



Calculating Compilers for Concurrency

PATRICK BAHR, IT University of Copenhagen, Denmark

GRAHAM HUTTON, University of Nottingham, UK

Choice trees have recently been introduced as a general structure for defining the semantics of programming languages with a wide variety of features and effects. In this article we focus on concurrent languages, and show how a codensity version of choice trees allows the semantics for such languages to be systematically transformed into compilers using equational reasoning techniques. The codensity construction is the key ingredient that enables a high-level, algebraic approach. As a case study, we calculate a compiler for a concurrent lambda calculus with channel-based communication.

CCS Concepts: • **Software and its engineering** → **Compilers; Formal software verification**; • **Theory of computation** → **Logic and verification; Program verification**.

Additional Key Words and Phrases: program calculation, concurrency, choice trees, codensity monad

ACM Reference Format:

Patrick Bahr and Graham Hutton. 2023. Calculating Compilers for Concurrency. *Proc. ACM Program. Lang.* 7, ICFP, Article 213 (August 2023), 28 pages. <https://doi.org/10.1145/3607855>

1 INTRODUCTION

Compilers are hard to write, and hard to get right. This is particularly so in the case of concurrent languages, where the addition of language primitives that introduce non-determinism make it significantly more challenging to develop and verify compilers.

One approach to compiler *verification* for concurrent languages is to define the semantics for both the source and target languages by translation into a lower-level concurrent language with suitable reasoning principles, such as bisimilarity and coinduction. This approach was pioneered by Wand [1995], who introduced the idea of translating into a process calculus, and has recently taken a step forward with the development of *choice trees* [Chappe et al. 2023], which provide a monadic language for expressing concurrency that supports modular, algebraic reasoning principles.

Such reasoning principles also make choice trees a suitable foundation for compiler *calculation*, a program synthesis technique that aims to derive correct-by-construction compilers from specifications of their correctness [Bahr and Hutton 2015, 2022]. The nature of the semantic reasoning principles is important for compiler calculation, because the aim is not only to produce correct compilers, but it also to *discover* compilation techniques. Simple, equational-style reasoning and a powerful (co-)induction principle are key features to enable this discovery.

In this article, we show how choice trees can be used as the semantic basis for compiler calculation for concurrent languages. In particular, the article makes the following contributions:

- We adapt the syntax and semantics of choice trees to enable the simple (co-)induction principle that powers the compiler calculation technique (Section 2).

Authors' addresses: Patrick Bahr, IT University of Copenhagen, Denmark; Graham Hutton, University of Nottingham, UK.



This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2023 Copyright held by the owner/author(s).

2475-1421/2023/8-ART213

<https://doi.org/10.1145/3607855>

- We identify a limitation of choice trees for defining the semantics of a simple concurrent language (Section 3), and show how the codensity construction can be applied to the choice tree monad to remove this limitation (Section 4).
- We present reasoning principles for the resulting *codensity choice trees* (Section 5), and show how they can be used to calculate a compiler for the simple concurrent language (Section 6).
- Finally, to demonstrate that our methodology scales to richly-featured concurrent languages, we show how to calculate a compiler for an untyped lambda calculus extended with concurrency and channel-based communication (Sections 7 and 8).

We use Haskell notation as our meta-language for accessibility, but assume that the language is total. Whereas in many articles calculations are often omitted or compressed for brevity, here they are the central focus so are typically presented in detail. All the calculations have been mechanically checked in Agda, and the proof scripts are available as online supplementary material [Bahr and Hutton 2023].

2 CHOICE TREES

In this section we introduce the basic concept of choice trees, show how they can be given a small-step operational semantics, and define the derived notion of parallel composition. To support the equational approach to reasoning that is used in compiler calculation, the syntax and semantics that we adopt for choice trees is different from the original article [Chappe et al. 2023]. As the full definitions for choice trees are developed in stages, we defer a discussion of these differences and their importance for compiler calculation until later on (Section 9).

2.1 Syntax

The type of choice trees $CTree\ e\ a$ represents non-deterministic computations that return values of type a and that may use algebraic effects defined by the type function $e :: * \rightarrow *$:

data $CTree\ e\ a$ **where**

```
Now :: a → CTree e a
(⊕) :: CTree e a → CTree e a → CTree e a
Zero :: CTree e a
Eff  :: e b → (b → CTree e a) → CTree e a
```

Informally, $Now\ v$ returns the value v without performing any effects, $p \oplus q$ makes a non-deterministic choice between two computations p and q , while $Zero$ is a computation that has terminated, and $Eff\ o\ c$ is a form of sequencing that feeds the result value produced by an effectful computation o into a continuation c . For example, if we wished to provide an effectful operation that prints an integer, this can be achieved by first defining a type function $PrintEff$ that provides a single constructor called $PrintInt$, which is then used to define a $print$ function:

data $PrintEff\ a$ **where**

```
PrintInt :: Int → PrintEff ()
print :: Int → CTree PrintEff ()
print n = Eff (PrintInt n) Now
```

Note that $print$ does not actually print an integer, but rather builds a choice tree that represents the action of printing. Choice trees form a monad, with $return$ and \gg defined as follows, which allows Haskell's **do** notation to be used to streamline the construction of choice trees:

$$\begin{array}{c}
\frac{}{\text{Now } v \xRightarrow{v} \text{Zero}} \qquad \frac{p \xRightarrow{l} p'}{p \oplus q \xRightarrow{l} p'} \qquad \frac{q \xRightarrow{l} q'}{p \oplus q \xRightarrow{l} q'} \\
\\
\frac{o :: e \ b \quad c :: b \rightarrow \text{CTree } e \ a}{\text{Eff } o \ c \xRightarrow{\uparrow o} c} \qquad \frac{c :: b \rightarrow \text{CTree } e \ a \quad i :: b}{c \xRightarrow{\downarrow i} c \ i}
\end{array}$$

Fig. 1. Transition semantics for (codensity) choice trees.

```

return :: a → CTree e a
return = Now
(⊗) :: CTree e a → (a → CTree e b) → CTree e b
Now v ⊗ f = f v
(p ⊕ q) ⊗ f = (p ⊗ f) ⊕ (q ⊗ f)
Zero ⊗ f = Zero
Eff o c ⊗ f = Eff o (λi → c i ⊗ f)

```

We will also make use of the functorial map function, which is derived from \otimes and *return*:

```

fmap :: (a → b) → CTree e a → CTree e b
fmap f p = p ⊗ (λx → return (f x))

```

Later on we will extend the notion of choice trees to support infinite (non-terminating) computations, but the above definitions will suffice for now.

2.2 Semantics

We define the semantics for choice trees by means of a labelled transition system. In our setting, a state for the transition system is either a choice tree $p :: \text{CTree } e \ a$, or a continuation $c :: b \rightarrow \text{CTree } e \ a$ that is waiting for an external input of type b in response to an effect of type $e \ b$. Transitions between states are labelled by one of four possible forms:

τ a silent transition	$\uparrow o$ an effect $o :: e \ b$ for some type b
v a return value $v :: a$	$\downarrow i$ an input $i :: b$ for some type b

Using these ideas, we define a labelled transition relation by the inference rules shown in Figure 1, which makes precise the informal meaning of choice trees from the previous section. Note that there is no rule for *Zero* because it represents a terminated computation that can make no further transitions. The silent transition τ plays no role yet, but will be used later on.

By way of example, the expression $\text{return } 1 \oplus (\text{print } 2 \otimes \lambda() \rightarrow \text{return } 3)$ expands to the choice tree $\text{Now } 1 \oplus \text{Eff } (\text{PrintInt } 2) (\lambda() \rightarrow \text{Now } 3)$, which has two possible transition sequences because of the use of the choice operator. In particular, it can either simply return the value 1,

$$\text{Now } 1 \oplus \text{Eff } (\text{PrintInt } 2) (\lambda() \rightarrow \text{Now } 3) \xRightarrow{1} \text{Zero}$$

or it can print 2, consume the resulting unit value $()$, and return the value 3:

$$\text{Now } 1 \oplus \text{Eff } (\text{PrintInt } 2) (\lambda() \rightarrow \text{Now } 3) \xRightarrow{\uparrow \text{PrintInt } 2} \lambda() \rightarrow \text{Now } 3 \xRightarrow{\downarrow ()} \text{Now } 3 \xRightarrow{3} \text{Zero}$$

As illustrated in this latter transition sequence, every effectful transition is always immediately followed by an input transition that consumes the resulting value:

$$Eff \circ c \xrightarrow{\uparrow o} c \xrightarrow{\downarrow i} c i$$

Hence, we could simplify the semantics by combining the two transitions into one:

$$Eff \circ c \xrightarrow{\uparrow o \downarrow i} c i$$

While this approach has the benefit of avoiding the need for two kinds of states in the semantics, it results in a notion of bisimilarity that is too coarse. We return to this issue in Section 9 in our comparison to the work of [Chappe et al. \[2023\]](#), who do use this simplified semantics. Nonetheless, as such transitions always occur in pairs, it is useful to define a relation that combines them:

$$p \xRightarrow{\uparrow o \downarrow i} q \quad \text{iff} \quad \exists c. p \xrightarrow{\uparrow o} c \wedge c \xrightarrow{\downarrow i} q$$

2.3 Parallelism

Choice trees do not have a built-in notion of parallelism, as this can be derived from the other primitives. In particular, we can define parallel composition using three auxiliary operators:

$(\parallel) :: CTree \ e \ a \rightarrow CTree \ e \ b \rightarrow CTree \ e \ (a, b)$

$$p \parallel q = (p \triangleleft q) \oplus (p \triangleright q) \oplus (p \bowtie q)$$

The first operator \triangleleft allows its left argument to perform an effectful computation:

$(\triangleleft) :: CTree \ e \ a \rightarrow CTree \ e \ b \rightarrow CTree \ e \ (a, b)$

$$Now \ v \ \triangleleft \ q = Zero$$

$$(p_1 \oplus p_2) \triangleleft q = (p_1 \triangleleft q) \oplus (p_2 \triangleleft q)$$

$$Zero \ \triangleleft \ q = Zero$$

$$Eff \ o \ c \ \triangleleft \ q = Eff \ o \ (\lambda i \rightarrow c \ i \parallel q)$$

The second operator \triangleright does the same for the right argument, and is defined symmetrically to \triangleleft . The final operator \bowtie allows both argument choice trees to perform computations simultaneously, with the resulting return values from each side being combined as a pair:

$(\bowtie) :: CTree \ e \ a \rightarrow CTree \ e \ b \rightarrow CTree \ e \ (a, b)$

$$(p_1 \oplus p_2) \bowtie q = (p_1 \bowtie q) \oplus (p_2 \bowtie q)$$

$$p \ \bowtie \ (q_1 \oplus q_2) = (p \bowtie q_1) \oplus (p \bowtie q_2)$$

$$Now \ v \ \bowtie \ Now \ w = Now \ (v, w)$$

$$_ \ \bowtie \ _ = Zero$$

The behaviour of \parallel can be concisely characterised by the following inference rules, which together express that parallel composition has the expected behaviour:

$$\frac{p \xrightarrow{\uparrow o \downarrow i} p'}{p \parallel q \xrightarrow{\uparrow o \downarrow i} p' \parallel q} \quad \frac{q \xrightarrow{\uparrow o \downarrow i} q'}{p \parallel q \xrightarrow{\uparrow o \downarrow i} p \parallel q'} \quad \frac{p \xrightarrow{v} Zero \quad q \xrightarrow{w} Zero}{p \parallel q \xrightarrow{(v,w)} Zero}$$

Moreover, these rules are complete, in the sense that any transition from $p \parallel q$ can be derived using them. Later we will consider more general situations in which both arguments may perform an effectful computation simultaneously, such as when one sends a message to the other.

It is also useful to define a variant of parallel composition that discards any result values produced by the left argument, hence this argument is only executed for its effects:

$$(\vec{\parallel}) :: CTree\ e\ a \rightarrow CTree\ e\ b \rightarrow CTree\ e\ b$$

$$p\ \vec{\parallel}\ q = fmap\ snd\ (p\ \parallel\ q)$$

Right-biased parallel composition $\vec{\parallel}$ can be characterised in a similar way to \parallel , except that the inference rule for values only propagates the right value w rather than the pair (v, w) .

3 EXAMPLE LANGUAGE

In this section, we introduce a simple concurrent language that we will use as an initial example for presenting our compiler calculation technique, and show how a semantics for this language can be defined in terms of choice trees. As we shall see, some care is required to ensure that the semantics correctly captures the intended concurrent behaviour.

