



CF-GKAT: Efficient Validation of Control-Flow Transformations

CHENG ZHANG*, University College London, United Kingdom

TOBIAS KAPPÉ†, LIACS, Leiden University, Netherlands

DAVID E. NARVÁEZ, Virginia Tech, USA

NICO NAUS, Open University of the Netherlands, Netherlands and Virginia Tech, USA

Guarded Kleene Algebra with Tests (GKAT) provides a sound and complete framework to reason about trace equivalence between simple imperative programs. However, there are still several notable limitations. First, GKAT is completely agnostic with respect to the meaning of primitives, to keep equivalence decidable. Second, GKAT excludes non-local control flow such as goto, break, and return. To overcome these limitations, we introduce *Control-Flow GKAT (CF-GKAT)*, a system that allows reasoning about programs that include non-local control flow as well as hardcoded values. CF-GKAT is able to soundly and completely verify trace equivalence of a larger class of programs, while preserving the nearly-linear efficiency of GKAT. This makes CF-GKAT suitable for the verification of control-flow manipulating procedures, such as decompilation and goto-elimination. To demonstrate CF-GKAT's abilities, we validated the output of several highly non-trivial program transformations, such as Erosa and Hendren's goto-elimination procedure and the output of Ghidra decompiler. CF-GKAT opens up the application of Kleene Algebra to a wider set of challenges, and provides an important verification tool that can be applied to the field of decompilation and control-flow transformation.

CCS Concepts: • Theory of computation → Automated reasoning; Logic and verification; Algebraic semantics; • Software and its engineering → Correctness.

Additional Key Words and Phrases: Program equivalence, Kleene algebra, control flow recovery

ACM Reference Format:

Cheng Zhang, Tobias Kappé, David E. Narváez, and Nico Naus. 2025. CF-GKAT: Efficient Validation of Control-Flow Transformations. *Proc. ACM Program. Lang.* 9, POPL, Article 21 (January 2025), 27 pages. <https://doi.org/10.1145/3704857>

1 Introduction

There are many notions of program equivalence with different granularity and computational complexity. At one end of the spectrum, syntactic equality compares two programs solely based on its syntax. Although decidable, this technique will not equate programs like the following:

$\text{if } t \text{ then } p \text{ else } q$

$\text{if } \neg t \text{ then } q \text{ else } p$

Conversely, input-output equivalence relates two programs if and only if they yield the same output on the same input. This equivalence is very powerful, but also well-known to be undecidable.

*Work performed at Boston University.

†Work performed at Open University of the Netherlands and ILLC, University of Amsterdam.

Authors' Contact Information: [Cheng Zhang](#), University College London, London, United Kingdom, czhang03@bu.edu; [Tobias Kappé](#), LIACS, Leiden University, Leiden, Netherlands, t.w.j.kappe@liacs.leidenuniv.nl; [David E. Narváez](#), Virginia Tech, Blacksburg, USA, david.narvaez@vt.edu; [Nico Naus](#), Open University of the Netherlands, Heerlen, Netherlands and Virginia Tech, Blacksburg, USA, nico.naus@ou.nl.



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

© 2025 Copyright held by the owner/author(s).

ACM 2475-1421/2025/1-ART21

<https://doi.org/10.1145/3704857>

Situated between these two extremes, *Guarded Kleene Algebra with Tests* [28, 41], or GKAT for short, reasons about *trace equivalence* between simple while-programs. By abstracting the meaning of the primitive tests and actions, it can focus on how tests determine which actions are performed. For example, the two programs above are equivalent in GKAT: they both execute p when t holds, and q when it does not. GKAT is also able to validate nontrivial equivalences like

$$\text{while } t \text{ do } p; \text{ while } s \text{ do } \{ q; \text{ while } t \text{ do } p \} \equiv \text{while } t \vee s \text{ do } \{ \text{if } t \text{ then } p \text{ else } q \}$$

Surprisingly, GKAT equivalence is decidable in nearly-linear time (assuming the set of test variables is fixed) [41], making it a reasonable compromise between complexity and granularity.

Nevertheless, the abstraction of GKAT comes at the price of not being able to validate some straightforward equivalences. First, as mentioned, GKAT disregards the meaning of primitive programs and tests. For instance, when given a program like

$$\text{if } y \neq 0 \text{ then } \{ x := 42; p \} \text{ else } \{ x := 42; q \}, \quad (1)$$

we can note that a change in the value of x does not have effect on whether $y \neq 0$. Hence, it should be possible to factor the assignment to x out of the branches, and obtain

$$x := 42; \text{ if } y \neq 0 \text{ then } p \text{ else } q. \quad (2)$$

GKAT does not admit this equivalence, precisely because it is agnostic with respect to the meaning of primitive actions. However, moving to a setting that accounts for the semantics of actions is hard, because Turing completeness – and by extension, undecidability – lurks nearby [4, 22, 29].

Second, GKAT excludes non-local control-flow constructs like *goto*, *break*, and *return*. In a general imperative language, this does not limit expressivity, as these constructs can be recovered using variables [16] – indeed, the Böhm-Jacopini theorem says that every type of control flow can be written using a single while-loop [7]. However, lacking both variables and non-local control structures, GKAT is not able to express all control flow in real-world programs.

As a concrete example, consider the programs below. The control flow in Program 3 is based purely on labels and *goto*. Meanwhile, Program 4 is structured as a loop with the option to terminate early using *break*. These programs happen to be trace equivalent (i.e., they always execute the same actions in the same order) but represent behavior not expressible in plain GKAT [39].

$$\text{label } \ell_0; \text{ if } \neg t \text{ then goto } \ell_1; p; \text{ if } t \text{ then goto } \ell_1; q; \text{ goto } \ell_0; \text{ label } \ell_1 \quad (3)$$

$$\text{while } t \text{ do } \{ p; \text{ if } \neg t \text{ then } q \text{ else break } \} \quad (4)$$

As it turns out, deciding equivalence between programs such as these is essential in validating control-flow manipulation procedures. Specifically, consider the control flow structuring phase of a decompiler [12], which is tasked with converting conditional and unconditional jumps into more conventional control flow constructs. Program 3 can be thought of as pseudo-assembly that models the input of this process, and Program 4 is a plausible outcome of decompilation. Thus, the control-flow structuring process is correct when Programs 3 and 4 are equivalent.

To overcome the limitations of GKAT, we propose control-flow GKAT (CF-GKAT), an extension that is capable of equating some interesting programs. Concretely, we extend GKAT in two ways.

First, we add *indicator variables*, which can be assigned and tested against hardcoded values, and do not appear in other primitive actions and tests. For example, assignments like $x := 42$ are allowed, but assignments like $x := y + 1$ are not. This strikes a delicate balance, by being weak enough to exclude general computation (thus keeping equivalence decidable), yet strong enough to model the equivalence of Programs 1 and 2. The addition of indicator variables empowers CF-GKAT to validate control-flow transformation algorithms that use them [16, 49]. We focus on a single indicator variable for brevity; an extension to multiple indicator variables should be straightforward.

Second, we extend GKAT with the non-local control-flow constructs `goto`, `break` and `return`. This poses a challenge, as the non-local nature of these commands prevents a compositional semantics — after all, the precise meaning of a statement like `break` depends on its context. To overcome this, we propose an intermediate *continuation semantics*. In this semantics, each trace is tagged with a “continuation” that can accept (terminate normally), `break`, `return`, or go to a label. Then, the trace semantics of the program can be obtained by resolving these continuations.

We were able to design an automaton model for CF-GKAT, where every CF-GKAT expression can be converted into a CF-GKAT automaton while preserving the continuation semantics. Furthermore, CF-GKAT automata and continuation semantics can be lowered into GKAT automata and trace semantics respectively, while preserving their semantic correspondence. We are thus able to reduce the problem of deciding trace equivalence of programs to deciding bisimilarity of two GKAT automata, which can be done efficiently. Consequently, CF-GKAT can soundly and completely validate trace equivalence of a larger class of programs, while preserving the nearly-linear efficiency of GKAT. For instance, it can automatically validate that Programs 3 and 4 are equivalent to each other, and also to their single-loop equivalent obtained via the Böhm-Jacopini theorem [7]:

```


$$\begin{aligned}
x := 1; \text{while } x \neq 0 \text{ do } \{ \\
\quad \text{if } x = 1 \wedge t \text{ then } \{ p; x := 2 \} \\
\quad \text{else if } x = 2 \wedge \neg t \text{ then } \{ q; x := 1 \} \\
\quad \text{else } x := 0 \}
\end{aligned} \tag{5}$$


```

To put this theory to work, we implemented an equivalence checker for CF-GKAT, with a front-end that can convert C code to CF-GKAT expressions by leveraging Clang’s parser. The resulting tool is able to automatically validate highly non-trivial program transformations like Erosa and Hendren’s classic `goto`-elimination procedure [16] and the control flow structuring of Ghidra.¹

Outline. The remainder of this article is organized as follows. In Section 2, we give an overview of CF-GKAT, including its syntax, continuation semantics, and trace semantics. In Section 3, we propose an automaton model for the continuation semantics, called CF-GKAT automata; we show how to translate a CF-GKAT expression to a CF-GKAT automaton, which in turn can be lowered to a GKAT automaton; ultimately, this gives rise to an algorithm for checking trace equivalence of CF-GKAT expressions. In Section 4, we report on an implementation of our algorithm, along with several experiments. We discuss related work in Section 5, and conclude in Section 6.

Our proofs are formalized [51] in Coq [14]; we provide proof sketches for the sake of intuition.

2 Overview

In this section, we introduce the language of CF-GKAT, and gradually develop its semantics. We begin by explaining the syntax of CF-GKAT; after that, we delve into the semantics of its tests. We then introduce the intermediate semantic model of (*labeled and indexed families of*) *guarded languages with continuations*, which is flattened into a model based on *guarded languages*. Having defined these tools, we then conclude by giving a semantics to CF-GKAT programs in this model. Along the way, we will single out and explain some of the finer points using examples.

2.1 Syntax

The syntax of CF-GKAT consists of two levels, similar to GKAT. At the bottom level, there are *tests*. These are Boolean assertions that can occur as guards inside conditional statements, or within

¹<https://ghidra-sre.org/>

assertions that occur in the program text. To model them, we fix a finite set of primitive tests T , which represent uninterpreted expressions that may or may not hold. The full syntax is as follows.

$$\text{BExp} \ni b, c ::= \text{false} \mid \text{true} \mid t \in T \mid x = i \ (i \in I) \mid b \vee c \mid b \wedge c \mid \neg b$$

Compared to GKAT, tests in CF-GKAT include the *indicator variable test* $x = i$ (highlighted in blue) for each *indicator value* i drawn from a finite but fixed set of possible indicator values I . As the notation suggests, this test holds when the indicator variable x currently has the value i .

The top level syntax of GKAT is built using a finite set of uninterpreted commands Σ (the *primitive actions*), as well as *assertions* of the form $\text{assert } b$, where $b \in \text{BExp}$ is a test. Expressions are composed using sequencing, if statements, and while loops. CF-GKAT extends the base elements of the syntax with indicator variable assignments $x := i$ (for each $i \in I$), which changes the value of the indicator variable x to i . In addition, it adds the non-local control flow commands `break` and `return`, as well as `goto` ℓ and `label` ℓ , where ℓ is taken from a fixed but finite set of labels L . The full syntax is given below (additions relative to GKAT highlighted in blue again).

$$\begin{aligned} \text{Exp} \ni e, f ::= & \text{assert } b \mid p \in \Sigma \mid x := i \ (i \in I) \mid e; f \mid \text{if } b \text{ then } e \text{ else } f \mid \\ & \text{while } b \text{ do } e \mid \text{break} \mid \text{return} \mid \text{goto } \ell \ (\ell \in L) \mid \text{label } \ell \ (\ell \in L) \end{aligned}$$

A *valid program*, or *program* for short, is an expression without (1) duplicate labels, (2) `goto` commands with an undefined label, or (3) `break` statements that occur outside a loop. For the sake of simplicity, we assume that the reader does not require a more formal definition of this notion.

Example 2.1. Any GKAT expression is a valid program. Also, programs 3 to 5 from the introduction are all valid CF-GKAT programs. The following expressions are *not* valid programs:

$$\begin{array}{ll} \text{label } \ell; (\text{if } t \text{ then label } \ell; p \text{ else } q) & (\text{the label } \ell \text{ is defined twice}) \\ (\text{while true do } p); \text{goto } \ell & (\text{the label } \ell \text{ is undefined}) \\ \text{if } t \text{ then break else } p & (\text{break appears outside a loop}) \end{array}$$

2.2 Boolean Semantics

To assign a semantics to CF-GKAT expressions, we first need to talk about the semantics of tests. Intuitively, each test in a CF-GKAT expression symbolically denotes a set of execution contexts in which it is true. But how do we model an execution context? Because the primitive tests from T are uninterpreted, we represent them by simply listing the ones that are true; the unlisted tests are then assumed to be false. This is in line with how the semantics of (G)KAT handles tests [27, 41]. As for the indicator tests, we include the current value of the indicator variable x in the execution context. This means that each execution context is of the form (i, α) , for $i \in I$ and $\alpha \subseteq T$.

Putting these ideas together, we can calculate the set of execution contexts that satisfy a given test $b \in \text{BExp}$ by induction. This set will be regarded as the semantics of b .

