Grokking the Sequent Calculus (Functional Pearl)

DAVID BINDER, University of Tübingen, Germany

MARCO TZSCHENTKE, University of Tübingen, Germany

MARIUS MÜLLER, University of Tübingen, Germany

KLAUS OSTERMANN, University of Tübingen, Germany

The sequent calculus is a proof system which was designed as a more symmetric alternative to natural deduction. The $\lambda\mu\tilde{\mu}$-calculus is a term assignment system for the sequent calculus and a great foundation for compiler intermediate languages due to its first-class representation of evaluation contexts. Unfortunately, only experts of the sequent calculus can appreciate its beauty. To remedy this, we present the first introduction to the $\lambda\mu\tilde{\mu}$-calculus which is not directed at type theorists or logicians but at compiler hackers and programming-language enthusiasts. We do this by writing a compiler from a small but interesting surface language to the $\lambda\mu\tilde{\mu}$-calculus as a compiler intermediate language.

CCS Concepts: • **Theory of computation** → **Lambda calculus**; • **Software and its engineering** → **Compilers**; *Control structures*.

Additional Key Words and Phrases: Intermediate representations, continuations, codata types, control effects

1 Introduction

Suppose you have just implemented your own small functional language. To test it, you write the following function which multiplies all the numbers contained in a list:

$$\textbf{def } \text{mult}(l) \coloneqq \textbf{case } l \textbf{ of } \{ \text{Nil} \Rightarrow 1, \text{Cons}(x, xs) \Rightarrow x * \text{mult}(xs) \}$$

What bugs you about this implementation is that you know an obvious optimization: The function should directly return zero if it encounters a zero in the list. There are many ways to achieve this, but you choose to extend your language with labeled expressions and a goto instruction. This allows you to write the optimized version:

$$\textbf{def } \text{mult}(l) \coloneqq \textbf{label } \alpha \, \{ \text{mult'}(l; \alpha) \}$$

$$\textbf{def } \text{mult'}(l; \alpha) \coloneqq \textbf{case } l \textbf{ of } \{ \text{Nil} \Rightarrow 1, \text{Cons}(x, xs) \Rightarrow \textbf{ifz}(x, \textbf{goto}(0; \alpha), x * \text{mult'}(xs; \alpha)) \}$$

You used **label** $\alpha \, \{\text{mult'}(l; \alpha)\}$ to introduce a label α around the call to the helper function mult' which takes this label as an additional argument (we use ; to separate the label argument from the other arguments), and **goto**$(0; \alpha)$ to jump to this label α with the expression 0 in the recursive helper function. But since your language now has control effects, you need to reconsider how you want to compile and optimize programs. In particular, you have to decide on an appropriate intermediate language which can express these control effects. In this paper, we introduce you to

Authors' Contact Information: David Binder, Department of Computer Science, University of Tübingen, Tübingen, Germany, david.binder@uni-tuebingen.de; Marco Tzschentke, Department of Computer Science, University of Tübingen, Tübingen, Germany, marco.tzschentke@uni-tuebingen.de; Marius Müller, Department of Computer Science, University of Tübingen, Tübingen, Germany, mari.mueller@uni-tuebingen.de; Klaus Ostermann, Department of Computer Science, University of Tübingen, Tübingen, Germany, klaus.ostermann@uni-tuebingen.de.

one such intermediate language: the sequent-calculus-based $\lambda\mu\tilde{\mu}$-calculus. The result of compiling the efficient multiplication function to the $\lambda\mu\tilde{\mu}$-calculus is:

$$\textbf{def } \text{mult}(l;\alpha) \coloneqq \text{mult'}(l;\alpha,\alpha)$$
$$\textbf{def } \text{mult'}(l;\alpha,\beta) \coloneqq$$
$$\langle l \mid \textbf{case } \{\text{Nil} \Rightarrow \langle 1 \mid \beta\rangle, \text{Cons}(x,xs) \Rightarrow \textbf{ifz}(x, \langle 0 \mid \alpha\rangle, \text{mult'}(xs;\alpha,\tilde{\mu}z. *(x,z;\beta)))\}\rangle$$

Here is how you read this snippet: Besides the list argument l, the definition $\textbf{def } \text{mult}(l;\alpha) \coloneqq \dots$ takes an argument α which indicates how the computation should continue once the result of the multiplication is computed (we again use ; to separate these two kinds of arguments). The helper function mult' takes a list argument l and two arguments α and β; the argument β indicates where the function should return to on a normal recursive call while α indicates the return point of a short-circuiting computation. In the body of mult' we use $\langle l \mid \textbf{case } \{\text{Nil} \Rightarrow \dots, \text{Cons}(x,xs) \Rightarrow \dots\}\rangle$ to perform a case split on the list l. If the list is Nil, then we use $\langle 1 \mid \beta\rangle$ to return 1 to β, which is the return for a normal recursive call. If the list has the form $\text{Cons}(x,xs)$ and x is zero, we return with $\langle 0 \mid \alpha\rangle$, where α is the return point which short-circuits the computation. If x isn't zero, then we have to perform the recursive call $\text{mult'}(xs;\alpha, \tilde{\mu}z. *(x,z;\beta))$, where we use $\tilde{\mu}z. *(x,z;\beta)$ to bind the result of the recursive call to the variable z before multiplying it with x and returning to β. Don't be discouraged if this looks complicated at the moment; the main part of this paper will cover everything in much more detail.

The $\lambda\mu\tilde{\mu}$-calculus that you have just seen was first introduced by Curien and Herbelin [2000] as a solution to a long-standing open question: What should a term language for the sequent calculus look like? The sequent calculus is one of two influential proof calculi introduced by Gentzen [1935a,b] in a single paper, the other calculus being natural deduction. The term language for natural deduction is the ordinary lambda calculus, but it was difficult to find a good term language for the sequent calculus. After it had been found, the $\lambda\mu\tilde{\mu}$-calculus was proposed as a better foundation for compiler intermediate languages, for example by Downen et al. [2016]. Despite this, most language designers and compiler writers are still unfamiliar with it. This is the situation that we hope to remedy with this pearl.

We frequently discuss ideas which involve the $\lambda\mu\tilde{\mu}$-calculus with students and colleagues and therefore have to introduce them to its central ideas. But we usually cannot motivate the $\lambda\mu\tilde{\mu}$-calculus as a term assignment system for the sequent calculus, since most of them are not familiar with it. We instead explain the $\lambda\mu\tilde{\mu}$-calculus on the whiteboard by compiling small functional programs into it. Such an introduction is regrettably still missing in the published literature; most existing presentations either presuppose knowledge of the sequent calculus or otherwise spend a lot of space introducing it first. We believe that if one can understand the lambda

Fig. 1. Screenshot of the online evaluator.

calculus without first learning about natural deduction proofs, then one should also be able to understand the $\lambda\mu\tilde{\mu}$-calculus without knowing the sequent calculus[1].

Why are we excited about the $\lambda\mu\tilde{\mu}$-calculus, and why do we think that more people should become familiar with its central ideas and concepts? The main feature which distinguishes the $\lambda\mu\tilde{\mu}$-calculus from the lambda calculus is its first-class treatment of evaluation contexts. An evaluation context is the remainder of the program which runs after the current subexpression we are focused on finishes evaluating.

This becomes clearer with an example: When we want to evaluate the expression $(2 + 3) * 5$, we first have to focus on the subexpression $2 + 3$ and evaluate it to its result 5. The remainder of the program, which will run after we have finished the evaluation, can be represented with the evaluation context $\square * 5$. We cannot bind an evaluation context like $\square * 5$ to a variable in the lambda calculus, but in the $\lambda\mu\tilde{\mu}$-calculus we can bind such evaluation contexts to covariables. Furthermore, the μ-operator gives direct access to the evaluation context in which the expression is currently evaluated. Having such direct access to the evaluation context is not always necessary for a programmer who wants to write an application, but it is often important for compiler implementors who write optimizations to make programs run faster. One solution that compiler writers use to represent evaluation contexts in the lambda calculus is called continuation-passing style. In continuation-passing style, an evaluation context like $\square * 5$ is represented as a function $\lambda x.x * 5$. This solution works, but the resulting types which are used to type a program in this style are arguably hard to understand. Being able to easily inspect these types can be very valuable, especially for intermediate representations, where terms tend to look complex. The promise of the $\lambda\mu\tilde{\mu}$-calculus is to provide the expressive power of programs in continuation-passing style without having to deal with the type-acrobatics that are usually associated with it.

The remainder of this paper is structured as follows:

- In Section 2 we introduce the surface language **Fun** and show how we can translate it into the sequent-calculus-based language **Core**. The surface language is mostly an expression-oriented functional programming language, but we have added some features such as codata types and control operators whose translations provide important insights into how the $\lambda\mu\tilde{\mu}$-calculus works. In this section, we also compare how redexes are evaluated in both languages.
- In Section 3 we discuss static and dynamic focusing, which are two closely related techniques for lifting subexpressions which are not values into a position where they can be evaluated.
- Section 4 introduces the typing rules for **Fun** and **Core** and proves standard results about typing and evaluation.
- We show why we are excited about the $\lambda\mu\tilde{\mu}$-calculus in Section 5. We present various programming language concepts which become much clearer when we present them in the $\lambda\mu\tilde{\mu}$-calculus: We show that let-bindings are precisely dual to control operators, that data and codata types are two perfectly dual ways of specifying types, and that the case-of-case transformation is nothing more than a μ-reduction. These insights are not novel for someone familiar with the $\lambda\mu\tilde{\mu}$-calculus, but not yet as widely known as they should be.
- Finally, in Section 6 we discuss related work and provide pointers for further reading. We conclude in Section 7.

This paper is accompanied by a Haskell implementation which we also make available as an interactive website (cf. Figure 1). You can run the examples presented in this paper in the online evaluator.

---

[1]For the interested reader, we show in Appendix A how the sequent calculus and the $\lambda\mu\tilde{\mu}$-calculus are connected.

2 Translating To Sequent Calculus

In this section, we introduce **Fun**, an expression-oriented functional programming language, together with its translation into the sequent-calculus-based intermediate language **Core**. We present both languages and the translation function $[\![-]\!]$ in multiple steps, starting with arithmetic expressions and adding more features in later subsections. We postpone the typing rules for both languages until Section 4.

2.1 Arithmetic Expressions

We begin with arithmetic expressions which consist of variables, integer literals, binary operators and **ifz**, a conditional expression which checks whether its first argument is equal to zero. The syntax of arithmetic expressions for **Fun** and **Core** is given in Definition 2.1.

Definition 2.1 (Arithmetic Expressions).

$$x, y, z, \ldots \in \text{Variables} \quad \bigstar, \alpha, \beta, \gamma, \ldots \in \text{Covariables} \quad \odot \in \{*, +, -\}$$

	Fun		**Core**	
	$t ::= x \mid \ulcorner n \urcorner \mid t \odot t \mid \mathbf{ifz}(t, t, t)$	p	$::= x \mid \ulcorner n \urcorner \mid \mu\alpha.s$	*Producer*
		c	$::= \alpha$	*Consumer*
		s	$::= \odot(p, p; c) \mid \mathbf{ifz}(p, s, s) \mid \langle p \mid c \rangle$	*Statement*

$$[\![x]\!] := x \qquad\qquad [\![t_1 \odot t_2]\!] := \mu\alpha. \odot ([\![t_1]\!], [\![t_2]\!]; \alpha) \qquad\qquad (\alpha \text{ fresh})$$

$$[\![\ulcorner n \urcorner]\!] := \ulcorner n \urcorner \qquad [\![\mathbf{ifz}(t_1, t_2, t_3)]\!] := \mu\alpha.\mathbf{ifz}([\![t_1]\!], \langle [\![t_2]\!] \mid \alpha \rangle, \langle [\![t_3]\!] \mid \alpha \rangle) \qquad (\alpha \text{ fresh})$$

In **Fun** there is only one syntactic category: terms t. These terms can either be variables x, literals $\ulcorner n \urcorner$, binary operators $t + t$, $t * t$ and $t - t$, or a conditional $\mathbf{ifz}(t, t_0, t_1)$. This conditional evaluates to t_0 if t evaluates to $\ulcorner 0 \urcorner$, or to t_1 otherwise. In contrast to this single category, **Core** uses three different syntactic categories: producers p, consumers c and statements s. These categories are directly inherited from the $\lambda\mu\tilde{\mu}$-calculus, and it is important to understand their differences:

Producers All constructs in **Core** which *construct* or *produce* an element of some type belong to the syntactic category of producers. In other words, producers correspond to "introduction forms" or "proof terms", and every term of the language **Fun** is translated to a producer in **Core**.

Consumers Consumers are probably less intuitive than producers since they do not correspond directly to any term of the language **Fun**. The basic idea is that if some consumer c has type τ, then c *consumes* or *destructs* a producer of type τ. If you have encountered evaluation contexts or continuations before, then it is helpful to think of consumers of type τ as continuations or evaluation contexts for a producer of type τ. And if you are familiar with the Curry-Howard correspondence, then you can think of consumers as refutations or direct evidence that a proposition is false.

Statements Statements are the ingredient which make computation *happen*; without statements, we would only have static objects without any dynamic behavior. Here is a non-exhaustive list of examples for statements: Every IO action which reads from or prints to the console or a file should be represented as a statement in **Core**. Computations on primitive types such as machine integers should be statements. Finally, everything which is a redex in an

expression-based language should also correspond to a statement in **Core**. Since statements themselves only compute and do not return anything they do not have a type.

After these general remarks, let us now look at how arithmetic expressions are represented in the language **Core**. Variables x and literals $\ulcorner n \urcorner$ both belong to the category of producers, but binary operators are represented as statements $\odot(p_1, p_2; c)$. First, let us explain why they are represented as statements instead of producers. The idea is that a binary operator on primitive integers has to be evaluated directly by the arithmetic logic unit (ALU) of the underlying machine. And any operation which directly invokes the machine should belong to the same syntactic category as a print or other IO instruction: statements. The machine does not return a result; rather, it reads inputs from registers and makes the result available in a register for further computation. This is also reflected in the second surprising aspect: the operator has three instead of two arguments. The two producers p_1 and p_2 correspond to the usual arguments, but the third consumer argument c says what should happen to the result once the binary operator has been evaluated. This is similar to the continuation argument of a function in continuation-passing style. Binary operators $\odot(p_1, p_2; c)$ also display a syntactic convention we use: whenever some construct has arguments of different syntactic categories, we use a semicolon instead of a comma to separate them.

We can immediately see that the result of $[\![p_1 + p_2]\!]$ should contain the statement $+([\![p_1]\!], [\![p_2]\!]; ?)$, but we still have to figure out which consumer to plug in at the third-argument place, and how to convert this statement into a producer. We can do this with a μ-abstraction in **Core**, which turns a statement into a producer while binding a covariable α: $\mu\alpha. + ([\![p_1]\!], [\![p_2]\!]; \alpha)$.

The statement **ifz** works similarly to binary operators: It is a computation which checks if the producer p is zero and then continues with one of its two branches. These branches are also statements, indicating which computation to run after the condition has been evaluated. In the language **Fun** the two branches were terms, so we now have to find a way to transform two producers into two statements. We can do this by using a cut $\langle p \mid c \rangle$ which combines a producer and a consumer of the same type to obtain a statement in each branch: **ifz**$([\![t_1]\!], \langle [\![t_2]\!] \mid ? \rangle, \langle [\![t_3]\!] \mid ? \rangle)$. We can then use the same covariable α in both statements to represent the fact that the we want the result in either branch to return to the same point in the program; we use a surrounding μ-binding again to bind this covariable: $\mu\alpha.$**ifz**$([\![t_1]\!], \langle [\![t_2]\!] \mid \alpha \rangle, \langle [\![t_3]\!] \mid \alpha \rangle)$.

Let us now see how arithmetic expressions are evaluated. Definition 2.2 introduces the syntax of values and covalues, and shows how to reduce immediate redexes. We use a simple syntactic convention here: The metavariable for a value of terms t is $\mathfrak{t}$, the values of producers p are written $\mathfrak{p}$ and the covalues which correspond to consumers c are written $\mathfrak{c}$. We use the symbol $\triangleright$ for reduction in both **Fun** and **Core** (and write $\triangleright^*$ when multiple steps are performed at once).

Definition 2.2 (Evaluation for Arithmetic Expressions).

Fun	**Core**
$\mathfrak{t} ::= \ulcorner n \urcorner$ *Values*	$\mathfrak{p} ::= \ulcorner n \urcorner$ *Values*
	$\mathfrak{c} ::= \alpha$ *Covalues*
ifz$(\ulcorner 0 \urcorner, t_1, t_2) \triangleright t_1$	**ifz**$(\ulcorner 0 \urcorner, s_1, s_2) \triangleright s_1$
ifz$(\ulcorner n \urcorner, t_1, t_2) \triangleright t_2$ (if $n \neq 0$)	**ifz**$(\ulcorner n \urcorner, s_1, s_2) \triangleright s_2$ (if $n \neq 0$)
$\ulcorner n \urcorner \odot \ulcorner m \urcorner \triangleright \ulcorner n \odot m \urcorner$	$\odot(\ulcorner n \urcorner, \ulcorner m \urcorner; c) \triangleright \langle \ulcorner n \odot m \urcorner \mid c \rangle$
	$\langle \mu\alpha.s \mid \mathfrak{c} \rangle \triangleright s[\mathfrak{c}/\alpha]$

Values and the evaluation of redexes in **Fun** is straightforward, the only noteworthy aspect is that the two rules for $\mathbf{ifz}(\cdot, t_1, t_2)$ do not require t_1 and t_2 to be values. Thus, let us proceed with the discussion of the language **Core**.

