optimization loop, giving a language to talk about join points for other optimizations to exploit.

⁶ We see the same code duplication issue arise during translation in Section 3; we avoid duplication with the same solution in both instances.

5.1 Analysis and Transformation

We divide the contification algorithm into two phases, an *analysis phase* and a *transformation phase*. The analysis phase finds functions that can be contified; then the transformation phase carries out the necessary rewrites in one sweep. See Section B in the appendix for a more detailed description of the algorithm.

In the analysis phase, we are interested in answering the all-important question: given a **let**-bound function f (a potential join point), is every call to f a saturated tail call? Sequent Core lets us state this condition precisely: all its calls must be of the form $\langle f \parallel \vec{v} \cdot \text{ret} \rangle$, where the *ret* is the same *ret* that is in scope at f 's binding site. This is another occasion on which it is helpful to think of *ret* as a lexically-scoped variable bound by μret .

To answer the question, the algorithm for the analysis phase gathers data bottom-up, similar to a free-variable analysis, and marks which **let**-bound functions may be replaced with labeled continuations. During traversal, the analysis determines the full set of variables that actually appear free in an expression, which we call the *free set*, as well as the subset of those variables that only appear as tail calls, which we call the *good set*. Three basic rules govern the upward propagation of these sets of variables:

- The head variable of a tail call—that is, any f in a command of the form $\langle f \parallel \vec{v} \cdot \text{ret} \rangle$ —is good so long as it is not free in the arguments.
- Terms cannot have any good variables, since labels cannot appear free in a term. The exception to this rule is a bound function that will be contified, since of course it won't be a term anymore; thus contification has a cascading effect.
- When considering other forms of expressions with several subexpressions, a variable is good on the whole if it is good in at least one subexpression and bad in none of them.

If a **let**-bound variable is considered good after analyzing its entire scope, then we mark it for contification. For a non-recursive binding $\text{let } f = v \text{ in } c$, the scope is just the body of the **let**, c . For a set of mutually recursive bindings $\text{let rec } \{f = \lambda \vec{x}. \mu\text{ret}. c\} \text{ in } c'$, the scope includes all the function bodies $\vec{c}$ as well as c' . Note that we can't contify only part of a recursive set; it's all or none. The reason for this restriction is that if we only contified some of the functions in a recursive binding, then those newly labeled continuations would be out of scope for the unconfited functions left behind, thus breaking the mutual recursion.

Once the analysis is complete, the transformation itself is straightforward. Note that we assume that all functions begin with a series of lambdas and then one μ ; this can always be arranged, since v and $\mu\text{ret}. \langle v \parallel \text{ret} \rangle$ are equivalent by η -conversion. At each binding $f = \lambda \vec{x}. \mu\text{ret}. c$ marked for contification, we pick a fresh label j and rewrite the binding as $j = \tilde{\mu}[\vec{x}]. c$. Then, at every call site of a contified f , we rewrite $\langle f \parallel \vec{v} \cdot \text{ret} \rangle$ as $\text{jump } j \vec{v}$.

This algorithm is similar to Kennedy's [20], except that we only contify functions called with the *ret* continuation rather than any single continuation. Clearly, this algorithm impacts fewer functions than the more general one. We also implemented a more comprehensive (and thus more expensive) contification pass for GHC, only to find it offered little impact on performance. Thus, the simpler algorithm seems to lie in a sweet spot between ease of implementation, cost of execution, and recovery of previously identified join points.

5.2 Discovering Join Points

By contifying aggressively enough, we may be able to discover—and exploit—additional join points that happen to arise naturally in the program. Consider the standard library function *find* that returns

the first value in a list satisfying a predicate, or *Nothing* if none satisfy it:

```
find : (a → Bool) → [a] → Maybe a
find = λp.λxs. let go = λys. case ys of
    []      → Nothing
    (y : ys') → case p y of
        True  → Just y
        False → goys'
in go xs
```

Here *go* is a join point, and our contification analysis can discover that fact.⁷ Now suppose that *find* is called as:

```
f = λxs. case find even xs of Nothing → 0
                        Just x   → x + 1
```

