

Beyond Trees: Calculating Graph-Based Compilers

PATRICK BAHR, IT University of Copenhagen, Denmark

GRAHAM HUTTON, University of Nottingham, UK

Bahr and Hutton recently developed an approach to compiler calculation that allows a wide range of compilers to be derived from specifications of their correctness. However, a limitation of the approach is that it results in compilers that produce tree-structured code. By contrast, realistic compilers produce code that is essentially graph-structured, where the edges in the graph represent jumps that transfer the flow of control to other locations in the code. In this article, we show how their approach can naturally be adapted to calculate compilers that produce graph-structured code, without changing the underlying calculational methodology, by using a higher-order abstract syntax representation of graphs.

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, graphs, higher-order abstract syntax

1 INTRODUCTION

The aim of *program calculation* [Backhouse 2003] is to develop programs whose correctness is guaranteed by the manner in which they are constructed. Typically, the process starts with a formal specification for the desired behaviour of a program, from which we then seek to derive an implementation that meets this specification using correctness-preserving reasoning.

In recent work, Bahr and Hutton have developed an approach to calculating *compilers*, which translate high-level programs into lower-level code that machines can execute. Their approach begins with an evaluation function that captures the semantics of the source language. For example, for a simple expression language an evaluation function (*eval*) might map an expression to its value. In contrast, a compilation function (*compile*) maps an expression into code to be executed by a suitable machine (*exec*). Compiler correctness is then expressed by a simple equation,

$$exec \circ compile = eval$$

which expresses that compilation followed by execution gives the same result as evaluation. Bahr and Hutton’s approach shows how to derive the function *compile* by solving such a correctness equation, just like we solve any equation in mathematics. In practice, the equation is usually extended with further concepts such as a stack or an environment, depending on the nature of the source and target languages, but the above captures the basic idea.

The methodology initially targeted stack-based machines [Bahr and Hutton 2015], and was used to calculate compilers for a wide range of language features and their combinations, including exceptions, state, lambda calculus, loops, non-determinism and interrupts. The approach was later adapted to register-based machines [Bahr and Hutton 2020], together with well-typed [Pickard and Hutton 2021], non-terminating [Bahr and Hutton 2022] and concurrent [Bahr and Hutton 2023] languages. It also allowed McCarthy and Painter’s seminal compiler verification [McCarthy and Painter 1967] to be reworked as a compiler calculation [Hutton and Bahr 2017].

However, a limitation of the above approach is that it results in compilers that produce *tree-structured* code. For example, when compiling a conditional expression *if cond then b1 else b2*, the code for the two branches become parameters for an instruction in the target language. In contrast, a traditional compiler would produce linear code in this case, by concatenating the code for the two branches, and inserting jumps to ensure the appropriate branch is executed. In this manner, we can think of traditional compilers as producing *graph-structured* code, where the edges

in the graph represent jumps that transfer the flow of control. Viewing code as graphs also allows a number of other control flow concepts to be naturally considered, such as sharing (multiple jumps to a single piece of code), and looping (backward jumps to repeat a piece of code.)

Unfortunately, graphs and program calculation are not usually regarded as being natural bedfellows. In particular, programming with graphs often requires the generation and management of some form of labels or names that are used to represent the edges in graphs, which can significantly complicate the progress of reasoning about the resulting programs. For example, the verification of a compiler for a simple language with exceptions in [Hutton and Wright 2004] requires a considerable number of lemmas regarding how jump labels are managed.

In this article, we avoid such issues by using a representation of graphs that does not require explicit labels. In particular, we utilise the approach of Oliveira and Cook [2012], in which graphs are represented using *parametric higher-order abstract syntax* [Chlipala 2008]. Using this approach, we can leverage the binding and naming mechanism of the host language to avoid having to explicitly deal with these concepts. As we shall see, this approach allows the compiler calculation methodology developed by Bahr and Hutton to be naturally adapted to compilers that produce graph-structured code, without changing the underlying calculational methodology.

Our approach is generally applicable and can be combined with other methods that derive tree-based compilers, such as Meijer [1992], Wand [1995], and Ager et al. [2003]. It also removes a restriction of previous work, by allowing the calculation of compilers that produce cyclic code, which cannot be represented with finite tree-structured code. In this manner, the graph-based approach that is presented does not invalidate previous work that produces tree-structured code, but rather complements it and demonstrates its relevance for practical compilers.

We introduce and illustrate our approach using three examples of increasing complexity, starting with a simple expression language with conditionals (Sections 2 and 3), then looping (Section 4), and finally state (Section 5). All our programs and calculations are written in Haskell notation, but for reasoning purposes we assume the language is total. All the calculations have been mechanically verified in Agda [Norell 2007], and are available as supplementary material.

2 CALCULATING A TREE-BASED COMPILER

To illustrate the issue with Bahr and Hutton’s compiler calculation methodology, we begin in this section by calculating a compiler for a simple expression language comprising integers, addition and conditionals, whose syntax and semantics is defined as follows:

```

data Expr = Val Int | Add Expr Expr | If Expr Expr Expr
eval :: Expr → Int
eval (Val n)    = n
eval (Add x y)  = eval x + eval y
eval (If x y z) = if eval x == 0 then eval z else eval y

```

For the purpose of conditionals, we follow the standard C-like convention where the logical value false is represented by the integer zero, and true is represented by any non-zero integer.

2.1 Compiler specification

Given the above definitions, we now aim to calculate a compilation function $compile :: Expr \rightarrow Code$ that translates expressions of the source language into code for our target language. We assume that the target language is a stack machine, whose semantics is given by a function $exec :: Code \rightarrow Stack \rightarrow Stack$, where stacks are simply represented as lists of integers:

```

type Stack = [Int]

```

Neither the target language *Code* nor its semantics *exec* are given up front, but will instead fall out of the calculation of the compiler. To simplify the calculation [Bahr and Hutton 2015], we generalise the compilation function *compile* to a function $comp :: Expr \rightarrow Code \rightarrow Code$ that takes an additional code argument to be appended to the code produced for the expression. Correctness of the generalised compiler can then be expressed by the following equation:

$$exec\ c\ (eval\ x\ :\ s) = exec\ (comp\ x\ c)\ s \quad (1)$$

That is, executing the result of compiling an expression followed by additional code gives the same result as executing the additional code with the value of the expression on top of the stack.

2.2 Compiler calculation

We now calculate the compiler by proving equation (1) by induction on the expression *x*. For each case, we aim to transform the left side of the equation, $exec\ c\ (eval\ x\ :\ s)$, into the form $exec\ c'\ s$ for some code *c'*. We can then define $comp\ x\ c = c'$ for that case, which by construction satisfies the correctness equation. The base case, $x = Val\ n$, proceeds as follows:

$$\begin{aligned} & exec\ c\ (eval\ (Val\ n)\ :\ s) \\ = & \{ \text{definition of } eval \} \\ & exec\ c\ (n\ :\ s) \\ = & \{ \text{define: } exec\ (PUSH\ n\ c)\ s = exec\ c\ (n\ :\ s) \} \\ & exec\ (PUSH\ n\ c)\ s \end{aligned}$$

To obtain a term of the required form, in the final step above we introduced a new constructor $PUSH :: Int \rightarrow Code \rightarrow Code$ for the *Code* type, along with a clause of the *exec* function, which defines the semantics of *PUSH*. The final term of the calculation is then of the form $exec\ c'\ s$ with $c' = PUSH\ n\ c$, which gives the first clause for the definition of the compiler:

$$comp\ (Val\ n)\ c = PUSH\ n\ c$$

The inductive case for addition proceeds in a similar manner, in which we introduce a new code constructor *ADD* that adds the top two values on the stack, this time motivated by the desire to transform the term into a form to which the induction hypotheses can be applied:

$$\begin{aligned} & exec\ c\ (eval\ (Add\ x\ y)\ :\ s) \\ = & \{ \text{definition of } eval \} \\ & exec\ c\ (eval\ x\ +\ eval\ y\ :\ s) \\ = & \{ \text{define: } exec\ (ADD\ c)\ (n\ :\ m\ :\ s) = exec\ c\ (m\ +\ n\ :\ s) \} \\ & exec\ (ADD\ c)\ (eval\ y\ :\ eval\ x\ :\ s) \\ = & \{ \text{induction hypothesis for } y \} \\ & exec\ (comp\ y\ (ADD\ c))\ (eval\ x\ :\ s) \\ = & \{ \text{induction hypothesis for } x \} \\ & exec\ (comp\ x\ (comp\ y\ (ADD\ c)))\ s \end{aligned}$$

The final term is now of the required form, which gives the next clause for the compiler:

$$comp\ (Add\ x\ y)\ c = comp\ x\ (comp\ y\ (ADD\ c))$$

Finally, the case for conditionals proceeds as follows, in which we use the fact that in a total language we can freely distribute functions over conditional expressions:

$$\begin{aligned} & exec\ c\ (eval\ (If\ x\ y\ z)\ :\ s) \\ = & \{ \text{definition of } eval \} \end{aligned}$$

$$\begin{aligned}
& \text{exec } c \text{ ((if eval } x == 0 \text{ then eval } z \text{ else eval } y) : s) \\
= & \quad \{ \text{distributivity} \} \\
& \text{exec } c \text{ (if eval } x == 0 \text{ then eval } z : s \text{ else eval } y : s) \\
= & \quad \{ \text{distributivity} \} \\
& \text{if eval } x == 0 \text{ then exec } c \text{ (eval } z : s) \text{ else exec } c \text{ (eval } y : s) \\
= & \quad \{ \text{induction hypotheses for } y \text{ and } z \} \\
& \text{if eval } x == 0 \text{ then exec (comp } z \text{ c) } s \text{ else exec (comp } y \text{ c) } s \\
= & \quad \{ \text{define: exec (JPZ } c' \text{ c) (n : s) = if n == 0 then exec } c' \text{ s else exec } c \text{ s} \} \\
& \text{exec (JPZ (comp } z \text{ c) (comp } y \text{ c)) (eval } x : s) \\
= & \quad \{ \text{induction hypothesis for } x \} \\
& \text{exec (comp } x \text{ (JPZ (comp } z \text{ c) (comp } y \text{ c))) } s
\end{aligned}$$

To apply the induction hypothesis for the expression x , in the penultimate step above we introduced a new code constructor JPZ that jumps to alternative code if the top of the stack is zero. The final term of the calculation then gives the final clause for the compiler:

$$\text{comp (If } x \text{ y } z) \text{ c} = \text{comp } x \text{ (JPZ (comp } z \text{ c) (comp } y \text{ c))}$$

We conclude this section by returning to the top-level compilation function $\text{compile} :: \text{Expr} \rightarrow \text{Code}$, whose correctness can be expressed by the following equation:

$$\text{eval } x : s = \text{exec (compile } x) \text{ s} \quad (2)$$

In a similar manner to equation (1) for comp , we aim to transform the left side of this equation, $\text{eval } x : s$, into the form $\text{exec } c \text{ s}$ for some code c , so that we can then define $\text{compile } x = c$. In this case we don't need induction, as introducing a new code constructor HALT that returns the current stack allows the correctness equation for comp to be used to obtain the required form:

$$\begin{aligned}
& \text{eval } x : s \\
= & \quad \{ \text{define: exec HALT } s = s \} \\
& \text{exec HALT (eval } x : s) \\
= & \quad \{ \text{correctness of comp - (1)} \} \\
& \text{exec (comp } x \text{ HALT) } s
\end{aligned}$$

In summary, we have calculated the definitions shown in Figure 1. The definition of exec also includes the catch-all clause $\text{exec } _ = []$, which we have added to make the definition total. The choice for this catch-all clause is not important as it plays no role in the calculation.