The first interesting aspect of the language **Core** is that there are both values and covalues. This can be explained by the role that values play in operational semantics: they specify the subset of terms that we are allowed to substitute for a variable. And since we have both variables which stand for producers and covariables which stand for consumers, we need both values and covalues as the respective subsets which we are allowed to substitute for a variable or covariable.

The second interesting aspect of the language **Core** is that only statements are reduced, not producers or consumers. This substantiates our remark from above that it is statements that introduce dynamism into the language by driving computation. It also contributes to the feeling that reduction in the language is close to the evaluation of an abstract machine and that the statements of **Core** correspond to the states of such an abstract machine.

We are still faced with a small problem when we want to show that a term of **Fun** evaluates to the same result as its translation into **Core**: We have only specified the reduction for statements but not for producers. We can easily solve this problem by introducing a special covariable $\bigstar$ which acts as the "top-level" consumer of an evaluation. Using $\bigstar$ we can then evaluate the statement $\langle [\![t]\!] \mid \bigstar \rangle$ instead of the producer $[\![t]\!]$.

Example 2.1. Consider the two terms $\ulcorner 2 \urcorner * \ulcorner 3 \urcorner$ and $\mathbf{ifz}(\ulcorner 2 \urcorner, \ulcorner 5 \urcorner, \ulcorner 10 \urcorner)$ of **Fun**. Their respective translations into **Core** are $\mu\alpha. * (\ulcorner 2 \urcorner, \ulcorner 3 \urcorner; \alpha)$ and $\mu\alpha.\mathbf{ifz}(\ulcorner 2 \urcorner, \langle \ulcorner 5 \urcorner \mid \alpha \rangle, \langle \ulcorner 10 \urcorner \mid \alpha \rangle)$. When we wrap them into a statement using the top-level continuation $\bigstar$, we observe the following evaluation:

$$\langle \mu\alpha. * (\ulcorner 2 \urcorner, \ulcorner 3 \urcorner; \alpha) \mid \bigstar \rangle \; \triangleright \; *(\ulcorner 2 \urcorner, \ulcorner 3 \urcorner; \bigstar) \; \triangleright \; \langle \ulcorner 6 \urcorner \mid \bigstar \rangle$$

$$\langle \mu\alpha.\mathbf{ifz}(\ulcorner 2 \urcorner, \langle \ulcorner 5 \urcorner \mid \alpha \rangle, \langle \ulcorner 10 \urcorner \mid \alpha \rangle) \mid \bigstar \rangle \; \triangleright \; \mathbf{ifz}(\ulcorner 2 \urcorner, \langle \ulcorner 5 \urcorner \mid \bigstar \rangle, \langle \ulcorner 10 \urcorner \mid \bigstar \rangle) \; \triangleright \; \langle \ulcorner 10 \urcorner \mid \bigstar \rangle$$

We have successfully evaluated the first term to the result $\ulcorner 6 \urcorner$ and the second term to the result $\ulcorner 10 \urcorner$.

In the following, we will often leave out the first reduction step in examples, thus silently replacing the covariable bound by the outermost μ-binding with the top-level consumer $\bigstar$.

Here is a bigger problem that we haven't addressed yet. The evaluation rules in the present section do not allow to evaluate nested expressions like $(\ulcorner 2 \urcorner * \ulcorner 4 \urcorner) + \ulcorner 5 \urcorner$ in **Fun** or its translation $\mu\alpha. + (\mu\beta. * (\ulcorner 2 \urcorner, \ulcorner 4 \urcorner; \beta), \ulcorner 5 \urcorner; \alpha)$ in **Core**. We will discuss this problem and its solution in more detail in Section 3.

2.2 Let Bindings

Let-bindings are important since we can use them to eliminate duplication and make code more readable. In this section we introduce let-bindings to **Fun** for an additional reason: they allow us to introduce the second construct which gives the $\lambda\mu\tilde{\mu}$-calculus its name: $\tilde{\mu}$-abstractions.

Definition 2.3 (Let-Bindings and $\tilde{\mu}$-abstractions).

Fun		**Core**	
$t \quad ::= \quad \dots \mid \mathbf{let}\ x = t\ \mathbf{in}\ t$		$c \quad ::= \quad \dots \mid \tilde{\mu}x.s$	
		$\mathfrak{c} \quad ::= \quad \dots \mid \tilde{\mu}x.s$	
$\mathbf{let}\ x = \mathfrak{t}\ \mathbf{in}\ t \triangleright t[\mathfrak{t}/x]$		$\langle \mathfrak{p} \mid \tilde{\mu}x.s \rangle \triangleright s[\mathfrak{p}/x]$	

$$[\![\mathbf{let}\ x = t_1\ \mathbf{in}\ t_2]\!] := \mu\alpha.\langle [\![t_1]\!] \mid \tilde{\mu}x.\langle [\![t_2]\!] \mid \alpha \rangle \rangle \quad (\alpha \text{ fresh})$$

The let-bindings in **Fun** are standard and are evaluated by substituting the *value* $\mathfrak{t}$ for the variable x in the body which is a term. The analogue of a let-binding in **Fun** is a $\tilde{\mu}$-binding in **Core** which also binds a variable, with the difference that the body of a $\tilde{\mu}$-binding is a statement. It can easily be seen that $\tilde{\mu}$-bindings are the precise dual of μ-bindings that we have already introduced.

With both μ- and $\tilde{\mu}$-bindings in **Core** we have to face a potential problem, namely statements of the form $\langle \mu\alpha.s_1 \mid \tilde{\mu}\alpha.s_2 \rangle$. Such a statement is called a *critical pair* since it can potentially be reduced to both $s_1[\tilde{\mu}x.s_2/\alpha]$ and $s_2[\mu\alpha.s_1/x]$ which can be a source of non-confluence. A closer inspection of the rules shows that we avoid this pitfall and always evaluate the statement to $s_1[\tilde{\mu}x.s_2/\alpha]$. We do not allow to reduce the statement to $s_2[\mu\alpha.s_1/x]$ since only values $\mathfrak{p}$ can be substituted for variables, and $\mu\alpha.s_1$ is not a value. This restriction precisely mirrors the restriction on the evaluation of let-bindings in **Fun**. In other words, we use call-by-value evaluation order. We will address the critical pair and how it relates to different evaluation orders again in Section 5.6.

Example 2.2. Consider the term **let** $x = \ulcorner 2 \urcorner * \ulcorner 2 \urcorner$ **in** $x * x$ whose translation into **Core** is the producer $\mu\alpha.\underline{\langle \mu\beta. * (\ulcorner 2 \urcorner, \ulcorner 2 \urcorner; \beta) \mid \tilde{\mu}x.\langle \mu\gamma. * (x, x; \gamma) \mid \alpha \rangle \rangle}$. This producer contains a critical pair which we have underlined. Because we are using call-by-value, we can observe how the following reduction steps resolve the critical pair by evaluating the μ-abstraction first.

$$\langle \mu\beta. * (\ulcorner 2 \urcorner, \ulcorner 2 \urcorner; \beta) \mid \tilde{\mu}x.\langle \mu\gamma. * (x, x; \gamma) \mid \bigstar \rangle \rangle \triangleright * (\ulcorner 2 \urcorner, \ulcorner 2 \urcorner; \tilde{\mu}x.\langle \mu\gamma. * (x, x; \gamma) \mid \bigstar \rangle) \triangleright$$

$$\langle \ulcorner 4 \urcorner \mid \tilde{\mu}x.\langle \mu\gamma. * (x, x; \gamma) \mid \bigstar \rangle \rangle \triangleright \langle \mu\gamma. * (\ulcorner 4 \urcorner, \ulcorner 4 \urcorner; \gamma) \mid \bigstar \rangle \triangleright * (\ulcorner 4 \urcorner, \ulcorner 4 \urcorner; \bigstar) \triangleright \langle \ulcorner 16 \urcorner \mid \bigstar \rangle$$

We can observe that the arithmetic expression $2 * 2$ has been evaluated only once, which is precisely what we expect from call-by-value.

2.3 Top-level Definitions

We introduce recursive top-level definitions to **Fun** and **Core** for two reasons. They allow us to write more interesting examples and they illustrate a difference in how recursive calls are handled. The extension is specified in Definition 2.4.

Definition 2.4 (Top-level Definitions). We assume for both languages that $f, g, h, \ldots \in$ NAMES.

Fun			**Core**		
F	$::=$	$\mathbf{def}\ f(\overline{x}; \overline{\alpha}) := t$ *Definitions*	F	$::=$	$\mathbf{def}\ f(\overline{x}; \overline{\alpha}) := s$ *Definitions*
P	$::=$	$\emptyset \mid F, P$ *Programs*	P	$::=$	$\emptyset \mid F, P$ *Programs*
t	$::=$	$\ldots \mid f(\overline{t}; \overline{\alpha})$ *Terms*	s	$::=$	$\ldots \mid f(\overline{p}; \overline{c})$ *Statements*

$$f(\overline{\mathfrak{t}}; \overline{\alpha}) \triangleright t[\overline{\mathfrak{t}}/\overline{x}, \overline{\alpha}/\overline{\beta}] \quad (\text{if } f(\overline{x}; \overline{\beta}) := t \in P) \qquad f(\overline{\mathfrak{p}}; \overline{\mathfrak{c}}) \triangleright s[\overline{\mathfrak{p}}/\overline{x}, \overline{\mathfrak{c}}/\overline{\alpha}] \quad (\text{if } f(\overline{x}; \overline{\alpha}) := s \in P)$$

$$\llbracket \mathbf{def}\ f(\overline{x}; \overline{\alpha}) := t \rrbracket := \mathbf{def}\ f(\overline{x}; \overline{\alpha}, \alpha) := \langle \llbracket t \rrbracket \mid \alpha \rangle \quad (\alpha \text{ fresh})$$

$$\llbracket f(\overline{t}; \overline{\alpha}) \rrbracket := \mu\alpha.f(\overline{\llbracket t \rrbracket}; \overline{\alpha}, \alpha) \quad (\alpha \text{ fresh})$$

Top-level definitions should not be confused with first-class functions which will be introduced later, since they cannot be passed as an argument or returned as a result. They are a part of a program that consists of a list of such top-level definitions. The top-level definitions in **Fun** curiously also take covariables as arguments even though the language does not contain consumers; you can ignore that for now. If you remember the example from the introduction, then you might recall that we use them for passing labels, but we will only formally introduce that construct in Section 2.6.

We evaluate the call of a top-level definition by looking up the body in the program and substituting the arguments of the call for the parameters in the body of the definition. The body of a

top-level definition is a term in **Fun** and a statement in **Core**. This difference explains why we have to add an additional parameter α to every top-level definition when we translate it; this parameter α also corresponds to the additional continuation argument when we ordinarily translate a function into continuation-passing style. We could also have specified that the body of a top-level definition in **Core** should be a producer. We don't do that because when we eventually translate **Core** to machine code we want every top-level definition to become the target of a jump with arguments *without building up a function call stack*. The following example shows how this works:

Example 2.3. Using a top-level definition, we can represent the factorial function in **Core**.

$$\mathbf{def} \; \mathrm{fac}(n; \alpha) \coloneqq \mathbf{ifz}(n, \langle \ulcorner 1 \urcorner \mid \alpha \rangle, -(n, \ulcorner 1 \urcorner; \tilde{\mu} x.\mathrm{fac}(x; \tilde{\mu} r. * (n, r; \alpha))))$$

For the argument $\ulcorner 1 \urcorner$ this evaluates in the following way:

$$\mathrm{fac}(\ulcorner 1 \urcorner, \bigstar) \rhd \mathbf{ifz}(\ulcorner 1 \urcorner, \langle \ulcorner 1 \urcorner \mid \bigstar \rangle, -(\ulcorner 1 \urcorner, \ulcorner 1 \urcorner; \tilde{\mu} x.\mathrm{fac}(x; \tilde{\mu} r. * (\ulcorner 1 \urcorner, r; \bigstar))))$$

$$\rhd -(\ulcorner 1 \urcorner, \ulcorner 1 \urcorner; \tilde{\mu} x.\mathrm{fac}(x; \tilde{\mu} r. * (\ulcorner 1 \urcorner, r; \bigstar)))$$

$$\rhd \langle \ulcorner 0 \urcorner \mid \tilde{\mu} x.\mathrm{fac}(x; \tilde{\mu} r. * (\ulcorner 1 \urcorner, r; \bigstar)) \rangle$$

$$\rhd \mathrm{fac}(\ulcorner 0 \urcorner; \tilde{\mu} r. * (\ulcorner 1 \urcorner, r; \bigstar)) \qquad\qquad (*)$$

$$\rhd \mathbf{ifz}(\ulcorner 0 \urcorner, \langle \ulcorner 1 \urcorner \mid \tilde{\mu} r. * (\ulcorner 1 \urcorner, r; \bigstar) \rangle, \ldots)$$

$$\rhd \langle \ulcorner 1 \urcorner \mid \tilde{\mu} r. * (\ulcorner 1 \urcorner, r; \bigstar) \rangle$$

$$\rhd * (\ulcorner 1 \urcorner, \ulcorner 1 \urcorner; \bigstar) \rhd \langle \ulcorner 1 \urcorner \mid \bigstar \rangle$$

At the point $(*)$ of the evaluation we can now see how the recursive call is evaluated. In **Fun** this recursive call would have the form $1 * \mathrm{fac}(0)$ and require a function stack, but in **Core** we can jump to the definition of fac with the consumer $\tilde{\mu} r. * (\ulcorner 1 \urcorner, r; \bigstar)$ as an additional argument which contains the information that the result of the recursive call should be bound to the variable r and then multiplied with $\ulcorner 1 \urcorner$. Note again that this consumer argument corresponds to a continuation in continuation-passing style (in that sense it might be viewed as a reified stack) and so the basic techniques used in CPS-based intermediate representations and compilers can be applied for its implementation.

2.4 Algebraic Data and Codata Types

We now extend **Fun** and **Core** with two new features: algebraic data and codata types. Algebraic data types are familiar from most typed functional programming languages. Algebraic codata types [Hagino 1989] are a little more unusual; they are defined by a set of observations or methods called destructors and are quite similar to interfaces in object-oriented programming [Cook 2009]. We introduce them both in the same section because they help to illustrate some of the deep theoretical dualities and symmetries of the sequent calculus and the $\lambda\mu\tilde{\mu}$-calculus.

To get acquainted with our syntax, let us first briefly look at two short examples in **Fun**. The following definition calculates the sum over a List it receives as input.

$$\mathbf{def} \; \mathrm{sum}(x) \coloneqq \mathbf{case} \; x \; \mathbf{of} \; \{\mathrm{Nil} \Rightarrow \ulcorner 0 \urcorner, \mathrm{Cons}(y, ys) \Rightarrow y + \mathrm{sum}(ys)\}$$

It does so by pattern matching using the **case** ... **of** $\{...\}$ construct which is entirely standard. As an example of codata types, consider this definition:

$$\mathbf{def} \; \mathrm{repeat}(x) \coloneqq \mathbf{cocase} \; \{\mathrm{hd} \Rightarrow x, \mathrm{tl} \Rightarrow \mathrm{repeat}(x)\}$$

It constructs an infinite Stream whose elements are all the same as the input x of the function. A Stream is defined by two destructors, hd yields the head of the stream and tl yields the remaining stream without the head. The stream is constructed by copattern matching [Abel et al. 2013] using the **cocase** $\{...\}$ construct.

Definition 2.5 (Algebraic Data and Codata Types).

Fun	**Core**

$$t \quad ::= \dots \mid K(\bar{t}) \mid \mathbf{case}\ t\ \mathbf{of}\ \{\overline{K(\overline{x}) \Rightarrow t}\}$$
$$\qquad \mid t.D(\bar{t}) \mid \mathbf{cocase}\ \{\overline{D(\overline{x}) \Rightarrow t}\}$$
$$t \quad ::= \dots \mid K(\bar{t}) \mid \mathbf{cocase}\ \{\overline{D(\overline{x}) \Rightarrow t}\}$$

$$p \quad ::= \dots \mid K(\overline{p};\overline{c}) \mid \mathbf{cocase}\ \{\overline{D(\overline{x};\overline{\alpha}) \Rightarrow s}\}$$
$$c \quad ::= \dots \mid D(\overline{p};\overline{c}) \mid \mathbf{case}\ \{\overline{K(\overline{x};\overline{\alpha}) \Rightarrow s}\}$$
$$\mathfrak{p} \quad ::= \dots \mid K(\overline{\mathfrak{p}};\overline{\mathfrak{c}}) \mid \mathbf{cocase}\ \{\overline{D(\overline{x};\overline{\alpha}) \Rightarrow s}\}$$
$$\mathfrak{c} \quad ::= \dots \mid D(\overline{p};\overline{c}) \mid \mathbf{case}\ \{\overline{K(\overline{x};\overline{\alpha}) \Rightarrow s}\}$$

$$\mathbf{case}\ K(\bar{t})\ \mathbf{of}\ \{K(\overline{x}) \Rightarrow t, \dots\} \rhd t[\bar{t}/\overline{x}]$$

$$\langle K(\overline{\mathfrak{p}};\overline{\mathfrak{c}}) \mid \mathbf{case}\ \{K(\overline{x};\overline{\alpha}) \Rightarrow s, \dots\}\rangle \rhd s[\overline{\mathfrak{p}}/\overline{x};\overline{\mathfrak{c}}/\overline{\alpha}]$$

$$\mathbf{cocase}\ \{D(\overline{x}) \Rightarrow t, \dots\}.D(\bar{t}) \rhd t[\bar{t}/\overline{x}]$$

$$\langle \mathbf{cocase}\ \{D(\overline{x};\overline{\alpha}) \Rightarrow s, \dots\} \mid D(\overline{\mathfrak{p}};\overline{\mathfrak{c}})\rangle \rhd s[\overline{\mathfrak{p}}/\overline{x};\overline{\mathfrak{c}}/\overline{\alpha}]$$

$$[\![K(t_1, \dots, t_n)]\!] := K([\![t_1]\!], \dots, [\![t_n]\!])$$

$$[\![\mathbf{case}\ t\ \mathbf{of}\ \{\overline{K_i(\overline{x_{i,j}}) \Rightarrow t_i}\}]\!] := \mu\alpha.\langle [\![t]\!] \mid \mathbf{case}\ \{\overline{K_i(\overline{x_{i,j}}) \Rightarrow \langle [\![t_i]\!] \mid \alpha\rangle}\}\rangle \qquad (\alpha\ \text{fresh})$$

$$[\![t.D(t_1, \dots, t_n)]\!] := \mu\alpha.\langle [\![t]\!] \mid D([\![t_1]\!], \dots, [\![t_n]\!];\alpha)\rangle \qquad (\alpha\ \text{fresh})$$

$$[\![\mathbf{cocase}\ \{\overline{D_i(\overline{x_{i,j}}) \Rightarrow t_i}\}]\!] := \mathbf{cocase}\ \{\overline{D_i(\overline{x_{i,j}};\alpha_i) \Rightarrow \langle [\![t_i]\!] \mid \alpha_i\rangle}\} \qquad (\overline{\alpha_i}\ \text{fresh})$$

The general syntax is given in Definition 2.5. We assume fixed sets of constructors K containing at least Nil, Cons and Tup and destructors D containing at least hd, tl, fst and snd. In **Fun** we use constructors K to define both terms $K(\bar{t})$ and case expressions $\mathbf{case}\ t\ \mathbf{of}\ \{\overline{K(\overline{x}) \Rightarrow t}\}$. Destructors D of codata types are used in destructor terms $t.D(\bar{t})$ and cocase expressions $\mathbf{cocase}\ \{\overline{D(\overline{x}) \Rightarrow t}\}$. The term t in $\mathbf{case}\ t\ \mathbf{of}\ \{\overline{K(\overline{x}) \Rightarrow t}\}$ and $t.D(\bar{t})$ is called the *scrutinee* in both cases.