Definition 2.2. Let At denote 2^T , the set of *atoms* of (the free Boolean algebra generated by) T . We define the *Boolean semantics* function $\llbracket - \rrbracket : \text{BExp} \rightarrow 2^{I \times \text{At}}$ inductively, as follows.

$$\begin{array}{lll} \llbracket \text{false} \rrbracket \triangleq \emptyset & \llbracket t \rrbracket \triangleq \{(i, \alpha) \mid i \in I, t \in \alpha\} & \llbracket b \vee c \rrbracket \triangleq \llbracket b \rrbracket \cup \llbracket c \rrbracket \\ \llbracket \text{true} \rrbracket \triangleq I \times \text{At} & \llbracket x = i \rrbracket \triangleq \{(i, \alpha) \mid \alpha \in \text{At}\} & \llbracket b \wedge c \rrbracket \triangleq \llbracket b \rrbracket \cap \llbracket c \rrbracket \\ & & \llbracket \neg b \rrbracket \triangleq I \times \text{At} \setminus \llbracket b \rrbracket \end{array}$$

Example 2.3. Take $T = \{t_1, t_2\}$ and $I = \{1, 2, 3\}$; then we can calculate that

$$\llbracket (t_1 \vee \neg t_2) \wedge (x = 2) \rrbracket = \{(\{t_1, t_2\}, 2), (\{t_1\}, 2), (\emptyset, 2)\}$$

In other words, the test above holds in execution contexts where the indicator has value 2, and either t_1 and t_2 are both true (first element), both false (last element), or where t_1 is true but t_2 is false (middle element). In contrast, $\llbracket x = 1 \wedge x = 3 \rrbracket = \emptyset$, which is to say that this test does not hold in any execution context, because the indicator variable x cannot be both 1 and 3 at the same time.

2.3 Guarded Language with Continuations

We can now turn our attention to the semantics of CF-GKAT. Like (G)KAT, the semantics of CF-GKAT is given in terms of *guarded languages* [27, 41], which are best thought of as sets of symbolic traces of the program. These traces record the initial, intermediate and final machine states observed during execution, as well as the actions that occurred between those states. Because the value of the indicator variable matters only for control flow, we do not consider indicators to be part of the machine state; hence, machine states in a guarded word are drawn from At .

Definition 2.4. A *guarded word* is a sequence of the form $\alpha_1 p_1 \alpha_2 p_2 \cdots \alpha_{n-1} p_n \alpha_n$, where $\alpha_i \in \text{At}$ and $p_i \in \Sigma$; that is to say, guarded words are elements of the regular language $\text{At} \cdot (\Sigma \cdot \text{At})^*$. We refer to sets of guarded words as *guarded languages*; the set of guarded languages is denoted \mathcal{G} .

Example 2.5. Let $T = \{t_1, t_2\}$ and $\Sigma = \{p_1, p_2\}$. Now the guarded word $\{t_1\} p_1 \{t_2\} p_2 \emptyset$ represents a program trace that starts out in a machine state where t_1 is true (but t_2 is not). The program then executes the action p_1 , after which t_2 is true (but t_1 is not). Finally, the program goes on to execute the action p_2 , and halts in a state where neither t_1 nor t_2 is true.

Our semantics of a CF-GKAT expression will ultimately be a guarded language. However, due to the non-local nature of our language, we cannot compute this semantics inductively, as the label in `goto` may occur in a different part of the program whose traces have not yet been computed. To address this challenge, we introduce an intermediate *continuation semantics*, which accounts for both the indicator variables and the non-local control flow. The remainder of this subsection explains the domain of continuation semantics, based on *guarded words with continuations*. Intuitively, these are guarded words equipped with a piece of information called a *continuation*, which indicates how control flow continues after the program ends. The inclusion continuation information at the end of a trace allows us to define a semantics of CF-GKAT expressions inductively. Once the continuation semantics of a CF-GKAT program is known, we can flatten it into a guarded language.

Definition 2.6. A *guarded word with continuation* is a pair $w \cdot c$, where w is a guarded word and c is a *continuation*, which can take on one of the following forms for $i \in I$ and $\ell \in L$:

| | | | |
|----------------------|----------------------|------------------|------------------------------|
| <code>acc</code> i | <code>brk</code> i | <code>ret</code> | <code>jmp</code> (ℓ, i) |
|----------------------|----------------------|------------------|------------------------------|

We write C for the set of all continuations. A set of guarded words with continuations is a *guarded language with continuations*; the set of guarded languages with continuations is written \mathcal{C} .

Intuitively, the different types of continuation may be interpreted as follows:

- The continuation `acc` i represents that the trace has successfully reached the end of this part of the program, with indicator value i . Execution can be picked up if the program is put in a larger context – e.g., if $w \cdot \text{acc } i$ is a trace of e , then it may be combined with a trace found when f is executed with indicator value i to compute the semantics of $e; f$.
- A continuation of the form `brk` i signals that the trace ends by halting the loop in which it occurs. Execution can resume only after this loop (with indicator value i). This kind of trace cannot be composed on the right, as is done for traces with accepting continuations, because we first need to enclose it in a loop to halt; it will then be converted into `acc` i .

- The continuation **ret** represents a trace that ends in the program halting completely. Traces of this kind will percolate upwards in the semantics, without changing their continuation. These are intended to model the **return** statement, which halts the program no matter how deeply it is nested. In this case, the indicator value does not matter any more.
- Finally, the continuation **jmp** (ℓ, i) is put on traces that will continue executing from label ℓ , with indicator value i . Like **brk** i and **ret**, these traces do not compose on the right, but unlike **brk** i this continuation does not change, as jump resolution happens only at the end, when the continuation semantics is known for the entire program.

Example 2.7. Let w be the guarded word from the previous example; the guarded word with continuation $w \cdot \text{jmp} (\ell_1, 2)$ represents a partial program trace that takes the steps in w , continues executing at the label ℓ_1 with indicator value 2. More examples appear in the next section.

2.4 Indexed Families and Sequencing

The continuation semantics of a CF-GKAT expression takes a starting indicator value, and produces a guarded language with continuations representing the traces of that program when started with the given indicator value. To properly encode this, we need the following notion.

Definition 2.8. An *indexed family* of guarded languages (with continuations), or “indexed family” for the sake of brevity, is function from I to guarded languages (with continuations). Typically, we use G and H to denote an indexed family. To lighten notation, we write G_i to denote $G(i)$.

Similar to guarded languages, indexed families can be composed in several ways. In particular, we are interested in the sequencing operation and the Kleene star operation of indexed families, because these will turn out to be useful when defining the continuation semantics of CF-GKAT.

When sequencing two families G and H , the traces in G_i with a continuation of the form **acc** j will be composed with traces in H_j ; traces with different continuations are copied over in full, because they do not compose on the right. Formally, this operation is defined as follows.

Definition 2.9. Let $G, H : I \rightarrow \mathcal{C}$. We write $G \diamond H$ for the *sequencing* (or *concatenation*) operation of G and H , which is defined as the smallest family of guarded languages with continuations (in the pointwise order) satisfying the following rules for all $i, j \in I$ as well as all $\ell \in L$:

$$\frac{w\alpha \cdot \text{acc} j \in G_i \quad \alpha x \cdot c \in H_j}{w\alpha x \cdot c \in (G \diamond H)_i} \quad \frac{w \cdot c \in G_j \quad c \in \{\text{brk } i, \text{acc } i, \text{jmp } (\ell, i)\}}{w \cdot c \in (G \diamond H)_j}$$

The first rule composes accepting traces in G with traces in H , picking up with the indicator value where the first trace left off. Note also that this rule requires the last atom in the trace on the left to match the first atom in the trace on the right, because we want the second trace to start from the machine state computed in the first trace. This mirrors the *coalesced product* used to define the sequential composition of guarded languages (without continuations) [27]. The last rule ensures that traces that encountered non-local control flow within G are preserved in $G \diamond H$.

Example 2.10. Let $I = \{1, 2\}$, and let G and H be indexed families given by:

$$\begin{array}{ll} G_1 = \{\alpha p \beta \cdot \text{brk } 1, \beta p \alpha \cdot \text{acc } 2\} & G_2 = \{\alpha q \beta \cdot \text{acc } 1\} \\ H_1 = \{\gamma q \beta \cdot \text{ret}\} & H_2 = \{\alpha r \beta \cdot \text{jmp } (\ell_1, 1)\} \end{array}$$

Then we can compute that the sequencing $G \diamond H$ is the following indexed family:

$$(G \diamond H)_1 = \{\alpha p \beta \cdot \text{brk } 1, \beta p \alpha r \beta \cdot \text{jmp } (\ell_1, 1)\} \quad (G \diamond H)_2 = \emptyset$$

Here, $(G \diamond H)_1$ contains $\alpha p \beta \cdot \text{brk } 1$ by the second rule, because G_1 does. Furthermore, the trace $\beta p \alpha \cdot \text{acc } 2$ in G_1 is composed with $\alpha r \beta \cdot \text{jmp } (\ell_1, 1)$ from H_2 to form $\beta p \alpha r \beta \cdot \text{jmp } (\ell_1, 1)$ in $(G \diamond H)_1$.

by the first rule. The set $(G \diamond H)_2$ is empty, because despite the fact that $\alpha q \beta \cdot \mathbf{acc} 1 \in G_2$, there is no trace in H_1 that starts with β , and so neither rule can apply.

2.5 Continuation Semantics

With the theory of indexed families in place, we can now define the continuation semantics $\llbracket e \rrbracket^\sharp$ of a CF-GKAT program e in terms of an indexed family. We start with the base cases.

Definition 2.11 (Continuation semantics, base). For all $i, j \in I$, we define the following sets:

$$\begin{aligned} \llbracket \mathbf{assert} b \rrbracket_i^\sharp &\triangleq \{\alpha \cdot \mathbf{acc} i \mid (i, \alpha) \in \llbracket b \rrbracket\} & \llbracket \mathbf{goto} \ell \rrbracket_i^\sharp &\triangleq \{\alpha \cdot \mathbf{jmp} (\ell, i) \mid \alpha \in \text{At}\} \\ \llbracket p \rrbracket_i^\sharp &\triangleq \{\alpha p \beta \cdot \mathbf{acc} i \mid \alpha, \beta \in \text{At}\} & \llbracket \mathbf{label} \ell \rrbracket_i^\sharp &\triangleq \{\alpha \cdot \mathbf{acc} i \mid \alpha \in \text{At}\} \\ \llbracket x := j \rrbracket_i^\sharp &\triangleq \{\alpha \cdot \mathbf{acc} j \mid \alpha \in \text{At}\} & \llbracket \mathbf{break} \rrbracket_i^\sharp &\triangleq \{\alpha \cdot \mathbf{brk} i \mid \alpha \in \text{At}\} \\ \llbracket \mathbf{return} \rrbracket_i^\sharp &\triangleq \{\alpha \cdot \mathbf{ret} \mid \alpha \in \text{At}\} \end{aligned}$$

Each of these base syntax elements yields a finite indexed family. For the constructs `return`, `goto`, and `break`, all traces terminate immediately in the corresponding continuation.

We inherit the semantics of assertions and primitive actions from (G)KAT [27, 41]. Assertions have traces that accept when their only atom satisfies the test. A primitive action p has traces of the form $\alpha p \beta \cdot \mathbf{acc} i$ for all $\alpha, \beta \in \text{At}$: because p is uninterpreted, we can reach any machine state by running p . Only the indicator value is retained, because primitive actions cannot interact with indicators. Assignments like $x := j$ accept immediately, without changing the machine state; however, each trace ends with indicator value j – regardless of the initial indicator value i .

Finally, labels are encoded as no-operations, which makes them neutral for sequencing operator, i.e., we have $\llbracket \mathbf{label} \ell \rrbracket^\sharp \diamond G = G = \llbracket \mathbf{label} \ell \rrbracket^\sharp \diamond G$ for all indexed families G . This is because labels serve only as potential starting points of execution; we will leverage them in the next subsection.

Remark. For soundness, it is important that the indicator variable x does not occur in any primitive test $t \in T$ or action $p \in \Sigma$. Concretely, primitive tests do not restrict the current value indicator variable, and primitive actions preserve the indicator value after their executions.

We now turn our attention to the program composition operators. These are generalizations of the guarded language semantics of GKAT [41]. First of all, the `if b then e else f` filters out traces in the semantics of the e that satisfy the guard b , as well as the traces in f that invalidate it.

Definition 2.12 (Continuation semantics, branching). Let $e, f \in \text{Exp}$. We set $\llbracket \text{if } b \text{ then } e \text{ else } f \rrbracket_i^\sharp$ to be the least indexed family that satisfies the following rules for all $i \in I$:

$$\frac{(i, \alpha) \in \llbracket b \rrbracket \quad \alpha w \cdot c \in \llbracket e \rrbracket_i^\sharp}{\alpha w \cdot c \in \llbracket \text{if } b \text{ then } e \text{ else } f \rrbracket_i^\sharp} \quad \frac{(i, \alpha) \notin \llbracket b \rrbracket \quad \alpha w \cdot c \in \llbracket f \rrbracket_i^\sharp}{\alpha w \cdot c \in \llbracket \text{if } b \text{ then } e \text{ else } f \rrbracket_i^\sharp}$$

The semantics of the sequencing operator is easy: it just composes the semantics of the operands with the sequencing operator we have for indexed families. For loops, some more care is needed because traces can be iterated, and we need to account for early termination.