3 CALCULATING A GRAPH-BASED COMPILER

Two characteristics of the code produced by the compiler calculated in the previous section make it unrealistic in practice. First of all, code for a realistic machine is linear, i.e. each instruction is followed by at most one other instruction to be executed next. This is true for the HALT , PUSH and ADD instructions, but not the JPZ instruction, which is followed by two possible pieces of code to be executed next, one for each branch of a conditional expression:

$$\text{JPZ} :: \text{Code} \rightarrow \text{Code} \rightarrow \text{Code}$$

In this manner, the compiler produces *tree-structured* code. In a realistic instruction set, we would expect the JPZ instruction to have type $\text{Loc} \rightarrow \text{Code} \rightarrow \text{Code}$, i.e. the first argument is not code itself but rather a location that points to a piece of code, to ensure that code is linear.

Secondly, when compiling a conditional expression, the compilation function comp duplicates code by passing the additional code c to two recursive calls of the compiler:

```

data Code = HALT | PUSH Int Code | ADD Code | JPZ Code Code
compile :: Expr → Code
compile x = comp x HALT
comp :: Expr → Code → Code
comp (Val n)    c = PUSH n c
comp (Add x y)  c = comp x (comp y (ADD c))
comp (If x y z) c = comp x (JPZ (comp z c) (comp y c))
exec :: Code → Stack → Stack
exec HALT      s      = s
exec (PUSH n c) s      = exec c (n : s)
exec (ADD c)   (n : m : s) = exec c (m + n : s)
exec (JPZ c' c) (n : s)   = if n == 0 then exec c' s else exec c s
exec _        _          = []

```

Fig. 1. Tree-based compiler for conditional expressions.

```

comp (If x y z) c = comp x (JPZ (comp z c) (comp y c))

```

In a realistic compiler, we would instead expect to find an instruction $\text{JMP} :: \text{Loc} \rightarrow \text{Code}$ that allows us to jump to the same piece of code multiple times, rather than duplicating code.

3.1 Graph-structured code

In light of the above observations, we could define a graph-based version of the code type,

```

data Codeg = HALTg | PUSHg Int Codeg | ADDg Codeg | JPZg Loc Codeg | JMPg Loc

```

where jump locations are simply integer addresses within the code, i.e. the first instruction is located at address zero, the second at address one, and so on:

```

type Loc = Int

```

However, the use of explicit jump locations complicates the compiler implementation, and makes reasoning about it difficult. Address calculation is delicate on its own, and it would be beneficial to separate this from the process of compiler calculation. To this end, we treat jump locations *abstractly* by using *parametric higher-order abstract syntax* [Chlipala 2008] to provide symbolic locations, which we call *labels*. We achieve this by abstracting the *Loc* type to a type variable l , and adding a new instruction $\text{LAB}_{\vec{g}}$ that brings a fresh label into scope using a function parameter:

```

data Codeg l = HALTg | PUSHg Int (Codeg l) | ADDg (Codeg l) | JPZg l (Codeg l) | JMPg l |
  LABg→ (l → Codeg l) (Codeg l)

```

Intuitively, $\text{JMP}_g l$ jumps to the code with label l , while $\text{LAB}_{\vec{g}} (\lambda l \rightarrow c)$ c' labels the code c' with a fresh label l and makes that label available in the code c so that it can jump to c' . For example, the following piece of linear code, written in an assembly-like notation,

```

    PUSH 2
    JPZ 1
    PUSH 3
    JMP 1
l:  PUSH 4
    HALT

```

can be represented by the following term of type $Code_g\ l$:

$$LAB_g^{\rightarrow} (\lambda l \rightarrow PUSH_g\ 2\ (JPZ_g\ l\ (PUSH_g\ 3\ (JMP_g\ l))))\ (PUSH_g\ 4\ HALT_g)$$

Graph-based code represented using the type $Code_g\ l$ has the desired linear structure. While the constructor LAB_g^{\rightarrow} may appear non-linear since it essentially takes two code arguments, code of the form $LAB_g^{\rightarrow} (\lambda l \rightarrow c)$ c' represents the *sequential composition* of two pieces of code as shown in the example above. Note that only forward jumps are possible at present, because in $LAB_g^{\rightarrow} (\lambda l \rightarrow c)$ c' the label l is only in scope in the first piece of code c , allowing jumps to be made to the second piece of code c' . The forward arrow in LAB_g^{\rightarrow} reflects this fact. We will introduce an instruction LAB_g^{\rightarrow} that permits backward jumps in Section 4 when we consider a language with looping.

The type $Code_g$ is an instance of the notion of a *structured graph* [Oliveira and Cook 2012]. Each structured graph type comes equipped with a canonical function $(\llbracket \cdot \rrbracket)$ that unravels a graph to its corresponding tree, which in our case is defined as follows:

$$\begin{aligned}
(\llbracket \cdot \rrbracket) &:: Code_g\ Code \rightarrow Code \\
(\llbracket HALT_g \rrbracket) &= HALT \\
(\llbracket PUSH_g\ n\ c \rrbracket) &= PUSH\ n\ (\llbracket c \rrbracket) \\
(\llbracket ADD_g\ c \rrbracket) &= ADD\ (\llbracket c \rrbracket) \\
(\llbracket JPZ_g\ l\ c \rrbracket) &= JPZ\ l\ (\llbracket c \rrbracket) \\
(\llbracket JMP_g\ l \rrbracket) &= l \\
(\llbracket LAB_g^{\rightarrow}\ f\ c \rrbracket) &= (\llbracket f\ (\llbracket c \rrbracket) \rrbracket)
\end{aligned}$$

That is, the function $(\llbracket \cdot \rrbracket)$ takes a graph of type $Code_g\ Code$ in which the free occurrences of abstract labels have already been replaced by concrete trees of type $Code$, and converts the graph into a tree. Note that in the case of $JMP_g\ l$ we simply return the code l that is supplied as a parameter, while for $LAB_g^{\rightarrow}\ f\ c$ we first unravel the graph c , then apply the function f to the resulting code, and finally unravel the resulting graph. For instance, applying $(\llbracket \cdot \rrbracket)$ to the example code graph above gives the following code tree, in which the code $PUSH\ 4\ HALT$ at label l occurs twice:

$$PUSH\ 2\ (JPZ\ (PUSH\ 4\ HALT)\ (PUSH\ 3\ (PUSH\ 4\ HALT)))$$

Structured graphs rely on parametricity in the label type l to ensure that labels can only be treated purely symbolically, i.e. they cannot be inspected. In our case, this means that we must only construct graphs of the polymorphic type $\forall l. Code_g\ l$. The unravelling function $(\llbracket \cdot \rrbracket)$ then instantiates this polymorphic type to $Code_g\ Code$, so that in the case of $(\llbracket LAB_g^{\rightarrow} (\lambda l \rightarrow c')\ c \rrbracket)$ it can use function application to substitute the free occurrences of label l in c' with the unravelled code $(\llbracket c \rrbracket)$. Parametricity ensures that the labels provided by LAB_g^{\rightarrow} can only be used as names [Atkey 2009], but otherwise doesn't play a role in our compiler calculations.

3.2 Compiler specification

Given the above definitions, our aim now is to calculate a compiler that produces graph-structured code of type $\forall l. Code_g\ l$. First of all, the semantics $exec_g$ of such code can be defined by:

$$\begin{aligned} exec_g &:: (\forall l. Code_g l) \rightarrow Stack \rightarrow Stack \\ exec_g c s &= exec (\llbracket c \rrbracket) s \end{aligned}$$

In this manner, the execution of graph-structured code comprises two separate parts: the first part, given by $\llbracket \cdot \rrbracket$, specifies how to execute jump instructions, while the second part, given by $exec$, specifies how to execute all other instructions. Note that the parametric type $\forall l. Code_g l$ of the argument c is implicitly instantiated to $Code_g Code$ before it is passed to $\llbracket \cdot \rrbracket$.

The correctness of a graph-based compilation function $comp_g :: Expr \rightarrow Code_g l \rightarrow Code_g l$ can now be expressed by the following equation, which has the same form as equation (1) for $comp$ in Section 2, except that it operates on code graphs rather than code trees:

$$exec_g c (eval x : s) = exec_g (comp_g x c) s \quad (3)$$

The above equation can then be simplified as follows:

$$\begin{aligned} exec_g c (eval x : s) &= exec_g (comp_g x c) s \\ \Leftrightarrow \{ \text{definition of } exec_g \} & \\ exec (\llbracket c \rrbracket) (eval x : s) &= exec (\llbracket comp_g x c \rrbracket) s \\ \Leftrightarrow \{ \text{correctness of } comp \} & \\ exec (comp x (\llbracket c \rrbracket)) s &= exec (\llbracket comp_g x c \rrbracket) s \\ \Leftarrow \{ \text{congruence} \} & \\ comp x (\llbracket c \rrbracket) &= (\llbracket comp_g x c \rrbracket) \end{aligned}$$

In conclusion, to establish the correctness of the compilation function $comp_g$ it suffices to satisfy the following simple equation, which provides a suitable specification from which we can calculate the graph-based compiler $comp_g$ from the tree-based compiler $comp$:

$$comp x (\llbracket c \rrbracket) = (\llbracket comp_g x c \rrbracket) \quad (4)$$

3.3 Compiler calculation

In a similar manner to Section 2, the calculation proceeds from equation (4) by induction on the expression x . For each case, we aim to transform the left side of the equation, $comp x (\llbracket c \rrbracket)$, into the form $\llbracket c' \rrbracket$ for some code c' . We can then define $comp_g x c = c'$ for that case, which by construction satisfies the correctness equation. Unlike the calculation in Section 2, the function $comp_g$ is the only unknown in equation (4), so no further definitions need to be discovered during the calculation. The cases for values and addition are straightforward:

$$\begin{aligned} &comp (Val n) (\llbracket c \rrbracket) \\ &= \{ \text{definition of } comp \} \\ &PUSH n (\llbracket c \rrbracket) \\ &= \{ \text{definition of } \llbracket \cdot \rrbracket \} \\ &\llbracket PUSH_g n c \rrbracket \end{aligned}$$

and

$$\begin{aligned} &comp (Add x y) (\llbracket c \rrbracket) \\ &= \{ \text{definition of } comp \} \\ &comp x (comp y (ADD (\llbracket c \rrbracket))) \\ &= \{ \text{definition of } \llbracket \cdot \rrbracket \} \\ &comp x (comp y (\llbracket ADD_g c \rrbracket)) \\ &= \{ \text{induction hypothesis for } y \} \\ &comp x (\llbracket comp_g y (ADD_g c) \rrbracket) \end{aligned}$$

$$= \{ \text{induction hypothesis for } x \} \\ \llbracket \text{comp}_g x (\text{comp}_g y (\text{ADD}_g c)) \rrbracket$$

The case for conditionals is more involved, as this is where the move to graph-based code becomes important. The calculation begins by applying *comp*, abstracting over the resulting duplicated term $\llbracket c \rrbracket$ to avoid code duplication, and then using the definition of $\llbracket \cdot \rrbracket$ to transform the term into a form to which the induction hypothesis for the two branches *y* and *z* can be applied:

$$\begin{aligned} & \text{comp } (\text{If } x \ y \ z) \llbracket c \rrbracket \\ = & \{ \text{definition of } \text{comp} \} \\ & \text{comp } x (\text{JPZ } (\text{comp } z \llbracket c \rrbracket) (\text{comp } y \llbracket c \rrbracket)) \\ = & \{ \text{abstract over } \llbracket c \rrbracket \} \\ & \text{comp } x ((\lambda l \rightarrow \text{JPZ } (\text{comp } z \ l) (\text{comp } y \ l)) \llbracket c \rrbracket) \\ = & \{ \text{definition of } \llbracket \cdot \rrbracket \} \\ & \text{comp } x ((\lambda l \rightarrow \text{JPZ } (\text{comp } z \ (\text{JMP}_g \ l)) (\text{comp } y \ (\text{JMP}_g \ l))) \llbracket c \rrbracket) \\ = & \{ \text{induction hypothesis for } y \text{ and } z \} \\ & \text{comp } x ((\lambda l \rightarrow \text{JPZ } (\llbracket \text{comp}_g z (\text{JMP}_g \ l) \rrbracket) (\llbracket \text{comp}_g y (\text{JMP}_g \ l) \rrbracket)) \llbracket c \rrbracket) \end{aligned}$$