3.1 Syntax

We consider a minimal expression language that comprises arithmetic expressions built up from integers values using an addition operator, extended with a primitive that prints the value of an expression, and a primitive that forks the evaluation of an expression:

```
data Expr = Val Int | Add Expr Expr | Print Expr | Fork Expr
```

Informally, *Fork* e starts evaluation of the expression e using a new concurrent process and immediately returns the result value 0, in a manner reminiscent of Haskell's *forkIO* primitive [Jones et al. 1996]. While the above language is not suitable for actual programming, it provides just what we need to explain our compiler calculation technique. In particular, the integers provide a simple notion of value, addition provides a simple notion of (sequential) computation, print provides a simple form of observable effect, and fork provides a simple form of concurrency.

3.2 Semantics

Using the choice tree machinery that was introduced in Section 2, a semantics for our simple expression language can be defined in terms of choice trees as follows:

$$eval :: Expr \rightarrow CTree\ PrintEff\ Int$$

$$eval\ (Val\ n) = return\ n$$

$$eval\ (Add\ x\ y) = do\ n \leftarrow eval\ x; m \leftarrow eval\ y; return\ (n + m)$$

$$eval\ (Print\ x) = do\ n \leftarrow eval\ x; print\ n; return\ n$$

$$eval\ (Fork\ x) = eval\ x\ \vec{\parallel}\ return\ 0$$

The first three cases are as we would expect, while the case for fork formalises the idea that the argument expression is evaluated in parallel with returning the result 0. Using right-biased parallel composition $\vec{\parallel}$ in the fork case is appropriate because our language both returns values and performs effects, and ensures that the argument expression x is evaluated purely for its effects.

While the above semantics for fork is simple, unfortunately it does not capture the desired behaviour. The problem is that the $\vec{\parallel}$ operator is *synchronous*, in the sense that it waits for both sides to complete. To illustrate the problem, consider the expression *Add* (*Fork* x) y with semantics:

$$(eval\ x\ \vec{\parallel}\ return\ 0) \gg \lambda n \rightarrow eval\ y \gg \lambda m \rightarrow return\ (n + m)$$

We would expect that x and y are evaluated in parallel – that is the point of using fork. However, in the above choice tree y is only evaluated after *eval* $x\ \vec{\parallel}\ return\ 0$ has completed, and hence only after x is evaluated. Instead, we would expect the semantics of *Add* (*Fork* x) y to be

$$eval\ x\ \vec{\parallel}\ (return\ 0 \gg \lambda n \rightarrow eval\ y \gg \lambda m \rightarrow return\ (n + m))$$

which then simplifies to $eval\ x \bar{\parallel} eval\ y$. Here, evaluation of x has been floated to the top-level, which ensures that it takes place *asynchronously*. In Haskell [Jones et al. 1996], this behaviour is realised by defining the semantics using an evaluation context that captures where the next step of evaluation takes place. In our setting, this gives the following semantics for `fork`:

$$eval\ (C[\text{Fork}\ x]) = eval\ x \bar{\parallel} eval\ (C[\text{Val}\ 0])$$

That is, if we are evaluating in a context C for which the next step is to fork an expression x , then we simply float evaluation of x to the top level, and continue with this expression replaced by the value zero. However, the above semantics is no longer compositional, because the semantics of $\text{Fork}\ x$ is no longer defined purely in terms of the semantics of x , but also involves the semantics for the residual expression $C[\text{Val}\ 0]$. Our compiler calculation methodology depends on the semantics being compositional, so we cannot use the contextual approach here.

Fortunately, we can achieve the same effect as the contextual semantics while retaining compositionality by rewriting the semantics in continuation-passing style. In particular, we take an additional argument c – the continuation – that is used to process of result of evaluating an expression, and hence captures the idea of what to do after the current evaluation:

$$\begin{aligned} eval &:: Expr \rightarrow (Int \rightarrow CTree\ PrintEff\ Int) \rightarrow CTree\ PrintEff\ Int \\ eval\ (\text{Val}\ n) &\quad c = c\ n \\ eval\ (\text{Add}\ x\ y) &\quad c = eval\ x\ (\lambda n \rightarrow eval\ y\ (\lambda m \rightarrow c\ (n + m))) \\ eval\ (\text{Print}\ x) &\quad c = eval\ x\ (\lambda n \rightarrow print\ n \gg \lambda () \rightarrow c\ n) \\ eval\ (\text{Fork}\ x) &\quad c = eval\ x\ return\ \bar{\parallel}\ c\ 0 \end{aligned}$$

This definition ensures that any continuation c that follows on from $\text{Fork}\ x$ is in evaluated in parallel with $eval\ x\ return$. For example, the expression $\text{Add}\ (\text{Fork}\ x)\ y$ has the semantics $eval\ x\ return\ \bar{\parallel}\ eval\ y\ c$, which ensures that x and y are evaluated in parallel as expected.

4 CODENSITY CHOICE TREES

While the continuation semantics in the previous section captures the intended behaviour, it is not so appealing as the simple, but incorrect, monadic semantics that we originally presented. More importantly, the explicit use of continuations would complicate the reasoning process, as we discuss later on (Section 6). However, we can regain both the simplicity of the monadic semantics and its reasoning principles by first generalising the return type of $eval$,

$$(Int \rightarrow CTree\ PrintEff\ Int) \rightarrow CTree\ PrintEff\ Int$$

from the specific case of integers to an arbitrary continuation input type a and result type r ,

$$(a \rightarrow CTree\ PrintEff\ r) \rightarrow CTree\ PrintEff\ r$$

and then observing that this is in fact the codensity monad for $CTree\ PrintEff$. The *codensity monad* [Voigtländer 2008] is similar to the familiar continuation monad, except that rather than having a fixed result type r , it has a variable (polymorphic) result type. Based on the generalised return type for $eval$, we can define a type of codensity choice trees $CTree_c\ e\ a$ as follows:

$$\text{type } CTree_c\ e\ a = \text{forall } r. (a \rightarrow CTree\ e\ r) \rightarrow CTree\ e\ r$$

Note that the return type r is universally quantified on the right-hand side of the declaration, in contrast to the continuation monad where r is a parameter on the left-hand side. Codensity choice trees form a monad, with the *return* and \gg operators defined as follows:

$$\begin{aligned} return &:: a \rightarrow CTree_c\ e\ a \\ return\ v &= \lambda c \rightarrow c\ v \end{aligned}$$

$$\begin{aligned}
(\gg) &:: CTree_c \ e \ a \rightarrow (a \rightarrow CTree_c \ e \ b) \rightarrow CTree_c \ e \ b \\
p \gg f &= \lambda c \rightarrow p (\lambda v \rightarrow f \ v \ c)
\end{aligned}$$

In turn, we can also define $CTree_c$ versions of the non-deterministic choice primitives, effect sequencing and printing, and the two versions of parallel composition:

$$\begin{aligned}
(\oplus_c) &:: CTree_c \ e \ a \rightarrow CTree_c \ e \ a \rightarrow CTree_c \ e \ a \\
p \oplus_c q &= \lambda c \rightarrow p \ c \oplus q \ c
\end{aligned}$$

$$\begin{aligned}
Zero_c &:: CTree_c \ e \ a \\
Zero_c &= \lambda c \rightarrow Zero
\end{aligned}$$

$$\begin{aligned}
Eff_c &:: e \ b \rightarrow (b \rightarrow CTree_c \ e \ a) \rightarrow CTree_c \ e \ a \\
Eff_c \ o \ k &= \lambda c \rightarrow Eff \ o \ (\lambda v \rightarrow k \ v \ c)
\end{aligned}$$

$$\begin{aligned}
print_c &:: Int \rightarrow CTree_c \ PrintEff \ () \\
print_c \ n &= Eff_c \ (PrintInt \ n) \ return
\end{aligned}$$

$$\begin{aligned}
(\parallel_c) &:: CTree_c \ e \ a \rightarrow CTree_c \ e \ b \rightarrow CTree_c \ e \ (a, b) \\
p \parallel_c q &= \lambda c \rightarrow (p \ return \ \parallel \ q \ return) \gg c
\end{aligned}$$

$$\begin{aligned}
(\vec{\parallel}_c) &:: CTree_c \ e \ a \rightarrow CTree_c \ e \ b \rightarrow CTree_c \ e \ b \\
p \vec{\parallel}_c q &= \lambda c \rightarrow p \ return \ \vec{\parallel} \ q \ c
\end{aligned}$$

Using these operations, it is now straightforward to redefine the continuation semantics from the previous section using the notion of codensity choice trees:

$$\begin{aligned}
eval &:: Expr \rightarrow CTree_c \ PrintEff \ Int \\
eval \ (Val \ n) &= return \ n \\
eval \ (Add \ x \ y) &= do \ n \leftarrow eval \ x; \ m \leftarrow eval \ y; \ return \ (n + m) \\
eval \ (Print \ x) &= do \ n \leftarrow eval \ x; \ print_c \ n; \ return \ n \\
eval \ (Fork \ x) &= eval \ x \ \vec{\parallel}_c \ return \ 0
\end{aligned}$$

This definition regains the simplicity of our original monadic definition of $eval$ from Section 3.2, but now correctly captures the intended semantics. Each $CTree_c$ represents a $CTree$, which can be obtained simply by passing $return$ for choice trees as the continuation:

$$\begin{aligned}
ctree &:: CTree_c \ e \ a \rightarrow CTree \ e \ a \\
ctree \ p &= p \ return
\end{aligned}$$

Using this translation function, the labelled transition system that defines the semantics for choice trees $CTree$ can be lifted to codensity choice trees $CTree_c$, and satisfies the same rules as Figure 1 with the operations replaced by the corresponding codensity versions.

5 BISIMILARITY AND LAWS

In this section we introduce the notion of bisimilarity for (codensity) choice trees, together with a number of laws for the operations on such trees that we will need for compiler calculation. We begin by considering choice trees, then extend to codensity choice trees.

Because the semantics of choice trees $CTree \ e \ a$ is defined using a labelled transition relation, the notion of bisimilarity \cong can be defined in the standard way. In particular, \cong is the largest relation on choice trees (and continuations) that satisfies the following two properties:

$$\begin{aligned}
&\text{If } p \cong q \text{ and } p \xrightarrow{l} p', \text{ then there is some } q' \text{ with } q \xrightarrow{l} q' \text{ and } p' \cong q' \\
&\text{If } p \cong q \text{ and } q \xrightarrow{l} q', \text{ then there is some } p' \text{ with } p \xrightarrow{l} p' \text{ and } p' \cong q'
\end{aligned}$$

$$\begin{array}{llll}
\text{return } x \gg f \cong f x & & & \\
p \gg \text{return} \cong p & & & \text{(monad laws)} \\
(p \gg f) \gg g \cong p \gg \lambda x \rightarrow (f x \gg g) & & & \\
\\
p \oplus p \cong p & & (\oplus\text{-idem}) & \\
p \oplus q \cong q \oplus p & & (\oplus\text{-comm}) & \\
\text{Zero} \oplus p \cong p & & (\oplus\text{-Zero}) & \\
(p \oplus q) \oplus r \cong p \oplus (q \oplus r) & & (\oplus\text{-assoc}) & \\
(p \oplus q) \gg f \cong (p \gg f) \oplus (q \gg f) & & (\oplus\text{-bind}) & \\
\text{Zero} \gg f \cong \text{Zero} & & (\text{Zero}\text{-bind}) & \\
\\
(\text{return } v) \bar{\parallel} p \cong p & & & (\bar{\parallel}\text{-return}) \\
(\text{fmap } f p) \bar{\parallel} q \cong p \bar{\parallel} q & & & (\bar{\parallel}\text{-fmap}) \\
(p \bar{\parallel} q) \bar{\parallel} r \cong p \bar{\parallel} (q \bar{\parallel} r) & & & (\bar{\parallel}\text{-assoc}) \\
(p \bar{\parallel} q) \bar{\parallel} r \cong (q \bar{\parallel} p) \bar{\parallel} r & & & (\bar{\parallel}\text{-comm})
\end{array}$$

Fig. 2. Laws for (codensity) choice trees.

Bisimilarity is an equivalence relation, i.e. reflexive, symmetric and transitive. Moreover, all operations on choice trees that we have defined satisfy congruence laws with respect to bisimilarity. For example, congruence for the monadic bind operator is stated as follows:

$$\frac{p \cong q \quad f x \cong g x \text{ for all } x}{p \gg f \cong q \gg g}$$

Other relevant laws for choice trees are given in Figure 2. In particular, choice trees form a monad, and an idempotent, commutative monoid under \oplus with *Zero* as the unit. Because $\bar{\parallel}$ only propagates result values for the right argument, using *return* and *fmap* in the left argument has no effect. As expected, $\bar{\parallel}$ also satisfies associativity and commutativity laws. However, because result values matter, commutativity only applies to the left argument of $\bar{\parallel}$, for which result values are discarded. An important observation is that \gg does not distribute over $\bar{\parallel}$, due to the synchronous nature of parallel composition for choice trees. In particular, with the expression $(p \bar{\parallel} q) \gg f$ the computation f can only start once both p and q have completed, whereas with $p \bar{\parallel} (q \gg f)$ the computation f can start as soon as q is complete and hence can run in parallel to p .