In GHC today, after inlining and performing some routine simplifications, we get

```
f = λxs. let go = λys. case ys of
    []      → Nothing
    (y : ys') → case y `mod` 2 of
        0 → Just y
        _ → go ys'
in case go xs of Nothing → 0
                Just x   → x + 1
```

And now *go* is no longer a join point, another example of the phenomenon we previously saw in Section 4.3. In Core, this is as far as we can go. In Sequent Core, however, *go* can be contified, and after the dust settles, we end up with the analog of:

```
f = λxs. let go = λys. case ys of
    []      → 0
    (y : ys') → case y `mod` 2 of
        0 → y + 1
        _ → go ys'
in go xs
```

This code is much better: the allocation of the function closure for *go* and the intermediate value (*Just y*) are both avoided, the outer **case** disappears, and the recursive call in *go* turns into a jump. In Sequent Core, once a join point is discovered by contification, it *stays* a join point.

5.3 Caveats

Since idiomatic Haskell code uses curried functions, we must be careful about what we consider a tail call. Note the rewrites above assumed the number of arguments in the tail calls matched the number of outer lambdas in the function. Thus, for our purposes, all tail calls must also be *exact*—each must provide exactly as many arguments as there are outer lambdas in the function definition. This restriction can be lifted so long as the calls are *consistent* in the number of arguments passed, also known as the *call arity* [7]. However, it is unclear that this is useful often enough to warrant the additional complexity. For the purpose of re-contification after round-trip translations, all relevant tail calls will already be exact.

Thus far, we have neglected to mention the impact of polymorphism on contification. Polymorphic functions are no problem, with one proviso: a function cannot be contified if it is polymorphic in its return type. To see why, suppose we have some function $f : \forall a : \star. \text{Int} \rightarrow a \rightarrow a$ and we generate the continuation:

$$j = \tilde{\mu}[a : \star, n : \text{Int}, x : a]. \langle f \parallel a \cdot n \cdot x \cdot \text{ret} \rangle$$

⁷ Note that *go* is recursive, but recursive join points are fine as long as they are properly tail-called.

We have a problem: the continuation here cannot be well-typed, because `ret` is a free variable that comes from the outside, so `ret` cannot be typed by *this* a since a was not in scope when `ret` was bound. Note that this is not an issue particular to Sequent Core; the same phenomenon arises in typed CPS languages. Contification in a typed setting must always be careful here.

Like arity, this is not a fatal issue. Just as we *can* call f with *True* by passing *Bool* as the type argument, we can contify f by fixing the appropriate return type in context. For example, if the `ret` in scope has type *Bool*, then the following is well-typed:

$$j = \tilde{\mu}[n: \text{Int}, x: \text{Bool}]. \langle f \parallel \text{Bool} \cdot n \cdot x \cdot \text{ret} \rangle$$

However, this situation appears to be vanishingly rare in practice. And in the case of re-contification, Sequent Core join points never have a polymorphic return type when translated to Core.⁸ Thus, while correctness demands that we at least *check* for polymorphic return types, re-contification can simply give up on them.

6. Related Work

6.1 Relation to Sequent Calculi

We might ask how and why Sequent Core differs from similar computational interpretations of the sequent calculus, like the $\lambda\mu\tilde{\mu}$ -calculus [9] or System L [26] for example. A primary difference is that in lieu of `let` bindings, these previous languages have continuations of the form $\tilde{\mu}x.c$, as mentioned in Section 2.2.2. As it turns out, they are not needed in Sequent Core. Indeed, a common reading for $\tilde{\mu}x.c$ at the top of a command is:

$$\langle v \parallel \tilde{\mu}x.c \rangle = (\text{let } x = v \text{ in } c)$$

where the $\tilde{\mu}$ is replaced with an explicit substitution via `let`. Then, using the call-by-name semantics for Sequent Core, all $\tilde{\mu}$ -abstractions can be *lifted*⁹ to the top of a command, so every $\tilde{\mu}$ -abstraction can be written as a `let`. Under a call-by-value semantics, $\tilde{\mu}$ -abstractions play a more crucial role as noted by Curien and Herbelin [9], but in that case they are exactly the same as a default `case` in Sequent Core: $\tilde{\mu}x.c = \text{case of } x \rightarrow c$. So again, the extra continuation is not needed. Considering the fact that the recursion so elegantly expressed by a `let` cannot be represented with $\tilde{\mu}$ -abstractions alone gives `let` its primary role in Sequent Core.