We then proceed by abstracting over the first argument of the *JPZ* instruction, which is motivated by the desire for this argument to become an abstract jump label, and then transform the term into a form to which the induction hypothesis for *x* can be applied:

$$\begin{aligned} & \text{comp } x ((\lambda l \rightarrow \text{JPZ } (\llbracket \text{comp}_g z (\text{JMP}_g \ l) \rrbracket) (\llbracket \text{comp}_g y (\text{JMP}_g \ l) \rrbracket)) \llbracket c \rrbracket) \\ = & \{ \text{abstract over } \llbracket \text{comp}_g z (\text{JMP}_g \ l) \rrbracket \} \\ & \text{comp } x ((\lambda l \rightarrow (\lambda l' \rightarrow \text{JPZ } l' (\llbracket \text{comp}_g y (\text{JMP}_g \ l) \rrbracket) (\llbracket \text{comp}_g z (\text{JMP}_g \ l) \rrbracket)) \llbracket c \rrbracket)) \\ = & \{ \text{definition of } \llbracket \cdot \rrbracket \} \\ & \text{comp } x ((\lambda l \rightarrow (\lambda l' \rightarrow (\text{JPZ}_g \ l' (\text{comp}_g y (\text{JMP}_g \ l))) (\llbracket \text{comp}_g z (\text{JMP}_g \ l) \rrbracket)) \llbracket c \rrbracket)) \\ = & \{ \text{definition of } \llbracket \cdot \rrbracket \} \\ & \text{comp } x ((\lambda l \rightarrow (\text{LAB}_g^\rightarrow (\lambda l' \rightarrow \text{JPZ}_g \ l' (\text{comp}_g y (\text{JMP}_g \ l)) (\llbracket \text{comp}_g z (\text{JMP}_g \ l) \rrbracket)) \llbracket c \rrbracket)) \\ = & \{ \text{definition of } \llbracket \cdot \rrbracket \} \\ & \text{comp } x (\llbracket \text{LAB}_g^\rightarrow (\lambda l \rightarrow \text{LAB}_g^\rightarrow (\lambda l' \rightarrow \text{JPZ}_g \ l' (\text{comp}_g y (\text{JMP}_g \ l)) (\llbracket \text{comp}_g z (\text{JMP}_g \ l) \rrbracket)) \llbracket c \rrbracket) \rrbracket) \\ = & \{ \text{induction hypothesis for } x \} \\ & \llbracket \text{comp}_g x (\text{LAB}_g^\rightarrow (\lambda l \rightarrow \text{LAB}_g^\rightarrow (\lambda l' \rightarrow \text{JPZ}_g \ l' (\text{comp}_g y (\text{JMP}_g \ l)) (\llbracket \text{comp}_g z (\text{JMP}_g \ l) \rrbracket)) \llbracket c \rrbracket)) \rrbracket \end{aligned}$$

The final term above now has the required form, which completes the calculation for this case. Note that in the two steps before the final step, we use the fact that the unravelling function $\llbracket \cdot \rrbracket$ for the case of labels can be beta-expanded to $\llbracket \text{LAB}_g^\rightarrow f \ c \rrbracket = (\lambda l \rightarrow (\llbracket f \ l \rrbracket)) \llbracket c \rrbracket$.

The top-level compilation function $\text{compile}_g :: \text{Expr} \rightarrow \text{Code}_g \ l$ can be calculated in a similar manner from the tree-based version $\text{compile} :: \text{Expr} \rightarrow \text{Code}$ using the specification $\text{compile } x = \llbracket \text{compile}_g \ x \rrbracket$. Figure 2 summarises all calculated definitions. In conclusion, we have derived a graph-based compiler that uses explicit labels and jumps to avoid the two issues with the previous version, namely tree-structured code and code duplication. For example, applying compile_g to the expression $\text{If } (\text{Val } 2) (\text{Val } 3) (\text{Val } 4)$ gives the code graph

$$\text{PUSH}_g \ 2 \ (\text{LAB}_g^\rightarrow (\lambda l \rightarrow \text{LAB}_g^\rightarrow (\lambda l' \rightarrow \text{JPZ}_g \ l' (\text{PUSH}_g \ 3 \ (\text{JMP}_g \ l)) (\text{PUSH}_g \ 4 \ (\text{JMP}_g \ l))) \ \text{HALT}_g)$$

which corresponds to the following code in assembly-like notation:

$$\begin{aligned}
\text{compile}_g &:: \text{Expr} \rightarrow \text{Code}_g \, l \\
\text{compile}_g \, x &= \text{comp}_g \, x \, \text{HALT}_g \\
\text{comp}_g &:: \text{Expr} \rightarrow \text{Code}_g \, l \rightarrow \text{Code}_g \, l \\
\text{comp}_g \, (\text{Val } n) \quad c &= \text{PUSH}_g \, n \, c \\
\text{comp}_g \, (\text{Add } x \, y) \quad c &= \text{comp}_g \, x \, (\text{comp}_g \, y \, (\text{ADD}_g \, c)) \\
\text{comp}_g \, (\text{If } x \, y \, z) \quad c &= \text{comp}_g \, x \, (\text{LAB}_g^{\rightarrow} \, (\lambda l \rightarrow \text{LAB}_g^{\rightarrow} \, (\lambda l' \rightarrow \\
&\quad \text{JPZ}_g \, l' \, (\text{comp}_g \, y \, (\text{JMP}_g \, l)))) \, (\text{comp}_g \, z \, (\text{JMP}_g \, l)) \, c)
\end{aligned}$$

Fig. 2. Graph-based compiler for conditional expressions.

```

PUSH 2
JPZ l'
PUSH 3
JMP l
l' : PUSH 4
    JMP l
l  : HALT

```

This code is both linear, i.e. each instruction is followed by at most one instruction, and avoids duplication, i.e. there are two jumps to the code at label l rather than this code being duplicated. For this simple example the code at label l is just HALT , but in general it could be an arbitrarily large piece of code, which would be duplicated in the original tree-based function compile .

3.4 Reflection

We conclude this section with some reflective remarks on our new methodology.

Explicit versus generic definitions. The first step in calculating a graph-based compiler is the definition of the type Code_g of code graphs, and the function $(\cdot|_)$ that unravels a code graph to the corresponding code tree. Both follow in a generic manner from the type Code of code trees. This can be made explicit by defining a generic tree type $\text{Tree } i$ and a corresponding graph type $\text{Graph } i \, l$, parameterised by a type i of instructions. For example, given the instruction type

```
data Inst l = PUSH Int | ADD | JPZ l
```

we can show that Code is isomorphic to $\text{Tree } \text{Inst}$ and that $\text{Code}_g \, l$ is isomorphic to $\text{Graph } \text{Inst } l$. The unravelling function $(\cdot|_)$ can be defined as a generic function of type $\forall l. \text{Graph } i \, l \rightarrow \text{Tree } i$, independent of the instruction type i . In turn, we can implement a generic function $\text{seqn} :: \forall l. \text{Graph } i \, l \rightarrow \text{Seqn } i$ that turns a graph into a sequence of instructions with explicit code pointers. By composing a graph-based compiler with the generic function seqn , we can obtain a compiler that produces linear code with explicit jumps. Details can be found in the accompanying Agda code. For simplicity of exposition, in this article we follow the approach of [Bahr and Hutton 2015] and use explicit definitions for the basic types and functions, rather than generic definitions.

Two-step calculation. The methodology [Bahr and Hutton 2015] that we used to calculate the tree-based compiler comp can be viewed as a fusion of three separate calculation steps: introducing a stack, introducing continuations, and defunctionalisation. It is thus natural to ask whether the calculation of the graph-based compiler comp_g from the tree-based version comp can be fused with the calculation of comp from the source language semantics eval , such that we can calculate comp_g directly from eval . However, our calculation technique for producing graph-based compilers

depends crucially on a two-step process, which prevents us from fusing the two steps together. We can see this by considering the correctness specification for $comp_g$:

$$exec_g c (eval x : s) = exec_g (comp_g x c) s$$

On its own, this equation does not express the desire for $comp_g$ to produce graph-structured code. Indeed, the above equation is equivalent to the correctness specification for $comp$,

$$exec c (eval x : s) = exec (comp x c) s$$

if the type $Code_g$ and function $exec_g$ are considered as unknowns, in the same way that $Code$ and $exec$ are unknowns for the latter equation. Instead, the desire for $comp_g$ to produce graph-shaped code is expressed explicitly in the definition of $Code_g$ and $exec_g$, which are given prior to calculation and are thus part of the specification. Moreover, as we have seen, these definitions follow in a systematic manner from the calculation of the tree-based compiler $comp$.

Branching control flow. In essence, the method that we presented here allows us to calculate compilers that produce code with branching control flow. This branching control flow manifests itself as the tree-structured code produced by the original compiler $comp$. However, the tree-structure itself is not necessarily the issue, but rather the underlying branching control flow. The syntax of the source language is also tree-structured, but its semantics is not necessarily branching. In particular, the semantics of addition in the source language is not branching and indeed $comp$ produces linear code for it. Wand [1982] and Gibbons [2021] have used this insight to derive a compiler for a simple expression language without branching control flow by linearising the tree structure of expressions using the associativity of function composition. However, this linearisation crucially depends on the fact that the control flow of the language is not branching.

4 COMPILING TO CYCLIC CODE

Code produced by the compiler in the previous section can jump *forward* to other code locations. However, for source languages that feature cyclic control flow, e.g. in the form of loops, we also need to produce code that can jump *backward*. For example, the following piece of code in an assembly-like notation repeatedly adds two to the top of the stack:

```
PUSH 0
1: PUSH 2
   ADD
   JMP 1
```

To illustrate such looping, we consider a simple expression language with an infinite repetition primitive, whose syntax and semantics can be defined as follows:

data $Expr = Val Int \mid Add Expr Expr \mid Repeat Expr$

$eval :: Expr \rightarrow Int$

$eval (Val n) = n$

$eval (Add x y) = eval x + eval y$

$eval (Repeat x) = \mathbf{let} n = eval x \mathbf{in} (eval (Repeat x))$

This language provides a degenerate form of looping, but suffices to illustrate the core reasoning principles prior to considering a more realistic example in Section 5. However, there is a problem with the above definition for the $eval$ function: we assume that our meta-language is total, but the above definition may not terminate and is hence partial.