Definition 2.13 (Continuation semantics, sequencing and loops). Let $e, f \in \text{Exp}$ and $b \in \text{BExp}$. We define $\llbracket e; f \rrbracket^\sharp \triangleq \llbracket e \rrbracket^\sharp \diamond \llbracket f \rrbracket^\sharp$. Also, $\llbracket \text{while } b \text{ do } e \rrbracket_i^\sharp$ is the least indexed family satisfying for $i \in I$:

$$\frac{(i, \alpha) \notin \llbracket b \rrbracket}{\alpha \cdot \mathbf{acc} i \in \llbracket \text{while } b \text{ do } e \rrbracket_i^\sharp} \quad \frac{(i, \alpha) \in \llbracket b \rrbracket \quad \alpha w \cdot c \in (\llbracket e \rrbracket^\sharp \diamond \llbracket \text{while } b \text{ do } e \rrbracket_i^\sharp)_i}{\alpha w \cdot \lfloor c \rfloor \in \llbracket \text{while } b \text{ do } e \rrbracket_i^\sharp}$$

The operation $\lfloor - \rfloor$ in the last rule is defined by $\lfloor c \rfloor = \mathbf{acc} i$ when $c = \mathbf{brk} i$, and $\lfloor c \rfloor = c$ otherwise.

The first rule accounts for traces that halt immediately because the loop guard is false. The second rule allows prepending traces from the loop body that satisfy the guard. Because of the way sequencing works, body traces that end in `brk i` may occur; the second rule converts their continuations to `acc i`, signaling that the loop has been exited and normal control flow can resume.

The semantics we have so far defines the traces of a program starting from the beginning. However, a CF-GKAT program can also resume from a label when a `goto` command is encountered. To obtain the traces for a given label ℓ , we must descend into the program until we encounter label ℓ . For the base cases, this is relatively simple to accomplish: just check if we start at the label.

Definition 2.14 (Continuation semantics starting from a label, base). Let $e \in \text{Exp}$. For each $\ell \in L$, we define the following guarded languages with continuations:

$$\begin{aligned} \llbracket \text{assert } b \rrbracket_i^\ell &\triangleq \emptyset & \llbracket \text{goto } \ell' \rrbracket_i^\ell &\triangleq \emptyset & \llbracket p \rrbracket_i^\ell &\triangleq \emptyset & \llbracket x := j \rrbracket_i^\ell &\triangleq \emptyset & \llbracket \text{break} \rrbracket_i^\ell &\triangleq \emptyset \\ \llbracket \text{return} \rrbracket_i^\ell &\triangleq \emptyset & \llbracket \text{label } \ell' \rrbracket_i^\ell &\triangleq \{\alpha \cdot \text{acc } i \mid \alpha \in \text{At}, \ell = \ell'\} \end{aligned}$$

Note how none of these cases has a trace, except the one for $\llbracket \text{label } \ell' \rrbracket_i^\ell$ when $\ell' = \ell$, which accepts immediately. With these cases covered, we can then treat the inductive step.

Definition 2.15 (Continuation semantics starting from a label, sequencing and branching). Let $e, f \in \text{Exp}$, $b \in \text{BExp}$ and $\ell \in L$. We define the following indexed families to cover the traces of CF-GKAT programs starting from the label ℓ when composed using branching or sequencing:

$$\llbracket \text{if } b \text{ then } e \text{ else } f \rrbracket_i^\ell \triangleq \llbracket e \rrbracket_i^\ell \cup \llbracket f \rrbracket_i^\ell \quad \llbracket e; f \rrbracket_i^\ell \triangleq (\llbracket e \rrbracket_i^\ell \diamond \llbracket f \rrbracket_i^\#)_i \cup \llbracket f \rrbracket_i^\ell$$

For branching, the semantics starting from ℓ disregards the guard and descends into the operands. The sequencing case is more interesting: here, we still need to account for the traces that start from the beginning of f after executing a trace in e starting from the label ℓ .

The only remaining case to cover is the loop. If we start execution from a label somewhere in the loop body, we may need to start the loop again after completing the loop body. On the other hand, early termination in the loop body still needs to be turned into an accepting trace.

Definition 2.16 (Continuation semantics starting from a label, loops). Let $e \in \text{Exp}$ and $b \in \text{BExp}$. We define the indexed family $\llbracket \text{while } b \text{ do } e \rrbracket_i^\ell$ below, where $\llbracket - \rrbracket$ is as in Definition 2.13:

$$\llbracket \text{while } b \text{ do } e \rrbracket_i^\ell = \{w \cdot \llbracket c \rrbracket \mid w \cdot c \in (\llbracket e \rrbracket_i^\ell \diamond \llbracket \text{while } b \text{ do } e \rrbracket_i^\#)_i\}$$

2.6 Guarded Language Semantics

The continuation semantics $\llbracket e \rrbracket_i^\#$ of a CF-GKAT program uses continuations to record how a trace ends. In particular, some traces may end in a continuation of the form `jmp` (ℓ, i) , signaling that computation needs to continue from the label ℓ . The semantics $\llbracket e \rrbracket_i^\ell$ provides the necessary information to resolve such jump continuations: we can resume `jmp` (ℓ, i) by concatenating with the traces in $\llbracket e \rrbracket_i^\ell$. We will end this section by doing just that.

To formalize our approach, we need a way to refer to the continuation CF-GKAT semantics of a program as a whole, i.e., for all indicator values, starting from either the beginning or some label.

Definition 2.17. A *labeled family of guarded languages (with continuations)*, or *labeled family* for short, is a function G from $L + \#$ to indexed families of guarded languages (with continuations), e.g., $G : L + \# \rightarrow I \rightarrow C$. We often use superscripts to denote the value at a given label $\#$, writing G^ℓ for $G(\ell)$. Note that under this convention, G^ℓ is an indexed family, which means that we may further unravel by writing G_i^ℓ to obtain the guarded language with continuations $G(\ell, i)$.

Crucially, we can retrofit the continuation semantics $\llbracket e \rrbracket$ as a labeled family; after all, $\llbracket e \rrbracket^\#$ is an indexed family, and so is $\llbracket e \rrbracket^\ell$ for each $\ell \in L$. We will treat $\llbracket e \rrbracket$ as such from this point on.

To resolve the jumps in a labeled family of guarded languages with continuations, we resolve the traces ending in $\text{jmp } (\ell, i)$ by looking up the traces that originate from label ℓ with indicator value i . We also remove the continuations $\text{acc } i$ and ret , because those come with traces that either reached the end of the program, or encountered a `return` statement respectively. Continuations of the form $\text{brk } i$ should not occur at the top level when computing the semantics of a program, so we can ignore them. The result is a labeled family of guarded languages (without continuations).

Definition 2.18. Let $G : L + \sharp \rightarrow I \rightarrow C$ be a labeled family of guarded languages with continuations. We write $G \downarrow$ for the (point-wise) least labeled family of guarded languages $G \downarrow$ such that the following rules are satisfied for all $k \in L + \sharp$, $\ell \in L$, and $i, j \in I$:

$$\frac{w \cdot \text{acc } i \in G_i^k}{w \in G \downarrow_i^k} \qquad \frac{w \cdot \text{ret} \in G_i^k}{w \in G \downarrow_i^k} \qquad \frac{w\alpha \cdot \text{jmp } (\ell, j) \in G_i^k \quad \alpha x \in G \downarrow_j^\ell}{w\alpha x \in G \downarrow_i^k}$$

The first two rules take care of flattening guarded words with continuations that end in acceptance. The third rule strings together guarded words continuations that jump to a different label.

Example 2.19. Let G be the labeled family of guarded languages with continuations defined by

$$\begin{array}{ll}
 G_1^\sharp = \{\alpha \cdot \mathbf{jmp}(\ell, 1)\} & G_2^\sharp = \emptyset \\
 G_1^\ell = \{\alpha p \alpha \cdot \mathbf{jmp}(\ell', 1), \beta \cdot \mathbf{acc} 1\} & G_2^\ell = \{\alpha \cdot \mathbf{jmp}(\ell', 2)\} \\
 G_1^{\ell'} = \{\alpha q \alpha \cdot \mathbf{jmp}(\ell, 1), \alpha r \beta \cdot \mathbf{jmp}(\ell, 1)\} & G_2^{\ell'} = \{\alpha \cdot \mathbf{jmp}(\ell, 1)\}
 \end{array}$$

Now $G \downarrow_1^\sharp$ contains, among other things, the guarded word $\alpha p \alpha q \alpha p \alpha r \beta$. Furthermore, $G \downarrow_2^\ell$ is empty, despite G_2^ℓ and $G_2^{\ell'}$ containing guarded words with mutually jumping continuations, as these can never be concatenated into one guarded word with continuation of the form $\text{acc } i$ or ret .

In total, we can then obtain the semantics of a CF-GKAT term as $\llbracket e \rrbracket \downarrow$, in the form of a labeled family of guarded languages. This concludes our discussion of the semantics of CF-GKAT.

3 Decision Procedure

To decide whether two CF-GKAT expressions denote the same guarded language, we propose CF-GKAT automata. We will show that every CF-GKAT expression can be converted to a CF-GKAT automaton with the same continuation semantics, by induction on the expression *à la* Thompson [45]. However, because CF-GKAT automata implement the continuation semantics, directly comparing them for language equivalence will not be sufficient. For instance, the following programs exhibit identical trace semantics, but differ in their continuation semantics:

$x := 1$ assert true.

The first program sets the indicator variable to 1, and the second program simply does nothing. These programs have the same trace semantics, because the trace semantics does not care about the final value of the indicator variable. Yet, their continuation semantics are different: guarded words in $\llbracket x := 1 \rrbracket_i^\#$ are always paired with the continuation $\text{acc } 1$ regardless of i , but $\llbracket \text{assert true} \rrbracket_i^\#$ preserves the starting indicator by outputting the continuation $\text{acc } i$. Furthermore, equivalence checking needs to account for jump resolution, by looking up the sequel of a trace ending in a continuation of the form $\text{jmp } (i, \ell)$ in the traces starting from label ℓ with indicator value i .

Thus, our decision procedure does not check equivalence of CF-GKAT automata directly; instead, we *lower* CF-GKAT automata into GKAT automata in a way that mirrors the move from the

continuation semantics to the guarded language semantics. This lets us reuse the decision procedure for GKAT automata equivalence [41]. The soundness and completeness of our decision procedure follows from the correctness of Thompson’s construction (Theorem 3.19), the correctness of the lowering (Theorem 3.13), and finally the correctness of GKAT automata equivalence checking [41].

3.1 GKAT Automata

To establish our decision procedure, we will first recap GKAT automata and their trace semantics.

Definition 3.1 (GKAT automata [28, 41]). A GKAT automaton $A \triangleq \langle S, \delta, \hat{s} \rangle$ consists of a set of states S , a transition function $\delta : S \rightarrow \text{At} \rightarrow \perp + \top + \Sigma \times S$, and a start state $\hat{s} \in S$.

Intuitively, given a state s and an atom α accounting for the truth value of each primitive test, a GKAT automaton will either *reject* the input, represented by $\delta(s, \alpha) = \perp$; *accept* the input, represented by $\delta(s, \alpha) = \top$; or to *transition* to a new state in S after executing an action from Σ , represented by $\delta(s, \alpha) \in \Sigma \times S$. A GKAT automaton induces a guarded language, by tracing all the possible execution paths reaching an accepting transition.

Definition 3.2. Given a GKAT automaton $A \triangleq \langle S, \delta, \hat{s} \rangle$, we define $\llbracket - \rrbracket_A : S \rightarrow \mathcal{G}$ as the (pointwise) smallest function satisfying the following rules for all $s \in S$ and $\alpha \in \text{At}$:

$$\frac{\delta(s, \alpha) = \top}{\alpha \in \llbracket s \rrbracket_A} \qquad \frac{\delta(s, \alpha) = (p, s') \quad w \in \llbracket s' \rrbracket_A}{\alpha p w \in \llbracket s \rrbracket_A}$$

Finally, we define the guarded language semantics of A by setting $\llbracket A \rrbracket = \llbracket \hat{s} \rrbracket_A$.

The trace equivalence of finite GKAT automata is decidable, which we record as follows.

THEOREM 3.3 (DECIDABILITY FOR GKAT [41]). *Given two finite GKAT automata A_0 and A_1 , it is decidable whether they represent the same guarded language, i.e., whether $\llbracket A_0 \rrbracket = \llbracket A_1 \rrbracket$. The algorithm to do this has a complexity that is nearly-linear² in the total number of states.*

3.2 CF-GKAT Automata

To leverage the efficient decision algorithm for GKAT automata, we will need to convert each CF-GKAT expression e and a starting indicator value i into a GKAT automaton that recognizes the guarded language $\llbracket e \rrbracket \downarrow_i^\sharp$. As discussed, this process is separated into two steps, where we use CF-GKAT automata as an intermediate between CF-GKAT expressions and GKAT automata. Like GKAT automata, CF-GKAT automata need a transition structure. Having a separate type for this transition structure will turn out to be useful, so we isolate it as follows.

Definition 3.4 (CF-GKAT dynamics). Given a set S , we write $G(S)$ for the set of possible CF-GKAT dynamics on S , which is given by the following function type, where C is as in Definition 2.6:

$$G(S) \triangleq I \times \text{At} \rightarrow \perp + C + \Sigma \times X \times I.$$

The elements of $G(S)$ represent transitions exiting a single (initial) state or label in a CF-GKAT automaton over a state set S . Given the current indicator value $i \in I$ and an atom α accounting for the truth value of each primitive test, a dynamics $\rho \in G(S)$ may either: *reject* the input ($\rho(i, \alpha) = \perp$); offer a *continuation* ($\rho(i, \alpha) \in C$); or an action, next state, and indicator value ($\rho(i, \alpha) \in \Sigma \times X \times I$).