2.4.1 Data Types. Let us consider another example to better understand the general syntax:

$$\mathbf{def}\ \mathrm{swap}(x) := \mathbf{case}\ x\ \mathbf{of}\ \{\mathsf{Tup}(y, z) \Rightarrow \mathsf{Tup}(z, y)\}$$

The function swap takes a Pair and swaps its elements. To do so, it pattern matches on its input using the $\mathbf{case}\ t\ \mathbf{of}\ \{\overline{K(\overline{x}) \Rightarrow t}\}$ construct, and constructs a tuple using a constructor $K(\bar{t})$, where K is specialized to Tup. Our syntax is quite general, so it is easy to extend it with new constructors; any such extension only requires that we also add corresponding typing rules (Section 4).

In **Core**, algebraic data types are mostly handled in the same way as in **Fun**. The main difference is that the scrutinee is no longer a part of a case expression. Instead, the case expression is a consumer and the scrutinee is a producer, which are then combined in a statement. This is exactly what is done in the translation. When a case and a constructor meet, there is an opportunity for computation, *consuming* the constructed term and continuing with the corresponding right-hand side of the case expression. This also explains our terminology of *producers* and *consumers*. Constructors create, or in other words, *produce* data structures while cases destroy, or *consume* them.

There is another difference, however. Constructors in **Core** can now also take consumers as arguments which is not the case in **Fun**. An example of this is the negation type of a type τ which can be formulated as a data type with one constructor taking a consumer of type τ as an argument. A program making use of this type can be found in section 7.2 of Ostermann et al. [2022].

Fun is a call-by-value language which manifests itself in that a value of an algebraic data type consists of a constructor applied to other values. A case expression $\mathbf{case}\ t\ \mathbf{of}\ \{\dots\}$ can only be evaluated if the scrutinee t is a value, so this means that it must be a constructor whose arguments are all values in the evaluation rule.

Evaluation in **Core** is done the same way, only with the scrutinee term changed to be the producer of a cut. Note that all consumers in **Core** are covalues (which is why the arguments of destructors in the definition of covalues are not in Fraktur font), so in order for a constructor term to be a value, only its producer arguments need to be values. This also means that the requirement for the consumer arguments of the constructor to be covalues is vacuously satisfied in the evaluation rule in **Core**.

Example 2.4. The translation of swap (including a simplification) is given by

$$\mathbf{def}\ \mathsf{swap}(x; \alpha) \coloneqq \langle x \mid \mathbf{case}\ \{\mathsf{Tup}(y, z) \Rightarrow \langle \mathsf{Tup}(z, y) \mid \alpha \rangle\}\rangle$$

Evaluating with an argument $\mathsf{Tup}(\ulcorner 2 \urcorner, \ulcorner 3 \urcorner)$ and $\bigstar$ then proceeds as we would expect

$$\langle \mathsf{Tup}(\ulcorner 2 \urcorner, \ulcorner 3 \urcorner) \mid \mathbf{case}\ \{\mathsf{Tup}(y, z) \Rightarrow \langle \mathsf{Tup}(z, y) \mid \bigstar \rangle\}\rangle \triangleright \langle \mathsf{Tup}(\ulcorner 3 \urcorner, \ulcorner 2 \urcorner) \mid \bigstar \rangle$$

2.4.2 *Codata Types.* To illustrate the syntax for codata types further, consider the definition

$$\mathbf{def}\ \mathsf{swap\_lazy}(x) \coloneqq \mathbf{cocase}\ \{\mathsf{fst} \Rightarrow x.\mathsf{snd}, \mathsf{snd} \Rightarrow x.\mathsf{fst}\}$$

swap_lazy takes a lazy pair (LPair), which is defined by its projections fst and snd, and swaps its elements. It does so with a copattern match **cocase** $\{\overline{D_i(\overline{x}) \Rightarrow t_i}\}$ which invokes the opposite destructor on the original pair in each branch. With a destructor invocation $t.D(\overline{t})$, where D is specialized to fst or snd, we can then obtain the corresponding component of the new pair.

For codata in **Fun**, the scrutinee is located in the destructor term instead of the cocase, inverse to data types. So now destructors are the consumers and cocases are the producers. This is mirrored in the translation which again separates the scrutinee, since in **Core** codata types and copattern matching are perfectly dual to data types and pattern matching.

All the destructors we have used here do not have producer parameters, but this is just due to the selection of examples. In the next section, we will see an example of a destructor with a producer parameter. Moreover, during the translation each destructor is endowed with an additional consumer parameter which again determines how execution continues after the destructor was invoked (and is thus bound by a surrounding μ). For constructors this is not necessary, as we can use the same consumer variable directly in each branch of a **case** (similar to **ifz**[2]) because the scrutinee and the **case** are in the same expression. Destructors (and also constructors) in **Core** can even have more than one consumer parameter. An example of this is given in Section 5.7.

Evaluation is done analogous to data types, with the roles of cases and constructors reversed for cocases and destructors. Note, however, that for evaluation in **Core** the producer arguments of the destructor also have to be values, so it is not sufficient for the destructor to be a covalue (which it always is). We will come back to this subtlety in Section 3.

Example 2.5. Translating swap_lazy is done analogously to swap.

$$\mathbf{def}\ \mathsf{swap\_lazy}(x; \alpha) \coloneqq \langle \mathbf{cocase}\ \{\mathsf{fst}(\beta) \Rightarrow \langle x \mid \mathsf{snd}(\beta) \rangle, \mathsf{snd}(\beta) \Rightarrow \langle x \mid \mathsf{fst}(\beta) \rangle\} \mid \alpha \rangle$$

Now take $p = \mathbf{cocase}\ \{\mathsf{fst}(\alpha) \Rightarrow \langle \ulcorner 1 \urcorner \mid \alpha \rangle, \mathsf{snd}(\alpha) \Rightarrow *(\ulcorner 2 \urcorner, \ulcorner 3 \urcorner; \alpha)\}$ and evaluate swap_lazy with snd to retrieve its first element:

$$\mathsf{swap\_lazy}(p; \mathsf{snd}(\bigstar)) \triangleright \langle \mathbf{cocase}\ \{\mathsf{fst}(\beta) \Rightarrow \langle p \mid \mathsf{snd}(\beta) \rangle, \mathsf{snd}(\beta) \Rightarrow \langle p \mid \mathsf{fst}(\beta) \rangle\} \mid \mathsf{snd}(\bigstar) \rangle$$
$$\triangleright \langle p \mid \mathsf{fst}(\bigstar) \rangle \triangleright \langle \ulcorner 1 \urcorner \mid \bigstar \rangle$$

Because **cocase**s are values regardless of their right-hand sides (in contrast to constructors), we can apply the destructor snd without first evaluating the product $*(\ulcorner 2 \urcorner, \ulcorner 3 \urcorner; \alpha)$. For pairs, we could

---

[2]We could have modelled **ifz** as a **case**, too, by modeling numbers as a data type. But since **ifz** corresponds to a machine instruction quite directly, it is natural to make it a statement, as explained in Section 2.1.

not do this, as $\mathsf{Tup}(\ulcorner 1 \urcorner, *(\ulcorner 2 \urcorner, \ulcorner 3 \urcorner; \alpha))$ is not a value, so its arguments have to be evaluated first. This is why this codata type is called *lazy pair*, as it allows to not evaluate its contents in contrast to regular pairs.

This section showed an important property of **Core** which does not hold for **Fun**. The data and codata types of **Core** are completely symmetric: the syntax for cases is the same as the syntax for cocases and the same is true for constructors and destructors. The reason for this deep symmetry is the same reason that makes the sequent calculus more symmetric than natural deduction, but in Definition 2.5 we can observe it in a programming language.

2.5 First-Class Functions

A core feature that we have omitted until now are first-class functions which are characterized by lambda abstractions $\lambda x.t$ and function applications $t_1 \ t_2$. But first-class functions do not add any expressive power to a language with codata types, since codata types are a more general concept which subsumes functions as a special case. We could therefore implement lambda abstractions and function applications as syntactic sugar in both **Fun** and **Core**. This is incidentally also what the developers of Java did when they introduced lambdas to the language [Goetz et al. 2014]. We introduce lambda abstractions and function application to the syntax of **Fun** and desugar them to cocases and destructors of a codata type with an ap destructor during the translation to **Core**.

Definition 2.6 (First-Class Functions).

$$
\begin{array}{c}
\textbf{Fun} \\
\begin{array}{lll}
t & ::= & \dots \mid \lambda x.t \mid t\ t \\
\mathsf{t} & ::= & \dots \mid \lambda x.t
\end{array} \\
(\lambda x.t)\ \mathsf{t} \rhd t[\mathsf{t}/x]
\end{array}
\qquad
\begin{array}{c}
\textbf{Core} \\
D \in \{\dots, \mathsf{ap}\}
\end{array}
$$

$$
\begin{aligned}
\llbracket \lambda x.t \rrbracket &:= \textbf{cocase}\ \{\mathsf{ap}(x; \alpha) \Rightarrow \langle \llbracket t \rrbracket \mid \alpha \rangle\} && (\alpha\ \text{fresh}) \\
\llbracket t_1\ t_2 \rrbracket &:= \mu\alpha.\langle \llbracket t_1 \rrbracket \mid \mathsf{ap}(\llbracket t_2 \rrbracket; \alpha)\rangle && (\alpha\ \text{fresh})
\end{aligned}
$$

Example 2.6. Consider the term $(\lambda x.x * x)\ \ulcorner 2 \urcorner$ in **Fun**. We can translate this term and evaluate it in **Core** as follows:

$$\langle \textbf{cocase}\ \{\mathsf{ap}(x, \beta) \Rightarrow \langle \mu\gamma. * (x, x; \gamma) \mid \beta \rangle\} \mid \mathsf{ap}(\ulcorner 2 \urcorner; \bigstar)\rangle \rhd \langle \mu\gamma. * (\ulcorner 2 \urcorner, \ulcorner 2 \urcorner; \gamma) \mid \bigstar \rangle \rhd^* \langle \ulcorner 4 \urcorner \mid \bigstar \rangle$$

2.6 Control Operators

Finally, we add the feature that we used in the motivating example in the introduction: labels and jumps. We have to extend **Fun** with **label** and **goto** constructs but since we can translate them locally to μ-bindings we don't have to add anything to **Core**.

Definition 2.7 (Control Operators).

$$t \quad ::= \quad \dots \mid \textbf{label}\ \alpha\ \{t\} \mid \textbf{goto}(t; \alpha)$$

$$\llbracket \textbf{label}\ \alpha\ \{t\} \rrbracket := \mu\alpha.\langle \llbracket t \rrbracket \mid \alpha \rangle \qquad \llbracket \textbf{goto}(t; \alpha) \rrbracket := \mu\beta.\langle \llbracket t \rrbracket \mid \alpha \rangle \quad (\beta\ \text{fresh})$$

A term **label** $\alpha\ \{t\}$ binds a covariable α in the term t and thereby provides a location to which a **goto** used within t can jump. Such a **goto**$(t; \alpha)$ takes the location α as an argument, as well as the term t that should be used to continue the computation at the location where α was bound. It is a

bit tricky to write down precisely how the evaluation of **label** and **goto** works, but the following two rules are a good approximation, where we assume that α does not occur free in t:

$$\textbf{label } \alpha \text{ \{t\}} \triangleright \text{t} \qquad\qquad \textbf{label } \alpha \text{ \{\dots \textbf{goto}(t; \alpha) \dots\}} \triangleright \text{t}$$

The left rule says that when the labeled term t can be evaluated to a value t without ever using a **goto**, then we can discard the surrounding **label**. The rule on the right says that if we do have a **goto** which jumps to the **label** α with a value t, then we discard everything between the **label** and the **goto** and continue the computation with this value t. In order to make this second rule precise, we have to make explicit what we only indicate with the ellipses separating the label from the jump; we will do so in Section 3.

Example 2.7. In the introduction, we used the example of a fast multiplication function which multiplies all the elements of a list and short-circuits the computation if it encounters a zero. As we have allowed top-level definitions to pass covariables as arguments, we can now write the example of the introduction.

> **def** $\text{mult}(l) := \textbf{label } \alpha \text{ \{mult'}(l; \alpha)\}$
>
> **def** $\text{mult'}(l; \alpha) := \textbf{case } l \textbf{ of } \{\text{Nil} \Rightarrow 1, \text{Cons}(x, xs) \Rightarrow \textbf{ifz}(x, \textbf{goto}(0; \alpha), x * \text{mult'}(xs; \alpha))\}$

When we translate to **Core** and simplify the resulting term, we get the result:

> **def** $\text{mult}(l; \alpha) := \text{mult'}(l; \alpha, \alpha)$
>
> **def** $\text{mult'}(l; \alpha, \beta) :=$
>
> $\qquad \langle l \mid \textbf{case } \{\text{Nil} \Rightarrow \langle 1 \mid \beta \rangle, \text{Cons}(x, xs) \Rightarrow \textbf{ifz}(x, \langle 0 \mid \alpha \rangle, *(x, \mu\gamma.\text{mult'}(xs; \alpha, \gamma); \beta))\}\rangle$

This is almost the result we have seen in the introduction. The only difference is that the recursive call to mult' is nested inside the multiplication. This is the same problem we have seen with nested arithmetic operations at the end of Section 2.1 and we will address it in the next section.

The **label**/**goto** control operator we have introduced in this subsection is of course named after the goto instructions and labels which can be found in many imperative programming languages. Our adaption to the context of functional programming languages is similar to classical control operators (see Section 5.3 for a more precise discussion) such as J [Landin 1965] or let/cc (also known as **escape**) [Reynolds 1972]; the programming language Scala also provides the closely related boundary/break[3] where a boundary marks a block of code to which the programmer can jump with a break instruction. One central property of this control effect is that it is lexically scoped, since the label names α are passed around lexically and can be shadowed. This distinguishes them from dynamically scoped control operators like the exception mechanisms found in many programming languages like Java or C++. (A dynamically scoped variant of our control operator would omit the label names, and the jump in **label** $\{\dots \textbf{goto}(t) \dots\}$ would return to the nearest enclosing label at runtime.) We follow the more recent reappraisal of lexically scoped control effects, for example by Zhang et al. [2016] in the case of exceptions or by Brachthäuser et al. [2020] in the case of effect handlers and delimited continuations.

3 Evaluation Within a Context

At the end of Section 2.1 we ran into the problem that we cannot yet fully evaluate the term $(\ulcorner2\urcorner * \ulcorner4\urcorner) + \ulcorner5\urcorner$ in **Fun** or its translation $\mu\alpha. +(\mu\beta. *(\ulcorner2\urcorner, \ulcorner4\urcorner; \beta), \ulcorner5\urcorner; \alpha)$ in **Core** with the rules that are available to us: we are stuck. In this section, we finally address this problem. We are going to show how we can evaluate subexpressions of **Fun** in Section 3.1, but since we are ultimately

---

[3]See scala-lang.org/api/3.3.0/scala/util/boundary$.html.

more interested in compiling programs into **Core** to optimize and reduce those programs, we are spending more time on the problem for **Core** in Section 3.2.

3.1 Evaluation Contexts for Fun

The problem with evaluating the term $(\ulcorner 2\urcorner * \ulcorner 4\urcorner) + \ulcorner 5\urcorner$ is that the available rules only allow to reduce direct redexes and not redexes that are nested somewhere within a term. Evaluation contexts solve this problem by specifying the locations within a term which are in evaluation position. In our example, the term $(\ulcorner 2\urcorner * \ulcorner 4\urcorner) + \ulcorner 5\urcorner$ can be factored into the evaluation context $\square + \ulcorner 5\urcorner$ and the redex $\ulcorner 2\urcorner * \ulcorner 4\urcorner$. We can then use the old rules to reduce this redex to $\ulcorner 8\urcorner$ and then plug this result back into the evaluation context, which yields the new term $\ulcorner 8\urcorner + \ulcorner 5\urcorner$. The syntax of evaluation contexts is given in Definition 3.1.