To express non-terminating computations, we require a means to explicitly represent them. The *partiality monad* [Capretta 2005; Danielsson 2012] provides this functionality:

codata $Partial\ a = Now\ a \mid Later\ (Partial\ a)$

Intuitively, $Now\ v$ represents a computation that immediately returns the value v , while $Later\ p$ delays the computation p by one time step. The partiality type is defined coinductively, which we indicate above by writing **codata**. In particular, this means that we can delay a computation infinitely often to represent a non-terminating computation, such as:

$never :: Partial\ a$
 $never = Later\ never$

While the definition for $never$ is non-terminating, it is well-defined as the recursive call is guarded by the coinductive constructor $Later$. The partiality type forms a monad under the following definitions, which allows the **do** notation to be used with this type:

$return :: a \rightarrow Partial\ a$
 $return = Now$
 $(\gg) :: Partial\ a \rightarrow (a \rightarrow Partial\ b) \rightarrow Partial\ b$
 $(Now\ x) \gg f = f\ x$
 $(Later\ p) \gg f = Later\ (p \gg f)$

We can now rewrite our semantics for expressions as a function of type $Expr \rightarrow Partial\ Int$. We ensure the definition is well-defined by making every recursive call either structurally recursive, as in the case for addition, or guarding it by $Later$, as in the case for repetition:

$eval :: Expr \rightarrow Partial\ Int$
 $eval\ (Val\ n) = return\ n$
 $eval\ (Add\ x\ y) = do\ m \leftarrow eval\ x; n \leftarrow eval\ y; return\ (m + n)$
 $eval\ (Repeat\ x) = do\ eval\ x; Later\ (eval\ (Repeat\ x))$

4.1 Compiler specification

The compiler calculation methodology extends in a natural manner to a monadic setting [Bahr and Hutton 2022]. First of all, because the source language semantics $eval$ uses the partiality monad to account for non-termination, so too does the target language semantics:

$exec :: Code \rightarrow Stack \rightarrow Partial\ Stack$

Our aim now is to calculate a compiler $comp :: Expr \rightarrow Code \rightarrow Code$ that satisfies the correctness property, which we formulate in a monadic style as follows:

$$do\ n \leftarrow eval\ x; exec\ c\ (n : s) \cong exec\ (comp\ x\ c)\ s \quad (5)$$

Both sides of the equation are computations of type $Partial\ Stack$, and we express their equivalence using the (strong) bisimilarity relation \cong . Intuitively, $p \cong q$ means that either p and q both diverge, or p and q both produce the same value in the same number of steps. To formalise this idea, we define a relation $p \Downarrow_i v$ which expresses that p terminates after at most i steps with result v :

$$\frac{}{Now\ v \Downarrow_i v} \qquad \frac{p \Downarrow_i v}{Later\ p \Downarrow_{i+1} v}$$

We can think of \Downarrow_i as capturing the idea of convergence using i units of ‘fuel’. The base case for $Now\ v$ uses index i rather than zero because we don’t need to use all the fuel. Given two

computations $p, q :: \text{Partial } a$, we say that p and q are bisimilar, written $p \cong q$, if they coincide in terms of their step-counting convergence behaviours, that is:

$$p \Downarrow_i v \quad \text{iff} \quad q \Downarrow_i v \quad \text{for all } v \text{ and } i$$

The semantics of the source language is defined by mixed induction and coinduction: the recursive calls for *Add* are structurally recursive, while the second recursive call for *Repeat* is guarded by *Later*. Correspondingly, in addition to structural induction over *Expr*, compiler calculation requires a coinductive reasoning principle for *Partial a* that is compatible with transitive reasoning. To this end, we define a step-indexed notion of bisimilarity. Given two computations $p, q :: \text{Partial } a$ and a natural number i , we say that p and q are i -bisimilar, written $p \cong_i q$, if the following holds:

$$p \Downarrow_j v \quad \text{iff} \quad q \Downarrow_j v \quad \text{for all } v \text{ and } j < i$$

Step-indexed bisimilarity is by definition downwards closed, i.e. $p \cong_i q$ implies $p \cong_j q$ for all $j \leq i$, which ensures that \cong_i is a congruence for the monadic bind operator. Moreover, by definition, we have $p \cong q$ iff $p \cong_i q$ for all step counts i . Hence, our compiler correctness equation (5) can be expressed in the following equivalent form using step-indexed bisimilarity:

$$\mathbf{do } n \leftarrow \text{eval } x; \text{exec } c (n : s) \cong_i \text{exec } (\text{comp } x \ c) \ s \quad \text{for all } i \quad (6)$$

Compiler correctness can then be established by proving this equation by simultaneous induction on x and i . During the proof, we can assume that for all step counts $j < i$, we have:

$$\mathbf{do } n \leftarrow \text{eval } x'; \text{exec } c' (n : s') \cong_j \text{exec } (\text{comp } x' \ c') \ s'$$

This inductive hypothesis can then be used by applying the following proof rule, which allows the induction hypothesis to be applied to any term that is guarded by a *Later*:

$$\frac{p \cong_j q \quad \text{for all } j < i}{\text{Later } p \cong_i \text{Later } q} \quad (7)$$

During the calculation, we also use the fact that *Partial* satisfies the monad laws up to bisimilarity, and hence the laws also hold for our notion of step-indexed bisimilarity:

$$\begin{aligned} \text{return } x \gg f &\cong f \ x \\ p \gg \text{return} &\cong p \\ (p \gg f) \gg g &\cong p \gg (\lambda x \rightarrow (f \ x \gg g)) \end{aligned}$$

4.2 Compiler calculation

The calculation aims to transform the left side $\mathbf{do } n \leftarrow \text{eval } x; \text{exec } c (n : s)$ of equation (6) into the form $\text{exec } c' \ s$ for some code c' , from which we can then define $\text{comp } x \ c = c'$ for this case. The cases for values and addition proceed in a similar manner to Section 2, except that we are now working in a monadic setting and use step-indexed bisimilarity. The *Val* calculation is given below by way of example, while the *Add* calculation can be found in Bahr and Hutton [2022]:

$$\begin{aligned} &\mathbf{do } v \leftarrow \text{eval } (\text{Val } n); \text{exec } c (v : s) \\ &= \{ \text{definition of } \text{eval} \} \\ &\mathbf{do } v \leftarrow \text{return } n; \text{exec } c (v : s) \\ &\cong_i \{ \text{monad laws} \} \\ &\text{exec } c (n : s) \\ &= \{ \text{define: } \text{exec } (\text{PUSH } n \ c) \ s = \text{exec } c (n : s) \} \\ &\text{exec } (\text{PUSH } n \ c) \ s \end{aligned}$$

The case for *Repeat* begins by applying definitions and the monad laws, which then allows proof rule (7) to be used to apply the induction hypothesis for all step counts $j < i$:

$$\begin{aligned}
& \mathbf{do} \ n \leftarrow \mathit{eval} \ (\mathit{Repeat} \ x); \mathit{exec} \ c \ (n : s) \\
= & \quad \{ \text{definition of } \mathit{eval} \} \\
& \mathbf{do} \ n \leftarrow (\mathbf{do} \ \mathit{eval} \ x; \mathit{Later} \ (\mathit{eval} \ (\mathit{Repeat} \ x))); \mathit{exec} \ c \ (n : s) \\
\cong_i & \quad \{ \text{monad laws} \} \\
& \mathbf{do} \ \mathit{eval} \ x; n \leftarrow \mathit{Later} \ (\mathit{eval} \ (\mathit{Repeat} \ x)); \mathit{exec} \ c \ (n : s) \\
= & \quad \{ \text{definition of } \succcurlyeq \} \\
& \mathbf{do} \ \mathit{eval} \ x; \mathit{Later} \ (\mathbf{do} \ n \leftarrow \mathit{eval} \ (\mathit{Repeat} \ x); \mathit{exec} \ c \ (n : s)) \\
\cong_i & \quad \{ \text{proof rule (7) and induction hypothesis for } j < i \} \\
& \mathbf{do} \ \mathit{eval} \ x; \mathit{Later} \ (\mathit{exec} \ (\mathit{comp} \ (\mathit{Repeat} \ x) \ c) \ s)
\end{aligned}$$

We could now continue the calculation to obtain a term of the desired form $\mathit{exec} \ c' \ s$, from which we can then define $\mathit{comp} \ (\mathit{Repeat} \ x) \ c = c'$. However, the resulting code c' would include a recursive call to $\mathit{comp} \ (\mathit{Repeat} \ x) \ c$ itself, which means that the compiler does not terminate in this case and is hence partial. To ensure that the compiler is total, we introduce a coinductive code constructor $\mathit{REC} :: \infty \ \mathit{Code} \rightarrow \mathit{Code}$ that can be used to guard non-terminating recursive calls, similarly to the use of *Later* in the definition of *eval*. Here we use ∞ to indicate that the argument of *REC* is coinductive, whereas arguments of the other code constructors are inductive. This is similar to the ∞ notation for mixed inductive-coinductive types in Agda [Danielsson and Altenkirch 2010], but to reduce clutter we assume conversions $\mathit{Delay} :: \mathit{Code} \rightarrow \infty \ \mathit{Code}$ and $\mathit{Force} :: \infty \ \mathit{Code} \rightarrow \mathit{Code}$ are inserted implicitly in the appropriate places as in Idris [Brady 2017].

We now resume our calculation. To ensure that the call to $\mathit{comp} \ (\mathit{Repeat} \ x) \ c$ in the current term of the calculation is guarded by *REC*, we can try to solve the following equation:

$$\mathit{exec} \ (\mathit{REC} \ (\mathit{comp} \ (\mathit{Repeat} \ x) \ c)) \ s = \mathit{exec} \ (\mathit{comp} \ (\mathit{Repeat} \ x) \ c) \ s$$

This can be solved by generalising from the specific term $\mathit{comp} \ (\mathit{Repeat} \ x) \ c$ to give

$$\mathit{exec} \ (\mathit{REC} \ c) \ s = \mathit{exec} \ c \ s$$

which can then be taken as a defining clause for *exec*. However, because *REC* is a coinductive constructor, c is not structurally smaller than $\mathit{REC} \ c$. To ensure well-definedness, the recursive call must be guarded by the coinductive constructor *Later*. This constructor already appears in the current term in our calculation, so we can make progress by instead solving the equation

$$\mathit{exec} \ (\mathit{REC} \ (\mathit{comp} \ (\mathit{Repeat} \ x) \ c)) \ s = \mathit{Later} \ (\mathit{exec} \ (\mathit{comp} \ (\mathit{Repeat} \ x) \ c) \ s)$$

which can be achieved by once again generalising from the specific term $\mathit{comp} \ (\mathit{Repeat} \ x) \ c$, and then taking the resulting equation as a defining clause for the function *exec*:

$$\mathit{exec} \ (\mathit{REC} \ c) \ s = \mathit{Later} \ (\mathit{exec} \ c \ s)$$

Using this idea, we can now continue the calculation, during which we introduce a new code constructor *POP* that simply removes the top value from the stack, motivated by the desire to transform the term into a form to which the induction hypothesis for x can be applied:

$$\begin{aligned}
& \mathbf{do} \ \mathit{eval} \ x; \mathit{Later} \ (\mathit{exec} \ (\mathit{comp} \ (\mathit{Repeat} \ x) \ c) \ s) \\
= & \quad \{ \text{define: } \mathit{exec} \ (\mathit{REC} \ c) \ s = \mathit{Later} \ (\mathit{exec} \ c \ s) \} \\
& \mathbf{do} \ \mathit{eval} \ x; \mathit{exec} \ (\mathit{REC} \ (\mathit{comp} \ (\mathit{Repeat} \ x) \ c)) \ s \\
= & \quad \{ \text{define: } \mathit{exec} \ (\mathit{POP} \ c) \ (_ : s) = \mathit{exec} \ c \ s \} \\
& \mathbf{do} \ n \leftarrow \mathit{eval} \ x; \mathit{exec} \ (\mathit{POP} \ (\mathit{REC} \ (\mathit{comp} \ (\mathit{Repeat} \ x) \ c))) \ (n : s)
\end{aligned}$$

```

data Code = HALT | PUSH Int Code | ADD Code | REC ( $\infty$  Code) | POP Code
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 (Repeat x) c = comp x (POP (REC (comp (Repeat x) c)))
exec :: Code  $\rightarrow$  Stack  $\rightarrow$  Partial Stack
exec HALT      s      = return s
exec (PUSH n c) s      = exec c (n : s)
exec (ADD c)   (m : n : s) = exec c ((n + m) : s)
exec (REC c)   s          = Later (exec c s)
exec (POP c)   (_ : s)    = exec c s
exec _         _          = return []

```

Fig. 3. Initial compiler for the repetition language.