Then the definition of CF-GKAT automaton is similar to GKAT automaton, except we will need a function $\lambda : L \rightarrow G(S)$ where $\lambda(\ell)$ provides a dynamics representing the “entry point” of label ℓ .

² $O(n \cdot \hat{\alpha}(n))$, where $\hat{\alpha}$ is the inverse Ackermann function; c.f. [43].

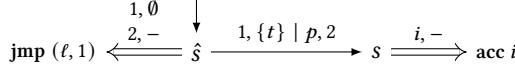
Definition 3.5. A *CF-GKAT automaton* $A \triangleq \langle S, \delta, \hat{s}, \lambda \rangle$ consists of a set of *states* S , a *transition function* $\delta : S \rightarrow G(S)$, a *start state* $\hat{s} \in S$, and a *jump map* $\lambda : L \rightarrow G(S)$.

The transition map δ assigns every state a dynamics, while the jump map λ assigns a dynamics to each label, indicating how to resume computation when a jump continuation is reached.

Example 3.6 (A simple CF-GKAT automaton). Let $I = \{1, 2\}$ and consider the following program:

if $t \wedge (x = 1)$ then $\{x := 2; \text{label } \ell; p\}$ else $\{x := 1; \text{goto } \ell\}$.

We can construct the following automaton with the same continuation semantics from state \hat{s} :



Here, $\hat{s} \xrightarrow[1, \{t\} \mid p, 2]{} s$ means that $\delta(\hat{s}, 1, \{t\}) = (p, s, 2)$ and $s \xrightarrow[i, -]{} \text{acc } i$ means that $\delta(s, i, \alpha) = \text{acc } i$ for $\alpha \in \text{At}$ and $i \in I$. The behavior of this automaton, when starting from \hat{s} , is as follows.

- If the input atom is $\{t\}$ – that is, the test t is true – and the indicator is 1, we transition to the state s , while setting the indicator to 2 and executing p (and skipping over label ℓ). Upon reaching s , the automaton accepts unconditionally, preserving the indicator.
- Otherwise, if the input atom is \emptyset – that is, the test t is false – or the indicator value is anything other than 1, the automaton will simply offer the continuation $\text{jmp } (\ell, 1)$, thereby telling execution to continue from label ℓ with indicator value 1.

As we can see, the behavior of \hat{s} indeed matches the behavior of the program when executing from the start. On the other hand, the entry dynamics for ℓ can be defined as follows:

$$\forall i \in I, \alpha \in \text{At}, \lambda(\ell, i, \alpha) \triangleq (p, s, i).$$

To put the above definition into words: when jumping to the label ℓ , we will reach the state s while executing p . Thus, the behavior of $\lambda(\ell)$ matches the behavior of the program when starting from ℓ .

Remark. We could have opted to instrument CF-GKAT automata with a *state* for each label ℓ , rather than a dynamics. This would simplify their definition, at the cost of complicating the Thompson construction (Section 3.4). An additional advantage of our approach is that we avoid creating unreachable states when lowering CF-GKAT automata to GKAT automata (Section 3.3).

To formalize the intuition of “behaviors”, as seen in the previous example, we will assign a continuation semantics to each CF-GKAT automaton. It is convenient to first assign a semantics to each dynamics $\rho \in G(S)$ in a CF-GKAT automaton $A \triangleq \langle S, \delta, \hat{s}, \lambda \rangle$, instead of every state.

Definition 3.7 (continuation semantics). Given an automaton $A \triangleq \langle S, \delta, \hat{s}, \lambda \rangle$, the continuation semantics of each dynamics $\rho \in G(S)$ is an indexed family $\llbracket \rho \rrbracket_A : I \rightarrow C$, defined as the (point-wise) smallest set satisfying the following rules for $i, j \in I, \alpha \in \text{At} \text{ and } c \in C$:

$$\begin{array}{c}
 \frac{\rho(i, \alpha) = \text{acc } j}{\alpha \cdot \text{acc } j \in (\llbracket \rho \rrbracket_A)_i} \quad \frac{\rho(i, \alpha) = \text{brk } j}{\alpha \cdot \text{brk } j \in (\llbracket \rho \rrbracket_A)_i} \quad \frac{\rho(i, \alpha) = \text{ret}}{\alpha \cdot \text{ret} \in (\llbracket \rho \rrbracket_A)_i} \\
 \\
 \frac{\rho(i, \alpha) = \text{jmp } (\ell, j)}{\alpha \cdot \text{jmp } (\ell, j) \in (\llbracket \rho \rrbracket_A)_i} \quad \frac{\rho(i, \alpha) = (p, s, j)}{p \cdot s \in (\llbracket \delta(s) \rrbracket_A)_j} \quad \frac{w \cdot c \in (\llbracket \delta(s) \rrbracket_A)_j}{\alpha p w \cdot c \in (\llbracket \rho \rrbracket_A)_i}
 \end{array}$$

Similar to the continuation semantics of expressions, the continuation semantics of automata are also labeled families of guarded languages with continuations: the semantics from the start $\llbracket A \rrbracket^\#$ is defined by the dynamics of the start state, and that of a label $\ell \in L$ is defined by the jump map:

$$\llbracket A \rrbracket^\# = \llbracket \delta(\hat{s}) \rrbracket_A, \quad \llbracket A \rrbracket^\ell = \llbracket \lambda(\ell) \rrbracket_A \text{ for } \ell \in L.$$

Example 3.8. Let the CF-GKAT automaton from Example 3.6 be A , we find that

$$\llbracket A \rrbracket_1^\sharp = \{\emptyset \cdot \mathbf{jmp}(\ell, 1), \{t\}p\emptyset \cdot \mathbf{acc} 2, \{t\}p\{t\} \cdot \mathbf{acc} 2\} \quad \llbracket A \rrbracket_2^\sharp = \{\emptyset \cdot \mathbf{jmp}(\ell, 1), \{t\} \cdot \mathbf{jmp}(\ell, 1)\}$$

$$\llbracket A \rrbracket_1^\ell = \{\alpha p \beta \cdot \mathbf{acc} 1 : \alpha, \beta \in \text{At}\} \quad \llbracket A \rrbracket_2^\ell = \{\alpha p \beta \cdot \mathbf{acc} 2 : \alpha, \beta \in \text{At}\}$$

3.3 Lowering CF-GKAT Automata to GKAT Automata

The process to lower a CF-GKAT automaton $\langle S, \delta, \hat{s}, \lambda \rangle$ into a GKAT automaton consists of two different components. We first “embed” the indicator values into the state set; the new state set then becomes $S \times I$. After that, we resolve all continuations in transitions using the jump map.

The second step requires some care. When $\delta(s, i, \alpha) = \mathbf{jmp}(\ell, j)$, the α -transition leaving the state (s, i) will take the behavior of ℓ starting from indicator j and atom α , that is $\lambda(\ell, j, \alpha)$. This by itself is alright, but we also need to account for subprograms such as label l ; $x := k$; $\text{goto } \ell'$, which entails that $\lambda(\ell, j, \alpha)$ points to a different label by returning $\mathbf{jmp}(\ell', k)$. In turn, $\lambda(\ell', k, \alpha)$ may also yield another jump, et cetera. To resolve these chained jumps, we will need to iterate the jump map, and terminate when either the result is no longer a jump, or an infinite loop is detected.

Definition 3.9 (iteration lifting). Given a function $h : X \rightarrow X + \perp + E$, where X is a finite set and $\perp + E$ specifies the “exit results”, we can iterate h until some exit value is reached, or a loop is detected. Formally, we define $\text{iter}(h)$ as the least function satisfying for all $x \in X$ that

$$\text{iter}(h)(x) = \begin{cases} \text{iter}(h)(h(x)) & h(x) \in X \\ h(x) & h(x) \in E \end{cases}$$

in the directed-complete partial order (DCPO) on functions from X to $X + \perp + E$ where $f \leq g$ when for all $x \in X$, we have that $f(x) = \perp$ or $f(x) = g(x)$. This makes $\text{iter}(h)$ the least fixed point of the Scott-continuous function on this DCPO that sends f to

$$x \mapsto \begin{cases} f(h(x)) & h(x) \in X \\ h(x) & h(x) \in E + \perp \end{cases}$$

By Kleene’s fixed point theorem this uniquely defines $\text{iter}(h)$.

Remark. When X is finite, we can directly compute $\text{iter}(h)$ by iterating h . More precisely, we can define a helper function $\text{iter}'(h) : 2^X \rightarrow X \rightarrow \perp + E$ with an additional parameter, as follows:

$$\text{iter}'(h)(M)(x) \triangleq \begin{cases} \perp & \text{if } x \in M \\ h(x) & \text{if } x \notin M \text{ and } h(x) \in \perp + E; \\ \text{iter}'(h)(M \cup \{x\})(h(x)) & \text{if } x \notin M \text{ and } h(x) \in X \end{cases}$$

It is then not too hard to show that $\text{iter}(h) = \text{iter}'(h)(\emptyset)$. Intuitively, M keeps track of explored values in X ; if an input has already been explored, then iter will return \perp . On the other hand, if $h(x)$ falls into the exit set $\perp + E$, then $\text{iter}(h)$ will exit and return $h(x)$. Finally, if $h(x)$ falls into X , then $\text{iter}(h)$ will continue to iterate with $h(x)$ as input, and mark x as explored. Note that $\text{iter}(h)$ is total because X is finite, so elements in X cannot be explored twice.

For the sake of clarity, when defining a function h to iterate, we will write **cont** for the injection $X \rightarrow X + \perp + E$, so that **cont**(x) indicates iteration continues with the value x ; and **exit** for the injection $\perp + E \rightarrow X + \perp + E$, so that **exit**(e) indicates iteration stops with the value e .

For jump resolution, iteration will continue when the result of λ is a jump, but exit otherwise.

Definition 3.10 (Jump resolution). Let S be a finite set, and let $\lambda : L \rightarrow G(S)$ be a jump function. We define the resolved jump map $\lambda \downarrow : L \rightarrow G(S)$, as follows:

$$\lambda \downarrow = \text{iter} \left((\ell, i, \alpha) \mapsto \begin{cases} \text{cont}(\ell', i', \alpha) & \lambda(\ell, i, \alpha) = \text{jmp } (\ell', i') \\ \text{exit}(\lambda(\ell, i, \alpha)) & \text{otherwise} \end{cases} \right)$$

Because of the way `iter` works, this means that a cycle of `jmp`-continuations (without actions in between) resolves to \perp . Indeed, GKAT treats unproductive infinite iterations in the `while` loops with the same strategy [41], identifying non-terminating behavior with behavior that rejects explicitly.

Note that $\lambda \downarrow$ resolves all internal jumps, i.e., $\lambda \downarrow(\ell, i, \alpha)$ is never a jump continuation, because the iteration map used to define $\lambda \downarrow$ keeps iterating on values of this form. We can then use the resolved jump map to tie together the loose ends indicated by jump continuations. In the process, we can reroute continuations of the form `brk` i to rejection, as these should not occur at the top level of a well-formed program (internal `brk`-continuations are handled in the next section).

Definition 3.11 (Lowering). Given a CF-GKAT automaton $A \triangleq \langle S, \delta, \hat{s}, \lambda \rangle$ and $i \in I$, we define the GKAT automaton $A \downarrow_i \triangleq \langle S \times I, \delta \downarrow, (\hat{s}, i) \rangle$, where $\delta \downarrow$ is defined in two steps:

$$\delta'((s, i), \alpha) \triangleq \begin{cases} \lambda \downarrow(\ell, i, \alpha) & \delta(s, i, \alpha) = \text{jmp } (\ell, i) \\ \delta(s, i, \alpha) & \text{otherwise} \end{cases}$$

$$\delta \downarrow((s, i), \alpha) \triangleq \begin{cases} \perp & \delta'(s, i, \alpha) = \text{brk } j \\ \top & \delta'(s, i, \alpha) = \text{ret} \text{ or } \delta'(s, i, \alpha) = \text{acc } j \\ \delta(s, i, \alpha) & \text{otherwise} \end{cases}$$

Note that $\delta \downarrow$ is well-defined: for all $s \in S$, $i \in I$ and $\alpha \in \text{At}$, it holds that $\delta \downarrow((s, i), \alpha) \in \perp + \top + \Sigma \times (S \times I)$, as expected for a GKAT automaton on state set $S \times I$.

Example 3.12. If A is the automaton from Example 3.6, then $A \downarrow_1$ is the following GKAT automaton:

$$\begin{array}{ccc} \emptyset, \{t\} & \xleftarrow{\emptyset \mid p} & (\hat{s}, 1) \xrightarrow{\{t\} \mid p} (s, 2) \xrightarrow{\quad} \emptyset, \{t\} \\ & \downarrow & \end{array}$$

Here, we use similar graphical conventions as before; for instance, $(\hat{s}, 1) \xrightarrow{\{t\} \mid p} (s, 2)$ means that $\delta \downarrow((\hat{s}, 1), \{t\}) = (p, (s, 2))$, and $(s, 2) \xrightarrow{\quad} \emptyset, \{t\}$ means that $\delta \downarrow((s, 2), \emptyset) = \delta \downarrow((s, 2), \{t\}) = \top$.

In this lowered automaton, the jump continuation originally reached from \hat{s} for indicator value 1 and t false (i.e., atom \emptyset) has been resolved via the jump map λ to continue in state s . This automaton also holds two copies of the state s , one for each possible indicator value; in this case, these states happen to have the same behavior, but in general they may be different.

Having defined our lowering operation, we can state its correctness as follows.