We now consider codensity choice trees. For our notion of bisimilarity on codensity choice trees we have two requirements. Firstly, it should imply bisimilarity of the corresponding choice trees, i.e. if codensity choice trees p and q are bisimilar, then so are the choice trees *ctree* p and *ctree* q . And secondly, it should satisfy the same reasoning principles as choice trees, i.e. congruence laws and the laws in Figure 2. To this end, we define bisimilarity for codensity choice trees as follows:

$$p \cong q \text{ iff } p c \cong q c \text{ for all } c$$

That is, two codensity choice trees are bisimilar precisely when they are bisimilar as choice trees for every continuation. Using this definition, it follows that if $p \cong q$ for two codensity choice trees, then their translations are bisimilar as choice trees, i.e. *ctree* $p \cong \text{ctree } q$.

In order to prove congruence laws and the laws in Figure 2, we restrict ourselves to codensity choice trees $p :: \text{CTree}_c e a$ that satisfy the following two well-formedness properties:

$$\text{if } c x \cong c' x \text{ for all } x, \text{ then } p c \cong p c' \text{ for all } c, c' :: a \rightarrow \text{CTree } e b \quad (\text{CTree}_c\text{-cont})$$

$$\text{fmap } f (p c) \cong p (\lambda r \rightarrow \text{fmap } f (c r)) \text{ for all } c :: a \rightarrow \text{CTree } e b, f :: b \rightarrow b' \quad (\text{CTree}_c\text{-fmap})$$

The first property states that if we supply two pointwise bisimilar continuations to a codensity choice tree, then we obtain two bisimilar choice trees. In turn, the second property states that *fmap*

distributes over the continuation of a codensity choice tree. All the operations on codensity choice trees that we present in this article preserve both well-formedness properties. That is, using these operators ensures that we can only construct well-formed codensity choice trees.

All laws in Figure 2 carry over to well-formed codensity choice trees, as do the congruence laws. In addition, we obtain the following distributivity law for parallel composition:

$$(p \bar{\parallel}_c q) \gg f \cong p \bar{\parallel}_c (q \gg f) \quad (\bar{\parallel}_c\text{-bind})$$

As observed above, this law does not hold for $p \bar{\parallel} q$ because the parallel computation is synchronous, whereas $p \bar{\parallel}_c q$ is asynchronous in that it produces a result as soon as q has completed, rather than also waiting for p . The $\bar{\parallel}_c$ -bind law captures this intuition formally. For example, this law can be used to show that $Add (Fork x) y$ has the intended semantics using codensity choice trees:

$$\begin{aligned} & eval (Add (Fork x) y) \\ \cong & \{ \text{applying } eval \} \\ & eval (Fork x) \gg (\lambda n \rightarrow eval y \gg \lambda m \rightarrow return (n + m)) \\ \cong & \{ \text{applying } eval \} \\ & (eval x \bar{\parallel}_c return 0) \gg (\lambda n \rightarrow eval y \gg \lambda m \rightarrow return (n + m)) \\ \cong & \{ \bar{\parallel}_c\text{-bind law} \} \\ & eval x \bar{\parallel}_c (return 0 \gg \lambda n \rightarrow eval y \gg \lambda m \rightarrow return (n + m)) \\ \cong & \{ \text{monad laws} \} \\ & eval x \bar{\parallel}_c eval y \end{aligned}$$

That is, x is evaluated in parallel with y , and the result produced by y is returned.

The crucial semantic difference between $\bar{\parallel}$ and $\bar{\parallel}_c$ discussed above also raises the question of whether we could define a version of $\bar{\parallel}$ on choice trees that *does* satisfy a distributivity law. Unfortunately, this is not possible. In general, with a parallel computation $p \bar{\parallel} q$ we expect that the effects of p may interact with the effects of q . Indeed, we extend parallel composition in Section 7 to allow such interaction. However, if the $\bar{\parallel}$ -bind law holds, then for an expression $(p \bar{\parallel} q) \gg f$ we should also expect that the effects of p may interact with the effects of $q \gg f$. Clearly, this cannot be the case for an operator $\bar{\parallel}$ that can only inspect p and q .

We conclude by noting that codensity choice trees are represented in a slightly different manner in our Agda formalisation. In particular, rather than defining $Ctree_c e a$ as the type $forall r. (a \rightarrow Ctree e r) \rightarrow Ctree e r$ subject to the well-formedness properties, it is defined as an inductive type with the operations $return$, \gg , \oplus_c , etc. as constructors, together with an interpretation function $sem :: Ctree_c e a \rightarrow (forall r. (a \rightarrow Ctree e r) \rightarrow Ctree e r)$ that performs the codensity construction as given in Section 4. In other words, the language of codensity choice trees is presented as a shallow embedding in this article, but as a deep embedding in our Agda formalisation. Using this approach allows us to prove that all our operations on codensity choice trees satisfy the required well-formedness properties, which simplifies the formalisation.

6 COMPILER CALCULATION

We have now defined a datatype $Expr$ that represents the syntax of a simple concurrent language, an evaluation function $eval$ that gives a semantics for the language in terms of codensity choice trees, and a notion of bisimilarity for such trees. In this section, we show how to specify the desired behaviour of a compiler for the simple language, and how such a specification can be used as the basis for systematically calculating an implementation of the compiler.

6.1 Specification

Our goal is to define a compilation function $comp :: Expr \rightarrow Code$ that translates an expression into code for an (as yet unspecified) target language. We assume the compiler targets a stack-based machine, whose semantics is given by a function $exec :: Code \rightarrow Stack \rightarrow Stack$, where $type\ Stack = [Int]$ is the stack type for the machine. However, because our evaluation function defines the semantics of expressions in terms of codensity choice trees,

$$eval :: Expr \rightarrow CTree_c\ PrintEff\ Int$$

we also generalise the type of the execution function to operate in the same monadic setting, i.e. within the codensity monad $CTree_c\ PrintEff$:

$$exec :: Code \rightarrow Stack \rightarrow CTree_c\ PrintEff\ Stack$$

The definitions for the $Code$ datatype and $exec$ function are not given up front, but will rather fall out naturally as part of the process of calculating the compiler itself.

Prior to specifying the desired behaviour of the compiler, we generalise the function $comp$ to take additional code to be executed after the compiled code. The addition of such a *code continuation* is a key aspect of the methodology [Bahr and Hutton 2015], and significantly simplifies the resulting calculations. Using this idea, our goal now is to establish the following compiler correctness property for the generalised compilation function $comp : Expr \rightarrow Code \rightarrow Code$:

$$do\ v \leftarrow eval\ e; exec\ c\ (v : s) \cong exec\ (comp\ e\ c)\ s$$

That is, compiling an expression and then executing the resulting code together with the supplied additional code should give the same result (up to bisimilarity) as executing the additional code with the value of the expression on top of the stack.

6.2 Calculation

The proof of the compiler correctness property proceeds by structural induction on the expression e . For each case, we start with the left-hand side of the property, and seek to transform it by equational reasoning using the bisimilarity laws from Section 5 into the form $exec\ c'\ s$ for some code c' . We then define $comp\ e\ c = c'$, which gives us a clause for the compiler in this case that is guaranteed by construction to satisfy the correctness property. During such calculations we will also discover new constructors for the $Code$ datatype and new clauses for the $exec$ function, driven by the desire to transform the term that is being manipulated into the required form.

The cases for Val and Add proceed in the same manner as Bahr and Hutton [2022], except that because our language doesn't yet include non-termination, simple bisimilarity suffices rather than the more refined notion of step-indexed bisimilarity. In the Val case, we first apply the definition of the evaluation function, and then simplify the resulting term using the monad laws:

$$\begin{aligned} & do\ v \leftarrow eval\ (Val\ n); exec\ c\ (v : s) \\ \cong & \quad \{ \text{definition of } eval \} \\ & do\ v \leftarrow return\ n; exec\ c\ (v : s) \\ \cong & \quad \{ \text{monad laws} \} \\ & exec\ c\ (n : s) \end{aligned}$$

Then, to complete the calculation, we need to arrive at a term of the form $exec\ c'\ s$. That is, we need to find some code c' that solves the following bisimilarity equation:

$$exec\ c'\ s \cong exec\ c\ (n : s)$$

Note that we cannot simply strengthen bisimilarity to equality and use the resulting equation $exec\ c'\ s = exec\ c\ (n : s)$ as a defining clause for the function $exec$, as the variables n and c would be unbound in the body of the definition. The solution is to package these two variables up in the code argument c' , which can freely be instantiated as it is existentially quantified. This can be achieved by adding a new constructor to the *Code* datatype that takes n and c as arguments,

$$PUSH :: Int \rightarrow Code \rightarrow Code$$

and defining a new clause for the function $exec$ as follows:

$$exec\ (PUSH\ n\ c)\ s = exec\ c\ (n : s)$$

That is, the code $PUSH\ n\ c$ is executed by pushing n onto the stack and then executing the code c , which motivates the name for the new code constructor. This definition solves the above equation, and allows us to complete the transformation into the required form:

$$\begin{aligned} & exec\ c\ (n : s) \\ \cong & \quad \{ \text{definition of } exec \} \\ & exec\ (PUSH\ n\ c)\ s \end{aligned}$$

In summary, via the above calculation we have discovered a new code constructor $PUSH$, a corresponding new clause for the $exec$ function, and a new clause for the compiler, namely $comp\ (Val\ n)\ c = PUSH\ n\ c$. In turn, the *Add* case proceeds as follows:

$$\begin{aligned} & \mathbf{do}\ v \leftarrow eval\ (Add\ x\ y); exec\ c\ (v : s) \\ \cong & \quad \{ \text{definition of } eval \} \\ & \mathbf{do}\ v \leftarrow \mathbf{do}\ \{ m \leftarrow eval\ x; n \leftarrow eval\ y; return\ (m + n) \}; exec\ c\ (v : s) \\ \cong & \quad \{ \text{monad laws} \} \\ & \mathbf{do}\ m \leftarrow eval\ x; n \leftarrow eval\ y; exec\ c\ ((m + n) : s) \\ \cong & \quad \{ \text{define: } exec\ (ADD\ c)\ (n : m : s) = exec\ c\ ((m + n) : s) \} \\ & \mathbf{do}\ m \leftarrow eval\ x; n \leftarrow eval\ y; exec\ (ADD\ c)\ (n : m : s) \\ \cong & \quad \{ \text{induction hypothesis for } y \} \\ & \mathbf{do}\ m \leftarrow eval\ x; exec\ (comp\ y\ (ADD\ c))\ (m : s) \\ \cong & \quad \{ \text{induction hypothesis for } x \} \\ & exec\ (comp\ x\ (comp\ y\ (ADD\ c)))\ s \end{aligned}$$

In the third step above, we introduce another code constructor ADD and clause for $exec$. In particular, in order to apply the induction hypothesis for y , we need to transform the term $exec\ c\ ((m + n) : s)$ into the form $exec\ c'\ (n : s')$ for some code c' and stack s' . We achieve this instantiating $c' = ADD\ c$ and $s' = m : s$, and defining a corresponding new clause for $exec$. Applying the two induction hypotheses then completes the transformation into the required form, and hence we have discovered another clause for the compiler, $comp\ (Add\ x\ y)\ c = comp\ x\ (comp\ y\ (ADD\ c))$.