The other differences between Sequent Core and previous sequent calculi are the labeled, multiple-input continuations and jumps. The exact formulations of these constructs were designed with the needs of a practical compiler in mind, but they do have a more theoretical reading as well. In particular, we could (and indeed at one point we did) consider adding general existential types to the language. That way a multiple-input continuation $\tilde{\mu}[\vec{a}, \vec{x}].c$ might be interpreted as just a `case` continuation `case of` $(\vec{a}, \vec{x}) \rightarrow c$, so that labels just refer to single-input continuations and jumps are just cut pairs $\langle (\vec{\sigma}, \vec{\tau}) \parallel j \rangle$. However, for this to represent the correct operational cost at run time, it is crucial that these existential tuples are *unboxed* [29], meaning that they are values of a truly positive type [26] unlike the normally boxed Haskell data types. Additionally, unboxed (existential) tuples give less restraint for labels and jumps than the syntactic limitations implicitly imposed in Figure 1. The result is an unfortunately heavy-handed encoding unless stricter measures are taken for these positive types, as in Zeilberger’s [46] interpretation of focalization.

⁸ Consider a term containing a local join point with no intervening μ :

$$\mu \text{ret}. \dots \text{let } j: \exists \vec{a}: \vec{\tau}. (\vec{\tau}) = \tilde{\mu}[\vec{a}: \vec{\tau}, \vec{x}: \vec{\tau}]. c \text{ in } \dots$$

Assuming the term has type σ , the μ -bound `ret` will also have type σ . After translation, j ’s type becomes $\forall \vec{a}: \vec{\tau}. \vec{\tau} \rightarrow \sigma$, where $\vec{a}$ cannot occur free in the return type σ due to the typical hygiene requirements of translation.

⁹ This lifting can be done by the ς reductions of [44] and [25].

6.2 CPS as an Intermediate Language

Though continuations had been actively researched for nearly a decade [34], the first use of CPS for compilation appeared in 1978, in the Rabbit compiler for Scheme [40]. Steele was interested in the connection between Scheme and the λ -calculus [43], and CPS was a way to “elucidate some important compilation issues,” such as evaluation order and intermediate results, while maintaining that connection. He also noted the ease of translation to an imperative machine language. Standard ML of New Jersey [1] is another prominent example; it even performs such low-level tasks as register assignment within the CPS language itself.

So what’s stopping GHC from just adopting CPS? One answer is “but *which* CPS?” Usually the “CPS” intermediate language refers to “the language produced by the call-by-value CPS transform.” Surely we would not use *this* “CPS” to compile a non-strict language like Haskell. Of course, there are call-by-name [17, 32] and call-by-need [27] CPS transforms, but (to our knowledge) they have not been used in compilers before, leaving us in unknown territory.

More importantly, the effect of *any* CPS transform is to fix an evaluation order in the syntax of the program, but GHC routinely exploits the ability to reorder calculations, like shifting between call-by-need and call-by-value, which gives more flexibility for optimizing a pure, lazy language like Haskell [30].

Hence an advantage of a sequent calculus for GHC: like the λ -calculus, *the syntax does not fix a specific evaluation order*. For example, we illustrated both call-by-name (Figure 2) and call-by-need (Figure 3) readings for the same Sequent Core programs, and call-by-value would be valid, too. So we can still reason about Haskell programs with call-by-name semantics while implementing them efficiently with call-by-need.

There is one more advantage shared by both Core and Sequent Core, but not CPS, which is critically important for GHC. Specifically, GHC allows for arbitrary rewrite rules that transform function calls [31], which enable user optimizations like stream fusion [8]. Both Core and Sequent Core make expressing and implementing these custom rules easy, since both languages make nested function call structure in expressions like `map f (map g xs)` apparent: either as a chain of applications or a call stack. Instead, CPS represents nested functional calls in the source *abstractly* as in:

$$\lambda k. \text{map } g (\lambda h. h \text{ xs } (\lambda y. \text{map } f (\lambda h'. h' y \text{ k})))$$

To understand the original call structure of the program requires chasing information through several indirections. Instead, Sequent Core represents function calls *structurally* as stacks that can be immediately inspected, so bringing continuations to GHC without getting in the way of what GHC already does. In this light, we can view the sequent calculus as a “strategically defunctionalized” [36] CPS language. There is an essential trade-off in the expressive capabilities of rigid structure versus free abstraction [35], and the additional structure provided by call stacks enables more optimizations, like rewrite rules, by making continuations scrutable.