THEOREM 3.13 (CORRECTNESS OF LOWERING). Let $A \triangleq \langle S, \delta, \hat{s}, \lambda \rangle$ be a CF-GKAT automaton. The translation from CF-GKAT automata to GKAT automata commutes with the semantic jump resolution operator, in the sense that for $i \in I$, it holds that $\llbracket A \downarrow_i \rrbracket = \llbracket A \rrbracket \downarrow_i^\sharp$.

PROOF SKETCH. More generally, we can prove that for any state s of A , it holds that $\llbracket (s, i) \rrbracket_{\mathcal{A}_i} = G \downarrow_i^\sharp$, in which G is the labeled family given by $G^\sharp = \llbracket \delta(s) \rrbracket_A$ and $G^\ell = \llbracket \lambda(\ell) \rrbracket_A$. This property, which implies the main claim, can be proved by induction on the length of guarded words. \square

3.4 Converting Expressions to CF-GKAT Automata

The final piece of our puzzle is to convert CF-GKAT expressions to CF-GKAT automata. To accomplish this, we generalize a construction proposed for GKAT, which turns a GKAT expression into a GKAT automaton in a trace-equivalent manner [41]. This construction, which was inspired by Thompson’s construction to obtain a non-deterministic finite automaton from a regular expression [45], proceeds by induction on the structure of the expression. In contrast to the original, however, the Thompson construction for GKAT produces a GKAT automaton with a *start dynamics*, also called an *initial pseudostate* [41], instead of an explicit start state. We adopt this shift in presentation to efficiently compose automata, avoiding the silent transitions in the original [45].

Definition 3.14. A CF-GKAT automaton with start dynamics $A \triangleq \langle S, \delta, \iota, \lambda \rangle$ consists of S , δ and λ as in a CF-GKAT automaton, in addition to a start dynamics $\iota \in G(S)$. The labeled family defined by a CF-GKAT automaton with start dynamics, also denoted $\llbracket A \rrbracket$, is simply given by $\llbracket A \rrbracket^\sharp = \llbracket \iota \rrbracket_A$ and $\llbracket A \rrbracket^\ell = \llbracket \lambda(\ell) \rrbracket_A$, where the right-hand sides are defined as for plain CF-GKAT automata.

It should be clear that CF-GKAT automata with start dynamics can easily be converted to plain CF-GKAT automata by adding a start state \hat{s} that takes the behavior of the start dynamics ι :

$$\langle S, \delta, \iota, \lambda \rangle \mapsto \langle S + \hat{s}, \delta_\iota, \hat{s}, \lambda \rangle, \quad \text{where} \quad \delta_\iota(s, i, \alpha) \triangleq \begin{cases} \iota(i, \alpha) & \text{if } s = \hat{s} \\ \delta(s, i, \alpha) & \text{if } s \neq \hat{s} \end{cases} \quad (6)$$

Our construction turns a CF-GKAT expression e into a CF-GKAT automaton with start dynamics, which we call the *Thompson automaton for e* . The following paragraphs describe the construction and intuition behind each case in the construction. We make the simplifying assumption that each ℓ appears at most once in a subexpression of the form `label` ℓ . Moreover, we write `!` for the unique function $! : \emptyset \rightarrow X$ (for any set X), and in inductive cases we denote the Thompson automata for e_1 and e_2 by $A_1 \triangleq \langle S_1, \delta_1, \iota_1, \lambda_1 \rangle$ and $A_2 \triangleq \langle S_2, \delta_2, \iota_2, \lambda_2 \rangle$ respectively. Finally, to compress our notation, we will use the compact syntax of GKAT [41] for `if`-statements and `while`-loops:

$$e_1 +_b e_2 \triangleq \text{if } b \text{ then } e_1 \text{ else } e_2, \quad e_1^{(b)} \triangleq \text{while } b \text{ do } e_1.$$

Converting break, return, goto, and indicator assignment: recall that the semantics of `break`, `return`, `goto` ℓ , and indicator assignments simply emit the corresponding continuations. Thus, the non-trivial part of these Thompson automata are the start dynamics ι , which output said continuations. Because the start dynamics does not need to transition anywhere, there are no further states; because these programs also do not contain any `label` primitives, there is no need to assign a dynamics for them either. We thus construct the following automata:

$$\begin{array}{llll} S_{\text{break}} \triangleq \emptyset & \delta_{\text{break}} \triangleq ! & \iota_{\text{break}}(i, \alpha) \triangleq \text{brk } i & \lambda_{\text{break}}(\ell, i, \alpha) \triangleq \perp \\ S_{\text{return}} \triangleq \emptyset & \delta_{\text{return}} \triangleq ! & \iota_{\text{return}}(i, \alpha) \triangleq \text{ret} & \lambda_{\text{return}}(\ell, i, \alpha) \triangleq \perp \\ S_{\text{goto } \ell} \triangleq \emptyset & \delta_{\text{goto } \ell} \triangleq ! & \iota_{\text{goto } \ell}(i, \alpha) \triangleq \text{jmp } (\ell, i) & \lambda_{\text{goto } \ell}(\ell, i, \alpha) \triangleq \perp \\ S_{x:=i} \triangleq \emptyset & \delta_{x:=i} \triangleq ! & \iota_{x:=i}(i, \alpha) \triangleq \text{acc } i & \lambda_{x:=i}(\ell, i, \alpha) \triangleq \perp \end{array}$$

Converting tests and primitive actions: the conversions of primitive tests and primitive actions largely inherit the Thompson construction for GKAT. The Thompson automaton for tests `assert` b contains only a start dynamics, which accepts the input indicator-atom pairs if and only if they satisfy b . The Thompson’s automata for primitive actions p contains a start dynamics that always

executes the action p before transitioning to the unique state, which accepts unconditionally.

$$\begin{aligned}
 S_{\text{assert } b} &\triangleq \emptyset & S_p &\triangleq \{\ast\} \\
 \delta_{\text{assert } b} &\triangleq ! & \delta_p(s, i, \alpha) &\triangleq \text{acc } i \\
 \iota_{\text{assert } b}(i, \alpha) &\triangleq \begin{cases} \text{acc } i & (i, \alpha) \in \llbracket b \rrbracket \\ \perp & (i, \alpha) \notin \llbracket b \rrbracket \end{cases} & \iota_p(i, \alpha) &\triangleq (p, \ast, i) \\
 \lambda_{\text{assert } b}(\ell, i, \alpha) &\triangleq \perp & \lambda_p(\ell, i, \alpha) &\triangleq \perp
 \end{aligned}$$

Converting labels: recall that label ℓ' is a does not affect the semantics from the start of the program, i.e., should be the same as the sequential identity assert true. However, the behavior of label ℓ' and assert true diverges when we consider the jump map: if $\ell = \ell'$, then executing label ℓ' starting from ℓ lets us reach the end of the program, whereas assert true does not contain any label. We thus end up with the following Thompson automaton for label ℓ' :

$$S_{\text{label } \ell'} \triangleq \emptyset \quad \delta_{\text{label } \ell'} \triangleq ! \quad \iota_{\text{label } \ell'}(i, \alpha) \triangleq \text{acc } i \quad \lambda_{\text{label } \ell'}(\ell, i, \alpha) \triangleq \begin{cases} \text{acc } i & \ell = \ell' \\ \perp & \ell \neq \ell' \end{cases}$$

Converting if statements: the Thompson automaton for if b then e_1 else e_2 is also similar to that of GKAT: if the input indicator-atom pair satisfies (resp. falsifies) b , then ι will enter the Thompson automaton of e_1 (resp. e_2) by taking on the behavior of ι_1 (resp. ι_2). The jump map λ assigns the entry point for label ℓ based on where it appears: if it appears in e_1 , then ℓ will take its entry point in A_1 ; similarly, if label ℓ appears in e_2 , ℓ will take its entry point in A_2 .

$$\begin{aligned}
 S_{e_1+b e_2} &\triangleq S_1 + S_2 & \delta_{e_1+b e_2}(s) &\triangleq \begin{cases} \delta_1(s) & \text{if } s \in S_1 \\ \delta_2(s) & \text{if } s \in S_2 \end{cases} \\
 \lambda_{e_1+b e_2}(\ell) &\triangleq \begin{cases} \lambda_1(\ell) & \text{label } \ell \text{ appears in } e_1 \\ \lambda_2(\ell) & \text{label } \ell \text{ appears in } e_2 \\ \perp & \text{otherwise} \end{cases} & \iota_{e_1+b e_2}(i, \alpha) &\triangleq \begin{cases} \iota_1(i, \alpha) & (i, \alpha) \in \llbracket b \rrbracket \\ \iota_2(i, \alpha) & (i, \alpha) \notin \llbracket b \rrbracket \end{cases}
 \end{aligned}$$

Converting Sequencing: Sequencing of automata can be defined by *uniform continuations* [41], which combine two dynamics $h_1, h_2 \in G(S)$ into a new dynamics $h_1[h_2]$. The latter acts like h_1 in almost all cases, except when h_1 accepts – then it will take on the behavior of h_2 . In other words, $h_1[h_2]$ connects all the accepting transition of h_1 to h_2 . Uniform continuation is typically used to compose two automata or add self-loops to an automaton; it can be formally defined as follows.

Definition 3.15 (Uniform Continuation). Let S be a set. Given two dynamics $h_1, h_2 \in G(S)$, their *uniform continuation* is the dynamics $h_1[h_2] \in G(S)$, defined as follows:

$$h_1[h_2](i, \alpha) \triangleq \begin{cases} h_2(i', \alpha) & \text{if } h_1(i, \alpha) = \text{acc } i' \\ h_1(i, \alpha) & \text{otherwise} \end{cases}$$

To construct the Thompson automaton for $e_1; e_2$, we will simply connect all the accepting transitions in A_1 to A_2 by applying uniform continuations on start dynamics ι_1 , transitions δ_1 , and jump map λ_1 , while preserving the dynamics in A_2 . Formally, this works out as follows:

$$\begin{aligned}
 S_{e_1; e_2} &\triangleq S_1 + S_2 & \iota_{e_1; e_2} &\triangleq \iota_1[\iota_2] \\
 \delta_{e_1; e_2}(s) &\triangleq \begin{cases} \delta_1(s)[\iota_2] & \text{if } s \in S_1 \\ \delta_2(s) & \text{if } s \in S_2 \end{cases} & \lambda_{e_1; e_2}(\ell, i, \alpha) &\triangleq \begin{cases} \lambda_1(\ell)[\iota_2] & \text{label } \ell \text{ appears in } e_1 \\ \lambda_2(\ell) & \text{label } \ell \text{ appears in } e_2 \\ \perp & \text{otherwise} \end{cases}
 \end{aligned}$$

Converting while loops: Like GKAT automata, CF-GKAT automata require every transition between states to execute a primitive action. This presents a unique challenge in defining the start dynamics for while loops. Namely, a loop may not immediately encounter a primitive action on its first iteration, but this may trigger a change in indicator value that causes a primitive action to be executed in the *next* iteration, or the one after that, et cetera. This is reminiscent of jump resolution, where we also have to chase through some indirection to arrive at the right transition.

As a concrete example of this phenomenon, consider the following CF-GKAT program:

```
while true do {if  $x = 0$  then  $x := 1$ 
else if  $x = 1$  then break else { assert true } }
```

(7)

If this program starts with the indicator 0, then the first continuation (`brk 1`) will be encountered on the second iteration of the loop. Even worse, when starting with an indicator value like $x = 2$, the program will enter an infinite loop and never encounter a primitive action or continuation.

Fortunately, these difficulties can be resolved by the `iter` function (Definition 3.9): we repeat the start of the loop body until we encounter some productive behavior, or we get stuck.

Definition 3.16 (Iterated Start Dynamics). Let S be a set, let $h \in G(S)$, and $b \in \text{BExp}$. We can use the `iter` function to define $h^b : G(S)$, as follows:

$$h^b \triangleq \text{iter} \left((i, \alpha) \mapsto \begin{cases} \text{exit}(\text{acc } i) & \text{if } (i, \alpha) \notin \llbracket b \rrbracket \\ \text{cont}(i', \alpha) & \text{if } (i, \alpha) \in \llbracket b \rrbracket \text{ and } h(i, \alpha) = \text{acc } i' \\ \text{exit}(h(i, \alpha)) & \text{otherwise} \end{cases} \right)$$

In the first case, the input (i, α) does not satisfy b , causing the while loop to terminate. In the second case, the loop body accepts (i, α) immediately and returns the exit indicator value i' , thus the iteration of loop body will continue with (i', α) . In the final case, the program executes an action or encounters a non-local control, which is when the iteration can be stopped.

Example 3.17 (Iterated Start Dynamics). Consider program 7 above with indicator set $\{0, 1, 2\}$, no primitive action, no label, and no primitive test. Then the only atom is \emptyset , and the Thompson automaton $A_1 \triangleq \langle S_1, \delta_1, \iota_1, \lambda_1 \rangle$ for the loop body can be computed as follows:

$$\begin{aligned} \iota_1(0, \emptyset) &\triangleq \text{acc } 1 \\ S_1 &\triangleq \emptyset & \iota_1(1, \emptyset) &\triangleq \text{brk } 1 & \delta_1 &\triangleq ! & \lambda(\ell, i, \alpha) &\triangleq \perp \\ && \iota_1(2, \emptyset) &\triangleq \text{acc } 2 \end{aligned}$$

Then we compute the iterated start dynamics ι_1^{true} with input $(0, \emptyset)$ and $(2, \emptyset)$:

$$\begin{aligned} \iota_1^{\text{true}}(0, \emptyset) &= \iota_1^{\text{true}}(1, \emptyset) & \text{because } (0, \emptyset) \in \llbracket \text{true} \rrbracket \text{ and } \iota_1(0, \emptyset) = \text{acc } 1 \\ &= \text{brk } 1 & & & & \text{because } \iota_1(1, \emptyset) = \text{brk } 1 \\ \iota_1^{\text{true}}(2, \emptyset) &= \iota_1^{\text{true}}(2, \emptyset) & \text{because } (2, \emptyset) \in \llbracket \text{true} \rrbracket \text{ and } \iota_1(2, \emptyset) = \text{acc } 2 \\ &= \perp & & & & \text{because the input } (2, \emptyset) \text{ is already explored} \end{aligned}$$

With the start dynamics defined, we still need to resolve structures within the loop body, like the break-continuation. To perform break-resolution, we extend the $\lfloor - \rfloor$ operator to dynamics.