The case for *Print* is straightforward, once again introducing a new clause for $exec$ that brings the term into the form that is required to apply the induction hypothesis:

$$\begin{aligned} & \mathbf{do}\ v \leftarrow eval\ (Print\ x); exec\ c\ (v : s) \\ \cong & \quad \{ \text{definition of } eval \} \\ & \mathbf{do}\ v \leftarrow \mathbf{do}\ \{ n \leftarrow eval\ x; print\ n; return\ n \}; exec\ c\ (v : s) \\ \cong & \quad \{ \text{monad laws} \} \\ & \mathbf{do}\ n \leftarrow eval\ x; print\ n; exec\ c\ (n : s) \\ \cong & \quad \{ \text{define: } exec\ (PRINT\ c)\ (n : s) = \mathbf{do}\ print\ n; exec\ c\ (n : s) \} \\ & \mathbf{do}\ n \leftarrow eval\ x; exec\ (PRINT\ c)\ (n : s) \end{aligned}$$

$$\cong \{ \text{induction hypothesis for } x \}$$

$$\text{exec } (\text{comp } x \text{ (PRINT } c)) \text{ } s$$

Finally, the case for *Fork* first exploits the laws for right-biased parallel composition to introduce a minimal new clause for *exec* that allows the induction hypothesis to be applied, and then introduces a further new clause to bring the resulting term into the required form:

$$\text{do } v \leftarrow \text{eval } (\text{Fork } x); \text{exec } c \text{ } (v : s)$$

$$\cong \{ \text{definition of eval} \}$$

$$\text{do } v \leftarrow (\text{eval } x \vec{\pi}_c \text{ return } 0); \text{exec } c \text{ } (v : s)$$

$$\cong \{ \text{do notation} \}$$

$$(\text{eval } x \vec{\pi}_c \text{ return } 0) \succcurlyeq \lambda v \rightarrow \text{exec } c \text{ } (v : s)$$

$$\cong \{ \vec{\pi}_c\text{-bind law} \}$$

$$\text{eval } x \vec{\pi}_c (\text{return } 0 \succcurlyeq \lambda v \rightarrow \text{exec } c \text{ } (v : s))$$

$$\cong \{ \text{monad laws} \}$$

$$\text{eval } x \vec{\pi}_c \text{exec } c \text{ } (0 : s)$$

$$\cong \{ \vec{\pi}_c\text{-fmap law} \}$$

$$\text{fmap } (\lambda v \rightarrow [v]) (\text{eval } x) \vec{\pi}_c \text{exec } c \text{ } (0 : s)$$

$$\cong \{ \text{definition of fmap} \}$$

$$(\text{do } v \leftarrow \text{eval } x; \text{return } [v]) \vec{\pi}_c \text{exec } c \text{ } (0 : s)$$

$$\cong \{ \text{define: exec HALT } s = \text{return } s \}$$

$$(\text{do } v \leftarrow \text{eval } x; \text{exec HALT } [v]) \vec{\pi}_c \text{exec } c \text{ } (0 : s)$$

$$\cong \{ \text{induction hypothesis for } x \}$$

$$(\text{exec } (\text{comp } x \text{ HALT}) []) \vec{\pi}_c \text{exec } c \text{ } (0 : s)$$

$$\cong \{ \text{define: exec (FORK } c' \text{ } c) \text{ } s = \text{exec } c' \text{ } [] \vec{\pi}_c \text{exec } c \text{ } (0 : s) \}$$

$$\text{exec } (\text{FORK } (\text{comp } x \text{ HALT}) \text{ } c) \text{ } s$$

Note how the $\vec{\pi}_c$ -*fmap* law is used above to transform *eval x* into **do** $v \leftarrow \text{eval } x; \text{return } [v]$, which places the result of evaluation into a singleton stack. This transformation is valid because $\vec{\pi}_c$ discards the result value produced by its left argument, and allows us to introduce a code constructor *HALT* that simply returns the current stack, and then apply the induction hypothesis.

To complete the compiler calculation, we consider a top-level function $\text{compile} :: \text{Expr} \rightarrow \text{Code}$ that compiles expressions into code, whose correctness is captured by the following property:

$$\text{do } v \leftarrow \text{eval } e; \text{return } (v : s) \cong \text{exec } (\text{compile } e) \text{ } s$$

Using the correctness of *comp*, it is straightforward to calculate the definition for *compile*:

$$\text{do } v \leftarrow \text{eval } e; \text{return } (v : s)$$

$$\cong \{ \text{definition of exec} \}$$

$$\text{do } v \leftarrow \text{eval } e; \text{exec HALT } (v : s)$$

$$\cong \{ \text{correctness of comp} \}$$

$$\text{exec } (\text{comp } e \text{ HALT}) \text{ } s$$

In summary, we have calculated the definitions in Figure 3.

6.3 Reflection

We conclude this section with some reflective remarks. First of all, note that the *exec* function is not total because addition and printing require stacks of specific forms, e.g. *ADD* requires a stack

Target language

data $Code = PUSH\ Int\ Code \mid ADD\ Code \mid PRINT\ Code \mid FORK\ Code\ Code \mid HALT$

Compiler

$compile :: Expr \rightarrow Code$
 $compile\ e = comp\ e\ HALT$
 $comp :: Expr \rightarrow Code \rightarrow Code$
 $comp\ (Val\ n)\ c = PUSH\ n\ c$
 $comp\ (Add\ x\ y)\ c = comp\ x\ (comp\ y\ (ADD\ c))$
 $comp\ (Print\ x)\ c = comp\ x\ (PRINT\ c)$
 $comp\ (Fork\ x)\ c = FORK\ (comp\ x\ HALT)\ c$

Virtual machine

$exec :: Code \rightarrow Stack \rightarrow CTree_c\ PrintEff\ Stack$
 $exec\ (PUSH\ n\ c)\ s = exec\ c\ (n : s)$
 $exec\ (ADD\ c)\ (n : m : s) = exec\ c\ ((m + n) : s)$
 $exec\ (PRINT\ c)\ (n : s) = \mathbf{do}\ print\ n ; exec\ c\ (n : s)$
 $exec\ (FORK\ c'\ c)\ s = exec\ c'\ [] \vec{\parallel}_c\ exec\ c\ (0 : s)$
 $exec\ HALT\ s = return\ s$

Fig. 3. Compiler and virtual machine for the simple expression language.

with at least two elements. To make it total, we can add the catch-all case $exec\ _ = Zero_c$, but the choice of semantics here is not important as compiler correctness does not depend on it.

Secondly, the $exec$ function returns a collection of parallel computations in which recursive calls are always tail calls, i.e. the final operation performed, which justifies referring to $exec$ as a (parallel) virtual machine. More precisely, using the transition semantics of codensity choice trees we can observe that all transition sequences starting from $exec\ c\ s$ are of the form

$$exec\ c\ s \xrightarrow{\uparrow o_1 \downarrow i_1} p_1 \xrightarrow{\uparrow o_2 \downarrow i_2} p_2 \xrightarrow{\uparrow o_3 \downarrow i_3} \dots \xrightarrow{s'} Zero_c$$

where each p_i is bisimilar to an expression of the form

$$exec\ c_1\ s_1 \vec{\parallel}_c\ exec\ c_2\ s_2 \vec{\parallel}_c\ \dots \vec{\parallel}_c\ exec\ c_{n+1}\ s_{n+1}$$

That is, the state of the virtual machine consists of $n + 1$ parallel threads of execution, each with its own code and stack. The result stack s' from the whole execution is produced by the rightmost thread, while the remaining n threads are only executed for their effects, and the order of these threads does not matter according to the $\vec{\parallel}_c$ -comm law. Using the $\vec{\parallel}_c$ -return law, $HALT$ kills the current thread, and using the $\vec{\parallel}_c$ -assoc law, $FORK$ spawns a new thread with an empty stack.

And finally, as shown in Section 3, we can capture the semantics of the source language using choice trees alone if we use a continuation-passing style, and indeed we can calculate a compiler based on this semantics. However, this approach has the drawback of having to prove ad-hoc lemmas about each semantics, e.g. congruence and distributivity properties such as:

$$\frac{c\ x \cong c'\ x \quad \text{for all } x}{eval\ e\ c \cong eval\ e\ c'} \quad \frac{}{fmap\ f\ (eval\ e\ c) \cong eval\ e\ (\lambda x \rightarrow fmap\ f\ (c\ x))}$$

These properties suggest that the continuation-passing style semantics $eval\ e\ c$ behaves similarly to $eval\ e\ \gg c$ for a suitably defined monad. Codensity choice trees make this idea precise, and provide an abstraction that satisfies these and other crucial structural properties *by construction*. In contrast, the continuation-passing style semantics makes essentially no use of the monadic structure of choice trees, because its only use of \gg can be replaced by continuation passing.

7 NON-TERMINATION AND EFFECT HANDLERS

For expository purposes, we have used a simplified version of (codensity) choice trees so far. We now give the full definition, which will allow us to consider non-terminating computations (Section 7.1). In addition, we introduce concurrent effect handlers on choice trees to allow us

to consider concurrent computations that can interact, e.g. by sending and receiving messages (Sections 7.2 and 7.3). The resulting semantic structure forms the basis for our compiler calculation for a concurrent lambda calculus with channel-based communication in Section 8.

7.1 Non-termination

To support non-termination, we extend *CTree* with an additional constructor *Step*, whose argument is a value of a coinductive type *CTreeInf* with a single constructor *Delay*:

data *CTree* *e a* **where**

...

Step :: *CTreeInf* *e a* → *CTree* *e a*

codata *CTreeInf* *e a* = *Delay* (*CTree* *e a*)

That is, *CTree* is still an inductive type, but is now defined mutually recursively with a coinductive type *CTreeInf*, which we indicate by writing **codata** instead of **data**. In this manner, a value of type *CTree* *e a* is a potentially infinite tree with nodes labelled by the constructors *Now*, *Step* and so on, such that every infinite path from the root must contain infinitely many nodes labelled *Step*. Subsequently, we write *Later* *p* as a shorthand for *Step* (*Delay* *p*) and don't use *Step* or *Delay* directly. For example, a computation that never terminates can be defined as follows:

never :: *CTree* *e a*

never = *Later* *never*

Despite being non-terminating, this definition is total because the recursive call is guarded by *Later*, and systems such as Agda will accept it. The transition semantics for choice trees, and hence the notion of bisimilarity, is extended to account for non-termination by adding the following transition rule, which expresses that the effect of *Later* is a silent transition τ :

$$\frac{}{\text{Later } p \xRightarrow{\tau} p}$$

For example, *never* gives the infinite transition sequence $\text{never} \xRightarrow{\tau} \text{never} \xRightarrow{\tau} \text{never} \xRightarrow{\tau} \dots$. In essence, *Later* is a more restrictive variant of *Eff* that can be used to express non-terminating behaviour. With this intuition in mind, we can extend monadic bind and parallel composition to take account of non-termination by adding the following clauses:

$\text{Later } p \gg f = \text{Later } (p \gg f)$

$\text{Later } p \triangleleft q = \text{Later } (p \parallel q)$

$p \triangleright \text{Later } q = \text{Later } (p \parallel q)$

To prove bisimilarity properties for choice trees defined co-recursively, such as *never*, we need a (co)-induction principle that is powerful enough to account for this. To this end, we follow the approach of Bahr and Hutton [2022] and use a step-indexed version of bisimilarity, denoted by \cong_i , indexed by a natural number *i* that counts the number of steps. While \cong is defined coinductively, \cong_i is defined inductively as the smallest relation such that $p \cong_0 q$ holds, and moreover $p \cong_{i+1} q$ holds if the following two conditions hold (where $j = i$ if $l = \tau$, and $j = i + 1$ otherwise):

If $p \xRightarrow{l} p'$ then there is some $q \xRightarrow{l} q'$ with $p' \cong_j q'$

If $q \xRightarrow{l} q'$ then there is some $p \xRightarrow{l} p'$ with $p' \cong_j q'$

We can show (classically, using the law of excluded middle) that $p \cong q$ iff $p \cong_i q$ for all i . In particular, this means that we can prove bisimilarity $p \cong q$ by proving $p \cong_i q$ by induction on i . To do this, we make use the following congruence law for *Later*:

$$\frac{p \cong_j q \quad \text{for all } j < i}{\text{Later } p \cong_i \text{ Later } q}$$

That is, whenever we ‘go under’ a *Later* we may use our induction hypothesis because we decrease the step index i . For example, suppose we define a variant of *never* that invokes *Later* twice before making a recursive call by $\text{never}' = \text{Later} (\text{Later } \text{never}')$. Then we can prove that never' and *never* are bisimilar by proving that $\text{never}' \cong_i \text{never}$ by induction on i , which proceeds as follows:

$$\begin{aligned} & \text{never}' \\ \cong_i & \quad \{ \text{definition of } \text{never}' \} \\ & \text{Later} (\text{Later } \text{never}') \\ \cong_i & \quad \{ \text{induction hypothesis, under two } \text{Later}s \} \\ & \text{Later} (\text{Later } \text{never}) \\ \cong_i & \quad \{ \text{definition of } \text{never}, \text{ applied twice} \} \\ & \text{never} \end{aligned}$$

The extended definition of choice trees carries over to the codensity construction. In particular, we have a Later_c operation on codensity choice trees CTree_c ,

$$\text{Later}_c :: \text{CTree}_c \ e \ a \rightarrow \text{CTree}_c \ e \ a$$

$$\text{Later}_c \ p = \lambda c \rightarrow \text{Later} \ (p \ c)$$

as well a step-indexed bisimilarity relation, defined by:

$$p \cong_i q \quad \text{iff} \quad p \ c \cong_i \ q \ c \quad \text{for all } c$$

All our previous laws for (codensity) choice trees, e.g. the congruence laws and those in Figure 2, also hold for step-indexed bisimilarity. In addition, we have the following laws for Later_c :

$$\frac{p \cong_j q \quad \text{for all } j < i}{\text{Later}_c \ p \cong_i \ \text{Later}_c \ q} \quad (\text{Later}_c\text{-CONG}) \qquad \frac{}{\text{Later}_c \ p \ \bowtie \ f \cong_i \ \text{Later}_c \ (p \ \bowtie \ f)} \quad (\text{Later}_c\text{-BIND})$$

Similarly to choice trees, $\text{Later}_c\text{-CONG}$ provides us with a powerful induction principle for codensity choice trees as we can prove $p \cong q$ by proving $p \cong_i q$ for all i .