6.3 ANF as an Intermediate Language

In 1992, Sabry and Felleisen [37] demonstrated that the actions of a CPS compiler could be understood in terms of the original source code. Hence, though a CPS transform could express call-by-value semantics for λ -terms, that same semantics could be expressed as reductions in the original term. The new semantics extended Plotkin’s call-by-value λ -calculus [32] with more reductions and hence further possible optimizations. Flanagan *et al.* [13] argued that the new semantics obviated the need for CPS in the compiler, since the same work would be performed by systematically applying the new reductions, putting the source terms into *administrative normal form* (ANF). Representations in ANF became popular in the following years, as its ease of implementation provides an obvious

benefit over CPS, but its costs took some time to appreciate in practical use. In 2007, Kennedy [20] outlined some of these issues, encouraging compiler writers to “give CPS a second chance.”

It is worthwhile to point out that Sequent Core does not suffer from the same difficulties as ANF described by Kennedy. Sequent Core does not require renormalization after routine transformations: the syntax is closed under reductions like inlining. Furthermore, the various *commutative conversions* in these direct-style representations—such as case-of-case discussed in Section 4.2—are uniformly represented through μ -abstraction and μ -reduction in Sequent Core [26], which is a strength shared with CPS. Likewise, labeled continuations in Sequent Core serve an analogous purpose to labeled continuations in CPS, which preserve code sharing during these commutative conversions. Therefore, we can say that the sequent calculus can do everything CPS can, so Sequent Core retains the advantages of continuations laid out by Kennedy.

6.4 Other Representations

Our focus has been on functional programming, but of course compilation in the imperative world has long been dominated by the *static single-assignment form* (SSA) [10]. While it is known that SSA can be converted to CPS [19], the flat structure of SSA may be more convenient for representing imperative programs wherein block layout is more natural and higher-order features are not used.

We don’t have to choose between flat and higher-order, however. Thorin [23] is a graph-based representation aiming to support both imperative and functional code by combining a flat structure for ease of code transformation and first-class closures for implementing higher-order languages. However, Thorin is still intended for use in strict languages with pervasive side effects; it remains to be seen whether such a representation could be adapted for high-level optimizations in a non-strict regime such as Haskell.

7. Reflections on Intermediate Languages

There are many different goals we have for an intermediate language in an optimizing compiler, some of which seem at odds with one another. In a perfect world, an intermediate representation would, among other things:

1. Have a simple grammar, which makes it easy to traverse and transform programs written in the language. For example, if the grammar is represented as data types in a functional language, there should be a minimal number of (mutually) recursive data types that represent the main workhorse of runtime behavior.
2. Have a simple operational meaning, which makes it easy to analyze and understand the performance and behavior of programs. For example, it should be easy to find the “next” step of the program from the top of the syntax tree, and language constructs should be easy to compile to machine code.
3. Be as flexible in evaluation order as the source language permits, to permit as many transformations and out-of-order reductions as possible during optimization.
4. Make it easy to express control flow and shared join points, to reduce code size without hurting performance.
5. Make it easy to apply arbitrary rewrite rules expressed in the source language, especially for curried function applications when they appear pervasively in the source language.

We summarize the trade-offs for different representation styles, using + for “good”, “−” for “not good”, and blank for neutral, in Figure 8. Is there a way to get the best of all worlds? Perhaps, just as Sabry and Felleisen showed that you can get the advantages of CPS by using direct-style ANF, we would be able to get the advantages of a sequent calculus in a direct-style variant of Core.