Definition 3.1 (Evaluation Contexts). Evaluation contexts E are defined as:

$$
\begin{aligned}
E \quad ::= \quad & \square \mid E \odot t \mid t \odot E \mid \mathbf{ifz}(E, t, t) \mid \mathbf{let}\ x = E\ \mathbf{in}\ t \mid f(\bar{t}, E, \bar{t}) \mid K(\bar{t}, E, \bar{t}) \\
\mid \quad & \mathbf{case}\ E\ \mathbf{of}\ \{\overline{K(\overline{x}) \Rightarrow t}\} \mid E\ t \mid t\ E \mid E.D(\bar{t}) \mid t.D(\bar{t}, E, \bar{t}) \mid \mathbf{label}\ \alpha\ \{E\} \mid \mathbf{goto}(E; \alpha)
\end{aligned}
$$

These evaluation contexts also allow us to specify formally the second approximate evaluation rule of the **label** and **goto** constructs from Section 2.6:

$$E[\mathbf{label}\ \alpha\ \{E'[\mathbf{goto}(t; \alpha)]\}] \triangleright E[t]$$

Here we again assume that α does not occur free in t and moreover that the inner evaluation context E' does not contain another **label** construct. For the full operational semantics of **label**/**goto** we also need to handle the cases where α can occur free in t and where E' can contain other **label**s. Otherwise, we could get stuck during evaluation even for closed and well-typed terms, i.e., the progress theorem (see Theorem 4.1 in Section 4.3) would not hold. As the full semantics is in essence that of other classical control operators (i.p., let/cc; also see the discussion in Section 5.3) and requires some more formalism, we do not give it here and instead refer the interested reader to the brief discussion in Appendix C.

With evaluation contexts, we finally have a working and precise operational semantics for **Fun** (apart from the approximate rules for **label** and **goto**) which we can use to reason about programs. Unfortunately, it is wildly inefficient to implement an evaluator which uses evaluation contexts in the way described above. The reason for this inefficiency is that we very elegantly specified how a term can be factored into an evaluation context and a redex, but the evaluator which implements this behavior has to search for the next redex after every single evaluation step. We will see in the next section that we have a better solution once our programs are compiled into **Core**.

3.2 Focusing on Evaluation in Core

Let us now come back to the problem in **Core** and find a solution for the stuck term $\mu\alpha. + (\mu\beta. * (\ulcorner 2\urcorner, \ulcorner 4\urcorner; \beta), \ulcorner 5\urcorner; \alpha)$. We know that we have to evaluate $\mu\beta. * (\ulcorner 2\urcorner, \ulcorner 4\urcorner; \beta)$ next and then somehow plug the intermediate result into the hole $[\cdot]$ in the producer $\mu\alpha. + ([\cdot], \ulcorner 5\urcorner; \alpha)$. If we give the intermediate result the name x and play around with cuts, μ-bindings and $\tilde{\mu}$ bindings, we might discover that we can recombine all these parts in the following way:

$$\mu\alpha.\langle \mu\beta. * (\ulcorner 2\urcorner, \ulcorner 4\urcorner; \beta) \mid \tilde{\mu}x. + (x, \ulcorner 5\urcorner; \alpha)\rangle$$

This term looks a bit mysterious, but the transformation corresponds roughly to what happens when we translate the term **let** $x = 2 * 4$ **in** $x + 5$ instead of $(2 * 4) + 5$ into **Core**. That is, we have lifted a subcomputation to the outside of the term we are evaluating. This kind of transformation is

called *focusing* [Andreoli 1992; Curien and Munch-Maccagnoni 2010] and we use it to solve the problem with stuck terms in **Core**. We can see that it worked in our example because the term now fully evaluates to its normal form.

Example 3.1. The producer $\mu\alpha.\langle\mu\beta. * (\ulcorner 2\urcorner, \ulcorner 4\urcorner; \beta) \mid \tilde{\mu}x. + (x, \ulcorner 5\urcorner; \alpha)\rangle$ reduces as follows:

$$\langle\mu\beta. * (\ulcorner 2\urcorner, \ulcorner 4\urcorner; \beta) \mid \tilde{\mu}x. + (x, \ulcorner 5\urcorner; \bigstar)\rangle \triangleright *(\ulcorner 2\urcorner, \ulcorner 4\urcorner; \tilde{\mu}x. + (x, \ulcorner 5\urcorner; \bigstar))$$
$$\triangleright \langle\ulcorner 8\urcorner \mid \tilde{\mu}x. + (x, \ulcorner 5\urcorner; \bigstar)\rangle$$
$$\triangleright +(\ulcorner 8\urcorner, \ulcorner 5\urcorner; \bigstar) \triangleright \langle\ulcorner 13\urcorner \mid \bigstar\rangle$$

Once we have settled on focusing, we have another choice to make: Do we want to use this trick during the evaluation of a statement or as a preprocessing step before we start with the evaluation? These two alternatives are called dynamic and static focusing.

Dynamic Focusing With dynamic focusing [Wadler 2003] we add additional evaluation rules, usually called ς-rules, to lift sub-computations to the outside of the statement we are evaluating.

Static Focusing For static focusing [Curien and Herbelin 2000] we perform a transformation on the code before we start evaluating it. This results in a focused normal form which is a subset of the syntax of **Core** that we have described so far.

Dynamic focusing is great for reasoning about the meaning of programs, but static focusing is more efficient if we are interested in compiling and running programs. For this reason, we only consider static focusing in what follows.

Definition 3.2 (Static Focusing). Static focusing is done using the following rules:

<div align="center">

Producers

$$\begin{aligned}
\mathcal{F}(\ulcorner n\urcorner) &:= \ulcorner n\urcorner \\
\mathcal{F}(x) &:= x \\
\mathcal{F}(\mu\alpha.s) &:= \mu\alpha.\mathcal{F}(s) \\
\mathcal{F}(K(\overline{\mathfrak{p}}, p, \overline{p}; \overline{c})) &:= \mu\alpha.\langle\mathcal{F}(p) \mid \tilde{\mu}x.\langle\mathcal{F}(K(\overline{\mathfrak{p}}, x, \overline{p}, \overline{c})) \mid \alpha\rangle\rangle \quad (p \text{ not a value}) \\
\overline{\mathcal{F}(K(\overline{\mathfrak{p}}; \overline{c}))} &:= K(\overline{\mathcal{F}(\mathfrak{p})}; \overline{\mathcal{F}(c)}) \\
\mathcal{F}(\mathbf{cocase}\ \{\overline{D(\overline{x}; \overline{\alpha}) \Rightarrow s}\}) &:= \mathbf{cocase}\ \{\overline{D(\overline{x}; \overline{\alpha}) \Rightarrow \mathcal{F}(s)}\}
\end{aligned}$$

Consumers

$$\begin{aligned}
\mathcal{F}(\alpha) &:= \alpha \\
\mathcal{F}(\tilde{\mu}x.s) &:= \tilde{\mu}x.\mathcal{F}(s) \\
\mathcal{F}(\mathbf{case}\ \{\overline{K(\overline{x}; \overline{\alpha}) \Rightarrow s}\}) &:= \mathbf{case}\ \{\overline{K(\overline{x}; \overline{\alpha}) \Rightarrow \mathcal{F}(s)}\} \\
\mathcal{F}(D(\overline{\mathfrak{p}}, p, \overline{p}, \overline{c})) &:= \tilde{\mu}y.\langle\mathcal{F}(p) \mid \tilde{\mu}x.\langle y \mid \mathcal{F}(D(\overline{\mathfrak{p}}, x, \overline{p}; \overline{c}))\rangle\rangle \quad (p \text{ not a value}) \\
\mathcal{F}(D(\overline{\mathfrak{p}}; \overline{c})) &:= D(\overline{\mathcal{F}(\mathfrak{p})}; \overline{\mathcal{F}(c)})
\end{aligned}$$

Statements

$$\begin{aligned}
\mathcal{F}(\langle p \mid c\rangle) &:= \langle\mathcal{F}(p) \mid \mathcal{F}(c)\rangle \\
\mathcal{F}(\odot(p_1, p_2, c)) &:= \langle\mathcal{F}(p_1) \mid \tilde{\mu}x.\mathcal{F}(\odot(x, p_2, c))\rangle \quad (p_1 \text{ not a value}) \\
\mathcal{F}(\odot(\mathfrak{p}, p, c)) &:= \langle\mathcal{F}(p) \mid \tilde{\mu}x.\mathcal{F}(\odot(\mathfrak{p}, x, c))\rangle \quad (p \text{ not a value}) \\
\mathcal{F}(\odot(\mathfrak{p}_1, \mathfrak{p}_2, c)) &:= \odot(\mathcal{F}(\mathfrak{p}_1), \mathcal{F}(\mathfrak{p}_2), \mathcal{F}(c)) \\
\mathcal{F}(\mathbf{ifz}(p, s_1, s_2)) &:= \langle\mathcal{F}(p) \mid \tilde{\mu}x.\mathbf{ifz}(x, s_1, s_2)\rangle \quad (p \text{ not a value}) \\
\mathcal{F}(\mathbf{ifz}(\mathfrak{p}, s_1, s_2)) &:= \mathbf{ifz}(\mathcal{F}(\mathfrak{p}), \mathcal{F}(s_1), \mathcal{F}(s_2)) \\
\mathcal{F}(\mathrm{f}(\overline{\mathfrak{p}}, p, \overline{p}; \overline{c})) &:= \langle\mathcal{F}(p) \mid \tilde{\mu}x.\mathcal{F}(\mathrm{f}(\overline{\mathfrak{p}}, x, \overline{p}; \overline{c}))\rangle \quad (p \text{ not a value}) \\
\mathcal{F}(\mathrm{f}(\overline{\mathfrak{p}}; \overline{c})) &:= \mathrm{f}(\overline{\mathcal{F}(\mathfrak{p})}, \overline{\mathcal{F}(c)})
\end{aligned}$$

</div>

The complete rules for static focusing are presented in Definition 3.2. Most of these rules are only concerned with performing the focusing transformation on all subexpressions, but some of the clauses where something interesting happens are the clauses for binary operators:

$$\mathcal{F}(\odot(p_1, p_2, c)) \coloneqq \langle \mathcal{F}(p_1) \mid \tilde{\mu}x.\mathcal{F}(\odot(x, p_2, c)) \rangle \quad (p_1 \text{ not a value})$$

$$\mathcal{F}(\odot(\mathfrak{p}, p, c)) \coloneqq \langle \mathcal{F}(p) \mid \tilde{\mu}x.\mathcal{F}(\odot(\mathfrak{p}, x, c)) \rangle \quad (p \text{ not a value})$$

$$\mathcal{F}(\odot(\mathfrak{p}_1, \mathfrak{p}_2, c)) \coloneqq \odot(\mathcal{F}(\mathfrak{p}_1), \mathcal{F}(\mathfrak{p}_2), \mathcal{F}(c))$$

The first two clauses look for the arguments of the binary operator $\odot$ which are not values and use the trick described above to lift them to the outside. Focusing is invoked recursively until the binary operator is only applied to values and the third clause comes into play. This third clause then applies the focusing transformation to all arguments of the binary operator. The clauses for constructors, destructors, **ifz** and calls to top-level definitions work in precisely the same way as those for binary operators. It is noteworthy that by focusing the producer arguments of destructors we guarantee that the evaluation rule for codata types can fire. If we had not required the producer arguments to be values in that rule (but only that the destructor is a covalue), we could easily introduce an unfocused term again by substituting a non-value for a variable.

The focusing transformation described in Definition 3.2 is not ideal since it creates a lot of administrative redexes. As an example, consider how the statement defining mult' from Example 2.7 is focused:

$$\mathcal{F}(\langle l \mid \textbf{case } \{\text{Nil} \Rightarrow \langle 1 \mid \beta \rangle, \text{Cons}(x, xs) \Rightarrow \textbf{ifz}(x, \langle 0 \mid \alpha \rangle, *(x, \mu\gamma.\text{mult'}(xs; \alpha, \gamma); \beta))\} \rangle)$$

$$= \langle l \mid \textbf{case } \{\text{Nil} \Rightarrow \langle 1 \mid \beta \rangle, \text{Cons}(x, xs) \Rightarrow \textbf{ifz}(x, \langle 0 \mid \alpha \rangle, \langle \mu\gamma.\text{mult'}(xs; \alpha, \gamma) \mid \tilde{\mu}z. * (x, z; \beta)\rangle)\} \rangle$$

Focusing has introduced the administrative redex $\langle \mu\gamma.\text{mult'}(xs; \alpha, \gamma) \mid \tilde{\mu}z. * (x, z; \beta) \rangle$ in the second statement of the **ifz**. After reducing this redex to mult'$(xs; \alpha, \tilde{\mu}z. * (x, z; \beta))$, we finally arrive at the result from the introduction. In the implementation, we solve this problem by statically reducing administrative redexes in a simplification step, but it is also possible to come up with a more elaborate definition of focusing which does not create them in the first place. Such an optimized focusing transformation is, however, much less transparent than the one we have described.

4 Typing Rules

In this section, we introduce the typing rules for **Fun** in Section 4.1 and for **Core** in Section 4.2. In Section 4.3 we state type soundness for both languages and prove that the translation from **Fun** to **Core** preserves the typeability of programs. We use the same constructors, destructors, types and typing contexts for both **Fun** and **Core**, which are summarized in Definition 4.1. Note that we distinguish between producer and consumer variables in the typing contexts, which we indicate with the prd and cns annotations.

Definition 4.1 (Types and Typing Contexts).

K	$\Coloneqq$	Nil | Cons | Tup	*Constructors*
D	$\Coloneqq$	hd | tl | fst | snd | ap	*Destructors*
τ	$\Coloneqq$	**Int** | List(τ) | Pair(τ, τ) | Stream(τ) | LPair(τ, τ) | $\tau \rightarrow \tau$	*Types*
Γ	$\Coloneqq$	$\emptyset$ | $\Gamma, x :^{\text{prd}} \tau$ | $\Gamma, \alpha :^{\text{cns}} \tau$	*Typing Contexts*

We specialize the rules for data types to the concrete types Pair and List, and the rules for codata types to LPair, Stream and functions $\sigma \rightarrow \tau$. A realistic programming language would use type declarations introduced by the programmer to typecheck data and codata types instead of using these special cases. But the formalization of such a general mechanism for specifying

data and codata types makes the typing rules less readable. This kind of mechanism for specifying algebraic data and codata types in sequent-calculus-based languages can be found in [Downen et al. 2015] or [Downen and Ariola 2020, section 8]. In all of the typing rules below we assume that we have a program environment which contains type declarations for all the definitions contained in the program, but don't explicitly thread this program environment through each of the typing rules.

4.1 Typing Rules for Fun

We don't discuss the typing rules for **Fun** in detail since they are mostly standard. Instead, we provide the full rules in Appendix B. The language **Fun** only has one syntactic category, terms, so we only need one typing judgment $\Gamma \vdash t : \tau$. This typing judgment says that in the context Γ (which contains type assignments for both variables and covariables) the term t has type τ. The only two interesting rules concern the control operators **label** and **goto**:

$$\frac{\Gamma, \alpha :^{\text{cns}} \tau \vdash t : \tau}{\Gamma \vdash \textbf{label } \alpha \: \{t\} : \tau} \: \text{Label} \qquad\qquad \frac{\Gamma \vdash t : \tau \qquad \alpha :^{\text{cns}} \tau \in \Gamma}{\Gamma \vdash \textbf{goto}(t; \alpha) : \tau'} \: \text{Goto}$$

In the rule Label we add the covariable $\alpha :^{\text{cns}} \tau$ to the typing context which is used to typecheck the term t. The labeled expression **label** $\alpha \: \{t\}$ can return in only one of two ways: either the term t is evaluated to a value and returned, or a jump instruction is used to jump to the label α. For this reason, the term t and the label α must have the same type τ, which is also the type for the labeled expression itself.

In the rule Goto we require that the covariable α is in the context with type τ, and that the term t can be typechecked with the same type. The term **goto**$(t; \alpha)$ itself can be used at any type τ' because it does not return to its immediately surrounding context.

4.2 Typing Rules for Core

The complete typing rules for **Core** are given in Figure 2, but we will present them step by step. We now have producers, consumers and statements as different syntactic categories. For each of these categories, we use a separate judgment form:

Producers The judgment $\Gamma \vdash p :^{\text{prd}} \tau$ says that the producer p has type τ in context Γ.
Consumers The judgment $\Gamma \vdash c :^{\text{cns}} \tau$ says that the consumer c has type τ in context Γ.
Statements The judgment $\Gamma \vdash s$ says that the statement s is well-typed in context Γ. In contrast to producers and consumers, statements do not have a type.

All typing judgments are also implicitly indexed by the program P containing the top-level definitions. However, as these definitions are only needed when typechecking their calls (rule Call), we usually omit the index from the presentation.