7.2 Concurrent Effect Handlers

While (codensity) choice trees have a parallel composition operator, there is no non-trivial interaction between parallel computations. In particular, there is no way for such computations to communicate with each other. To support this, we need to extend the definition of the \bowtie operator from Section 2.3, which describes how two parallel computations interact. With the current definition, only very simple interactions are possible, namely that if both computations finish with a result value then their parallel composition finishes with the combined result value:

$$\text{Now } v \bowtie \text{Now } w = \text{Now} \ (v, w)$$

To allow effects of the two computations to interact, we parameterise the parallel composition operator with a type class that provides a concurrent effect handler:

class *Concurrent* e **where**

$$(\varpi) :: e \ a \rightarrow e \ b \rightarrow \text{CTree} \ e \ (a, b)$$

The definition of \bowtie is then generalised so that it uses this concurrent effect handler, which is achieved by adding the following clause to its definition:

$$\mathit{Eff} \ e_1 \ c_1 \bowtie \mathit{Eff} \ e_2 \ c_2 = (e_1 \overset{\tau}{\rightleftharpoons} e_2) \succcurlyeq \lambda(x, y) \rightarrow c_1 \ x \parallel c_2 \ y$$

To ensure associativity of parallel composition \parallel , as well as the right-biased version $\bar{\parallel}$, we require that choice trees $e_1 \overset{\tau}{\rightleftharpoons} e_2$ can only have transitions of the form

$$e_1 \overset{\tau}{\rightleftharpoons} e_2 \xrightarrow{\tau} \mathit{return} \ v$$

which means that $e_1 \overset{\tau}{\rightleftharpoons} e_2$ is (bisimilar to) a sum of terms of the form $\mathit{Later} \ (\mathit{return} \ (v, w))$, i.e. two concurrent effects e_1 and e_2 are handled by a silent transition τ that provides return values v and w for the two effects. If the sum is empty, the two effects do not interact concurrently, whereas if there is more than one summand, their interaction is non-deterministic. To ensure commutativity, we also require that $\overset{\tau}{\rightleftharpoons}$ is commutative in the following way:

$$e_1 \overset{\tau}{\rightleftharpoons} e_2 \xrightarrow{\tau} \mathit{return} \ (v, w) \quad \text{implies} \quad e_2 \overset{\tau}{\rightleftharpoons} e_1 \xrightarrow{\tau} \mathit{return} \ (w, v)$$

The resulting generalised \parallel operator specialises to the previous version if $\overset{\tau}{\rightleftharpoons}$ always returns Zero . For example, the instance declaration for the printing effect is simply:

instance *Concurrent PrintEff* **where**

$$_ \overset{\tau}{\rightleftharpoons} _ = \mathit{Zero}$$

For a more interesting example, let us consider an effect type *CommEff* that supports communication between computations in the form of sending and receiving integer values:

data *CommEff* **a** **where**

$$\mathit{Send} \quad \text{::} \ \mathit{Int} \rightarrow \mathit{CommEff} \ ()$$

$$\mathit{Receive} \text{::} \ \mathit{CommEff} \ \mathit{Int}$$

$$\mathit{send} \text{::} \ \mathit{Int} \rightarrow \mathit{CTree} \ \mathit{CommEff} \ ()$$

$$\mathit{send} \ n = \mathit{Eff} \ (\mathit{Send} \ n) \ \mathit{return}$$

$$\mathit{receive} \text{::} \ \mathit{CTree} \ \mathit{CommEff} \ \mathit{Int}$$

$$\mathit{receive} = \mathit{Eff} \ \mathit{Receive} \ \mathit{return}$$

For the two operations to interact in the desired way, we define $\overset{\tau}{\rightleftharpoons}$ as follows:

instance *Concurrent CommEff* **where**

$$\mathit{Send} \ n \overset{\tau}{\rightleftharpoons} \mathit{Receive} = \mathit{Later} \ (\mathit{return} \ ((), n))$$

$$\mathit{Receive} \overset{\tau}{\rightleftharpoons} \mathit{Send} \ n = \mathit{Later} \ (\mathit{return} \ (n, ()))$$

$$_ \overset{\tau}{\rightleftharpoons} _ = \mathit{Zero}$$

For example, we have the following transitions:

$$\mathit{send} \ 1 \ \bar{\parallel} \ \mathit{receive} \xrightarrow{\tau} \mathit{return} \ () \ \bar{\parallel} \ \mathit{return} \ 1 \xrightarrow{1} \mathit{Zero}$$

All previously presented laws for $\bar{\parallel}$ and $\bar{\parallel}_c$ (see Figure 2) carry over to this extension with concurrent effect handlers. Moreover, the rules characterising the transitions of $\bar{\parallel}$ and $\bar{\parallel}_c$ (see Section 2.3) carry over to this generalisation. However, to make the rules complete, i.e. any transition starting from $p \ \bar{\parallel} \ q$ can be derived using them, we additionally also prove the following rules:

$$\frac{p \xrightarrow{\tau} p'}{p \ \bar{\parallel} \ q \xrightarrow{\tau} p \ \bar{\parallel} \ q} \quad \frac{q \xrightarrow{\tau} q'}{p \ \bar{\parallel} \ q \xrightarrow{\tau} p \ \bar{\parallel} \ q'} \quad \frac{p \xrightarrow{\uparrow e} c \quad p' \xrightarrow{\uparrow e'} c' \quad e \overset{\tau}{\rightleftharpoons} e' \xrightarrow{\tau} \mathit{return} \ (v, v')}{p \ \bar{\parallel} \ p' \xrightarrow{\tau} c \ v \ \bar{\parallel} \ c' \ v'}$$

These account for both the addition of the *Later* constructor and concurrent effect handlers.

7.3 Effect Handlers

While the generalised parallel composition *allows* communication, it does not prevent communication effects to trigger independently. For example, we also have the transition

$$\text{send } 1 \ \bar{u} \ \text{receive} \xrightarrow{\uparrow() \downarrow 2} \text{send } 1 \ \bar{u} \ \text{return } 2$$

in which the value 2 is received from the outside context independently of the parallel computation *send* 1. To restrict such communication to a local context, we follow [Chappe et al. \[2023\]](#) and use an effect handling function *interp*, defined as follows:

$$\begin{aligned} \text{interp} &:: (\text{forall } b. e \ b \rightarrow \text{CTree } f \ b) \rightarrow \text{CTree } e \ a \rightarrow \text{CTree } f \ a \\ \text{interp } \text{han} \ (\text{Now } v) &= \text{Now } v \\ \text{interp } \text{han} \ (\text{Later } p) &= \text{Later } (\text{interp } \text{han } p) \\ \text{interp } \text{han} \ (p \oplus q) &= \text{interp } \text{han } p \oplus \text{interp } \text{han } q \\ \text{interp } \text{han} \ \text{Zero} &= \text{Zero} \\ \text{interp } \text{han} \ (\text{Eff } e \ c) &= \text{han } e \ \succ \lambda i \rightarrow \text{interp } \text{han} \ (c \ i) \end{aligned}$$

The argument *han* handles each effect from the effect type *e* by interpreting it as a choice tree with a potentially different effect type *f*. We restrict effects to a local context by simply removing all effects, which is achieved by interpreting them by the choice tree *Zero*:

$$\begin{aligned} \text{restrict} &:: \text{CTree } \text{CommEff} \ a \rightarrow \text{CTree } e \ a \\ \text{restrict} &= \text{interp} \ (\lambda_ \rightarrow \text{Zero}) \end{aligned}$$

Note that the result type of *restrict* is polymorphic in the effect type *e*. In particular, we could choose the empty effect type, which witnesses the fact that there are indeed no observable effects other than τ transitions. Using restriction, we now only have a single transition from the choice tree *restrict* (*send* 1 \bar{u} *receive*), namely the τ transition to *restrict* (*return* () \bar{u} *return* 1). That is, values are now prevented from being sent to and received from the outside context.

We define the corresponding variant of *interp* for codensity choice trees as follows:

$$\begin{aligned} \text{interp}_c &:: (\text{forall } b. e \ b \rightarrow \text{CTree}_c \ f \ b) \rightarrow \text{CTree}_c \ e \ a \rightarrow \text{CTree}_c \ f \ a \\ \text{interp}_c \ f \ p \ c &= \text{interp} \ (\lambda e \rightarrow \text{ctree} \ (f \ e)) \ (\text{ctree} \ p) \ \succ \ c \end{aligned}$$

In this definition, the bind operator for choice trees is used to apply the continuation. Following the extension of other choice tree operators to codensity choice trees, we may have expected the following definition, in which the continuation is passed directly to *p*:

$$\text{interp}'_c \ f \ p \ c = \text{interp} \ (\lambda e \rightarrow \text{ctree} \ (f \ e)) \ (p \ c)$$

However, this definition suffers from two drawbacks. First of all, its type would be restricted to the case where *f* = *e*, which means that it would not be applicable for effect handlers that change the set of effects. For example, the type of the *restrict* function would no longer reflect the absence of observable effects. And secondly, *interp}'_c* satisfies the following law:

$$\text{interp}'_c \ h \ p \ \succ \ f \ \cong_i \ \text{interp}'_c \ h \ (p \ \succ \ f)$$

This law is similar to the \bar{u}_c -bind law and could be useful for calculation. However, it also reveals that the scope of *interp}'_c* extends arbitrarily to the right of a bind operator, which would make it unsuitable for defining a restriction function with a delimited scope. Instead of the above undesirable law, both *interp* and *interp}_c* satisfy the following restricted version:

$$\text{fmap } f \ (\text{interp } h \ p) \ \cong_i \ \text{interp } h \ (\text{fmap } f \ p) \quad (\text{interp-fmap})$$

In addition, both *interp* and *interp}_c* also satisfy congruence laws.

Finally, we note that effect handlers can also be generalised so that they can use an internal state, by means of the following interpretation functions:

$$\begin{aligned}
\text{interpSt} &:: s \rightarrow (\text{forall } b. s \rightarrow e b \rightarrow \text{CTree } f (b, s)) \rightarrow \text{CTree } e a \rightarrow \text{CTree } f a \\
\text{interpSt } s f (\text{Now } v) &= \text{Now } v \\
\text{interpSt } s f (\text{Later } p) &= \text{Later } (\text{interpSt } s f p) \\
\text{interpSt } s f (p \oplus q) &= \text{interpSt } s f p \oplus \text{interpSt } s f q \\
\text{interpSt } s f \text{Zero} &= \text{Zero} \\
\text{interpSt } s f (\text{Eff } e c) &= f s e \gg (\lambda(x, s') \rightarrow \text{interpSt } s' f (c x)) \\
\text{interpSt}_c &:: s \rightarrow (\text{forall } b. s \rightarrow e b \rightarrow \text{CTree}_c f (b, s)) \rightarrow \text{CTree}_c e a \rightarrow \text{CTree}_c f a \\
\text{interpSt}_c s f p c &= \text{interpSt } s (\lambda st e \rightarrow \text{ctree } (f st e)) (\text{ctree } p) \gg c
\end{aligned}$$

For example, instead of using *restrict* to prevent communication with the outside context, we could simulate a context that stores sent values and returns them back when asked:

$$\begin{aligned}
\text{parrot} &:: \text{Int} \rightarrow \text{CTree } \text{CommEff } a \rightarrow \text{CTree } e a \\
\text{parrot } s &= \text{interpSt } s \text{han } \mathbf{where} \\
\text{han} &:: \text{Int} \rightarrow \text{CommEff } b \rightarrow \text{CTree } e (b, \text{Int}) \\
\text{han } s (\text{Send } n) &= \text{Later } (\text{return } ((), n)) \\
\text{han } s \text{Receive} &= \text{Later } (\text{return } (s, s))
\end{aligned}$$

For example, we have the following transitions:

$$\text{parrot } 0 (\text{send } 1 \gg \text{receive}) \xRightarrow{\tau} \text{parrot } 1 \text{receive} \xRightarrow{\tau} \text{return } 1$$

8 CONCURRENT LAMBDA CALCULUS

To demonstrate that codensity choice trees and their associated reasoning principles scale to more sophisticated concurrent languages, we show how to calculate a compiler for an untyped lambda calculus extended with concurrency and channel-based communication.