\cong_i { induction hypothesis for x }
 $exec (comp x (POP (REC (comp (Repeat x) c)))) s$

The final term above now has the required form, which completes the calculation for this case. Finally, the top-level compilation function $compile :: Expr \rightarrow Code$ can be calculated in a similar manner to Section 2, starting from the following specification:

$$\mathbf{do} \ n \leftarrow eval \ x; \mathbf{return} \ (n : s) \cong \ exec \ (compile \ x) \ s \quad (8)$$

Figure 3 summarises all the definitions we have calculated, along with a catch-all clause to make $exec$ total, which as previously can be chosen arbitrarily as it plays no role in the calculation.

4.3 Cyclic graphs

The compiler we have calculated for the repetition language produces linear code, but suffers from another problem that makes it unrealistic. In particular, the coinductive constructor REC allows code sequences to be infinite, which is utilised when compiling repetition:

$$comp (Repeat x) c = comp x (POP (REC (comp (Repeat x) c)))$$

However, all is not lost as the compiler only produces *cyclic sequences* of code, which can be represented finitely if we allow cyclic jumps. To this end, we define a graph-based version of the code type with explicit labels and jumps, as we did in Section 3. However, that code type only permitted acyclic control flow, i.e. code could jump forward, but not backward.

To allow backward jumps, we include a new instruction LAB_g^- that takes a function parameter of type $l \rightarrow Code_g \ l$, such that a term of the form $LAB_g^- (\lambda l \rightarrow c)$ labels the code c with a fresh label l and makes that label available in c so that it can jump backward to itself. As we now only need to represent finite code sequences, the coinductive constructor REC is no longer required. Taken together, our graph-based code type is defined as follows:

$$\mathbf{data} \ Code_g \ l = HALT_g \mid PUSH_g \ Int \ (Code_g \ l) \mid ADD_g \ (Code_g \ l) \mid POP_g \ (Code_g \ l) \mid JMP_g \ l \mid LAB_g^+ \ (l \rightarrow Code_g \ l) \ (Code_g \ l) \mid LAB_g^- \ (l \rightarrow Code_g \ l)$$

For example, the code fragment from the beginning of this section,

```
PUSH 0
1: PUSH 2
  ADD
  JMP 1
```

can be represented by the following term of type $Code_g$ l :

$$PUSH_g 0 (LAB_g^- (\lambda l \rightarrow PUSH_g 2 (ADD_g (JMP_g l))))$$

The function (\cdot) that unravels a code graph to a code sequence is then defined as follows:

$$\begin{aligned} (\cdot) &:: Code_g \ Code \rightarrow Code \\ (HALT_g) &= HALT \\ (PUSH_g \ n \ c) &= PUSH \ n \ (\cdot c) \\ (ADD_g \ c) &= ADD \ (\cdot c) \\ (POP_g \ c) &= POP \ (\cdot c) \\ (JMP_g \ l) &= l \\ (LAB_g^+ \ f \ c) &= (\cdot (f (\cdot c))) \\ (LAB_g^- \ f) &= (\cdot (f (REC (\cdot (LAB_g^- \ f))))) \end{aligned}$$

Note that for a backward jump LAB_g^- one of the recursive calls is not structurally recursive, and hence we need to guard it by the coinductive constructor REC . For instance, applying (\cdot) to the example code graph above gives the following infinite, but cyclic, code sequence:

$$PUSH \ 0 \ c \ \mathbf{where} \ c = PUSH \ 2 \ (ADD \ (REC \ c))$$

4.4 Compiler calculation

We can now calculate a graph-based compiler for the repetition language in a similar manner to Section 3. As previously, we begin by defining the semantics of code graphs by unravelling:

$$\begin{aligned} exec_g &:: (\forall l. Code_g \ l) \rightarrow Stack \rightarrow Partial \ Stack \\ exec_g \ c \ s &= exec \ (\cdot c) \ s \end{aligned}$$

The correctness of a graph-based compilation function $comp_g :: Expr \rightarrow Code_g \ l \rightarrow Code_g \ l$ can then be expressed by the following bisimilarity equation:

$$\mathbf{do} \ n \leftarrow eval \ x; \ exec_g \ c \ (n : s) \cong \ exec_g \ (comp_g \ x \ c) \ s \quad (9)$$

Similarly to the approach in Section 3, we can use the definition of $exec_g$ and the correctness equation (5) for the original compiler $comp$ to strengthen this equation to:

$$comp \ x \ (\cdot c) \cong (\cdot (comp_g \ x \ c)) \quad (10)$$

This equation uses a bisimilarity relation \cong on the $Code$ type. Similarly to bisimilarity on the coinductive type $Partial$, bisimilarity on code can be expressed equivalently by a step-indexed relation \cong_i , where for two pieces of code $c, d :: Code$, writing $c \cong_i d$ intuitively means that we cannot distinguish between c and d if we traverse both code sequences but may only look under at most i nested occurrences of the coinductive constructor REC . Hence, equation (10) can be expressed in the following equivalent form using step-indexed bisimilarity:

$$comp \ x \ (\cdot c) \cong_i (\cdot (comp_g \ x \ c)) \quad \text{for all } i \quad (11)$$

The step-indexed relation \cong_i is a congruence for all inductive constructors of the *Code* type. In addition, in a similar manner to the coinductive constructor *Later* of the *Partial* monad, we have the following proof rule for the coinductive constructor *REC*:

$$\frac{c \cong_j d \quad \text{for all } j < i}{REC \ c \ \cong_i \ REC \ d} \quad (12)$$

From the above congruence properties, we also obtain the following congruence property for *comp*, which can be proved by a straightforward induction on the expression x :

$$\frac{c \cong_i \ d}{comp \ x \ c \ \cong_i \ comp \ x \ d} \quad (13)$$

Calculation of the compiler $comp_g$ now proceeds from equation (11) by simultaneous induction on x and i . For each case, we aim to transform the left side of the equation, $comp \ x \ (\llbracket c \rrbracket)$, into the form $\llbracket c' \rrbracket$ for some code c' , so that we can then define $comp_g \ x \ c = c'$. The calculations for values and addition are exactly as in Section 3. The calculation for repetition is given below. Its key step is abstracting over the term starting with *REC* that is motivated by the desire to apply the definition clause of $\llbracket \cdot \rrbracket$ in which such a term appears as argument to a function.

$$\begin{aligned} & comp \ (Repeat \ x) \ (\llbracket c \rrbracket) \\ = & \quad \{ \text{definition of } comp \} \\ & comp \ x \ (POP \ (REC \ (comp \ (Repeat \ x) \ (\llbracket c \rrbracket)))) \\ \cong_i & \quad \{ \text{proof rule (11) and induction hypothesis for } j < i \} \\ & comp \ x \ (POP \ (REC \ (\llbracket comp_g \ (Repeat \ x) \ c \rrbracket))) \\ = & \quad \{ \text{abstract over } REC \ (\llbracket comp_g \ (Repeat \ x) \ c \rrbracket) \} \\ & (\lambda l \rightarrow comp \ x \ (POP \ l)) \ (REC \ (\llbracket comp_g \ (Repeat \ x) \ c \rrbracket)) \\ = & \quad \{ \text{definition of } \llbracket \cdot \rrbracket \} \\ & (\lambda l \rightarrow comp \ x \ (POP \ (\mathcal{J}MP_g \ l))) \ (REC \ (\llbracket comp_g \ (Repeat \ x) \ c \rrbracket)) \\ = & \quad \{ \text{definition of } \llbracket \cdot \rrbracket \} \\ & (\lambda l \rightarrow comp \ x \ (\llbracket POP_g \ (\mathcal{J}MP_g \ l) \rrbracket)) \ (REC \ (\llbracket comp_g \ (Repeat \ x) \ c \rrbracket)) \\ \cong_i & \quad \{ \text{induction hypothesis for } x \} \\ & (\lambda l \rightarrow \llbracket comp_g \ x \ (POP_g \ (\mathcal{J}MP_g \ l)) \rrbracket) \ (REC \ (\llbracket comp_g \ (Repeat \ x) \ c \rrbracket)) \\ = & \quad \{ \text{definition of } \llbracket \cdot \rrbracket \} \\ & (\llbracket LAB_g^- \ (\lambda l \rightarrow comp_g \ x \ (POP_g \ (\mathcal{J}MP_g \ l))) \rrbracket) \end{aligned}$$

The final term is now of the required form, which gives the following clause for $comp_g$:

$$comp_g \ (Repeat \ x) \ c = LAB_g^- \ (\lambda l \rightarrow comp_g \ x \ (POP_g \ (\mathcal{J}MP_g \ l)))$$

To understand the final step of the calculation, it is helpful to expand out the term $comp_g \ (Repeat \ x) \ c$ using this clause, from which it is then clear that the final step uses the definition of $\llbracket \cdot \rrbracket$ for backward labels, in the beta-expanded form $\llbracket LAB_g^- \ f \rrbracket = (\lambda l \rightarrow \llbracket f \ l \rrbracket) \ (REC \ (\llbracket LAB_g^- \ f \rrbracket))$:

$$\begin{aligned} & (\lambda l \rightarrow \llbracket comp_g \ x \ (POP_g \ (\mathcal{J}MP_g \ l)) \rrbracket) \ (REC \ (\llbracket LAB_g^- \ (\lambda l \rightarrow comp_g \ x \ (POP_g \ (\mathcal{J}MP_g \ l))) \rrbracket)) \\ = & \quad \{ \text{definition of } \llbracket \cdot \rrbracket \} \\ & (\llbracket LAB_g^- \ (\lambda l \rightarrow comp_g \ x \ (POP_g \ (\mathcal{J}MP_g \ l))) \rrbracket) \end{aligned}$$

Figure 4 gives the full definition for $comp_g$ that we have calculated, along with the top-level compilation function $compile_g$ that can be calculated in a similar manner to previously.

$$\begin{aligned}
& \text{compile}_g :: \text{Expr} \rightarrow \text{Code}_g \text{ l} \\
& \text{compile}_g x = \text{comp}_g x \text{ HALT}_g \\
& \text{comp}_g :: \text{Expr} \rightarrow \text{Code}_g \text{ l} \rightarrow \text{Code}_g \text{ l} \\
& \text{comp}_g (\text{Val } n) \quad c = \text{PUSH}_g n c \\
& \text{comp}_g (\text{Add } x \ y) \quad c = \text{comp}_g x (\text{comp}_g y (\text{ADD}_g c)) \\
& \text{comp}_g (\text{Repeat } x) \quad c = \text{LAB}_g^- (\lambda l \rightarrow \text{comp}_g x (\text{POP}_g (\text{JMP}_g l)))
\end{aligned}$$

Fig. 4. Graph-based compiler for the repetition language.

In conclusion, we have derived a graph-based compiler that produces finite code by representing the infinite code sequences produced by *compile* using backward jumps. For example, applying *compile_g* to the expression *Repeat (Add (Val 2) (Val 3))* gives the code graph

$$\text{LAB}_g^- (\lambda l \rightarrow \text{PUSH}_g 2 (\text{PUSH}_g 3 (\text{ADD}_g (\text{POP}_g (\text{JMP}_g l))))))$$

which corresponds to the following code in assembly-like notation:

```

l:  PUSH 2
    PUSH 3
    ADD
    POP
    JMP l

```

In particular, this code avoids the need for infinite code by using a label and jump to capture the cyclic behaviour of the repetition primitive in the source language.