Definition 3.18. Let S be a set, and let $h \in G(S)$. We define $\lfloor h \rfloor \in G(S)$ by lifting h via $\lfloor - \rfloor$ (c.f. Definition 2.13) when it returns a continuation, that is to say:

$$\lfloor h \rfloor(i, \alpha) = \begin{cases} \lfloor h(i, \alpha) \rfloor & \text{if } h(i, \alpha) \in C \\ h(i, \alpha) & \text{otherwise} \end{cases}$$

Finally, the transition function δ and jump map λ can be defined by first connecting δ_1 and λ_1 back to the start dynamics $\iota^{(b)}$, forming a loop in the automaton, and then resolving the **brk** i continuations using $\lfloor - \rfloor$. Formally, the Thompson automaton for loop $e^{(b)}$ is defined as follows:

$$S_{e_1^{(b)}} \triangleq S_1 \quad \delta_{e_1^{(b)}} \triangleq \lfloor \delta_1(s)[\iota_1^{(b)}] \rfloor \quad \iota_{e_1^{(b)}} \triangleq \lfloor \iota_1^{(b)} \rfloor \quad \lambda_{e_1^{(b)}}(\ell) \triangleq \lfloor \lambda(\ell)[\iota_1^{(b)}] \rfloor$$

Note how in the case of the jump map, we resolve **brk**-continuations *after* the uniform continuation with the iterated start dynamics. This way, a jump to a label inside the loop that hits a break statement immediately will continue executing execution after the loop, as is to be expected.

The correctness of the Thompson construction for CF-GKAT can be formulated as follows:

THEOREM 3.19 (THOMPSON'S CONSTRUCTION PRESERVES THE CONTINUATION SEMANTICS). *Given an expression $e \in \text{Exp}$, let A_e be the Thompson automaton for e . Now $\llbracket e \rrbracket = \llbracket A_e \rrbracket$.*

PROOF SKETCH. By induction on e . This proof is somewhat tedious, but ultimately doable. The inductive cases are very similar to the ones for the Thompson construction in plain GKAT [41]; the only difference is that, this time, the semantics of the jump map needs to be taken into account. \square

3.5 Algorithm, Completeness, and Complexity

With the definitions of lowering and Thompson's construction established, the decision procedure mostly follows. Nevertheless, it remains essential to define the alphabets Σ, T, I , and L , representing the set of primitive actions, primitive tests, indicator values, and labels, respectively. We may safely restrict primitive actions, primitive tests, and labels to those explicitly present in the expression, as expanding the alphabet beyond these will preserve (the equational theory of) the trace semantics.

The set of indicator values, however, occupies a unique position. If the initial indicator value is absent from the program, the program's traces may diverge from traces starting from the present indicator values. Program 7 is one of the witnesses of this phenomenon: if the initial indicator value is 0 or 1, the program terminates; however, when starting from an indicator value that does not appear in the program, the program will loop indefinitely. An even simpler example is $\text{assert}(x = 1); p$, which executes p if the initial indicator value is 1, but rejects when started with indicator values that are not present in the program. Fortunately, given an expression e , it is not hard to show that if neither i nor i' appears in e , then the behaviors for both values coincide, i.e., $\forall \ell, \llbracket e \rrbracket_{i'}^\ell = \llbracket e \rrbracket_i^\ell$. Thus, when gathering the indicator values, it suffices to take those that appear explicitly in the program augmented with a fresh value $*$ that does not appear in the program.

We summarize our decision procedure as follows:

- (1) Given two CF-GKAT programs e, f , we first collect their alphabets Σ, T, I , and L . We gather the sets of primitive actions Σ , primitive tests T , and labels L that are present in either e or f . Additionally, we identify the set of indicator values I , encompassing those found in either e or f , along with an additional indicator $*$ that is exclusive to both programs.
- (2) We then proceed to compute the Thompson automata of e and f and convert them into CF-GKAT automata, denoted as A_e and A_f . It is noteworthy that these automata preserve the continuation semantics (Theorem 3.19). Formally,

$$\llbracket A_e \rrbracket = \llbracket e \rrbracket, \quad \llbracket A_f \rrbracket = \llbracket f \rrbracket.$$

- (3) Subsequently, we lower both A_e and A_f to GKAT automata $A_e \downarrow_i$ and $A_f \downarrow_i$ for each $i \in I$. By Theorem 3.13, the latter exhibits the same traces as e and f starting from i :

$$\llbracket A_e \downarrow_i \rrbracket = \llbracket A_e \rrbracket \downarrow_i^\# = \llbracket e \rrbracket \downarrow_i^\#, \quad \llbracket A_f \downarrow_i \rrbracket = \llbracket A_f \rrbracket \downarrow_i^\# = \llbracket f \rrbracket \downarrow_i^\#$$

- (4) Finally, we run a GKAT automata equivalence check [41] on $A_e \downarrow_i$ and $A_f \downarrow_i$ for each $i \in I$, and return true when all GKAT automata equivalence checks return true.

Soundness and completeness: The soundness and completeness of this algorithm now follow as a corollary of the corresponding properties for the decision procedure in GKAT. We denote the algorithm introduced above as $\text{equiv}_{\text{CFGKAT}}$, while the decision algorithm for equivalence of GKAT automata is denoted as $\text{equiv}_{\text{GKAT}}$. Thus, we establish the equivalence:

$$\begin{aligned}
 \text{equiv}_{\text{CFGKAT}}(e, f) &\iff \forall i \in I, \text{equiv}_{\text{GKAT}}(A_e \downarrow_i, A_f \downarrow_i) && (\text{def. equiv}_{\text{CFGKAT}}) \\
 &\iff \forall i \in I, \llbracket A_e \downarrow_i \rrbracket = \llbracket A_f \downarrow_i \rrbracket && (\text{corr. equiv}_{\text{GKAT}}) \\
 &\iff \llbracket A_e \rrbracket \downarrow^\# = \llbracket A_f \rrbracket \downarrow^\# && (\text{Theorem 3.13}) \\
 &\iff \llbracket e \rrbracket \downarrow^\# = \llbracket f \rrbracket \downarrow^\# && (\text{Theorem 3.19})
 \end{aligned}$$

Therefore e and f are trace equivalent if and only if $\text{equiv}_{\text{GKAT}}(e, f)$ returns true.

Algorithm complexity: We now give a rough account of the computational cost for deciding equivalence of CF-GKAT programs. Like GKAT, we consider T to be fixed for the purpose of analyzing complexity; otherwise, deciding equivalence of (CF-)GKAT programs is co-NP hard [41].

Starting with an expression e , we observe that the number of states in the Thompson automaton A_e is bounded by $|e|$, which is the size of e as a term. While computing this automaton we must bear in mind that deriving the iterated start dynamics of a loop using `iter` may take up to $|I|$ steps (assuming we use memoization to compute the output for each input), and this happens at most $|e|$ times. After lowering, each GKAT automaton $A_e \downarrow$ contains at most $|I| \times |e|$ states, and computing the jump resolution using `iter` takes on the order of $|L| \times |I|$ steps (again using memoization).

For each $i \in I$, determining the equivalence between $A_e \downarrow_i$ and $A_f \downarrow_i$ is accomplished in (nearly³) linear time relative to the number of states, so this check takes about $|I| \times (|e| + |f|)$ steps. To verify trace equivalence between e and f , this check is required for $A_e \downarrow_i$ and $A_f \downarrow_i$ across all $i \in I$.

Therefore, the overall time spent on the equivalence checks is on the order of $|I|^2 \times (|e| + |f|)$, while computing the automata can take roughly $|L| \times |I| + |I| \times (|e| + |f|)$ time. This implies that the algorithm's complexity scales (nearly) linearly with the sizes of e and f , linearly in the size of L , and quadratically with respect to the number of indicator values $|I|$.

4 Control Flow Validation

We hypothesize that CF-GKAT can be a useful tool to check whether two programs have the same control flow – that is, under the same circumstances, they execute the same sequence of actions. In this section, we report on the performance of CF-GKAT in two particular use cases.

The first use case concerns the *control flow structuring* stage of a decompiler. Briefly put, a decompiler is a program that infers a high-level language representation of a binary executable file. In earlier stages, the decompiler builds a *control-flow graph* from the binary [46], which represents how control transfers from one part to the next. The control flow structuring pass is tasked with inferring an equivalent representation of this control flow in terms of constructs like `if-then-else` and `while-do`. The second use case is about refactoring operations aimed at making code more readable by eliminating `goto` instructions, which may introduce indicator variables [10, 16].

CF-GKAT can be used to validate that the control flow of the input (graph or code) corresponds to that of the output code. In general, such algorithms can never be *complete* with respect to input-output equivalences, as they cannot automatically validate the correctness of refactorings depend on the meaning of primitive commands. However, CF-GKAT should be applicable to refactoring operations that rearrange the code for the sake of improving the presentation of the control flow.

³We omit the factor coming from the inverse Ackermann function $\hat{\alpha}(n)$, which is at most 5 for any realistic number of states, out of consideration for the sake of simplicity.

4.1 Implementation

We target C, as it is a widely used programming language that allows for non-local control flow. Our implementation is written in OCaml and is divided into a front-end and a back-end [51].

The *front-end* converts a function defined in a C file to a CF-GKAT expression, using the `clangML`⁴ OCaml bindings for `clang`.⁵ The conversion first determines which variables, if any, qualify as indicator variables and picks (at most) one. Next, it lifts the C syntax tree to a CF-GKAT program, mapping (1) all assignments and comparisons of the indicator variable to their CF-GKAT counterparts, (2) all other primitive statements to uninterpreted actions, and (3) all control flow constructs to the corresponding CF-GKAT operator whenever possible.

The *back-end* is responsible for compiling a CF-GKAT expression to a GKAT automaton, and for comparing two GKAT automata. This GKAT automaton is derived using the construction from Section 3; equivalence is decided by checking bisimulation [41].

These parts combine into a tool that accepts two C files, and checks whether their functions (paired by name) have the same control flow. Not all C programs can be converted to CF-GKAT expressions, and so our front-end is not complete; instead, we aimed to support a sizeable portion of the code in GNU Coreutils to perform the experiments described in the next section.

Remark. In particular, our front-end internally converts `for`-loops to `while`-loops, and `do-while` loops to (partially unrolled) `while` loops. The former transformation is sound; the latter is admissible when the loop body does not contain `break` or a label, which is the case for our experiment below.

4.2 A Case Study from GNU Coreutils

We used GNU Coreutils as a source of C code containing non-local control flow. As is customary with projects of this size, the code is organized into several files that collectively define hundreds of functions, ranging from very simple to very complex. Throughout this section, we focus on the function `mp_factor_using_pollard_rho` in `factor.c`. The code in Figure 1a shows an abridged version of this function in Coreutils 9.5. We now discuss the two transformations we targeted.

Compilation-decompilation. We hypothesize that CF-GKAT should be usable to validate the output of control flow structuring algorithms in decompilers. To fully test this hypothesis, we would need to have access to the internal representations used before and after control flow structuring, and convert those to GKAT automata and CF-GKAT expressions, respectively. Unfortunately, doing this would entail a substantial engineering effort. As a more feasible but slightly less rigorous benchmark, we *compiled* C code to x86 binary code, and then *decompiled* the result using Ghidra.¹ This transformation can be implemented without modifying existing codebases, and should still give insight into decompiler correctness by letting us compare the decompiled source to the original.

However, we immediately face the challenge of pairing the primitive actions from the decompiled code with actions in the source code. For instance, a primitive action `i += 1` in the source code could be decompiled to `i++`. Detecting *all* such transformations would be outside the scope of our project. To address this, we make the compiler *blind* to the nature of the primitive actions and primitive tests by replacing them with calls to new functions `void pact(int)` and `bool pbool(int)`, respectively. The parameter distinguishes the primitives, and the correspondence in the decompiled code can be inferred from it. Figure 1b shows the (abbreviated) *blinded* version of our case study function. This transformation does not alter control flow, but the blinded code can be longer, as the blinder also expands preprocessor macros. Crucially, blinding depends on indicator variable detection, since indicator tests and assignments need to remain in the blinded code.