The three different judgments can be illustrated by the rules for variables, covariables and cuts. In the rules Var_1 and Var_2 we check that a variable or covariable is contained in the typing context Γ and then type the variable as a producer or the covariable as a consumer. The rule Cut combines a producer p and consumer c of the same type τ into the statement $\langle p \mid c \rangle$ which does not have a type.

$$\frac{x :^{\text{prd}} \tau \in \Gamma}{\Gamma \vdash x :^{\text{prd}} \tau} \: \text{Var}_1 \qquad\qquad \frac{\alpha :^{\text{cns}} \tau \in \Gamma}{\Gamma \vdash \alpha :^{\text{cns}} \tau} \: \text{Var}_2 \qquad\qquad \frac{\Gamma \vdash p :^{\text{prd}} \tau \qquad \Gamma \vdash c :^{\text{cns}} \tau}{\Gamma \vdash \langle p \mid c \rangle} \: \text{Cut}$$

The two unusual constructs which are central to **Core** and give the $\lambda\mu\tilde{\mu}$-calculus its name are the μ- and $\tilde{\mu}$-abstractions. A μ-abstraction $\mu\alpha.s$ abstracts over a consumer α of type τ in the statement s and is typed as a producer of type τ. A $\tilde{\mu}$-abstraction $\tilde{\mu}x.s$ abstracts over a producer x of type τ and is typed as a consumer of type τ, which can be seen in the following two rules.

$$\frac{\Gamma, \alpha :^{\text{cns}} \tau \vdash s}{\Gamma \vdash \mu\alpha.s :^{\text{prd}} \tau} \; \mu \qquad \frac{\Gamma, x :^{\text{prd}} \tau \vdash s}{\Gamma \vdash \tilde{\mu}x.s :^{\text{cns}} \tau} \; \tilde{\mu} \qquad \frac{x :^{\text{prd}} \tau \in \Gamma}{\Gamma \vdash x :^{\text{prd}} \tau} \; \text{Var}_1 \qquad \frac{\alpha :^{\text{cns}} \tau \in \Gamma}{\Gamma \vdash \alpha :^{\text{cns}} \tau} \; \text{Var}_2$$

$$\frac{\Gamma \vdash p :^{\text{prd}} \tau \qquad \Gamma \vdash c :^{\text{cns}} \tau}{\Gamma \vdash \langle p \mid c \rangle} \; \text{Cut} \qquad \frac{\Gamma \vdash p :^{\text{prd}} \textbf{Int} \qquad \Gamma \vdash s_1 \qquad \Gamma \vdash s_2}{\Gamma \vdash \textbf{ifz}(p, s_1, s_2)} \; \text{IfZ}$$

$$\frac{}{\Gamma \vdash \ulcorner n \urcorner :^{\text{prd}} \textbf{Int}} \; \text{Lit} \qquad \frac{\Gamma \vdash p_1 :^{\text{prd}} \textbf{Int} \qquad \Gamma \vdash p_2 :^{\text{prd}} \textbf{Int} \qquad \Gamma \vdash c :^{\text{cns}} \textbf{Int}}{\Gamma \vdash \odot(p_1, p_2; c)} \; \text{binop}$$

$$\frac{\textbf{def } f(\overline{x_i :^{\text{prd}} \tau_i}; \overline{\alpha_j :^{\text{cns}} \tau_j}) \in P \qquad \overline{\Gamma \vdash p_i :^{\text{prd}} \tau_i} \qquad \overline{\Gamma \vdash c_j :^{\text{cns}} \tau_j}}{\Gamma \vdash_P f(\overline{p_i}; \overline{c_j})} \; \text{Call}$$

$$\frac{\Gamma \vdash s_1 \qquad \Gamma, x :^{\text{prd}} \tau, xs :^{\text{prd}} \text{List}(\tau) \vdash s_2}{\Gamma \vdash \textbf{case } \{\text{Nil} \Rightarrow s_1, \text{Cons}(x, xs) \Rightarrow s_2\} :^{\text{cns}} \text{List}(\tau)} \; \text{Case-List}$$

$$\frac{}{\Gamma \vdash \text{Nil} :^{\text{prd}} \text{List}(\tau)} \; \text{Nil} \qquad \frac{\Gamma \vdash t_1 :^{\text{prd}} \tau \qquad \Gamma \vdash t_2 :^{\text{prd}} \text{List}(\tau)}{\Gamma \vdash \text{Cons}(t_1, t_2) :^{\text{prd}} \text{List}(\tau)} \; \text{Cons}$$

$$\frac{\Gamma \vdash t_1 :^{\text{prd}} \tau_1 \qquad \Gamma \vdash t_2 :^{\text{prd}} \tau_2}{\Gamma \vdash \text{Tup}(t_1, t_2) :^{\text{prd}} \text{Pair}(\tau_1, \tau_2)} \; \text{Tup} \qquad \frac{\Gamma, x :^{\text{prd}} \tau_1, y :^{\text{prd}} \tau_2 \vdash s}{\Gamma \vdash \textbf{case } \{\text{Tup}(x, y) \Rightarrow s\} :^{\text{cns}} \text{Pair}(\tau_1, \tau_2)} \; \text{Case-Pair}$$

$$\frac{\Gamma \vdash k :^{\text{cns}} \tau}{\Gamma \vdash \text{hd}(k) :^{\text{cns}} \text{Stream}(\tau)} \; \text{Hd} \qquad \frac{\Gamma \vdash k :^{\text{cns}} \text{Stream}(\tau)}{\Gamma \vdash \text{tl}(k) :^{\text{cns}} \text{Stream}(\tau)} \; \text{Tl}$$

$$\frac{\Gamma, \alpha :^{\text{cns}} \tau \vdash s_1 \qquad \Gamma, \beta :^{\text{cns}} \text{Stream}(\tau) \vdash s_2}{\Gamma \vdash \textbf{cocase } \{\text{hd}(\alpha) \Rightarrow s_1, \text{tl}(\beta) \Rightarrow s_2\} :^{\text{prd}} \text{Stream}(\tau)} \; \text{Cocase-Stream}$$

$$\frac{\Gamma \vdash k :^{\text{cns}} \tau_1}{\Gamma \vdash \text{fst}(k) :^{\text{cns}} \text{LPair}(\tau_1, \tau_2)} \; \text{Fst} \qquad \frac{\Gamma \vdash k :^{\text{cns}} \tau_2}{\Gamma \vdash \text{snd}(k) :^{\text{cns}} \text{LPair}(\tau_1, \tau_2)} \; \text{Snd}$$

$$\frac{\Gamma, \alpha :^{\text{cns}} \tau_1 \vdash s_1 \qquad \Gamma, \beta :^{\text{cns}} \tau_2 \vdash s_2}{\Gamma \vdash \textbf{cocase } \{\text{fst}(\alpha) \Rightarrow s_1, \text{snd}(\beta) \Rightarrow s_2\} :^{\text{prd}} \text{LPair}(\tau_1, \tau_2)} \; \text{Cocase-LPair}$$

$$\frac{\Gamma \vdash p :^{\text{prd}} \sigma \qquad \Gamma \vdash c :^{\text{cns}} \tau}{\Gamma \vdash \text{ap}(p, c) :^{\text{cns}} \sigma \to \tau} \; \text{Ap} \qquad \frac{\Gamma, x :^{\text{prd}} \sigma, \alpha :^{\text{cns}} \tau \vdash s}{\Gamma \vdash \textbf{cocase } \{\text{ap}(x, \alpha) \Rightarrow s\} :^{\text{prd}} \sigma \to \tau} \; \text{Cocase-Fun}$$

$$\frac{}{\vdash \emptyset \; \text{Ok}} \; \text{Wf-Empty} \qquad \frac{\vdash P \; \text{Ok} \qquad \overline{x :^{\text{prd}} \tau_i}, \overline{\alpha :^{\text{cns}} \tau_j} \vdash_{P, \textbf{def } f(\overline{x_i :^{\text{prd}} \tau_i}; \overline{\alpha_j :^{\text{cns}} \tau_j}) := s} \; s}{\vdash P, \textbf{def } f(\overline{x_i :^{\text{prd}} \tau_i}, \overline{\alpha_j :^{\text{cns}} \tau_j}) := s \; \text{Ok}} \; \text{Wf-Cons}$$

Fig. 2. Typing rules of **Core**.

$$\frac{\Gamma, \alpha :^{\text{cns}} \tau \vdash s}{\Gamma \vdash \mu\alpha.s :^{\text{prd}} \tau} \; \mu \qquad\qquad\qquad \frac{\Gamma, x :^{\text{prd}} \tau \vdash s}{\Gamma \vdash \tilde{\mu}x.s :^{\text{cns}} \tau} \; \tilde{\mu}$$

4.2.1 Data and Codata Types. Figure 2 contains the typing rules for both Pair and List; since their rules are so similar we only discuss those of Pair explicitly:

$$\frac{\Gamma \vdash t_1 :^{\text{prd}} \tau_1 \qquad \Gamma \vdash t_2 :^{\text{prd}} \tau_2}{\Gamma \vdash \mathsf{Tup}(t_1, t_2) :^{\text{prd}} \mathsf{Pair}(\tau_1, \tau_2)} \; \text{Tup} \qquad \frac{\Gamma, x :^{\text{prd}} \tau_1, y :^{\text{prd}} \tau_2 \vdash s}{\Gamma \vdash \mathbf{case} \; \{\mathsf{Tup}(x, y) \Rightarrow s\} :^{\text{cns}} \mathsf{Pair}(\tau_1, \tau_2)} \; \text{Case-Pair}$$

In the rule Tup we type a pair constructor Tup applied to two arguments as a producer, and in the rule Case-Pair we type the case, which pattern-matches on this constructor and brings two variables into scope, as a consumer.

The typing rules for codata types look exactly the same, only the roles of producers and consumers are swapped.

$$\frac{\Gamma \vdash k :^{\text{cns}} \tau}{\Gamma \vdash \mathsf{hd}(k) :^{\text{cns}} \mathsf{Stream}(\tau)} \; \text{Hd} \qquad \frac{\Gamma, \alpha :^{\text{cns}} \tau \vdash s_1 \qquad \Gamma, \beta :^{\text{cns}} \mathsf{Stream}(\tau) \vdash s_2}{\Gamma \vdash \mathbf{cocase} \; \{\mathsf{hd}(\alpha) \Rightarrow s_1, \mathsf{tl}(\beta) \Rightarrow s_2\} :^{\text{prd}} \mathsf{Stream}(\tau)} \; \text{Cc-Str}$$

Most of the other rules directly correspond to a similar rule for **Fun**. When typing arithmetic expressions, for example, we only have to make sure all subterms have type **Int**.

We typecheck programs using the two rules Wf-Empty and Wf-Cons. The former is used to typecheck an empty program, and the rule Wf-Cons extends a typechecked program with a new top-level definition. When we typecheck the body of this top-level definition that we are about to add, we extend the program with this definition so that it can refer to itself recursively.

4.3 Type Soundness

In this section, we discuss the soundness of the type systems for both **Fun** and **Core** and show that the translation $[\![-]\!]$ preserves the typeability of terms. We follow Wright and Felleisen [1994] in presenting type soundness as the combination of a progress and a preservation theorem.

THEOREM 4.1 (PROGRESS, **FUN**). *Let t be a closed term in **Fun**, such that $\vdash t : \tau$ for some type τ. Then either t is a value or there is some term t' such that $t \triangleright t'$.*

This can easily be proved with an induction on typing derivations. Due to the presence of the **label**/**goto** construct, the standard formulation of the (strong) preservation theorem does not immediately hold for **Fun** (also see the discussion in Appendix C). The following weak form of preservation can again be easily proved by induction.

THEOREM 4.2 ((WEAK) PRESERVATION, **FUN**). *Let t, t' be terms in **Fun** such that $t \triangleright t'$, Γ an environment and P a program such that $\Gamma \vdash_P t : \tau$. Then there is a type τ' such that $\Gamma \vdash_P t' : \tau'$.*

The usual strong preservation theorem requires $\tau' = \tau$. But in fact, a slight variation of this strong form can be proved for **Fun** by adapting the technique found in Section 6 in [Wright and Felleisen 1994]. Thus, strong type soundness still does hold.

Before we can state the analogous theorems for **Core**, we will need an additional definition as a termination condition for evaluation.

Definition 4.3 (Terminal statement). If $\mathfrak{p}$ is a producer value in **Core** and $\bigstar$ a covariable which does not appear free in $\mathfrak{p}$, then $\langle \mathfrak{p} \mid \bigstar \rangle$ is called a *terminal statement*.

Terminal statements in **Core** have the same role as values in **Fun**. Some sequent-calculus-based languages use a special statement **Done** instead of terminal statements for this purpose.

THEOREM 4.4 (PROGRESS, **CORE**). *Let s be a focused statement in **Core** such that $\vdash s$. Then either s is a terminal statement, or there is some s' such that $s \triangleright s'$.*

For this theorem, we require s to be focused, in contrast to **Fun**, where progress holds for any (well-typed) term. This is because we have used static focusing for full evaluation in **Core**. If we used dynamic focusing instead, this requirement could be dropped, corresponding to using evaluation contexts (dynamic) in **Fun**, instead of a translation to normal form (static).

The Preservation theorem for **Core** is analogous to **Fun**.

THEOREM 4.5 (PRESERVATION, **CORE**). *Let s, s' be statements in* **Core** *with $s \triangleright s'$, Γ an environment and P a program. If $\Gamma \vdash_P s$, then $\Gamma \vdash_P s'$.*

This theorem can also be proven with a straightforward induction on typing derivations. Of course, this preservation theorem does not make any assertion about result types, as statements do not return anything that could be typed. However, if evaluation starts with a statement $\langle p \mid \bigstar \rangle$ where $\bigstar$ does not occur free in p and ends in a terminal statement s, then $s = \langle \mathfrak{p} \mid \bigstar \rangle$ for some producer value $\mathfrak{p}$. This is because no reduction step can introduce a free variable, so the final one must be the same as the initial one. Hence, by well-typedness, if p :$^{\mathrm{prd}} \tau$, then also $\mathfrak{p}$:$^{\mathrm{prd}} \tau$, because $\bigstar$:$^{\mathrm{cns}} \tau$.

Lastly, we come to an important property of the translation between these languages:

THEOREM 4.6 (TYPE PRESERVATION OF TRANSLATION). *Let t be a term in* **Fun***, Γ an environment and P a program. If $\Gamma \vdash_P t : \tau$ for some type τ, then $\Gamma \vdash_{[\![P]\!]} [\![t]\!]$:$^{\mathrm{prd}} \tau$ where $[\![P]\!]$ denotes the translation of all definitions in P.*

PROOF. Most cases are straightforward in the proof which proceeds by a structural induction on the typing derivation. The interesting cases are when the typing derivation types a control operator. The only rule in **Fun** with **label** $\alpha \{t_1\}$ in the conclusion is LABEL. This rule has the premise Γ, α :$^{\mathrm{cns}} \tau \vdash t_1 : \tau$, and applying the induction hypothesis gives Γ, α :$^{\mathrm{cns}} \tau \vdash [\![t_1]\!]$:$^{\mathrm{prd}} \tau$. Then we can derive $[\![t]\!] = \mu\alpha.\langle [\![t_1]\!] \mid \alpha \rangle$:$^{\mathrm{prd}} \tau$:

$$\frac{\dfrac{\text{(Induction Hypothesis)}}{\Gamma, \alpha \text{ :}^{\mathrm{cns}} \tau \vdash [\![t_1]\!] \text{ :}^{\mathrm{prd}} \tau} \quad \dfrac{}{\Gamma, \alpha \text{ :}^{\mathrm{cns}} \tau \vdash \alpha \text{ :}^{\mathrm{cns}} \tau} \text{VAR}_2}{\dfrac{\Gamma, \alpha \text{ :}^{\mathrm{cns}} \tau \vdash \langle [\![t_1]\!] \mid \alpha \rangle}{\Gamma \vdash \mu\alpha.\langle [\![t_1]\!] \mid \alpha \rangle \text{ :}^{\mathrm{prd}} \tau} \mu} \text{CUT}$$

The only rule with **goto**$(t_1; \alpha)$ in the conclusion is GOTO, which has premises $\Gamma \vdash t_1 : \tau$ and α :$^{\mathrm{cns}} \tau \in \Gamma$. Applying the induction hypothesis gives $\Gamma \vdash [\![t_1]\!]$:$^{\mathrm{prd}} \tau$, and we can therefore type the translation of the translation as follows (where we implicitly use weakening, which is allowed since β is fresh)

$$\frac{\dfrac{\text{(Induction Hypothesis)}}{\Gamma, \beta \text{ :}^{\mathrm{cns}} \sigma \vdash [\![t_1]\!] \text{ :}^{\mathrm{prd}} \tau} \quad \dfrac{\alpha \text{ :}^{\mathrm{cns}} \tau \in \Gamma, \beta \text{ :}^{\mathrm{cns}} \sigma}{\Gamma, \beta \text{ :}^{\mathrm{cns}} \sigma \vdash \alpha \text{ :}^{\mathrm{cns}} \tau} \text{VAR}_2}{\dfrac{\Gamma, \beta \text{ :}^{\mathrm{cns}} \sigma \vdash \langle [\![t_1]\!] \mid \alpha \rangle}{\Gamma \vdash \mu\beta.\langle [\![t_1]\!] \mid \alpha \rangle \text{ :}^{\mathrm{prd}} \tau} \mu} \text{CUT}$$

$\square$

5 Insights

In the previous section, we have explained *what* the $\lambda\mu\tilde{\mu}$-calculus is, and *how* it works. Now that we know the what and how we can explain *why* this calculus is so interesting. This section is therefore a small collection of independent insights. To be clear, these insights are obvious to those who are deeply familiar with the $\lambda\mu\tilde{\mu}$-calculus, but we can still recall how surprising they were for us when we first learned about them.