5 WHILE LOOPS AND STATE

For our final example, we consider a more realistic language that features while loops, along with primitives for reading and writing a mutable reference cell that stores an integer value. In particular, the reference cell allows loops to be constructed that have any desired number of iterations, both finite and infinite. The syntax of the language is defined as follows:

data *Expr* = *Val Int* | *Add Expr Expr* | *Get* | *Put Expr Expr* | *While Expr Expr*

Informally, *Get* returns the current value of the reference cell, *Put x y* sets the cell to the value of the expression *x* and then behaves as the expression *y*, and *While x y* repeatedly evaluates *y* as long as *x* has a non-zero value. For example, the following expression sets the reference cell to ten and then repeatedly decrements the reference until it becomes zero:

$$\text{Put } (\text{Val } 10) (\text{While } \text{Get } (\text{Put } (\text{Add } \text{Get } (\text{Val } (-1))) \text{Get}))$$

Because while loops may not always terminate, we again use the partiality monad to define the semantics of the language. In addition, we also need to keep track of the state of the reference cell, so the semantics of an expression becomes a state transformer:

type *State* = *Int*

$$\text{eval} :: \text{Expr} \rightarrow \text{State} \rightarrow \text{Partial } (\text{Int}, \text{State})$$

$$\text{eval } (\text{Val } n) \quad q = \text{return } (n, q)$$

$$\text{eval } (\text{Add } x \ y) \quad q = \text{do } (m, q_1) \leftarrow \text{eval } x \ q; (n, q_2) \leftarrow \text{eval } y \ q_1; \text{return } (m + n, q_2)$$

$$\text{eval } \text{Get} \quad q = \text{return } (q, q)$$

$$\text{eval } (\text{Put } x \ y) \quad q = \text{do } (n, _) \leftarrow \text{eval } x \ q; \text{eval } y \ n$$

$$\begin{aligned} \text{eval } (\text{While } x \ y) \ q &= \mathbf{do} \ (n, q_1) \leftarrow \text{eval } x \ q \\ &\quad \mathbf{if} \ n == 0 \ \mathbf{then} \ \text{return} \ (0, q_1) \\ &\quad \mathbf{else} \ \mathbf{do} \ (_, q_2) \leftarrow \text{eval } y \ q_1; \text{Later} \ (\text{eval} \ (\text{While } x \ y) \ q_2) \end{aligned}$$

As we already use the symbol s for stacks, in this section we use the symbol q for states. Note that the above semantics combines the partiality monad with the state monad. However, for calculation purposes we don't treat the state monad abstractly, but instead explicitly manipulate the state as shown above. This approach allows us to discover an appropriate mechanism to compile the state monad effect to a lower-level implementation in the target machine.

5.1 Tree-based compiler

For the tree-based compiler, we follow the methodology of Bahr and Hutton [2015], in which the target machine is specified as a function $\text{exec} :: \text{Code} \rightarrow \text{Conf} \rightarrow \text{Conf}$, where Conf is the type of machine configurations that comprise a stack and a stack. However, as in the previous section, to account for non-termination we amend the type of exec to use the partiality monad:

type $\text{Conf} = (\text{Stack}, \text{State})$

$\text{exec} :: \text{Code} \rightarrow \text{Conf} \rightarrow \text{Partial Conf}$

The type of the compiler remains the same as previously, $\text{comp} :: \text{Expr} \rightarrow \text{Code} \rightarrow \text{Code}$, while the compiler correctness property can be adapted for the presence of state as follows:

$$\mathbf{do} \ (n, q') \leftarrow \text{eval } x \ q; \text{exec } c \ (n : s, q') \cong \text{exec} \ (\text{comp } x \ c) \ (s, q) \quad (14)$$

As before, however, for the purposes of calculating the compiler we utilise the following equivalent version that is expressed using step-indexed bisimilarity:

$$\mathbf{do} \ (n, q') \leftarrow \text{eval } x \ q; \text{exec } c \ (n : s, q') \cong_i \text{exec} \ (\text{comp } x \ c) \ (s, q) \quad \text{for all } i \quad (15)$$

We now proceed with the calculation, by induction on x and i . In a similar manner to previously, we start with the left side of equation (15) and aim to transform it into the form $\text{exec } c' \ (s, q)$ for some code c' , so that we can then define $\text{comp } x \ c = c'$. We give the calculations for *Get* and *While* below. The other cases are similar and can be found in the accompanying Agda formalisation.

The case for *Get* applies the definition of eval , simplifies the resulting term using the monad laws, and finally introduces a new code constructor *LOAD* to obtain a term of the desired form:

$$\begin{aligned} &\mathbf{do} \ (n, q') \leftarrow \text{eval } \text{Get } q; \text{exec } c \ (n : s, q') \\ &= \quad \{ \text{definition of } \text{eval} \} \\ &\mathbf{do} \ (n, q') \leftarrow \text{return} \ (q, q); \text{exec } c \ (n : s, q') \\ &\cong_i \quad \{ \text{monad laws} \} \\ &\quad \text{exec } c \ (q : s, q) \\ &= \quad \{ \text{define } \text{exec} \ (\text{LOAD } c) \ (s, q) = \text{exec } c \ (q : s, q) \} \\ &\quad \text{exec} \ (\text{LOAD } c) \ (s, q) \end{aligned}$$

The case for *While* begins in a manner similar to *Repeat* from Section 4.2:

$$\begin{aligned} &\mathbf{do} \ (m, q') \leftarrow \text{eval} \ (\text{While } x \ y) \ q; \text{exec } c \ (m : s, q') \\ &= \quad \{ \text{definition of } \text{eval} \} \\ &\mathbf{do} \ (m, q') \leftarrow (\mathbf{do} \ (n, q_1) \leftarrow \text{eval } x \ q; \mathbf{if} \ n == 0 \ \mathbf{then} \ \text{return} \ (0, q_1) \\ &\quad \quad \quad \mathbf{else} \ \mathbf{do} \ (_, q_2) \leftarrow \text{eval } y \ q_1 \\ &\quad \quad \quad \text{Later} \ (\text{eval} \ (\text{While } x \ y) \ q_2)) \\ &\quad \text{exec } c \ (m : s, q') \\ &\cong_i \quad \{ \text{monad laws} \} \end{aligned}$$

$$\begin{aligned}
& \mathbf{do} (n, q_1) \leftarrow \mathit{eval} \ x \ q; \mathbf{if} \ n == 0 \ \mathbf{then} \ \mathit{exec} \ c \ (0 : s, q_1) \\
& \quad \mathbf{else} \ \mathbf{do} \ (_, q_2) \leftarrow \mathit{eval} \ y \ q_1 \\
& \quad \quad (n, q_3) \leftarrow \mathit{Later} \ (\mathit{eval} \ (\mathit{While} \ x \ y) \ q_2) \\
& \quad \quad \mathit{exec} \ c \ (n : s, q_3) \\
= & \quad \{ \text{definition of } \succcurlyeq \} \\
& \mathbf{do} (n, q_1) \leftarrow \mathit{eval} \ x \ q; \mathbf{if} \ n == 0 \ \mathbf{then} \ \mathit{exec} \ c \ (0 : s, q_1) \\
& \quad \mathbf{else} \ \mathbf{do} \ (_, q_2) \leftarrow \mathit{eval} \ y \ q_1 \\
& \quad \quad \mathit{Later} \ (\mathbf{do} \ (n, q_3) \leftarrow \mathit{eval} \ (\mathit{While} \ x \ y) \ q_2; \mathit{exec} \ c \ (n : s, q_3)) \\
\cong_i & \quad \{ \text{proof rule (7) and induction hypothesis for } j < i \} \\
& \mathbf{do} (n, q_1) \leftarrow \mathit{eval} \ x \ q; \mathbf{if} \ n == 0 \ \mathbf{then} \ \mathit{exec} \ c \ (0 : s, q_1) \\
& \quad \mathbf{else} \ \mathbf{do} \ (_, q_2) \leftarrow \mathit{eval} \ y \ q_1 \\
& \quad \quad \mathit{Later} \ (\mathit{exec} \ (\mathit{comp} \ (\mathit{While} \ x \ y) \ c) \ (s, q_2))
\end{aligned}$$

Similarly to the calculation for *Repeat*, to ensure that the resulting compiler is total, we now seek to guard the call $\mathit{comp} \ (\mathit{While} \ x \ y) \ c$ in the final term above by a coinductive code constructor $\mathit{REC} :: \infty \mathit{Code} \rightarrow \mathit{Code}$. That is, we seek to solve the following equation:

$$\mathit{exec} \ (\mathit{REC} \ (\mathit{comp} \ (\mathit{While} \ x \ y) \ c)) \ (s, q_2) = \mathit{Later} \ (\mathit{exec} \ (\mathit{comp} \ (\mathit{While} \ x \ y) \ c) \ (s, q_2))$$

This can be solved by generalising from the specific term $\mathit{comp} \ (\mathit{While} \ x \ y) \ c$ to give

$$\mathit{exec} \ (\mathit{REC} \ c) \ (s, q_2) = \mathit{Later} \ (\mathit{exec} \ c \ (s, q_2))$$

which can then be taken as defining clause for exec . Using this idea, we can now continue the calculation, during which we introduce new code constructors *POP* and *JPZ* in a similar manner to previously, to allow the induction hypotheses for the argument expressions to be applied:

$$\begin{aligned}
& \mathbf{do} (n, q_1) \leftarrow \mathit{eval} \ x \ q; \mathbf{if} \ n == 0 \ \mathbf{then} \ \mathit{exec} \ c \ (0 : s, q_1) \\
& \quad \mathbf{else} \ \mathbf{do} \ (_, q_2) \leftarrow \mathit{eval} \ y \ q_1; \mathit{Later} \ (\mathit{exec} \ (\mathit{comp} \ (\mathit{While} \ x \ y) \ c) \ (s, q_2)) \\
= & \quad \{ \text{define: } \mathit{exec} \ (\mathit{REC} \ c) \ (s, q) = \mathit{Later} \ (\mathit{exec} \ c \ (s, q)) \} \\
& \mathbf{do} (n, q_1) \leftarrow \mathit{eval} \ x \ q; \mathbf{if} \ n == 0 \ \mathbf{then} \ \mathit{exec} \ c \ (0 : s, q_1) \\
& \quad \mathbf{else} \ \mathbf{do} \ (_, q_2) \leftarrow \mathit{eval} \ y \ q_1; \mathit{exec} \ (\mathit{REC} \ (\mathit{comp} \ (\mathit{While} \ x \ y) \ c)) \ (s, q_2) \\
= & \quad \{ \text{define: } \mathit{exec} \ (\mathit{POP} \ c) \ (_ : s, q) = \mathit{exec} \ c \ (s, q) \} \\
& \mathbf{do} (n, q_1) \leftarrow \mathit{eval} \ x \ q; \mathbf{if} \ n == 0 \ \mathbf{then} \ \mathit{exec} \ c \ (0 : s, q_1) \\
& \quad \mathbf{else} \ \mathbf{do} \ (m, q_2) \leftarrow \mathit{eval} \ y \ q_1 \\
& \quad \quad \mathit{exec} \ (\mathit{POP} \ (\mathit{REC} \ (\mathit{comp} \ (\mathit{While} \ x \ y) \ c))) \ (m : s, q_2) \\
\cong_i & \quad \{ \text{induction hypothesis for } y \} \\
& \mathbf{do} (n, q_1) \leftarrow \mathit{eval} \ x \ q; \mathbf{if} \ n == 0 \ \mathbf{then} \ \mathit{exec} \ c \ (0 : s, q_1) \\
& \quad \mathbf{else} \ \mathit{exec} \ (\mathit{comp} \ y \ (\mathit{POP} \ (\mathit{REC} \ (\mathit{comp} \ (\mathit{While} \ x \ y) \ c)))) \ (s, q_1) \\
= & \quad \{ \text{define: } \mathit{exec} \ (\mathit{JPZ} \ c' \ c) \ (n : s, q) = \mathbf{if} \ n == 0 \ \mathbf{then} \ \mathit{exec} \ c' \ (0 : s, q) \ \mathbf{else} \ \mathit{exec} \ c \ (s, q) \} \\
& \mathbf{do} (n, q_1) \leftarrow \mathit{eval} \ x \ q; \mathit{exec} \ \mathit{JPZ} \ c \ (\mathit{comp} \ y \ (\mathit{POP} \ (\mathit{REC} \ (\mathit{comp} \ (\mathit{While} \ x \ y) \ c)))) \ (n : s, q_1) \\
\cong_i & \quad \{ \text{induction hypothesis for } x \} \\
& \mathit{exec} \ (\mathit{comp} \ x \ (\mathit{JPZ} \ c \ (\mathit{comp} \ y \ (\mathit{POP} \ (\mathit{REC} \ (\mathit{comp} \ (\mathit{While} \ x \ y) \ c)))))) \ (s, q)
\end{aligned}$$