8.1 Syntax

The syntax for our language is defined as follows, in which bound variables are represented using de Bruijn indices, and channel names are represented as integers:

$$\begin{aligned}
\mathbf{data} \text{ Expr} &= \text{Val } \text{Int} \mid \text{Add } \text{Expr } \text{Expr} \mid \\
&\quad \text{Var } \text{Int} \mid \text{Abs } \text{Expr} \mid \text{App } \text{Expr } \text{Expr} \mid \\
&\quad \text{Send } \text{Expr } \text{Expr} \mid \text{Receive } \text{Expr} \mid \text{Fork } \text{Expr}
\end{aligned}$$

Informally, *Var* i is the variable with de Bruijn index $i \geq 0$, *Abs* x constructs an abstraction over the expression x , *App* $x y$ applies the abstraction that results from evaluation of x to the value of y , *Send* $c x$ sends the integer that results from evaluation of x on the channel c , and *Receive* c receives an integer on the channel c . Finally, *Fork* x spawns a new concurrent process to evaluate x , creates a new channel c that is passed to this process, and returns c as the result value.

In this lambda calculus, we can express complex concurrent programs that fork processes, create communication channels, and pass integers on those channels. For example, using a standard syntax that can readily be translated into the *Expr* type, the following program forks a new process that increments an integer n received on the newly created channel c :

$$\text{fork } (\lambda c. \mathbf{let } n = \text{receive } c \mathbf{in } \text{send } c (n + 1))$$

8.2 Semantics

We begin by defining an effect type *ChanEff* for channels that support sending and receiving integers and creating new channels, where channels themselves are simply integers:

```

type Chan = Int
data ChanEff a where
  SendInt  :: Chan → Int → ChanEff ()
  ReceiveInt :: Chan → ChanEff Int
  NewChan  :: ChanEff Chan

send :: Chan → Int → CTreec ChanEff ()
send c n = Effc (SendInt c n) return

receive :: Chan → CTreec ChanEff Int
receive c = Effc (ReceiveInt c) return

newChan :: CTreec ChanEff Chan
newChan = Effc NewChan return

```

To use the parallel composition operator, we also provide a concurrent effect handler \rightleftharpoons that expresses the interaction of send and receive effects. In particular, *SendInt c n* causes any concurrent effect *ReceiveInt c* on the same channel *c* to evaluate to the integer *n*:

```

instance Concurrent ChanEff where
  SendInt c n  ⇔ ReceiveInt c' | c ≡ c' = Later (return ((), n))
  ReceiveInt c' ⇔ SendInt c n  | c ≡ c' = Later (return (n, ()))
  -           ⇔ -               = Zero

```

Using the effect type for channels, our aim now is to define the semantics of the language as a function that evaluates an expression to a value in a given environment:

```
eval : Expr → Env → CTreec ChanEff Value
```

Because the language now has first-class functions, it no longer suffices to use integers as the value domain for the semantics, so we define a value type that also includes closures, which comprise an unevaluated expression and an environment that captures its free variables:

```
data Value = Num Int | Clo Expr Env
```

In turn, an environment can be represented simply as a list of values,

```
type Env = [ Value ]
```

where the value of the variable with de Bruijn index *i* is given by indexing into the list at position *i* using a lookup function that terminates execution if the variable is not found:

```

lookup :: Int → [ a ] → CTreec e a
lookup _ []      = Zeroc
lookup 0 (v : vs) = return v
lookup i (v : vs) = lookup (i - 1) vs

```

Using these ideas, the semantics for expressions can now be defined as follows:

```

eval :: Expr → Env → CTreec ChanEff Value
eval (Val n)    e = return (Num n)
eval (Add x y) e = do Num n ← eval x e; Num m ← eval y e; return (Num (n + m))

```

```

eval (Var n)    e = lookup n e
eval (Abs x)   e = return (Clo x e)
eval (App x y) e = do Clo x' e' ← eval x e; v ← eval y e; Laterc (eval x' (v : e'))
eval (Send x y) e = do Num c ← eval x e; Num n ← eval y e; send c n; return (Num n)
eval (Receive x) e = do Num c ← eval x e; n ← receive c; return (Num n)
eval (Fork x)  e = do c ← newChan; eval x (Num c : e)  $\overline{\text{!}}_c$  return (Num c)

```

There are a number of points to note about the above semantics. First of all, it uses a syntactic notion of closures, as in our original work on compiler calculation [Bahr and Hutton 2015]. Secondly, to make the definition well-founded, we need to guard the final recursive call in the *App* case with a *Later_c*. All other recursive calls are on structurally smaller expressions. Thirdly, some cases make use of non-exhaustive pattern matching. For example, the results of the recursive calls in the *Add* case must be of the form *Num n* and *Num m*. Here we assume that if pattern matching fails, e.g. by attempting to add closures, then the whole expression evaluates to *Zero_c*. In Haskell, this can be achieved by implementing the *fail* method of the *MonadFail* type class for *CTree_c*:

```

fail :: String → CTreec e a
fail _ = Zeroc

```

And finally, the *Send*, *Receive* and *Fork* cases are defined using the operations from the effect type *ChanEff*. For example, *Fork x* creates a new channel using *newChan*, and then spawns a new concurrent process to evaluate *x*, with the new channel being passed to this process by adding it to the environment, and to the current process by returning it as the result.

We conclude this section by defining the top-level semantics of expressions. This semantics must cover three aspects: i) providing the initial environment; ii) disallowing *SendInt* effects that have not been handled by a concurrent *ReceiveInt* effect and vice versa; and iii) giving the semantics for the *NewChan* effect. The first is achieved by simply providing the empty environment, while the latter two are taken care of by a suitable stateful effect handler:

```

data NoEff a
evaluate :: Expr → CTreec NoEff Value
evaluate x = interpStc 0 hanChan (eval x [])
hanChan :: Chan → ChanEff a → CTreec NoEff (a, Chan)
hanChan c (SendInt _ _) = Zeroc
hanChan c (ReceiveInt _) = Zeroc
hanChan c NewChan      = return (c, c + 1)

```

The effect handler *hanChan* handles the effect from *ChanEff* without using any uninterpreted effects, which is indicated by the empty effect type *NoEff* that provides no operations. *SendInt* and *ReceiveInt* effects are simply handled by *Zero_c*, whereas the *NewChan* effect is handled by using a state of type *Chan*, which is initialised to 0 and incremented every time the effect is used.

8.3 Compiler Specification

Following the approach of Bahr and Hutton [2022], our goal to define a compiler $comp :: Expr \rightarrow Code \rightarrow Code$ that produces code for a stack machine $exec :: Code \rightarrow Conf \rightarrow CTree_c ChanEff Conf$, where *Conf* is the type of configurations for the machine. Because the source language semantics now requires an environment, the configuration type also includes an environment:

```

type Conf = (Stack, Env')

```

However, the machine may require a different form of environment to the source language, so we use a new type Env' for this purpose, defined as a list of machine values of type $Value'$:

type $Env' = [Value']$

To convert between source language and machine values, we assume a conversion function $conv :: Value \rightarrow Value'$, which is lifted to environments by simply mapping over the list of values:

$conv_E :: Env \rightarrow Env'$
 $conv_E = map\ conv$

Similarly to $comp$, $Code$ and $exec$, the definitions for $Value'$ and $conv$ are not given in advance, but will be derived during the compiler calculation. Finally, a stack is initially defined as a list of machine values, with the element type being extended as required during the calculation:

type $Stack = [Elem]$
data $Elem = VAL\ Value'$

Using these definitions, compiler correctness for $comp$ can be specified as follows:

$$\begin{array}{l} \mathbf{do}\ v \leftarrow \mathit{eval}\ x\ e \\ \mathit{exec}\ c\ (VAL\ (conv\ v) : s, conv_E\ e) \end{array} \cong \mathit{exec}\ (comp\ x\ c)\ (s, conv_E\ e)$$

This property has the same form as for the simple language in Section 6, except that the machine now operates on configurations comprising a stack and environment, and we need to take account of the different value and environment types used by the source and target languages. In turn, for the top-level semantics $evaluate$ of expressions, our aim is to define a top-level compilation function $compile :: Expr \rightarrow Code$ and a top-level execution function $execute :: Code \rightarrow Conf \rightarrow CTree_c\ NoEff\ Conf$ that satisfy the following correctness property:

$$\begin{array}{l} \mathbf{do}\ v \leftarrow \mathit{evaluate}\ x \\ \mathit{return}\ ([VAL\ (conv\ v)], []) \end{array} \cong \mathit{execute}\ (compile\ x)\ ([], [])$$

That is, compiling an expression and then executing the resulting code using an empty stack and environment should result in a stack that contains the value of the expression.

8.4 Compiler Calculation

Using the fact that $p \cong q$ iff $p \cong_i q$ for all step counts i , to prove the correctness property for $comp$ it suffices to prove the following by induction on the step count i and expression x :

$$\begin{array}{l} \mathbf{do}\ v \leftarrow \mathit{eval}\ x\ e \\ \mathit{exec}\ c\ (VAL\ (conv\ v) : s, conv_E\ e) \end{array} \cong_i \mathit{exec}\ (comp\ x\ c)\ (s, conv_E\ e)$$

For each case of x , we start on the left-side of the property and seek to transform it into the form $\mathit{exec}\ c' (s, conv_E\ e)$ for some code c' , which then gives a clause $comp\ x\ c = c'$ for the compiler in this case. As previously, the calculation is driven by the desire to transform the term being manipulated into the require form using the induction hypotheses.

The calculation proceeds in a similar manner to the pure lambda calculus [Bahr and Hutton 2022], with the cases for the extra concurrency primitives $Send$, $Receive$ and $Fork$ being similar to the cases for $Print$ and $Fork$ in Section 6. The supplementary material for the articles includes full details of the calculations. One notable difference compared to previous calculations is that the source language now has both externally observable effects and the possibility of getting stuck because of an error, such as trying to add non-numeric values or looking up the value of an unbound variable.

Due to the presence of effects, it may be observable precisely when a computation gets stuck. To illustrate the consequences, consider the calculation for the *Add* case, which proceeds as follows:

$$\begin{aligned}
& \mathbf{do} \ v \leftarrow \mathit{eval} \ (\mathit{Add} \ x \ y) \ e \\
& \quad \mathit{exec} \ c \ (\mathit{VAL} \ (\mathit{conv} \ v) : s, \mathit{conv}_E \ e) \\
\cong_i & \quad \{ \text{definition of } \mathit{eval} \} \\
& \mathbf{do} \ v \leftarrow \mathbf{do} \ \{ \mathit{Num} \ n \leftarrow \mathit{eval} \ x \ e; \mathit{Num} \ m \leftarrow \mathit{eval} \ y \ e; \mathit{return} \ (\mathit{Num} \ (n + m)) \} \\
& \quad \mathit{exec} \ c \ (\mathit{VAL} \ (\mathit{conv} \ v) : s, \mathit{conv}_E \ e) \\
\cong_i & \quad \{ \text{monad laws} \} \\
& \mathbf{do} \ \mathit{Num} \ n \leftarrow \mathit{eval} \ x \ e; \mathit{Num} \ m \leftarrow \mathit{eval} \ y \ e \\
& \quad \mathit{exec} \ c \ (\mathit{VAL} \ (\mathit{conv} \ (\mathit{Num} \ (n + m))) : s, \mathit{conv}_E \ e) \\
\cong_i & \quad \{ \text{define: } \mathit{conv} \ (\mathit{Num} \ n) = \mathit{Num}' \ n \} \\
& \mathbf{do} \ \mathit{Num} \ n \leftarrow \mathit{eval} \ x \ e; \mathit{Num} \ m \leftarrow \mathit{eval} \ y \ e \\
& \quad \mathit{exec} \ c \ (\mathit{VAL} \ (\mathit{Num}' \ (n + m)) : s, \mathit{conv}_E \ e) \\
\cong_i & \quad \left\{ \begin{array}{l} \text{define: } \mathit{exec} \ (\mathit{ADD} \ c) \ (\mathit{VAL} \ (\mathit{Num}' \ m) : \mathit{VAL} \ (\mathit{Num}' \ n) : s, e) = \\ \mathit{exec} \ c \ (\mathit{VAL} \ (\mathit{Num}' \ (n + m)) : s, e) \end{array} \right\} \\
& \mathbf{do} \ \mathit{Num} \ n \leftarrow \mathit{eval} \ x \ e; \mathit{Num} \ m \leftarrow \mathit{eval} \ y \ e \\
& \quad \mathit{exec} \ (\mathit{ADD} \ c) \ (\mathit{VAL} \ (\mathit{Num}' \ m) : \mathit{VAL} \ (\mathit{Num}' \ n) : s, \mathit{conv}_E \ e) \\
\cong_i & \quad \{ \text{define: } \mathit{exec} \ (\mathit{ADD} \ c) \ _ = \mathit{Zero}_c \} \\
& \mathbf{do} \ \mathit{Num} \ n \leftarrow \mathit{eval} \ x \ e; \ v \leftarrow \mathit{eval} \ y \ e \\
& \quad \mathit{exec} \ (\mathit{ADD} \ c) \ (\mathit{VAL} \ (\mathit{conv} \ v) : \mathit{VAL} \ (\mathit{Num}' \ n) : s, \mathit{conv}_E \ e) \\
\cong_i & \quad \{ \text{induction hypothesis for } y \} \\
& \mathbf{do} \ \mathit{Num} \ n \leftarrow \mathit{eval} \ x \ e \\
& \quad \mathit{exec} \ (\mathit{comp} \ y \ (\mathit{ADD} \ c)) \ (\mathit{VAL} \ (\mathit{Num}' \ n) : s, \mathit{conv}_E \ e) \\
\cong_i & \quad \left\{ \begin{array}{l} \text{define: } \mathit{exec} \ (\mathit{ISNUM} \ c) \ (\mathit{VAL} \ (\mathit{Num}' \ n) : s, e) = \mathit{exec} \ c \ (\mathit{VAL} \ (\mathit{Num}' \ n) : s, e) \\ \mathit{exec} \ (\mathit{ISNUM} \ c) \ _ = \mathit{Zero}_c \end{array} \right\} \\
& \mathbf{do} \ v \leftarrow \mathit{eval} \ x \ e \\
& \quad \mathit{exec} \ (\mathit{ISNUM} \ (\mathit{comp} \ y \ (\mathit{ADD} \ c))) \ (\mathit{VAL} \ (\mathit{conv} \ v) : s, \mathit{conv}_E \ e) \\
\cong_i & \quad \{ \text{induction hypothesis for } x \} \\
& \quad \mathit{exec} \ (\mathit{comp} \ x \ (\mathit{ISNUM} \ (\mathit{comp} \ y \ (\mathit{ADD} \ c)))) \ (s, \mathit{conv}_E \ e)
\end{aligned}$$

In the above calculation, in addition to the *ADD* instruction that adds together two numbers on top of the stack, we introduce an *ISNUM* instruction that checks if the top of the stack is a numeric value. The introduction of this latter instruction is driven by the need to manipulate the term so that we can apply the induction hypothesis for *x*. Intuitively, including *ISNUM* after the compiled code for *x* is required because the generated code must have the same semantics as the source expression *Add x y* and hence fail as early as possible. Otherwise, the generated code would exhibit the computational effects of the expression *y* even though the source expression *Add x y* would not. Importantly, we did not have to make this observation, as the need for an instruction with the semantics of *ISNUM* falls out naturally as part of the calculation.

In turn, we can calculate the top-level function *compile* from its correctness property. In this case we don't need step-counting or induction, as a simple equational reasoning suffices:

$$\begin{aligned}
& \mathbf{do} \ v \leftarrow \mathit{evaluate} \ x; \ \mathit{return} \ ([\mathit{VAL} \ (\mathit{conv} \ v)], []) \\
\cong & \quad \{ \text{definition of } \mathit{evaluate} \} \\
& \mathbf{do} \ v \leftarrow \mathit{interpSt}_c \ 0 \ \mathit{hanChan} \ (\mathit{eval} \ x \ []); \ \mathit{return} \ ([\mathit{VAL} \ (\mathit{conv} \ v)], []) \\
\cong & \quad \{ \mathit{interpSt}_c\text{-fmap law} \}
\end{aligned}$$

$$\begin{aligned}
& \text{interpSt}_c \ 0 \ \text{hanChan} \ (\mathbf{do} \ v \leftarrow \text{eval} \ x \ []; \ \text{return} \ ([\text{VAL} \ (\text{conv} \ v)], \ [])) \\
& \cong \ \{ \text{define: } \text{exec} \ \text{HALT} \ (s, e) = \text{return} \ (s, e) \} \\
& \text{interpSt}_c \ 0 \ \text{hanChan} \ (\mathbf{do} \ v \leftarrow \text{eval} \ x \ []; \ \text{exec} \ \text{HALT} \ ([\text{VAL} \ (\text{conv} \ v)], \ [])) \\
& \cong \ \{ \text{compiler correctness for } \text{comp} \} \\
& \text{interpSt}_c \ 0 \ \text{hanChan} \ (\text{exec} \ (\text{comp} \ x \ \text{HALT}) \ ([], \ [])) \\
& \cong \ \{ \text{define: } \text{execute} \ c \ (s, e) = \text{interpSt}_c \ 0 \ \text{hanChan} \ (\text{exec} \ c \ (s, e)) \} \\
& \text{execute} \ (\text{comp} \ x \ \text{HALT}) \ ([], \ [])
\end{aligned}$$

In summary, we have calculated the definitions in Figure 4.

8.5 Reflection

We conclude this section with some reflective remarks. First of all, note that the above calculation is based on a *strong* notion of bisimilarity in which silent transitions τ introduced by *Later* must be preserved. Strong bisimilarity supports the equational style of reasoning that underpins our approach to compiler calculation for non-terminating languages [Bahr and Hutton 2022].

And secondly, the configuration of the virtual machine for our concurrent lambda calculus is similar to that of the virtual machine calculated in Section 6, with the difference that the machine now also has a global state ch that denotes the next available channel, and each thread has an environment in addition to a stack. An execution sequence of the machine has the form

$$\text{execute} \ c \ (s, e) \xrightarrow{\tau} ch_1 \otimes p_1 \xrightarrow{\tau} ch_2 \otimes p_2 \xrightarrow{\tau} \dots$$

where $ch \otimes p$ abbreviates $\text{interpSt}_c \ ch \ \text{hanChan} \ p$, and each p_i is bisimilar to an expression:

$$\text{exec} \ c_1 \ (s_1, e_1) \ \vec{\parallel}_c \ \text{exec} \ c_2 \ (s_2, e_2) \ \vec{\parallel}_c \ \dots \ \vec{\parallel}_c \ \text{exec} \ c_{n+1} \ (s_{n+1}, e_{n+1})$$

Note that because all effects are handled by interpSt_c , the only externally observable effects are silent τ transitions. However, in a similar manner to Section 6, we could have included a *Print* primitive in the lambda calculus, which the effect handler would have left uninterpreted and which therefore would appear as *PrintInt* transitions in addition to the τ transitions.

9 RELATED WORK

The codensity construction is a common trick in the functional programming literature, which is typically used to improve efficiency; for example, see Hinze [2012] for an overview. Curiously, the form of codensity choice tree without *Later*_c that was presented in Section 4 is very close to the definition of an efficient parallel parser by Claessen [2004]. However, because it implements a parser, Claessen's definition only supports one effect, namely reading a symbol, and to improve efficiency the *Now* and \oplus constructors are fused together.

As noted in Section 4, the semantics of Haskell's *forkIO* primitive [Jones et al. 1996] provided the initial inspiration for our continuation-passing style semantics, which in turn motivated the use of the codensity construction. In the remainder of this section, we review related work on compiler verification, and compare the original notion of choice trees with our version.

9.1 Compiler Verification

Wand [1982a] pioneered a compiler verification technique that uses a suitably expressive lambda calculus as a common language in which to define both the source and target language of the compiler. This idea has its origins in Reynolds' seminal work on definitional interpreters [Reynolds 1972], and has proved fruitful for deriving correct-by-construction compilers [Ager et al. 2003a,b; Bahr and Hutton 2015, 2020; Gibbons 2021; Wand 1982a,b]. While this line of work is limited to sequential programming languages, Wand [1995] later demonstrated how to prove compiler

compiler [Leroy 2006, 2009]. This work inspired many further projects on compiler verification [Paterson and Ahmed 2019], including Chlipala’s [2010] verified compiler, CakeML [Kumar et al. 2014] and DeepSpec [2023]. CompCert and its correctness proof have since been extended to support concurrency [Ševčík et al. 2013]. Rather than an intermediate language like HOCC or choice trees, Ševčík et al. specify the semantics of source and target languages directly using a small-step semantics. A key challenge for the verification of realistic compilers for concurrent languages is devising a suitable memory model [Kang et al. 2017].

9.2 Choice Trees

Chappe et al. [2023] introduced choice trees as an extension of interaction trees [Xia et al. 2019] with a distinguished effect for non-determinism, in order to obtain important equational reasoning principles such the idempotent, commutative monoid laws for non-deterministic choice. In turn, interaction trees extend the freer monad construction of Kiselyov and Ishii [2015] with an explicit non-termination effect, in the style of Capretta’s [2005] delay monad. Again, the purpose of treating non-termination differently from other effects is to obtain reasoning principles for non-termination, such as coinduction or step-indexed reasoning. Rivas et al. [2018] give a systematic account of extending computational effects with non-determinism, and in particular show how to construct a free non-determinism monad as a free near-semiring. Because the freer monad construction is the composition of the left Kan extension followed by the free monad construction [Kiselyov and Ishii 2015], we can also think of choice trees as the free near-semiring applied to the left Kan extension of an effect signature extended with non-termination in the style of Capretta [2005].

Choice trees [Chappe et al. 2023] offer an alternative language to serve as the common semantic domain for reasoning, and Chappe et al. have demonstrated how this language can be used to model concurrent languages. Unlike Wand’s higher-order process calculus HOCC, choice trees don’t explicitly feature parallelism or higher-order constructs, but these features can be encoded [Chappe et al. 2023; Danielsson 2012; Xia et al. 2019]. Indeed, the concurrent lambda calculus from Section 8 is very similar to HOCC and the latter can be given a semantics in terms of (codensity) choice trees. The only notable differences are that HOCC also allows sending and receiving closed lambda terms, and that communication is via thread identifiers rather than channel identifiers.

The idea to use the effect handler operation *interpSt* to model the restriction of effects to a local context, and the parallelism operator \parallel for choice trees, are both due to [Chappe et al. 2023]. Our addition of concurrent effect handlers is a natural generalisation. However, our syntax and semantics for choice trees is different from Chappe et al.’s in crucial ways that enable the equational reasoning that underlies our methodology. We discuss these differences in turn below.

Syntax. First of all, there are two superficial syntactic differences: instead of a binary choice operator \oplus with a unit *Zero*, Chappe et al. [2023] use a choice operator *brS* of arbitrary finite arity, and instead of *Later*, they have an operator *brD* that is the composition of *brS* and *Later*. However, *brS* and *brD* are interdefinable with \oplus , *Zero* and *Later*. Adjusting for these differences and using our notation, Chappe et al.’s definition of choice trees is equivalent to the following:

codata $CTree' e a$ where

Now $:: a \rightarrow CTree' e a$

Later $:: CTree' e a \rightarrow CTree' e a$

(\oplus) $:: CTree' e a \rightarrow CTree' e a \rightarrow CTree' e a$

Zero $:: CTree' e a$

Eff $:: e b \rightarrow (b \rightarrow CTree' e a) \rightarrow CTree' e a$

This definition looks almost identical to ours, with one key difference: instead of an inductive definition with a nested coinductive definition for *Later*, the entire type is defined coinductively. As a result, a $CTree'$ might be infinite even though it contains no *Later*, which is different from $CTree$, where all infinite behaviour must be due to *Later*. This means that while on the surface $CTree'$ only permits finite non-determinism, it can in fact encode infinite non-deterministic choice:

$$\begin{aligned} \mathit{infChoice} &:: (\mathit{Int} \rightarrow \mathit{CTree}' e a) \rightarrow \mathit{CTree}' e a \\ \mathit{infChoice} \ ps &= \mathit{ps} \ 0 \oplus \mathit{infChoice} (\lambda n \rightarrow \mathit{ps} (n + 1)) \end{aligned}$$

This definition doesn't work for $CTree$ because \oplus is an inductive constructor for $CTree$ and hence $\mathit{infChoice}$ would not be well-founded. Having infinite non-deterministic choice for $CTree'$ makes the notion of step-indexed bisimilarity \cong_i that underpins our methodology unsound, in the sense that $p \cong_i q$ for all i no longer implies $p \cong q$. For example, consider the choice trees $p, q :: \mathit{CTree}' e a$ defined by $p = \mathit{Later} \ p$ and $q = \mathit{infChoice} \ qs$, where $qs \ 0 = \mathit{Zero}$ and $qs \ n = \mathit{Later} (qs (n - 1))$. Then we have $p \oplus q \cong_i q$ for all i , but not $p \oplus q \cong q$. Failure of the latter can be seen by the fact that $p \oplus q$ has non-terminating behaviour whereas q does not, while the former can be proved by first showing that $p \cong_i qs \ i$ by induction on i and then using this result to show $p \oplus q \cong_i q$.