5.1 Evaluation Contexts are First Class

A central feature of the $\lambda\mu\tilde{\mu}$-calculus is the treatment of evaluation contexts as first-class objects, as we have mentioned before. For example, consider the term $(\ulcorner 2 \urcorner * \ulcorner 3 \urcorner) * \ulcorner 4 \urcorner$ in **Fun**. When we want to evaluate this, we have to use the evaluation context $\Box * \ulcorner 4 \urcorner$ to evaluate the subterm $(\ulcorner 2 \urcorner * \ulcorner 3 \urcorner)$ and get $\ulcorner 6 \urcorner * \ulcorner 4 \urcorner$ which we can then evaluate to $\ulcorner 24 \urcorner$. Translating this term into **Core** gives $\mu\alpha. * (\mu\beta. * (\ulcorner 2 \urcorner, \ulcorner 3 \urcorner; \beta), \ulcorner 4 \urcorner; \alpha)$. To evaluate this term, we first need to focus it giving

$$\mu\alpha.\langle \mu\beta. * (\ulcorner 2 \urcorner, \ulcorner 3 \urcorner; \beta) \mid \tilde{\mu}x. * (x, \ulcorner 4 \urcorner; \alpha) \rangle$$

When we now start evaluating with $\bigstar$, the steps are the same as in **Fun**. Using call-by-value, the μ-abstraction is evaluated first, giving $*(\ulcorner 2 \urcorner, \ulcorner 3 \urcorner; *(\tilde{\mu}x. * (x, \ulcorner 4 \urcorner; \bigstar)))$. This now has the form where the product can be evaluated to $\langle \ulcorner 6 \urcorner \mid \tilde{\mu}x. * (x, \ulcorner 4 \urcorner; \bigstar) \rangle$, after which $\ulcorner 6 \urcorner$ is substituted for x. The term $*(\ulcorner 6 \urcorner, \ulcorner 4 \urcorner; \bigstar)$ can then be directly evaluated to $\ulcorner 24 \urcorner$.

After focusing, we can see how β is a variable that stands for the evaluation context in **Fun**. The term $\tilde{\mu}x. * (x, \ulcorner 4 \urcorner; \alpha)$ is the first-class representation of the evaluation context $\Box * \ulcorner 4 \urcorner$. We first evaluate the subexpression $*(\ulcorner 2 \urcorner, \ulcorner 3 \urcorner; \beta)$ and then insert the result into $*(x, \ulcorner 4 \urcorner; \bigstar)$ to finish the evaluation, as we did in **Fun**. In other words, the $\Box$ of an evaluation context in **Fun**, corresponds to a continuation β in **Core**, and similarly determines in which order subexpressions are evaluated.

5.2 Data is Dual to Codata

The sequent calculus clarifies the relation between data and codata as being exactly dual to each other. When looking at the typing rules in Figure 2, we can see that data and codata types are completely symmetric. The two are not symmetric in languages based on natural deduction: A pattern match on data types includes the scrutinee but there is no corresponding object in the construction of codata. Similarly, invoking a destructor D of a codata type always includes the codata object x to be destructed, e.g., $x.D(\ldots)$, whereas the invocation of the constructor of a data type has no corresponding object.

This asymmetry is fixed in the sequent calculus. Destructors (such as fst) are first-class and don't require a scrutinee, which repairs the symmetry to constructors. Similarly, pattern matches (**case** $\{\ldots\}$) do not require an object to destruct, which makes them completely symmetrical to copattern matches. This duality reduces the conceptual complexity and opens the door towards shared design and implementation of features of data and codata types.

5.3 Let-Bindings are Dual to Control Operators

The **label** construct in **Fun** is translated to a μ-binding in **Core**. Also, when considering the typing rule for **label** α $\{t\}$ in Section 4.1, we can see that it directly corresponds to typing a μ-binding with the label α being the bound covariable. Similarly, a **let**-binding is translated to a $\tilde{\mu}$-binding and typing a **let**-binding in **Fun** closely corresponds to typing a $\tilde{\mu}$-term in **Core**. This way, **label**s and **let**-bindings are dual to each other, the same way μ and $\tilde{\mu}$ are. The duality can be extended to other control operators such as call/cc.

As it turns out, the **label** construct is very closely related to call/cc. There are in fact only two differences. First, **label** α $\{t\}$ has the binder α for the continuation built into the construct, just as the variation of call/cc named let/cc (which Reynolds [1972] called **escape**). The second, and more important difference is that the invocation of the continuation captured by **label** α $\{t\}$ happens through an explicit language construct **goto**$(t; \alpha)$. This makes it easy to give a translation to **Core** as we can simply insert another μ-binding to discard the remaining continuation at exactly the place where the captured continuation is invoked. In contrast, with call/cc and let/cc the continuation is applied in the same way as a normal function, making it necessary to redefine the

variable the captured continuation is bound to when translating to **Core**. This obscures the duality to **let**-bindings which is so evident for **label** and **goto**.

To see this, here is a translation of let/cc k t to **Core**

$$[\![\texttt{let/cc } k \ t]\!] := \mu\alpha.\langle \mathbf{cocase} \ \{\mathsf{ap}(x, \beta) \Rightarrow \langle x \mid \alpha\rangle\} \mid \tilde{\mu}k.\langle[\![t]\!] \mid \alpha\rangle\rangle$$

The essence of the translation still is that the current continuation is captured by the outer μ and bound to α. But now we also have to transform this α into a function (the **cocase** here) which discards its context (here bound to β) and bind this function to k, which is done using $\tilde{\mu}$. For call/cc, the duality is even more obscured, as there the binder for the continuation is hidden in the function which call/cc is applied to. For the translation, this function must then be applied to the above **cocase** and the captured continuation α, resulting in the following term (cf. also [Miquey 2019]).

$$[\![\texttt{call/cc } f]\!] := \mu\alpha.\langle[\![f]\!] \mid \mathsf{ap}(\mathbf{cocase} \ \{\mathsf{ap}(x, \beta) \Rightarrow \langle x \mid \alpha\rangle\}, \alpha)\rangle$$

Other control operators for undelimited continuations can be translated in a similar way. For example, consider Felleisen's C [Felleisen et al. 1987]. The difference to call/cc is that C discards the current continuation if it is not invoked somewhere in the term C is applied to, whereas call/cc leaves it in place and thus behaves as a no-op if the captured continuation is never invoked. The only change that needs to be made in the translation to **Core** is that the top-level continuation ★ has to be used for the outer cut instead of using the captured continuation. This is most easily seen for a variation of C which has the binder for the continuation built into the operator and where the invocation of the continuation is explicit, similar to **label**/**goto**. Calling this variation **label**$_C$, we obtain the following translation

$$[\![\mathbf{label}_C \ \alpha \ \{t\}]\!] := \mu\alpha.\langle[\![t]\!] \mid \bigstar\rangle$$

Here the duality to **let**-bindings is evident again. The translation for C itself is then obtained in the same way as for call/cc

$$[\![C \ f]\!] := \mu\alpha.\langle[\![f]\!] \mid \mathsf{ap}(\mathbf{cocase} \ \{\mathsf{ap}(x, \beta) \Rightarrow \langle x \mid \alpha\rangle\}, \bigstar)\rangle$$

5.4 The Case-of-Case Transformation

One important transformation in functional compilers is the case-of-case transformation. Maurer et al. [2017] give the following example of this transformation. The term

$$\mathbf{if} \ (\mathbf{if} \ e_1 \ \mathbf{then} \ e_2 \ \mathbf{else} \ e_3) \ \mathbf{then} \ e_4 \ \mathbf{else} \ e_5$$

can be replaced by the term

$$\mathbf{if} \ e_1 \ \mathbf{then} \ (\mathbf{if} \ e_2 \ \mathbf{then} \ e_4 \ \mathbf{else} \ e_5) \ \mathbf{else} \ (\mathbf{if} \ e_3 \ \mathbf{then} \ e_4 \ \mathbf{else} \ e_5).$$

Logicians call these kinds of transformations *commutative conversions*, and they play an important role in the study of the sequent calculus. But as Maurer et al. [2017] show, they are also important for compiler writers who want to generate efficient code.

In the $\lambda\mu\tilde{\mu}$-calculus, commuting conversions don't have to be implemented as a special compiler pass. They fall out *for free* as a special instance of μ-reductions! Let us illustrate this point by translating Maurer et al.'s example into the $\lambda\mu\tilde{\mu}$-calculus. First, let us translate the two examples using pattern-matching syntax:

$$\mathbf{case} \ (\mathbf{case} \ e_1 \ \mathbf{of} \ \{\mathsf{T} \Rightarrow e_2; \mathsf{F} \Rightarrow e_3\}) \ \mathbf{of} \ \{\mathsf{T} \Rightarrow e_4; \mathsf{F} \Rightarrow e_5\}$$

$$\mathbf{case} \ e_1 \ \mathbf{of} \ \{\mathsf{T} \Rightarrow \mathbf{case} \ e_2 \ \mathbf{of} \ \{\mathsf{T} \Rightarrow e_4; \mathsf{F} \Rightarrow e_5\}; \mathsf{F} \Rightarrow \mathbf{case} \ e_3 \ \mathbf{of} \ \{\mathsf{T} \Rightarrow e_4; \mathsf{F} \Rightarrow e_5\}\}$$

Let us now translate these two terms into the $\lambda\mu\tilde{\mu}$-calculus:

$\mu\alpha.\langle\mu\beta.\langle[\![e_1]\!] \mid \textbf{case } \{\mathsf{T} \Rightarrow \langle[\![e_2]\!] \mid \beta\rangle; \mathsf{F} \Rightarrow \langle e_3 \mid \beta\rangle\}\rangle \mid \textbf{case } \{\mathsf{T} \Rightarrow \langle[\![e_4]\!] \mid \alpha\rangle, \mathsf{F} \Rightarrow \langle[\![e_5]\!] \mid \alpha\rangle\}\rangle$

$\mu\alpha.\langle[\![e_1]\!]|\textbf{case } \{$

$\qquad \mathsf{T} \Rightarrow \langle\mu\beta.\langle[\![e_2]\!] \mid \textbf{case } \{\mathsf{T} \Rightarrow \langle[\![e_4]\!] \mid \beta\rangle, \mathsf{F} \Rightarrow \langle[\![e_5]\!] \mid \beta\rangle\}\rangle \mid \alpha\rangle$

$\qquad \mathsf{F} \Rightarrow \langle\mu\beta.\langle[\![e_3]\!] \mid \textbf{case } \{\mathsf{T} \Rightarrow \langle[\![e_4]\!] \mid \beta\rangle, \mathsf{F} \Rightarrow \langle[\![e_5]\!] \mid \beta\rangle\}\rangle \mid \alpha\rangle\}\rangle$

We can see that just by reducing all of the underlined redexes we reduce both of these examples to the same term.

5.5 Direct and Indirect Consumers

As mentioned in the introduction, a natural competitor of sequent calculus as an intermediate representation is continuation-passing style (CPS). In CPS, reified evaluation contexts are represented by functions. This makes the resulting types of programs in CPS arguably harder to understand. There is, however, another advantage of sequent calculus over CPS as described by Downen et al. [2016]. The first-class representation of consumers in sequent calculus allows us to distinguish between two different kinds of consumers: direct consumers, i.e., destructors, and indirect consumers. In particular, this allows to chain direct consumers in **Core** in a similar way as in **Fun**.

Suppose we have a codata type with destructors get and set for getting and setting the value of a reference. Now consider the following chain of destructor calls on a reference r in **Fun**

$$r.\mathsf{set}(3).\mathsf{set}(4).\mathsf{get}()$$

A compiler could use a user-defined custom rewrite rule for rewriting two subsequent calls to set into only the second call. In **Core** the above example looks as follows:

$$\mu\alpha.\langle r \mid \mathsf{set}(3; \mathsf{set}(4; \mathsf{get}(\alpha))\rangle$$

We still can immediately see the direct chaining of destructors and thus apply essentially the same rewrite rule. In CPS, however, the example would rather become

$$\lambda k.\, r.\mathsf{set}(3; \lambda s.\, s.\mathsf{set}(4; \lambda t.\, t.\mathsf{get}(k)))$$

The chaining of the destructors becomes obfuscated by the indirections introduced by representing the continuations for each destructor as a function. To apply the custom rewrite rule mentioned above, it is necessary to see through the lambdas, i.e. the custom rewrite rule has to be transformed to be applicable.

5.6 Call-By-Value, Call-By-Name and Eta-Laws

In Section 2.2 we already pointed out the existence of statements $\langle\mu\alpha.s_1 \mid \tilde{\mu}x.s_2\rangle$ which are called *critical pairs* because they can a priori be reduced to either $s_1[\tilde{\mu}x.s_2/\alpha]$ or $s_2[\mu\alpha.s_1/x]$. These critical pairs were already discussed by Curien and Herbelin [2000] when they introduced the $\lambda\mu\tilde{\mu}$-calculus. One solution is to pick an evaluation order, either call-by-value (cbv) or call-by-name (cbn), that determines to which of the two statements we should evaluate, and in this paper we chose to always use the call-by-value evaluation order. The difference between these two choices has also been discussed by Wadler [2003]. Note that this freedom for the evaluation strategy is another advantage of sequent calculus over continuation-passing style, as the latter always fixes an evaluation strategy.

Which evaluation order we choose has an important consequence for the optimizations we are allowed to perform in the compiler. If we choose call-by-value, then we are not allowed to use all η-equalities for codata types, and if we use call-by-name, then we are not allowed to use

all η-equalities for data types. Let us illustrate the problem in the case of codata types with the following example:

$$\langle \mathbf{cocase}\ \{\mathsf{ap}(x;\alpha) \Rightarrow \langle \mu\beta.s_1 \mid \mathsf{ap}(x;\alpha)\rangle\} \mid \tilde{\mu}x.s_2\rangle \equiv_\eta \langle \mu\beta.s_1 \mid \tilde{\mu}x.s_2\rangle$$

We assume that x and α do not appear free in s_1. The η-transformation is just the ordinary η-law for functions but applied to the representation of functions as codata types. The statement on the left-hand side reduces the $\tilde{\mu}$ first under both call-by-value and call-by-name evaluation order, i.e.

$$\langle \mathbf{cocase}\ \{\mathsf{ap}(x;\alpha) \Rightarrow \langle \mu\beta.s_1 \mid \mathsf{ap}(x;\alpha)\rangle\} \mid \tilde{\mu}x.s_2\rangle \quad \begin{array}{l} \vartriangleright_{\mathrm{cbv}} \quad s_2[\mathbf{cocase}\ \{\ldots\}/x] \\ \vartriangleright_{\mathrm{cbn}} \quad s_2[\mathbf{cocase}\ \{\ldots\}/x] \end{array}$$

The right-hand side of the η-equality, however, reduces the μ first under call-by-value evaluation order, i.e.

$$\langle \mu\beta.s_1 \mid \tilde{\mu}x.s_2\rangle \quad \begin{array}{l} \vartriangleright_{\mathrm{cbv}} \quad s_1[\tilde{\mu}x.s_2/\beta] \\ \vartriangleright_{\mathrm{cbn}} \quad s_2[\mu\beta.s_1/x] \end{array}$$

Therefore, the η-equality is only valid under call-by-name evaluation order. This example shows that the validity of applying this η-rule as an optimization depends on whether the language uses call-by-value or call-by-name. If we instead used a data type such as Pair, a similar η-reduction would only give the same result as the original statement when using call-by-value.

5.7 Linear Logic and the Duality of Exceptions

We have introduced the data type $\mathsf{Pair}(\sigma, \tau)$ and the codata type $\mathsf{LPair}(\sigma, \tau)$ as two different ways to formalize tuples. The data type $\mathsf{Pair}(\sigma, \tau)$ is defined by the constructor Tup whose arguments are evaluated eagerly, so this type corresponds to strict tuples in languages like ML or OCaml. The codata type $\mathsf{LPair}(\sigma, \tau)$ is a lazy pair which is defined by its two projections fst and snd, and only when we invoke the first or second projection do we start to compute its contents. This is closer to how tuples behave in a lazy language like Haskell.

Linear logic [Girard 1987; Wadler 1990] adds another difference to these types. In linear logic we consider arguments as resources which cannot be arbitrarily duplicated or discarded; every argument to a function has to be used exactly once. If we follow this stricter discipline, then we have to distinguish between two different types of pairs: In order to use a pair $\sigma \otimes \tau$ (pronounced "times" or "tensor"), we have to use both the σ and the τ, but if we want to use a pair $\sigma \mathbin{\&} \tau$ (pronounced "with"), we must choose to either use the σ or the τ. It is now easy to see that the type $\sigma \otimes \tau$ from linear logic corresponds to the data type $\mathsf{Pair}(\sigma, \tau)$, since when we pattern match on this type we get *two* variables in the context, one for σ and one for τ. The type $\sigma \mathbin{\&} \tau$ similarly corresponds to the type $\mathsf{LPair}(\sigma, \tau)$ which we use by invoking either the first or the second projection, consuming the whole pair.