	Core	Sequent Core	CPS
Simple grammar	+		
Operational reading	+	++	++
Flexible eval order	+	+	−
Control flow	−	++	++
Rewrite rules	+	+	−

Figure 8. Advantages of different representation styles

In particular, the killer advantage of Sequent Core has turned out to be its treatment of join points, which are different from both functions and continuations, and the more powerful case-of-case transformations that they support (Section 4.2). Informed by this experience, we speculate that it should be possible to augment Core with explicit join points, rather than treat them as ordinary bindings the way GHC does now. We are actively exploring this line of work, using Sequent Core as our model that gives us a foundation for the theory and design of a direct-style λ -calculus with join points. Such a λ -calculus would sit between Core and Sequent Core by having both a simple syntactic structure while also preserving the extra information about control flow. Thus, Sequent Core can currently be seen as a laboratory for compiler intermediate representations.

In developing Sequent Core, we had a love/hate relationship with purity—specifically, with the absence of control effects. On the one hand, keeping Sequent Core “pure” lets us easily leverage the existing technology for compiling the λ -calculus efficiently. The restrictions on continuation variables and jumps create the direct-style correspondence between the two, enabling the same techniques for simplification and call-by-need evaluation (in contrast to [5, 6]). On the other hand, the sequent calculus gives rise to a language of first-class control effects in its natural state, equivalent to adding callcc to the λ -calculus. Thus, the *classical* sequent calculus is more expressive [11], and lets us collapse otherwise distinct concepts—like control flow and data flow, or functions and data structures—into symmetrical dual pairs. Here, we chose to restrain Sequent Core and maintain the connection with Core. However, it still remains to be seen how an unrestrained, and thus more cohesive and simpler, classic sequent calculus would fare as the intermediate language of a compiler.

Looking back to the table of trade-offs, we see that Sequent Core strikes a middle ground between Core and CPS. Beside the point about simple grammar—for which it is hard to improve upon the elegance of the λ -calculus—Sequent Core manages to combine the advantages of both direct and continuation-passing styles. Clearly, the focus of our comparison was between Core and Sequent Core, for which we conclude that the sequent calculus shows us how to bring control flow and continuation-passing-style optimizations to GHC without getting in the way of what GHC already does well. But this is a two-way road: the sequent calculus can also teach us how to bring flexibility and direct-style optimizations, like rewrite rules, to CPS compilers by bringing the *structures* underlying continuations out of the *abstractions*. We chalk this up as another in a long line of wins for the Curry-Howard isomorphism: in the debate between direct and continuation-passing style compilers, the logic tells us how we might have our cake and eat it too.

Acknowledgements

We would like to thank Iavor S. Diatchki for his help and input on the design and implementation of Sequent Core, and to thank Matthias Felleisen and Olin Shivers for discussions on the merits of continuations passing style. Paul Downen and Zena M. Ariola have been supported by NSF grant CCF-1423617.