Semantics. As we have observed in Section 2, every effectful transition is always immediately followed by an input transition that consumes the resulting value:

$$\mathit{Eff} \ o \ c \xRightarrow{\uparrow o} c \xRightarrow{\downarrow i} c \ i$$

In contrast, [Chappe et al. \[2023\]](#) combine the two transitions into one:

$$\mathit{Eff} \ o \ c \xRightarrow{\uparrow o \downarrow i} c \ i$$

This has the benefit of avoiding the need for two kinds of states in the transition semantics, namely choice trees and continuations. However, the resulting notion of bisimilarity is too coarse, which in turn causes the congruence property for *interp* to fail. For example, suppose we have an effect that flips a coin, and returns the result as a value of type *Bool*:

data *CoinEff* *b* **where**

Flip :: *CoinEff* *Bool*

Then with the semantics of [Chappe et al.](#) the following two choice trees are bisimilar

ignore = $\mathit{Eff} \ \mathit{Flip} (\lambda b \rightarrow \mathit{return} \ \mathit{True}) \oplus \mathit{Eff} \ \mathit{Flip} (\lambda b \rightarrow \mathit{return} \ \mathit{False})$

negate = $\mathit{Eff} \ \mathit{Flip} (\lambda b \rightarrow \mathit{return} \ b) \oplus \mathit{Eff} \ \mathit{Flip} (\lambda b \rightarrow \mathit{return} \ (\neg b))$

Both examples start with a non-deterministic choice, and each choice component starts with a coin flip. However, the two components of *ignore* both ignore the result of the coin flip and return the fixed results *True* and *False*, respectively. On the other hand, the two components of *negate* return the result of the coin flip unchanged and negated, respectively. The two examples have different observable behaviours and hence should not be considered bisimilar, and indeed they are not bisimilar with the semantics for choice trees that keeps $\uparrow o$ and $\downarrow i$ transitions separate.

However, with the transition semantics in the combined style of [Chappe et al. \[2023\]](#), the two examples are bisimilar because they have the same four possible transitions:

$$\mathit{ignore} \xRightarrow{\uparrow \mathit{Flip} \downarrow b} \mathit{return} \ b' \quad \text{and} \quad \mathit{negate} \xRightarrow{\uparrow \mathit{Flip} \downarrow b} \mathit{return} \ b' \quad \text{for all } b, b' :: \mathit{Bool}$$

As a result of this coarser notion of bisimilarity, the effect handler operations *interp* and *interpSt* do not satisfy the congruence property, which does hold for our notion of bisimilarity. Instead, [Chappe et al. \[2023\]](#) prove congruence for *interp* and *interpSt* only for effect handlers *han* with

$han\ e \cong\ return\ v$ or $han\ e \cong\ Eff\ f\ return$, which excludes the effect handler we used in Section 8. Indeed, if we define $han\ Flip = Later\ (Eff\ Flip\ return)$, then congruence fails using Chappe et al.’s semantics, because $interp\ han\ ignore \not\equiv interp\ han\ negate$ even though $ignore \cong negate$.

10 CONCLUSION AND FURTHER WORK

In recent years, interaction trees and choice trees have proved to be a flexible, expressive and modular approach to mechanising programming language meta-theory [Chappe et al. 2023; Hur et al. 2020; Xia et al. 2019; Yoon et al. 2022]. In this article, we showed how the notion of choice trees can also be adapted to admit an equational reasoning style that supports the derivation of correct-by-construction compilers for concurrent languages. In particular, in combination with subtle changes in the syntax and semantics of choice trees, the use of a codensity construction allows the semantics of concurrent languages to be concisely captured in a monadic style, and supports a high-level, algebraic approach to transforming the resulting semantics into compilers.

This article builds upon our recent work [Bahr and Hutton 2022] on compiler calculation for non-terminating languages in a number of aspects. First of all, it shows how the standard notion of choice trees can be adapted to support compiler calculation. Secondly, it shows how the resulting form of choice trees can be used to extend our methodology to handle concurrency and general effects. And finally, it demonstrates the practical application of the extended methodology by calculating a compiler for a concurrent lambda calculus with channel-based communication.

In terms of further work, the use of (codensity) choice trees with explicit effect types opens up the opportunity for deriving multi-stage compilers, where each stage compiles away some effects and leaves others untouched. The concurrent lambda calculus compiler derived in Section 8 provides an example of this, where the print effect remains uninterpreted in the type $Ctree_c\ PrintEff$ of the codensity choice tree that is used for both the *Expr* and *Code* languages. Thus one could see this compiler as the first stage of a multi-stage compiler. The semantics of *Code* could be refined by an effect handler that handles the print effect so that a second compiler calculation could produce the second stage compiler that translates *Code* to a lower-level language.

Finally, we note that the use of explicit step-indexing in our methodology could be replaced by the use of guarded type theory [Bizjak et al. 2016], e.g. using Guarded Cubical Agda [Kristensen et al. 2022; Veltri and Vezzosi 2023], which simplifies the formalisation and also has the potential to enable some extensions to the results, e.g. to support higher-order effects.

ACKNOWLEDGEMENTS

We would like to thank the reviewers for many useful comments and suggestions.

REFERENCES

- Mads Sig Ager, Dariusz Biernacki, Olivier Danvy, and Jan Midtgaard. 2003a. A Functional Correspondence Between Evaluators and Abstract Machines. In *Proceedings of the International Conference on Principles and Practice of Declarative Programming*.
- Mads Sig Ager, Dariusz Biernacki, Olivier Danvy, and Jan Midtgaard. 2003b. *From Interpreter to Compiler and Virtual Machine: A Functional Derivation*. Technical Report RS-03-14. BRICS, Department of Computer Science, University of Aarhus.
- Patrick Bahr and Graham Hutton. 2015. Calculating Correct Compilers. *Journal of Functional Programming* 25 (2015).
- Patrick Bahr and Graham Hutton. 2020. Calculating Correct Compilers II: Return of the Register Machines. *Journal of Functional Programming* 30 (2020).
- Patrick Bahr and Graham Hutton. 2022. Monadic Compiler Calculation. *Proceedings of the ACM on Programming Languages* 6, ICFP (2022).
- Patrick Bahr and Graham Hutton. 2023. Supplementary Material for “Calculating Compilers for Concurrency”. <https://doi.org/10.5281/zenodo.8124116>
- Aleš Bizjak, Hans Grathwohl, Randal Clouston, Rasmus Møgelberg, and Lars Birkedal. 2016. Guarded Dependent Type Theory with Coinductive Types. In *Foundations of Software Science and Computation Structures*.

- Venanzio Capretta. 2005. General Recursion via Coinductive Types. *Logical Methods in Computer Science* 1, 2 (2005).
- Nicolas Chappe, Paul He, Ludovic Henrio, Yannick Zakowski, and Steve Zdancewic. 2023. Choice Trees: Representing Nondeterministic, Recursive, and Impure Programs in Coq. *Proceedings of the ACM on Programming Languages* 4, POPL (2023).
- Adam Chlipala. 2010. A Verified Compiler for an Impure Functional Language. In *Proceedings of the Symposium on Principles of Programming Languages*.
- Koen Claessen. 2004. Functional Pearl: Parallel Parsing Processes. *Journal of Functional Programming* 14, 6 (2004).
- Nils Anders Danielsson. 2012. Operational Semantics Using the Partiality Monad. In *Proceedings of the International Conference on Functional Programming*.
- DeepSpec 2023. The Science of Deep Specification. (2023). <https://deepspec.org/>.
- Jeremy Gibbons. 2021. Continuation-Passing Style, Defunctionalization, Accumulations, and Associativity. *The Art, Science, and Engineering of Programming* 6, 2 (2021).
- Ralf Hinze. 2012. Kan Extensions for Program Optimisation Or: Art and Dan Explain an Old Trick. In *Proceedings of the International Conference on Mathematics of Program Construction*.
- Chung-kil Hur, Paul He, Yannick Zakowski, and Steve Zdancewic. 2020. An Equational Theory for Weak Bisimulation via Generalized Parameterized Coinduction. In *Proceedings of the International Conference on Certified Programs and Proofs*.
- Simon Peyton Jones, Andrew Gordon, and Sigbjorn Finne. 1996. Concurrent Haskell. In *Proceedings of the Symposium on Principles of Programming Languages*.
- Jeehoon Kang, Chung-Kil Hur, Ori Lahav, Viktor Vafeiadis, and Derek Dreyer. 2017. A Promising Semantics for Relaxed-Memory Concurrency. In *Proceedings of the Symposium on Principles of Programming Languages*.
- Oleg Kiselyov and Hiromi Ishii. 2015. Freer Monads, More Extensible Effects. In *Proceedings of the 2015 ACM SIGPLAN Symposium on Haskell*.
- Magnus Kristensen, Rasmus Mogelberg, and Andrea Vezzosi. 2022. Greatest HITS: Higher Inductive Types in Coinductive Definitions via Induction Under Clocks. In *Proceedings of the Symposium on Logic in Computer Science*.
- Ramana Kumar, Magnus O. Myreen, Michael Norrish, and Scott Owens. 2014. CakeML: A Verified Implementation of ML. In *Proceedings of the Symposium on Principles of Programming Languages*.
- Xavier Leroy. 2006. Formal Certification of a Compiler Back-End or: Programming a Compiler with a Proof Assistant. In *Proceedings of the Symposium on Principles of Programming Languages*.
- Xavier Leroy. 2009. Formal Verification of a Realistic Compiler. *Commun. ACM* 52, 7 (2009).
- Daniel Patterson and Amal Ahmed. 2019. The Next 700 Compiler Correctness Theorems. *Proceedings of the ACM on Programming Languages* 3, ICFP (2019).
- John C Reynolds. 1972. Definitional Interpreters for Higher-Order Programming Languages. In *Proceedings of the ACM Annual Conference*.
- Exequiel Rivas, Mauro Jaskelioff, and Tom Schrijvers. 2018. A Unified View of Monadic and Applicative Non-Determinism. *Science of Computer Programming* 152 (Jan. 2018).
- Niccolò Veltri and Andrea Vezzosi. 2023. Formalizing CCS and π -calculus in Guarded Cubical Agda. *Journal of Logical and Algebraic Methods in Programming* 131 (2023).
- Janis Voigtländer. 2008. Asymptotic Improvement of Computations over Free Monads. In *Proceedings of the International Conference on Mathematics of Program Construction*.
- Mitchell Wand. 1982a. Deriving Target Code as a Representation of Continuation Semantics. *Transactions on Programming Languages and Systems* 4, 3 (1982).
- Mitchell Wand. 1982b. Semantics-Directed Machine Architecture. In *Proceedings of the Symposium on Principles of Programming Languages*.
- Mitchell Wand. 1995. Compiler Correctness for Parallel Languages. In *Proceedings of International Conference on Functional Programming Languages and Computer Architecture*.
- Li-yao Xia, Yannick Zakowski, Paul He, Chung-Kil Hur, Gregory Malecha, Benjamin C. Pierce, and Steve Zdancewic. 2019. Interaction Trees: Representing Recursive and Impure Programs in Coq. *Proceedings of the ACM on Programming Languages* 4, POPL (2019).
- Irene Yoon, Yannick Zakowski, and Steve Zdancewic. 2022. Formal Reasoning About Layered Monadic Interpreters. *Proceedings of the ACM on Programming Languages* 6, ICFP (2022).
- Jaroslav Ševčík, Viktor Vafeiadis, Francesco Zappa Nardelli, Suresh Jagannathan, and Peter Sewell. 2013. CompCertTSO: A Verified Compiler for Relaxed-Memory Concurrency. *J. ACM* 60, 3 (2013).

Received 2023-03-01; accepted 2023-06-27