In addition to these two different kinds of conjunction, we also have two different kinds of disjunction. These two disjunctions are written $\sigma \oplus \tau$ (pronounced "plus") and $\sigma \mathbin{\Im} \tau$ (pronounced "par") and correspond to two different ways to handle errors in programming languages. Their typing rules in **Core** are:

$$\frac{\Gamma \vdash t :^{\mathrm{prd}} \sigma}{\Gamma \vdash \mathrm{Inl}(t) :^{\mathrm{prd}} \sigma \oplus \tau} \qquad\qquad \frac{\Gamma \vdash t :^{\mathrm{prd}} \tau}{\Gamma \vdash \mathrm{Inr}(t) :^{\mathrm{prd}} \sigma \oplus \tau}$$

$$\frac{\Gamma, x :^{\mathrm{prd}} \sigma \vdash s_1 \qquad \Gamma, y :^{\mathrm{prd}} \tau \vdash s_2}{\Gamma \vdash \mathbf{case}\ \{\mathrm{Inl}(x) \Rightarrow s_1, \mathrm{Inr}(y) \Rightarrow s_2\} :^{\mathrm{cns}} \sigma \oplus \tau}$$

$$\frac{\Gamma \vdash c_1 :^{\mathrm{cns}} \sigma \qquad \Gamma \vdash c_2 :^{\mathrm{cns}} \tau}{\Gamma \vdash \mathrm{Par}(c1, c2) :^{\mathrm{cns}} \sigma \mathbin{\Im} \tau} \qquad \frac{\Gamma, \alpha :^{\mathrm{cns}} \sigma, \beta :^{\mathrm{cns}} \tau \vdash s}{\Gamma \vdash \mathbf{cocase}\ \{\mathrm{Par}(\alpha, \beta) \Rightarrow s\} :^{\mathrm{prd}} \sigma \mathbin{\Im} \tau}$$

Languages like Rust and Haskell use $\sigma \oplus \tau$ for error handling, which corresponds to the "Either" and "Result" types in those languages. This corresponds to the calling convention that the function returns a tagged result which indicates whether an error has occurred or not, and the caller of the function has to check this tag. The type $\sigma \mathbin{\rotatebox[origin=c]{180}{$\&$}} \tau$ behaves differently: A function which returns a value of type $\sigma \mathbin{\rotatebox[origin=c]{180}{$\&$}} \tau$ has to be called with two continuations, one for the possibility that the function returns successfully and one for the possibility that the function throws an error. And the function itself decides which continuation to call, so there is no overhead for checking the result of a function call. This is quite similar to how some functions in Javascript are called with an "onSuccess" continuation and an "onFailure" continuation and different to the exception model of, e.g., Java, where the exception handler is dynamically scoped instead of lexically passed as an argument. This duality between the two different ways of handling exceptions can be seen most clearly in the sequent calculus; more details on this duality can be found in section 3.4 of [Spiwack 2014] or in section 7.1 of [Ostermann et al. 2022].

6 Related Work

The central ideas of the calculi that we have presented in this pearl are not novel: the $\lambda\mu\tilde{\mu}$-calculus is by now over 20 years old. We chose a variant of this calculus that can be used as a starting point to explore all the variants that have been described in the literature. This related work section is therefore intended to provide suggestions for further reading and the chance to dive deeper into specific topics that we have only touched upon.

6.1 The Sequent Calculus

The basis of our language **Core** is a term assignment system for the sequent calculus, an alternative logical system to natural deduction. The sequent calculus was originally introduced by Gentzen in the articles Gentzen [1935a,b, 1969]. For a more thorough introduction to the sequent calculus as a logical system, we can recommend the books by Negri and Von Plato [2001] and Troelstra and Schwichtenberg [2000] which introduce the sequent calculus and show how it differs from the natural deduction systems that are more commonly taught.

6.2 Term Assignment for the Sequent Calculus

The original article which introduced the $\lambda\mu\tilde{\mu}$-calculus as a term assignment system for the sequent calculus was by Curien and Herbelin [2000]. Before we list some of the other articles, we should preface them with the following remark on notation:

Remark 1 (Alternative Notation). Our notation for producers, consumers and statements follows the established conventions in the literature. However, we diverge in the way that we write typing judgments from the example of Curien and Herbelin [2000] which is followed by most other authors. We use one typing context Γ which binds both variables $x :^{\text{prd}} \tau$ and covariables $\alpha :^{\text{cns}} \tau$, whereas Curien and Herbelin [2000] use two contexts; a context Γ which contains bindings for all variables and a context Δ which contains the bindings for all covariables. The following table summarizes the difference between their notation and the notation used in our paper.

Judgment Form	Our notation	Curien and Herbelin [2000]
Typing Producers	$\Gamma \vdash p :^{\text{prd}} \tau$	$\Gamma \vdash p : \tau \mid \Delta$
Typing Consumers	$\Gamma \vdash c :^{\text{cns}} \tau$	$\Gamma \mid c : \tau \vdash \Delta$
Typing Statements	$\Gamma \vdash s$	$s : (\Gamma \vdash \Delta)$

The reasons for this divergence are easily explained. The notation of Curien and Herbelin [2000] with its two contexts Γ and Δ perfectly illustrates the correspondence to the sequent calculus which operates with sequents $\Gamma \vdash \Delta$ which contain multiple formulas on the left- and right-hand side of the turnstile. This close correspondence to the sequent calculus is less important for us. We found that splitting the context in this way often makes it more difficult to write down rules in their full generality when we extend the language with other features. Features which introduce a dependency of later bindings on earlier bindings within a typing context, for example when we add parametric polymorphism, don't fit easily into the format of Curien and Herbelin [2000].

With these remarks out of the way, we can recommend the articles by Zeilberger [2008], Downen and Ariola [2014, 2018b, 2020], Munch-Maccagnoni [2009] and Spiwack [2014] which were very helpful to us when we learned about the $\lambda\mu\tilde{\mu}$-calculus.

6.3 Codata Types

Codata types were originally invented by Hagino [1989]. They had the most success in proof assistants such as Agda where they help circumvent certain technical problems that arise when we try to model coinductive types. Copattern matching as a way to create producers of codata types was popularized by Abel et al. [2013], although the basic idea of the concept had been around before that, see, e.g., [Zeilberger 2008]. But probably the best starting point to learn more about codata types is an article written by Downen et al. [2019].

6.4 Control Operators and Classical Logic

The **label**/**goto** construct that we are using in **Fun** is an example of a control operator, of which Landin's operator J [Felleisen 1987; Landin 1965; Thielecke 1998] likely is the oldest. Their translation into **Core** uses μ-abstractions, which are also a form of control operator that was originally introduced by Parigot [1992] before it became a part of the $\lambda\mu\tilde{\mu}$-calculus of Curien and Herbelin [2000]. Control operators have an important relationship to classical logic via the Curry-Howard isomorphism. This relationship was discovered by Griffin [1989]; a more thorough introduction can be found in Sørensen and Urzyczyn [2006].

6.5 Different Evaluation Orders

We have already talked about the evaluation strategies call-by-value and call-by-name, and how their difference can be explained by different choices of how a critical pair should be evaluated. This duality between call-by-value and call-by-name has already been observed by Filinski [1989] and has been explored in more detail by Wadler [2003, 2005]. We have also seen in Section 5.6 how η-reduction only works with data types in call-by-value and with codata types in call-by-name. A lot of people therefore conclude that the choice of an evaluation order should maybe not be a global decision, but should instead depend on the type. This approach requires tracking the polarity of types and providing additional shift connectives which help mediate between the different evaluation orders; the article by Downen and Ariola [2018a] is a good entry point for pursuing these kinds of questions which are discussed in detail in [Zeilberger 2009] and [Munch-Maccagnoni 2013]. A well-known example of mixing evaluation orders is the call-by-push-value paradigm [Levy 1999] which distinguishes value types and computation types and subsumes both call-by-value and call-by-name.

7 Conclusion

In this functional pearl, we have presented the $\lambda\mu\tilde{\mu}$-calculus in the way we introduce it to our colleagues and students on the whiteboard; by compiling small examples of functional programs.

We think this is a better way to introduce programming-language enthusiasts and compiler writers to the $\lambda\mu\tilde{\mu}$-calculus, since it doesn't require prior knowledge of the sequent calculus. We have also shown *why* we are excited about this calculus, by giving examples of how it allows us to express aspects like strict vs. lazy evaluation or compiler optimizations like case-of-case in an extremely clear way. We want to share our enthusiasm for the sequent calculus and languages built on it with more people, and with this pearl, we hope that others will start to write their own little compilers to the sequent calculus and explore the exciting possibilities it offers.

Data Availability Statement

Acknowledgments

A   The Relationship to the Sequent Calculus

In the main part of the paper we introduced the $\lambda\mu\tilde{\mu}$-calculus without any references to the sequent calculus, because we think it is not essential to understand the latter in order to understand the former. In this appendix, we provide the details which help make the connection between the logical calculus and the term system clear. We only discuss a very simple sequent calculus which contains two logical connectives: the two conjunctions $A \otimes B$ and $A \& B$ which correspond to the strict and lazy pairs that we have seen in **Core**. We use X for propositional variables.

$$A, B ::= X \mid A \otimes B \mid A \& B$$

In the (classical) sequent calculus both the premises and the conclusion of a derivation rule consist of *sequents* $\Gamma \vdash \Delta$. Both Γ and Δ are multisets of formulas; that is, it is important how often a formula occurs on the left or the right, but not in which order the formulas occur. In the sequent calculus, we only have *introduction rules*. This means that the logically complex formula $A \otimes B$ or $A \& B$ only occurs in the conclusion of the rules that define it, and not in one of the premises. Every connective comes with a set of rules which introduce the connective on the left and the right of the turnstile. In our case, the rules look like this:

$$\frac{}{A \vdash A} \text{ Axiom} \qquad \qquad \frac{\Gamma_1 \vdash \Delta_1, A \qquad A, \Gamma_2 \vdash \Delta_2}{\Gamma_1, \Gamma_2 \vdash \Delta_1, \Delta_2} \text{ Cut}$$

$$\frac{\Gamma, A_1, A_2 \vdash \Delta}{\Gamma, A_1 \otimes A_2 \vdash \Delta} \otimes\text{-L} \qquad \qquad \frac{\Gamma_1 \vdash A_1, \Delta_1 \qquad \Gamma_2 \vdash A_2, \Delta_2}{\Gamma_1, \Gamma_2 \vdash A_1 \otimes A_2, \Delta_1, \Delta_2} \otimes\text{-R}$$

$$\frac{\Gamma, A_1 \vdash \Delta}{\Gamma, A_1 \& A_2 \vdash \Delta} \&\text{-L}_1 \qquad \frac{\Gamma, A_2 \vdash \Delta}{\Gamma, A_1 \& A_2 \vdash \Delta} \&\text{-L}_2 \qquad \frac{\Gamma \vdash A_1, \Delta \qquad \Gamma \vdash A_2, \Delta}{\Gamma \vdash A_1 \& A_2, \Delta} \&\text{-R}_2$$

The rule Cut is the only rule which destroys the so-called *subformula property*. This property says that every formula which occurs anywhere in a derivation is a subformula of a formula occurring in the conclusion of the derivation. Proof theorists therefore try to show that we can *eliminate the cuts*; if every sequent which can be derived using the Cut rule can also be derived without using it, we say that the calculus enjoys the *cut-elimination property*. The Curry-Howard correspondence for the sequent calculus relates this cut-elimination procedure to the computations that we have seen in the paper.

The first step from the sequent calculus towards the $\lambda\mu\tilde{\mu}$-calculus consists in marking at most one of the formulas in each of the sequents as *active*. We mark a formula as active by enclosing it in a pair of brackets. This yields two versions of the rule Axiom, one where we mark the formula on the left and one where we mark the formula on the right. If we want to translate every derivation using the original rules to a derivation in the new variant we also have to add special rules which *activate* and *deactivate* formulas both on the left and on the right. This yields the following new set of rules:

$$\frac{}{[A] \vdash A} \text{ Axiom-L} \qquad \frac{}{A \vdash [A]} \text{ Axiom-R} \qquad \frac{\Gamma_1 \vdash \Delta_1, [A] \qquad [A], \Gamma_2 \vdash \Delta_2}{\Gamma_1, \Gamma_2 \vdash \Delta_1, \Delta_2} \text{ Cut}$$

$$\frac{\Gamma, A_1, A_2 \vdash \Delta}{\Gamma, [A_1 \otimes A_2] \vdash \Delta} \otimes\text{-L} \qquad \frac{\Gamma_1 \vdash [A_1], \Delta_1 \qquad \Gamma_2 \vdash [A_2], \Delta_2}{\Gamma_1, \Gamma_2 \vdash [A_1 \otimes A_2], \Delta_1, \Delta_2} \otimes\text{-R}$$

$$\frac{\Gamma, [A_1] \vdash \Delta}{\Gamma, [A_1 \,\&\, A_2] \vdash \Delta} \&\text{-L}_1 \qquad \frac{\Gamma, [A_2] \vdash \Delta}{\Gamma, [A_1 \,\&\, A_2] \vdash \Delta} \&\text{-L}_2 \qquad \frac{\Gamma \vdash A_1, \Delta \qquad \Gamma \vdash A_2, \Delta}{\Gamma \vdash [A_1 \,\&\, A_2], \Delta} \&\text{-R}_2$$

$$\frac{\Gamma, A \vdash \Delta}{\Gamma, [A] \vdash \Delta} \text{ Act-L} \qquad \frac{\Gamma, [A] \vdash \Delta}{\Gamma, A \vdash \Delta} \text{ Deact-L} \qquad \frac{\Gamma \vdash A, \Delta}{\Gamma \vdash [A], \Delta} \text{ Act-R} \qquad \frac{\Gamma \vdash [A], \Delta}{\Gamma \vdash A, \Delta} \text{ Deact-R}$$

We can now begin to assign terms to derivations in this calculus by associating every non-active formula A in the *left* side of the turnstile with a producer variable $x :^{\text{prd}} A$, and every non-active formula B on the *right* side of the turnstile with a consumer variable $\alpha :^{\text{cns}} B$. As discussed in Section 6.2 we write both producer and consumer variables in a joint context Γ on the left-hand side of typing rules. We have to distinguish three different sequents, depending on whether a formula is active, and if so, on which side the active formula occurs. If there is no active formula, then we assign a *statement* to the sequent, if the formula on the right is active, we assign a *producer*, and if the formula on the left is active, we assign a *consumer*. For most rules the correspondence is clear: The rule Axiom-R corresponds to the typing rule Var$_1$ (and Axiom-L to Var$_2$). The rule Tup corresponds to the rule $\otimes$-R, and Case-Pair to $\otimes$-L. The rules Fst and Snd correspond to the rules $\&$-L$_1$ and $\&$-L$_2$, and Cocase-LPair to $\&$-R. The activation rules correspond to the rules μ and $\tilde{\mu}$, and deactivation can be expressed as a cut with a variable.

B Typing Rules for Fun

Given a term t, an environment Γ and a program P, if t has type τ in environment Γ and program P, we write $\Gamma \vdash_P t : \tau$. As P is only used for typing calls to top-level definitions (rule Call), we usually leave it implicit in the typing rules. To make sure programs P are well-formed, we have additional checking rules for programs $\emptyset$-ok and P-Ok. If a program is well-formed, we write $\vdash P$ Ok.

$$\frac{x :^{\mathrm{prd}} \tau \in \Gamma}{\Gamma \vdash x : \tau} \; \text{Var} \qquad\qquad \frac{}{\Gamma \vdash \ulcorner n \urcorner : \mathbf{Int}} \; \text{Lit} \qquad\qquad \frac{\Gamma \vdash t_1 : \mathbf{Int} \quad \Gamma \vdash t_2 : \mathbf{Int}}{\Gamma \vdash t_1 \odot t_2 : \mathbf{Int}} \; \text{Op}$$

$$\frac{\Gamma \vdash n : \mathbf{Int} \quad \Gamma \vdash t_1 : \tau \quad \Gamma \vdash t_2 : \tau}{\Gamma \vdash \mathbf{ifz}(n, t_1, t_2) : \tau} \; \text{Ifz} \qquad\qquad \frac{\Gamma \vdash t_1 : \tau_1 \quad \Gamma, x :^{\mathrm{prd}} \tau_1 \vdash t_2 : \tau_2}{\Gamma \vdash \mathbf{let}\; x = t_1 \;\mathbf{in}\; t_2 : \tau_2} \; \text{Let}$$

$$\frac{\mathbf{def}\; f(\overline{x_i :^{\mathrm{prd}} \tau_i}; \overline{\alpha_j :^{\mathrm{cns}} \tau_j}) : \tau \in P \qquad \overline{\Gamma \vdash t_i : \tau_i} \qquad \overline{\Gamma \vdash \alpha_j :^{\mathrm{cns}} \tau_j}}{\Gamma \vdash_P f(\overline{t_i}; \overline{\alpha_j}) : \tau} \; \text{Call}$$

$$\frac{\Gamma \vdash t : \mathsf{List}(\tau') \quad \Gamma \vdash t_1 : \tau \quad \Gamma, y :^{\mathrm{prd}} \tau', z :^{\mathrm{prd}} \mathsf{List}(\tau') \vdash t_2 : \tau}{\Gamma \vdash \mathbf{case}\; t\; \mathbf{of}\; \{\mathsf{Nil} \Rightarrow t_1, \mathsf{Cons}(y, z) \Rightarrow t_2\} : \tau} \; \text{Case-List}$$

$$\frac{\Gamma \vdash t_1 : \tau \quad \Gamma \vdash t_2 : \mathsf{List}(\tau)}{\Gamma \vdash \mathsf{Cons}(t_1, t_2) : \mathsf{List}(\tau)} \; \text{Cons} \qquad\qquad \frac{}{\Gamma \vdash \mathsf{Nil} : \mathsf{List}(\tau)} \; \text{Nil}$$

$$\frac{\Gamma \vdash t : \mathsf{Pair}(\tau_1, \tau_2) \quad \Gamma, x :^{\mathrm{prd}} \tau_1, y :^{\mathrm{prd}} \tau_2 \vdash t : \tau}{\Gamma \vdash \mathbf{case}\; t\; \mathbf{of}\; \{\mathsf{Tup}(x, y) \Rightarrow t\} : \tau} \; \text{Case-Pair} \qquad \frac{\Gamma \vdash t_1 : \tau_1 \quad \Gamma \vdash t_2 : \tau_2}{\Gamma \vdash \mathsf{Tup}(t_1, t_2) : \mathsf{Pair}(\tau_1, \tau_2)} \; \text{Tup}$$

$$\frac{\Gamma \vdash t : \mathsf{Stream}(\tau)}{\Gamma \vdash t.\mathsf{hd} : \tau} \; \text{Hd} \qquad\qquad \frac{\Gamma \vdash t : \mathsf{Stream}(\tau)}{\Gamma \vdash t.\mathsf{tl} : \mathsf{Stream}(\tau)} \; \text{Tl}$$

$$\frac{\Gamma \vdash t_1 : \tau \quad \Gamma \vdash t_2 : \mathsf{Stream}(\tau)}{\Gamma \vdash \mathbf{cocase}\; \{\mathsf{hd} \Rightarrow t_1, \mathsf{tl} \Rightarrow t_2\} : \mathsf{Stream}(\tau)} \; \text{Stream}$$

$$\frac{\Gamma \vdash t : \mathsf{LPair}(\tau_1, \tau_2)}{\Gamma \vdash t.\mathsf{fst} : \tau_1} \; \text{Fst} \qquad\qquad \frac{\Gamma \vdash t : \mathsf{LPair}(\tau_1, \tau_2)}{\Gamma \vdash t.\mathsf{snd} : \tau_2} \; \text{Snd}$$

$$\frac{\Gamma \vdash t_1 : \tau_1 \quad \Gamma \vdash t_2 : \tau_2}{\Gamma \vdash \mathbf{cocase}\; \{\mathsf{fst} \Rightarrow t_1, \mathsf{snd} \Rightarrow t_2\} : \mathsf{LPair}(\tau_1, \tau_2)} \; \text{LPair}$$

$$\frac{\Gamma \vdash t_1 : \tau_1 \rightarrow \tau_2 \quad \Gamma \vdash t_2 : \tau_1}{\Gamma \vdash t_1\, t_2 : \tau_2} \; \text{App} \qquad\qquad \frac{\Gamma, x :^{\mathrm{prd}} \tau_1 \vdash t : \tau_2}{\Gamma \vdash \lambda x.t : \tau_1 \rightarrow \tau_2} \; \text{Lam}$$

$$\frac{\Gamma \vdash t : \tau \quad \alpha :^{\mathrm{cns}} \tau \in \Gamma}{\Gamma \vdash \mathbf{goto}(t; \alpha) : \tau'} \; \text{Goto} \qquad\qquad \frac{\Gamma, \alpha :^{\mathrm{cns}} \tau \vdash t : \tau}{\Gamma \vdash \mathbf{label}\; \alpha\; \{t\} : \tau} \; \text{Label}$$