⁴<https://memcad.gitlabpages.inria.fr/clangml/>

⁵<https://clang.llvm.org/>

```

void mp_factor_using_pollard_rho() {
    pact(197);
    ...
    do {
        if (pbool(83)) {
            pact(194);
        }
    }

    } while (0);
    pact(193);
    ...

static void mp_factor_using_pollard_rho(... {
    mpz_t x, z, y, P;
    mpz_t t, t2;
    devmsg("...", a);
    ...
    while (mpz_cmp_ui(n, 1) != 0) {
        for (;;) {
            do {
                mpz_mul(t, x, x);
                ...
                if (k % 32 == 1) {
                    mpz_gcd(t, P, n);
                    if (mpz_cmp_ui(t, 1) != 0)
                        goto factor_found;
                    mpz_set(y, x);
                }
            } while (--k != 0);
            mpz_set(z, x);
            ...
            for (unsigned long long int i = 0; i < k; i++)
                ...
                mpz_mul(t, x, x);
            }
            mpz_set(y, x);
        }
        factor_found:
        do {
            pact(173);
            ...
            if (!pbool(165)) {
                do {
                    if (pbool(83)) {
                        pact(164);
                    }
                } while (0);
                pact(163);
            } else {
                pact(161);
            }
            if (pbool(160)) {
                pact(159);
                break;
            }
            pact(158);
            ...
        }
        mpz_clears(P, t2, t, z, x, y, nullptr);
    }
}

```

(a) Original code of the function.

(b) “Blinded” code of the function.

Fig. 1. Different versions of `mp_factor_using_pollard_rho` in `factor.c`, part of GNU Coreutils.

```

void mp_factor_using_pollard_rho() {
    int factor_found = 0;
    pact(197);
    ...
    do {
        if (pbool(83)) {
            pact(194);
        }
    } while (0);
    pact(193);
    ...
    while (pbool(83)) {
        for (; factor_found == 0;) {
            do {
                pact(176);
                ...
                if (pbool(183)) {
                    pact(182);
                    if (pbool(83)) {
                        factor_found = 1;
                    }
                    if (factor_found == 0)
                        pact(173);
                }
            } while ((factor_found == 0) && pbool(181));
            if (factor_found == 0) {
                pact(180);
                ...
                for (pact(177); pbool(138); pbool(59)) {
                    pact(176);
                    ...
                }
                pact(173);
            }
        }
        factor_found = 0;
        do {
            pact(172);
            ...
        } while (pbool(83));
        pact(166);
        if (!pbool(165)) {
            do {
                if (pbool(83)) {
                    pact(164);
                }
            } while (0);
            pact(163);
        } else {
            pact(161);
        }
        if (pbool(160)) {
            pact(159);
            break;
        }
        pact(158);
        ...
    }
    pact(155);
}

```

(a) Ghidra's decompilation of the function.

(b) The original, with gotos removed by Calipso.

Fig. 2. Different versions of mp_factor_using_pollard_rho in factor.c, part of GNU Coreutils.

In this experiment, we use Ghidra as our decompiler, and clang as our compiler. The C code obtained from compiling and then decompiling the blinded code (Figure 1b) and manually removing decompilation artifacts (see remark below) is shown in Figure 2a. The code produced is markedly different from the source – it has three more gotos with as many labels – yet our implementation is able to validate, in a fraction of a second,⁶ that this code is equivalent to the original.

Remark. The following manual modifications were needed. (a) Ghidra assigns the return value of pbool to a variable before reading it. This would be interpreted as a primitive action by our conversion. We manually propagate the assignments to avoid this. (b) In conditional expressions, Ghidra masks values with 1 and then compares the result to 0, e.g.: $(pbool(n) \& 1) != 0$. Again, our conversion would interpret this as a primitive action and not a test. We simplify this to $pbool(n)$. We believe that these steps could be automated, or Ghidra could be adapted to avoid them, but opted to perform them manually to keep our case study simple.

Goto-elimination. In general, goto statements can be eliminated by introducing additional indicator variables to guide control flow [7, 16, 49]. This idea underpins Erosa and Hendren’s goto-elimination algorithm [16]. Calipso [10] provides an improved implementation⁷ of their algorithm.

We ran Calipso on the blinded code in Figure 1b, and manually adjusted the output (see remark below). The result again fundamentally different from the blinded source, as it contains a new indicator variable that is not present in the input, and does not contain any gotos. Our tool confirms, again in a fraction of a second, that the code thus obtained is equivalent to the blinded input. Indicator detection is crucial: if the newly-introduced variable `factor_found` is not detected as an indicator, its assignments and tests will be converted into primitive actions and tests respectively.

Remark. The following manual modifications were needed for this experiment. (a) We changed the type of the newly introduced variable `factor_found` from `char` to `int`. (b) We changed instances where `factor_found` is used as a Boolean into comparisons (e.g., `if (factor_found == 0)`). As before, we consider these adaptations to be relatively minor; a future goto-elimination algorithm could prevent them from being necessary, or they could be performed automatically.

Blinder coverage. Overall, we consider our blinder to be quite promising: despite the fact that our current prototype does not support constructs like `switch`, we are able to blind 807 of the 1425 functions in Coreutils. We use the cyclomatic complexity number (CCN) [33] as reported by Lizard⁸ to provide an overview of how intricate these functions are. The majority of the blinded functions have CCN up to 4, but we are able to blind some rather complex functions with the maximum CCN being 44. Figure 3 provides a plot of the frequency of blinded functions per CCN.

CCN alone is not the discriminating factor that prevents us from blinding functions: a run of Lizard on the original source code of Coreutils shows that, while we are unable to blind functions with very high CCN (the maximum being 281), the median and mode CCN of functions we are not able to blind are 9 and 2, respectively, which shows there are relatively simple functions with respect to the CCN metric that we cannot handle. In order to improve our coverage the focus should be on supporting more C constructs; for example, `switch` statements, `continue` statement, and ternary operators. Such an extension requires some additions to our theory, as would support for other constructs. Nevertheless, it is quite possible to embed special cases in the current theory.

⁶The experiments ran on a system with an Intel Core i7-9750H CPU @ 2.60GHz and 16 GB of RAM.

⁷<https://github.com/BinaryAnalysisPlatform/FrontC>

⁸<http://www.lizard.ws/>

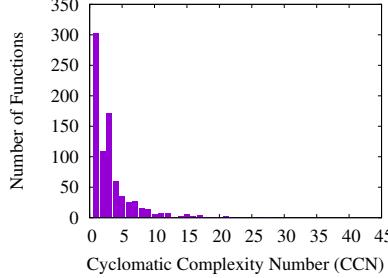


Fig. 3. Plot of the number of blinded functions per cyclomatic complexity number (CCN). The maximum CCN found was 44, yet the majority of the blinded functions have low CCN.

5 Related Work

5.1 (De)compiler Verification and Validation

Because we anticipate that the primary use case of CF-GKAT is the verification and validation of (de)compilers, we briefly go over existing work in this area, and relate it to ours.

Compiler verification and validation is a well-developed research area. The most rigorous approach is to use a proof assistant to verify the compilation algorithm relative to a formal semantics of the source and target language [31, 36, 42]. This method offers a foundational guarantee of correctness without imposing runtime overhead, but can be challenging and time-consuming to adapt on a large scale. We envision that CF-GKAT could be used to automate correctness proofs of program transformations (provided that CF-GKAT equivalence is sound w.r.t. the language semantics).

Compiler testing feeds synthesized programs into different or differently configured compilers, and compares the resulting programs [11, 30, 50]. While this approach does not provide any formal guarantee of correctness, it has been effective in finding compiler bugs [50]. Within this framework, CF-GKAT could be used as an equivalence checking oracle; however, because our work is limited to *control flow* equivalence, a blinding pass (similar to our experiments above) is likely necessary.

Translation validation [18, 20, 21, 35, 37, 40, 52] scales well *and* offers formal guarantees. This is achieved through the use of symbolic evaluations and bisimulation to generate verification conditions for SMT solvers, allowing the validator to handle both data-flow and control-flow transformations simultaneously, albeit at the cost of decidability. Some approaches also employ graph isomorphism checks [5], which can be more computationally expensive than bisimulation. Compared to other efforts, our focus is exclusively on control flow transformations. This allows us to work with a minimal language, sufficient to cover well-known control-flow transformations [7, 16, 19, 49], and establish meta-properties such as correctness of our decision procedure and low complexity — all of which are properties out of reach for unabsttracted (Turing-complete) languages.

Predicate abstraction [6] can help focus on relevant parts of a program prior to analysis, which makes it similar to blinding. However, predicate abstraction requires input on *how* to abstract the program, whereas blinding abstracts away all parts that are not relevant to flow control.

Decompiler verification and validation is, by comparison, not as explored as compiler verification and validation. Again, a distinction can be made between verification [46], validation [15] and testing [9, 32]. Current efforts in verification and validation focus on *lifting*, i.e., decoding machine code into assembly instructions or some other intermediate representation [15, 46, 47]. There has been some preliminary work on developing a sound decompiler that can produce C code, although control flow structuring is limited [48]. This constitutes a prime use case for CF-GKAT.

5.2 Kleene Algebra and Control Flow Verification

Existing work has explored non-local control-flow structures and indicator variables within the framework of KAT, albeit with a number of differences from our work.

Kozen characterized the semantics for programs with non-local control-flow structures as a family of KAT expressions [23]. This approach yields a decision procedure for program equivalences. In contrast, CF-GKAT takes a more explicit approach by defining the continuation semantics, and the equivalence is computed by converting programs directly into automata. This resolves an open question by Kozen [23], on whether non-local control flow structures can be treated “directly”, without being converted into KAT expressions. The interested reader is referred to [44] for an investigation of the limits of compositionally specified deterministic control flow.

Grathwohl et al. [19] proposed *KAT+B!* an extension of KAT with “mutable tests”, which can be thought of as indicator variables with $I = \{\text{true}, \text{false}\}$. Concretely, their setter $b!$ equates to indicator assignment $x = \text{true}$ and $\bar{b}!$ to $x = \text{false}$; similarly, their tester $b?$ corresponds to primitive indicator tests. Although this is a special case of indicator variables, *KAT+B!* can simulate indicator variable over a finite set I with $|I|$ mutable tests. We opt to treat indicator assignments and tests as primitives, rather than restricting ourselves to (boolean-valued) mutable tests.

Our treatment of indicator variables also draws inspiration from NetKAT [2]. Specifically, NetKAT can be seen as a special case of KAT with indicator variables, the only primitive action is `dup`. KATch [34], a fast symbolic equivalence checker for NetKAT, is of particular interest to us in this space. In the future, we hope to scale our equivalence checker using similar symbolic techniques.

To reason about variable assignment beyond indicator variables, Schematic KAT [3] provides a fine-grained algebraic theory for assignments over uninterpreted functions. Later, Schematic KAT was also extended to reason about local variables [1]. Neither work covers the complexity of the equivalence problem for schematic KAT and its extensions. Kleene algebra can also be extended with nominal techniques [17, 24, 25], which may help to reason about potentially infinite data domains, although the inclusion of tests to nominal Kleene algebra has not yet been investigated.

Separately, Kozen and Patron [26] have applied KAT to verify compiler correctness. Their system directly uses postulated equalities for parts of their verification task. In contrast, our framework is based on the trace semantics, a commonly accepted semantics for while-programs.

5.3 Complexity and Expressivity

Finally, unlike the systems above, our system is based on GKAT, instead of KAT. This enhances the scalability of our equivalence checking algorithms to accommodate larger programs: whereas equivalence of KAT expressions is PSPACE-complete [13], equivalence of CF-GKAT expressions can be verified in polynomial time for a fixed number of tests [41]. Symbolic techniques previously applied to (Net)KAT may also provide better ways of mitigating the complexity of tests [34, 38].

Our system is the first to integrate both indicator variables and non-local control flow in a unified framework, which enables the verification of complex control-flow transformations that leverage both indicator variables and non-local control flow [49]. Our notion of trace equivalence is also coarser than previous systems; CF-GKAT equates programs ending in different indicator values:

```
x := true; if x = true then print(1) else print(2)
x := false; if x = false then print(1) else print(2)
```

The two programs above are equivalent to an outside observer, as both of them will print 1; yet the assignment of x is different at the end of the so, thus systems like *KAT+B!* [19] will not equate them. On the other hand, our equivalence is not a congruence. For example, concatenating `assert (x = true)` to the two equivalent programs above will yield inequivalent programs.

6 Conclusion and Future Work

In this paper, we introduced CF-GKAT (Control-Flow GKAT), a system that extends GKAT with non-local control flow and indicator variables to validate control-flow transformation programs. We formalized two semantics for CF-GKAT. The first is the continuation semantics, where each trace, represented as a guarded word, is augmented with a continuation. The second is the trace semantics, which is obtained by resolving the continuations in the continuation semantics.

We proposed CF-GKAT automata as the operational model for CF-GKAT programs, and showed how to obtain these automata from CF-GKAT expressions while preserving their continuation semantics. By lowering these automata into GKAT automata, we can then build an equivalence checker for CF-GKAT programs that scales nearly linearly with respect to the size of the input CF-GKAT programs. Thus, our work provides an efficient validation algorithm for various control-flow transformations that utilize indicator variables and non-local control flow [16, 49].

While we successfully addressed one of Kozen’s questions [23] by presenting an algorithm to directly convert CF-GKAT programs into automata, we have yet to develop a coalgebraic perspective on non-local control flow utilizing Brzozowski derivatives [8]. Such an approach could streamline several proofs, such as trace preservation of the lowering (Theorem 3.13) and the correctness of the operational semantics (Theorem 3.19), and lead to a memory-efficient on-the-fly algorithm for trace equivalences between CF-GKAT programs. A coalgebraic checker could also make use of symbolic techniques [38] to prevent explicit calculations based on the atoms of a Boolean algebra.

Additional future work could be the inclusion of the `continue` command within loops, as well as other types of control flow found in modern programming languages such as `do-while` and `switch`. In terms of the theory’s extensibility, it would be beneficial to separate the treatment of indicator variables and non-local control. Currently, both components are integrated into the CF-GKAT automata signature as a unified entity. While this approach provides a compact definition of operational semantics, it also introduces complexities when incorporating other non-local controls like `continue`. Specifically, we will need to pass the indicator value in the continuation for `continue`, despite none of the non-local control-flow structures changing the indicator variable.