The final term now has the required form, from which we obtain the definition:

$$\mathit{comp} \ (\mathit{While} \ x \ y) \ c = \mathit{comp} \ x \ (\mathit{JPZ} \ c \ (\mathit{comp} \ y \ (\mathit{POP} \ (\mathit{REC} \ (\mathit{comp} \ (\mathit{While} \ x \ y) \ c))))$$

Figure 5 presents all definitions that are calculated in this manner, along with the top-level compilation function $\mathit{compile}$, and a catch-all clause for exec to make it total.

```

data Code = HALT | PUSH Int Code | ADD Code | LOAD Code | STORE Code |
           REC (∞ Code) | POP Code | JPBZ Code Code

compile :: Expr → Code
compile e = comp e HALT

comp :: Expr → Code → Code
comp (Val n)    c = PUSH n c
comp (Add x y)  c = comp x (comp y (ADD c))
comp Get       c = LOAD c
comp (Put x y)  c = comp x (STORE (comp y c))
comp (While x y) c = comp x (JPBZ c (comp y (POP (REC (comp (While x y) c)))))

exec :: Code → Conf → Partial Conf
exec HALT      conf      = return conf
exec (PUSH n c) (s, q)    = exec c (n : s, q)
exec (ADD c)    (m : n : s, q) = exec c ((n + m) : s, q)
exec (LOAD c)   (s, q)       = exec c (q : s, q)
exec (STORE c)  (n : s, _)   = exec c (s, n)
exec (REC c)    conf        = Later (exec c conf)
exec (POP c)    (_, s, q)    = exec c (s, q)
exec (JPBZ c' c) (n : s, q) = if n == 0 then exec c' (0 : s, q) else exec c (s, q)
exec _         _            = return ([], 0)

```

Fig. 5. Tree-based compiler for the while language.

5.2 Graph-based compiler

The compiler we have calculated for the repetition language is unrealistic for three reasons, all concerning while loops. First of all, the compiler produces tree-shaped code because the `JPBZ` instruction takes two pieces of code as arguments. Secondly, it duplicates the additional code `c` when compiling while loops. And finally, it produces infinite code because while loops are compiled by unrolling the loop. However, the compiler in fact only produces *regular trees*, i.e. trees with only finitely many subtrees, which may be represented finitely if we allow cyclic jumps.

To calculate a graph-based compiler from the tree-based version in Figure 5, we proceed in the same manner as the repetition language in the previous section. First of all, we turn the `Code` type into a corresponding structured graph type `Codeg l` that supports forward and backward jumps, and define its canonical unravelling function $\llbracket \cdot \rrbracket$, as shown in Figure 6.

Next, we define the semantics of the graph-based target language by unravelling:

```

execg :: (∀ l. Codeg l) → Conf → Partial Conf
execg c s = exec (llbracket c llbracket) s

```

The correctness of our graph-based compilation function $comp_g :: Expr \rightarrow Code_g l \rightarrow Code_g l$ can then be expressed by the following bisimilarity relation in the partiality monad:

$$exec_g c (eval x : s, q) \cong exec_g (comp_g x c) (s, q) \quad (16)$$

data $Code_g l$ $= HALT_g$ $ PUSH_g Int (Code_g l)$ $ ADD_g (Code_g l)$ $ LOAD_g (Code_g l)$ $ STORE_g (Code_g l)$ $ POP_g (Code_g l)$ $ JPZ_g l (Code_g l)$ $ JMP_g l$ $ LAB_g^+ (l \rightarrow Code_g l) (Code_g l)$ $ LAB_g^- (l \rightarrow Code_g l)$	$(\cdot) :: Code_g Code \rightarrow Code$ $(\langle HALT_g \rangle) = HALT$ $(\langle PUSH_g n c \rangle) = PUSH n (\langle c \rangle)$ $(\langle ADD_g c \rangle) = ADD (\langle c \rangle)$ $(\langle LOAD_g c \rangle) = LOAD (\langle c \rangle)$ $(\langle STORE_g c \rangle) = STORE (\langle c \rangle)$ $(\langle POP_g c \rangle) = POP (\langle c \rangle)$ $(\langle JPZ_g l c \rangle) = JPZ l (\langle c \rangle)$ $(\langle JMP_g l \rangle) = l$ $(\langle LAB_g^+ f c \rangle) = (\langle f (\langle c \rangle) \rangle)$ $(\langle LAB_g^- f \rangle) = (\langle f (REC (\langle LAB_g^- f \rangle)) \rangle)$
--	---

Fig. 6. Graph-based code and its unravelling for the while language.

Using the definition of $exec_g$ and the correctness equation (14) for the tree-based compiler $comp$, this equation can be strengthened to a bisimilarity relation on the $Code$ type,

$$comp\ x\ (\langle c \rangle) \cong (\langle comp_g\ x\ c \rangle) \quad (17)$$

which can then be formulated equivalently using a step-indexed bisimilarity relation on $Code$:

$$comp\ x\ (\langle c \rangle) \cong_i (\langle comp_g\ x\ c \rangle) \quad \text{for all } i \quad (18)$$

The calculation of $comp_g$ proceeds from equation (18) by induction on x and i . The cases for *Val*, *Add*, *Get* and *Put* are straightforward: first apply the definition of $comp$, then turn all tree-based terms into graph-based versions in a syntax-directed manner by applying the definition of unravelling, or an induction hypothesis. The interesting case is the calculation for *While*, which is shown in Figure 7. This calculation combines the ideas that were used for conditionals in Section 3 and repetition in Section 4. In particular, the key steps in the calculation are abstracting over the REC term to introduce a backward jump that avoids infinite code, and abstracting over the term $(\langle c \rangle)$ to introduce a forward jump that avoids tree-shaped code and code duplication. The result of the calculation is the graph-based compiler implementation shown in Figure 8.

6 RELATED WORK

We have discussed related work on deriving correct-by-construction compilers throughout the paper. We supplement that discussion with related work in other relevant areas below.

Graphs. There are many approaches to representing graph in functional languages. Most use explicit names to identify vertices in a graph, often represented as integers [Kashiwagi and Wise 1991; King and Launchbury 1995; Martin Erwig 2001]. Such representations are often motivated by efficiency considerations. However, efficient graph representations can still admit powerful high-level reasoning principles as King and Launchbury [1995] have shown. More recently, Mokhov [2017] introduced a high-level combinator library to construct graphs, and the structured graphs of [Oliveira and Cook 2012] make use of higher-order abstract syntax to avoid explicit names.

Reasoning. To the best of our knowledge, there is little work on the use of structured graphs for formal reasoning. Oliveira and Cook [2012] present some equational laws for *map* and *fold* for a simple type of structured graphs, namely cyclic streams. Bahr [2014] uses structured graphs to implement a graph-based compiler and prove it correct in two stages: first, prove a tree-based

$$\begin{aligned}
& \text{comp} (\text{While } x \ y) \ (\llbracket c \rrbracket) \\
= & \quad \{ \text{definition of } \text{comp} \} \\
& \text{comp } x \ (\text{JPZ} \ (\llbracket c \rrbracket) \ (\text{comp } y \ (\text{POP} \ (\text{REC} \ (\text{comp} \ (\text{While } x \ y) \ (\llbracket c \rrbracket)))))) \\
\cong_i & \quad \{ \text{induction hypothesis for } j < i \} \\
& \text{comp } x \ (\text{JPZ} \ (\llbracket c \rrbracket) \ (\text{comp } y \ (\text{POP} \ (\text{REC} \ (\llbracket \text{comp}_g \ (\text{While } x \ y) \ c \rrbracket)))))) \\
= & \quad \{ \text{abstract over } \text{REC} \ (\llbracket \text{comp}_g \ (\text{While } x \ y) \ c \rrbracket) \} \\
& (\lambda l \rightarrow \text{comp } x \ (\text{JPZ} \ (\llbracket c \rrbracket) \ (\text{comp } y \ (\text{POP} \ l)))) \ (\text{REC} \ (\llbracket \text{comp}_g \ (\text{While } x \ y) \ c \rrbracket)) \\
= & \quad \{ \text{definition of } (\cdot) \} \\
& (\lambda l \rightarrow \text{comp } x \ (\text{JPZ} \ (\llbracket c \rrbracket) \ (\text{comp } y \ (\text{POP} \ (\llbracket \text{JMP}_g \ l \rrbracket)))) \ (\text{REC} \ (\llbracket \text{comp}_g \ (\text{While } x \ y) \ c \rrbracket)) \\
= & \quad \{ \text{definition of } (\cdot) \} \\
& (\lambda l \rightarrow \text{comp } x \ (\text{JPZ} \ (\llbracket c \rrbracket) \ (\text{comp } y \ (\text{POP}_g \ (\text{JMP}_g \ l)))) \ (\text{REC} \ (\llbracket \text{comp}_g \ (\text{While } x \ y) \ c \rrbracket)) \\
\cong_i & \quad \{ \text{induction hypothesis for } y \} \\
& (\lambda l \rightarrow \text{comp } x \ (\text{JPZ} \ (\llbracket c \rrbracket) \ (\llbracket \text{comp}_g \ y \ (\text{POP}_g \ (\text{JMP}_g \ l) \rrbracket)))) \ (\text{REC} \ (\llbracket \text{comp}_g \ (\text{While } x \ y) \ c \rrbracket)) \\
= & \quad \{ \text{abstract over } (\llbracket c \rrbracket) \} \\
& (\lambda l \rightarrow \text{comp } x \ ((\lambda l' \rightarrow \text{JPZ } l' \ (\llbracket \text{comp}_g \ y \ (\text{POP}_g \ (\text{JMP}_g \ l) \rrbracket)) \ (\llbracket c \rrbracket))) \ (\text{REC} \ (\llbracket \text{comp}_g \ (\text{While } x \ y) \ c \rrbracket)) \\
= & \quad \{ \text{definition of } (\cdot) \} \\
& (\lambda l \rightarrow \text{comp } x \ ((\lambda l' \rightarrow (\text{JPZ}_g \ l' \ (\text{comp}_g \ y \ (\text{POP}_g \ (\text{JMP}_g \ l) \rrbracket)) \ (\llbracket c \rrbracket))) \ (\text{REC} \ (\llbracket \text{comp}_g \ (\text{While } x \ y) \ c \rrbracket)) \\
= & \quad \{ \text{definition of } (\cdot) \} \\
& (\lambda l \rightarrow \text{comp } x \ (\llbracket \text{LAB}_g^+ \ (\lambda l' \rightarrow \text{JPZ}_g \ l' \ (\text{comp}_g \ y \ (\text{POP}_g \ (\text{JMP}_g \ l) \rrbracket)) \ c \rrbracket) \ (\text{REC} \ (\llbracket \text{comp}_g \ (\text{While } x \ y) \ c \rrbracket)) \\
\cong_i & \quad \{ \text{induction hypothesis for } x \} \\
& (\lambda l \rightarrow (\llbracket \text{comp}_g \ x \ (\text{LAB}_g^+ \ (\lambda l' \rightarrow \text{JPZ}_g \ l' \ (\text{comp}_g \ y \ (\text{POP}_g \ (\text{JMP}_g \ l) \rrbracket)) \ c \rrbracket) \rrbracket) \ (\text{REC} \ (\llbracket \text{comp}_g \ (\text{While } x \ y) \ c \rrbracket)) \\
= & \quad \{ \text{definition of } (\cdot) \} \\
& (\llbracket \text{LAB}_g^- \ (\lambda l \rightarrow \text{comp}_g \ x \ (\text{LAB}_g^+ \ (\lambda l' \rightarrow \text{JPZ}_g \ l' \ (\text{comp}_g \ y \ (\text{POP}_g \ (\text{JMP}_g \ l) \rrbracket)) \ c)) \rrbracket)
\end{aligned}$$