References

- [1] A. W. Appel. *Compiling with Continuations*. Cambridge University Press, New York, NY, USA, 1992. ISBN 0-521-41695-7.
- [2] Z. M. Ariola and H. Herbelin. Control reduction theories: The benefit of structural substitution. *Journal of Functional Programming*, 18(3):373–419, May 2008. ISSN 0956-7968. doi: 10.1017/S0956796807006612. URL <http://dx.doi.org/10.1017/S0956796807006612>.
- [3] Z. M. Ariola, J. Maraist, M. Odersky, M. Felleisen, and P. Wadler. A call-by-need lambda calculus. In *Proceedings of the 22nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '95, pages 233–246, New York, NY, USA, 1995. ACM. ISBN 0-89791-692-1. doi: 10.1145/199448.199507. URL <http://doi.acm.org/10.1145/199448.199507>.
- [4] Z. M. Ariola, A. Bohannon, and A. Sabry. Sequent calculi and abstract machines. *ACM Transactions on Programming Languages and Systems*, 31(4):13:1–13:48, May 2009. ISSN 0164-0925. doi: 10.1145/1516507.1516508. URL <http://doi.acm.org/10.1145/1516507.1516508>.
- [5] Z. M. Ariola, H. Herbelin, and A. Saurin. Classical call-by-need and duality. In *Proceedings of the 10th International Conference on Typed Lambda Calculi and Applications*, TLCA'11, pages 27–44, Berlin, Heidelberg, 2011. Springer-Verlag. ISBN 978-3-642-21690-9. URL <http://dl.acm.org/citation.cfm?id=2021953>. 2021961.
- [6] Z. M. Ariola, P. Downen, H. Herbelin, K. Nakata, and A. Saurin. Classical call-by-need sequent calculi: The unity of semantic artifacts. In *Proceedings of the 11th International Conference on Functional and Logic Programming*, FLOPS'12, pages 32–46, Berlin, Heidelberg, 2012. Springer-Verlag. ISBN 978-3-642-29821-9. doi: 10.1007/978-3-642-29822-6_6. URL http://dx.doi.org/10.1007/978-3-642-29822-6_6.
- [7] J. Breitner. Call arity. In *Trends in Functional Programming - 15th International Symposium, TFP 2014, Soesterberg, The Netherlands, May 26-28, 2014. Revised Selected Papers*, pages 34–50, 2014. doi: 10.1007/978-3-319-14675-1_3. URL http://dx.doi.org/10.1007/978-3-319-14675-1_3.
- [8] D. Coutts, R. Leshchinskiy, and D. Stewart. Stream fusion: From lists to streams to nothing at all. In *Proceedings of the 12th ACM SIGPLAN International Conference on Functional Programming*, ICFP '07, pages 315–326, New York, NY, USA, 2007. ACM. ISBN 978-1-59593-815-2. doi: 10.1145/1291151.1291199. URL <http://doi.acm.org/10.1145/1291151.1291199>.
- [9] P. Curien and H. Herbelin. The duality of computation. In *Proceedings of the Fifth ACM SIGPLAN International Conference on Functional Programming (ICFP '00), Montreal, Canada, September 18-21, 2000*, pages 233–243, 2000. doi: 10.1145/351240.351262. URL <http://doi.acm.org/10.1145/351240.351262>.
- [10] 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. doi: 10.1145/115372.115320. URL <http://doi.acm.org/10.1145/115372.115320>.
- [11] M. Felleisen. On the expressive power of programming languages. *Science of Computer Programming*, 17(1-3):35–75, Dec. 1991. ISSN 0167-6423. doi: 10.1016/0167-6423(91)90036-W. URL [http://dx.doi.org/10.1016/0167-6423\(91\)90036-W](http://dx.doi.org/10.1016/0167-6423(91)90036-W).
- [12] M. Felleisen and D. P. Friedman. Control operators, the SECD-machine, and the λ -calculus. In M. Wirsing, editor, *Formal Description of Programming Concepts III*, pages 193–219. North Holland Press, Amsterdam, 1986.
- [13] C. Flanagan, A. Sabry, B. F. Duba, and M. Felleisen. The essence of compiling with continuations. In *Proceedings of the ACM SIGPLAN 1993 Conference on Programming Language Design and Implementation*, PLDI '93, pages 237–247, New York, NY, USA, 1993. ACM. ISBN 0-89791-598-4. doi: 10.1145/155090.155113. URL <http://doi.acm.org/10.1145/155090.155113>.
- [14] G. Gentzen. Investigations into logical deduction. In M. Szabo, editor, *Collected papers of Gerhard Gentzen*, pages 68–131. North-Holland, 1969.