Acknowledgments

C. Zhang was supported by the US National Science Foundation under awards no. CNS 1845803 and CNS 2040249. T. Kappé was partially supported by the European Union’s Horizon 2020 research and innovation programme under grant no. 101027412 (VERLAN), and partially by the Dutch research council (NWO) under grant no. VI.Veni.232.286 (ChEOPs). N. Naus and D. E. Narváez are supported by the Defense Advanced Research Projects Agency (DARPA) and the Naval Information Warfare Center Pacific (NIWC Pacific) under contract no. N66001-21-C-4028. Additionally, N. Naus is also supported by the US National Science Foundation under award no. CNS 2234257.

We would like to thank Todd Schmid for the insightful discussions while preparing this work and the anonymous POPL reviewers for their thoughtful comments.

References

- [1] Kamal Aboul-Hosn and Dexter Kozen. 2008. Local variable scoping and Kleene algebra with tests. *J. Log. Algebraic Methods Program.* 76, 1 (2008), 3–17. <https://doi.org/10.1016/J.JLAP.2007.10.007>
- [2] Carolyn Jane Anderson, Nate Foster, Arjun Guha, Jean-Baptiste Jeannin, Dexter Kozen, Cole Schlesinger, and David Walker. 2014. NetKAT: semantic foundations for networks. In *POPL*. 113–126. <https://doi.org/10.1145/2535838.2535862>
- [3] Allegra Angus and Dexter Kozen. 2001. *Kleene Algebra with Tests and Program Schematology*. Technical Report TR2001-1844. Cornell University.
- [4] Cheng Zhang Arthur Azevedo de Amorim and Marco Gaboardi. 2024. *Kleene Algebra with Commutativity Conditions Is Undecidable*. <https://hal.science/hal-04534715v2>

- [5] Payas Awadhutkar, Ahmed Tamrawi, Ryan Goluch, and Suresh C. Kothari. 2022. Control flow equivalence method for establishing sanctity of compiling. *Comput. Secur.* 115 (2022), 102608. <https://doi.org/10.1016/J.COSE.2022.102608>
- [6] Thomas Ball, Rupak Majumdar, Todd D. Millstein, and Sriram K. Rajamani. 2001. Automatic Predicate Abstraction of C Programs. In *PLDI*. 203–213. <https://doi.org/10.1145/378795.378846>
- [7] Corrado Böhm and Giuseppe Jacopini. 1966. Flow diagrams, Turing machines and languages with only two formation rules. *Commun. ACM* 9, 5 (1966), 366–371. <https://doi.org/10.1145/355592.365646>
- [8] Janusz A. Brzozowski. 1964. Derivatives of Regular Expressions. *J. ACM* 11, 4 (1964), 481–494. <https://doi.org/10.1145/321239.321249>
- [9] Ying Cao, Runze Zhang, Ruigang Liang, and Kai Chen. 2024. Evaluating the Effectiveness of Decompilers. In *ISSTA*. ACM, 491–502. <https://doi.org/10.1145/3650212.3652144>
- [10] Hugues Cassé, Louis Féraud, Christine Rochange, and Pascal Sainrat. 2002. Une approche pour réduire la complexité du flot de contrôle dans les programmes C. *Tech. Sci. Informatiques* 21, 7 (2002), 1009–1032. <http://tsi.revuesonline.com/article.jsp?articleId=3831>
- [11] Junjie Chen, Jibesh Patra, Michael Pradel, Yingfei Xiong, Hongyu Zhang, Dan Hao, and Lu Zhang. 2021. A Survey of Compiler Testing. *ACM Comput. Surv.* 53, 1 (2021), 4:1–4:36. <https://doi.org/10.1145/3363562>
- [12] Cristina Cifuentes. 1994. *Reverse compilation techniques*. Ph.D. Dissertation. Queensland University of Technology.
- [13] Ernie Cohen, Dexter Kozen, and Frederick Smith. 1996. *The Complexity of Kleene Algebra with Tests*. Technical Report TR96-1598.
- [14] Coq Development Team. 2022. *The Coq Reference Manual, version 8.15*. Available electronically at <http://coq.inria.fr/doc>.
- [15] Sandeep Dasgupta, Sushant Dinesh, Deepan Venkatesh, Vikram S. Adve, and Christopher W. Fletcher. 2020. Scalable validation of binary lifters. In *PLDI*. 655–671. <https://doi.org/10.1145/3385412.3385964>
- [16] Ana M. Erosa and Laurie J. Hendren. 1994. Taming Control Flow: A Structured Approach to Eliminating Goto Statements. In *ICCL*. IEEE Computer Society, 229–240. <https://doi.org/10.1109/ICCL.1994.288377>
- [17] Murdoch James Gabbay and Vincenzo Ciancia. 2011. Freshness and Name-Restriction in Sets of Traces with Names. In *FoSSaCS*. 365–380. https://doi.org/10.1007/978-3-642-19805-2_25
- [18] Benjamin Goldberg, Lenore D. Zuck, and Clark W. Barrett. 2004. Into the Loops: Practical Issues in Translation Validation for Optimizing Compilers. In *COCV@ETAPS*. 53–71. <https://doi.org/10.1016/J.ENTCS.2005.01.030>
- [19] Niels Bjørn Bugge Grathwohl, Dexter Kozen, and Konstantinos Mamouras. 2014. KAT + B!. In *CSL-LICS*. ACM, 44:1–44:10. <https://doi.org/10.1145/2603088.2603095>
- [20] Theodoros Kasampalis. 2021. *Translation validation for compilation verification*. Ph.D. Dissertation. University of Illinois Urbana-Champaign, USA. <https://hdl.handle.net/2142/110460>
- [21] Theodoros Kasampalis, Daejun Park, Zhengyao Lin, Vikram S. Adve, and Grigore Rosu. 2021. Language-parametric compiler validation with application to LLVM. In *ASPLOS*. ACM, 1004–1019. <https://doi.org/10.1145/3445814.3446751>
- [22] Dexter Kozen. 1996. Kleene algebra with tests and commutativity conditions. In *TACAS*. 14–33. https://doi.org/10.1007/3-540-61042-1_35
- [23] Dexter Kozen. 2008. Nonlocal Flow of Control and Kleene Algebra with Tests. In *LICS*. 105–117. <https://doi.org/10.1109/LICS.2008.32>
- [24] Dexter Kozen, Konstantinos Mamouras, Daniela Petrisan, and Alexandra Silva. 2015. Nominal Kleene Coalgebra. In *ICALP*. 286–298. https://doi.org/10.1007/978-3-662-47666-6_23
- [25] Dexter Kozen, Konstantinos Mamouras, and Alexandra Silva. 2017. Completeness and incompleteness in nominal Kleene algebra. *J. Log. Algebraic Methods Program.* 91, 17–32. <https://doi.org/10.1016/J.JLAMP.2017.06.002>
- [26] Dexter Kozen and Maria-Christina Patron. 2000. Certification of Compiler Optimizations Using Kleene Algebra with Tests. In *CL*. 568–582. https://doi.org/10.1007/3-540-44957-4_38
- [27] Dexter Kozen and Frederick Smith. 1996. Kleene Algebra with Tests: Completeness and Decidability. In *CSL*. 244–259. https://doi.org/10.1007/3-540-63172-0_43
- [28] Dexter Kozen and Wei-Lung Dustin Tseng. 2008. The Böhm-Jacopini Theorem Is False, Propositionally. In *MPC*. 177–192. https://doi.org/10.1007/978-3-540-70594-9_11
- [29] Stepan L. Kuznetsov. 2023. On the Complexity of Reasoning in Kleene Algebra with Commutativity Conditions. In *ICTAC*. 83–99. https://doi.org/10.1007/978-3-031-47963-2_7
- [30] Vu Le, Mehrdad Afshari, and Zhendong Su. 2014. Compiler validation via equivalence modulo inputs. In *PLDI*. ACM, 216–226. <https://doi.org/10.1145/2594291.2594334>
- [31] Xavier Leroy, Sandrine Blazy, Daniel Kästner, Bernhard Schommer, Markus Pister, and Christian Ferdinand. 2016. CompCert – A Formally Verified Optimizing Compiler. In *ERTS*.
- [32] Zhibo Liu and Shuai Wang. 2020. How far we have come: testing decompilation correctness of C decompilers. In *ISSTA*. ACM, 475–487. <https://doi.org/10.1145/3395363.3397370>
- [33] T.J. McCabe. 1976. A Complexity Measure. *IEEE Transactions on Software Engineering* SE-2, 4 (1976), 308–320. <https://doi.org/10.1109/TSE.1976.233837>

- [34] Mark Moeller, Jules Jacobs, Olivier Savary Bélanger, David Darais, Cole Schlesinger, Steffen Smolka, Nate Foster, and Alexandra Silva. 2024. KATch: A Fast Symbolic Verifier for NetKAT. In *PLDI*. 1905–1928. <https://doi.org/10.1145/3656454>
- [35] George C. Necula. 2000. Translation validation for an optimizing compiler. In *PLDI*. ACM, 83–94. <https://doi.org/10.1145/349299.349314>
- [36] Georg Neis, Chung-Kil Hur, Jan-Oliver Kaiser, Craig McLaughlin, Derek Dreyer, and Viktor Vafeiadis. 2015. Pilsner: a compositionally verified compiler for a higher-order imperative language. In *ICFP*. ACM, 166–178. <https://doi.org/10.1145/2784731.2784764>
- [37] Amir Pnueli, Michael Siegel, and Eli Singerman. 1998. Translation Validation. In *TACAS*. 151–166. <https://doi.org/10.1007/BF0054170>
- [38] Damien Pous. 2015. Symbolic Algorithms for Language Equivalence and Kleene Algebra with Tests. In *POPL*. 357–368. <https://doi.org/10.1145/2676726.2677007>
- [39] Todd Schmid, Tobias Kappé, Dexter Kozen, and Alexandra Silva. 2021. Guarded Kleene Algebra with Tests: Coequations, Coinduction, and Completeness. , 142:1–142:14 pages. <https://doi.org/10.4230/LIPICS.ICALP.2021.142>
- [40] Thomas Arthur Leck Sewell, Magnus O. Myreen, and Gerwin Klein. 2013. Translation validation for a verified OS kernel. In *PLDI*. ACM, 471–482. <https://doi.org/10.1145/2491956.2462183>
- [41] Steffen Smolka, Nate Foster, Justin Hsu, Tobias Kappé, Dexter Kozen, and Alexandra Silva. 2020. Guarded Kleene algebra with tests: verification of uninterpreted programs in nearly linear time. In *POPL*. 61:1–61:28. <https://doi.org/10.1145/3371129>
- [42] Yong Kiam Tan, Magnus O. Myreen, Ramana Kumar, Anthony C. J. Fox, Scott Owens, and Michael Norrish. 2016. A new verified compiler backend for CakeML. In *ICFP*. ACM, 60–73. <https://doi.org/10.1145/2951913.2951924>
- [43] Robert Endre Tarjan. 1975. Efficiency of a Good But Not Linear Set Union Algorithm. *J. ACM* 22, 2 (1975), 215–225. <https://doi.org/10.1145/321879.321884>
- [44] Balder ten Cate and Tobias Kappé. 2025. Algebras for Deterministic Computation are Inherently Incomplete. In *POPL*. 25:1–25:27. <https://doi.org/10.1145/3704861>
- [45] Ken Thompson. 1968. Regular Expression Search Algorithm. *Commun. ACM* 11, 6 (1968), 419–422. <https://doi.org/10.1145/363347.363387>
- [46] Freek Verbeek, Joshua A. Bockenek, Zhouhai Fu, and Binoy Ravindran. 2022. Formally verified lifting of C-compiled x86-64 binaries. In *PLDI*. ACM, 934–949. <https://doi.org/10.1145/3519939.3523702>
- [47] Freek Verbeek, Nico Naus, and Binoy Ravindran. 2024. Verifiably Correct Lifting of Position-Independent x86-64 Binaries to Symbolized Assembly. In *CCS*.
- [48] Freek Verbeek, Pierre Olivier, and Binoy Ravindran. 2020. Sound C Code Decompilation for a Subset of x86-64 Binaries. In *SEFM*. 247–264. https://doi.org/10.1007/978-3-030-58768-0_14
- [49] Khaled Yakdan, Sebastian Eschweiler, Elmar Gerhards-Padilla, and Matthew Smith. 2015. No More Gotos: De-compilation Using Pattern-Independent Control-Flow Structuring and Semantic-Preserving Transformations. In *NDSS*. <https://www.ndss-symposium.org/ndss2015/no-more-gotos-decompilation-using-pattern-independent-control-flow-structuring-and-semantics>
- [50] Xuejun Yang, Yang Chen, Eric Eide, and John Regehr. 2011. Finding and understanding bugs in C compilers. In *PLDI*. ACM, 283–294. <https://doi.org/10.1145/1993498.1993532>
- [51] Cheng Zhang, Tobias Kappé, David E. Narváez, and Nico Naus. 2024. *CF-GKAT: Efficient Validation of Control-Flow Transformations (Artifact)*. <https://doi.org/10.5281/zenodo.13938565>
- [52] Yiji Zhang and Lenore D. Zuck. 2018. Formal Verification of Optimizing Compilers. In *ICDCIT*. 50–65. https://doi.org/10.1007/978-3-319-72344-0_3

Received 2024-07-11; accepted 2024-11-07