Fig. 7. Calculation of comp_g for while loops.

$$\begin{aligned}
\text{compile}_g &:: \text{Expr} \rightarrow \text{Code}_g \ l \\
\text{compile}_g \ e &= \text{comp}_g \ e \ \text{HALT}_g \\
\text{comp}_g &:: \text{Expr} \rightarrow \text{Code}_g \ l \rightarrow \text{Code}_g \ l \\
\text{comp}_g \ (\text{Val } n) & \quad c = \text{PUSH}_g \ n \ c \\
\text{comp}_g \ (\text{Add } x \ y) & \quad c = \text{comp}_g \ x \ (\text{comp}_g \ y \ (\text{ADD}_g \ c)) \\
\text{comp}_g \ \text{Get} & \quad c = \text{LOAD}_g \ c \\
\text{comp}_g \ (\text{Put } x \ y) & \quad c = \text{comp}_g \ x \ (\text{STORE}_g \ (\text{comp}_g \ y \ c)) \\
\text{comp}_g \ (\text{While } x \ y) & \quad c = \text{LAB}_g^- \ (\lambda l \rightarrow \text{comp}_g \ x \ (\text{LAB}_g^+ \ (\lambda l' \rightarrow \text{JPZ}_g \ l' \ (\text{comp}_g \ y \ (\text{POP}_g \ (\text{JMP}_g \ l) \rrbracket)) \ c))
\end{aligned}$$

Fig. 8. Graph-based compiler for the while language.

compiler correct, essentially the correctness property (1) of comp ; and secondly, prove that the graph-based compiler is equivalent to the tree-based compiler, essentially the correctness property (4) of comp_g . The main differences to the present work is that Bahr proves correctness post hoc, whereas we calculate compilers from their specification, and his work is limited to acyclic graphs and terminating languages, whereas we considered cyclic graphs and non-termination.

Cyclic control-flow. Cyclic graphs mark a qualitative difference in the expressive power of the target language compared to acyclic graphs, as graphs can no longer be represented by finite trees if they contain cycles. Some previous work does consider languages with cyclic control-flow, e.g. Bahr and Hutton [2022] calculate a compiler for a language with while loops. However, as the

target language is limited to finite trees, the looping behaviour is instead imitated by repeatedly copying the code for a while loop to the stack and retrieving it later when it has to be executed again. Meijer [1992] also calculates a compiler for a language with while loops, starting with a least fixed-point semantics of such loops. In Meijer’s approach, the syntax of the target language is not formally defined but is instead simply made up of the calculated combinators that are used to describe the source language semantics. Depending on the semantics of the meta language, one could interpret the target code as infinite trees. But it is not clear how the least fixed-points of the semantics give rise to a co-inductive syntax for the target language.

Control-flow graphs. The code produced by the graph-based compiler $compile_g$ corresponds to a control-flow graph. In compiler implementations, such graphs typically use concrete labels, i.e. unique identifiers, to specify targets of branching control flow [Ramsey and Dias 2006; Ramsey et al. 2010]. We also find this representation of control-flow using concrete labels in large-scale compiler verification efforts, e.g. CompCert [Leroy 2009] and CakeML [Kumar et al. 2014].

7 CONCLUSION AND FURTHER WORK

We have shown how existing work on deriving correct-by-construction compilers that produce tree-structured code can be extended to instead produce graph-structured code. While our method is calculational and hence naturally complements a calculational approach such as Meijer [1992] and Bahr and Hutton [2015], it also readily applies to other methods such as Wand [1995] and Ager et al. [2003]. By combining our method with existing tree-based calculation methods, we can calculate graph-based compilers for lambda calculi [Bahr and Hutton 2022], concurrent source languages [Bahr and Hutton 2023] and register-based target languages [Bahr and Hutton 2020].

The key issue when calculating graph-based compilers is dealing with names, i.e. labels that denote code pointers. We used parametric higher-order abstract syntax [Chlipala 2008] because it requires little additional machinery, both for manual pen and paper proofs, and when using a mechanical proof assistant. A drawback of parametric higher-order abstract syntax is that reasoning about parametricity is underdeveloped, and often not possible in proof assistants. For example, such reasoning is necessary if we wished to calculate a direct implementation of the graph-based machine $exec_g$ by fusing together the $exec$ and (\cdot) functions in the definition $exec_g\ c\ s = exec\ (\cdot c)\ s$. In the accompanying Agda formalisation, we give an example of such a calculation. The formalisation also gives an alternative representation of $Code_g$, which uses de Bruijn indices [de Bruijn 1972] rather than parametric higher-order abstract syntax. It may also be useful to consider other name representations, particularly if good support for them is available in a proof assistant.

Our graph-based code type $Code_g$ has two labelling constructors LAB_g^{\rightarrow} and LAB_g^{\leftarrow} for forward and for backward jumps, respectively. Some graphs are not expressible using only these constructors, such as those with overlapping jumps, but so far we have not found the need for more general graph constructors in our compiler calculations. However, LAB_g^{\rightarrow} and LAB_g^{\leftarrow} can be combined into a more general constructor $LAB_g :: (Vec\ l\ n \rightarrow Vec\ (Code_g\ l)\ n) \rightarrow Code_g\ l$ similarly to the original work on structured graphs [Oliveira and Cook 2012] to extend expressiveness if needed.

Our use of step-indexing for equational reasoning about coinductive types, namely *Partial* and *Code*, is also motivated by the fact that it requires relatively little formal machinery, and can be readily used in many proof assistants. More advanced methods for coinductive definitions and reasoning, such as guarded recursion [Nakano 2000] and sized types [Hughes et al. 1996], could allow compiler calculations to be streamlined in a proof assistant that supports these features. In particular if combined with cubical type theory [Birkedal et al. 2016], guarded types support powerful equational reasoning principles for coinductive types [Møgelberg and Veltri 2019].

REFERENCES

- Mads Sig Ager, Dariusz Biernacki, Olivier Danvy, and Jan Midtgaard. 2003. *From Interpreter to Compiler and Virtual Machine: A Functional Derivation*. Technical Report RS-03-14. Department of Computer Science, University of Aarhus.
- Robert Atkey. 2009. Syntax for Free: Representing Syntax with Binding Using Parametricity. In *Typed Lambda Calculi and Applications*. Lecture Notes in Computer Science, Vol. 5608.
- Roland Backhouse. 2003. *Program Construction: Calculating Implementations from Specifications*. John Wiley and Sons, Inc.
- Patrick Bahr. 2014. Proving Correctness of Compilers Using Structured Graphs. In *Functional and Logic Programming (Lecture Notes in Computer Science, Vol. 8475)*.
- 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, Article 93 (2022).
- Patrick Bahr and Graham Hutton. 2023. Calculating Compilers for Concurrency. *Proceedings of the ACM on Programming Languages* 7, ICFP, Article 213 (2023).
- Lars Birkedal, Aleš Bizjak, Ranald Clouston, Hans Bugge Grathwohl, Bas Spitters, and Andrea Vezzosi. 2016. Guarded Cubical Type Theory: Path Equality for Guarded Recursion. Schloss-Dagstuhl - Leibniz Zentrum für Informatik.
- Edwin Brady. 2017. *Type-Driven Development with Idris*. Manning Publications.
- Venanzio Capretta. 2005. General Recursion via Coinductive Types. *Logical Methods in Computer Science* 1, 2 (2005).
- Adam Chlipala. 2008. Parametric Higher-Order Abstract Syntax for Mechanized Semantics. In *Proceedings of the International Conference on Functional Programming*.
- Nils Anders Danielsson. 2012. Operational Semantics Using the Partiality Monad. In *Proceedings of the International Conference on Functional Programming*.
- Nils Anders Danielsson and Thorsten Altenkirch. 2010. Subtyping, Declaratively. In *Proceedings of the International Conference on Mathematics of Program Construction*. Vol. 6120. Lecture Notes in Computer Science.
- N.G de Bruijn. 1972. Lambda Calculus Notation with Nameless Dummies, A Tool for Automatic Formula Manipulation, with Application to the Church-Rosser Theorem. *Indagationes Mathematicae* 75, 5 (1972).
- Jeremy Gibbons. 2021. Continuation-Passing Style, Defunctionalization, Accumulations, and Associativity. *The Art, Science, and Engineering of Programming* 6, 2 (2021).
- John Hughes, Lars Pareto, and Amr Sabry. 1996. Proving the Correctness of Reactive Systems Using Sized Types. In *Proceedings of the 23rd Symposium on Principles of Programming Languages*.
- Graham Hutton and Patrick Bahr. 2017. Compiling a 50-Year Journey. *Journal of Functional Programming* 27 (2017).
- Graham Hutton and Joel Wright. 2004. Compiling Exceptions Correctly. In *Proceedings of the International Conference on Mathematics of Program Construction (Lecture Notes in Computer Science, Vol. 3125)*.
- Yugo Kashiwagi and David Wise. 1991. *Graph Algorithms in a Lazy Functional Programming Language*.
- David King and John Launchbury. 1995. Structuring Depth-First Search Algorithms in Haskell. In *Proceedings of the 22nd Symposium on Principles of Programming Languages*.
- Ramana Kumar, Magnus O. Myreen, Michael Norrish, and Scott Owens. 2014. CakeML: A Verified Implementation of ML. *ACM SIGPLAN Notices* 49, 1 (2014).
- Xavier Leroy. 2009. Formal Verification of a Realistic Compiler. *Commun. ACM* 52, 7 (2009).
- Martin Erwig. 2001. Inductive graphs and functional graph algorithms. *Journal of Functional Programming* 11, 5 (2001).
- John McCarthy and James Painter. 1967. Correctness of a compiler for arithmetic expressions. In *Mathematical Aspects of Computer Science (Proceedings of Symposia in Applied Mathematics, Vol. 19)*. American Mathematical Society, 33–41.
- Erik Meijer. 1992. *Calculating Compilers*. PhD Thesis. Katholieke Universiteit Nijmegen.
- Andrey Mokhov. 2017. Algebraic Graphs with Class (Functional Pearl). In *Proceedings of the 10th Symposium on Haskell*.
- Rasmus Ejlers Møgelberg and Niccolò Veltri. 2019. Bisimulation As Path Type for Guarded Recursive Types. *Proceedings of the ACM on Programming Languages* 3, POPL (2019).
- Hiroshi Nakano. 2000. A Modality for Recursion. In *Proceedings of the 15th Annual IEEE Symposium on Logic in Computer Science*.
- Ulf Norell. 2007. *Towards a Practical Programming Language Based on Dependent Type Theory*. PhD Thesis. Chalmers University of Technology.
- Bruno Oliveira and William Cook. 2012. Functional Programming with Structured Graphs. In *Proceedings of the International Conference on Functional Programming*.
- Mitchell Pickard and Graham Hutton. 2021. Calculating Dependently-Typed Compilers. *Proceedings of the ACM on Programming Languages* 5, ICFP, Article 82 (2021).
- Norman Ramsey and João Dias. 2006. An Applicative Control-Flow Graph Based on Huet’s Zipper. In *Proceedings of the Workshop on ML*.

- Norman Ramsey, João Dias, and Simon Peyton Jones. 2010. Hoopl: A Modular, Reusable Library for Dataflow Analysis and Transformation. In *Proceedings of the Third Symposium on Haskell*.
- Mitchell Wand. 1982. Deriving Target Code as a Representation of Continuation Semantics. *ACM Transactions on Programming Languages and Systems* 4, 3 (1982).
- Mitchell Wand. 1995. Compiler Correctness for Parallel Languages. In *Proceedings of the Seventh International Conference on Functional Programming Languages and Computer Architecture*.