- [15] H. Herbelin. A lambda-calculus structure isomorphic to gentzen-style sequent calculus structure. In *Computer Science Logic, 8th International Workshop, CSL '94, Kazimierz, Poland, September 25-30, 1994, Selected Papers*, pages 61–75, 1994. doi: 10.1007/BFb0022247. URL <http://dx.doi.org/10.1007/BFb0022247>.
- [16] H. Herbelin. Explicit substitutions and reducibility. *Journal of Logic and Computation*, 11(3):431–451, 2001. doi: 10.1093/logcom/11.3.431. URL <http://logcom.oxfordjournals.org/content/11/3/431.abstract>.
- [17] M. Hofmann and T. Streicher. Continuation models are universal for lambda-mu-calculus. In *Proceedings, 12th Annual IEEE Symposium on Logic in Computer Science, Warsaw, Poland, June 29 - July 2, 1997*, pages 387–395, 1997. doi: 10.1109/LICS.1997.614964. URL <http://dx.doi.org/10.1109/LICS.1997.614964>.
- [18] G. Huet. The zipper. *Journal of Functional Programming*, 7(5):549–554, Sept. 1997. ISSN 0956-7968. doi: 10.1017/S0956796897002864. URL <http://dx.doi.org/10.1017/S0956796897002864>.
- [19] R. A. Kelsey. A correspondence between continuation passing style and static single assignment form. In *Papers from the 1995 ACM SIGPLAN Workshop on Intermediate Representations*, IR '95, pages 13–22, New York, NY, USA, 1995. ACM. ISBN 0-89791-754-5. doi: 10.1145/202529.202532. URL <http://doi.acm.org/10.1145/202529.202532>.
- [20] A. Kennedy. Compiling with continuations, continued. In *Proceedings of the 12th ACM SIGPLAN International Conference on Functional Programming*, ICFP '07, pages 177–190, New York, NY, USA, 2007. ACM. ISBN 978-1-59593-815-2. doi: 10.1145/1291151.1291179. URL <http://doi.acm.org/10.1145/1291151.1291179>.
- [21] D. A. Kranz, R. Kelsey, J. Rees, P. Hudak, and J. Philbin. ORBIT: an optimizing compiler for Scheme. In *Proceedings of the 1986 SIGPLAN Symposium on Compiler Construction, Palo Alto, California, USA, June 25-27, 1986*, pages 219–233, 1986. doi: 10.1145/12276.13333. URL <http://doi.acm.org/10.1145/12276.13333>.
- [22] J. Launchbury. A natural semantics for lazy evaluation. In *Proceedings of the 20th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '93, pages 144–154, New York, NY, USA, 1993. ACM. ISBN 0-89791-560-7. doi: 10.1145/158511.158618. URL <http://doi.acm.org/10.1145/158511.158618>.
- [23] R. Leißa, M. Köster, and S. Hack. A graph-based higher-order intermediate representation. In *Proceedings of the 13th Annual IEEE/ACM International Symposium on Code Generation and Optimization*, CGO '15, pages 202–212, Washington, DC, USA, 2015. IEEE Computer Society. ISBN 978-1-4799-8161-8. URL <http://dl.acm.org/citation.cfm?id=2738600>. 2738626.
- [24] A. Meyer and S. Cosmadakis. Semantical Paradigms: Notes for an Invited Lecture. Technical report, MIT Laboratory for Computer Science, 545 Technology Square, Cambridge, MA 02139, July 1988.
- [25] G. Munch-Maccagnoni. Focalisation and classical realisability. In *Proceedings of the 23rd CSL International Conference and 18th EACSL Annual Conference on Computer Science Logic*, CSL'09/EACSL'09, pages 409–423, Berlin, Heidelberg, 2009. Springer-Verlag. ISBN 3-642-04026-8, 978-3-642-04026-9. URL <http://dl.acm.org/citation.cfm?id=1807662>. 1807695.
- [26] G. Munch-Maccagnoni. *Syntax and Models of a non-Associative Composition of Programs and Proofs*. PhD thesis, Univ. Paris Diderot, 2013.
- [27] C. Okasaki, P. Lee, and D. Tarditi. Call-by-need and continuation-passing style. *Lisp and Symbolic Computation*, 7(1):57–82, Jan. 1994. ISSN 0892-4635. doi: 10.1007/BF01019945. URL <http://dx.doi.org/10.1007/BF01019945>.
- [28] M. Parigot. Lambda-my-calculus: An algorithmic interpretation of classical natural deduction. In *Proceedings of the International Conference on Logic Programming and Automated Reasoning*, LPAR '92, pages 190–201, London, UK, UK, 1992. Springer-Verlag. ISBN 3-540-55727-X. URL <http://dl.acm.org/citation.cfm?id=645706>. 663989.
- [29] S. L. Peyton Jones and J. Launchbury. Unboxed values as first class citizens in a non-strict functional language. In *Proceedings*