To check a program, we start with the empty program, which we know is well-formed (Wf-Empty), and then add one definition at a time. A definition is then well-typed if there are types $\overline{\tau_i}$ and $\overline{\tau_j}$ for its arguments such that its body is well-typed. Because we explicitly allow recursive definitions, the body t might contain the name f as well. Thus, while we typecheck t, we add the definition of f to the program and assume it is well-typed. After finding $\overline{\tau_i}$, $\overline{\tau_j}$ and τ, these are added to the program as well, that is, well-formed programs contain type annotations while definitions

on their own do not. This way, these types can be used while checking types of calls (in rule Call).

$$\frac{}{\vdash \emptyset \ \text{Ok}} \ \text{Wf-Empty} \qquad \frac{\vdash P \ \text{Ok} \qquad \overline{x :^{\text{prd}} \tau_i}, \overline{\alpha :^{\text{cns}} \tau_j} \vdash_{P, \text{def } f(\overline{x_i : \tau_i}, \overline{\alpha_j :^{\text{cns}} \tau_j}) : \tau := t} \ t : \tau}{\vdash P, \text{def } f(\overline{x_i : \tau_i}, \overline{\alpha_j :^{\text{cns}} \tau_j}) : \tau := t \ \text{Ok}} \ \text{Wf-Cons}$$

C Operational Semantics of label/goto

The full operational semantics for the **label**/**goto** construct is in essence the same as for `let`/`cc`. To make it precise, we promote evaluation contexts to runtime values.

We first repeat Definition 3.1 of evaluation contexts with one change: **label** α $\{E\}$ is not an evaluation context. We reduce a **label** as soon as it comes into evaluation position.

$$\begin{aligned} E \quad ::= \quad & \Box \mid E \odot t \mid \mathsf{t} \odot E \mid \mathbf{ifz}(E, t, t) \mid \mathbf{let} \ x = E \ \mathbf{in} \ t \mid f(\overline{\mathsf{t}}, E, \overline{t}) \mid K(\overline{\mathsf{t}}, E, \overline{t}) \\ & | \quad \mathbf{case} \ E \ \mathbf{of} \ \{\overline{K(\overline{x}) \Rightarrow t}\} \mid E \ t \mid \mathsf{t} \ E \mid E.D(\overline{t}) \mid \mathsf{t}.D(\overline{\mathsf{t}}, E, \overline{t}) \mid \mathbf{goto}(E; \alpha) \end{aligned}$$

Now we add them as another form of value

$$\mathsf{t} ::= \dots \mid E$$

Note that these values only exist at runtime, that is, they cannot appear in expressions before evaluation has started. They are typed as consumers, which means that they are the only values with a consumer type. This makes sure that we can substitute them for covariables. Their typing can be captured by the following rule.

$$\frac{x :^{\text{prd}} \tau \vdash E[x] : \tau_0}{\vdash E :^{\text{cns}} \tau} \ \text{Ctx}$$

The rule means that if the hole of a context E expects an expression of type τ to be plugged in, then we have $E :^{\text{cns}} \tau$.

Now we can give the evaluation rules for **label** and **goto**:

$$E[\mathbf{label} \ \alpha \ \{t\}] \rhd E[t[E/\alpha]] \qquad E'[\mathbf{goto}(\mathsf{t}; E)] \rhd E[\mathsf{t}]$$

In the rule for **label**, the surrounding evaluation context E is reified as a value and then substituted for the covariable α in the body t. Note that E is not removed, i.e., evaluation continues in this context. In particular, if α does not occur free in t, then the **label** is effectively a no-op. We can also see that the types are correct: If t has type τ, then so does **label** α $\{t\}$ and consequently we have $E :^{\text{cns}} \tau$ which is the same type as that of α. In the rule for **goto** the covariable must have already been replaced by an evaluation context, which is ensured if the evaluated term was closed and well-typed, because the only way to introduce a covariable is through a **label**. The evaluation step then removes and discards the surrounding context E' and continues evaluation by plugging the value t into the previously reified context E. Note that E' cannot contain **label**s, as they are not evaluation contexts. This ensures that there is no risk of removing a binder for a free variable in t. Together these two rules also allow us to simulate the approximate rule from Section 3.1:

$$E[\mathbf{label} \ \alpha \ \{E'[\mathbf{goto}(\mathsf{t}; \alpha)]\}] \rhd E[E'[\mathbf{goto}(\mathsf{t}; E)]] \rhd E[\mathsf{t}]$$

The rule for **goto** also is the reason why the theorem of strong preservation does not immediately hold (see the discussion in Section 4.3). The problem is that from this rule and the given typing rules for **Fun** it is not immediate that the evaluation contexts E' and E yield a term of the same type when filling their holes, so that the overall type of the term may not be preserved. But this cannot actually happen, because all other reduction rules preserve the overall type and hence all evaluation contexts that are reified by the rule for **label** must yield a term of that same overall type when their holes are filled. Therefore, also the rule for **goto** is type-preserving. This can be made

precise by explicitly tracking the overall type in the type system (see, e.g., Section 6 in [Wright and Felleisen 1994]).

References

Andreas Abel, Brigitte Pientka, David Thibodeau, and Anton Setzer. 2013. Copatterns: Programming Infinite Structures by Observations. In *Proceedings of the 40th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (Rome, Italy) *(POPL '13)*. Association for Computing Machinery, New York, NY, USA, 27–38. https://doi.org/10.1145/2480359.2429075

Jean-Marc Andreoli. 1992. Logic Programming with Focusing Proofs in Linear Logic. *Journal of Logic and Computation* 2 (1992), 297–347. Issue 3. https://doi.org/10.1093/logcom/2.3.297

Jonathan Immanuel Brachthäuser, Philipp Schuster, and Klaus Ostermann. 2020. Effects as capabilities: effect handlers and lightweight effect polymorphism. *Proc. ACM Program. Lang.* 4, OOPSLA, Article 126 (nov 2020), 30 pages. https://doi.org/10.1145/3428194

William R. Cook. 2009. On Understanding Data Abstraction, Revisited. In *Proceedings of the Conference on Object-Oriented Programming, Systems, Languages and Applications: Onward! Essays* (Orlando). Association for Computing Machinery, New York, NY, USA, 557–572. https://doi.org/10.1145/1640089.1640133

Pierre-Louis Curien and Hugo Herbelin. 2000. The Duality of Computation. In *Proceedings of the Fifth ACM SIGPLAN International Conference on Functional Programming (ICFP '00)*. Association for Computing Machinery, New York, NY, USA, 233–243. https://doi.org/10.1145/357766.351262

Pierre-Louis Curien and Guillaume Munch-Maccagnoni. 2010. The Duality of Computation under Focus. In *Theoretical Computer Science*, Cristian S. Calude and Vladimiro Sassone (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 165–181.

Paul Downen and Zena M. Ariola. 2014. The Duality of Construction. In *Proceedings of the 23rd European Symposium on Programming Languages and Systems - Volume 8410 (ESOP '14)*. Springer, Berlin, Heidelberg, 249–269. https://doi.org/10.1007/978-3-642-54833-8_14

Paul Downen and Zena M. Ariola. 2018a. Beyond Polarity: Towards a Multi-Discipline Intermediate Language with Sharing. In *27th EACSL Annual Conference on Computer Science Logic (CSL 2018) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 119)*, Dan Ghica and Achim Jung (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 21:1–21:23. https://doi.org/10.4230/LIPIcs.CSL.2018.21

Paul Downen and Zena M. Ariola. 2018b. A tutorial on computational classical logic and the sequent calculus. *Journal of Functional Programming* 28 (2018). https://doi.org/10.1017/S0956796818000023

Paul Downen and Zena M. Ariola. 2020. Compiling With Classical Connectives. *Logical Methods in Computer Science* Volume 16, Issue 3 (Aug. 2020). https://doi.org/10.23638/LMCS-16(3:13)2020

Paul Downen, Philip Johnson-Freyd, and Zena M. Ariola. 2015. Structures for structural recursion. In *Proceedings of the 20th ACM SIGPLAN International Conference on Functional Programming* (Vancouver, BC, Canada) *(ICFP 2015)*. Association for Computing Machinery, New York, NY, USA, 127–139. https://doi.org/10.1145/2784731.2784762

Paul Downen, Luke Maurer, Zena M Ariola, and Simon Peyton Jones. 2016. Sequent calculus as a compiler intermediate language. In *Proceedings of the 21st ACM SIGPLAN International Conference on Functional Programming*. 74–88.

Paul Downen, Zachary Sullivan, Zena M. Ariola, and Simon Peyton Jones. 2019. Codata in Action. In *European Symposium on Programming (ESOP '19)*. Springer, 119–146. https://doi.org/10.1007/978-3-030-17184-1_5

Matthias Felleisen. 1987. Reflections on Landin's J-operator: A Partly Historical Note. *Computer Languages* 12, 3 (1987), 197–207. https://doi.org/10.1016/0096-0551(87)90022-1

Matthias Felleisen, Daniel P. Friedman, Eugene Kohlbecker, and Bruce Duba. 1987. A syntactic theory of sequential control. *Theoretical Computer Science* 52, 3 (1987), 205–237. https://doi.org/10.1016/0304-3975(87)90109-5

Andrzej Filinski. 1989. Declarative Continuations: an Investigation of Duality in Programming Language Semantics. In *Category Theory and Computer Science*. Springer-Verlag, Berlin, Heidelberg, 224–249.

Gerhard Gentzen. 1935a. Untersuchungen über das logische Schließen. I. *Mathematische Zeitschrift* 35 (1935), 176–210.

Gerhard Gentzen. 1935b. Untersuchungen über das logische Schließen. II. *Mathematische Zeitschrift* 39 (1935), 405–431.

Gerhard Gentzen. 1969. *The collected papers of Gerhard Gentzen*. North-Holland Publishing Co., Amsterdam.

Jean-Yves Girard. 1987. Linear Logic. *Theoretical Computer Science* 50, 1 (1987), 1–101. https://doi.org/10.1016/0304-3975(87)90045-4

Brian Goetz et al. 2014. *JSR 335: Lambda Expressions for the Java Programming Language*. https://jcp.org/en/jsr/detail?id=335

Timothy G. Griffin. 1989. A Formulae-as-Type Notion of Control. In *Proceedings of the 17th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (San Francisco, California, USA) *(POPL '90)*. Association for Computing Machinery, New York, NY, USA, 47–58. https://doi.org/10.1145/96709.96714

Tatsuya Hagino. 1989. Codatatypes in ML. *Journal of Symbolic Computation* 8, 6 (1989), 629–650. https://doi.org/10.1016/S0747-7171(89)80065-3

Peter John Landin. 1965. Correspondence between ALGOL 60 and Church's Lambda-notation: part I. *Commun. ACM* 8, 2 (feb 1965), 89–101. https://doi.org/10.1145/363744.363749

Paul Blain Levy. 1999. Call-by-Push-Value: A Subsuming Paradigm. In *Proceedings of the 4th International Conference on Typed Lambda Calculi and Applications (TLCA '99)*. Springer-Verlag, Berlin, Heidelberg, 228–242.

Luke Maurer, Paul Downen, Zena M. Ariola, and Simon Peyton Jones. 2017. Compiling without Continuations. In *Proceedings of the 38th ACM SIGPLAN Conference on Programming Language Design and Implementation* (Barcelona, Spain) *(PLDI 2017)*. Association for Computing Machinery, New York, NY, USA, 482–494. https://doi.org/10.1145/3062341.3062380

Étienne Miquey. 2019. A Classical Sequent Calculus with Dependent Types. *ACM Trans. Program. Lang. Syst.* 41, 2, Article 8 (mar 2019), 47 pages. https://doi.org/10.1145/3230625

Guillaume Munch-Maccagnoni. 2009. Focalisation and Classical Realisability. In *Computer Science Logic: 23rd international Workshop, CSL 2009, 18th Annual Conference of the EACSL* (Coimbra, Portugal) *(CSL '09)*, Erich Grädel and Reinhard Kahle (Eds.). Springer, Berlin, Heidelberg, 409–423. https://doi.org/10.1007/978-3-642-04027-6_30

Guillaume Munch-Maccagnoni. 2013. *Syntax and Models of a non-Associative Composition of Programs and Proofs*. Ph. D. Dissertation. Univ. Paris Diderot.

Sara Negri and Jan Von Plato. 2001. *Structural Proof Theory*. Cambridge University Press. https://doi.org/10.1017/CBO9780511527340

Klaus Ostermann, David Binder, Ingo Skupin, Tim Süberkrüb, and Paul Downen. 2022. Introduction and Elimination, Left and Right. *Proc. ACM Program. Lang.* 6, ICFP, Article 106 (2022), 28 pages. https://doi.org/10.1145/3547637

Michel Parigot. 1992. $\lambda\mu$-Calculus: An algorithmic interpretation of classical natural deduction. In *Logic Programming and Automated Reasoning*, Andrei Voronkov (Ed.). Springer, Berlin, Heidelberg, 190–201.

John Charles Reynolds. 1972. Definitional Interpreters for Higher-Order Programming Languages. In *ACMConf* (Boston). Association for Computing Machinery, New York, NY, USA, 717–740. https://doi.org/10.1145/800194.805852

Arnaud Spiwack. 2014. A Dissection of L. (2014). Unpublished draft.

Morten Heine Sørensen and Paweł Urzyczyn. 2006. *Lectures on the Curry-Howard Isomorphism*. Studies in Logic and the Foundations of Mathematics, Vol. 149. Elsevier.

Hayo Thielecke. 1998. An Introduction to Landin's "A Generalization of Jumps and Labels". *Higher Order Symbol. Comput.* 11, 2 (sep 1998), 117–123. https://doi.org/10.1023/A:1010060315625

Anne Sjerp Troelstra and Helmut Schwichtenberg. 2000. *Basic Proof Theory, Second Edition.* Cambridge University Press.

Philip Wadler. 1990. Linear Types Can Change the World!. In *Programming Concepts and Methods*. North-Holland.

Philip Wadler. 2003. Call-by-value is dual to call-by-name. In *Proceedings of the Eighth ACM SIGPLAN International Conference on Functional Programming* (Uppsala, Sweden) *(ICFP '03)*. Association for Computing Machinery, New York, NY, USA, 189–201. https://doi.org/10.1145/944705.944723

Philip Wadler. 2005. Call-by-Value Is Dual to Call-by-Name - Reloaded. In *Term Rewriting and Applications, 16th International Conference, RTA 2005, Nara, Japan, April 19-21, 2005, Proceedings (Lecture Notes in Computer Science, Vol. 3467)*, Jürgen Giesl (Ed.). Springer, 185–203. https://doi.org/10.1007/978-3-540-32033-3_15

Andrew K. Wright and Matthias Felleisen. 1994. A Syntactic Approach to Type Soundness. *Information and Computation* 115, 1 (11 1994), 38–94. https://doi.org/10.1006/inco.1994.1093

Noam Zeilberger. 2008. On the Unity of Duality. *Annals of Pure and Applied Logic* 153, 1-3 (2008), 66–96. https://doi.org/10.1016/j.apal.2008.01.001

Noam Zeilberger. 2009. *The Logical Basis of Evaluation Order and Pattern-Matching*. Ph. D. Dissertation. Carnegie Mellon University, USA. Advisor(s) Pfenning, Frank and Lee, Peter.

Yizhou Zhang, Guido Salvaneschi, Quinn Beightol, Barbara Liskov, and Andrew C. Myers. 2016. Accepting Blame for Safe Tunneled Exceptions. In *Proceedings of the 37th ACM SIGPLAN Conference on Programming Language Design and Implementation* (Santa Barbara, CA, USA) *(PLDI '16)*. Association for Computing Machinery, New York, NY, USA, 281–295. https://doi.org/10.1145/2908080.2908086