- of the 5th ACM Conference on Functional Programming Languages and Computer Architecture, pages 636–666, London, UK, UK, 1991. Springer-Verlag. ISBN 3-540-54396-1. URL <http://dl.acm.org/citation.cfm?id=645420.652528>.
- [30] S. L. Peyton Jones and A. L. M. Santos. A transformation-based optimiser for Haskell. *Science of Computer Programming*, 32(1-3): 3–47, Sept. 1998. ISSN 0167-6423. doi: 10.1016/S0167-6423(97)00029-4. URL [http://dx.doi.org/10.1016/S0167-6423\(97\)00029-4](http://dx.doi.org/10.1016/S0167-6423(97)00029-4).
 - [31] S. L. Peyton Jones, A. Tolmach, and T. Hoare. Playing by the rules: rewriting as a practical optimisation technique in GHC. In *2001 Haskell Workshop*. ACM SIGPLAN, September 2001. URL <http://research.microsoft.com/apps/pubs/default.aspx?id=74064>.
 - [32] G. D. Plotkin. Call-by-name, call-by-value and the λ -calculus. *Theoretical Computer Science*, 1(2):125–159, 1975. doi: 10.1016/0304-3975(75)90017-1.
 - [33] J.-C. Raoult and J. Vuillemin. Operational and semantic equivalence between recursive programs. *Journal of the Association for Computing Machinery*, 27(4):772–796, Oct. 1980. ISSN 0004-5411. doi: 10.1145/322217.322229. URL <http://doi.acm.org/10.1145/322217.322229>.
 - [34] J. C. Reynolds. The discoveries of continuations. *Lisp and Symbolic Computation*, 6(3-4):233–248, 1993.
 - [35] J. C. Reynolds. User-defined types and procedural data structures as complementary approaches to data abstraction. In C. A. Gunter and J. C. Mitchell, editors, *Theoretical Aspects of Object-oriented Programming*, pages 13–23. MIT Press, Cambridge, MA, USA, 1994. ISBN 0-262-07155-X. URL <http://dl.acm.org/citation.cfm?id=186677.186680>.
 - [36] J. C. Reynolds. Definitional interpreters for higher-order programming languages. *Higher-Order and Symbolic Computation*, 11(4):363–397, Dec. 1998. ISSN 1388-3690. doi: 10.1023/A:1010027404223. URL <http://dx.doi.org/10.1023/A:1010027404223>.
 - [37] A. Sabry and M. Felleisen. Reasoning about programs in continuation-passing style. In *LISP and Functional Programming*, pages 288–298, 1992. doi: 10.1145/141471.141563. URL <http://doi.acm.org/10.1145/141471.141563>.
 - [38] A. L. M. Santos. *Compilation by Transformation in Non-Strict Functional Languages*. PhD thesis, University of Glasgow, 1995.
 - [39] T. Standish, D. Harriman, D. Kibler, and J. Neighbors. *The Irvine program transformation catalogue*. 1976.
 - [40] G. L. Steele, Jr. RABBIT: A compiler for SCHEME. Technical Report AITR-474, Massachusetts Institute of Technology, Artificial Intelligence Laboratory, 1978.
 - [41] G. L. Steele, Jr. and G. J. Sussman. Lambda: The ultimate declarative. Memo AIM-379, Massachusetts Institute of Technology, Artificial Intelligence Laboratory, 1976.
 - [42] G. L. Steele, Jr. and G. J. Sussman. Lambda: The ultimate imperative. Memo AIM-353, Massachusetts Institute of Technology, Artificial Intelligence Laboratory, 1976.
 - [43] G. J. Sussman and G. L. Steele, Jr. SCHEME: An interpreter for untyped lambda-calculus. Memo AIM-349, Massachusetts Institute of Technology, Artificial Intelligence Laboratory, 1975.
 - [44] P. Wadler. Call-by-value is dual to call-by-name. In *Proceedings of the Eighth ACM SIGPLAN International Conference on Functional Programming*, ICFP '03, pages 189–201, New York, NY, USA, 2003. ACM. ISBN 1-58113-756-7. doi: 10.1145/944705.944723. URL <http://doi.acm.org/10.1145/944705.944723>.
 - [45] P. Wadler. Propositions as types. *Communications of the ACM*, 58(12): 75–84, Nov. 2015. ISSN 0001-0782. doi: 10.1145/2699407. URL <http://doi.acm.org/10.1145/2699407>.
 - [46] N. Zeilberger. *The Logical Basis of Evaluation Order and Pattern-Matching*. PhD thesis, Carnegie Mellon University, 2009.